

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 CAN CONFIDENCE ESTIMATES DECIDE WHEN CHAIN-OF-THOUGHT IS NECESSARY FOR LLMs?

Anonymous authors

Paper under double-blind review

ABSTRACT

Chain-of-thought (CoT) prompting has emerged as a common technique for enhancing the reasoning abilities of large language models (LLMs). While extended reasoning can boost accuracy on complex tasks, it is often unnecessary and substantially increases token usage, limiting the practicality of reasoning models in many scenarios. Recent models, such as GPT-OSS and Qwen3, expose controls that enable users to adjust the length of CoT or determine whether it is used at all. Yet, it remains unclear when CoT should be used: on some tasks it improves performance, while on others it provides little benefit or even harms performance. We address this challenge with confidence-gated CoT, where a model invokes reasoning only when confidence in its direct answer is low. To this end, we present the first systematic study of training-free confidence estimation methods for CoT gating. Specifically, we evaluate four training-free confidence estimation methods and compare them to a random baseline and an oracle that always knows when CoT is needed. Through extensive experiments, we show that existing training-free confidence measures can reduce redundant CoT and outperform randomly invoked CoT. However, the utility of individual confidence measures is inconsistent, varying with both the dataset and the model, underscoring the difficulty of deploying confidence-gated CoT in practice. By analysing both strengths and failure modes, our study highlights the potential and limitations of current methods and paves the way toward more reliable adaptive gating of CoT.¹

1 INTRODUCTION

Chain-of-thought (CoT) prompting (Wei et al., 2022; Guo et al., 2025) has become a cornerstone for improving the reasoning capabilities of large language models (LLMs). By encouraging models to generate step-by-step explanations before producing an answer, CoT consistently improves accuracy on tasks requiring multi-step reasoning, such as mathematics, symbolic reasoning, and scientific question answering (Wei et al., 2022; Guo et al., 2025; Qwen Team, 2025). However, extended reasoning is not always beneficial. For many queries, additional reasoning provides limited benefit and sometimes harms accuracy, while substantially increasing token usage and latency (Liu et al., 2024; Sprague et al., 2025). This inefficiency can limit the practicality of reasoning-augmented LLMs where efficiency is important.

Recent models such as GPT-OSS (OpenAI, 2025) and Qwen3 (Qwen Team, 2025) provide a *hybrid thinking mode* that lets users control when and how much reasoning the model produces. However, deciding whether CoT is necessary falls on the user, who must anticipate the difficulty of each query. Adaptive reasoning methods aim to relieve this burden by automatically adjusting reasoning depth. Most past work relies on reinforcement learning or classifiers to predict when CoT helps (Yue et al., 2025; Jiang et al., 2025; Chuang et al., 2025a). These are powerful, but they require additional training. Other work explores training-free indicators such as perplexity (Lu et al., 2025). *Our work generalises this idea under the broader notion of confidence-gating (Figure 1), where confidence signals are used to decide whether the model should answer directly or switch to CoT.*

Confidence scores give a simple signal of how reliable a model’s answer is (Kadavath et al., 2022; Kuhn et al., 2023; Farquhar et al., 2024). They can be verbalised directly by the model (Tian et al.,

¹Our anonymous code is available on <https://anonymous.4open.science/r/cgr-DDCE>

054 2023) or derived from its output probabilities Kadavath et al. (2022). They have already been used
 055 in model routing (Ramirez et al., 2024; Chuang et al., 2025b), where easy queries are sent to smaller
 056 models and harder ones to larger models. This motivates our central questions: *can self-assessed*
 057 *confidence guide LLMs in deciding when to invoke CoT reasoning?*

058 Our objective is to activate CoT only when necessary, reducing redundant tokens while preserving
 059 accuracy. We call this approach confidence-gated CoT. To evaluate this, we benchmark four
 060 representative self-assessed confidence methods across diverse reasoning benchmarks within our
 061 confidence-gated CoT. We frame this as a gating problem, where each query is routed either to di-
 062 rect answer or to CoT reasoning. To put the results in context we compare against two baselines: the
 063 expected performance of random gating and an oracle that always knows when CoT is required. The
 064 four approaches we evaluate are: asking the model to state its own certainty (verbalised confidence)
 065 (Tian et al., 2023), using the answer’s perplexity, asking whether its answer is correct($P(\text{True})$) (Ka-
 066 davath et al., 2022), and comparing the probabilities of the top two tokens (margin) Ramirez et al.
 067 (2024). We measure accuracy and token cost.

068 Our main findings are:

- 070 • Training-free confidence measures can reduce re-
 071 dundant CoT and can consistently outperform ran-
 072 dom gating, but fall short of the oracle.
- 073 • Larger models benefit more from confidence-based
 074 gating, but performance varies across tasks.
- 075 • In realistic settings, confidence-gated CoT can
 076 preserve accuracy while reducing CoT usage by
 077 25–30%, therefore substantially reducing overall to-
 078 ken cost.
- 079 • Oracle policies highlight a big opportunity for im-
 080 provement: saving more than 500 tokens on average
 081 per query and achieving 4–5% higher accuracy.

082 We identify both strengths and failure cases with qual-
 083 itative analysis, underscoring the potential challenges of
 084 deploying confidence-gated CoT in practice. Overall, our
 085 results provide the first empirical foundation for under-
 086 standing how self-assessed confidence can support adap-
 087 tive reasoning. We outline clear directions for develop-
 088 ing more reliable strategies to decide when LLMs should
 089 “think step-by-step” and when they can answer directly.

