

000 WHERE DO REASONING MODELS MAKE A DIFFER- 001 ENCE? FOLLOW THE REASONING LEADER FOR EF- 002 FICIENT DECODING 003

004 **Anonymous authors**
 005
 006

007 Paper under double-blind review
 008
 009

010 ABSTRACT 011

012 Large reasoning models (LRMs) achieve strong reasoning performance by emitting
 013 long chains of thought. Yet, these verbose traces slow down inference and often
 014 drift into unnecessary detail, known as the overthinking phenomenon. To better
 015 understand LRM’s behavior, we systematically analyze the token-level misalign-
 016 ment between reasoning and non-reasoning models. While it is expected that their
 017 primary difference lies in the stylistic “thinking cues”, LRM’s uniquely exhibit
 018 two pivotal, previously under-explored phenomena: a *Global Misalignment Re-*
 019 *bound*, where their divergence from non-reasoning models persists or even grows
 020 as response length increases, and more critically, a *Local Misalignment Dimin-*
 021 *ish*, where the misalignment concentrates at the “thinking cues” each sentence
 022 starts with but rapidly declines in the remaining of the sentence. Motivated by the
 023 *Local Misalignment Diminish*, we propose *FoReaL-Decoding*, a collaborative fast-
 024 slow thinking decoding method for cost-quality trade-off. In FoReaL-Decoding,
 025 a Leading model leads the first few tokens for each sentence, and then a weaker
 026 draft model completes the following tokens to the end of each sentence. FoReaL-
 027 Decoding adopts a stochastic gate to smoothly interpolate between the small and
 028 the large model. On four popular math-reasoning benchmarks (AIME24, GPQA-
 029 Diamond, MATH500, AMC23), FoReaL-Decoding reduces theoretical FLOPs by
 030 30 - 50% and trims CoT length by up to 40%, while preserving 86 - 100% of model
 031 performance. These results establish FoReaL-Decoding as a simple, plug-and-play
 032 route to controllable cost-quality trade-offs in reasoning-centric tasks.
 033

034 1 INTRODUCTION 035

036 Reasoning has become a pivotal capability of large language models (LLMs), driving rapid progress
 037 in mathematical problem solving, code generation, and commonsense question answering (Huang &
 038 Chang, 2023; Ahn et al., 2024; Wang et al., 2024b; 2025b). Contemporary Large Reasoning Models
 039 (LRMs) such as OpenAI’s GPT-o1 (OpenAI, 2024) and the open-source DeepSeek-R1 (DeepSeek-AI
 040 et al., 2025) demonstrate this trend by producing explicit long chains of thought (CoT) (Wei et al.,
 041 2023) that markedly improve performance on challenging tasks in mathematics (Xiong et al., 2025;
 042 Xia et al., 2025b), programming (Liu et al., 2024a), and other complex domains. These deeper, longer,
 043 and more precise reasoning trajectories represent the advanced “slow-thinking” patterns (Kahneman,
 044 2011; Li et al., 2024d; 2025b). Although these slow-thinking LRM’s showcase impressive reasoning
 045 skills, communities are increasingly concerned about the efficiency and fidelity of their often-lengthy
 046 chains of thought, a phenomenon known as “overthinking” (Chen et al., 2025c; Fan et al., 2025),
 047 where excessive computational resources are allocated for simple problems with minimal benefit.

048 To alleviate overthinking and improve efficiency, a series of methods has been proposed¹. Most
 049 of these, however, require further post-training or manipulate the LRM’s distribution itself, adding
 050 complexity or computational overhead. Motivated by Speculative Decoding (Leviathan et al., 2023)
 051 and the distinctions between fast and slow thinking, we ask: **Where do reasoning models make a**
 052 **difference?** To answer this, we first seek to pinpoint what truly differentiates strong reasoning models
 053 from standard instruction-following LLMs at the token level. For instruction-following models,

¹The detailed related works are shown in Appendix C.

054
055
056
057
058
059
060
061
062
063
064
065
066
067
068
069
070
071
072
073
074
075
076
077
078
079
080
081
082
083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099
100
101
102
103
104
105
106
107

Problem

A regular hexagon can be divided into six equilateral triangles. If the perimeter of one of the triangles is 21 inches, what is the perimeter, in inches, of the regular hexagon?

Response

Okay, so I have this problem about a regular hexagon and **some equilateral triangles**. Let me try to understand what it's asking. It says a regular hexagon can be divided into six equilateral triangles. Hmm, I remember that a regular hexagon has all sides equal and all angles equal. If you divide it into six equilateral triangles, that must mean each triangle has **sides equal** to the sides of the hexagon. The problem gives me the perimeter of one of these triangles as 21 inches. I need to find the perimeter of the regular hexagon. Alright, let's break this down step by step. First, let me recall that the perimeter of a shape is the **sum of all its sides**. So, if each triangle is equilateral, all three sides are equal. That means **each side of the triangle is the same length**. If the perimeter is 21 inches, I can find the length of one side by ... Wait, let me make sure I got that right. If the hexagon is divided into six equilateral triangles, **each triangle has sides equal to the sides of the hexagon**. So, **each triangle's side is the same as the hexagon's side**. That makes sense because if you connect the center of the ...

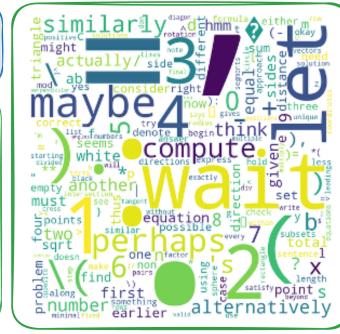


Figure 1: **Left:** An example comparing the token distribution alignment between *DeepSeek-R1-Distill-Qwen-32B* and *Qwen2.5-1.5B-Instruct*, qualitatively showing that the misaligned tokens (colored in red) are related to thinking patterns, and probably appear at the start of sentences. **Right:** The WordCloud of the misaligned tokens calculated on a mix of math datasets, quantitatively showing the high-frequency misaligned tokens like “*wait*”, “*perhaps*”, “*maybe*”, “*let*”, and “*alternatively*”.

LIMA (Zhou et al., 2023) proposes the “superficial alignment” hypothesis, in which it shows that most of the knowledge has been learned in the pretraining and only a small amount of data is needed for alignment. Lin et al. (2023) further verifies this hypothesis from token-level analysis between the base model and the aligned model.

Leveraging the diagnostic framework of (Lin et al., 2023), our systematic analysis of misalignment across various model types (large reasoning, small reasoning, instruction-following, and pretrained model) reveals critical insights. We observe a “superficial alignment” phenomenon similar to (Lin et al., 2023), where misaligned tokens are predominantly stylistic (e.g., “*Wait*”, “*Let me check*”) rather than content-specific, often related to explicit thinking patterns. More strikingly, while previous work showed that misalignment between instruction-following and base models diminishes with longer context, we find this does not hold for reasoning models. Instead, we identify a ***Global Misalignment Rebound***, where overall misalignment between reasoning and non-reasoning models can slightly grow with response length, suggesting that increasing the length cannot reduce the misalignment. This indicates that the reasoning abilities are *not* as superficial as instruction-following. Crucially, despite this global trend, we uncover a corresponding ***Local Misalignment Diminish*** phenomenon: most token misalignments occur at the *beginning of each sentence*, then rapidly decrease until the next sentence starts. These findings reveal a novel *periodical, sentence-level misalignment diminishing pattern* unique to LRM, driven by thinking-pattern indicators concentrated at sentence openings, shedding light on a better understanding of token-level divergences of these two types of models.

Based on this core insight that the reasoning pattern of LRM is often front-loaded in each sentence, we hypothesize that strategic, limited intervention by a strong LRM can guide a weaker model, balancing reasoning quality with efficiency. To this end, we propose ***Follow the Reasoning Leader (FoReAL-Decoding)***, an efficient collaborative decoding method. In FoReAL-Decoding, a strong Leading model generates the initial few tokens of each sentence (capturing the potentially misaligned “thinking cues”), after which a weaker Draft model completes the sentence. To further mitigate potential overthinking from the Leading model (e.g., endlessly generating “*Wait*”), we introduce a stochastic binary gate that controls whether the Leading model intervenes. These two control knobs, lead token count and lead probability, allow FoReAL-Decoding to smoothly interpolate between the Draft and Leading models, offering strong controllability over the cost-quality spectrum.

Contributions. In summary, our primary contributions can be illustrated as follows:

- We conduct a systematic token-level analysis comparing LRM with non-reasoning models, identifying two pivotal, under-explored phenomena: (1) ***Global Misalignment Rebound***, where the token distribution of LRM diverges from that of non-reasoning models and their gap even increases with longer responses; (2) ***Local Misalignment Diminish***, where LRM only make noticeable difference on generating stylistic “thinking-patterns” at the very beginning of each sentence. But such divergence from non-reasoning models rapidly drops on subsequent tokens within the sentence. This periodical *sentence-level misalignment diminishing pattern* has not been explored previously. These two discoveries significantly advance the understanding of LRM.

108 • Leveraging these analytical insights (particularly the *Local Misalignment Diminish*), we propose
 109 **FoReal-Decoding**, a training-free, collaborative algorithm that mixes the strength of a “slow-
 110 thinking” LRM (as Leading model) with the efficiency of a “fast-thinking”, weaker model (as
 111 draft model). FoReal-Decoding is designed to be plug-and-play, offering strong controllability to
 112 balance the cost and quality under diverse budgets of tokens.
 113 • Experimental results on several reasoning-heavy math tasks (AIME24, GPQA-Diamond,
 114 MATH500, AMC23) demonstrate that FoReal-Decoding reduces FLOPS by 30-55% and CoT
 115 length by up to 40%, while preserving 86-100% of the leading model’s performance, effectively
 116 mitigating “overthinking”.

118 2 TOKEN DISTRIBUTIONS OF REASONING VS. NON-REASONING MODELS

120 Large-scale reasoning models (LRMs) often outperform smaller instruction-tuned models on complex
 121 reasoning-heavy tasks, yet how their generation behavior differs from instruction models within
 122 the same model family remains unclear. (Lin et al., 2023) proposes an analytical method through
 123 the lens of token-distribution shifts and finds that alignments between instruction-following and
 124 base pretrained models are often superficial. This phenomenon is supported by nearly identical
 125 decoded tokens in the majority of token positions under the same input contexts, with distribution
 126 shifts occurring mainly with stylistic tokens like discourse markers. However, the critical question
 127 remains: “*Does this superficial alignment finding on instruction-following LLMs still hold for today’s*
 128 *capable LRM*s?” Thus, our work systematically investigates token misalignment across various
 129 model combinations involving LRMs.

130 **Experimental Setup & Metric.** In this analysis, we utilize *DeepSeek-R1-Distill-Qwen-32B* as
 131 the targeting LRM, which we denote as the Leading model $P_L(\cdot)$. The corresponding small models,
 132 within the same family, that are used for comparison are noted as the Draft models $P_D(\cdot)$. The
 133 Draft models can be (i) the pretrained base model (*Qwen2.5-1.5B*), (ii) the instruction-following
 134 model (*Qwen2.5-1.5B-Instruct*), or (iii) the small reasoning model (*DeepSeek-R1-Distill-Qwen-1.5B*)
 135 in our analysis and method. For a user query q , the output response generated greedily from the
 136 Leading model can be denoted as $y = \{y_1, \dots, y_T\}$, where T represents the length of the response.
 137 This response serves as the target for calculating the token distribution for the Draft model. At each
 138 position t , the context for predicting this token can be presented as $c_t = \langle q; y_{<t} \rangle$, where $\langle; \rangle$
 139 represents the concatenation operation.

