Beyond Introspection: Reinforcing Thinking via Externalist Behavioral Feedback

Diji Yang[1*] **Linda Zeng**[2*†] **Kezhen Chen**[3‡] **Yi Zhang**[1]
[1]University of California, Santa Cruz
[2]The Harker School, [3]Meta
{dyang39, yiz}@ucsc.edu
{lindazeng979, kzchen0204}@gmail.com

Abstract

While inference-time thinking allows Large Language Models (LLMs) to address complex problems, the extended thinking process can be unreliable or inconsistent because of the model's probabilistic nature, especially near its knowledge boundaries. Existing approaches attempt to mitigate this by having the model critique its own reasoning to make corrections. However, such self-critique inherits the same biases of the original output, known as the introspection illusion. Moving beyond such introspection and inspired by core methodologies in ethology, we propose an externalist three-step framework **D**istillation-**R**einforcement-**R**easoning (**DRR**). Rather than relying on a model's introspection, DRR evaluates its observable behaviors to provide corrective feedback. DRR first distills the reasoner's behavioral traces, then trains a lightweight, external Discriminative Model (DM). At inference time, this DM acts as a critic, identifying and rejecting suspicious reasoning steps. This external feedback compels the LLM to discard flawed pathways and explore alternatives, thereby enhancing reasoning quality without altering the base model. Experiments on multiple reasoning benchmarks show that our framework significantly outperforms prominent self-critique methods. Benefiting from a lightweight and annotation-free design, DRR offers a scalable and adaptable solution for improving the reliability of reasoning in a wide range of LLMs.

1 Introduction

Articulating a thinking process, often termed chain-of-thought, shows great success in enhancing the reasoning capabilities of Large Language Models (LLMs) [54, 36]. However, the extended token sequences required for such thinking are vulnerable to error accumulation [6, 39]. This vulnerability stems from the probabilistic nature of autoregressive models, which can produce divergent outputs when reasoning over uncertain content (i.e., near the knowledge boundary) [22].

A seemingly straightforward approach to this problem would be to expand the model's knowledge boundaries through further training. However, this strategy is fundamentally limited. Knowledge boundaries are an inherent property of any model and intelligence; for any given state of training, a new frontier of uncertainty will inevitably emerge when faced with novel problems. Therefore, simply fine-tuning on more data merely shifts the boundary rather than eliminating the core issue of unreliability at its edge.

This insight suggests that a more robust solution lies not in attempting to create an omniscient model, but in equipping it with the ability to accurately assess the quality of its reasoning outputs

[*]Equal contribution.
[†]Work done during internship at UCSC.
[‡]Work done as external collaboration.

at inference time, i.e., knowing what is unknown. The critical capability, therefore, is to reliably identify moments of uncertainty as they occur. Building on this concept, recent research attempted self-critique mechanisms, enabling models to assess and refine their outputs by leveraging their pre-trained knowledge or by estimating confidence scores [38]. However, these strategies have limitations, linked to the introspection illusion [40], due to the models' inherent biases and tendencies towards hallucinations [45, 21, 46]. Moreover, instead of providing faithful representations of internal states, LLMs often perform mere mimicry or role-playing of human introspective data [12]. Additionally, methods that rely on accessing internal model states are impractical for closed-source LLMs, restricting their applicability.

Inspired by core methodologies in ethology, where behavior is analyzed externally without inferring unmeasurable internal states [49, 41], we propose an externalist approach to augmenting LLM reasoning. We introduce Distillation-Reinforcement-Reasoning (DRR), a framework that treats the LLM as a fixed reasoner and uses an external module to empower the LLM with a multi-step self-correction mechanism through an in-context reinforcement learning process [35]. DRR first distills the reasoner's natural problem-solving attempts into behavioral traces. It then uses this data to train a lightweight, external Discriminative Model (DM) that learns to identify patterns of successful versus flawed reasoning. At inference time, this DM acts as a critic, assessing the LLM's reasoning steps and rejecting those it deems suspicious. This explicit external verbal feedback reward [43] compels the LLM to abandon flawed pathways and explore alternative lines of reasoning—a crucial mechanism for overcoming the cognitive fixation that often limits self-critique. The resulting interaction forms a feedback loop analogous to reinforcement in behaviorism, where external feedback guides an agent to improve its decision-making over time. This dynamic fosters a more robust and deliberative thinking process, thereby enhancing the quality and reliability of the final answer without altering the base model's weights. Moreover, the entire procedure is supported by an automated data-generation pipeline, eliminating the need for expensive, human-labeled process supervision data.

Our contributions can be summarized as follows:

- We introduce and formalize an externalist paradigm for reinforcing LLM reasoning at inference time. The DRR framework serves as a concrete realization of this paradigm, designed to overcome the fundamental limitations and inherited biases of internal self-critique.
- We demonstrate a model-agnostic design for this paradigm, while its automated data distillation process eliminates the need for expensive manual annotation, ensuring practical scalability.
- We provide extensive empirical validation of the externalist approach, which is shown to be beneficial for both open-source and closed-source LLMs across various benchmarks. [1]

2 Related Work

The probabilistic nature of LLMs, which are optimized for statistical plausibility over factual accuracy, is a fundamental source of potential unreliability [6]. Stochasticity introduced during decoding is amplified at the model's knowledge boundaries, where output distributions flatten. During extended generation, such as inference-time thinking, this allows plausible-sounding errors to propagate and compound with each subsequent token. This phenomenon of compounding errors, identified as the compositionality gap, can undermine the logical integrity of a long chain of thought [39, 27]. Consequently, error accumulation can lead to a final output that is logically flawed despite appearing coherent [6]. The straightforward way to combat this problem is to apply intervention on the model's reasoning chain directly after discovering the flaw, but discovering the flaw itself is challenging, as it requires the system to know what it does not know, which is a capability often referred to as meta-cognition [16, 29] or self-awareness through self-criticism [38]. Existing research on system-level self-criticism can be categorized into two approaches: feedback generated by the LLM itself (LLM Self-critics) and feedback provided by additional components within the AI system (External Feedback). [2]

[1] All works are open sourced at https://github.com/dyang39/DRR

[2] In this work, while we explicitly distinguish between LLM Self-critics and External Feedback, we note that from the system's perspective, all critic and feedback (whether from the LLM or external modules) remains internal to the AI system.