091 2 RELATED WORK

093 2.1 ADAPTIVE REASONING

095 In order to mitigate overthinking, adaptive reasoning aims to enable LLMs to dynamically adjust
 096 the depth or length of their reasoning processes based on certain indicators (Yue et al., 2025). Adaptive
 097 reasoning methods typically adopt reinforcement learning (RL) frameworks, where carefully-
 098 designed reward mechanisms guide LLMs to learn strategies under varying conditions (Jiang et al.,
 099 2025; Wang et al., 2025; Luo et al., 2025; Chung et al., 2025). Cheng et al. (2025) propose the Adaptive
 100 Cognition Policy Optimisation (ACPO) framework and an online token length budget (TLB) to
 101 enable dynamic switches between fast and slow thinking based on the estimated task difficulty. Liu
 102 et al. (2025) propose a classification-based method, which leverages features of the token probabili-
 103 ty distribution, to predict whether CoT will provide gains and switch between direct answers and
 104 CoT. The papers mentioned above rely on RL, while there are other works employ training-free
 105 estimators. Zhu et al. (2025) use entropy and token probability to decide if CoT is necessary to gen-
 106 erate each line of code during code generation. Lu et al. (2025) introduce Certainty-based Adaptive
 107 Reasoning (CAR) that uses the perplexity of a direct answer to decide if the model should think for
 longer. However, there is no research on systematically evaluating which training-free estimator is
 best suited for adaptive reasoning across diverse tasks.




Figure 1: Confidence-gated CoT controls if a query is answered directly or with reasoning: high-confidence queries are answered directly, while low-confidence ones trigger reasoning.

108 2.2 MODEL CASCADES AND ROUTING
109

110 Different from making a specific LLMs adapt to multiple reasoning modes, model cascades and
111 routing dynamically switching between multiple models. Ong et al. (2025) propose to decide when
112 to route based on a win prediction model that estimates the probability of a strong model win over
113 a weak model for a given query. Feng et al. (2025) predict the effect and cost of potential edges
114 in a graph where the task, query, and LLM are modelled as heterogeneous nodes. Ramirez et al.
115 (2024) find that simple confidence measures can effectively route harder queries to stronger models
116 compared to trained routing models. Chuang et al. (2025b) investigates a comprehensive set of
117 self-assessed confidence estimation methods for model routing.

118 2.3 CONFIDENCE AND UNCERTAINTY ESTIMATION IN LLMs
119

120 Confidence and uncertainty are two closely related concepts which are often used for gauging the
121 trustworthiness of responses generated by LLMs (Zhu et al., 2025; Chuang et al., 2025b). Lin et al.
122 (2024) give the following differentiation: uncertainty reflects the variability of a model’s predictions,
123 while confidence estimates the probability that a specific prediction is correct. Uncertainty is often
124 estimated by sampling multiple responses and measuring their semantic diversity (Kuhn et al., 2023;
125 Farquhar et al., 2024). Semantic entropy clusters equivalent answers and computes entropy over
126 aggregated probabilities, outperforming logit-based baselines such as $P(\text{True})$ (Kuhn et al., 2023;
127 Farquhar et al., 2024) at detecting hallucinations, though it requires multiple samples. To reduce
128 this cost, Kossen et al. (2025) predict semantic entropy directly from model activations. Shifting
129 Attention to Relevance (SAR) reweights token entropy by their importance to the final answer, also
130 relying on multiple sampling to get token importance scores (Duan et al., 2024). As our focus is
131 on single-pass, training-free methods with small overhead, we do not include these sampling-based
132 approaches in our evaluation.

133 3 CONFIDENCE-GATED CHAIN-OF-THOUGHT
134

136 We propose confidence-gated CoT, where a model selectively triggers reasoning based on its self-
137 assessed confidence. Each query is first answered directly. If the confidence score is low, the model
138 re-runs the query with CoT enabled. We systematically evaluate using four confidence estimation
139 methods: *perplexity*, $P(\text{True})$, *margin sampling*, and *verbalised confidence*.

140 3.1 PROBLEM DEFINITION
141

143 We study the decision of whether a model should stop after a direct answer or answer with CoT rea-
144 soning. For each input x_i , the model first generates a direct answer. A direct answer and confidence
145 score $s(x_i; \theta)$ is then derived from a model parametrised by θ . If the score is above the threshold τ ,
146 the direct answer is accepted; otherwise, the model answers the question with CoT enabled:

$$147 \quad 148 \quad 149 \quad \text{gate}(x_i; \tau, \theta) = \begin{cases} \text{CoT}(x_i; \theta), & s(x_i; \theta) < \tau \\ \text{DIRECT}(x_i; \theta), & s(x_i; \theta) \geq \tau, \end{cases}$$

150 This differs from early-exit methods, which require generating partial reasoning before deciding to
151 stop (Yang et al., 2025). In our formulation, reasoning is skipped entirely when the confidence in
152 the direct answer is sufficient. These two approaches are complementary since confidence gating
153 selects when to trigger reasoning and early exiting can still be applied once CoT has been selected.

155 **Chain-of-Thought:** This mode triggers the model to generate an explicit intermediate reasoning
156 trace before emitting a concise final answer. Specifically, we use the thinking mode of Qwen3
157 or GPT-OSS, which triggers multi-step reasoning by inserting a special instruction in the prompt
158 (Qwen Team, 2025; OpenAI, 2025).

159 **Direct:** The model is instructed to output only the final answer without generating intermediate
160 reasoning. To enforce this behaviour, we append a concise instruction such as “Answer:” to the
161 prompt, which reliably elicits a short response with no CoT or explanation.

162 3.2 SELF-ASSESSED CONFIDENCE
163

164 In this study, we limit the scope within self-assessed confidence, where the confidence scores are
165 produced by the model itself or computed based on its outputs without using another predictor. All
166 strategies we study can be generated without sampling answers multiple times and without additional
167 training. These methods have low inference overhead and are broadly applicable.

168 **Perplexity:** In our study, we view the perplexity of the generated direct answer as a measure of the
169 LLM’s confidence in it. Given a direct answer sequence $y = (y_1, \dots, y_T)$ with T tokens, perplexity
170 is defined as:

$$172 \quad 173 \quad \text{PPL}(y | x_i) = \exp\left(-\frac{1}{T} \sum_{t=1}^T \log p(y_t | y_{<t}, x_i)\right).$$

174 A higher perplexity indicates lower confidence in the generated answer.
175

176 ***P*(True) (Kadavath et al., 2022):** This approach first generates an answer via direct prompting.
177 Then, we ask the LLM whether the generated answer is (A) True or (B) False in a second forward
178 pass. We then extract the probability of generating the token “A”. Full prompt details are found in
179 Appendix D.

