

REASONING UNDER CONSTRAINT: HOW BATCH PROMPTING SUPPRESSES OVERTHINKING IN REASONING MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Large Reasoning Models (LRMs) achieve strong performance through explicit chain-of-thought reasoning but suffer from *overthinking*: generating excessive reasoning tokens even for trivial queries. Beyond inflating cost, overthinking can be self-defeating: models enter recursive self-doubt loops that exhaust token budgets without producing an answer, causing API timeouts that directly hurt accuracy. We present an empirical study showing that **batch prompting**, originally introduced for throughput optimization, effectively suppresses overthinking at inference time. Across 13 diverse benchmarks with DeepSeek-R1 and OpenAI-o1, batch prompting reduces reasoning tokens by 76% (2,950→710), on average, while preserving or improving accuracy. Through behavioral analysis, we find that batching induces three beneficial effects: (1) it reduces per-query reasoning effort when multiple queries share a context; (2) it enables pattern induction, where models generalize from earlier examples to solve later ones; and (3) it suppresses hedging behavior (e.g., “wait,” “let me double-check”) that signals metacognitive loops. We also show that explicit prompt constraints (“Use no more than 100 tokens in thinking.”) fail to reduce overthinking; models either ignore them or sacrifice accuracy. These findings reframe batch prompting as more than a cost optimization: it is a practical inference-time technique that improves efficiency and reliability without model modification¹.

1 INTRODUCTION

Large Reasoning Models (LRMs) have achieved state-of-the-art performance in mathematics, code generation, and logical planning by producing explicit chains-of-thought (CoT) that expose intermediate reasoning steps (Zhong et al., 2024; Guo et al., 2025). However, this deliberate reasoning comes at a cost: models frequently **overthink**, generating thousands of reasoning tokens even for trivial queries (Cuadron et al., 2025; Huang et al., 2025; Chen et al., 2025; Srivastava & Yao, 2025). For API-based deployments where reasoning tokens are billed at premium rates, this behavior renders LRMs economically impractical at scale. Worse, overthinking is not merely wasteful; it is accuracy-degrading. Models can enter recursive self-doubt loops that exhaust their token budget without producing an answer, causing API timeouts that directly reduce task accuracy.

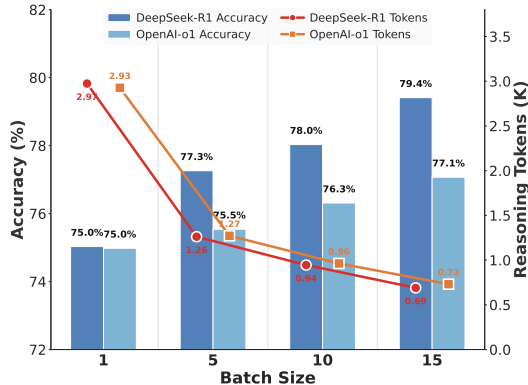


Figure 1: Results for DeepSeek-R1 and OpenAI-o1 (averaged over 13 datasets) demonstrate that *batching dramatically reduces reasoning tokens without hurting accuracy*. Increasing batch size improves accuracy while achieving a 76% reduction in reasoning tokens.

¹Code, prompts, and all model outputs will be made publicly available upon acceptance.

Prior approaches to curbing overthinking require model-internal access: self-training methods teach early stopping (Wei et al., 2025), while activation-steering (Huang et al., 2025) suppresses token generation by modifying hidden states. These strategies are inapplicable to closed-weight models such as OpenAI-o1. This leaves a fundamental gap: *can we reduce overthinking entirely at inference time, without touching the model?*

We answer affirmatively with a surprising finding: batch prompting (Cheng et al., 2023), originally introduced for throughput optimization, acts as an effective *implicit regularizer* that suppresses overthinking. When multiple queries are presented together, the model distributes its *reasoning effort* across all questions rather than exhaustively analyzing each one, analogous to how humans abbreviate reasoning when juggling multiple tasks under time pressure. We hypothesize that this arises because the shared context creates a *soft behavioral constraint*: the model allocates fewer reasoning tokens per query as the number of concurrent queries increases.

Figure 1 quantifies this effect. Increasing the batch size from 1 to 15 reduces the reasoning tokens by **76%** (2,950 \mapsto 710 tokens), on average, while preserving or improving accuracy. This $\sim 4.5\times$ reduction in reasoning tokens is achieved with *zero* model modification, only a change in how queries are presented. The effect is consistent across both DeepSeek-R1 (Guo et al., 2025) and OpenAI-o1 (Zhong et al., 2024), suggesting a consistent pattern in how reasoning models allocate computation under shared context.

Beyond efficiency, we observe that batching induces several beneficial behavioral changes. Models exhibit *pattern induction*: when solving similar queries together, they generalize patterns from earlier examples to harder ones in the same batch. Batching also *suppresses hedging*, tokens like “wait,” “let me double-check” that signal metacognitive loops, drop dramatically (from 21 occurrences \mapsto 1 in our analysis). These effects suggest that batching acts as more than a cost optimization: it behaviorally regularizes the reasoning process.

We also find that, in certain cases, single-query overthinking escalates to the point of complete failure: the model exhausts its token budget re-deriving and second-guessing an already-correct answer, producing no output before the API times out. This partly explains why single-query accuracy can fall below that of batched inference; batch prompting prevents such runaway generation by distributing reasoning effort across queries (Section 5).

We validate these findings through a comprehensive study across 13 diverse benchmarks spanning arithmetic reasoning, question answering, structured extraction, and scientific tasks. Our contributions are as follows:

1. **Batch prompting suppresses overthinking:** We show that batch prompting, beyond its known efficiency benefits, effectively suppresses overthinking in LRMs at inference time, reducing reasoning tokens by 76% while preserving or improving accuracy, without model modification.
2. **Comprehensive empirical validation:** We validate across 13 benchmarks, 2 frontier models (DeepSeek-R1, OpenAI-o1), and 4 batch sizes, establishing batch prompting as a robust, model-agnostic technique for efficient LRM deployment.
3. **Behavioral analysis:** We provide a detailed analysis of *how* batching regularizes reasoning: through pattern induction, hedging suppression, and adaptive token allocation that preserves accuracy on difficult instances.
4. **Mechanistic hypotheses:** We propose candidate explanations for why batching suppresses overthinking, including shared-context pressure on reasoning depth, sequential anchoring from earlier responses, and implicit difficulty calibration, and discuss how our empirical findings relate to each.

2 RELATED WORKS

Overthinking in LRMs. Large Reasoning Models frequently generate unnecessarily long reasoning traces, a phenomenon termed “overthinking” by Chen et al. (2025). They show diminishing returns from extended chains-of-thought, with models producing $10\times$ more tokens than necessary on simple tasks. Cuadron et al. (2025) demonstrates that overthinking actively harms performance through error accumulation, while Huang et al. (2025) proposes activation-steering to suppress excessive

generation. However, these model-internal approaches are inapplicable to closed-weight APIs like OpenAI-o1. Most closely related to our work, Yan et al. (2025) observes that batch processing induces implicit resource competition in reasoning models, leading to more efficient reasoning allocation. They leverage this insight to generate preference data from batch outputs and train models via reinforcement learning to replicate this adaptive behavior during single-question inference. While conceptually similar in recognizing the benefits of batching, their approach requires model fine-tuning and access to model weights. Our work demonstrates that the benefits of batch-induced reasoning compression can be obtained *directly at inference time* without any training, making it immediately applicable to closed API settings.

Batch Prompting. Cheng et al. (2023) introduced batch prompting as a throughput optimization, showing that grouping queries reduces per-query cost through prompt overhead amortization. Their focus was on standard Large Language Models (LLMs) and cost efficiency. We discover an orthogonal benefit: for *reasoning* models specifically, batching serves as an implicit regularizer that suppresses overthinking—an effect not observed or characterized in prior work.

Prompt-Based Efficiency. Zero-shot prompting techniques like self-consistency (Wang et al., 2023) and least-to-most prompting (Zhou et al., 2023) aim to improve accuracy, but often incur added generation cost. More recent work reduces CoT length via output constraints or refinement. Our work operates in a black-box prompt-only setting, demonstrating that batching provides a natural, effective mechanism for reasoning brevity without explicit token constraints.