140 In the analysis, the aligned positions are defined as those token steps where the Draft model, when
 141 conditioned on the Leading model’s history, would greedily generate exactly the same token as the
 142 Leading model, which means that *the two models have the same most probable behavior under the*
 143 *same context, indicating the alignment*.

144 Suppose \mathcal{V} is the vocabulary for next-token prediction, then the aligned token indices are:

$$145 \quad \mathcal{A} = \left\{ t \in \{1, \dots, T\} : \arg \max_{y \in \mathcal{V}} P_D(y | c_t) = \arg \max_{y \in \mathcal{V}} P_L(y | c_t) \right\}, \quad (1)$$

147 which collects exactly those positions where the Draft model’s top-1 prediction matches the Leading
 148 model’s emitted token under the shared causal context c_t . Thus, the aligned and misaligned tokens
 149 can be defined:

$$150 \quad \mathbf{y}_{\mathcal{A}} = \{y_t | t \in \mathcal{A}\} \quad \mathbf{y}_{\mathcal{A}^c} = \{y_t | t \notin \mathcal{A}\} \quad (2)$$

152 **Qualitative Analysis on Misaligned Tokens.** Figure 1 (left) shows a qualitative example (truncated)
 153 from MATH500, comparing the token distribution alignment between *DeepSeek-R1-Distill-Qwen-32B*
 154 as the Leading model and *Qwen2.5-1.5B-Instruct* as the Draft model. The shown response y is
 155 generated by the Leading model, the aligned tokens $\mathbf{y}_{\mathcal{A}}$ are colored in blue, and misaligned tokens
 156 $\mathbf{y}_{\mathcal{A}^c}$ are colored in red. Through the example, it can be intuitively perceived that the misaligned
 157 tokens are mostly stylistic tokens related to thinking patterns, and the beginning of each sentence²
 158 has a larger probability of being misaligned. To further quantitatively investigate what exactly these
 159 misaligned tokens are, we extract all the misaligned tokens from the mix of AIME24, AMC23,
 160 GPQA, and MATH datasets, count their frequencies, and generate the corresponding WordCloud
 161 shown in Figure 1 (right). From the WordCloud, it is observed that most of the high-frequency

²Sentences are defined as tokens separated by a period, question mark, exclamation mark, or newline symbol.

misaligned tokens are related to thinking patterns of LRM, like “*wait*”, “*perhaps*”, “*maybe*”, “*let*”, and “*alternatively*”, which shows a similar but different superficial phenomenon than previous instruction-following LLMs: While misalignment in both types of models is primarily stylistic rather than content-based, those in LRM are distinctively characterized by tokens reflecting their overt reasoning or self-correction patterns. Thus, our qualitative exploration reveals that LRM misalignment is characterized by stylistic “thinking cues” concentrated at sentence beginnings, prompting a more detailed quantitative analysis of their underlying distribution patterns.

Global Misalignment Rebound. Existing analysis on token distribution shifts between instruct and base models has identified that such shifts will gradually diminish over time during the decoding process due to the more comprehensive context given, as shown in Figure 2 (upper, red line). In the figure, the y-axis represents the average misalignment rate at each token position, while the x-axis represents the token position within the whole response (upper panel) or sentence (lower panel). As shown, the red line, representing misalignment between the instruct model and base model, decreases and remains at a low rate. This implies that providing longer context can gradually compensate for the misalignment between instruct and base models.

However, this response-level misalignment diminishing phenomenon does not strictly hold for LRM. As illustrated in Figure 2 (upper), lines corresponding to LRM as the Leading model exhibit different behaviors. When the Draft models are instruct (blue line) or base (orange line) models, the misalignment rates initially decrease dramatically to around 0.2, then rebound and persist around 0.3. In contrast, the green line, representing misalignment between large and small reasoning models (which belong to the same family and are trained on similar data), shows consistently low misalignment from the beginning, indicating a distinct trend. We term the observed persistent or rebounding divergence between LRM and non-reasoning models the ***Global Misalignment Rebound*** phenomenon. This phenomenon, characteristic of LRM comparisons with non-reasoning models, is mainly caused by LRM continuously generating thinking patterns at the beginning of sentences, partly to prevent premature conclusion of the generation process. This finding demonstrates that merely extending context length is insufficient to resolve the misalignment between reasoning and non-reasoning models, indicating that reasoning capability is *not* as superficial as instruction-following.

Local Misalignment Diminish. It is uncommon that a longer context does not benefit the alignment. Thus, to further understand this behavior, we conduct the sentence-level analysis by calculating the token misalignment rate at each sentence-level position. In the response, sentences can be separated by periods, question marks, exclamation marks, and the newline symbol. Specifically, for any position x , we first collect every sentence that is at least x tokens long. Mark the x -th token in each of those sentences as 1 if it is misaligned and 0 if it is aligned. The average of these 0-1 indicators across all selected sentences is the misalignment rate for position x .

As shown in Figure 1 (lower), for the red line, there is no obvious misalignment decrease that can be observed. It means that between the instruct and the base model, the misalignment occurs relatively evenly across the whole sentence. On the contrary, for LRM-involved model combinations, the blue, orange, and green lines, the misalignment rates drop dramatically at the first several tokens,

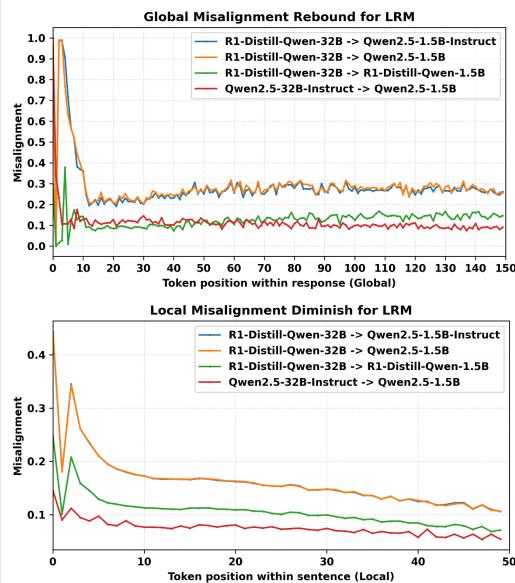


Figure 2: **Top:** Response-level misalignment changes with response length. **Bottom:** sentence-level misalignment changes with response length. The y-axis represents the average misalignment rate at each token position, the x-axis represents the token position within the whole response or sentence. We reveal the novel Global Misalignment Rebound and Local Misalignment Diminish phenomenon that only occurs on current LRM, shown as the blue, orange, and green lines of the upper figure. This phenomenon does not hold for the previous alignment between the instruction-following and base models, shown in the red line.

4

e.g., from 0.4 to 0.15, and then keep diminishing, indicating a totally different behavior. Thus, we term this phenomenon the ***Local Misalignment Diminish*** phenomenon for reasoning models. These findings reveal a novel periodical, sentence-level misalignment diminish pattern unique to LRM, driven by thinking-pattern indicators concentrated at sentence openings, shedding light on a better understanding of token-level divergences of these two types of models.

Findings. From this section, several key findings can be concluded:

- LRM misalignment with non-reasoning models, while largely superficial and characterized by stylistic “thinking cues”, uniquely exhibits a ***Global Misalignment Rebound***. Unlike instruct models that increasingly align with more context, token divergence at the response level can persist or even grow, underscoring deeper, ingrained differences in their generative behavior.
- LRM distinctively display a ***Local Misalignment Diminish***. This manifests as a novel, periodical sentence-level pattern where high misalignment, driven by “thinking cues” concentrated at sentence beginnings, rapidly decreases as the sentence progresses. This predictable intra-sentence dynamic is a crucial insight for developing LRM-guided decoding and understanding LRM patterns.

3 FOREAL-DECODING

Motivated by the above token divergence analysis, we propose a collaborative fast-slow thinking decoding method for cost-quality Trade-off, ***Follow the Reasoning Leader (FoReal-Decoding)***, a plug-and-play training-free method that mixes the strength of a slow but highly capable large reasoning model with the speed of a small model. The central idea is to let the strong, large (*Leading*) model lead at the beginning of sentences, and allow the weaker, small (*Draft*) model to complete the rest of the tokens. This decoding algorithm is of strong controllability, which can smoothly transfer into the Leading model only or downgrade to the Draft model only, by controlling the probability and the number of tokens to lead³.

Preliminaries. The two control knobs that govern the trade-off between cost and quality: *Required lead count* $n \in \mathbb{N}$: the minimum number of tokens the Leading model generates before yielding control to the Draft model. *Lead probability* $p \in [0, 1]$: probability that a sentence is led by the Leading model. When $p = 0$, the decoding system degenerates to pure Draft model decoding; when $p = 1$ and n exceeds the sentence length, it transfers to Leading model decoding. Intermediate settings form a *continuity* of compute–accuracy trade-offs.

In addition, let $t \in \mathbb{N}$ represent the global token index in the response, and $s \in \mathbb{N}$ represent the sentence index. $g_s \sim \text{Bernoulli}(p)$ represents the sentence-level gate to decide what model to start the sentence s : the sentence will be led by the Leading model if $g_s = 1$. τ_s represents the global position of the first token in s . $s(t) = \max\{s : \tau_s \leq t\}$ is the function that maps the token t to the sentence index that t belongs to. $\lambda_t = t - \tau_{s(t)} + 1$ is the local position of token t within its sentence.

Intra-Sentence Lead Within a sentence s , the generation of each token at position t is governed by the token-level policy,

$$\pi_t = \begin{cases} L & g_{s(t)} = 1 \wedge [\lambda_t \leq n \vee t < H_{s(t)}^{\text{hit}}], \\ D & \text{otherwise.} \end{cases} \quad (3)$$

$g_{s(t)} = 1$ represents this sentence $s(t)$ should be led by the Leading model, decided by the gate. L and D represent the Leading model and Draft model, respectively. $\lambda_t \leq n$ represents the index of this token within this sentence that is smaller than the required lead count n , thus should be generated by the Leading model. H_s^{hit} is the first token index within s where the top-1 token generated by the Draft model matches that of the Leading model for k consecutive steps:

$$H_s^{\text{hit}} = \min\{t : s(t) = s, \lambda_t > n, h_t = k\}, \quad (4)$$

where h_t represents the number of consecutive hits within the max sliding window of k :

$$h_t = \sum_{i=0}^{k-1} \delta_{t-i}, \quad \delta_t = \mathbf{1}\left\{\arg \max_{y \in \mathcal{V}} P_D(\cdot | c_t) = \arg \max_{y \in \mathcal{V}} P_L(\cdot | c_t)\right\} \quad (5)$$

³The detailed pseudo code is shown in Appendix A.

Put it simply, for each sentence, if the Bernoulli gate decides to let P_L lead the sentence with the probability p , P_L will generate the first n tokens. Then, P_D begins the generation process as well, with the purpose of measuring the alignment between the two models. When the top-1 predictions of these two models aligned with each other for k times, the generation process is yielded to P_D , otherwise, P_L generates the whole sentence. On the contrary, if the gate decides not to let P_L lead, then the whole sentence will be completely generated by P_D .