180 **Margin Sampling:** This method measures the difference of the probabilities between the most
181 likely and second most likely predictions produced by the model for a given input. Margin sampling
182 has been used with success for model cascades (Ramirez et al., 2024).

183 **Verbalised Confidence:** This approach prompts off-the-shelf LLMs to self-evaluate and express
184 its confidence as part of its response (Yang et al., 2024). Following prior work (Yang et al., 2024;
185 Tian et al., 2023), we ask the model to output a confidence score between 0.0 and 1.0 after its answer,
186 which has shown to provide good calibration. Full prompt details are found in Appendix D.

187 3.3 BUDGETS AND PARETO-OPTIMAL THRESHOLDS
188

189 We define the CoT budget as the proportion of queries that trigger CoT. This reflects scenarios where
190 only a limited fraction of queries can be allocated to the more costly reasoning mode. To vary this
191 budget, we sweep percentiles of the confidence score distribution, which provides a fixed fraction
192 of queries to be routed to CoT. This allows us to trace accuracy-efficiency trade-offs across different
193 budgets, plotting accuracy against average token cost or CoT usage.

194 We also explore a practical method for identifying Pareto-optimal thresholds. A threshold is Pareto-
195 optimal if no other setting achieves equal or higher accuracy at lower token cost. The set of such
196 thresholds forms the Pareto front, which traces the best accuracy-cost trade-offs. In practice, we are
197 interested in finding the point in this front with the lowest token cost whose accuracy is within a
198 tolerance ϵ of the use CoT all the time:

$$201 \quad 202 \quad \tau^* = \arg \min_{\tau} \text{Tok}(\tau) \quad \text{s.t. } \text{Acc}(\tau) \geq \text{Acc}_{\text{All-CoT}} - \epsilon.$$

203 To simulate realistic deployment, thresholds are estimated from a calibration set. We sweep per-
204 centiles, construct the Pareto front, and select τ^* . The chosen threshold is then applied to the held-
205 out test set. To account for variability in calibration splits, we repeat the procedure multiple times
206 with random fixed-size calibration/test partitions via Monte Carlo cross-validation (Xu & Liang,
207 2001). We report the mean and standard deviation of accuracy and average tokens per query across
208 runs. This tests if confidence gating can realistically preserve accuracy while reducing cost.

209 **Online vs Offline Evaluation** We consider both offline and online settings for estimating per-
210 centile thresholds. In the offline case, all direct answers and confidence scores are computed first,
211 giving access to the full distribution of confidence scores before any decision is made. This allows
212 thresholds to be set exactly at chosen percentiles. In the online case, we simulate streaming input
213 queries so thresholds must be decided on the fly without access to the overall confidence score dis-
214 tribution. We follow the dynamic percentile method introduced by Ramirez et al. (2024). After
215 each query t , the threshold τ_t is set to the p -th percentile of $\{s(x_1), \dots, s(x_{t-1})\}$. We randomise




Figure 2: **Offline accuracy–efficiency trade-offs under percentile budgets.** Accuracy vs. CoT usage (top) and vs. average tokens (bottom), aggregated over all datasets for GPT-OSS-20B (medium effort), Qwen3-32B, and Qwen3-8B. Curves show verbalised, perplexity, $P(\text{True})$, and margin vs. the random baseline; stars denote the oracle. Full GPT-OSS results for low/medium/high effort are in Appendix E.

dataset order and use a short warm-up phase (the first 20 queries answered directly) to initialise the observations, and report the mean and standard deviation over 10 runs.

4 EXPERIMENTAL SETUP

4.1 MODELS

Hybrid-reasoning models allow the user to choose between **thinking** and **non-thinking modes** (Qwen Team, 2025). We extend this definition to GPT-OSS, which allows the user to choose between low, medium and high reasoning effort (CoT length) in the prompt (OpenAI, 2025). *We focus on hybrid models such as Qwen3 and GPT-OSS, which natively support confidence-based gating.* In contrast, non-hybrid models without controllable modes require additional instructions or mechanisms to selectively enable reasoning at inference time. GPT-OSS supports three CoT effort settings (*low/medium/high*) controlled via prompt. Unless specified, the CoT results of GPT-OSS were generated with the *medium* setting. We provide the results of the other effort levels in Appendix E.

4.2 DATASETS

The experiments include seven datasets (statistics in Appendix D Table 3) from four reasoning types (Sprague et al., 2025): (1) *commonsense reasoning* including CommonsenseQA (CSQA) (Talmor et al., 2019) and StrategyQA (Geva et al., 2021); (2) *knowledge-based reasoning* using MMLU-redu (Gema et al., 2025); (3) *mathematical and scientific reasoning* on GPQA (Rein et al., 2024) and GSM8k (Cobbe et al., 2021); and (4) *soft reasoning* using LSAT-AGI (Zhong et al., 2024) and MUSR (Sprague et al., 2025). Following (Sprague et al., 2025), these are multiple choice or short answer tasks as CoT is not used as frequently for long-form responses. This wide range of reasoning types allows us to test datasets where reasoning has shown different levels of effectiveness.

4.3 BASELINES

Expected Random Baseline. For a given CoT usage budget $r \in [0, 1]$, we report the expected accuracy and token cost: $\text{Acc}_r = (1 - r) \text{Acc}_{\text{Direct}} + r \text{Acc}_{\text{CoT}}$, $\text{Tok}_r = (1 - r) \text{Tok}_{\text{Direct}} + r \text{Tok}_{\text{CoT}}$.




Figure 3: **Task-level accuracy–efficiency trade-offs.** Representative datasets (CSQA, StrategyQA, GSM8K, GPQA) comparing confidence-gating to random and oracle across models.