3 METHODOLOGY AND FORMAL SETUP

We consider the problem of minimizing inference-time cost for multi-step reasoning tasks without compromising model accuracy. Let \mathcal{M} be a reasoning-oriented language model (e.g., DeepSeek-R1 or OpenAI-o1) that consumes prompt tokens and emits a *response* comprising reasoning (thinking) tokens and output (answer) tokens. For a single query q , the total token cost consists of fixed prompt overhead T_{fix} , query tokens T_q , reasoning tokens T_r , and output tokens T_o . Thus, the total token usage is:

$$C(q) = T_{\text{fix}} + T_q + T_r + T_o.$$

This cost is incurred independently for each query in traditional inference, leading to inefficiencies, especially in high-throughput applications. Crucially, for reasoning-oriented models, $T_r \gg T_o$ in practice, i.e., reasoning traces often consume 5–10× more tokens than the final answer, making T_r the primary target for cost reduction.

3.1 BATCH PROMPTING FOR EFFICIENCY

To reduce this cost, we apply **batch prompting** (Cheng et al., 2023), which groups b queries into a single prompt (Figure 2):

$$P = [\text{Instruction}] \parallel q_1 \parallel q_2 \parallel \cdots \parallel q_b,$$

where \parallel denotes concatenation. The model processes P in a single forward pass, producing responses $\{(r_1, o_1), \dots, (r_b, o_b)\}$ where each response comprises reasoning tokens r_i and output tokens o_i . The total cost for b queries is then:

$$C_{\text{batch}} = T_{\text{fix}} + \sum_{i=1}^b (T_{q_i} + T_{r_i} + T_{o_i}).$$

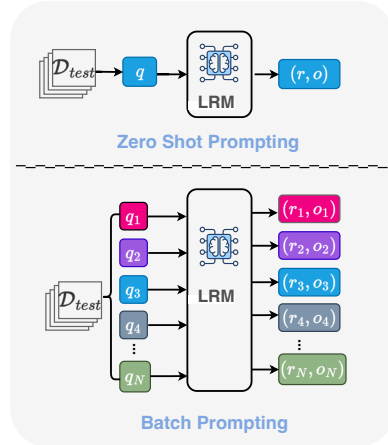


Figure 2: Batch prompting mechanism. Multiple queries ($q = [q_1, \dots, q_N]$) are processed jointly, forcing the model to distribute attention and computation across all items. This implicit constraint suppresses overthinking without explicit instruction. The final responses are produced with a combination of reasoning and output tokens (r_i, o_i) .

The effective per-query cost becomes:

$$\frac{C_{\text{batch}}}{b} = \frac{T_{\text{fix}}}{b} + \frac{1}{b} \sum_{i=1}^b (T_{q_i} + T_{r_i} + T_{o_i}),$$

showing that prompt overhead is amortized as $1/b$, yielding substantial gains for large b . In practice, batching leads to $\sim 76\%$ reduction in token usage on average with minimal loss in accuracy (see Figure 1).

3.2 WHY DOES BATCHING REDUCE REASONING? CANDIDATE HYPOTHESES

The empirical observation that batching reduces reasoning tokens while preserving accuracy calls for explanation. We propose three candidate hypotheses that are consistent with our findings. We emphasize that these are *conjectures* informed by our experiments, not proven mechanisms; disentangling their individual contributions remains an open question.

Hypothesis 1: Shared-Context Pressure. We conjecture that when b queries share a single context window, the model faces implicit pressure to produce shorter per-query reasoning traces. In single-query inference, the model can devote its full generation budget to a single query, but with batching it must produce b responses within the same generation, which may discourage extensive deliberation on any individual item. The precise mechanism behind this effect remains unclear and could involve learned generation-length priors, context-window utilization patterns, or other factors. Nevertheless, the empirical trend is consistent: per-query reasoning tokens decrease monotonically with batch size across both models (Figure 1).

Hypothesis 2: Sequential Anchoring. When generating responses for batched queries, the model processes them sequentially within a single generation. Let \hat{y}_i denote the full response to query q_i . The conditional generation follows:

$$P(\hat{y}_i \mid q_1, \dots, q_b, \hat{y}_1, \dots, \hat{y}_{i-1}).$$

One possible explanation for the observed brevity is that if \hat{y}_1 adopts a concise reasoning style, subsequent responses $\hat{y}_2, \dots, \hat{y}_b$ are biased toward similar brevity through in-context pattern matching. This is consistent with our observation that batching enables pattern induction (Section 5), though we have not isolated this effect from the shared-context hypothesis above.

Hypothesis 3: Implicit Difficulty Calibration. We hypothesize that the presence of multiple queries provides implicit calibration of task difficulty. When a model sees query q_i alongside other queries $\{q_j\}_{j \neq i}$, the batch context may signal that exhaustive reasoning is not required for any single item. This is analogous to how humans adjust effort when faced with multiple tasks under time pressure. However, we note that our experiments use semantically homogeneous batches, which limits our ability to test this hypothesis against heterogeneous mixtures.

Regardless of the underlying mechanism, the empirical effect is robust. Let $R(q)$ denote the reasoning trace length. We observe:

$$\mathbb{E}[|R(q)|_{\text{batch}}] \ll \mathbb{E}[|R(q)|_{\text{single}}],$$

with the reduction factor increasing with batch size (Figure 1). Importantly, this compression appears *adaptive*: models allocate more tokens to harder queries within a batch while abbreviating easier ones, preserving accuracy despite reduced total computation.

4 EXPERIMENTS

Datasets and Metrics. We evaluate across 13 diverse benchmarks spanning arithmetic (GSM8k (Mirzadeh et al., 2024), Math500 (Hendrycks et al., 2021), Game of 24 (Yao et al., 2023a)), factual and epistemic QA (GPQA (Rein et al., 2024), SciTab (Lu et al., 2023), Epistemic (Srivastava et al., 2023), StrategyQA (Geva et al., 2021)), structured tasks (Event Extraction (Doddington et al., 2004), Object Count (Srivastava et al., 2023), Last Letter Concatenation (Srivastava et al., 2023)),

All results reported with <i>OpenAI-o1</i>																	
Batch Size	Info Extraction				Reasoning / QA						Structured / Scientific					Average	
	TI	TC	AI	AC	Penguins	GSM8K	StratQA	IIT	GPQA	Last	Epis.	ObjCnt	Math	BBEH	SciTab		Game24
Accuracies (%)																	
1	58.40	54.70	31.80	29.80	99.00	97.00	89.00	87.00	79.00	95.00	84.00	99.00	88.00	42.00	70.00	96.00	74.98
5	55.80	50.60	25.20	23.00	100.00	96.00	91.00	93.00	83.00	97.00	91.00	100.00	91.00	45.00	70.00	97.00	75.54
10	59.70	57.10	29.40	25.70	99.00	97.00	92.00	91.00	81.00	97.00	90.00	99.00	92.00	46.00	69.00	96.00	76.31
15	64.40	59.10	33.90	30.70	99.00	96.00	91.00	91.00	82.00	96.00	91.00	99.00	91.00	43.00	71.00	96.00	77.07
Reasoning Tokens																	
1	1570.80	1694.56	1830.50	1876.89	1393.14	1792.57	1166.58	3701.48	7559.94	1788.95	1229.70	2680.54	9815.81	4531.56	1825.05	2361.40	2926.22
5	459.45	486.67	585.45	583.36	340.66	575.46	288.92	3178.64	3223.71	499.13	577.85	1206.04	4668.14	2527.85	512.15	637.53	1271.94
10	430.58	440.67	559.56	560.56	260.36	562.37	279.35	2929.45	2333.25	271.89	418.56	934.24	2959.47	1829.60	278.98	352.47	962.58
15	427.78	432.67	430.76	427.89	229.57	505.24	268.45	2735.75	1313.87	170.59	384.45	734.35	2184.48	1030.26	175.04	246.44	731.10
Output Tokens																	
1	430.80	449.54	451.67	458.56	242.07	264.18	239.73	742.87	8135.57	201.18	241.28	395.58	615.22	504.48	232.29	407.17	875.76
5	205.56	208.56	210.67	211.43	144.07	179.31	149.09	337.40	275.18	134.34	172.60	179.58	211.33	173.29	161.14	202.01	197.22
10	198.25	196.49	192.22	189.45	127.52	172.57	142.23	318.46	205.74	108.94	157.58	173.65	205.74	168.71	139.92	157.58	178.44
15	175.27	171.21	165.45	161.34	127.58	169.35	135.73	294.36	119.74	74.61	129.43	133.21	194.25	159.28	107.12	131.21	153.07

Table 1: Accuracies, reasoning tokens, and output tokens across 13 benchmarks grouped by task type. TI, TC, AI and AC represent Trigger Identification, Trigger Classification, Argument Identification and Argument Classification tasks on ACE’05.

