



SIFT: Grounding LLM Reasoning in Contexts via Stickers

Anonymous ACL submission

Abstract

This paper identifies the misinterpretation of the context can be a significant issue during the reasoning process of large language models, spanning from smaller models like Llama3.2-3B-Instruct to cutting-edge ones like DeepSeek-R1. For example, in the phrase “10 dollars per kilo,” Llama3.2-3B-Instruct might not recognize that “per” means “for each,” leading to calculation errors. We introduce a novel, post-training approach called **Stick to the Facts (SIFT)** to tackle this. SIFT leverages increasing inference-time compute to ground LLM reasoning in contexts. At the core of SIFT lies the *Sticker*, which is generated by the model itself to explicitly emphasize the key information within the context. Given the curated Sticker, SIFT generates two predictions—one from the original query and one from the query augmented with the Sticker. If they differ, the Sticker is sequentially refined via *forward* optimization (to better align the extracted facts with the query) and *inverse* generation (to conform with the model’s inherent tendencies) for more faithful reasoning outcomes. Studies across diverse models (from 3B to 100B+) and benchmarks (e.g., GSM8K, MATH-500) reveal consistent performance improvements. Notably, SIFT improves the pass@1 accuracy of DeepSeek-R1 on AIME2024 from 78.33% to **85.67%**, establishing a new state-of-the-art in the open-source community. The code will be public after acceptance.

1 Introduction

Recent advancements in large language models (LLMs) (Dubey et al., 2024; Yang et al., 2024; Liu et al., 2024) have significantly advanced the field of natural language processing. Techniques including Chain-of-Thought (CoT) Prompting (Wei et al., 2022b; Kojima et al., 2022) and Self-Consistency (Wang et al., 2023b), as

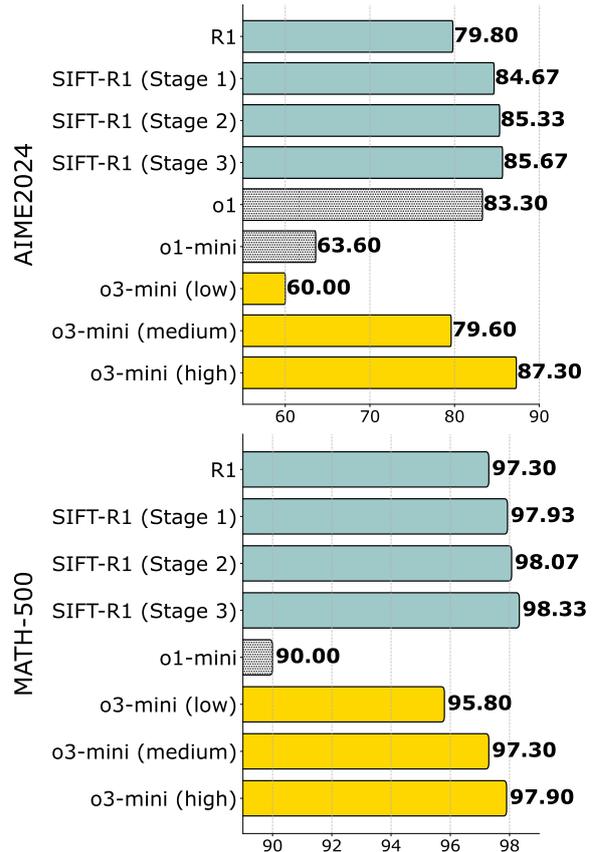


Figure 1: Applying SIFT to DeepSeek-R1 establishes state-of-the-art reasoning performance on both AIME2024 and MATH-500. (pass@1 accuracy)

well as reasoning-enhanced models, e.g., OpenAI-o1 (Jaech et al., 2024), DeepSeek-R1 (Guo et al., 2025), and KIMI-k1.5 (Team et al., 2025), have all contributed to improvements in multi-step reasoning for solving hard problems.

Recent discussions in the community suggest that advanced reasoning capabilities in LLMs mainly stem from two factors: (i) foundational knowledge acquisition through massive pretraining on diverse data (Dubey et al., 2024; Lin et al., 2025), and (ii) strategic refinement via post-training interventions like supervised fine-tuning

Query
Josh decides to try flipping a house. He buys a house for \$80,000 and then puts in \$50,000 in repairs. This increased the value of the house by 150%. How much profit did he make?

Sticker
<p>Conditions:</p> <ol style="list-style-type: none"> 1. Josh buys a house for \$80,000. 2. He spends \$50,000 on repairs. 3. The value of the house increases by 150%. <p>Question:</p> <p>What is the total profit Josh made from flipping the house?</p>

Figure 2: An example of a query and its Sticker.

(SFT) (Chung et al., 2022) or reinforcement learning (RL) (Guo et al., 2025), which optimize the model’s ability to select contextually relevant reasoning pathways. However, our studies reveal a critical lacuna in this framework: LLMs of varying sizes systematically misinterpret, overlook, or hallucinate key information in the query context—an emergent vulnerability we term *factual drift*. For example, Llama3.2-3B-Instruct (Dubey et al., 2024) might incorrectly interpret “per” as “total” instead of “for each” in the phrase “10 dollars per kilo,” leading to reasoning errors even with the logical steps being correct. As a result, while current research prioritizes *optimizing reasoning mechanisms* in LLMs (Zelikman et al., 2022, 2024; Wu et al., 2024; Zhang et al., 2024b), we argue equal attention should also be placed on *whether LLMs are reasoning about the correct problem*.

We note that advanced reasoning models, such as DeepSeek-R1 (Guo et al., 2025), can partially mitigate factual drift during its reasoning process via *self-verification*. For example, the model dynamically paraphrases critical constraints (e.g., converting “at least 3 days” to “minimum duration ≥ 72 hours”) to implicitly perform error-checking. This helps correct prior misunderstandings of the context and leads to better-aligned reasoning results. However, such self-verification operates as a stochastic safeguard rather than a systematic protocol—it is not guaranteed to be triggered in various reasoning scenarios. Namely, the risk of *factual drift* remains and it can be significant considering the results in Figure 1.