We compute these analytically rather than by randomly selecting for each point, which provides a fairer and more stable baseline.

Oracle. To assess the ceiling of confidence-based gating, we include an oracle method that triggers CoT whenever the direct answer is incorrect. This setting assumes perfect knowledge of correctness and therefore represents the maximum performance that any confidence signal could achieve. The oracle thus serves as an upper bound on the potential of confidence-guided CoT routing.

5 EXPERIMENTAL RESULTS

First, we look at the offline setting, where thresholds are chosen with access to the score distribution over the entire datasets. This provides a clear view of the trade-offs between accuracy and efficiency at different CoT budgets. We then turn to the online setting to see if these trade-offs hold on streaming inputs. Finally, the Pareto-optimal analysis identifies settings that maintain accuracy while lowering token cost.

CoT Budget–Accuracy Trade-offs. We evaluate accuracy–efficiency curves by sweeping percentile budgets as defined in §3.3. At each budget level, we report both accuracy and average token usage. Figure 2 shows aggregate results for GPT-OSS-20B, Qwen3-32B, and Qwen3-8B, comparing confidence-based gating against random selection and the oracle.

For both GPT-OSS-20B and Qwen3-32B, there are confidence methods that achieve clear wins over the random baseline. Specifically, *margin* and *perplexity* consistently outperform random for GPT-OSS-20B, while $P(\text{True})$ is most effective for Qwen3-32B. Using these methods, both models can match the accuracy of always using CoT while invoking it roughly 30–40% less often, showing that confidence can cut token usage effectively at different budgets. For Qwen3-8B, sometimes *margin* and *perplexity* outperform random at low percentiles and $P(\text{True})$ at higher ones, but no method consistently beats random across all budgets. The oracle highlights that large efficiency improvements are possible, for example, for GPT-OSS-20B the oracle achieves 5% higher accuracy while invoking CoT on less than half of the queries.

Efficiency Gains Vary Across Tasks. *Commonsense, soft reasoning, and knowledge tasks benefit the most from confidence-based gating.* In Figure 3, we show representative examples including CSQA and Strategy QA. On datasets such as MMLU, StrategyQA, and MUSR, both GPT-OSS and Qwen3-32B can achieve the same accuracy as always using CoT while reducing token usage by 30–50%. In some cases, such as StrategyQA and MUSR with GPT-OSS and Qwen3-32B, performance even improves slightly at certain budgets while using fewer tokens. Figure 3 also shows high potential for these tasks, with the oracle using about 75% less CoT for CSQA and StrategyQA. In contrast, mathematical and scientific tasks show limited benefit. For GSM8K, direct answering without CoT has very low accuracy, making it difficult to save tokens without hurting performance. This is also clear from the oracle, which shows less room for improvement (Figure 3). Similarly, on GPQA, some confidence methods (e.g., *perplexity* for GPT-OSS-20B) perform better than random and yield modest savings, but the efficiency gains are much less pronounced. The oracle highlights that there is headroom for efficiency on GPQA, but current models are not effective at separating

324 correct from incorrect answers for these challenging questions. Full results across all datasets can
 325 be found in Appendix E.
 326


327 **No Confidence Method Dominates.** As seen in
 328 Figure 2, the effectiveness of confidence signals
 329 varies strongly across models. For GPT-OSS-20B,
 330 *margin* and *perplexity* can outperform random gat-
 331 ing across all budgets. In contrast, both $P(\text{True})$
 332 and *verbalised* confidence often perform worse than
 333 random. For Qwen3-32B, $P(\text{True})$ emerges as the
 334 most effective method, outperforming other signals
 335 across a wide range of budgets. *Margin* and *per-
 336 perplexity* achieve above-random performance only at
 337 low budgets, but quickly saturate: the distributions
 338 collapse to narrow ranges, limiting their separating
 339 power. Finally, for Qwen3-8B, no method consis-
 340 tently outperforms random gating across all budgets.
 341

342 **Scale and Calibration Effects.** To better understand why confidence gating is more effective in
 343 larger models, we look at the calibration of each confidence signal. Figure 4 shows reliability di-
 344 agrams for GPT-OSS-20B, where margin sampling achieves the highest AUROC. Broadly, both
 345 GPT-OSS-20B and Qwen3-32B achieve higher AUROC across methods compared to Qwen3-8B
 346 (Appendix F). This suggests that, at a larger scale, LLMs are better calibrated and can more reli-
 347 ably separate correct and incorrect predictions. This finding is consistent with prior findings that
 348 calibration improves with model size (Kadavath et al., 2022). It also demonstrates why we observe
 349 positive confidence gating results from the larger models. Notably, Qwen3-8B generally produces
 350 longer CoT with an average length of 1,269 tokens compared to 625 for GPT-OSS-20B (high) and
 351 884 for Qwen3-32B. Although Qwen3-8B stands to benefit the most from effective gating, its weak
 352 calibration prevents it from achieving these gains.
 353

354 5.1 REALISTIC CoT DEPLOYMENT

355 In realistic settings, a model must decide when
 356 to invoke CoT given a budget, without access to
 357 the full confidence distribution. As outlined in
 358 §3.3, we address this with dynamic percentile
 359 thresholding (Ramirez et al., 2024), which up-
 360 dates thresholds online from past scores. We
 361 first examine budget–accuracy trade-offs under
 362 this setting, and then turn to Pareto-optimal
 363 thresholds, which approximate realistic deploy-
 364 ment by selecting accuracy-preserving operat-
 365 ing points from a calibration set.
 366

367 **Online CoT Budget–Accuracy Trade-offs.**
 368 We implement the dynamic percentile thresh-
 369 olding procedure from §3.3 to enforce CoT
 370 budgets in the online setting. Figure 5 shows
 371 that the online curves broadly mirror the of-
 372 fline ones, confirming that CoT budgets can re-
 373 main effective under realistic deployment con-
 374 ditions. The main difference is increased vari-
 375 ance. GPT-OSS-20B remains stable across
 376 budgets with behaviour very similar to the of-
 377 fline setting. However, for the Qwen3 models,
 378 *margin* and *perplexity* show noticeably higher variabil-
 379 ity at mid-to-high budgets, reflecting instabil-
 380 ity from a loss of separability in their scores. In contrast, $P(\text{True})$ on Qwen3-32B remains stable