All results reported with <i>DeepSeek-R1</i>																	
Batch Size	Info Extraction				Reasoning / QA						Structured / Scientific					Average	
	TI	TC	AI	AC	Penguins	GSM8K	StratQA	IIT	GPQA	Last	Epis.	ObjCnt	Math	BBEH	SciTab		Game24
Accuracies (%)																	
1	59.17	55.14	32.66	30.50	99.00	98.00	90.00	88.00	80.00	96.00	85.00	97.00	42.00	80.00	71.00	97.00	75.03
5	56.63	51.23	25.96	23.35	100.00	97.00	92.00	93.00	84.00	98.00	91.00	100.00	45.00	84.00	98.00	97.00	77.26
10	59.83	57.58	29.63	26.37	98.00	96.00	92.00	92.00	81.00	97.00	96.00	99.00	44.00	86.00	98.00	96.00	78.03
15	65.11	59.88	34.36	31.27	98.00	96.00	92.00	91.00	91.00	100.00	96.00	97.00	44.00	83.00	97.00	95.00	79.41
Reasoning Tokens																	
1	1470.80	1607.56	1817.50	1845.89	1310.14	1777.57	1079.58	3666.48	7489.94	1767.95	1148.70	2580.54	9766.81	5759.42	2184.45	2301.40	2973.42
5	412.45	440.67	537.45	499.36	301.66	550.46	248.92	3166.64	3148.71	426.13	553.85	1115.04	4588.14	3105.60	502.36	585.53	1261.44
10	335.58	417.67	504.56	506.56	205.36	562.37	254.35	2885.45	2266.25	206.89	401.56	893.24	2887.47	2235.22	243.90	310.47	944.81
15	387.78	335.67	387.76	335.89	165.57	432.24	217.45	2701.75	1233.87	137.59	334.45	657.35	2151.48	1216.98	162.20	171.44	689.34
Output Tokens																	
1	413.80	431.54	401.67	366.56	221.07	213.18	165.73	714.87	8073.57	142.18	154.28	384.58	532.22	452.39	163.05	381.17	825.74
5	151.56	133.56	150.67	106.07	112.31	127.09	274.40	213.36	258.18	56.34	159.60	153.58	144.33	122.68	118.77	155.01	152.34
10	171.25	156.49	122.22	166.45	59.64	153.57	117.23	255.46	120.46	8.94	102.29	77.94	168.74	143.43	61.18	95.58	123.80
15	139.27	159.21	142.45	64.34	79.58	116.35	83.73	213.36	65.74	3.61	34.43	109.45	159.25	135.36	20.92	77.21	100.27

Table 2: Accuracies, reasoning tokens, and output tokens across 13 benchmarks grouped by task type. TI, TC, AI and AC represent Trigger Identification, Trigger Classification, Argument Identification and Argument Classification tasks on ACE’05.

visual or perceptual reasoning (Penguins (Srivastava et al., 2023)), and challenging benchmarks such as IIT-JEE (Arora et al., 2023) and BBEH (Kazemi et al., 2025). We use 100 samples from each dataset, consistent with standard LLM evaluation practice (Zhong et al., 2024). Accuracy is measured via exact match or task-specific metrics, while efficiency is tracked using model-reported reasoning and output token counts. We conduct all experiments with DeepSeek-R1 (deepseek-reasoner, accessed via API) and OpenAI-o1 (o1-2024-12-17), using zero-shot templated prompting (see Appendix A for full prompt template). For both models, we use default API parameters. Batches are constructed by randomly sampling b queries from each dataset and grouping them into a single prompt; all queries within a batch come from the same dataset (homogeneous batching). We evaluate batch sizes $b \in \{1, 5, 10, 15\}$. Token counts (reasoning and output) are obtained from the API response metadata. If a model exhausts its token budget without producing a parseable answer, we score that instance as incorrect. We provide additional details in Appendix B.

Statistical Methodology. For token count and accuracy comparisons, we report means with 95% confidence intervals (CIs) computed as $\bar{x} \pm 1.96 \cdot \frac{s}{\sqrt{n}}$. Confidence intervals are calculated over pooled evaluation instances ($n = 1300$ per model per batch size; 13 datasets \times 100 samples each), since batching is applied uniformly and all metrics are defined at the instance level. Pooling across instances provides a stable estimate of the population-level effect of batching while preserving statistical validity. We observe that reasoning token reductions are consistent in direction across all 13 datasets and both models, indicating robustness beyond any single-dataset comparison. Full confidence interval tables and additional statistical details are provided in Appendix C.

4.1 MAIN RESULTS

Batching Improves Accuracy Across Tasks. We evaluate batch prompting across 13 diverse benchmarks spanning question answering, arithmetic reasoning, structured extraction, and perception-style tasks (Tables 1-2). Compared to single-query prompting (batch size = 1), larger batches (up to 15) generally improve, though some tasks show non-monotonic behavior. For instance, on OpenAI-o1, ACE’05 trigger identification accuracy increases from 58.4 to 64.4, while on tasks prone to overthinking (Srivastava et al., 2024; Rein et al., 2024), batch prompting yields notable boosts (e.g., GPQA: 79 \mapsto 83; Last Letter: 95 \mapsto 97 on o1). Averaged across both models, accuracy rises from 75.0% to 78.2% at batch size 15. We hypothesize this gain stems from both token efficiency and the inductive signal offered by seeing similar queries together, reinforcing the right reasoning patterns.

Batching Reduces Reasoning Cost with No Output Penalty. Beyond accuracy, batching substantially reduces reasoning tokens. Averaged across both models, reasoning cost drops from 2,950 to just 710 tokens (4.5 \times reduction), with a corresponding decrease in output tokens (851 \mapsto 127). Even in reasoning-heavy tasks like GSM8K or SciTab, batching reduces internal deliberation while preserving performance, consistent with the hypothesis that batching discourages excessive internal deliberation. This effect is especially prominent in tasks where models tend to hedge or overexplore (e.g., Game of 24, Epistemic QA). Together, these results demonstrate that batch prompting is a practical and general-purpose technique for improving LRM efficiency without model or system changes.

5 ANALYSIS

Batch Prompting Regularizes Overthinking. Despite their strength on multi-step tasks, LRMs frequently overthink even trivial problems, generating verbose and unstable reasoning traces. Table 3 (Part A) illustrates this phenomenon across three simple queries: a decimal comparison, a letter-counting task, and an arithmetic word problem. In each case, DeepSeek-R1 produces the correct answer early in its reasoning trace but then spirals into self-doubt (“But maybe the question is a trick?”), redundant rechecking (“Actually careful...”), and irrelevant tangents. We identify four distinct overthinking patterns: *hedging*, *rechecking*, *repetitive re-derivation*, and *irrelevant tangents*, that inflate token counts without improving answer quality.

Overthinking Hurts Accuracy, Not Just Efficiency. In extreme cases, overthinking causes outright failure. Table 3 (Part B) shows a last-letter concatenation query where DeepSeek-R1 derives the correct answer within the first few hundred tokens, but then re-derives, second-guesses, and explores tangential concerns for over 5,000 tokens until the API times out with *no answer generated*. This partly explains why single-query accuracy can fall below batched inference on certain tasks: overthinking can cause complete answer failures on instances the model is otherwise capable of solving. Batch prompting eliminates this failure mode: when the same query appears alongside similar examples (Table 3 Part B, bottom), the model answers all queries concisely without entering any reasoning spiral.

Batching Reduces Overall Token Usage.

Figure 3 decomposes per-query token consumption into reasoning and output components across batch sizes. Both models exhibit a monotonic decline in total tokens, but the savings are overwhelmingly driven by reasoning compression. For DeepSeek-R1, reasoning tokens drop from 2,973 (BS=1) to 689 (BS=15), a 76.8% reduction, while output tokens decline from 826 to 100 (87.9%). OpenAI-o1 follows the same trend: reasoning falls from 2,926 to 731 (75.0%), output from 876 to 153 (82.5%). In absolute terms, reasoning accounts for \sim 76% of the total token reduction at every

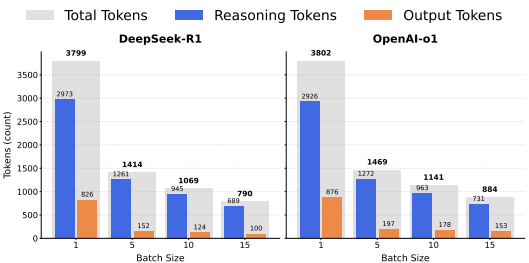


Figure 3: Token breakdown for DeepSeek-R1 and OpenAI-o1. Larger batch sizes reduce total tokens mainly by shortening reasoning traces, with minimal effect on output accuracies.