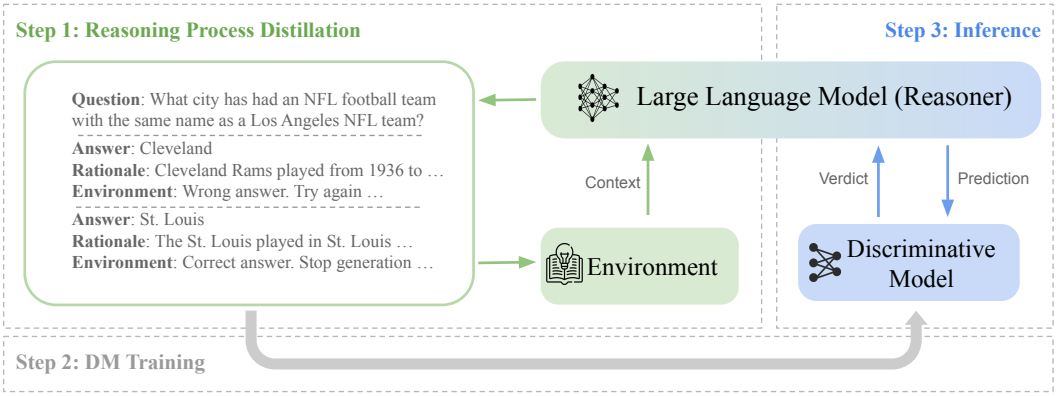

Figure 1: Overview of the three-step Distllation-Reinforcement-Reasoning (DRR) pipeline.

LLM Self-critics LLM self-critic mechanisms empower language models to assess and refine their own outputs [25, 59, 3, 1]. There are two predominant strategies for LLMs to self-present their feedback [38]. A straightforward approach leverages the model's pre-trained knowledge to critique and enhance its responses [52, 56, 31, 53, 60]. While this method utilizes the model's inherent capabilities, it may not be able to reliably identify its own errors due to its inherent biases and tendencies to hallucinations [45]. The second approach involves using confidence scores or uncertainty estimates to evaluate the quality of the model's outputs [57, 15, 8, 61]. This strategy typically requires manually crafted criteria, which can be labor-intensive and lack a theoretical foundation for optimal configuration, thereby limiting its reliability and generalizability [38]. Moreover, the need to access the internal state of the model closes off the opportunity for its application to good-performing closed-source LLM APIs such as ChatGPT and Gemini.

Inspired by the fact that ethologists draw conclusions about animals through observable behavior and avoid making assumptions about internal states that can hardly be directly measured [49, 41], our approach emphasizes the analysis of real observable behavior (i.e., the outputs of LLM) rather than defining internal states (e.g., confidence thresholds).

External Feedback Recent studies have shown that external modules can provide meaningful feedback during the generation process of LLMs, thereby enhancing system performance [19]. These modules encompass a diverse range of variations, including tools like Code Interpreters [9], Search Engines [50], and other specialized software [26], as well as carefully designed rules [58]. MemPrompt [30] predefines a dictionary of possible LLM outputs and their corresponding scores to offer feedback at inference time, similar to our analysis of LLM behavior. While these approaches effectively respond to LLM behavior across various benchmarks, however, these non-deep learning tools tend to limit the flexibility of the system.

On the other hand, since RLHF [37], using external deep learning models as a reward provider has also been widely studied recently [5, 17]. These works focus on using the feedback provided by the reward model to fine-tune the LLM instead of directly participating in the generation at inference time. Consequently, these resource-intensive approaches require large amounts of high-quality labeled data to train the reward models and substantial computational resources to perform LLM tuning.

In contrast, our training only involved a small classifier fine-tuning on the LLM's behavioral data. This data can be easily generated from any raw data with only the required initial input and final output, which are typically available for most datasets.

Hallucination Detection Hallucination detection has emerged as a pivotal task in the era of large language models, aiming to differentiate between hallucinated and accurate content using binary classification [23]. While our work does not specifically focus on hallucination detection, it shares certain methodological similarities with one prominent approach in that field: training hallucination classifiers to detect hallucinations based on the model's internal states, such as hidden layer representations [4, 47]. However, this method requires direct access to the model's internal

workings, making it unsuitable for closed-source systems. A recent approach, RelD [10], utilizes hallucination data from LLMs and designs complex training objectives to align a hallucination detector with human evaluation metrics. Inspired by this line of research, our approach circumvents the complexities of human alignment and intricate training designs. Instead, we rely entirely on behavioral data from model outputs, which allows our DM to prevent hallucinations. More importantly, since our focus extends beyond hallucination detection, our system considers this as a mid-step signal and analyzes how this signal can be used to further optimize the reasoning chains of LLMs, thereby improving the overall performance in general natural language generation tasks.

3 Approach

In this section, we introduce our main framework for enhancing LLM inference-time reasoning. As shown in Figure 1, it consists of three steps: behavioral data generation via Reasoning Process Distillation, where LLM reasoning patterns are collected as training data (Section 3.1); training of the DM, where the DM learns to assess LLM behavior (Section 3.2); and system deployment for inference, during which the LLM and the DM interact over multiple iterations in order to produce a final answer or abstention (Section 3.3).

3.1 Reasoning Process Distillation

Algorithm 1 Reasoning Process Distillation Algorithm

```
Notation:
R: Reasoner LLM
Q: Question, A: Ground-truth Answer, C: Context (initialized as empty)
r: Generated rationale, A': LLM-generated answer
for each data point do
    C ← ∅
    while true do
        (A', r) ← R(Q, C)
        if A' = A then
            Record {Q, C, A', r} with verdict Accept
            break {Exit the while loop}
        else
            Record {Q, C, A', r} with verdict Reject
        end if
        C ← C ∪ {A', r}
        exit while loop if maximum turns reached.
    end while
end for
```

Algorithm 1 depicts our semi-supervised Reasoning Process Distillation algorithm, which collects data for DM training. This algorithm simulates the LLM's inference-time behavior in multiple reasoning iterations using raw training data. For each question Q, the Reasoner LLM $\mathcal{R}$ generates an answer A' along with its rationale r. A binary verdict label is assigned to indicate whether the DM should accept or reject the response, depending on whether A' matches the ground-truth answer A. For all incorrect answers, $\mathcal{R}$ is re-prompted with the question Q and past history (i.e., context C) to explore alternative reasoning paths. This iterative process continues until $\mathcal{R}$ generates the correct answer A, signaling that the DM should accept the response and stop further reasoning. If no accurate response is generated after a predefined number of iterations, the process terminates to avoid an infinite loop (more discussions in Appendix A). The generated data reflect a diverse range of LLM behaviors that the DM may encounter across various turns during inference, allowing the DM to directly learn the unconscious reasoning patterns that LLM tends to generate when leading to a correct or incorrect answer. By leveraging LLM's natural behaviors with its environment, this generation method streamlines the feedback collection process, relying solely on the input and label from source data, which are typically available for most NLP and AI datasets.