381 Figure 4: Reliability diagram for GPT-OSS-
 382 20B with *margin* confidence.
 383

384 Table 1: Results for all datasets with Pareto-
 385 optimal thresholds ($\epsilon = 1\%$). Accuracy remains
 386 within 1% of All CoT; differences are in CoT us-
 387 age and tokens saved per query.
 388

Method	Acc. \uparrow	Δ Acc \uparrow	CoT (%) \downarrow	Avg. Tok. saved \uparrow
All CoT	79.9	0.0	100.0	0.0
GPT-OSS-20B	All Direct	54.1	-25.9	483.3
	$P(\text{True})$	79.2 ± 0.5	-0.7	95.5 ± 2.3
	Verbalised	79.7 ± 0.1	-0.2	99.2 ± 0.0
	Margin	79.1 ± 0.4	-0.8	68.1 ± 3.8
	Perplexity	78.9 ± 0.5	-1.0	70.6 ± 7.9
Oracle	85.0	+5.1	45.9	187.2
Qwen3-32B	All CoT	83.8	0.0	100.0
	All Direct	67.8	-16.0	0.0
	$P(\text{True})$	82.8 ± 0.5	-1.0	73.8 ± 5.6
	Verbalised	83.7 ± 0.1	-0.1	98.9 ± 0.0
	Margin	83.8 ± 0.1	0.0	100.0 ± 0.0
Oracle	87.9	+4.1	32.2	446.7
Qwen3-8B	All CoT	79.1	0.0	100.0
	All Direct	57.8	-21.3	0.0
	$P(\text{True})$	78.4 ± 0.5	-0.7	90.8 ± 4.9
	Verbalised	79.0 ± 0.3	-0.1	100.0 ± 0.5
	Margin	79.1 ± 0.2	0.0	100.0 ± 0.0
Oracle	83.8	+4.0	42.2	563.8

389




Figure 5: Online Accuracy vs. CoT deferral rate (top) and average tokens (bottom) across all datasets in the **online** setting. Stars show oracle performance.

and very close to its offline performance. Beyond online budget trade-offs, we also consider whether we can find a Pareto-optimal threshold using a calibration set.

Pareto-optimal thresholds. We implement the procedure from §3.3 using a 10% calibration split, $\epsilon = 1\%$, and 100 repeats sampling a different calibration split each time (Xu & Liang, 2001). We report the mean and standard deviation of accuracy and cost across these runs. Table 1 reports these results. We see that for Qwen3-8B, $P(\text{True})$ maintains accuracy within 1% of the full CoT baseline while reducing CoT usage by around 10% and saving 89 tokens per query on average. This shows that although no method on Qwen3-8B consistently outperforms the random baseline across the full budget sweep, confidence signals can still identify thresholds that deliver useful savings without hurting accuracy. The larger models also yield Pareto-optimal thresholds that preserve accuracy while lowering token cost, with GPT-OSS-20B achieving reductions of 30-35% in CoT usage and Qwen3-32B showing meaningful savings under $P(\text{True})$. These results confirm that even when the overall trade-off curves appear modest, calibration can highlight settings where performance is maintained and efficiency improves.

6 ANALYSIS

To better understand how confidence-gated CoT operates in practice, we examine qualitative examples of both successful and unsuccessful gating cases for a maximum accuracy Pareto-optimal threshold obtained following the method described in the previous section. Then, we separate outcome types from cases where CoT is genuinely needed to those where it adds little or no value (i.e., direct answers that are already correct). Table 2 shows the distribution of these outcomes, defined by the policy’s decision to use CoT and the correctness of the final answer.

Outcome breakdown. Table 2 reports average outcomes at the Pareto-optimal thresholds selected in the previous results, using the best-performing confidence method for each model. The largest share of cases for Qwen3-8B (50.8%) and Qwen3-32B (44.4%) falls into *Excess CoT*, where the direct answer was already correct but the policy still used reasoning. GPT-OSS-20B is lower at 26.5%. The *Direct* category, where the policy chose to answer directly and that answer was correct, improves with scale, from 7.0% (Qwen3-8B) to 23.5% (Qwen3-32B) and 27.2% (GPT-OSS-20B). This trend is expected since knowing when to use direct requires the model to have well calibrated confidence scores, and Qwen3-8B shows the poorest calibration. *CoT Fixed* accounts for 24.7% of queries on Qwen3-8B, 18.8% on Qwen3-32B, and 29.9% on GPT-OSS-20B, capturing the cases where the direct answer would have been wrong but enabling CoT corrected it. *Both Fail*, cases where gating couldn’t help, remain around 12–16% across models. We present representative examples of some of these categories below.

432 **Example 1: CoT Fixed.** Across models, 18–30% of queries fall into this category. Good examples
 433 of this come from GSM8K, where direct answering often fails but CoT achieves high accuracy
 434 (Sprague et al., 2025). In Example 1, the direct answer was wrong, but the CoT reasoning solved
 435 the problem correctly.

437 **Example 1**

439 **Question:** Martha is planning her Christmas party. She invited 2 families with 6 people and 3 families with 4 people. 8 people
 440 couldn't come due to illness, and 1/4 that number had previous commitments. How many people show up for Martha's party?

441 **Direct answer:** 25 (incorrect).

442 **CoT reasoning (excerpt):** “Okay, let me try to figure out how many people are coming to Martha’s Christmas party... 2 families
 443 times 6 people = 12. 3 families times 4 people = 12. Adding together gives 24 invited. Now, 8 people couldn’t come due to illness.
 Then, 1/4 of that number (2) had previous commitments... So total not coming = 8 + 2 = 10. Therefore, 24 - 10 = 14. ...Yes, that
 444 seems correct. **The answer is 14.**”

445 **Ground Truth:** 14.