Inspired by that humans usually use sticky notes to externalize critical elements when handling complex tasks, we propose the **Stick to the Facts (SIFT)** method to explicitly ground LLM reasoning in contexts using Stickers generated by the model itself. SIFT is a post-training approach, leverag-

ing inference-time compute to improve generation quality yet without reliance on reward models as in Best-of-N (BoN) (Brown et al., 2024; Snell et al., 2024) and Monte-Carlo tree search (MCTS) (Qi et al., 2024; Zhang et al., 2025). Concretely, SIFT lets the target LLM summarize key facts within the input query, including essential *conditions* and the core *question*, into a structured *Sticker* (see Figure 2), and make two predictions based on the Sticker alone and the query augmented with the Sticker, respectively. If they differ, the Sticker is refined through bidirectional optimization—a *forward* one to better align the Sticker with the query and an *inverse* one to conform to the model’s reasoning preference—for more faithful reasoning.

Experiments demonstrate that SIFT can consistently improve the reasoning performance across various LLMs and benchmarks. Notably, for DeepSeek-R1 (Guo et al., 2025), SIFT achieves a 1.03% accuracy improvement over the vanilla CoT (97.3%) on MATH-500 (Lightman et al., 2023). Additionally, on AIME2024 (of America, 2024), it brings a significant 7.34% accuracy improvement (see Figure 1), establishing a new state-of-the-art in the open-source community. We also witness a striking performance improvement for small-to-medium-sized models including Llama3.2-3B-Instruct (Dubey et al., 2024), Llama3.1-8B-Instruct (Dubey et al., 2024), and Qwen2.5-7B-Instruct (Yang et al., 2024).

2 Related Work

Reasoning has long been a significant challenge for LLMs. Several approaches aim to improve the reasoning capabilities of LLMs. These methods can be broadly categorized into techniques that align reasoning through training, enhance reasoning through search and planning, or augment reasoning during inference.

Some approaches focus on aligning the reasoning path of LLMs through Supervised Fine-Tuning (SFT) or Reinforcement Learning (RL). STaR (Zelikman et al., 2022) enables the model to use reject sampling and learn from its mistakes by rationalizing its outputs, progressively enhancing its reasoning capabilities. Quiet-STaR (Zelikman et al., 2024) generates multiple rationales in parallel before each output token, thereby improving the model’s ability to predict subsequent tokens. V-STaR (Hosseini et al., 2024) employs a dual-system framework where the generator creates preference

pairs to train the verifier, which then scores the candidate solutions.

Additionally, a significant body of work aims to enhance model reasoning abilities through search and planning. Q* (Wang et al., 2024) formalizes multi-step reasoning as a Markov Decision Process (MDP) and uses the A* algorithm to guide the model in selecting the optimal next step. rStar (Qi et al., 2024) employs Monte Carlo Tree Search (MCTS) to enhance the model’s reasoning exploration and uses Mutual Verification to evaluate the reasoning paths. SR-MCTS (Zhang et al., 2024a) combines Self-Refinement and MCTS to iteratively improve and optimize newly discovered reasoning paths. MCTS-DPO (Xie et al., 2024) leverages MCTS to collect step-level preference data and uses Decision-Policy Optimization (DPO) to refine the model’s policy through multiple iterations. ReST-MCTS* (Zhang et al., 2025) takes a broader approach in evaluating reasoning paths, considering not only the correctness of the results but also the quality of the reasoning process, such as the shortest path and error-free intermediate steps. CoRe (Zhu et al., 2022) constructs a dual-system approach with System 1 for generation and System 2 for verification, training, and reasoning simultaneously to simulate human-like reasoning processes. AlphaMath (Chen et al., 2024) treats the output of the LLM as an action and integrates a value model and a policy model, iteratively training the model to enhance its reasoning capabilities.

There are also methods that focus on enhancing reasoning abilities during inference. Innovations in prompt engineering have contributed to advancements in reasoning capabilities. Chain-of-Thought (CoT) prompting (Wei et al., 2022a; Kojima et al., 2022) guides models in stepwise reasoning, such as by manually annotating natural language rationales or appending “Let’s think step by step” after questions. Auto-CoT (Zhang et al., 2022) clusters questions and uses zero-shot Chain-of-Thought to generate reasoning chains, which are then used as prompts to guide the model’s answers. ToT (Yao et al., 2023) removes the constraints of chain structures by incorporating tree structures and search algorithms, allowing models to explore widely during reasoning. The seminal Self-Consistency method (Wang et al., 2023a) aggregates answers through majority voting over multiple reasoning paths, while Madaan et al. (2024) introduces iterative self-correction via feedback loops.

However, these methods primarily focus on re-

(i) Factual drift occurs within query \Rightarrow Sticker.

Query
Carla is downloading a 200 GB file. Normally she can download 2 GB/minute, but 40% of the way through the download. Windows forces a restart to install updates, which takes 20 minutes. Then Carla has to restart the download from the beginning. How long does it take to download the file?

LLM

Incorrect Sticker: Key constraints neglected (underline above)

Conditions:
1. Carla is downloading a 200 GB file.
2. Normally she can download 2 GB/minute.
3. Windows forces a restart to install updates, which takes 20 minutes.
4. Then Carla has to restart the download from the beginning.
Question:
How long does it take to download the file?

(ii) Factual drift occurs within Sticker \Rightarrow prediction.

Query
(...) However, she has to choose between the boots and two pairs of high heels that together cost five dollars less than the boots (...) how many dollars are the boots?

LLM

Correct Sticker

Conditions:
(...) The two pairs of high heels together cost five dollars less than the boots. (...)
Question:
How many dollars are the boots?

LLM

Incorrect Prediction: Misinterpretation (underline above)

(...) The boots cost five dollars less than the two pairs of heels, (...)

Figure 3: Factual drift occurs during (i) Sticker generation and (ii) prediction generation from Sticker.

fining *how* models reason rather than ensuring that they address the *correct problem*. Our approach differs by prioritizing factual comprehension before answer generation, ensuring proper problem understanding.