3.2 Discriminative Model Training

Given a question, its past context, and an LLM-generated answer with rationale, the goal of the DM is to predict whether the new answer should be accepted or rejected. In other words, the DM acts as a binary classification model, and its training on a distribution over (x, y) is alignment between the input tuple $x_i = \{Q, C, A', r\}$ and the label $y \in \{\texttt{Accept}, \texttt{Reject}\}$. Through observing both the LLM's previous responses and its new rationale, which may include corrections to past reasoning or further explorations into internal knowledge, the DM learns to be both a coach, supervising the LLM's behavior across multiple iterations, and a judge, determining the reliability of its final response. This enables the DM to determine if the LLM's response is influenced by hallucinations and when it should stop.

Specifically, the DM predicts logits $\mathbf{z} \in \mathbb{R}^2$, representing unnormalized probabilities for the two classes. To account for the higher cost of false-positive errors (i.e., Accepting an incorrect answer), the weight w_1 for class `Accept` is set higher than w_0, the weight for class `Reject`. The weighting design encourages DM to adopt stricter acceptance criteria, prioritizing the reliability of the final decision and reducing the likelihood of harmful false-positive predictions (see Section 6 for details). The training objective is to minimize a weighted cross-entropy loss as shown in Equation 1, where $P(y|x;\theta)$ is the predicted probability of the class y, and w_y represents class-specific weights.

$$\mathcal{L}_{\text{DM}}(\theta) = -\,\mathbb{E}_{(x,y)\sim\mathcal{D}}\big[w_y \,\log P\big(y \mid x;\theta\big)\big] \tag{1}$$

By utilizing the behavioral data described in Section 3.1, the DM is exposed to diverse scenarios reflective of inference-time conditions. This approach eliminates the reliance on external reward signals, commonly collected through human feedback, as the DM inherently models these features by learning directly from the dataset.

3.3 Inference

At inference, the system features an iterative line of exchanges between the Reasoner LLM and the DM. While the LLM's goal is to provide a final answer to a given question, the DM assesses LLM reasoning and provides feedback signals through accepting or rejecting its responses.

Similar to the data generation process in Section 3.1, the LLM produces answers and rationales across multiple iterations. However, during inference, the stopping condition is based on the DM's feedback rather than comparisons to the ground truth. The DM evaluates each LLM-generated answer and rationale in real-time, predicting `Accept` or `Reject` verdicts based on the question, context, answer, and rationale of each turn. This process guides the system toward a reliable final response through recognizing patterns in responses that indicate hallucinations or errors in reasoning. In general, this architecture offers flexibility by supporting adaptive reasoning iterations. The DM dynamically decides whether further exploration is necessary based on the behavioral data it was trained on. This allows the system to efficiently balance between preventing overextended reasoning and addressing incomplete or incorrect responses. This adaptability mirrors human problem-solving, where some answers are instinctively clear, while others demand deeper reasoning.

4 Experiment

4.1 Environment Setup

Tasks and Datasets Our approach is evaluated on commonsense and knowledge-intensive reasoning tasks. The former involves general, everyday scenarios, while the latter requires applying specialized factual knowledge to solve complex problems. We assess commonsense reasoning with CommonsenseQA [48] and WinoGrande [42], and knowledge-intensive reasoning with OpenBookQA [33] and PIQA [7]. To generate mid-step behavioral data for DM training, we use QA pairs from the training sets. We evaluate OpenBookQA on the test set and the other datasets on the validation sets.

Evaluation We use standard exact matching accuracy (Acc) as our primary measure; however, it may not fully capture user satisfaction, since an incorrect or misleading answer can be more disappointing than no answer. Therefore, we also include Formula Score (FS) [13], inspired by

		Commonsense						Knowledge Intensive						Overall		
		CommonSenseQA			WinoGrande			OpenBookQA			PIQA			Combined		
		Acc	FS(1)	FS(3)	Acc	FS(1)	FS(3)	Acc	FS(1)	FS(3)	Acc	FS(1)	FS(3)	Acc	FS(1)	FS(3)
Open Source, ZS	Self-Talk (Llama3)*	70.6	-	-	-	-	-	72.2	-	-	77.2	-	-	-	-	-
	CoT (Llama3)	73.5	47.0	-6.1	59.0	17.9	-64.2	70.0	40.0	-20.0	78.1	56.3	12.5	72.5	48.5	-9.9
	Abstain (Llama3)	73.1	46.3	-7.5	59.3	18.9	-62.9	70.2	40.4	-19.2	66.5	33.1	-33.8	66.5	35.6	-33.4
Open Source, Learned	Crystal (Llama3)*	75.1	-	-	-	-	-	72.6	-	-	78.2	-	-	-	-	-
	SFT-LoRA (Llama3)	77.2	54.5	8.9	76.4	52.8	5.6	80.6	61.2	22.4	82.0	64.0	28.0	79.8	54.5	19.0
	DRR (Llama3)	**82.3**	**65.5**	**33.6**	**77.0**	**54.1**	**8.1**	**81.2**	**63.0**	**26.6**	**83.6**	**67.3**	**34.8**	**81.8**	**65.5**	**28.3**
Closed Source	CoT (GPT-4)	83.5	67.1	34.2	74.6	49.2	-1.7	**91.2**	82.4	64.8	91.1	82.2	64.3	86.1	65.5	44.5
	Abstain (GPT-4)	82.6	65.1	30.2	74.4	48.9	-2.3	85.6	71.2	42.4	85.4	70.8	33.0	82.6	59.3	25.8
	Self-Critic (GPT-4)	74.6	49.2	-1.6	77.0	54.1	8.1	88.8	77.6	55.2	88.7	77.4	54.7	83.4	60.5	33.7
	DRR (GPT-4)	**86.2**	**74.4**	**52.2**	**79.5**	**59.2**	**19.9**	**91.2**	**82.8**	**67.6**	**91.5**	**82.9**	**66.3**	**87.9**	**69.4**	**53.9**