446 **Example 2: Direct.** This represents a positive case where accuracy is preserved and tokens are
 447 saved by directly answering. This example saved 284 tokens by choosing to skip CoT.

449 **Example 2**

451 **Question:** Would a Nike shoebox be too small to fit a swan in?

452 **Direct answer:** Yes ($P(\text{True}) = 0.99$).

453 **Ground Truth:** Yes.

455 **Example 3: Excess CoT.** In this example, the direct answer was already correct, but the policy
 456 still used CoT, leading to redundant tokens.

458 **Example 3**

460 **Question:** Where would you put a glass after drinking from it?

461 **Answer choices:** (A) ocean, (B) water cooler, (C) cabinet, (D) dishwasher, (E) dining room.

462 **Direct answer:** (D) ($P(\text{True}) = 0.59$).

463 **CoT reasoning (excerpt):** “Option A doesn’t make sense... Option D, dishwasher, is correct. Therefore, the answer is D.”

464 These examples highlight both the promise and
 465 the limitations of confidence-gated CoT. On the
 466 positive side, gating can recover accuracy when
 467 CoT is required (as in GSM8K) and preserve ac-
 468 curacy while saving tokens when direct answers
 469 are sufficient. At the same time, unnecessary
 470 CoT remains common, with unnecessary reason-
 471 ing the single largest category in our breakdown
 472 (Table 2). This underlines that while training-free
 473 confidence signals can guide useful savings, they
 474 are inconsistent in practice, and stronger, more
 475 consistent gating indicators will be needed for re-
 476 liable CoT deployment.

Table 2: Distribution of outcome categories across three models. Values are averages over calibration runs with standard deviations shown.

Category	Qwen8B	Qwen32B	OSS20B
CoT Fixed	24.7% _{±0.8}	18.8% _{±0.6}	29.9% _{±0.6}
Direct	7.0% _{±3.6}	23.5% _{±4.6}	27.2% _{±2.5}
Excess CoT	50.8% _{±3.6}	44.4% _{±4.6}	26.5% _{±2.5}
Missed Fix	1.2% _{±0.8}	1.3% _{±0.6}	1.3% _{±0.6}
Both fail	16.2% _{±0.1}	12.1% _{±0.1}	15.0% _{±0.1}

477 **7 CONCLUSION**

479 To our knowledge, we conducted the first systematic study of confidence-guided CoT gating in
 480 LLMs. Our results show that training-free confidence signals preserve accuracy and cut redundant
 481 reasoning by 25–30%, thereby lowering overall token cost. These findings imply that LLMs already
 482 possess useful self-assessment signals that can make reasoning more efficient, especially at scale, but
 483 current confidence estimation methods are too brittle for robust deployment. The challenge ahead
 484 is to develop models that are not only capable of reasoning but also calibrated in terms of when
 485 to reason. Progress would lower inference cost and latency while improving reliability, making
 adaptive CoT a practical tool for large-scale, real-world systems.

486 REFERENCES
487

488 Xiaoxue Cheng, Junyi Li, Zhenduo Zhang, Xinyu Tang, Wayne Xin Zhao, Xinyu Kong, and
489 Zhiqiang Zhang. Incentivizing dual process thinking for efficient large language model reasoning,
490 2025. URL <https://arxiv.org/abs/2505.16315>.

491 Yu-Neng Chuang, Prathusha Kameswara Sarma, Parikshit Gopalan, John Boccio, Sara Bolouki,
492 Xia Hu, and Helen Zhou. Learning to route LLMs with confidence tokens. In *Forty-second*
493 *International Conference on Machine Learning*, 2025a. URL <https://openreview.net/forum?id=U08mUogGDM>.

494

495 Yu-Neng Chuang, Leisheng Yu, Guanchu Wang, Lizhe Zhang, Zirui Liu, Xuanting Cai, Yang Sui,
496 Vladimir Braverman, and Xia Hu. Confident or seek stronger: Exploring uncertainty-based on-
497 device llm routing from benchmarking to generalization, 2025b. URL <https://arxiv.org/abs/2502.04428>.

498

499 Stephen Chung, Wenyu Du, and Jie Fu. Thinker: Learning to think fast and slow, 2025. URL
500 <https://arxiv.org/abs/2505.21097>.

501

502 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
503 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John
504 Schulman. Training verifiers to solve math word problems, 2021. URL <https://arxiv.org/abs/2110.14168>.

505

506 Jinhao Duan, Hao Cheng, Shiqi Wang, Alex Zavalny, Chenan Wang, Renjing Xu, Bhavya Kailkhura,
507 and Kaidi Xu. Shifting attention to relevance: Towards the predictive uncertainty quantification
508 of free-form large language models. In Lun-Wei Ku, Andre Martins, and Vivek Sriku-
509 mar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Lin-
510 guistics (Volume 1: Long Papers)*, pp. 5050–5063, Bangkok, Thailand, August 2024. Asso-
511 ciation for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.276. URL <https://aclanthology.org/2024.acl-long.276/>.

512

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

515

516 Tao Feng, Yanzhen Shen, and Jiaxuan You. Graphrouter: A graph-based router for LLM selections.
517 In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=eU39PDsZtT>.

518

519 Aryo Pradipta Gema et al. Are we done with MMLU? In Luis Chiruzzo, Alan Ritter, and
520 Lu Wang (eds.), *Proceedings of the 2025 Conference of the Nations of the Americas Chapter*
521 *of the Association for Computational Linguistics: Human Language Technologies (Volume 1:
522 Long Papers)*, pp. 5069–5096, Albuquerque, New Mexico, April 2025. Association for Compu-
523 tational Linguistics. ISBN 979-8-89176-189-6. doi: 10.18653/v1/2025.naacl-long.262. URL
524 <https://aclanthology.org/2025.naacl-long.262/>.