Sentence-level likelihood. For sentence s with token span $Y_s = (y_{\tau_s}, \dots, y_{\tau_{s+1}-1})$ and length L_s , the conditional likelihood under FoReaL-Decoding is:

$$P_{\text{CoL}}(Y_s \mid g_s) = \prod_{i=0}^{L_s-1} P_{\pi_{\tau_s+i}}(y_{\tau_s+i} \mid c_{\tau_s+i}), \quad (6)$$

Whenever $\pi_t = L$, the factor draws its probability from the distribution P_L of the Leader model; otherwise from the Draft model of distribution P_D .

Inter-Sentence Transfer Upon encountering a sentence boundary at the token t , i.e., the sentence is complete, we execute the inter-sentence update by resetting the hit counter and resampling the gate for the next sentence.

$$s \leftarrow s + 1, \quad g_s \sim \text{Bernoulli}(p), \quad h_t \leftarrow 0 \quad (7)$$

4 EXPERIMENTS

4.1 IMPLEMENTATION DETAILS

Models, Datasets, and Setup. To assess the effects of FoReaL-Decoding, extensive experiments are conducted for different model combinations in the Qwen2.5 family, including reasoning models like *R1-Distill-Qwen-32B* (DeepSeek-AI et al., 2025), *R1-Distill-Qwen-1.5B* (DeepSeek-AI et al., 2025), non-reasoning instruct models like *Qwen2.5-7B-Instruct* (Team, 2024), *Qwen2.5-1.5B-Instruct* (Team, 2024), and base models like *Qwen2.5-1.5B* (Team, 2024). To cover a wide scope of potential trade-offs, we utilize the reasoning models as the Leading models, while any of the above types as the Draft models. Moreover, our extensive experiments on the recently released *Qwen3* (Team, 2025) series further verify the generalizability of our method. We evaluate our method on relatively hard, reasoning-heavy math datasets, including AIME2024 (AI-MO, 2024a), GPQA-Diamond (Rein et al., 2024), AMC23 (AI-MO, 2024b), and MATH500 (Lightman et al., 2023). All experiments were conducted on NVIDIA A100 GPUs (80G), utilizing the Huggingface Transformers package. During the generation, we follow the recommended generation configuration from R1-Distill models as `temperature=0.6, top_p=0.95, top_k=40` for all the experiments. During the generation, we always let the Leading model generate the first paragraph, and we fix the required hits for generation transfer as $k = 5$ for all the experiments.

4.2 MAIN RESULTS

Table 1 presents the comparisons between accuracy and efficiency (TFLOPs) of FoReaL-Decoding on commonly used reasoning-heavy math problem tasks. We provide some different configurations as controls to show the wide trade-off scopes of our method. We also present the reported results of the concurrent work, Speculative Thinking (Yang et al., 2025), for better comparison. The accuracies on each line are compared with the Draft model, and the TFLOPs are compared with the Leading models: better values are colored in green, otherwise red. We utilize the theoretically estimated TFLOPs⁴ as the efficiency measurement since it takes the generation length into account, different from the estimated speed. In the main comparison, we focus on 3 collaborative settings. Across four benchmarks, FoReaL-Decoding cuts inference cost by 30 – 55% relative to Leader-only decoding while retaining 86 - 100% of its accuracy. The detailed statistics, including response length and leading ratios on AIME24, can be found in Table 2 for better understanding.

R1-Distill-Qwen-32B for Leading, R1-Distill-Qwen-1.5B for Draft. This collaborative setting yields the highest accuracies for all of the math reasoning datasets. In this setting, the larger 32B reasoning model takes charge of the leading of the sentences, while the smaller 1.5B reasoning model needs to

⁴The estimated TFLOPs are calculated based on the detailed formulas presented in Appendix B.

324
325
326
327
328
329
Table 1: Comparisons of Accuracy and Efficiency (TFLOPs) of FoReAL-Decoding on commonly
used reasoning-heavy math problem tasks. To further show the wide trade-off scopes of our method,
we provide some different configurations as the control. The results of Speculative Thinking are the
reported results. The accuracies are better with higher (\uparrow) values, while the TFLOPs are better with
lower (\downarrow) values. The accuracies on each line are compared with the Draft model, and the TFLOPs
are compared with the Leading models: better values are colored in green, otherwise red.

330	331	Model	AIME24		GPQA-D		MATH500		AMC23																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																									
332	333	334	335	336	337	338	339	340	341	342	343	344	345	346	347	348	349	350	351	352	353	354	355	356	357	358	359	360	361	362	363	364	365	366	367	368	369	370	371	372	373	374	375	376	377	378	379	380	381	382	383	384	385	386	387	388	389	390	391	392	393	394	395	396	397	398	399	400	401	402	403	404	405	406	407	408	409	410	411	412	413	414	415	416	417	418	419	420	421	422	423	424	425	426	427	428	429	430	431	432	433	434	435	436	437	438	439	440	441	442	443	444	445	446	447	448	449	450	451	452	453	454	455	456	457	458	459	460	461	462	463	464	465	466	467	468	469	470	471	472	473	474	475	476	477	478	479	480	481	482	483	484	485	486	487	488	489	490	491	492	493	494	495	496	497	498	499	500	501	502	503	504	505	506	507	508	509	510	511	512	513	514	515	516	517	518	519	520	521	522	523	524	525	526	527	528	529	530	531	532	533	534	535	536	537	538	539	540	541	542	543	544	545	546	547	548	549	550	551	552	553	554	555	556	557	558	559	560	561	562	563	564	565	566	567	568	569	570	571	572	573	574	575	576	577	578	579	580	581	582	583	584	585	586	587	588	589	590	591	592	593	594	595	596	597	598	599	600	601	602	603	604	605	606	607	608	609	610	611	612	613	614	615	616	617	618	619	620	621	622	623	624	625	626	627	628	629	630	631	632	633	634	635	636	637	638	639	640	641	642	643	644	645	646	647	648	649	650	651	652	653	654	655	656	657	658	659	660	661	662	663	664	665	666	667	668	669	670	671	672	673	674	675	676	677	678	679	680	681	682	683	684	685	686	687	688	689	690	691	692	693	694	695	696	697	698	699	700	701	702	703	704	705	706	707	708	709	710	711	712	713	714	715	716	717	718	719	720	721	722	723	724	725	726	727	728	729	730	731	732	733	734	735	736	737	738	739	740	741	742	743	744	745	746	747	748	749	750	751	752	753	754	755	756	757	758	759	760	761	762	763	764	765	766	767	768	769	770	771	772	773	774	775	776	777	778	779	780	781	782	783	784	785	786	787	788	789	790	791	792	793	794	795	796	797	798	799	800	801	802	803	804	805	806	807	808	809	810	811	812	813	814	815	816	817	818	819	820	821	822	823	824	825	826	827	828	829	830	831	832	833	834	835	836	837	838	839	840	841	842	843	844	845	846	847	848	849	850	851	852	853	854	855	856	857	858	859	860	861	862	863	864	865	866	867	868	869	870	871	872	873	874	875	876	877	878	879	880	881	882	883	884	885	886	887	888	889	890	891	892	893	894	895	896	897	898	899	900	901	902	903	904	905	906	907	908	909	910	911	912	913	914	915	916	917	918	919	920	921	922	923	924	925	926	927	928	929	930	931	932	933	934	935	936	937	938	939	940	941	942	943	944	945	946	947	948	949	950	951	952	953	954	955	956	957	958	959	960	961	962	963	964	965	966	967	968	969	970	971	972	973	974	975	976	977	978	979	980	981	982	983	984	985	986	987	988	989	990	991	992	993	994	995	996	997	998	999	1000	1001	1002	1003	1004	1005	1006	1007	1008	1009	1010	1011	1012	1013	1014	1015	1016	1017	1018	1019	1020	1021	1022	1023	1024	1025	1026	1027	1028	1029	1030	1031	1032	1033	1034	1035	1036	1037	1038	1039	1040	1041	1042	1043	1044	1045	1046	1047	1048	1049	1050	1051	1052	1053	1054	1055	1056	1057	1058	1059	1060	1061	1062	1063	1064	1065	1066	1067	1068	1069	1070	1071	1072	1073	1074	1075	1076	1077	1078	1079	1080	1081	1082	1083	1084	1085	1086	1087	1088	1089	1090	1091	1092	1093	1094	1095	1096	1097	1098	1099	1100	1101	1102	1103	1104	1105	1106	1107	1108	1109	1110	1111	1112	1113	1114	1115	1116	1117	1118	1119	1120	1121	1122	1123	1124	1125	1126	1127	1128	1129	1130	1131	1132	1133	1134	1135	1136	1137	1138	1139	1140	1141	1142	1143	1144	1145	1146	1147	1148	1149	1150	1151	1152	1153	1154	1155	1156	1157	1158	1159	1160	1161	1162	1163	1164	1165	1166	1167	1168	1169	1170	1171	1172	1173	1174	1175	1176	1177	1178	1179	1180	1181	1182	1183	1184	1185	1186	1187	1188	1189	1190	1191	1192	1193	1194	1195	1196	1197	1198	1199	1200	1201	1202	1203	1204	1205	1206	1207	1208	1209	1210	1211	1212	1213	1214	1215	1216	1217	1218	1219	1220	1221	1222	1223	1224	1225	1226	1227	1228	1229	1230	1231	1232	1233	1234	1235	1236	1237	1238	1239	1240	1241	1242	1243	1244	1245	1246	1247	1248	1249	1250	1251	1252	1253	1254	1255	1256	1257	1258	1259	1260	1261	1262	1263	1264	1265	1266	1267	1268	1269	1270	1271	1272	1273	1274	1275	1276	1277	1278	1279	1280	1281	1282	1283	1284	1285	1286	1287	1288	1289	1290	1291	1292	1293	1294	1295	1296	1297	1298	1299	1300	1301	1302	1303	1304	1305	1306	1307	1308	1309	1310	1311	1312	1313	1314	1315	1316	1317	1318	1319	1320	1321	1322	1323	1324	1325	1326	1327	1328	1329	1330	1331	1332	1333	1334	1335	1336	1337	1338	1339	1340	1341	1342	1343	1344	1345	1346	1347	1348	1349	1350	1351	1352	1353	1354	1355	1356	1357	1358	1359	1360	1361	1362	1363	1364	1365	1366	1367	1368	1369	1370	1371	1372	1373	1374	1375	1376	1377	1378	1379	1380	1381	1382	1383	1384	1385	1386	1387	1388	1389	1390	1391	1392	1393	1394	1395	1396	1397	1398	1399	1400	1401	1402	1403	1404	1405	1406	1407	1408	1409	1410	1411	1412	1413	1414	1415	1416	1417	1418	1419	1420	1421	1422	1423	1424	1425	1426	1427	1428	1429	1430	1431	1432	1433	1434	1435	1436	1437	1438	1439	1440	1441	1442	1443	1444	1445	1446	1447	1448	1449	1450	1451	1452	1453	1454	1455	1456	1457	1458	1459	1460	1461	1462	1463	1464	1465	1466	1467	1468	1469	1470	1471	1472	1473	1474	1475	1476	1477	1478	1479	1480	1481	1482	1483	1484	1485	1486	1487	1488	1489	1490	1491	1492	1493	1494	1495	1496	1497	1498	1499	1500	1501	1502	1503	1504	1505	1506	1507	1508	1509	1510	1511	1512	1513	1514	1515	1516	1517	1518	1519	1520	1521	1522	1523	1524	1525	1526	1527	1528	1529	1530	1531	1532	1533	1534	1535	1536	1537	1538	1539	1540	1541	1542	1543	1544	1545	1546	1547	1548	1549	1550	1551	1552	1553	1554	1555	1556	1557	1558	1559	1560	1561	1562	1563	1564	1565	1566	1567	1568	1569	1570	1571	1572	1573	1574	1575	1576	1577	1578	1579	1580	1581	1582	1583	1584	1585	1586	1587	1588	1589	1590	1591	1592	1593	1594	1595	1596	1597	1598	1599	1600	1601	1602	1603	1604	1605	1606	1607	1608	1609	1610	1611	1612	1613	1614	1615	1616	1617	1618	1619	1620	1621	1622	1623	1624	1625	1626	1627	1628	1629	1630