Table 1: Results on four question answering datasets. Each method is evaluated by Accuracy (Acc) and two Formula Scores FS(1) and FS (3). The Overall (Combined) results aggregate all data points across datasets into one unified set for holistic evaluation. Methods marked with * indicate results reported by prior works.[3]

scoring systems in mathematical competitions [32], which assign negative scores to incorrect answers and no penalties for abstentions. Specifically, each correct answer yields S_{correct}, each incorrect answer yields $S_{\text{incorrect}}$, and an abstention yields S_{abstain}. n our experiments, these are set to 1, -1, and 0, respectively, which we denote as FS(1). This scoring scheme can be generalized to adjust the relative penalties, depending on the importance of reliability. For instance, FS(3) increases the penalty for errors to -3, placing greater emphasis on avoiding incorrect answers. By reporting both accuracy and FS, we provide a more balanced evaluation of QA performance, especially in settings where reliability and safety are critical (e.g., medical or legal advice).

Additionally, we compute the critic-decision accuracy Acc(D), which measures the accuracy of a system's decisions to give an answer or abstain (i.e. how well it decides to abstain or not). For DRR, this refers to the DM accuracy at the last reasoning turn before an answer is given. For zero-shot methods, this refers to the percentage of LLM's successful answers and abstentions.

4.2 Implementation Details

Model Choice We use Llama3-8b [14] and GPT-4 as the LLM, and we finetune the Flan-T5-783M [11] as the DM. Our choice of LLMs reflects both open-source and closed-source models to demonstrate the flexibility of our system. Although the DM can take on any classifier architecture, we adopt Flan-T5 for its strong trade-off between performance and size on binary classification tasks. Smaller discriminative models are particularly practical, as they require fewer parameters, retrain efficiently, and operate independently of the main LLM. Prior work, such as Tyen et al. [51], has also shown the effectiveness of lightweight classifiers for detecting LLM errors. Building on these insights, our design leverages the LLM's generative capacity and the DM's discriminative specialization to form an efficient self-correction framework, reducing the training burden on the LLM while constructing a simple classification task that a small DM can perform.

Fine-tuning Details The DM is fine-tuned on the behavioral data generated for all datasets from the first Reasoning Process Distillation step of DRR. Note that our method does not require any optimization for individual datasets; that is, for all experiments, we train one classifier for all four datasets. To account for the more serious consequences of false-positive errors (i.e., accepting an incorrect answer; see Section 6 for more explanation), the weight factor w_1 for class `Accept` is set three times higher than w_0, the weight for class `Reject`. During loss calculation, this class-specific weight factor is multiplied by the cross-entropy shown in Equation 1. For reproduction concern, we detail the data preparation process, training hyperparameters, prompt configurations, and computing cost in Appendices A and B.

[3]The original works use models different from Llama3. For a fair comparison, we report the performance from recent reproduction [34]. Due to limited access to the prediction file, the FS and Overall score are unavailable.

5 Results and Analysis

Baseline methods Table 1 categorizes baselines by model availability, distinguishing between Open Source and Closed Source methods.

For baselines using an open-source model, we first include three zero-shot methods. Self-Talk [44], as one of the pioneering works leveraging model self-awareness (for information seeking) in a zero-shot setting, has recently been enhanced with the support of a LLM (Llama3) and updated prompting methods, establishing it as a strong baseline [34]. Then, we evaluate two widely used prompting-based methods: CoT used a well-tuned QA prompt to encourage LLM to think step by step, while Abstain included an abstain [55] option to give up answering when LLM thinks it is not confident. Notably, DRR includes abstentions by treating any instances where the DM consistently rejects responses until the maximum number of attempts is reached as an abstention action. In addition to ZS approaches, we evaluate two Learned methods. Crystal [28] is a system-level self-critic approach similar to Self-Talk, but it tunes the model for self-feedback ability via reinforcement learning. We also conduct a supervised fine-tuning baseline (SFT-LoRA) to enable a more direct and fair comparison to our approach using the same amount of training data. To match the trained parameters used in our system, we apply Low-Rank Adaptation (LoRA) [20] for parameter-efficient fine-tuning. This baseline highlights the comparison between the standard outcome fine-tuning approach and DRR.

For methods supporting Closed Source models, we include CoT and Abstain under the same settings as above but with strong backbone LLM (GPT-4). Additionally, we built a prompt-based self-reasoning baseline (Self-Critic) following the implementation by Huang et al. [21]. This evaluates the LLM's self-awareness to refine its responses without external feedback.

Results Table 1 reports the experimental results. DRR outperforms all baselines under both Open Source and Closed Source settings.

For the Open Source setting, DRR with Llama3 demonstrates consistent improvements in all metrics from the three ZS baselines. Raising overall accuracy from the CoT baseline of 72.5% to 81.8%, DRR demonstrates its ability to enhance self-correction via iterative reasoning. While the Abstain allows the LLM to opt-out, it performs worse than CoT, even in FS(3), indicating that LLMs struggle with the self-awareness to abstain at the right time. Compared to learned methods, DRR surpasses Crystal on three available datasets, benefiting from external DM feedback without requiring expensive RL. Unlike SFT-LoRA, which does not provide reasoning chains during inference time, DRR generates a mid-step thinking path, improving interpretability and achieving a 2.5% accuracy gain. These gains are even more pronounced in FS, where DRR outperforms the Abstain baseline by 29.9% for FS(1) and 61.7% FS(3), showing its effectiveness in mitigating incorrect responses through selective abstention.