3 Method

This section first presents the factual drift issue during LLM reasoning and then elaborates on the proposed Stick to the Facts (SIFT) approach.

3.1 Factual Drift in LLM Reasoning

We define *factual drift* as the phenomenon where the LLM reasoning fails due to misaligned comprehension of the query context rather than flawed reasoning logic. This occurs when LLMs neglect key constraints, misinterpret semantic relationships, or hallucinate non-existent conditions during reasoning procedures.

We show that factual drift can be a systematic failure mode of general LLM problem-solving processes beyond reasoning. Taking the task of applying Stickers to Llama3.2-3B-Instruct (Dubey et al.,

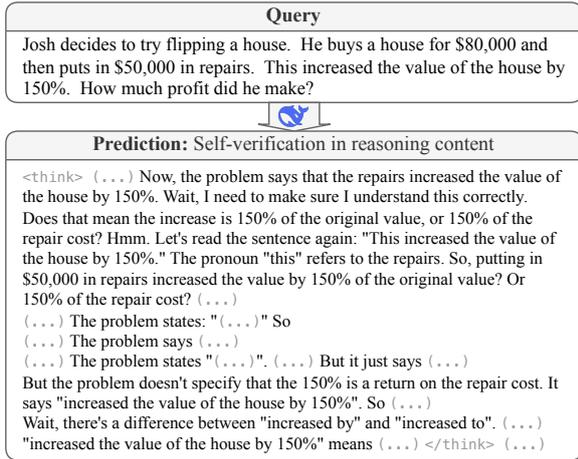


Figure 4: Self-verification occurs during DeepSeek-R1’s reasoning, where the model revisiting the query, focusing on key information, and paraphrasing it.

2024) on GSM8K test set (Cobbe et al., 2021) as an example, we curate Stickers with the model, based on which predictions are made. We observe extensive factual drift errors, with typical examples displayed in Figure 3. As shown, when mapping the query to Stickers, LLMs may neglect the original constraints. Moreover, even when the Sticker is correct, LLMs may still misunderstand it, especially when the question is complex or uses less familiar phrasing. The above observations also highlight that *more optimization mechanisms regarding the Sticker are required to make it 1) more aligned with the query and 2) able to be easily understood and leveraged by the target LLM.*

Self-verification of Advanced Reasoning Models. We note that, for advanced models like DeepSeek-R1 (Guo et al., 2025), the reasoning process sometimes involves *self-verification*—revisiting the original problem, focusing on key information, and paraphrasing it. As illustrated in Figure 4, DeepSeek-R1 often states, “Let’s read the sentence again: ...” or “Wait, the problem states: ...” as part of its thought process, helping to deepen its understanding of the context or self-correct.

The excellent performance of such advanced reasoning models underscores the efficacy of mitigating factual drift to make the model better respect the context. Nevertheless, this self-verification functions more as a stochastic safeguard than a systematic protocol—it may not always be activated across different reasoning scenarios. Consequently, the risk of factual drift persists. We consequently develop the novel SIFT framework to address this.

Algorithm 1: LLM reasoning with SIFT

```

Input :Query  $Q$ 
Output :Final result of  $Q$ 

 $S_1 \leftarrow \text{SG}(Q)$ ; // Sticker generation
 $P_1 \leftarrow \text{CP}(Q, S_1)$ ;
if  $P_1 \not\sim$  then
  return  $P_1$ ; // Exit if consensus
else
  // Forward
   $S_2 \leftarrow \text{FO}(Q, S_1)$ ,  $P_2 \leftarrow \text{CP}(Q, S_2)$ ;
  if  $P_2 \not\sim$  then
    return  $P_2$ 
  else
    // Inverse
     $S_3 \leftarrow \text{FO}(Q, \text{IG}(P_2, S_2))$ ;
     $P_3 \leftarrow \text{CP}(Q, S_3)$ ;
    return  $P_3$  if  $P_3 \not\sim$  else  $\text{LLM}(Q)$ 
  end
end

```

Algorithm 2: Consensus Prediction (CP)

```

Input :Query  $Q$ , Sticker  $S$ 
Output :Prediction from  $Q$  &  $S$ , or  $\sim$  (unequal)

 $P_S \leftarrow \text{LLM}(S)$ ; // Sticker-only
 $P_{Q,S} \leftarrow \text{LLM}(Q, S)$ ; // Query+Sticker
if  $\text{EQUIVALENT}(P_S, P_{Q,S})$  then
  // Consensus validation
  return  $P_{Q,S}$ 
else
  return  $\sim$ 
end

```

3.2 Stick to the Facts (SIFT)

SIFT includes four core operations (see Figure 5): (i) Sticker Generation (SG), which extracts the Sticker from the original query; (ii) Consensus Prediction (CP), which validates the alignment between predictions from the Sticker and the query augmented with the Sticker; (iii) Forward Optimization (FO), which refines the Sticker to improve its alignment with the facts in the query; (iv) Inverse Generation (IG), which generates the Sticker based on the prediction inversely.

The full procedure of SIFT is shown in Algorithm 1 with the details of Consensus Prediction in Algorithm 2. All prompts used can be found in Appendix A. We explain some rationales below.

Consensus Prediction: Beyond Answer Aggregation. Traditional self-consistency methods sample diverse reasoning paths to aggregate answers (Wang et al., 2023a), focusing on *how* models reason. In contrast, our Consensus Prediction verifies *whether models reason about the same problem* with dual representations: (i) the *Sticker-Only* one, which forces the model to solve the problem using only the key conditions and the core