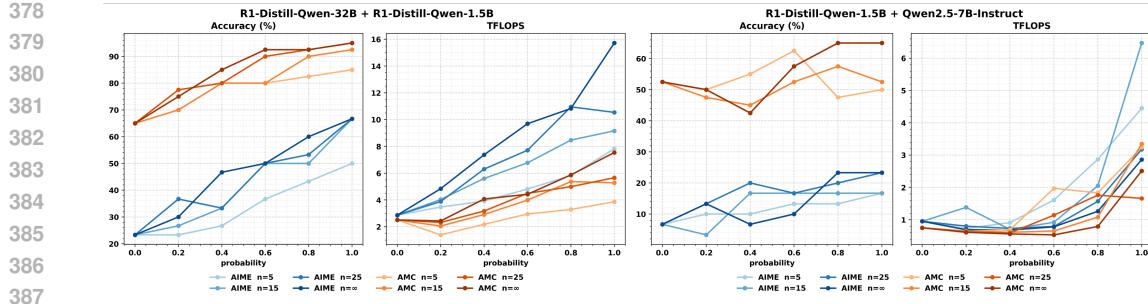


Figure 3: Effects of lead count and lead probability on AIME24 and AMC23 datasets, based on 2 collaborative configurations. FoReAL-Decoding provides a smooth cost-quality trade-off, making the transition from the weak Draft model to the strong Leading model smooth and controllable.

R1-Distill-Qwen-1.5B for Leading, Qwen2.5-7B-Instruct for Draft. Different from the above settings, in which a strong but large reasoning model is used as the Leading model, this setting considers a different and most efficient scenario, utilizing a small reasoning model for leading and a slightly larger instruct model for Draft. In this setting, the efficiencies are reduced to an extremely low level, even faster than directly utilizing the small reasoning models. As shown in Table 2, FoReAL-Decoding largely reduces the length required for the problem, thus largely reducing the computation required. On AIME24 and AMC23, our method reaches the same accuracy as the Leading model with similar or less computation. On GPQA, our method reaches an intermediate accuracy, since the abnormal situation where a non-reasoning model has better performance than the reasoning model.

4.3 EFFECTS OF LEAD COUNT AND LEAD PROBABILITY

Figure 3 sweeps the two hyperparameters that govern the controllability of FoReAL-Decoding, lead count n and lead probability p on AIME24 and AMC23 datasets, based on 2 collaborative configurations, *DeepSeek-R1-Distill-Qwen-32B + DeepSeek-R1-Distill-Qwen-1.5B* and *DeepSeek-R1-Distill-Qwen-1.5B + Qwen2.5-7B-Instruct*, representing the high-performance and high-efficiency settings, respectively. For each model combination, we run experiments on $n \in \{5, 15, 25, +\infty\}$, $p \in \{0, 0.2, 0.4, 0.6, 0.8, 1.0\}$. When $p = 0$, FoReAL-Decoding utilizes the Draft model only, and utilizes the Leading model only when $p = 1$ and $n = +\infty$. According to the figure, FoReAL-Decoding provides a smooth cost-quality trade-off, making the transition from the weak Draft model to the strong Leading model smooth and controllable. For any fixed n , increasing the probability p of the Leader intervention shifts the operating point up and to the right: accuracy rises while estimated TFLOPs grows almost linearly. The resulting curve is smooth, allowing practitioners to trade latency for quality by adjusting (n, p) . The jump from $n = 5$ to $n = 15$ yields large accuracy gains at a modest cost increase. Further enlarging the Leader count ($n \geq 25$) adds little accuracy yet inflates compute up a lot, indicating that sentence-level guidance already captures most of the benefit of slow reasoning.

4.4 TRADE-OFF CURVES

Figure 4 plots the trade-off curves between accuracy and estimated TFLOPs for every (n, p) configuration tested on AIME24 (left) and AMC23 (right), according to our experiment scopes on Qwen2.5 family. Blue markers correspond to our variants, red circles denote the corresponding LRM, and the dashed line is the empirically computed Pareto frontier. On both benchmarks, every LRM point is Pareto dominated: an alternative FoReAL-Decoding setting always achieves higher accuracy at lower cost. Moreover, we find that the frontier rises sharply between 0.5 and 2 estimated TFLOPs, as each additional estimated TFLOPs yields 10-15 percentage points of accuracy. However, beyond ≈ 5 estimated TFLOPs, the curve flattens; extra compute buys only marginal improvements up to the ceiling.

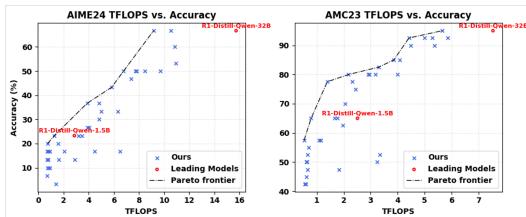


Figure 4: The trade-off curves between accuracy and estimated TFLOPs. Blue markers correspond to our variants, red circles denote the corresponding LRM, and the dashed line is the empirically computed Pareto frontier. On both benchmarks, every LRM point is Pareto dominated.

432

4.5 RESULTS ON QWEN3 FAMILIES

433 To further verify the effectiveness and generalizability of FoReaL-Decoding on other models, additional experiments are conducted on the Qwen3 series of models, including *Qwen3-32B*, *Qwen3-1.7B*, and *Qwen3-0.6B*, due to the various sizes of models provided in the family. Since Qwen3 family models have both reasoning and non-reasoning modes, we utilize both modes for these models and follow exactly the same generation configuration as our main experiments. The detailed experimental results are shown in Table 3. As shown in the table, FoReaL-Decoding shows promising performances on all the configurations. Most notably, in the *Qwen3-32B (reasoning)* + *Qwen3-0.6B (non-reasoning)* configuration, FoReaL-Decoding reaches a similar accuracy (66.6% to 73.3%) with only half of the estimated TFLOPs and with less response length. Compared with the utilizing Qwen2.5 family models using reasoning plus non-reasoning models, Qwen3 family models achieve much better performance, these results not only present the effectiveness and generalizability of FoReaL-Decoding, but also shows the potential of our method using stronger models.

456

4.6 COMPARED WITH SPECULATIVE DECODING

457 In this section, we compare and analyze FoReaL-Decoding with Speculative Decoding (SD) (Leviathan et al., 2023) on the Qwen3 family models. Technically, FoReaL-Decoding can be viewed as a variant of SD, but with a crucial difference in the interaction pattern between the strong and weak models. In standard SD, the weaker model first generates a fixed number of tokens m , after which the stronger model verifies and potentially corrects them. In contrast, FoReaL-Decoding has a stronger model that generates the initial tokens, those most likely to cause divergence, before handing off to the weaker model to complete the remainder of the sequence. As shown in the Table 3, FoReaL-Decoding achieves similar accuracy as SD but has lower estimated TFLOPs, which is mainly caused by the reduction on response lengths, i.e., alleviate the phenomenon of overthinking. As mentioned in the motivation, SD has the same distribution with the verifier model, promising the performance but also keeping the overthinking problem.

468 For qualitative comparison, FoReaL-Decoding is *particularly advantageous in scenarios where*
 469 *misaligned tokens are clustered together*, which is common in current reasoning models, as shown
 470 in our findings on token divergence. For instance, if there are n consecutive misaligned tokens,
 471 FoReaL-Decoding requires the strong model to generate only these n tokens. In contrast, SD with a
 472 draft length of m would require the weak model to generate m tokens, followed by verification and
 473 correction by the strong model, repeating this process until all n misaligned tokens are covered. This
 474 results in $n \times m$ tokens generated by the weak model and n tokens by the strong model, an $n \times m$
 475 increase in weak model computation compared to FoReaL-Decoding.

476

5 CONCLUSION

477 Our systematic token-level analysis comparing Large Reasoning Models (LRMs) with non-reasoning
 478 models has uncovered two pivotal, previously under-explored divergence phenomena: ***Global Mis-***
479 alignment Rebound and ***Local Misalignment Diminish***. This pattern reveals a novel, periodical

480 sentence-level pattern wherein LRM-specific stylistic “thinking cues” cause high token divergence at
 481 the very beginning of sentences, after which this misalignment rapidly decreases within the sentence.
 482 Building upon this insight, we proposed FoReaL-Decoding, a training-free, plug-and-play collabora-
 483 tive decoding algorithm. It allows a strong LRM to lead the crucial initial tokens of sentences
 484 (capturing these divergent “thinking cues”), while a lightweight Draft model efficiently completes the
 485 subsequent, more aligned portions. Our experiments demonstrate that FoReaL-Decoding achieves
 good cost-quality trade-off on reasoning-heavy math tasks.

433 Table 3: The detailed results of Qwen3 series models on AIME, including the comparison with Speculative Decoding, where m denotes the draft length. FoReaL-Decoding shows promising performance in this additional family.

Model		AIME24			
Method	Config	ACC (%)	Length	Ratio	TFLOPs
Base Models					
Qwen3-32B	—	76.6	13 275	—	15.75
Qwen3-1.7B	—	40.0	14 990	—	2.81
Qwen3-0.6B	—	13.3	15 839	—	1.14
Qwen3-32B (reasoning) + Qwen3-1.7B (reasoning)					
FoReaL-Decoding	$n=15, p=0.4$	60.0	14 840	0.272	7.20
FoReaL-Decoding	$n=15, p=0.6$	73.3	14 110	0.412	8.83
FoReaL-Decoding	$n=15, p=0.8$	73.3	15 081	0.536	11.43
Qwen3-32B (reasoning) + Qwen3-0.6B (reasoning)					
FoReaL-Decoding	$n=15, p=0.4$	36.7	17 782	0.281	7.44
FoReaL-Decoding	$n=15, p=0.6$	63.0	14 279	0.410	8.18
FoReaL-Decoding	$n=15, p=0.8$	60.0	15 478	0.560	11.01
Qwen3-32B (reasoning) + Qwen3-0.6B (non-reasoning)					
FoReaL-Decoding	$n=15, p=0.6$	60.0	8 762	0.472	5.24
FoReaL-Decoding	$n=15, p=0.8$	66.6	9 468	0.558	6.73
FoReaL-Decoding	$n=15, p=1.0$	73.3	10 267	0.673	8.32
Speculative Decoding: Qwen3-32B (reasoning) + Qwen3-0.6B (non-reasoning)					
Speculative Decoding	$m = 20$	73.3	15 374	0.427	10.25
Speculative Decoding	$m = 10$	73.3	14 672	0.403	9.02

486 THE USE OF LARGE LANGUAGE MODELS
487

488 In preparing this manuscript, we utilized large language models exclusively as writing assistants to
489 enhance clarity and presentation. The models were employed to suggest improvements to sentence
490 structure, readability, and stylistic consistency throughout the document. Following each set of
491 AI-generated suggestions, we carefully reviewed, evaluated, and selectively incorporated or modified
492 the proposed changes to ensure accuracy and appropriateness. Importantly, large language models
493 did not contribute to any aspect of research conceptualization, methodological design, experimental
494 implementation, data analysis, or interpretation of results.