For the Closed Source setting, DRR with GPT-4 surpasses CoT, Abstain, and Self-Critic by 1.8%, 5.3%, and 4.5% in overall accuracy, respectively. DRR's high improvement over Self-Critic in FS(3) (53.9% vs. 33.7%) emphasizes the advantage of external feedback over self-reasoning, especially for closed-source models where fine-tuning is infeasible. Notably, Self-Critic performs worse than CoT, dropping 3.1% in accuracy, aligning with Huang et al. [21] on LLM self-awareness limitations. With heightened punishment to errors, Formula Score further illustrates DRR's advantage: while accuracy gains 5.3% from Abstain, FS(1) and FS(3) increase by 10.1% and 28.1%, demonstrating its effectiveness in scenarios where reliability is critical and abstaining is preferred over hallucinating.

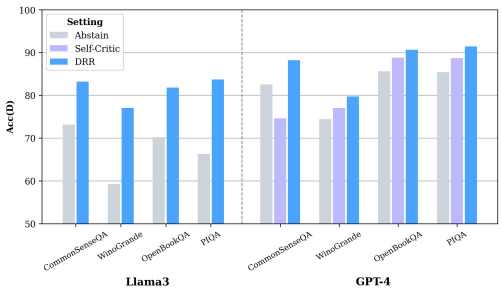

Figure 2: Critic-decision accuracy Acc(D) for Abstain, Self-Critic, and DRR settings using Llama3 and GPT-4 as Reasoner.

Critic-Decision Accuracy Figure 2 compares DRR's critic-decision accuracy against Abstain and Self-Critic baselines, demonstrating the DM's positive impact on inference-time decision-making. As

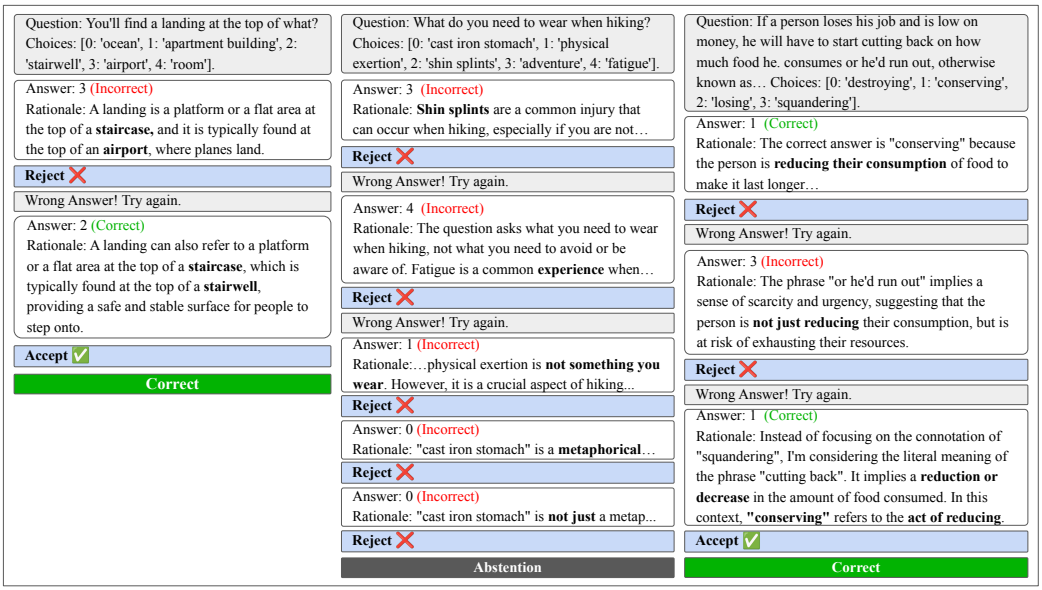

Figure 3: Examples from CommonsenseQA and OpenBookQA with DRR. *Left*: successful correction and acceptance; *Center*: continued rejection leading to abstention; *Right*: false rejection followed by successful acceptance.

previously discussed, the critic's decision ability inherently depends on the model's self-awareness of its knowledge. Using Llama3 as the Reasoner, DRR achieves 81.4% accuracy across four datasets, surpassing Abstain's 67.2%. For GPT-4, DRR reaches 87.5%, outperforming Abstain (82.0%) and Self-Critic (82.3%) across all datasets. By achieving better Acc(D) compared to baselines where the LLM makes decisions independently, the use of the DM facilitates more accurate answers and meaningful abstentions during inference-time reasoning.

6 Discussion

Qualitative Examples Figure 3 displays various examples of the common behavior of our system at inference time. The first example shows optimal behavior, where the LLM responds incorrectly with disjoint reasoning in its rationale on the first turn, the DM identifies this hallucination and rejects the response, and the LLM retries on the second turn, adequately correcting its answer and rationale, leading to an accepted correct answer. We also observe this behavior occurring flexibly across more than two turns. The second example shows our system's ability to mitigate cases where LLM can never reach a correct answer. The DM consistently rejects faulty responses, where the LLM answer does not match an otherwise correct rationale (Turn 1), where the rationale contradicts itself (Turn 2), or where the rationales do not directly answer the question (subsequent turns), until the maximum turns are reached, leading to abstention. While these abstentions are considered incorrect answers for the accuracy score, they mitigate the negative effects measured by FS, highlighting the metric's importance. The third example demonstrates the system's ability to mitigate errors by correcting itself, both on the LLM and DM sides. Initially, the DM incorrectly rejects the answer, but later corrects itself and accepts the correct response. This illustrates how our system can prompt the LLM to reconsider its initial answer and provide a more logical and well-explained rationale in subsequent turns. This approach reduces the impact of false negatives from the DM, making such errors less costly to overall system performance.

Generative vs. Discriminative Our system relies on a lightweight DM to identify faulty reasoning steps rather than a single generative model handling both answer production and error detection. We argue that discriminative decisions, such as binary classification of correctness, can be less complex than open text generation. The DM only needs to learn whether the LLM's output is acceptable or

flawed, rather than produce the entire answer. In practice, as illustrated in Figure 3, identifying a hallucinated statement or a gap in reasoning is often more straightforward than generating a complete solution from scratch.

This observation aligns with adversarial training research in image generation, where the discriminator often outperforms the generator in classification, sometimes leading to mode collapse [18, 2]. Although we differ from image-generation settings, this suggests that a specialized classifier can outpace a generative model at spotting errors. Similarly, recent self-critique work shows that LLMs are not inherently strong in discriminative tasks like self-correction unless trained for them [24]. In our approach, the DM handles binary classification, freeing the LLM to focus on generative reasoning.