PART A: OVERTHINKING PATTERNS — STANDARD VS. BATCH PROMPTING	
Standard Prompting (BS = 1)	<p>► Q1: Which is larger, 0.9 or 0.11?</p> <p>▷ Reasoning: We need to compare 0.9 and 0.11. Obviously 0.9 is larger because (...) So $0.9 > 0.11$. But maybe the question is a trick? Sometimes people might (...) But no, it is decimal. So the answer is 0.9. (...)</p> <p>⇒ Output: 0.9 is larger than 0.11. When comparing decimals, 0.9 is equivalent to 0.90, which is greater than 0.11.</p>
	<p>► Q2: How many Rs in “Strawberry”?</p> <p>▷ Reasoning: We need to count the number of ‘R’s (...). Let’s spell it: S-t-r-a-w-b-e-r-r-r-y (...). Actually careful: The word is “Strawberry”. Typically (...) Actually strawberry: (...) However, lets double-check (...) Actually strawberry So just count the (...).</p> <p>⇒ Output: There are 5 Rs in “Strawberry”.</p>
	<p>► Q3: A fort has enough food for 100 soldiers for 120 days. 10 days later, 20 more soldiers join. How many more days will the food last?</p> <p>▷ Reasoning: We are given: “A fort has enough food (...) But the question (...) new soldiers join. But sometimes they want integer? (...) Perhaps they mean “how many more”? (...) But maybe they want “more days” (...) Lets read carefully “10 days later (...)</p> <p>⇒ Output: After 10 days, the food consumed by 100 soldiers is $100 \times 10 = 1000$ soldier-days (...)</p>
Batch Prompting (BS = 5)	<p>► Prompt: Solve these questions: Q1: How many Rs are in “Strawberry”? Q2: Fort food problem (same as above). Q3: Which is larger, 0.9 or 0.11? ...truncated...</p> <p>▷ Reasoning: We need to answer three questions (...) “Strawberry” — “Straw” + “berry” (...) Actually strawberry - typical spelling (...)</p> <p>▷ But careful: “Strawberry” - it might (...) So answer: 5. Q2: A fort has (...) But perhaps they (...); Q3: 0.9 or 0.11? (...).</p> <p>⇒ Output: Q1: 5 Rs. Q2: Total food = $100 \times 120 = 12,000$ soldier-days (...). Q3: 0.9 is larger. ...truncated...</p>
PART B: OVERTHINKING-INDUCED API TIMEOUT	
Standard Prompting (BS = 1)	<p>► Q: Take the last letters of each word in “Oswaldo José Luis Sheldon Tara” and concatenate them.</p> <p>▷ Reasoning: “Oswaldo”: O-s-w-a-l-d-o. Last letter is “o”. But is it lowercase or uppercase? For concatenation, case might matter (...) “José”: J-o-s-é. Last character is ‘é’. In string processing, ‘é’ is a single Unicode character (...) Perhaps I should confirm the words. For each word: Oswaldo→‘o’, José→‘é’, Luis→‘s’, Sheldon→‘n’, Tara→‘a’.</p> <p>— correct answer “oésna” first derived here —</p> <p>In some contexts, people might consider the base letter (...) Let me double-check. “Oswaldo”: O,s,w,a,l,d,o → ‘o’; “José”: J,o,s,é → ‘é’ (U+00E9). In Python: words=[...] I think for this context, it’s fine. (...) but the word is given (...)</p> <p>continues for 5,000+ reasoning tokens</p> <p>⇒ API Timeout — No answer generated.</p>
Batch Prompting (BS = 5)	<p>► Q: Solve these questions: Q1: Take the last letters of each words in “Oswaldo José Luis Sheldon Tara” and concatenate them. Q2: Take the last letters of each words in “Alec Arianna Corina JuanCarlos” and concatenate them. Q3: Take the last letters of each words in “Leandro Maricela Genevieve Lesly” and concatenate them. ...truncated...</p> <p>▷ Reasoning: We need to solve three questions. Each asks to take the last letters (...) Words Oswaldo, José, Luis, Sheldon, Tara (...). But we need to consider if they want it as a single string. (...) Words: Alec, Arianna, Corina, and JuanCarlos. (...) Words: Leandro, Maricela, Genevieve, Lesly (...) But we need to be careful: The phrase in Q1 (...) Thus final answer Q1: oésna; Q2: caas; Q3: oaey. ...truncated...</p> <p>⇒ Output: Q1: oésna; Q2: caas; Q3: oaey ...truncated...</p>

Legend: Self-doubt / hedging Redundant rechecking Repetitive re-derivation Irrelevant tangents

Table 3: Qualitative comparison of reasoning behaviour in DeepSeek-R1 under standard vs. batch prompting. **Part A:** Single-query prompting triggers four distinct overthinking patterns (see legend); batch prompting produces concise reasoning for the *same* queries with no model modification. **Part B:** In extreme cases, the model derives the correct answer early but continues to spiral through repetition, tangents, and self-doubt until the API times out—directly explaining lower accuracy at BS = 1. Batched examples are *truncated* showing only 3 examples to save space.

batch size, confirming that batching primarily compresses the model’s internal deliberation rather than truncating its answers. The resulting total cost drops from $\sim 3,800$ tokens per query (BS=1) to ~ 840 (BS=15), a $4.5\times$ reduction, while accuracy improves (Tables 1–2). While both components decrease substantially, reasoning tokens dominate the absolute savings because they constitute the majority of tokens at baseline ($\sim 78\%$ of total). This is consistent with the view that batch prompting primarily compresses the model’s internal deliberation, though we note that output tokens also decrease significantly (83–88%).

Batching Encourages Pattern Induction. Beyond regularizing verbosity, batching structurally similar examples appears to facilitate inductive generalization. In isolation, models often struggle with formatting-sensitive tasks, e.g., when asked to extract and concatenate the last letters of words, DeepSeek-R1 produces the correct reasoning but outputs “nn y0” (with a spurious space), failing under exact match. However, when the same example is placed in a batch of similar queries, the

model not only solves each instance correctly but also aligns to a more consistent output schema. This suggests that batched inference enables in-context pattern discovery, allowing the model to leverage structural similarity and prior completions to guide subsequent reasoning. In effect, later examples benefit from earlier ones, even without gradient updates, supporting the view that batches can serve as implicit demonstrations and reasoning anchors.

Batching Suppresses Hedging and Metacognitive Loops.

One of the hallmarks of overthinking in LRMs is the emergence of verbose, recursive self-dialog, often marked by tokens like “wait,” “let me double-check,” or “hold on.” We quantify this by measuring the normalized frequency of four overthinking patterns across baseline (BS=1) and batched (BS=5) inference: *hedging*, *rechecking*, *re-derivation*, and *irrelevant tangents*. As shown in Figure 4, batching uniformly suppresses all four patterns across both models. DeepSeek-R1 exhibits the highest baseline frequencies (hedging: 0.42, rechecking: 0.55, re-derivation: 0.38, tangents: 0.31), which drop to 0.18, 0.22, 0.16, and 0.12, respectively, under batching, a 59% reduction on average.

OpenAI-o1 shows a similar pattern: baseline frequencies of 0.33, 0.47, 0.29, and 0.24 fall to 0.15, 0.19, 0.13, and 0.11, representing a 56% average reduction. Rechecking behavior is most prevalent at baseline and also most effectively suppressed, dropping by 60% for DeepSeek-R1 and 60% for OpenAI-o1. For instance, on the Game of 24 task, single-query inference with DeepSeek-R1 produced 21 instances of the token “wait” across five queries, reflecting repeated self-corrections; the same five questions in a batch yielded only one occurrence, focused solely on the most difficult query. This suggests that batching acts as a soft interrupt mechanism: by distributing attention across multiple queries, the model avoids getting *stuck* in internal loops, implicitly curtailing metacognitive wandering.

Explicit Prompt Constraints Also Fail.