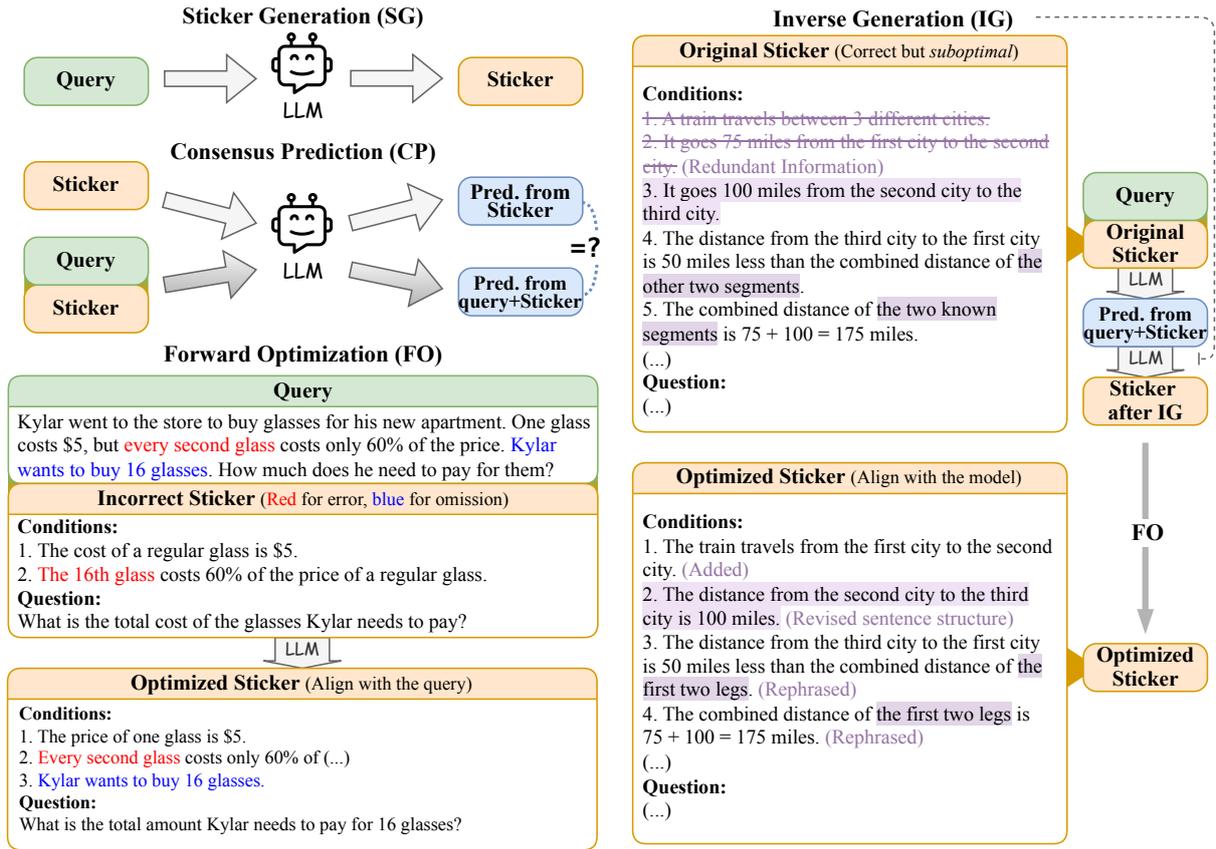


Figure 5: Four core operations in SIFT: (1) Sticker Generation (SG), (2) Consensus Prediction (CP), (3) Forward Optimization (FO), (4) Inverse Generation (IG).

question, and (ii) the *Query+Sticker* one, which provides richer contexts. This way, the model explores *semantic invariance* rather than sampling diversity when reasoning about the answers.

CP does not require sampling and operates with greedy decoding by default. However, it remains compatible with stochastic sampling, as demonstrated in Table 1. Besides, the CP operates only based on the current Sticker, preventing contamination from historical reasoning traces. As illustrated in Algorithm 2, consensus between representations acts as a factual invariant—a necessary (though not sufficient) condition for correctness. This design intentionally avoids conflating factual grounding with reasoning quality assessment.

Forward Optimization: Anchoring Stickers to Source Semantics. As discussed in Section 3.1, the SG process can also inevitably suffer from factual drift, where the original constraints are misrepresented or misunderstood. To address this, we combine the generated Sticker with the query to produce a refined Sticker. For example, it can correct misinterpretations, such as changing “the 16th glass” to “every second glass” in Figure 5.

Inverse Generation: Aligning Stickers to Model Reasoning Preference. It is frequently observed that for LLM reasoning, contexts with the same semantics but different presentations can yield distinct results. This implies that, after doing FO, it can be beneficial to further refine the Sticker based on the LLM’s reasoning process. Given this insight, we use the LLM to inversely infer a new Sticker given the model prediction. We further invoke FO once again to the new Sticker to avoid factual drift. This step makes the Sticker respect the internal reasoning preferences of the model for representing facts, arranging conditions, or structuring questions (see Figure 5). This also helps the model recognize the difference between the Sticker from IG and the original question, to enable the model to capture overlooked information and generate a more comprehensive Sticker.

4 Experiments

In this section, we first validate the effectiveness and generalization of SIFT (Section 4.1). Next, we explore several variants (Section 4.2 & 4.3).