The Impact of False Positives False positives, where the DM incorrectly accepts a wrong LLM answer, are more costly compared to DM rejections in our system because acceptances immediately end the reasoning process and provide no opportunity for future correction. On the other hand, false negatives, which are rejections of correct answers, proceed to another round of reasoning, where the DM or the LLM can correct its behavior (i.e., the third example in Figure 3). Motivated by the idea that the DM should be more cautious in choosing to accept an answer than to reject one, we introduce weighted training to the DM training (detailed in Section 3.2). This weighting design encourages DM to adopt stricter acceptance criteria, prioritizing the reliability of the final decision and reducing the number of harmful false-positive predictions.

7 Limitations

Inference-time latency is one of the key limitations of the DRR framework. Unlike single-pass approaches with one LLM call, such as standard Chain-of-Thought prompting, DRR requires multiple interactions between the LLM and the DM. Similar to other models with inference-time thinking capability like OpenAI-o1 [36], the thinking period makes the method less suitable for real-time or latency-sensitive applications where efficiency is a primary concern.

To preserve pre-trained LLM's well-tuned alignments and to support the use of closed-source models for real-world usability, this paper does not explore modifying model weights or accessing internal states. Although this makes DRR broadly applicable, it also limits the system to the pre-trained capability of LLM. Training the LLM could improve both the LLM's reasoning and the quality of behavioral data for the DM, potentially enhancing inference-time reasoning. Studying DRR in trainable settings remains a meaningful direction for future work.

While DRR is designed to be model-independent and can be used in a wide range of domains, its generalizability to specialized domains such as law, medicine, or scientific reasoning remains to be explored. Due to space constraints, we leave the extensive domain adaptation studies to future work. Exploring strategies to improve robustness in different domains, such as incorporating domain-specific cues or augmenting training data, could further improve the applicability of DRR in real-world scenarios.

8 Conclusion

LLMs often exhibit unreliable inference-time reasoning, where self-critique fails due to inherited biases. Furthermore, LLMs face the pressing challenge of diminishing fresh training data as they need ever-increasing high-quality inputs to refine their advanced capabilities. To address these challenges, we proposed Distillation-Reinforcement-Reasoning (DRR), a framework that enhances LLM's inference-time reasoning via in-context reinforcement learning supervised by a lightweight discriminative model. This model is trained from synthetic data generated from the LLM through a novel reasoning process distillation algorithm, providing an avenue for generating fresh behavioral training data without relying on introspection. Empirical results on multiple standard QA benchmarks demonstrate the effectiveness of this framework over self-critique approaches. We expect this low-cost and externalist solution to provide a practical path for not only LLM, but also more complex agentic systems [62] to continually improve their reasoning quality and reliability at inference time without relying on flawed introspection.

References

[1] Alfonso Amayuelas, Liangming Pan, Wenhu Chen, and William Wang. Knowledge of knowledge: Exploring known-unknowns uncertainty with large language models. *arXiv preprint arXiv:2305.13712*, 2023.

[2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 214–223, 2017.

[3] Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. Self-rag: Learning to retrieve, generate, and critique through self-reflection. In *The Twelfth International Conference on Learning Representations*, 2023.

[4] Amos Azaria and Tom Mitchell. The internal state of an llm knows when it's lying. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 967–976, 2023.

[5] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.

[6] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623, 2021.

[7] Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7432–7439, 2020.

[8] Chao Chen, Kai Liu, Ze Chen, Yi Gu, Yue Wu, Mingyuan Tao, Zhihang Fu, and Jieping Ye. Inside: Llms' internal states retain the power of hallucination detection. In *The Twelfth International Conference on Learning Representations*, 2024.

[9] Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. Teaching large language models to self-debug. In *The Twelfth International Conference on Learning Representations*, 2023.

[10] Yuyan Chen, Qiang Fu, Yichen Yuan, Zhihao Wen, Ge Fan, Dayiheng Liu, Dongmei Zhang, Zhixu Li, and Yanghua Xiao. Hallucination detection: Robustly discerning reliable answers in large language models. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 245–255, 2023.

[11] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53, 2024.

[12] Iulia M Comsa and Murray Shanahan. Does it make sense to speak of introspection in large language models? *arXiv preprint arXiv:2506.05068*, 2025.

[13] Frederick B Davis. A note on the correction for chance success. *The Journal of Experimental Education*, 35(3):42–47, 1967.

[14] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

[15] Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017):625–630, 2024.

[16] John H Flavell. Metacognition and cognitive monitoring: A new area of cognitive–developmental inquiry. *American psychologist*, 34(10):906, 1979.

[17] Amelia Glaese, Nat McAleese, Maja Trębacz, John Aslanides, Vlad Firoiu, Timo Ewalds, Maribeth Rauh, Laura Weidinger, Martin Chadwick, Phoebe Thacker, et al. Improving alignment of dialogue agents via targeted human judgements. *arXiv preprint arXiv:2209.14375*, 2022.

[18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

[19] Zhibin Gou, Zhihong Shao, Yeyun Gong, Yujiu Yang, Nan Duan, Weizhu Chen, et al. Critic: Large language models can self-correct with tool-interactive critiquing. In *The Twelfth International Conference on Learning Representations*, 2024.

[20] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2021.

[21] Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. Large language models cannot self-correct reasoning yet. In *The Twelfth International Conference on Learning Representations*, 2023.

[22] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, 43(2):1–55, 2025.

[23] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38, 2023.

[24] Dongwei Jiang, Jingyu Zhang, Orion Weller, Nathaniel Weir, Benjamin Van Durme, and Daniel Khashabi. Self-[in] correct: Llms struggle with refining self-generated responses. *arXiv preprint arXiv:2404.04298*, 2024.

[25] Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*, 2022.

[26] Geunwoo Kim, Pierre Baldi, and Stephen McAleer. Language models can solve computer tasks. *Advances in Neural Information Processing Systems*, 36, 2024.

[27] Tamera Lanham, Anna Chen, Ansh Radhakrishnan, Benoit Steiner, Carson Denison, Danny Hernandez, Dustin Li, Esin Durmus, Evan Hubinger, Jackson Kernion, et al. Measuring faithfulness in chain-of-thought reasoning. *arXiv preprint arXiv:2307.13702*, 2023.