A natural baseline is to ask whether simple prompt-level constraints can induce comparable reductions in overthinking. We evaluate two such interventions, explicit token limits (e.g., “use no more than 100 tokens for reasoning”) and time limits (e.g., “think for at most 10 seconds”), on GAME OF 24 and LAST LETTER. As shown in Table 4, these approaches yield limited and inconsistent gains. Models frequently acknowledge the constraint in natural language but continue to generate long internal reasoning traces. For instance, on GAME OF 24, DeepSeek-R1 under an explicit token limit still produces 1,256 reasoning tokens on average, corresponding to only a 45% reduction relative to the baseline (2,301 tokens), far below the 93% reduction achieved by batch prompting. OpenAI-o1 is even less responsive: token constraints reduce reasoning tokens by only 17% while substantially degrading accuracy (96% → 45%).

These results highlight a key limitation of instructional constraints: they rely on the model to voluntarily comply, and can be ignored or superficially satisfied without materially reducing computation. In contrast, batch prompting creates a *behavioral* constraint; the shared context appears to induce shorter reasoning traces without requiring explicit compliance, yielding large reductions in reasoning tokens while preserving, or even improving, task accuracy.

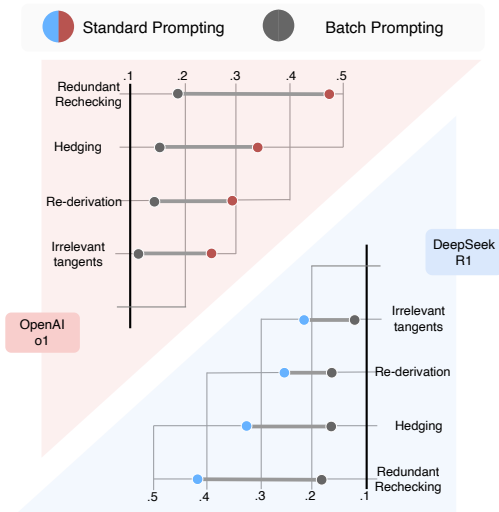


Figure 4: Normalized frequency of four overthinking patterns under baseline (BS=1) vs. batch prompting (BS=5) for DeepSeek-R1 and OpenAI-o1. Batching uniformly suppresses all pattern types, with the largest reductions in rechecking and re-derivation.

6 IMPLICATIONS FOR FUTURE RESEARCH

Our findings that batch prompting compresses reasoning traces while preserving accuracy have implications beyond the evaluation setting studied in this paper. We outline several concrete directions where batched inference could yield practical gains.

Compact Chain-of-Thought Generation for Distillation. A major bottleneck in chain-of-thought (CoT) distillation (Hsieh et al., 2023) is that teacher models produce verbose reasoning traces, making the distillation data expensive to generate and harder for student models to learn from. Our results show that batch prompting produces shorter yet equally accurate reasoning traces (76% fewer tokens at BS=15 with no accuracy loss). These compressed traces could serve as higher-quality supervision for training smaller models: the student receives the same correct reasoning signal in a fraction of the tokens, reducing both the cost of generating distillation data and the sequence length the student must learn to produce. Concretely, distilling from batched DeepSeek-R1 traces (averaging 689 reasoning tokens per query at BS=15) rather than single-query traces (2,973 tokens) would reduce the teacher-generation cost by $\sim 4.5\times$ while providing more focused, less noisy supervision.

Batched Prompt Optimization. In prompt optimization frameworks such as DSPy (Khattab et al., 2024) and OPRO (Yang et al., 2024), an LLM evaluates candidate prompts by running inference on individual examples and using the outputs as a proxy gradient signal to update the prompt. This single-instance evaluation is both costly and noisy: each gradient step requires a separate API call, and the feedback from a single example may not generalize. Batch prompting offers a natural improvement: by evaluating a candidate prompt on a batch of b examples in a single call, the optimizer receives b data points of feedback simultaneously. This not only reduces the per-example API cost (as shown in our experiments), but may also yield more informative gradient signals, observing a prompt’s behavior across multiple examples in one pass exposes systematic failure patterns that single-instance evaluation would miss. Recent work on prompt optimization with reasoning models (Srivastava & Yao, 2025) has highlighted the cost barrier of iterative prompt refinement with LRMs; batched evaluation directly addresses this by reducing the number of API calls needed per optimization step by a factor of b .

Cost-Efficient Agentic Workflows. Multi-step agentic systems (Yao et al., 2023b; Shinn et al., 2023) invoke reasoning models repeatedly, for planning, tool selection, self-reflection, and error correction, making reasoning token costs a dominant expense. When an agent must process multiple independent sub-tasks (e.g., evaluating several candidate actions, verifying multiple facts, or generating plans for parallel sub-goals), these

Dataset	Method	Accuracy	Avg. Tokens
DeepSeek-R1			
Game of 24	Baseline (BS=1)	97.0%	2301
	+ Token Limit	85.0% $\downarrow 12$	1256 $\downarrow 45\%$
	+ Time Limit	96.0% $\downarrow 1$	1424 $\downarrow 38\%$
	Batch (BS=5)	97.0% \approx	586 $\downarrow 75\%$
	Batch (BS=10)	96.0% $\downarrow 1$	310 $\downarrow 87\%$
	Batch (BS=15)	95.0% $\downarrow 2$	171 $\downarrow 93\%$
Last Letter	Baseline (BS=1)	96.0%	1768
	+ Token Limit	92.0% $\downarrow 4$	852 $\downarrow 52\%$
	+ Time Limit	94.0% $\downarrow 2$	974 $\downarrow 45\%$
	Batch (BS=5)	98.0% $\uparrow 2$	426 $\downarrow 76\%$
	Batch (BS=10)	97.0% $\uparrow 1$	207 $\downarrow 88\%$
	Batch (BS=15)	100.0% $\uparrow 4$	138 $\downarrow 92\%$
OpenAI-o1			
Game of 24	Baseline (BS=1)	96.0%	2361
	+ Token Limit	45.0% $\downarrow 51$	1952 $\downarrow 17\%$
	+ Time Limit	89.0% $\downarrow 7$	2125 $\downarrow 10\%$
	Batch (BS=5)	97.0% $\uparrow 1$	638 $\downarrow 73\%$
	Batch (BS=10)	96.0% \approx	352 $\downarrow 85\%$
	Batch (BS=15)	96.0% \approx	246 $\downarrow 90\%$
Last Letter	Baseline (BS=1)	95.0%	1789
	+ Token Limit	77.0% $\downarrow 18$	1607 $\downarrow 10\%$
	+ Time Limit	92.0% $\downarrow 3$	1680 $\downarrow 6\%$
	Batch (BS=5)	97.0% $\uparrow 2$	499 $\downarrow 72\%$
	Batch (BS=10)	97.0% $\uparrow 2$	272 $\downarrow 85\%$
	Batch (BS=15)	96.0% $\uparrow 1$	171 $\downarrow 90\%$

Table 4: Comparison of prompt-based solutions for reducing overthinking. Explicit constraints modestly reduce tokens but *can hurt accuracy* (red arrows), while batch prompting achieves large token reduction while *preserving or improving accuracy* (green/gray arrows). Subscripts show change vs. baseline (BS=1).

sub-tasks can be batched into a single prompt. Our finding that batching reduces reasoning tokens by $\sim 4.5\times$ without accuracy loss suggests that batched sub-task processing could substantially reduce the cost of agentic pipelines. For instance, a ReAct-style agent (Yao et al., 2023b) that must evaluate 5 candidate tool calls could batch these evaluations into a single prompt, reducing reasoning cost while obtaining all evaluations simultaneously.

Test-Time Compute Allocation. Recent work has shown that scaling test-time compute, allocating more inference-time reasoning, can be more effective than scaling model parameters (Snell et al., 2024). However, current test-time scaling strategies treat each query independently, allocating compute without cross-query awareness. Our observation that batching induces adaptive token allocation (Section 3.2), where models spend more tokens on harder queries and fewer on easier ones within a batch, suggests a complementary approach: batch prompting could serve as a lightweight mechanism for adaptive test-time compute allocation across a workload, without requiring explicit difficulty classifiers or routing logic. Combining batching with explicit compute-allocation strategies (e.g., adaptive batch sizes based on estimated query difficulty) is a promising direction for efficient deployment of reasoning models at scale.

7 CONCLUSION

We present an empirical study demonstrating that batch prompting effectively suppresses overthinking in Large Reasoning Models at inference time. Across 13 diverse benchmarks with DeepSeek-R1 and OpenAI-o1, batching reduces reasoning tokens by 76% while maintaining or improving accuracy, a $4\times$ efficiency gain with no model modification. Through behavioral analysis, we find that batching suppresses hedging and rechecking patterns, enables pattern induction from earlier examples, and adaptively allocates tokens to harder queries.