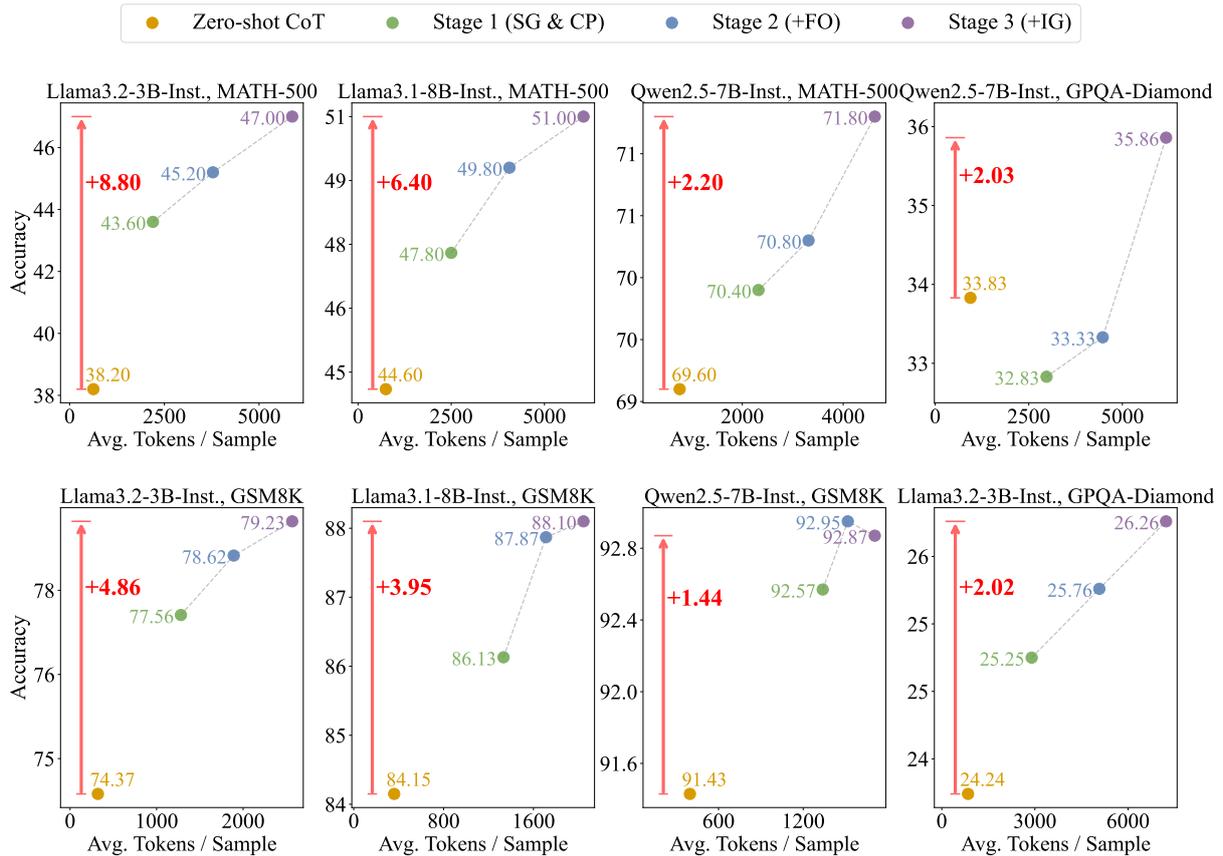


Figure 6: Comparison of SIFT and traditional Zero-shot CoT across multiple models and datasets. We divide SIFT into three stages: Stage 1 only uses SG & CP, while Stage 2 and Stage 3 optimize the Sticker through forward (+FO) and inverse (+IG) direction, respectively. The bidirectional arrows in the figure highlight the performance gap between Zero-shot CoT and the complete SIFT (i.e., Stage 3). We see that in nearly all scenarios, SIFT leads to a significant performance improvement.

Finally, we include ablation studies to gain further insights into our approach (Section 4.4).

4.1 Enhancing LLM Reasoning with SIFT

Models & Datasets. We test SIFT on a diverse set of state-of-the-art LLMs, including Llama3.2-3B-Instruct (Dubey et al., 2024), Llama3.1-8B-Instruct (Dubey et al., 2024), Qwen2.5-7B-Instruct (Yang et al., 2024), and DeepSeek-R1 (Guo et al., 2025). These models cover a range of sizes, architectures (Mixture-of-Experts (MoE) vs. dense), and reasoning capabilities. We select well-established reasoning benchmarks, including GSM8K (Cobbe et al., 2021), MATH-500 (Lightman et al., 2023), GPQA-Diamond (Rein et al., 2023), and AIME2024 (of America, 2024).

Test Protocol. To isolate the effect of SIFT from the influence of sampling, all tests are conducted using greedy decoding, except for DeepSeek-R1. Because the default settings of the used Volcengine API (temperature=1.0, top-p=0.7) cannot be mod-

ified, the SIFT on DeepSeek-R1 is based on sampling. Specifically, for DeepSeek-R1 on MATH-500, we perform 3 sampling runs and report average results. For AIME2024, due to its small size, we perform 10 sampling runs and report the average. Additionally, we divide the entire SIFT process into three stages: (i) Stage 1: Only SG and CP are used. (ii) Stage 2: Building upon Stage 1, FO is used to optimize the Sticker. (iii) Stage 3: The complete process outlined in Algorithm 1. The accuracy after each stage is measured: If the CP results are not aligned (\sim), the model’s direct answer to the query is used instead. All evaluations are performed on OpenCompass (Contributors, 2023).

Main Results. The results are shown in Figures 1 and 6. As observed, SIFT consistently delivers robust and significant performance improvements compared to traditional Zero-shot CoT across all settings. From a methodological perspective, as the stages increase—i.e., with the forward and inverse optimization of Sticker—the average num-

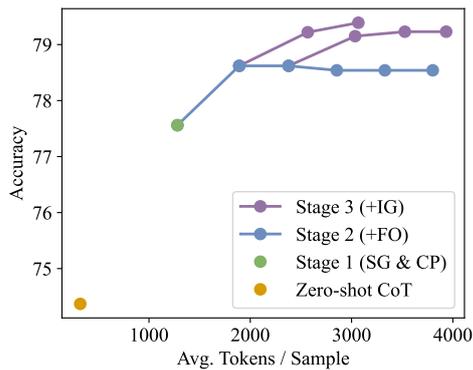


Figure 7: Iterative optimization results for SIFT. The performance improves as the number of tokens per sample increases across different stages. Significant gains are observed in the first repeats of Stage 2 and Stage 3.

ber of tokens used per sample rises, and accuracy shows an upward trend as well. From a model standpoint, SIFT demonstrates notable effectiveness across various scales (ranging from several billion to hundreds of billions of parameters), architectures (both dense and MoE), and paradigms (traditional and reasoning models). Particularly noteworthy is its significant impact on DeepSeek-R1. For instance, on MATH-500, it achieves a 1.03% absolute accuracy improvement over an already exceptionally high baseline of 97.3%. On AIME2024, it also brings a substantial absolute accuracy increase of 7.34%. These results indicate that even for advanced reasoning models like DeepSeek-R1, sticking to the facts remains crucial for optimal performance.

4.2 Iterative Optimization

In this section, we explore whether the Sticker can be continually optimized in SIFT.