[28] Jiacheng Liu, Ramakanth Pasunuru, Hannaneh Hajishirzi, Yejin Choi, and Asli Celikyilmaz. Crystal: Introspective reasoners reinforced with self-feedback. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 11557–11572, 2023.

[29] Li Liu, Diji Yang, Sijia Zhong, Kalyana Suma Sree Tholeti, Lei Ding, Yi Zhang, and Leilani H Gilpin. Right this way: Can vlms guide us to see more to answer questions? In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.

[30] Aman Madaan, Niket Tandon, Peter Clark, and Yiming Yang. Memory-assisted prompt editing to improve gpt-3 after deployment. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2833–2861, 2022.

[31] Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36, 2024.

[32] Mathematical Association of America. Amc 10, 2024. Accessed: 2024-10.

[33] Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2381–2391, 2018.

[34] Francesco Molfese, Simone Conia, Riccardo Orlando, and Roberto Navigli. Zebra: Zero-shot example-based retrieval augmentation for commonsense question answering. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 22429–22444, 2024.

[35] Giovanni Monea, Antoine Bosselut, Kianté Brantley, and Yoav Artzi. Llms are in-context reinforcement learners. *arXiv preprint arXiv:2410.05362*, 2024.

[36] OpenAI. Openai o1, September 2024.

[37] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.

[38] Liangming Pan, Michael Saxon, Wenda Xu, Deepak Nathani, Xinyi Wang, and William Yang Wang. Automatically correcting large language models: Surveying the landscape of diverse self-correction strategies. *arXiv preprint arXiv:2308.03188*, 2023.

[39] Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah A Smith, and Mike Lewis. Measuring and narrowing the compositionality gap in language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5687–5711, 2023.

[40] Emily Pronin. The introspection illusion. *Advances in experimental social psychology*, 41:1–67, 2009.

[41] Iyad Rahwan, Manuel Cebrian, Nick Obradovich, Josh Bongard, Jean-François Bonnefon, Cynthia Breazeal, Jacob W Crandall, Nicholas A Christakis, Iain D Couzin, Matthew O Jackson, et al. Machine behaviour. *Nature*, 568(7753):477–486, 2019.

[42] Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adversarial winograd schema challenge at scale. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8732–8740, 2020.

[43] Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36, 2024.

[44] Vered Shwartz, Peter West, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Unsupervised commonsense question answering with self-talk. *arXiv preprint arXiv:2004.05483*, 2020.

[45] Kaya Stechly, Matthew Marquez, and Subbarao Kambhampati. Gpt-4 doesn't know it's wrong: An analysis of iterative prompting for reasoning problems. *arXiv preprint arXiv:2310.12397*, 2023.

[46] Kaya Stechly, Karthik Valmeekam, and Subbarao Kambhampati. On the self-verification limitations of large language models on reasoning and planning tasks. *arXiv preprint arXiv:2402.08115*, 2024.

[47] Weihang Su, Changyue Wang, Qingyao Ai, Yiran Hu, Zhijing Wu, Yujia Zhou, and Yiqun Liu. Unsupervised real-time hallucination detection based on the internal states of large language models. *arXiv preprint arXiv:2403.06448*, 2024.

[48] Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 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, 2019.

[49] Nikolaas Tinbergen. On aims and methods of ethology. *Zeitschrift für tierpsychologie*, 20(4):410–433, 1963.

[50] Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10014–10037, 2023.

[51] Gladys Tyen, Hassan Mansoor, Peter Chen, Tony Mak, and Victor Cărbune. Llms cannot find reasoning errors, but can correct them! *arXiv preprint arXiv:2311.08516*, 2023.

[52] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*, 2023.

[53] Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. Self-instruct: Aligning language models with self-generated instructions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13484–13508, 2023.

[54] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.

[55] Bingbing Wen, Jihan Yao, Shangbin Feng, Chenjun Xu, Yulia Tsvetkov, Bill Howe, and Lucy Lu Wang. Know your limits: A survey of abstention in large language models. *arXiv preprint arXiv:2407.18418*, 2024.

[56] Yixuan Weng, Minjun Zhu, Fei Xia, Bin Li, Shizhu He, Shengping Liu, Bin Sun, Kang Liu, and Jun Zhao. Large language models are better reasoners with self-verification. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023.

[57] Yuxi Xie, Kenji Kawaguchi, Yiran Zhao, James Xu Zhao, Min-Yen Kan, Junxian He, and Michael Xie. Self-evaluation guided beam search for reasoning. *Advances in Neural Information Processing Systems*, 36, 2024.

[58] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural Information Processing Systems*, 36, 2024.

[59] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations*, 2022.

[60] Zhangyue Yin, Qiushi Sun, Qipeng Guo, Jiawen Wu, Xipeng Qiu, and Xuan-Jing Huang. Do large language models know what they don't know? In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8653–8665, 2023.

[61] Tianhang Zhang, Lin Qiu, Qipeng Guo, Cheng Deng, Yue Zhang, Zheng Zhang, Chenghu Zhou, Xinbing Wang, and Luoyi Fu. Enhancing uncertainty-based hallucination detection with stronger focus. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023.

[62] Bingxi Zhao, Lin Geng Foo, Ping Hu, Christian Theobalt, Hossein Rahmani, and Jun Liu. Llm-based agentic reasoning frameworks: A survey from methods to scenarios. *arXiv preprint arXiv:2508.17692*, 2025.

A Training Details

Data Preparation To prepare the raw generated data in Section 3.1 for DM training, we down-sample data points with the `Reject` label while keeping all `Accept` data points constant for each dataset. We observe that generated data naturally consists of more rejections than acceptances of LLM answers because the condition for generation to continue is a rejection. As a result, examples of rejections are present in many turns for a question (e.g., an answer is repeatedly rejected), but acceptances are only present in the final turn. When down-sampling, we directly remove data points without keeping data points of different turns of conversations with the same original question (same ID) together. We make this choice because the past history of answers provided in the task content of data points already shows the full context and does not require all turns to be in the training data. Then, for each dataset, we randomly split the generated data into a train and dev set for DM training using an 80/20 split, keeping data points of different turns with same ID together to avoid data leakage. Last, we concatenate the prepared train sets and the dev sets of each dataset together. Table 2 displays the final size of the train and dev sets for Llama3 and GPT-4 experiments, as well as the relative size of each dataset.