We also find that explicit prompt constraints (token/time limits) fail to reliably reduce reasoning tokens while often hurting accuracy, highlighting the advantage of batching’s implicit approach. These findings reframe batch prompting as more than a throughput optimization: it is a practical, model-agnostic tool for efficient LRM deployment. Understanding the precise mechanisms behind batch-induced reasoning compression remains an open direction for future work.

8 LIMITATIONS

Our work has several limitations. First, we focus on batches of *independent* queries; dependent queries (e.g., multi-turn dialogues) may benefit less from batching as models must maintain state across queries. Second, extreme batch sizes or highly heterogeneous query mixtures may degrade coherence; our experiments use semantically homogeneous batches. Third, while we use 100 samples per dataset (a standard for LLM evaluation), larger-scale validation would strengthen statistical claims; the consistency across 13 datasets partially addresses this. Finally, batching lacks explicit control over token budgets; future work should explore adaptive methods combining batching with dynamic reasoning control (Srivastava et al., 2024; 2025), as well as extensions to multi-modal settings (Dey et al., 2025).

REFERENCES

- Daman Arora, Himanshu Singh, and Mausam. Have LLMs advanced enough? a challenging problem solving benchmark for large language models. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 7527–7543, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.468. URL <https://aclanthology.org/2023.emnlp-main.468/>.
- Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong Yu. Do not think that much for $2+3=?$ on the overthinking of o1-like llms, 2025. URL <https://arxiv.org/abs/2412.21187>.

- 540 Zhoujun Cheng, Jungo Kasai, and Tao Yu. Batch prompting: Efficient inference with large lan-
541 guage model APIs. In Mingxuan Wang and Imed Zitouni (eds.), *Proceedings of the 2023 Con-*
542 *ference on Empirical Methods in Natural Language Processing: Industry Track*, pp. 792–810,
543 Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.
544 emnlp-industry.74. URL [https://aclanthology.org/2023.emnlp-industry.](https://aclanthology.org/2023.emnlp-industry.74/)
545 [74/](https://aclanthology.org/2023.emnlp-industry.74/).
- 546 Alejandro Cuadron, Dacheng Li, Wenjie Ma, Xingyao Wang, Yichuan Wang, Siyuan Zhuang, Shu
547 Liu, Luis Gaspar Schroeder, Tian Xia, Huanzhi Mao, Nicholas Thumiger, Aditya Desai, Ion
548 Stoica, Ana Klimovic, Graham Neubig, and Joseph E. Gonzalez. The danger of overthinking:
549 Examining the reasoning-action dilemma in agentic tasks, 2025. URL [https://arxiv.org/](https://arxiv.org/abs/2502.08235)
550 [abs/2502.08235](https://arxiv.org/abs/2502.08235).
- 551 Abhishek Dey, Aabha Bothera, Samhita Sarikonda, Rishav Aryan, Sanjay Kumar Podishetty, Ak-
552 shay Havalgi, Saurabh Srivastava, and Gaurav Singh. Multimodal event detection: Current ap-
553 proaches and defining the new playground through llms and vlms. In *International Conference*
554 *on Applications of Natural Language to Information Systems*, pp. 457–471. Springer, 2025.
- 555 George R Doddington, Alexis Mitchell, Mark A Przybocki, Lance A Ramshaw, Stephanie M
556 Strassel, Ralph M Weischedel, et al. The automatic content extraction (ace) program-tasks, data,
557 and evaluation. In *LREC*, volume 2. Lisbon, 2004.
- 558 Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. Did aristotle
559 use a laptop? a question answering benchmark with implicit reasoning strategies. *Transactions of*
560 *the Association for Computational Linguistics*, 9:346–361, 2021.
- 561 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
562 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
563 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- 564 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
565 and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In
566 *Advances in Neural Information Processing Systems*, volume 34, 2021.
- 567 Cheng-Yu Hsieh, Chun-Liang Li, Chih-Kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alexander Rat-
568 ner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. Distilling step-by-step! Outperforming
569 larger language models with less training data and smaller model sizes. In *Findings of the Asso-*
570 *ciation for Computational Linguistics: ACL 2023*, pp. 8003–8017, 2023.
- 571 Yao Huang, Huanran Chen, Shouwei Ruan, Yichi Zhang, Xingxing Wei, and Yinpeng Dong. Mit-
572 igating overthinking in large reasoning models via manifold steering, 2025. URL [https:](https://arxiv.org/abs/2505.22411)
573 [//arxiv.org/abs/2505.22411](https://arxiv.org/abs/2505.22411).
- 574 Mehran Kazemi, Bahare Fatemi, Hritik Bansal, John Palowitch, Chrysovalantis Anastasiou, San-
575 ket Vaibhav Mehta, Lalit K Jain, Virginia Aglietti, Disha Jindal, Yuanzhu Peter Chen, et al. Big-
576 bench extra hard. In *Proceedings of the 63rd Annual Meeting of the Association for Computational*
577 *Linguistics (Volume 1: Long Papers)*, pp. 26473–26501, 2025.
- 578 Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vard-
579 hamanan, Saiful Haq, Ashutosh Sharma, Thomas T Joshi, Hanna Mober, et al. DSPy: Compiling
580 declarative language model calls into self-improving pipelines. In *International Conference on*
581 *Learning Representations*, 2024.
- 582 Xinyuan Lu, Liangming Pan, Qian Liu, Preslav Nakov, and Min-Yen Kan. Scitab: A challenging
583 benchmark for compositional reasoning and claim verification on scientific tables. *arXiv preprint*
584 *arXiv:2305.13186*, 2023.
- 585 Iman Mirzadeh, Keivan Alizadeh, Hooman Shahrokhi, Oncel Tuzel, Samy Bengio, and Mehrdad
586 Farajtabar. Gsm-symbolic: Understanding the limitations of mathematical reasoning in large
587 language models. *arXiv preprint arXiv:2410.05229*, 2024.
- 588 David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Di-
589 rani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a bench-
590 mark. In *First Conference on Language Modeling*, 2024.
- 591