525

526 Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. Did aristotle
527 use a laptop? a question answering benchmark with implicit reasoning strategies. *Transactions*
528 *of the Association for Computational Linguistics*, 9:346–361, 2021. doi: 10.1162/tacl_a_00370.
529 URL <https://aclanthology.org/2021.tacl-1.21/>.

530

531 Daya Guo et al. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*,
532 645(8081):633–638, Sep 2025. ISSN 1476-4687. doi: 10.1038/s41586-025-09422-z. URL
533 <https://doi.org/10.1038/s41586-025-09422-z>.

534

535 Lingjie Jiang, Xun Wu, Shaohan Huang, Qingxiu Dong, Zewen Chi, Li Dong, Xingxing Zhang,
536 Tengchao Lv, Lei Cui, and Furu Wei. Think only when you need with large hybrid-reasoning
537 models, 2025. URL <https://arxiv.org/abs/2505.14631>.

538

539 Saurav Kadavath et al. Language models (mostly) know what they know, 2022. URL <https://arxiv.org/abs/2207.05221>.

540 Jannik Kossen, Jiatong Han, Muhammed Razzak, Lisa Schut, Shreshth A Malik, and Yarin Gal.
 541 Semantic entropy probes: Robust and cheap hallucination detection in LLMs, 2025. URL
 542 <https://openreview.net/forum?id=YQvvJjLWX0>.

543 Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. Semantic uncertainty: Linguistic invariances
 544 for uncertainty estimation in natural language generation. In *The Eleventh International Conference*
 545 on *Learning Representations*, 2023. URL <https://openreview.net/forum?id=VD-AYtP0dve>.

546 Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. Generating with confidence: Uncertainty quantification
 547 for black-box large language models. *Transactions on Machine Learning Research*, 2024.
 548 ISSN 2835-8856. URL <https://openreview.net/forum?id=DWkJCSxKU5>.

549 Peijie Liu, Fengli Xu, and Yong Li. Token signature: Predicting chain-of-thought gains with token
 550 decoding feature in large language models. In *Forty-second International Conference on Machine*
 551 *Learning*, 2025. URL <https://openreview.net/forum?id=UfLJqcEle6>.

552 Ryan Liu, Jiayi Geng, Addison J. Wu, Ilia Sucholutsky, Tania Lombrozo, and Thomas L. Griffiths.
 553 Mind your step (by step): Chain-of-thought can reduce performance on tasks where thinking
 554 makes humans worse, 2024. URL <https://arxiv.org/abs/2410.21333>.

555 Jinghui Lu et al. Prolonged reasoning is not all you need: Certainty-based adaptive routing for
 556 efficient llm/mllm reasoning, 2025. URL <https://arxiv.org/abs/2505.15154>.

557 Haotian Luo, Haiying He, Yibo Wang, Jinluan Yang, Rui Liu, Naiqiang Tan, Xiaochun Cao,
 558 Dacheng Tao, and Li Shen. Ada-r1: Hybrid-cot via bi-level adaptive reasoning optimization,
 559 2025. URL <https://arxiv.org/abs/2504.21659>.

560 Isaac Ong, Amjad Almahairi, Vincent Wu, Wei-Lin Chiang, Tianhao Wu, Joseph E. Gonzalez,
 561 M Waleed Kadous, and Ion Stoica. RouteLLM: Learning to route LLMs from preference
 562 data. In *The Thirteenth International Conference on Learning Representations*, 2025. URL
 563 <https://openreview.net/forum?id=8sSqNntaMr>.

564 OpenAI. gpt-oss-120b gpt-oss-20b model card, 2025. URL <https://arxiv.org/abs/2508.10925>.

565 Qwen Team. Qwen3 technical report, 2025. URL <https://arxiv.org/abs/2505.09388>.

566 Guillem Ramirez, Alexandra Birch, and Ivan Titov. Optimising calls to large language models
 567 with uncertainty-based two-tier selection. In *Proceedings of the 2024 Conference on Language*
 568 *Modeling*, July 2024. URL <https://colmweb.org/>. Conference on Language Modeling,
 569 COLM 2024 ; Conference date: 07-10-2024 Through 09-10-2024.

570 David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien
 571 Dirani, Julian Michael, and Samuel R. Bowman. GPQA: A graduate-level google-proof q&a
 572 benchmark. In *First Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=Ti67584b98>.

573 Zayne Rea Sprague, Xi Ye, Kaj Bostrom, Swarat Chaudhuri, and Greg Durrett. MuSR: Testing the
 574 limits of chain-of-thought with multistep soft reasoning. In *The Twelfth International Conference*
 575 on *Learning Representations*, 2024. URL <https://openreview.net/forum?id=jenyYQzuel>.

576 Zayne Rea Sprague, Fangcong Yin, Juan Diego Rodriguez, Dongwei Jiang, Manya Wadhwa,
 577 Prasann Singhal, Xinyu Zhao, Xi Ye, Kyle Mahowald, and Greg Durrett. To cot or not to cot?
 578 chain-of-thought helps mainly on math and symbolic reasoning. In *The Thirteenth International*
 579 *Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=w6nlcS8Kkn>.

580 Alon Talmor, Jonathan Herzig, and Jonathan Berant. CommonsenseQA: A question answering
 581 challenge targeting commonsense knowledge. In Jill Burstein, Christy Doran, and Thamar Solorio
 582 (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for*
 583 *Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*,
 584 pp. 4149–4158, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
 585 doi: 10.18653/v1/N19-1421. URL <https://aclanthology.org/N19-1421/>.

594 Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea
 595 Finn, and Christopher Manning. Just ask for calibration: Strategies for eliciting calibrated con-
 596 fidence scores from language models fine-tuned with human feedback. In Houda Bouamor,
 597 Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Meth-
 598 ods in Natural Language Processing*, pp. 5433–5442, Singapore, December 2023. Associa-
 599 tion for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.330. URL <https://aclanthology.org/2023.emnlp-main.330/>.