Table 2: DM training and development dataset sizes across four QA benchmarks, with Llama3 and GPT-4 as the Reasoner.

Dataset	**Train**		**Dev**	
	Llama3	GPT-4	Llama3	GPT-4
CommonSenseQA	11416	11124	2855	2783
Winogrande	10969	10529	2744	2668
PIQA	7785	3654	1950	911
OpenBookQA	5496	4909	1365	1230
Overall	**35666**	**30216**	**8914**	**7592**

Training Setting Prior to training, we map `Accept` labels to 1, `Reject` labels to 0, and include a basic instruction to the DM. We fine-tune Flan-T5 with an Adam optimizer ($\eta = 3e^{-5}, \lambda = 1e^{-4}$) with 500 warmup steps. The DM is fine-tuned using the Adam optimizer with the loss function described in Equation 1. All experiments were conducted on two NVIDIA A6000 GPUs.

Maximum Turns While our ideal system could run for any number of necessary reasoning turns, our experiments are constrained by time and computational resources, and after a large amount of iterations, there is a higher possibility of an infinite state where the LLM never generates a correct answer. Consequently, we set a maximum number of turns of four for generation and five for inference, after which wrong answers are counted as abstention during inference.

B Prompts

Zero-Shot Baselines For zero-shot baselines, we designed three types of prompts: (a) *Standard QA Prompt*: a straightforward question-answering prompt; (b) *Abstain QA Prompt*: an extended prompt allowing the model to abstain from answering if none of the options are correct; and (c) *Self-Critic Prompt*: a prompt adapted from [21] where the model critiques its own response and iteratively improves it. Zero-shot predictions are generated by evaluating the first turn of our system, ignoring DM outputs. Abstain zero-shot predictions are run separately, with temperature 0.1. Note that the self-critic prompt directly from [21] may produce discrepancies, such as self-critic scores being lower than zero-shot scores using our system prompts. Self-critic baselines use temperature 0, following [21].

DRR Prompt Settings DRR employs two distinct prompt strategies. The first strategy uses a direct instruction-based approach. The first turn uses the *Standard QA Prompt* as a system prefix, mimicking zero-shot performance. Subsequent turns replace the standard QA prompt with a *Exploration Prompt*, instructing the LLM to explore new reasoning paths based on feedback such as "Wrong Answer! Try again." Temperature settings are 0.1 (Top-P 0.9) for the first turn to ensure consistent zero-shot responses and 0.6 (Top-P 0.7) for later turns to increase diversity. The second strategy adopts a gradual prompting approach. All turns use the *Standard QA Prompt* as a system prefix. The appended

context includes an environment message that emphasizes gradual improvement and encourages exploring new reasoning paths without explicitly labeling past answers as wrong. Temperature settings are 0.6 (Top-P 0.9) for all turns. For a given experiment, the same strategy is used for both the generation and inference steps. The first strategy is used for Llama models, while the second strategy is applied to GPT models.

Prompt Examples Examples of each prompt are displayed as following.

Standard QA Prompt (Llama3 Zero-Shot):

> You are a knowledgeable question-answering assistant, specializing in multiple-choice questions. Based on the question and the list of choices provided, select the best answer. Carefully evaluate each option before deciding. Provide your choice (e.g., 0, 1, 2, etc) along with a brief explanation of your reasoning.
>
> Respond only with the following format, nothing else: Answer: [Provide the answer here] Rationale: [Provide the rationale here]
>
> Do not include any additional text, headers, or explanations outside this format.

Standard QA Prompt (GPT-4 Zero-Shot):

> You are a highly knowledgeable assistant skilled in multi-step reasoning for multiple-choice question answering. Based on the question and the list of choices provided, select the best answer. Carefully evaluate each option before deciding. Provide your choice (e.g., 0, 1, 2, etc) along with a brief explanation of your reasoning. Respond only with the following format, nothing else: Answer: [Provide the answer here] Rationale: [Provide the rationale here]
>
> Do not include any additional text, headers, or explanations outside this format.

Abstain QA Prompt (Llama3/GPT-4 Zero-Shot):

> You are a highly knowledgeable assistant skilled in multi-step reasoning for multiple-choice question answering. Based on the question and the list of choices provided, select the best answer. Carefully evaluate each option before deciding. Provide your choice (e.g., 0, 1, 2, etc) along with a brief explanation of your reasoning. Respond only with the following format, nothing else: Answer: [Provide answer here or 'none of the above'] Rationale: [Provide the rationale here]
>
> Do not include any additional text, headers, or explanations outside this format.

Self-Critic QA Prompt (GPT-4 Self-Critic):

> Q: ... Choices: ...
> Explain your reasoning. You must choose only one option from above. Your final answer should be a single number (e.g., 0, 1, 2, etc), in the form (answer), at the end of your response.
> ...
> Review your previous answer and find problems with your answer.
> ...
> Based on the problems you found, improve your answer. You must choose only one option from above. Please reiterate your answer, with your final answer a single number (e.g., 0, 1, 2, etc), in the form (answer).

Exploration Prompt (Llama3 DRR):

> You are an expert assistant specializing in multiple-choice questions, dedicated to exploring multiple ways of thinking to provide accurate answers. Below, you will see an LLM's previous answer, including the choice it selected and its reasoning, followed by the feedback: 'Wrong answer! Try again.'
>
> Your task is to **think outside the box** and use a **completely different line of reasoning** to approach the question. Carefully reassess each option, explore alternative interpretations, and **avoid repeating the same ideas**. Focus on providing fresh insights and explain your reasoning in a distinct way.

Respond only with the following format, nothing else: Answer: [Provide answer here or 'none of the above'] Rationale: [Provide the rationale here]

Do not include any additional text, headers, or explanations outside this format.

Appended Environment Message to Context (Llama3 DRR):

Previous LLM Answer: Answer:...
Rationale:...
Wrong Answer! Try again.

Appended Environment Message to Context (GPT-4 DRR):

LLM Answer 1:...
Rationale 1:...
Env: The previous response was insufficient; explore a new line of reasoning to approach a more accurate answer.