- 594 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion:
595 Language agents with verbal reinforcement learning. *Advances in Neural Information Processing*
596 *Systems*, 36, 2023.
- 597 Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling LLM test-time compute opti-
598 mally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*,
599 2024.
- 600
- 601 Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam
602 Fisch, Adam R. Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, Agnieszka Kluska,
603 Aitor Lewkowycz, Akshat Agarwal, Alethea Power, Alex Ray, Alex Warstadt, Alexander W.
604 Kocurek, Ali Safaya, Ali Tazarv, Alice Xiang, Alicia Parrish, Allen Nie, Aman Hussain,
605 Amanda Askeel, Amanda Dsouza, Ambrose Slone, Ameet Rahane, Anantharaman S. Iyer, An-
606 ders Andreassen, Andrea Madotto, Andrea Santilli, Andreas Stuhlmüller, Andrew Dai, An-
607 drew La, Andrew Lampinen, Andy Zou, Angela Jiang, Angelica Chen, Anh Vuong, Animesh
608 Gupta, Anna Gottardi, Antonio Norelli, Anu Venkatesh, Arash Gholamidavoodi, Arfa Tabas-
609 sum, Arul Menezes, Arun Kirubarajan, Asher Mullokandov, Ashish Sabharwal, Austin Her-
610 rick, Avia Efrat, Aykut Erdem, Ayla Karakaş, B. Ryan Roberts, Bao Sheng Loe, Barret Zoph,
611 Bartłomiej Bojanowski, Batuhan Özyurt, Behnam Hedayatnia, Behnam Neyshabur, Benjamin
612 Inden, Benno Stein, Berk Ekmekci, Bill Yuchen Lin, Blake Howald, Bryan Orinion, Cameron
613 Diao, Cameron Dour, Catherine Stinson, Cedrick Argueta, César Ferri Ramírez, Chandan Singh,
614 Charles Rathkopf, Chenlin Meng, Chitta Baral, Chiyu Wu, Chris Callison-Burch, Chris Waites,
615 Christian Voigt, Christopher D. Manning, Christopher Potts, Cindy Ramirez, Clara E. Rivera,
616 Clemencia Siro, Colin Raffel, Courtney Ashcraft, Cristina Garbacea, Damien Sileo, Dan Gar-
617 rette, Dan Hendrycks, Dan Kilman, Dan Roth, Daniel Freeman, Daniel Khashabi, Daniel Levy,
618 Daniel Moseguí González, Danielle Perszyk, Danny Hernandez, Danqi Chen, Daphne Ippolito,
619 Dar Gilboa, David Dohan, David Drakard, David Jurgens, Debajyoti Datta, Deep Ganguli, De-
620 nis Emelin, Denis Kleyko, Deniz Yuret, Derek Chen, Derek Tam, Dieuweke Hupkes, Diganta
621 Misra, Dilyar Buzan, Dimitri Coelho Mollo, Diyi Yang, Dong-Ho Lee, Dylan Schrader, Eka-
622 terina Shutova, Ekin Dogus Cubuk, Elad Segal, Eleanor Hagerman, Elizabeth Barnes, Eliza-
623 beth Donoway, Ellie Pavlick, Emanuele Rodola, Emma Lam, Eric Chu, Eric Tang, Erkut Erdem,
624 Ernie Chang, Ethan A. Chi, Ethan Dyer, Ethan Jerzak, Ethan Kim, Eunice Engefu Manyasi, Ev-
625 genii Zheltonozhskii, Fanyue Xia, Fatemeh Siar, Fernando Martínez-Plumed, Francesca Happé,
626 Francois Chollet, Frieda Rong, Gaurav Mishra, Genta Indra Winata, Gerard de Melo, Germán
627 Kruszewski, Giambattista Parascandolo, Giorgio Mariani, Gloria Wang, Gonzalo Jaimovitch-
628 López, Gregor Betz, Guy Gur-Ari, Hana Galijasevic, Hannah Kim, Hannah Rashkin, Hannaneh
629 Hajishirzi, Harsh Mehta, Hayden Bogar, Henry Shevlin, Hinrich Schütze, Hiromu Yakura, Hong-
630 ming Zhang, Hugh Mee Wong, Ian Ng, Isaac Noble, Jaap Jumelet, Jack Geissinger, Jackson
631 Kernion, Jacob Hilton, Jaehoon Lee, Jaime Fernández Fisac, James B. Simon, James Koppel,
632 James Zheng, James Zou, Jan Kocoń, Jana Thompson, Janelle Wingfield, Jared Kaplan, Jarema
633 Radom, Jascha Sohl-Dickstein, Jason Phang, Jason Wei, Jason Yosinski, Jekaterina Novikova,
634 Jelle Bosscher, Jennifer Marsh, Jeremy Kim, Jeroen Taal, Jesse Engel, Jesujoba Alabi, Ji-
635 acheng Xu, Jiaming Song, Jillian Tang, Joan Waweru, John Burden, John Miller, John U. Balis,
636 Jonathan Batchelder, Jonathan Berant, Jörg Frohberg, Jos Rozen, Jose Hernandez-Orallo, Joseph
637 Boudeman, Joseph Guerr, Joseph Jones, Joshua B. Tenenbaum, Joshua S. Rule, Joyce Chua,
638 Kamil Kanclerz, Karen Livescu, Karl Krauth, Karthik Gopalakrishnan, Katerina Ignatyeva, Katja
639 Markert, Kaustubh D. Dhole, Kevin Gimpel, Kevin Omondi, Kory Mathewson, Kristen Chia-
640 fullo, Ksenia Shkaruta, Kumar Shridhar, Kyle McDonell, Kyle Richardson, Laria Reynolds, Leo
641 Gao, Li Zhang, Liam Dugan, Lianhui Qin, Lidia Contreras-Ochando, Louis-Philippe Morency,
642 Luca Moschella, Lucas Lam, Lucy Noble, Ludwig Schmidt, Luheng He, Luis Oliveros Colón,
643 Luke Metz, Lütfi Kerem Şenel, Maarten Bosma, Maarten Sap, Maartje ter Hoeve, Maheen Fa-
644 rooqi, Manaal Faruqui, Mantas Mazeika, Marco Baturan, Marco Marelli, Marco Maru, Maria
645 Jose Ramírez Quintana, Marie Tolkiehn, Mario Giulianelli, Martha Lewis, Martin Potthast,
646 Matthew L. Leavitt, Matthias Hagen, Mátyás Schubert, Medina Orduna Baitemirova, Melody
647 Arnaud, Melvin McElrath, Michael A. Yee, Michael Cohen, Michael Gu, Michael Ivanitskiy,
648 Michael Starritt, Michael Strube, Michał Swedrowski, Michele Bevilacqua, Michihiro Yasunaga,
649 Mihir Kale, Mike Cain, Mimeo Xu, Mirac Suzgun, Mitch Walker, Mo Tiwari, Mohit Bansal,
650 Moin Aminnaseri, Mor Geva, Mozhddeh Gheini, Mukund Varma T, Nanyun Peng, Nathan A.
651 Chi, Nayeon Lee, Neta Gur-Ari Krakover, Nicholas Cameron, Nicholas Roberts, Nick Doiron,
652 Nicole Martinez, Nikita Nangia, Niklas Deckers, Niklas Muennighoff, Nitish Shirish Keskar,

- 648 Niveditha S. Iyer, Noah Constant, Noah Fiedel, Nuan Wen, Oliver Zhang, Omar Agha, Omar El-
649 baghdadi, Omer Levy, Owain Evans, Pablo Antonio Moreno Casares, Parth Doshi, Pascale Fung,
650 Paul Pu Liang, Paul Vicol, Pegah Alipoormolabashi, Peiyuan Liao, Percy Liang, Peter Chang, Pe-
651 ter Eckersley, Phu Mon Htut, Pinyu Hwang, Piotr Miłkowski, Piyush Patil, Pouya Pezeshkpour,
652 Priti Oli, Qiaozhu Mei, Qing Lyu, Qinlang Chen, Rabin Banjade, Rachel Etta Rudolph, Raefer
653 Gabriel, Rahel Habacker, Ramon Risco, Raphaël Millière, Rhythm Garg, Richard Barnes, Rif A.
654 Saurous, Riku Arakawa, Robbe Raymaekers, Robert Frank, Rohan Sikand, Roman Novak, Ro-
655 man Sitelew, Ronan LeBras, Rosanne Liu, Rowan Jacobs, Rui Zhang, Ruslan Salakhutdinov,
656 Ryan Chi, Ryan Lee, Ryan Stovall, Ryan Teehan, Rylan Yang, Sahib Singh, Saif M. Moham-
657 mad, Sajant Anand, Sam Dillavou, Sam Shleifer, Sam Wiseman, Samuel Gruetter, Samuel R.
658 Bowman, Samuel S. Schoenholz, Sanghyun Han, Sanjeev Kwatra, Sarah A. Rous, Sarik Ghaz-
659 arian, Sayan Ghosh, Sean Casey, Sebastian Bischoff, Sebastian Gehrmann, Sebastian Schus-
660 ter, Sepideh Sadeghi, Shadi Hamdan, Sharon Zhou, Shashank Srivastava, Sherry Shi, Shikhar
661 Singh, Shima Asaadi, Shixiang Shane Gu, Shubh Pachchigar, Shubham Toshniwal, Shyam Upad-
662 hyay, Shyamolima, Debnath, Siamak Shakeri, Simon Thormeyer, Simone Melzi, Siva Reddy,
663 Sneha Priscilla Makini, Soo-Hwan Lee, Spencer Torene, Sriharsha Hatwar, Stanislas Dehaene,
664 Stefan Divic, Stefano Ermon, Stella Biderman, Stephanie Lin, Stephen Prasad, Steven T. Pianta-
665 dosi, Stuart M. Shieber, Summer Mishnerghi, Svetlana Kiritchenko, Swaroop Mishra, Tal Linzen,
666 Tal Schuster, Tao Li, Tao Yu, Tariq Ali, Tatsu Hashimoto, Te-Lin Wu, Théo Desbordes, Theodore
667 Rothschild, Thomas Phan, Tianle Wang, Tiberius Nkinyili, Timo Schick, Timofei Kornev, Ti-
668 tus Tunduny, Tobias Gerstenberg, Trenton Chang, Trishala Neeraj, Tushar Khot, Tyler Shultz,
669 Uri Shaham, Vedant Misra, Vera Demberg, Victoria Nyamai, Vikas Raunak, Vinay Ramasesh,
670 Vinay Uday Prabhu, Vishakh Padmakumar, Vivek Srikumar, William Fedus, William Saun-
671 ders, William Zhang, Wout Vossen, Xiang Ren, Xiaoyu Tong, Xinran Zhao, Xinyi Wu, Xudong
672 Shen, Yadollah Yaghoobzadeh, Yair Lakretz, Yangqiu Song, Yasaman Bahri, Yejin Choi, Yichi
673 Yang, Yiding Hao, Yifu Chen, Yonatan Belinkov, Yu Hou, Yufang Hou, Yuntao Bai, Zachary
674 Seid, Zhuoye Zhao, Zijian Wang, Zijie J. Wang, Zirui Wang, and Ziyi Wu. Beyond the im-
675 itation game: Quantifying and extrapolating the capabilities of language models, 2023. URL
676 <https://arxiv.org/abs/2206.04615>.
- 675 Saurabh Srivastava and Ziyu Yao. Revisiting prompt optimization with large reasoning models-a
676 case study on event extraction. *arXiv preprint arXiv:2504.07357*, 2025.
- 677 Saurabh Srivastava, Chengyue Huang, Weiguo Fan, and Ziyu Yao. Instances need more care:
678 Rewriting prompts for instances with LLMs in the loop yields better zero-shot performance. In
679 Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for Com-
680 putational Linguistics: ACL 2024*, pp. 6211–6232, Bangkok, Thailand, August 2024. Associ-
681 ation for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.371. URL <https://aclanthology.org/2024.findings-acl.371/>.
- 682 Saurabh Srivastava, Sweta Pati, and Ziyu Yao. Instruction-tuning LLMs for event extraction with an-
683 notation guidelines. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher
684 Pilehvar (eds.), *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 13055–
685 13071, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-
686 89176-256-5. doi: 10.18653/v1/2025.findings-acl.677. URL <https://aclanthology.org/2025.findings-acl.677/>.
- 687 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-
688 ury, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models.
689 In *International Conference on Learning Representations*, 2023.
- 690 Zihao Wei, Liang Pang, Jiahao Liu, Wenjie Shi, Jingcheng Deng, Shicheng Xu, Zenghao Duan, Fei
691 Sun, Huawei Shen, and Xueqi Cheng. The evolution of thought: Tracking llm overthinking via
692 reasoning dynamics analysis, 2025. URL <https://arxiv.org/abs/2508.17627>.
- 693 Kaiwen Yan, Xuanqing Shi, Hongcheng Guo, Wenxuan Wang, Zhuosheng Zhang, and Chengwei
694 Qin. Drqa: Dynamic reasoning quota allocation for controlling overthinking in reasoning large
695 language models, 2025. URL <https://arxiv.org/abs/2508.17803>.
- 696 Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V Le, Denny Zhou, and Xinyun Chen.
697 Large language models as optimizers. In *International Conference on Learning Representations*,
698 2024.