600

601 Yunhao Wang, Yuhao Zhang, Tinghao Yu, Can Xu, Feng Zhang, and Fengzong Lian. Adaptive deep
 602 reasoning: Triggering deep thinking when needed, 2025. URL <https://arxiv.org/abs/2505.20101>.

603

604 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi,
 605 Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language
 606 models. In *Proceedings of the 36th International Conference on Neural Information Processing
 607 Systems*, NIPS '22, Red Hook, NY, USA, 2022. Curran Associates Inc. ISBN 9781713871088.

608

609 Qing-Song Xu and Yi-Zeng Liang. Monte carlo cross validation. *Chemometrics and Intel-
 610 ligent Laboratory Systems*, 56(1):1–11, 2001. ISSN 0169-7439. doi: [https://doi.org/10.1016/S0169-7439\(00\)00122-2](https://doi.org/10.1016/S0169-7439(00)00122-2). URL <https://www.sciencedirect.com/science/article/pii/S0169743900001222>.

611

612

613 Chenxu Yang, Qingyi Si, Yongjie Duan, Zheliang Zhu, Chenyu Zhu, Qiaowei Li, Zheng Lin, Li Cao,
 614 and Weiping Wang. Dynamic early exit in reasoning models, 2025. URL <https://arxiv.org/abs/2504.15895>.

615

616

617 Daniel Yang, Yao-Hung Hubert Tsai, and Makoto Yamada. On verbalized confidence scores for
 618 llms, 2024. URL <https://arxiv.org/abs/2412.14737>.

619

620 Linan Yue, Yichao Du, Yizhi Wang, Weibo Gao, Fangzhou Yao, Li Wang, Ye Liu, Ziyu Xu, Qi Liu,
 621 Shimin Di, and Min-Ling Zhang. Don't overthink it: A survey of efficient r1-style large reasoning
 622 models, 2025. URL <https://arxiv.org/abs/2508.02120>.

623

624 Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied,
 625 Weizhu Chen, and Nan Duan. AGIEval: A human-centric benchmark for evaluating foundation
 626 models. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Findings of the Association
 627 for Computational Linguistics: NAACL 2024*, pp. 2299–2314, Mexico City, Mexico, June 2024.
 628 Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-naacl.149. URL
<https://aclanthology.org/2024.findings-naacl.149/>.

629

630 Yuqi Zhu, Ge Li, Xue Jiang, Jia Li, Hong Mei, Zhi Jin, and Yihong Dong. Uncertainty-guided
 631 chain-of-thought for code generation with llms, 2025. URL <https://arxiv.org/abs/2503.15341>.

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646


647

648 **A REPRODUCIBILITY STATEMENT**
649650 Our code to reproduce all experiments is available on an anonymous GitHub repository:
651 <https://anonymous.4open.science/r/cgr-DDCE>. This repository will remain accessible until the
652 ICLR 2026 decision notification date: Jan 22, 2026 (AOE). All inference hyperparameters are spec-
653 ified in Appendix B. Experiments were run on a combination of Nvidia A100 (80GB) and H100
654 (80GB) GPUs. We report results as the mean and standard deviation across repeated experiments.
655 We will also provide all generated outputs, including CoT traces and confidence scores. All datasets
656 are publicly available.
657658 **B MODEL INFERENCE SETTINGS**
659660 We use Hugging Face Transformers for all inference. For Qwen models (8B and 32B), we follow
661 the recommended decoding settings from the model cards, using temperature 0.6 and top-p 0.95
662 to avoid degenerate repetition. For GPT-OSS-20B, we use the default sampling configuration with
663 temperature 1.0 and top-p 1.0. In all setting we set a maximum limit of 7000 thinking tokens and
664 insert text that prompts the model to answer after this limit has been reached.
665666 **C LLM USAGE**
667668 The writing of this paper received proofreading and language polishing suggestions using LLMs. In
669 addition, parts of our experimental code were drafted or refactored with the assistance of GitHub
670 Copilot; all final text and code was manually reviewed and verified by the authors.
671672 **D PROMPTS AND DATASET STATISTICS**
673674 **Verbalised Prompt**675 Please directly provide your best guess of the answer to the question and give the probability that you
676 think it is correct (0.0 to 1.0). Take your uncertainty in the prompt, the task difficulty, your knowledge
677 availability, and other sources of uncertainty into account.
678679 Give only the guess and probability, with no other words or explanation.
680681 Format your final response as:
682 Answer: <your_best_guess>
683 Probability: <score between 0.0 and 1.0>684 **$P(\text{True})$ Prompt**685 User:
686 Is this answer:
687 (A) True
688 (B) False
689690 Assistant:
691 The answer is:692 **Table 3: Dataset statistics.**
693694
695

Dataset	# Samples
CommonsenseQA (CSQA)	1221
StrategyQA	2290
MMLU-redux	3000
GSM8K	1319
GPQA	448
LSAT-AGI	1009
MUSR	756

E PER DATASET TRADE-OFF PLOTS


Figure 6: CSQA: Accuracy vs. CoT use (top) and average tokens (bottom) across models



756
757
758
759
760
761
762
763


764 Figure 7: GPQA: Accuracy vs. CoT use (top) and average tokens (bottom) across models

765



810
811
812
813
814
815
816
817


Figure 8: GSM8K: Accuracy vs. CoT use (top) and average tokens (bottom) across models



864
865
866
867
868
869
870
871


Figure 9: LSAT-All: Accuracy vs. CoT use (top) and average tokens (bottom) across models

872
873



918
919
920
921
922
923
924
925


Figure 10: MuSR-All: Accuracy vs. CoT use (top) and average tokens (bottom) across models



1026
1027
1028
1029
1030
1031
1032
1033

Figure 12: StrategyQA: Accuracy vs. CoT use (top) and average tokens (bottom) across models

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



F RELIABILITY DIAGRAMS




Figure 13: Reliability diagrams for **Qwen3-8B**. Bars darken with bin count; dashed line is perfect calibration.