495
496 REFERENCES
497

498 Pranjal Aggarwal and Sean Welleck. L1: Controlling how long a reasoning model thinks with
499 reinforcement learning, 2025. URL <https://arxiv.org/abs/2503.04697>.

500 Janice Ahn, Rishu Verma, Renze Lou, Di Liu, Rui Zhang, and Wenpeng Yin. Large language models
501 for mathematical reasoning: Progresses and challenges. In Neele Falk, Sara Papi, and Mike
502 Zhang (eds.), *Proceedings of the 18th Conference of the European Chapter of the Association
503 for Computational Linguistics: Student Research Workshop*, pp. 225–237, St. Julian’s, Malta,
504 March 2024. Association for Computational Linguistics. URL <https://aclanthology.org/2024.eacl-srw.17/>.

505 AI-MO. AIME 2022–2024 Validation Set, 2024a. URL <https://huggingface.co/datasets/AI-MO/aimo-validation-aime>.

506 AI-MO. AMC 12 2023 Integer-Answer Validation Set, 2024b. URL <https://huggingface.co/datasets/AI-MO/aimo-validation-amc>.

507 Maciej Besta, Julia Barth, Eric Schreiber, Ales Kubicek, Afonso Catarino, Robert Gerstenberger,
508 Piotr Nyczyk, Patrick Iff, Yueling Li, Sam Houlston, Tomasz Sternal, Marcin Copik, Grzegorz
509 Kwaśniewski, Jürgen Müller, Łukasz Flis, Hannes Eberhard, Hubert Niewiadomski, and Torsten
510 Hoefler. Reasoning language models: A blueprint, 2025. URL <https://arxiv.org/abs/2501.11223>.

511 Alexander Bukharin and Tuo Zhao. Data diversity matters for robust instruction tuning, 2023.

512 Tianle Cai, Yuhong Li, Zhengyang Geng, Hongwu Peng, Jason D Lee, Deming Chen, and Tri Dao.
513 Medusa: Simple llm inference acceleration framework with multiple decoding heads. *arXiv
514 preprint arXiv:2401.10774*, 2024.

515 Charlie Chen, Sebastian Borgeaud, Geoffrey Irving, Jean-Baptiste Lespiau, Laurent Sifre, and John
516 Jumper. Accelerating large language model decoding with speculative sampling. *arXiv preprint
517 arXiv:2302.01318*, 2023a.

518 Lichang Chen, Shiyang Li, Jun Yan, Hai Wang, Kalpa Gunaratna, Vikas Yadav, Zheng Tang, Vijay
519 Srinivasan, Tianyi Zhou, Heng Huang, and Hongxia Jin. Alpaganus: Training a better alpaca with
520 fewer data, 2023b.

521 Qiguang Chen, Libo Qin, Jinhao Liu, Dengyun Peng, Jiannan Guan, Peng Wang, Mengkang Hu,
522 Yuhang Zhou, Te Gao, and Wanxiang Che. Towards reasoning era: A survey of long chain-of-
523 thought for reasoning large language models, 2025a. URL <https://arxiv.org/abs/2503.09567>.

524 Runjin Chen, Gabriel Jacob Perin, Xuxi Chen, Xilun Chen, Yan Han, Nina ST Hirata, Junyuan Hong,
525 and Bhavya Kailkhura. Extracting and understanding the superficial knowledge in alignment.
526 *arXiv preprint arXiv:2502.04602*, 2025b.

527 Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuwei
528 Liu, Mengfei Zhou, Zhusong Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong Yu.
529 Do not think that much for $2+3=?$ on the overthinking of o1-like llms, 2025c. URL <https://arxiv.org/abs/2412.21187>.

540 Yushuo Chen, Tianyi Tang, Erge Xiang, Linjiang Li, Wayne Xin Zhao, Jing Wang, Yunpeng Chai,
 541 and Ji-Rong Wen. Towards coarse-to-fine evaluation of inference efficiency for large language
 542 models. *arXiv preprint arXiv:2404.11502*, 2024.

543

544 Alejandro Cuadron, Dacheng Li, Wenjie Ma, Xingyao Wang, Yichuan Wang, Siyuan Zhuang, Shu
 545 Liu, Luis Gaspar Schroeder, Tian Xia, Huanzhi Mao, Nicholas Thumiger, Aditya Desai, Ion Stoica,
 546 Ana Klimovic, Graham Neubig, and Joseph E. Gonzalez. The danger of overthinking: Examining
 547 the reasoning-action dilemma in agentic tasks, 2025. URL <https://arxiv.org/abs/2502.08235>.

548

549 Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu
 550 Yu, Qixin Xu, Weize Chen, Jiarui Yuan, Huayu Chen, Kaiyan Zhang, Xingtai Lv, Shuo Wang,
 551 Yuan Yao, Xu Han, Hao Peng, Yu Cheng, Zhiyuan Liu, Maosong Sun, Bowen Zhou, and Ning
 552 Ding. Process reinforcement through implicit rewards, 2025. URL <https://arxiv.org/abs/2502.01456>.

553

554 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu,
 555 Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, and etc. Deepseek-
 556 r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025. URL <https://arxiv.org/abs/2501.12948>.

557

558 Qianlong Du, Chengqing Zong, and Jiajun Zhang. Mods: Model-oriented data selection for instruction
 559 tuning, 2023.

560

561 Chenrui Fan, Ming Li, Lichao Sun, and Tianyi Zhou. Missing premise exacerbates overthinking: Are
 562 reasoning models losing critical thinking skill? *arXiv preprint arXiv:2504.06514*, 2025.

563

564 Tingxu Han, Zhenting Wang, Chunrong Fang, Shiyu Zhao, Shiqing Ma, and Zhenyu Chen. Token-
 565 budget-aware llm reasoning. *arXiv preprint arXiv:2412.18547*, 2024.

566

567 Xiaotian Han. Reproduce the inference-time scaling experiment, 2024. URL <https://ahxt.github.io/blog/2024-12-30-inference-time-scaling-exp/>.

568

569 Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong
 570 Tian. Training large language models to reason in a continuous latent space. *arXiv preprint
 571 arXiv:2412.06769*, 2024.

572

573 Jie Huang and Kevin Chen-Chuan Chang. Towards reasoning in large language models: A survey,
 574 2023. URL <https://arxiv.org/abs/2212.10403>.

575

576 Lifeng Jin, Baolin Peng, Linfeng Song, Haitao Mi, Ye Tian, and Dong Yu. Collaborative decoding of
 577 critical tokens for boosting factuality of large language models. *arXiv preprint arXiv:2402.17982*,
 578 2024.

579

580 Daniel Kahneman. *Thinking, fast and slow*. macmillan, 2011.

581

582 Abhinav Kumar, Jaechul Roh, Ali Naseh, Marzena Karpinska, Mohit Iyyer, Amir Houmansadr,
 583 and Eugene Bagdasarian. Overthink: Slowdown attacks on reasoning llms, 2025. URL <https://arxiv.org/abs/2502.02542>.

584

585 Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative
 586 decoding. In *International Conference on Machine Learning*, pp. 19274–19286. PMLR, 2023.

587

588 Jianwei Li and Jung-Eun Kim. Superficial safety alignment hypothesis. *arXiv preprint
 589 arXiv:2410.10862*, 2024.

590

591 Ming Li, Lichang Chen, Juhai Chen, Shuai He, and Tianyi Zhou. Reflection-tuning: Recycling data
 592 for better instruction-tuning. In *NeurIPS 2023 Workshop on Instruction Tuning and Instruction
 593 Following*, 2023.

594

595 Ming Li, Han Chen, Chenguang Wang, Dang Nguyen, Dianqi Li, and Tianyi Zhou. Ruler: Improving
 596 llm controllability by rule-based data recycling. *arXiv preprint arXiv:2406.15938*, 2024a.

594 Ming Li, Lichang Chen, Juhai Chen, Shuai He, Jiuxiang Gu, and Tianyi Zhou. Selective reflection-
 595 tuning: Student-selected data recycling for LLM instruction-tuning. In Lun-Wei Ku, Andre
 596 Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics ACL*
 597 2024, pp. 16189–16211, Bangkok, Thailand and virtual meeting, August 2024b. Association for
 598 Computational Linguistics. URL <https://aclanthology.org/2024.findings-acl-958>.

600 Ming Li, Pei Chen, Chenguang Wang, Hongyu Zhao, Yijun Liang, Yupeng Hou, Fuxiao Liu, and
 601 Tianyi Zhou. Mosaic-it: Free compositional data augmentation improves instruction tuning. *arXiv*
 602 preprint [arXiv:2405.13326](https://arxiv.org/abs/2405.13326), 2024c.

603 Ming Li, Yanhong Li, and Tianyi Zhou. What happened in llms layers when trained for fast vs. slow
 604 thinking: A gradient perspective. *arXiv preprint arXiv:2410.23743*, 2024d.

605 Ming Li, Yong Zhang, Shuai He, Zhitao Li, Hongyu Zhao, Jianzong Wang, Ning Cheng, and
 606 Tianyi Zhou. Superfiltering: Weak-to-strong data filtering for fast instruction-tuning. In Lun-
 607 Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting*
 608 of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 14255–14273,
 609 Bangkok, Thailand, August 2024e. Association for Computational Linguistics. URL <https://aclanthology.org/2024.acl-long.769>.

610 Ming Li, Yong Zhang, Zhitao Li, Juhai Chen, Lichang Chen, Ning Cheng, Jianzong Wang, Tianyi
 611 Zhou, and Jing Xiao. From quantity to quality: Boosting LLM performance with self-guided
 612 data selection for instruction tuning. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.),
 613 *Proceedings of the 2024 Conference of the North American Chapter of the Association for Com-
 614 putational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 7595–
 615 7628, Mexico City, Mexico, June 2024f. Association for Computational Linguistics. URL
 616 <https://aclanthology.org/2024.naacl-long.421>.

617 Ming Li, Yanhong Li, Ziyue Li, and Tianyi Zhou. How instruction and reasoning data shape post-
 618 training: Data quality through the lens of layer-wise gradients. *arXiv preprint arXiv:2504.10766*,
 619 2025a.

620 Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke
 621 Zettlemoyer, and Mike Lewis. Contrastive decoding: Open-ended text generation as optimization.
 622 *arXiv preprint arXiv:2210.15097*, 2022.