702 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik
703 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Ad-
704 vances in Neural Information Processing Systems*, 36, 2023a.

706 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao.
707 ReAct: Synergizing reasoning and acting in language models. In *International Conference on
708 Learning Representations*, 2023b.

710 Tianyang Zhong, Zhengliang Liu, Yi Pan, Yutong Zhang, Yifan Zhou, Shizhe Liang, Zihao Wu,
711 Yanjun Lyu, Peng Shu, Xiaowei Yu, et al. Evaluation of openai o1: Opportunities and challenges
712 of agi. *arXiv preprint arXiv:2409.18486*, 2024.

714 Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuur-
715 mans, Claire Cui, Olivier Bousquet, Quoc Le, et al. Least-to-most prompting enables complex
716 reasoning in large language models. In *International Conference on Learning Representations*,
717 2023.

719 A PROMPT TEMPLATE

Prompt Template

723 You are a highly intelligent and capable AI assistant. Your job is to solve the user queries given
724 below.

725 Answer Format: {answer_format}

726 You can reason about all your questions together in <reasoning>...</reasoning>
727 block.

728 Then, for each question, your job is to solve each question and provide your response in the
729 following format: <answers>

```
730 [
731   {{"example1": {"reason": "Let's think step by step...", "answer": "answer1"}},
732   {{"example2": {"reason": "Let's think step by step...", "answer": "answer2"}},
733   ...
734   {{"exampleN": {"reason": "Let's think step by step...", "answer": "answerN"}}}
735 ]
```

736 </answers>

737 For each example, start with "Let's think step by step. ..." in the reason field. For each answer,
738 provide the final answer in the answer field and nothing else.

739 Here are the questions you need to answer: {batch_prompt}

740 —> Remember, you have to solve {num_questions} questions in total so plan in advance
741 to answer cheaply and efficiently. {constrain} <—

Constraints

```
744 time constrain: "Think for only 10 seconds."
745 token constrain: "Use no more than 100 tokens in
746 thinking."
```

749 B EXPERIMENTAL DETAILS

752 **Models and API Configuration.** We use two reasoning-oriented LLMs: (1) **DeepSeek-R1**
753 (model ID: deepseek-reasoner), accessed via the DeepSeek API, and (2) **OpenAI-o1** (model
754 ID: o1-2024-12-17), accessed via the OpenAI API. Both models expose reasoning (thinking)
755 tokens separately from output tokens in their API responses. Temperature and top-p are not user-
configurable for these reasoning endpoints; we use default settings throughout.

Model	Batch Size	Accuracy (%) \uparrow	Reasoning Tokens \downarrow	Total Tokens \downarrow
DeepSeek-R1	1	75.0 \pm 2.1	2973 \pm 190	3799 \pm 220
	5	77.3 \pm 1.9	1261 \pm 110	1414 \pm 130
	10	78.0 \pm 1.8	945 \pm 75	1069 \pm 85
	15	79.4 \pm 1.7	689 \pm 55	790 \pm 65
OpenAI-o1	1	75.0 \pm 2.2	2926 \pm 210	3802 \pm 230
	5	75.5 \pm 2.0	1272 \pm 120	1469 \pm 140
	10	76.3 \pm 1.9	963 \pm 80	1141 \pm 95
	15	77.1 \pm 1.8	731 \pm 60	884 \pm 70

Table 5: Mean \pm 95% confidence intervals for accuracy, reasoning tokens, and total tokens across batch sizes. Confidence intervals are computed over pooled evaluation instances ($n = 1300$ per model per batch size). Batch prompting yields statistically significant reductions in reasoning tokens (paired two-sided t-test, $p < 0.01$) while preserving or modestly improving accuracy.

Batch Construction. For each dataset, we randomly sample 100 instances and partition them into batches of size $b \in \{1, 5, 10, 15\}$. All queries within a batch are drawn from the same dataset (homogeneous batching). The order of queries within each batch is randomized. For batch size 1, we run standard single-query inference as the baseline.

Output Parsing and Scoring. Batched model outputs are parsed by extracting the structured JSON response for each query (delimited by the `<answers>` tags in the prompt template). If a query’s response cannot be parsed or the model times out without producing a complete answer, that instance is scored as incorrect. Accuracy is computed via exact match against gold labels (or task-specific metrics for extraction tasks).

Hedging Pattern Measurement. The overthinking pattern frequencies reported in Figure 4 are measured via manual annotation over 500 reasoning traces randomly sampled across all datasets combined. Annotators identified and counted occurrences of four pattern types: hedging markers (e.g., “wait,” “hold on”), rechecking phrases (e.g., “let me double-check,” “actually”), re-derivation indicators (e.g., repeated formula re-statements), and irrelevant tangents (topic shifts unrelated to the query). Counts are normalized by reasoning trace length to enable fair comparison across batch sizes.

C ADDITIONAL EXPERIMENTAL RESULTS

Statistical Significance of Token Reduction. To quantify the reliability of the observed reductions, we report mean \pm 95% confidence intervals (CIs) for accuracy, reasoning tokens, and total tokens across batch sizes (Table 5). Confidence intervals are computed over pooled evaluation instances ($n = 1300$ per model per batch size, corresponding to 13 datasets \times 100 samples each) using $\bar{x} \pm 1.96 \cdot \frac{s}{\sqrt{n}}$. Pooling across instances is appropriate because batching is applied uniformly and all metrics are defined at the instance level. Across both DeepSeek-R1 and OpenAI-o1, reasoning token reductions between $b = 1$ and $b \geq 10$ are well outside sampling variability, with highly significant differences under paired two-sided t-tests ($p < 0.01$). Accuracy remains stable or improves modestly, with differences within narrow confidence bounds, confirming that batching reduces internal deliberation without sacrificing performance.