Setup. We test with Llama3.2-3B-Instruct (Dubey et al., 2024) on the GSM8K dataset (Cobbe et al., 2021). Specifically, we conduct multiple optimization repeats for Stage 2 and Stage 3. The other settings are the same as in Section 4.1.

Results. The experimental results are shown in Figure 7. We observe that SIFT shows a test-time scaling, with the performance improving as the average number of tokens per sample increases. For Stage 2, the saturation is rapid, but adding Stage 3 can result in an additional, noticeable performance boost. Nevertheless, the most significant gains are observed at the first repeat. One possible explanation is that extracting the optimal Sticker for GSM8K is relatively easy. In more complex conditions, however, extracting a good Sticker may be

Consistency Dimension	Stage 1	Stage 2	Stage 3
Greedy	77.56	78.62	79.23
(i) Sticker	78.85	79.65	80.29
(ii) Prediction	85.37	86.20	86.28
(iii) SIFT	-	-	87.25

Table 1: Performance comparison of different consistency integration strategies for SIFT across multiple stages. The results show that integrating SIFT with Self-Consistency (Wang et al., 2023a) leads to significant performance improvements, with SIFT-Consistency achieving the highest accuracy boost.

harder, requiring more repeats to achieve optima. Additionally, since we use a training-free approach for SIFT, a model trained to exclusively optimize Sticker could lead to better iterative results.

4.3 Sample Augmentation

In this section, we explore the use of Self-Consistency (SC) (Wang et al., 2023a) to enhance SIFT, demonstrating how SIFT and SC can be effectively coupled together.

Specifically, SIFT and SC can be integrated in three ways: (i) **Sticker-Consistency**: Multiple Sticker samples are drawn, and consistency is applied to the predictions generated by each Sticker or by the query combined with each Sticker. (ii) **Prediction-Consistency**: Consistency is applied separately to predictions generated using *Sticker* alone and those generated with *Query + Sticker*, considering their respective samples. (iii) **SIFT-Consistency**: End-to-end sampling is conducted across the entire SIFT to ensure consistency. We test Llama3.2-3B-Instruct (Dubey et al., 2024) on GSM8K (Cobbe et al., 2021) with a temperature of 0.6, a top-p of 0.9, and 10 sampling iterations.

The results of these configurations are presented in Table 1. It is observed that our method can be combined with SC to achieve better performance. Specifically, Integrating SIFT across all dimensions results in performance improvements. Notably, SIFT-Consistency provides the most significant boost, demonstrating that the simplest sampling method—end-to-end—can lead to substantial performance gains for SIFT.

4.4 Ablation

Evolution of Consensus Across Optimization Stages. The efficacy of SIFT hinges on improving agreement between predictions derived from

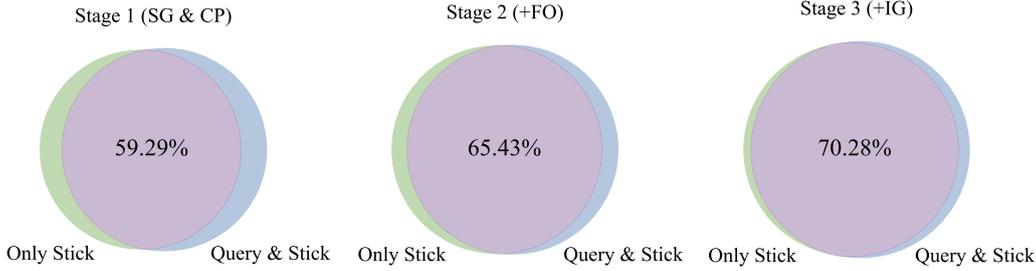


Figure 8: Venn diagrams illustrating the accuracy of predictions obtained from the “Only Sticker” and “Query & Sticker” representations at each stage. The percentages represent the accuracy where both methods correctly predict the same outcomes. From Stage 1 to Stage 2, the accuracy increases by 6.14%, and from Stage 2 to Stage 3, it increases by 4.85%. The results show the significant impact of Forward Optimization (FO) and Inverse Generation (IG) in improving prediction alignment from the two representations.

Model	Stage 1	Stage 2	Stage 3	Stage 3 from Stage 1
Llama	77.56	78.62	79.23	74.07
Qwen	92.57	92.95	92.87	90.90

Table 2: Performance comparison of Llama3.2-3B-Instruct and Qwen2.5-7B-Instruct on GSM8K, with and without Stage 2. The results show a performance drop when skipping directly from Stage 1 to Stage 3.

Strategy	Accuracy
$P_{Q,S}$ if $P_{Q,S}=P_S$ else P_Q	77.56
P_S if $P_S=P_Q$ else $P_{Q,S}$	77.02
P_Q if $P_Q=P_{Q,S}$ else P_S	76.04

Table 3: Performance comparison of various CP strategies. Here, P_Q , P_S , and $P_{Q,S}$ represent the predictions generated from query, Sticker, and query augmented with Sticker, respectively. The first row of the table represents the strategy used in SIFT, which is shown to be the optimal approach.

430 *Sticker-only* and *Query + Sticker* representations
 431 through iterative refinement. To quantify this
 432 alignment, We select Llama3.2-3B-Instruct (Dubey
 433 et al., 2024) on the GSM8K dataset (Cobbe et al.,
 434 2021). We plot the accuracy of predictions ob-
 435 tained using “Only Sticker” and “Query & Sticker”
 436 after each stage, visualized in the Venn diagram
 437 in Figure 8. As shown, both FO and IG signif-
 438 icantly improve the alignment of the predictions
 439 from the two representations. Specifically, the ac-
 440 curacy where both methods correctly predict the
 441 same outcomes increased by 6.14% from Stage 1
 442 to Stage 2, and by an additional 4.85% from Stage
 443 2 to Stage 3.