623 Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian
 624 Xu, Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, Yingying Zhang, Fei Yin, Jiahua Dong, Zhiwei
 625 Li, Bao-Long Bi, Ling-Rui Mei, Junfeng Fang, Zhijiang Guo, Le Song, and Cheng-Lin Liu.
 626 From system 1 to system 2: A survey of reasoning large language models, 2025b. URL <https://arxiv.org/abs/2502.17419>.

627 Baohao Liao, Yuhui Xu, Hanze Dong, Junnan Li, Christof Monz, Silvio Savarese, Doyen Sahoo, and
 628 Caiming Xiong. Reward-guided speculative decoding for efficient llm reasoning. *arXiv preprint*
 629 *arXiv:2501.19324*, 2025.

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

633 Bill Yuchen Lin, Abhilasha Ravichander, Ximing Lu, Nouha Dziri, Melanie Sclar, Khyathi Chandu,
 634 Chandra Bhagavatula, and Yejin Choi. The unlocking spell on base llms: Rethinking alignment
 635 via in-context learning. *arXiv preprint arXiv:2312.01552*, 2023.

636 Changshu Liu, Shizhuo Dylan Zhang, Ali Reza Ibrahimzada, and Reyhaneh Jabbarvand. Codemind:
 637 A framework to challenge large language models for code reasoning, 2024a. URL <https://arxiv.org/abs/2402.09664>.

638 Wei Liu, Weihao Zeng, Keqing He, Yong Jiang, and Junxian He. What makes good data for
 639 alignment? a comprehensive study of automatic data selection in instruction tuning. *arXiv preprint*
 640 *arXiv:2312.15685*, 2023.

648 Xiaoxuan Liu, Cade Daniel, Langxiang Hu, Woosuk Kwon, Zhuohan Li, Xiangxi Mo, Alvin Cheung,
 649 Zhijie Deng, Ion Stoica, and Hao Zhang. Optimizing speculative decoding for serving large
 650 language models using goodput. *arXiv preprint arXiv:2406.14066*, 2024b.

651

652 Yue Liu, Jiaying Wu, Yufei He, Hongcheng Gao, Hongyu Chen, Baolong Bi, Jiaheng Zhang, Zhiqi
 653 Huang, and Bryan Hooi. Efficient inference for large reasoning models: A survey, 2025. URL
 654 <https://arxiv.org/abs/2503.23077>.

655

656 Haotian Luo, Li Shen, Haiying He, Yibo Wang, Shiwei Liu, Wei Li, Naiqiang Tan, Xiaochun Cao,
 657 and Dacheng Tao. O1-pruner: Length-harmonizing fine-tuning for o1-like reasoning pruning, 2025.
 658 URL <https://arxiv.org/abs/2501.12570>.

659

660 Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke
 661 Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time
 662 scaling, 2025. URL <https://arxiv.org/abs/2501.19393>.

663

664 OpenAI. OpenAI o1 System Card, December 2024. URL <https://cdn.openai.com/o1-system-card-20241205.pdf>.

665

666 Xiaoye Qu, Yafu Li, Zhaochen Su, Weigao Sun, Jianhao Yan, Dongrui Liu, Ganqu Cui, Daizong
 667 Liu, Shuxian Liang, Junxian He, Peng Li, Wei Wei, Jing Shao, Chaochao Lu, Yue Zhang, Xian-
 668 Sheng Hua, Bowen Zhou, and Yu Cheng. A survey of efficient reasoning for large reasoning
 669 models: Language, multimodality, and beyond, 2025. URL <https://arxiv.org/abs/2503.21614>.

670

671 Mohit Raghavendra, Vaskar Nath, and Sean Hendryx. Revisiting the superficial alignment hypothesis.
 672 *arXiv preprint arXiv:2410.03717*, 2024.

673

674 David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani,
 675 Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a benchmark. In
 676 *First Conference on Language Modeling*, 2024.

677

678 Matthew Renze and Erhan Guven. The benefits of a concise chain of thought on problem-solving in
 679 large language models. In *2024 2nd International Conference on Foundation and Large Language
 680 Models (FLLM)*, pp. 476–483. IEEE, 2024.

681

682 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 683 Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of
 684 mathematical reasoning in open language models, 2024. URL <https://arxiv.org/abs/2402.03300>.

685

686 Shannon Zejiang Shen, Hunter Lang, Bailin Wang, Yoon Kim, and David Sontag. Learning to decode
 687 collaboratively with multiple language models. *arXiv preprint arXiv:2403.03870*, 2024.

688

689 Xuan Shen, Yizhou Wang, Xiangxi Shi, Yanzhi Wang, Pu Zhao, and Jiuxiang Gu. Efficient reasoning
 690 with hidden thinking. *arXiv preprint arXiv:2501.19201*, 2025a.

691

692 Zhenyi Shen, Hanqi Yan, Linhai Zhang, Zhanghao Hu, Yali Du, and Yulan He. Codi: Compressing
 693 chain-of-thought into continuous space via self-distillation. *arXiv preprint arXiv:2502.21074*,
 694 2025b.

695

696 Hanshi Sun, Momin Haider, Ruiqi Zhang, Huitao Yang, Jiahao Qiu, Ming Yin, Mengdi Wang, Peter
 697 Bartlett, and Andrea Zanette. Fast best-of-n decoding via speculative rejection. *arXiv preprint
 698 arXiv:2410.20290*, 2024.

699

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

702

703 Qwen Team. Qwen2.5: A party of foundation models, September 2024. URL <https://qwenlm.github.io/blog/qwen2.5/>.

704

705 Qwen Team. Qwen3, April 2025. URL <https://qwenlm.github.io/blog/qwen3/>.

702 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay
 703 Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, and etc. Llama 2: Open foundation
 704 and fine-tuned chat models, 2023. URL <https://arxiv.org/abs/2307.09288>.

705 Guangya Wan, Yuqi Wu, Jie Chen, and Sheng Li. Dynamic self-consistency: Leveraging reasoning
 706 paths for efficient llm sampling. *arXiv preprint arXiv:2408.17017*, 2024.

708 Xiya Wang, Zhengyuan Yang, Linjie Li, Hongjin Lu, Yuancheng Xu, Chung-Ching Lin, Kevin
 709 Lin, Furong Huang, and Lijuan Wang. Scaling inference-time search with vision value model for
 710 improved visual comprehension. *arXiv preprint arXiv:2412.03704*, 2024a.

711 Xiya Wang, Yuhang Zhou, Xiaoyu Liu, Hongjin Lu, Yuancheng Xu, Feihong He, Jaehong Yoon,
 712 Taixi Lu, Gedas Bertasius, Mohit Bansal, et al. Mementos: A comprehensive benchmark for multi-
 713 modal large language model reasoning over image sequences. *arXiv preprint arXiv:2401.10529*,
 714 2024b.

715 Xiya Wang, Zhengyuan Yang, Chao Feng, Hongjin Lu, Linjie Li, Chung-Ching Lin, Kevin Lin,
 716 Furong Huang, and Lijuan Wang. Sota with less: Mcts-guided sample selection for data-efficient
 717 visual reasoning self-improvement. *arXiv preprint arXiv:2504.07934*, 2025a.

718 Yaoting Wang, Shengqiong Wu, Yuecheng Zhang, Shuicheng Yan, Ziwei Liu, Jiebo Luo, and Hao
 719 Fei. Multimodal chain-of-thought reasoning: A comprehensive survey, 2025b. URL <https://arxiv.org/abs/2503.12605>.

720 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le,
 721 and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models, 2023.
 722 URL <https://arxiv.org/abs/2201.11903>.

723 Yuyang Wu, Yifei Wang, Tianqi Du, Stefanie Jegelka, and Yisen Wang. When more is less: Under-
 724 standing chain-of-thought length in llms. *arXiv preprint arXiv:2502.07266*, 2025.

725 Heming Xia, Yongqi Li, Chak Tou Leong, Wenjie Wang, and Wenjie Li. Tokenskip: Controllable
 726 chain-of-thought compression in llms. *arXiv preprint arXiv:2502.12067*, 2025a.

727 Shijie Xia, Xuefeng Li, Yixin Liu, Tongshuang Wu, and Pengfei Liu. Evaluating mathematical
 728 reasoning beyond accuracy, 2025b. URL <https://arxiv.org/abs/2404.05692>.

729 Wei Xiong, Hanning Zhang, Chenlu Ye, Lichang Chen, Nan Jiang, and Tong Zhang. Self-
 730 rewarding correction for mathematical reasoning, 2025. URL <https://arxiv.org/abs/2502.19613>.

731 Silei Xu, Wenhao Xie, Lingxiao Zhao, and Pengcheng He. Chain of draft: Thinking faster by writing
 732 less. *arXiv preprint arXiv:2502.18600*, 2025.

733 Xiaohan Xu, Ming Li, Chongyang Tao, Tao Shen, Reynold Cheng, Jinyang Li, Can Xu,
 734 Dacheng Tao, and Tianyi Zhou. A survey on knowledge distillation of large language models.
 735 *ArXiv*, abs/2402.13116, 2024. URL <https://api.semanticscholar.org/CorpusID:267760021>.

736 Wang Yang, Xiang Yue, Vipin Chaudhary, and Xiaotian Han. Speculative thinking: Enhancing
 737 small-model reasoning with large model guidance at inference time, 2025. URL <https://arxiv.org/abs/2504.12329>.

738 Yixin Ye, Zhen Huang, Yang Xiao, Ethan Chern, Shijie Xia, and Pengfei Liu. Limo: Less is more for
 739 reasoning, 2025. URL <https://arxiv.org/abs/2502.03387>.

740 Ping Yu, Jing Xu, Jason Weston, and Ilia Kulikov. Distilling system 2 into system 1. *arXiv preprint
 741 arXiv:2407.06023*, 2024.

742 Jintian Zhang, Yuqi Zhu, Mengshu Sun, Yujie Luo, Shuofei Qiao, Lun Du, Da Zheng, Huajun
 743 Chen, and Ningyu Zhang. Lightthinker: Thinking step-by-step compression. *arXiv preprint
 744 arXiv:2502.15589*, 2025.

745 Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia
 746 Efrat, Ping Yu, Lili Yu, et al. Lima: Less is more for alignment. *Advances in Neural Information
 747 Processing Systems*, 36:55006–55021, 2023.

756	TABLE OF CONTENTS FOR APPENDIX	
757		
758	A Pseudo Code	16
759		
760	B FLOPs Calculation	17
761		
762		
763	C Related Works	19
764		
765	D Detailed Results	21
766		
767		
768		
769		
770		
771		
772		
773		
774		
775		
776		
777		
778		
779		
780		
781		
782		
783		
784		
785		
786		
787		
788		
789		
790		
791		
792		
793		
794		
795		
796		
797		
798		
799		
800		
801		
802		
803		
804		
805		
806		
807		
808		
809		

864 B FLOPs CALCULATION

866 Empirical latency depends on vendor-specific kernel fusion and memory layouts, so a timing measured
 867 on one backend may not transfer to another. Counting floating-point operations (FLOPs) provides
 868 a hardware-agnostic yardstick that isolates algorithmic differences. The performance figures we
 869 report are presented in TeraFLOPs (TFLOPs), where one TFLOP equals 10^{12} FLOPs. The generation
 870 process for the speculative decoding (SD)-type methods can be categories in three modes, prefill,
 871 decode, and prefix:

- 873 • *Prefill*: This stage corresponds to loading the entire prompt into the model without any
 KV cache. The cost here is determined by the length of the input (e.g., question and chat
 874 template) and is incurred only once per inference.
- 876 • *Decode*: This stage involves the continuous generation of new tokens, where each token
 877 is generated sequentially with the benefit of the KV cache. This is typically the most
 878 computationally intensive part of inference.
- 879 • *Prefix*: This stage occurs whenever there is a switch between models (e.g., from Drafting
 880 to Leading or vice versa), it mainly used in the SD-type methods. Here, the model must
 881 process a sequence of tokens given the KV cache of previous tokens. In SD, this happens
 882 each time the verifier model checks tokens generated by the draft model, and each time the
 883 draft model resumes generation after verification. In FoReal-Decoding, prefix computation
 884 is triggered whenever control is transferred between the two models. Notably, the prefix
 885 stage incurs a nearly fixed cost (largely independent of the number of tokens), which is a key
 886 reason why SD-based methods can be efficient even though both models process all tokens.

887 The total FLOPs for a single inference is thus computed as: $\text{Prefill}(\text{Drafting}) + \text{Prefill}(\text{Leading}) +$
 888 $\text{Decode}(\text{Drafting}) + \text{Decode}(\text{Leading}) + \text{Prefix}(\text{Drafting} \rightarrow \text{Leading}) + \text{Prefix}(\text{Leading} \rightarrow \text{Drafting})$.
 889 When GPU memory is sufficient, profiling shows that producing multiple tokens during the prefix
 890 phase costs almost the same as decoding a single token. Therefore, we upper-bound the prefix cost
 891 by the single-token decode cost, following Speculative Thinking (Yang et al., 2025). We calculate
 892 the precise total FLOPs using the detailed formulas presented below, which is based on (Chen et al.,
 893 2024; Han, 2024; Yang et al., 2025). The resulting total FLOPs are then converted to estimated
 894 TFLOPs for reporting.

895 The variables involved are defined as:

- 897 • s : Represents the sequence length.
 - 898 – For the prefill stage ($\text{FLOPs}_{\text{prefill}}(s)$), s is the length of the input prompt, denoted as p_l .
 - 899 – For the decode stage ($\text{FLOPs}_{\text{decode}}(s)$), s is the current length of the context (prompt +
 900 tokens generated so far) that the model attends to via its Key-Value (KV) cache.
- 902 • h : The hidden size of the model.
- 903 • h' : The intermediate size of the feed-forward network (FFN).
- 904 • n : The number of attention heads.
- 906 • p_l : The length of the initial problem prompt.
- 907 • d_l : The number of tokens to be generated in the solution.

909 It is noted that the hidden size h relates to the number of attention heads n and the size of each
 910 attention head d by $h = n \cdot d$.

911 The FLOPs for the prefill stage, which processes the initial input prompt of length $s = p_l$, is given by
 912 Equation 8:

$$913 \text{FLOPs}_{\text{prefill}}(s) = 8sh^2 + 16sh + 4s^2h + 4s^2n + 6shh' + 2sh' \quad (8)$$

915 The FLOPs for the decode stage, which generates a single token when the current KV cache has a
 916 length of s , is given by Equation 9:

$$917 \text{FLOPs}_{\text{decode}}(s) = 8h^2 + 16h + 4sh + 4sn + 6hh' + 2h' \quad (9)$$

918 The total FLOPs to generate d_l tokens from a prompt of length p_l combines the prefill cost for the
 919 prompt and the sum of decode costs for each generated token, as shown in Equation 10:
 920

$$921 \quad \text{FLOPs}_{\text{total}} = \text{FLOPs}_{\text{prefill}}(p_l) + \sum_{i=0}^{d_l-1} \text{FLOPs}_{\text{decode}}(p_l + i) \quad (10)$$

924 In this formula, for the i -th token being generated (0-indexed), the argument to $\text{FLOPs}_{\text{decode}}$ is $p_l + i$,
 925 representing the sequence length in the KV cache at that generation step.

926
 927
 928
 929
 930
 931
 932
 933
 934
 935
 936
 937
 938
 939
 940
 941
 942
 943
 944
 945
 946
 947
 948
 949
 950
 951
 952
 953
 954
 955
 956
 957
 958
 959
 960
 961
 962
 963
 964
 965
 966
 967
 968
 969
 970
 971

972 C RELATED WORKS
973974 C.1 LARGE REASONING MODELS
975

976 Recent advances in large language models (LLMs) have spurred a surge of work aimed at strengthening
977 their reasoning abilities (Ahn et al., 2024; Besta et al., 2025; Chen et al., 2025a). Core reasoning
978 skills are already instilled during pre-training, where models absorb commonsense and mathematical
979 patterns from vast text corpora (Touvron et al., 2023; OpenAI, 2024). Researchers have therefore
980 concentrated on post-training techniques to further polish these skills. One prominent direction
981 employs reinforcement learning to nudge models toward more effective chains of thought (Shao et al.,
982 2024; Xiong et al., 2025; Cui et al., 2025; Wang et al., 2025a). Another line shows that carefully
983 curated instruction-tuning data can likewise deliver sizable gains in reasoning accuracy (Ye et al.,
984 2025; Muennighoff et al., 2025; Wang et al., 2024a).

985 Despite the impressive benchmark scores of recent Reasoning Language Models, several studies
986 have begun to probe the quality and efficiency of the reasoning they generate. (Xia et al., 2025b)
987 conduct a broad assessment and reveal substantial redundancy in many model-produced solutions.
988 Follow-up investigations (Chen et al., 2025c; Cuadron et al., 2025; Qu et al., 2025; Liu et al., 2025;
989 Fan et al., 2025) underscore an “overthinking” phenomenon, whereby models craft unduly verbose
990 derivations even for simple problems. Capitalizing on this trait, (Kumar et al., 2025) demonstrate
991 a slowdown attack: small input perturbations can trigger excessive reasoning, markedly degrading
992 inference speed.

993 To alleviate overthinking and improve efficiency for reasoning models, a series of efficient reasoning
994 methods has been proposed. For example, (Yu et al., 2024; Team et al., 2025; Aggarwal & Welleck,
995 2025; Xia et al., 2025a; Luo et al., 2025) utilize model-based methods that either add further
996 constraints on RL rewards or SFT on diverse lengths of CoTs, (Hao et al., 2024; Shen et al., 2025b;a;
997 Zhang et al., 2025) utilize latent-space reasoning methods that transfer the massive tokens into the
998 embedding space, (Han et al., 2024; Xu et al., 2025; Renze & Guven, 2024) utilize the prompt-based
999 methods, (Sun et al., 2024; Wan et al., 2024; Wu et al., 2025) utilize the sampling methods. Most of
1000 these methods either require further post-training or manipulating the distribution of LRM itself.

1001 C.2 ALIGNMENT AND TOKEN PATTERN ANALYSIS
1002

1003 A key empirical foundation for LLM Alignment is LIMA (Zhou et al., 2023), which demonstrated
1004 that just 1,000 carefully curated instruction–response pairs are already enough for LLM alignment,
1005 crystallizing the “superficial alignment” hypothesis. While a line of work directly follows the
1006 hypotheses by introducing data selection or alignment methods (Chen et al., 2023b; Li et al., 2024f;
1007 2023; 2024b; Du et al., 2023; Li et al., 2024a; Bukharin & Zhao, 2023; Liu et al., 2023; Li et al.,
1008 2024e;c; 2025a; Xu et al., 2024), there are also works that try to further investigate this phenomenon.

1009 (Lin et al., 2023) provides a comprehensive token-level evidence by comparing the top-k token
1010 distributions of base models and their chat-tuned counterparts. The authors show that almost all
1011 divergence concentrates on discourse markers, politeness phrases, and safety disclaimers, while core
1012 content tokens remain unchanged. (Chen et al., 2025b) dissects which prompt-level cues are sufficient
1013 (and which are not) for alignment, showing that reasoning gaps emerge precisely where superficial
1014 patterns end. The debate has sparked push-back as well: (Raghavendra et al., 2024) demonstrates
1015 systematic performance gains when the amount of post-training data scales up, arguing that some
1016 deeper representational changes do accrue beyond mere style. Researchers are also probing where
1017 superficial signals live: (Li & Kim, 2024) argues that data curation, not extra optimization steps, is the
1018 primary lever: filtering for safety disclaimers yields larger alignment jumps than adding thousands
1019 of generic examples. Together, these works paint a nuanced picture: much of the alignment gap
1020 after pre-training is indeed “superficial”, residing in a narrow band of stylistic tokens that can be
1021 manipulated through tiny prompts, judicious data selection. However, in this paper, we show that *the
1022 reasoning capabilities might not be as superficial as previous findings*.

1023 C.3 SPECULATIVE DECODING AND COLLABORATIVE DECODING
1024

1025 Speculative decoding, inaugurated by (Leviathan et al., 2023), uses a small “draft” model to pro-
pose several tokens that the large “target” model then verifies in one batch, yielding 2–3× latency

1026 reductions with provably identical output distributions. Follow-up work, such as (Chen et al., 2023a)
 1027 extends the idea to 70 B-parameter models and confirms similar speed-ups, while (Cai et al., 2024)
 1028 replaces the external draft model with extra decoding heads to remove system complexity System-
 1029 level schedulers like (Liu et al., 2024b) dynamically adapt draft length to traffic conditions and push
 1030 end-to-end gains beyond $3\times$ in production settings.

1031 Collaborative decoding improves text quality by letting multiple models cooperate during generation.
 1032 (Li et al., 2022) runs a weak “amateur” model alongside a strong “expert” and selects tokens that
 1033 maximize their likelihood gap, sharply reducing repetition and incoherence without retraining. (Jin
 1034 et al., 2024) introduces a critical-token strategy that switches to the pretrained base model whenever
 1035 factual precision is needed, cutting hallucinations in instruction-tuned LLMs. At an even finer grain,
 1036 (Shen et al., 2024) treats “who should emit the next token” as a latent variable, enabling on-the-fly
 1037 delegation between a generalist LLM and domain specialists and outperforming any single model on
 1038 cross-domain tasks.

1039
 1040 **Comparisons with related methods.** In speculative decoding (Leviathan et al., 2023), the final
 1041 text provably matches what the large model alone would have produced. However, our ***FoReal-
 1042 Decoding*** focuses on the reasoning-heavy scenarios where the responses generated by the LRM itself
 1043 are not desirable due to the overthinking. Thus, our method serves as a deliberate mixture of two
 1044 distributions, aiming at reducing the overthinking problem of LRM by inserting the distribution
 1045 from weaker models, and at the same time increasing the efficiency. A recent work, RSD (Liao
 1046 et al., 2025), also aims at reducing computation cost by utilizing speculative decoding. However, it
 1047 introduces an additional process reward model as the judge, while our method focuses on utilizing
 1048 the collaborative models themselves only, thus, it is largely different from our settings. Another
 1049 concurrent work, Speculative Thinking (Yang et al., 2025), also shares similar motivation as ours,
 1050 in which a “small-writes, large-fixes” mechanism is utilized, which differs from our “large-leads,
 1051 small-follows”. Moreover, ***FoReal-Decoding*** provides a smooth transition from the small to the
 1052 large model, representing wider trade-off scopes.