444 **FO Required Before Adding IG.** We investigate
 445 whether it is possible to skip directly from Stage
 446 1 to Stage 3. We select Llama3.2-3B-Instruct and
 447 Qwen2.5-7B-Instruct on GSM8K. All settings re-
 448 main the same as in Section 4.1, except for skipping
 449 directly to Stage 3 after Stage 1. The results are
 450 shown in Table 2. As observed, skipping Stage
 451 2 leads to a significant performance drop. This
 452 indicates that during the initial optimization of
 453 Sticker, FO is essential to align Sticker with the
 454 query, followed by aligning it with model cogni-
 455 tion. This is consistent with our experience, where

456 the effectiveness of Sticker depends primarily on
 457 its correctness—ensuring no *factual drift*—before
 458 considering its alignment with the model.

459 **Optimal Consensus Prediction Strategy.** CP pro-
 460 cess, our strategy involves comparing predictions
 461 from *Sticker* and *query + Sticker*. If the predictions
 462 are consistent, we adopt the prediction from Query
 463 + Sticker; otherwise, we use the prediction directly
 464 from *query*. We validate this as the optimal strat-
 465 egy. Several alternative strategies were evaluated
 466 using Stage 1 results of Llama3.2-3B-Instruct on
 467 the GSM8K dataset, as shown in Table 3. The re-
 468 sults demonstrate that our CP strategy is effective,
 469 aligning with the prior analysis in Section 3.2.

470 5 Conclusion

471 This study presents Stick to the Facts (SIFT), a
 472 training-free framework that anchors LLM reason-
 473 ing to contextual facts through iterative self-
 474 refinement. This approach enhances the reliability
 475 of LLM reasoning, providing a practical solution
 476 for factually grounded reasoning without the need
 477 for additional data or training.

478 Limitations

479 This work focuses on the training-free setting. In
480 the future, SIFT could be internalized into small
481 LLMs through dedicated training, enabling more
482 efficient on-device reasoning. Separately, SIFT can
483 be applied to reduce the output token length of rea-
484 soning models, improving computational efficiency
485 without compromising accuracy. Additionally, In-
486 verse Generation in SIFT offers new inspiration for
487 data generation in inverse synthesis tasks. Further
488 studies are needed to generalize its effectiveness
489 across a wider range of tasks.

490 References

491 Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald
492 Clark, Quoc V Le, Christopher Ré, and Azalia Mirho-
493 seini. 2024. Large language monkeys: Scaling infer-
494 ence compute with repeated sampling. *arXiv preprint*
495 *arXiv:2407.21787*.

496 Guoxin Chen, Minpeng Liao, Chengxi Li, and Kai Fan.
497 2024. Alphamath almost zero: process supervision
498 without process. *arXiv preprint arXiv:2405.03553*.

499 Hyung Won Chung, Le Hou, Shayne Longpre, Barret
500 Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi
501 Wang, Mostafa Dehghani, Siddhartha Brahma, Al-
502 bert Webson, Shixiang Shane Gu, Zhuyun Dai,
503 Mirac Suzgun, Xinyun Chen, Aakanksha Chowdh-
504 ery, Alex Castro-Ros, Marie Pellat, Kevin Robinson,
505 Dasha Valter, Sharan Narang, Gaurav Mishra, Adams
506 Yu, Vincent Zhao, Yanping Huang, Andrew Dai,
507 Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Ja-
508 cob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le,
509 and Jason Wei. 2022. *Scaling instruction-finetuned*
510 *language models*. *Preprint*, arXiv:2210.11416.

511 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian,
512 Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias
513 Plappert, Jerry Tworek, Jacob Hilton, Reiichiro
514 Nakano, Christopher Hesse, and John Schulman.
515 2021. Training verifiers to solve math word prob-
516 lems. *arXiv preprint arXiv:2110.14168*.

517 OpenCompass Contributors. 2023. Opencompass:
518 A universal evaluation platform for foundation
519 models. [https://github.com/open-compass/](https://github.com/open-compass/opencompass)
520 [opencompass](https://github.com/open-compass/opencompass).

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

526 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song,
527 Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma,
528 Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: In-
529 centivizing reasoning capability in llms via reinforce-
530 ment learning. *arXiv preprint arXiv:2501.12948*.

Arian Hosseini, Xingdi Yuan, Nikolay Malkin, Aaron
Courville, Alessandro Sordoni, and Rishabh Agar-
wal. 2024. V-star: Training verifiers for self-taught
reasoners. *arXiv preprint arXiv:2402.06457*.

Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richard-
son, Ahmed El-Kishky, Aiden Low, Alec Helyar,
Aleksander Madry, Alex Beutel, Alex Carney, et al.
2024. Openai o1 system card. *arXiv preprint*
arXiv:2412.16720.

Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yu-
taka Matsuo, and Yusuke Iwasawa. 2022. Large lan-
guage models are zero-shot reasoners. *Advances in*
neural information processing systems, 35:22199–
22213.

Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri
Edwards, Bowen Baker, Teddy Lee, Jan Leike,
John Schulman, Ilya Sutskever, and Karl Cobbe.
2023. Let’s verify step by step. *arXiv preprint*
arXiv:2305.20050.

Zhenghao Lin, Zhibin Gou, Yeyun Gong, Xiao Liu,
Yelong Shen, Ruochen Xu, Chen Lin, Yujiu Yang,
Jian Jiao, Nan Duan, and Weizhu Chen. 2025. *Rho-*
1: Not all tokens are what you need. *Preprint*,
arXiv:2404.07965.

Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang,
Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi
Deng, Chenyu Zhang, Chong Ruan, et al. 2024.
Deepseek-v3 technical report. *arXiv preprint*
arXiv:2412.19437.

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

Mathematical Association of America. 2024. *American*
invitational mathematics examination - aime 2024.
Accessed: 2024-02-11.

Zhenting Qi, Mingyuan Ma, Jiahang Xu, Li Lina Zhang,
Fan Yang, and Mao Yang. 2024. Mutual reasoning
makes smaller llms stronger problem-solvers. *arXiv*
preprint arXiv:2408.06195.

David Rein, Betty Li Hou, Asa Cooper Stickland, Jack-
son Petty, Richard Yuanzhe Pang, Julien Dirani, Ju-
lian Michael, and Samuel R Bowman. 2023. Gpqa: A
graduate-level google-proof q&a benchmark. *arXiv*
preprint arXiv:2311.12022.

Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Ku-
mar. 2024. Scaling llm test-time compute optimally
can be more effective than scaling model parameters.
arXiv preprint arXiv:2408.03314.

Kimi Team, Angang Du, Bofei Gao, Bowei Xing,
Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun
Xiao, Chenzhuang Du, Chonghua Liao, et al. 2025.
Kimi k1. 5: Scaling reinforcement learning with llms.
arXiv preprint arXiv:2501.12599.

587	Chaojie Wang, Yanchen Deng, Zhiyi Lyu, Liang Zeng, Jujie He, Shuicheng Yan, and Bo An. 2024. Q*: Improving multi-step reasoning for llms with deliberative planning. <i>arXiv preprint arXiv:2406.14283</i> .	643
588		644
589		645
590		646
591	Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023a. Self-consistency improves chain of thought reasoning in language models. <i>The Eleventh International Conference on Learning Representations</i> .	647
592		
593		
594		
595		
596		
597	Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023b. Self-consistency improves chain of thought reasoning in language models. In <i>The Eleventh International Conference on Learning Representations</i> .	648
598		649
599		650
600		651
601		652
602		
603	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022a. Chain-of-thought prompting elicits reasoning in large language models . In <i>Advances in Neural Information Processing Systems</i> , volume 35, pages 24824–24837. Curran Associates, Inc.	653
604		654
605		655
606		656
607		
608		
609		
610	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022b. Chain-of-thought prompting elicits reasoning in large language models. <i>Advances in neural information processing systems</i> , 35:24824–24837.	657
611		658
612		659
613		660
614		
615	Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, and Yiming Yang. 2024. Inference scaling laws: An empirical analysis of compute-optimal inference for problem-solving with language models . <i>Preprint</i> , arXiv:2408.00724.	661
616		662
617		663
618		664
619		665
620	Yuxi Xie, Anirudh Goyal, Wenyue Zheng, Min-Yen Kan, Timothy P Lillicrap, Kenji Kawaguchi, and Michael Shieh. 2024. Monte carlo tree search boosts reasoning via iterative preference learning. <i>arXiv preprint arXiv:2405.00451</i> .	
621		
622		
623		
624		
625	An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024. Qwen2. 5 technical report. <i>arXiv preprint arXiv:2412.15115</i> .	
626		
627		
628		
629	Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. <i>Advances in neural information processing systems</i> , 36:11809–11822.	
630		
631		
632		
633		
634	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. <i>arXiv preprint arXiv:2403.09629</i> .	
635		
636		
637		
638		
639	Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Goodman. 2022. Star: Bootstrapping reasoning with reasoning. <i>Advances in Neural Information Processing Systems</i> , 35:15476–15488.	
640		
641		
642		

666
667
668
669

A Prompting for SIFT

In this section, we present the complete prompt formats used in the SIFT process (see Figures 9 to 12 for details).

Query ⇒ Prediction
{Query} Please reason step by step, and put your final answer within \boxed{}
Sticker ⇒ Prediction
{Sticker} Please reason step by step, and put your final answer within \boxed{}
Query + Sticker ⇒ Prediction
{Query} {Sticker} Please reason step by step, and put your final answer within \boxed{}

Figure 9: Prompt format for generating predictions.

Query ⇒ Sticker
Extract fundamental elements from the following query using atomic decomposition methodology.
Requirements: 1. Conditions: Clearly list all the given information. Write each condition on a separate line, numbered sequentially. 2. Question: Summarize what is being asked in one clear sentence. Remove all known conditions.
Output Format: ` **Conditions:** 1. [Condition 1] 2. [Condition 2] ...(add more conditions as needed) **Question:** [Clearly state what is being asked.]`
Example: Query:(...) Expected Output:(...)
Query to Process: ` {question}`
Please provide your output strictly following the output format without other unnecessary words.

Figure 10: Prompt format for generating a Sticker from the query.

Query + Sticker ⇒ Sticker
Given a query and a candidate abstract (which includes conditions and a question), output an optimized abstract.
Requirements: 1. Definitions of Conditions and Question: * Conditions: Clearly list all the given information. Write each condition on a separate line, numbered sequentially. * Question: Summarize what is being asked in one clear sentence. Remove all known conditions. 2. Focus of Optimization: Compare the Original Query with the candidate Abstract. Identify and fix: * Missing/incorrect/redundant conditions * Imprecise question phrasing * Mathematical/logical inconsistencies * Output format error
Output Format: ` ..` **Conditions:** 1. [optimized Condition 1] 2. [optimized Condition 2] ...(add more conditions as needed) **Question:** [Optimized question phrasing. Clearly state what is being asked.]`
Some Examples:(...)
Input to Process: ` Original Query: {question} Candidate Abstract: {abstract}`
Please provide your output strictly following the output format without other unnecessary words.

Figure 11: Prompt format for forward optimization of the Sticker.

Prediction ⇒ Sticker
<p>Given the prediction provided below, reverse-engineer the abstract that led to it. The abstract should include both the conditions and the question.</p> <hr/> <p>Abstract Format:</p> <pre> \ **Conditions:** 1. [Condition 1] 2. [Condition 2] ...(add more conditions as needed) **Question:** [Clearly state what is being asked.] \ </pre> <hr/> <p>Requirements:</p> <ol style="list-style-type: none"> 1. Conditions: <ul style="list-style-type: none"> - Clearly list all the given information. - Write each condition on a separate line, numbered sequentially. - EACH CONDITION MUST BE ATOMIC AND INDIVISIBLE (i.e., it cannot be divided into two sub-conditions). - DO NOT INCLUDE ANY PART OF THE REASONING PROCESS!!! 2. Question: <ul style="list-style-type: none"> - Summarize what is being asked in one clear sentence. - Remove all known conditions. <hr/> <p>Example:</p> <p>Prediction:(...)</p> <p>Expected Output:(...)</p> <hr/> <p>Prediction to Process:</p> <pre> \ {prediction} </pre> <hr/> <p>Please provide your output strictly following the ABSTRACT FORMAT without other unnecessary words.</p>

Figure 12: Prompt format for generating a Sticker inversely from the prediction.