1053
 1054
 1055
 1056
 1057
 1058
 1059
 1060
 1061
 1062
 1063
 1064
 1065
 1066
 1067
 1068
 1069
 1070
 1071
 1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079

1080 D DETAILED RESULTS

1081 Table 4 and Table 5 show the detailed results of different settings of our method.

1082 Table 4: The detailed results of different collaborative settings on AIME24, GPQA-D, MATH500,
1083 and AMC23, including length and ratio.

1087	Model	AIME24			GPQA-D			MATH500			AMC23							
		1088 Method	Config	1089 ACC (%)	Length	Ratio	1090 TFLOPs	ACC (%)	Length	Ratio	1091 TFLOPs	ACC (%)	Length	Ratio	1092 TFLOPs			
DeepSeek-R1-Distill-Qwen-32B + DeepSeek-R1-Distill-Qwen-1.5B																		
1089	DeepSeek-R1-Distill-Qwen-32B	66.7	13035	-	15.72	59.6	6602	-	8.09	93.6	3542	-	4.13	95.0	6243	-	7.54	
1090	DeepSeek-R1-Distill-Qwen-1.5B	23.3	18021	-	2.86	22.2	8696	-	1.13	81.4	6704	-	1.14	65.0	13311	-	2.51	
1091	FoReal-Decoding	$n=15, p=0.4$	33.3	11876	0.272	5.60	43.3	5841	0.294	2.47	90.2	3402	0.312	1.45	80.0	6043	0.304	2.91
1092	FoReal-Decoding	$n=15, p=0.6$	50.0	10934	0.401	6.77	48.2	7007	0.431	4.50	91.4	3995	0.452	2.40	80.0	6460	0.429	3.99
1093	FoReal-Decoding	$n=15, p=0.8$	50.0	11532	0.527	8.47	54.6	6110	0.570	4.69	93.4	3658	0.590	2.70	90.0	7037	0.571	5.37
1094	FoReal-Decoding	$n=15, p=1.0$	66.7	10617	0.666	9.16	56.6	6796	0.692	6.21	93.2	3655	0.726	3.14	92.5	5942	0.708	5.28
1095	FoReal-Decoding	$n=25, p=0.8$	53.3	12081	0.676	10.95	57.7	6223	0.702	5.65	92.6	3585	0.719	3.13	92.5	5529	0.710	4.99
1096	FoReal-Decoding	$n=25, p=1.0$	66.7	11116	0.683	10.54	57.6	6065	0.882	6.68	94.5	3403	0.890	3.50	95.0	5422	0.872	5.66
DeepSeek-R1-Distill-Qwen-32B + Qwen2.5-1.5B-Instruct																		
1097	DeepSeek-R1-Distill-Qwen-32B	66.7	13035	-	15.72	59.6	6602	-	8.09	93.6	3542	-	4.13	95.0	6243	-	7.54	
1098	Qwen2.5-1.5B-Instruct	0.0	998	-	0.12	23.7	923	-	0.12	49.2	747	-	0.09	15.0	818	-	0.10	
1099	FoReal-Decoding	$n=15, p=0.8$	20.0	12584	0.571	9.05	47.5	7013	0.587	5.63	76.2	3792	0.614	2.85	65.0	7629	0.514	5.22
1100	FoReal-Decoding	$n=15, p=1.0$	20.0	14188	0.588	11.19	47.5	6294	0.737	5.86	85.9	3894	0.750	3.28	65.0	7673	0.707	6.15
1101	FoReal-Decoding	$n=25, p=0.8$	36.7	11575	0.710	9.58	56.7	4718	0.719	4.37	82.0	3025	0.729	2.52	72.5	5415	0.649	4.65
1102	FoReal-Decoding	$n=25, p=1.0$	40.0	11239	0.813	11.00	57.1	5944	0.887	6.27	90.8	3403	0.894	3.36	92.5	6989	0.867	6.88
DeepSeek-R1-Distill-Qwen-1.5B + Qwen2.5-7B-Instruct																		
1103	DeepSeek-R1-Distill-Qwen-1.5B	23.3	18021	-	2.86	22.2	8696	-	1.13	81.4	6704	-	1.14	65.0	13311	-	2.51	
1104	Qwen2.5-7B-Instruct	6.7	1243	-	0.95	38.4	1054	-	0.89	76.0	773	-	0.61	52.5	994	-	0.75	
1105	FoReal-Decoding	$n=15, p=0.8$	16.7	4120	0.545	2.05	34.3	2130	0.602	1.07	76.4	1341	0.634	0.57	57.5	2515	0.580	1.08
1106	FoReal-Decoding	$n=15, p=1.0$	16.7	14132	0.651	6.47	29.8	7913	0.703	3.08	79.6	3480	0.735	1.42	52.5	7330	0.686	3.35
1107	FoReal-Decoding	$n=25, p=0.8$	20.0	4474	0.693	1.57	33.1	1801	0.718	0.80	78.6	1498	0.736	0.55	65.0	3778	0.683	1.76
1108	FoReal-Decoding	$n=25, p=1.0$	23.3	11436	0.841	3.18	29.3	6800	0.863	2.53	79.2	3586	0.891	1.04	60.0	5721	0.865	1.66

1134
 1135
 1136
 1137
 1138
 1139
 1140
 1141
 1142
 1143

Table 5: The detailed results of different collaborative settings on AIME24 and AMC23, including length and ratio.

Model		AIME24				AMC23			
Method	Config	ACC (%)	Length	Ratio	TFLOPs	ACC (%)	Length	Ratio	TFLOPs
DeepSeek-R1-Distill-Qwen-32B + DeepSeek-R1-Distill-Qwen-1.5B									
DeepSeek-R1-Distill-Qwen-32B		66.7	13035	-	15.72	95.0	6243	-	7.54
DeepSeek-R1-Distill-Qwen-1.5B		23.3	18021	-	2.86	65.0	13311	-	2.51
FoReal-Decoding	$n=5, p=0.2$	23.3	12926	0.076	3.47	77.5	5634	0.089	1.39
FoReal-Decoding	$n=5, p=0.4$	26.7	11590	0.145	3.92	80.0	6549	0.157	2.18
FoReal-Decoding	$n=5, p=0.6$	36.7	11560	0.202	4.81	80.0	7081	0.228	2.95
FoReal-Decoding	$n=5, p=0.8$	43.3	11907	0.270	5.82	82.5	6399	0.294	3.29
FoReal-Decoding	$n=5, p=1.0$	50.0	13750	0.328	7.82	85.0	6916	0.355	3.86
FoReal-Decoding	$n=15, p=0.2$	26.7	12457	0.138	4.03	70.0	6680	0.154	2.05
FoReal-Decoding	$n=15, p=0.4$	33.3	11876	0.272	5.60	80.0	6043	0.303	2.91
FoReal-Decoding	$n=15, p=0.6$	50.0	10934	0.401	6.77	80.0	6460	0.429	3.99
FoReal-Decoding	$n=15, p=0.8$	50.0	11532	0.527	8.47	90.0	7037	0.571	5.37
FoReal-Decoding	$n=15, p=1.0$	66.7	10617	0.666	9.16	92.5	5942	0.708	5.28
FoReal-Decoding	$n=25, p=0.2$	36.7	10805	0.178	3.88	77.5	6798	0.193	2.32
FoReal-Decoding	$n=25, p=0.4$	33.3	11428	0.347	6.30	80.0	5929	0.362	3.17
FoReal-Decoding	$n=25, p=0.6$	50.0	10816	0.515	7.71	90.0	6169	0.537	4.49
FoReal-Decoding	$n=25, p=0.8$	53.3	12081	0.675	10.95	92.5	5529	0.710	4.99
FoReal-Decoding	$n=25, p=1.0$	66.7	11117	0.683	10.54	95.0	5422	0.872	5.66
FoReal-Decoding	$n=\infty, p=0.2$	30.0	12241	0.204	4.84	75.0	6502	0.216	2.43
FoReal-Decoding	$n=\infty, p=0.4$	46.7	11906	0.417	7.37	85.0	6719	0.423	4.07
FoReal-Decoding	$n=\infty, p=0.6$	50.0	11515	0.605	9.69	92.5	5671	0.607	4.42
FoReal-Decoding	$n=\infty, p=0.8$	60.0	10538	0.798	10.83	92.5	5925	0.797	5.87
FoReal-Decoding	$n=\infty, p=1.0$	66.7	13035	1.000	15.72	95.0	6244	1.000	7.54
DeepSeek-R1-Distill-Qwen-1.5B + Qwen2.5-7B-Instruct									
DeepSeek-R1-Distill-Qwen-1.5B		23.3	18021	-	2.86	65.0	13311	-	2.51
Qwen2.5-7B-Instruct		6.7	1243	-	0.95	52.5	994	-	0.75
FoReal-Decoding	$n=5, p=0.2$	10.0	1047	0.170	0.73	50.0	923	0.179	0.64
FoReal-Decoding	$n=5, p=0.4$	10.0	1381	0.230	0.91	55.0	1065	0.244	0.69
FoReal-Decoding	$n=5, p=0.6$	13.3	2377	0.306	1.61	62.5	2574	0.302	1.97
FoReal-Decoding	$n=5, p=0.8$	13.3	4203	0.345	2.87	47.5	2897	0.373	1.83
FoReal-Decoding	$n=5, p=1.0$	16.7	7236	0.382	4.45	50.0	5614	0.428	3.24
FoReal-Decoding	$n=15, p=0.2$	3.3	1936	0.208	1.38	47.5	985	0.224	0.65
FoReal-Decoding	$n=15, p=0.4$	16.7	1189	0.339	0.70	45.0	1055	0.360	0.61
FoReal-Decoding	$n=15, p=0.6$	16.7	1793	0.455	0.92	52.5	1307	0.482	0.65
FoReal-Decoding	$n=15, p=0.8$	16.7	4120	0.545	2.05	57.5	2515	0.580	1.08
FoReal-Decoding	$n=15, p=1.0$	16.7	14132	0.651	6.47	52.5	7330	0.686	3.35
FoReal-Decoding	$n=25, p=0.2$	13.3	1243	0.249	0.80	50.0	958	0.231	0.62
FoReal-Decoding	$n=25, p=0.4$	20.0	1317	0.389	0.73	42.5	1077	0.405	0.59
FoReal-Decoding	$n=25, p=0.6$	16.7	1743	0.536	0.79	57.5	2047	0.560	1.14
FoReal-Decoding	$n=25, p=0.8$	20.0	4474	0.693	1.57	65.0	3778	0.683	1.76
FoReal-Decoding	$n=25, p=1.0$	23.3	11436	0.841	3.18	65.0	5721	0.865	1.66
FoReal-Decoding	$n=\infty, p=0.2$	13.3	1072	0.260	0.69	50.0	986	0.290	0.61
FoReal-Decoding	$n=\infty, p=0.4$	6.7	1276	0.420	0.68	42.5	1140	0.467	0.56
FoReal-Decoding	$n=\infty, p=0.6$	10.0	1914	0.614	0.78	57.5	1324	0.618	0.53
FoReal-Decoding	$n=\infty, p=0.8$	23.3	4244	0.788	1.26	65.0	2854	0.817	0.79
FoReal-Decoding	$n=\infty, p=1.0$	23.3	18021	1.000	2.86	65.0	13311	1.000	2.51

1181
 1182
 1183
 1184
 1185
 1186
 1187