© TOKUR: TOKEN-LEVEL UNCERTAINTY ESTIMATION FOR LARGE LANGUAGE MODEL REASONING

Anonymous authors

Paper under double-blind review

ABSTRACT

While Large Language Models (LLMs) have demonstrated impressive capabilities, their output quality remains inconsistent across various application scenarios, making it difficult to identify trustworthy responses, especially in complex tasks requiring multi-step reasoning. In this paper, we propose a **Token**-level Uncertainty estimation framework for **Reasoning** (**Tokur**) that enables LLMs to self-assess and self-improve their responses in mathematical reasoning. Specifically, we introduce low-rank random weight perturbation during LLM decoding to generate predictive distributions for token-level uncertainty estimation, and we aggregate these uncertainty quantities to capture the semantic uncertainty of generated responses. Experiments on mathematical reasoning datasets of varying difficulty demonstrate that Tokur exhibits a strong correlation with answer correctness and model robustness, and the uncertainty signals produced by Tokur can be leveraged to enhance the model's reasoning performance at test time. These results highlight the effectiveness of Tokur as a principled and scalable approach for improving the reliability and interpretability of LLMs in challenging reasoning tasks.

1 Introduction

Large Language Models (LLMs) have demonstrated remarkable capabilities in various reasoning tasks (Wei et al., 2022a; Wang et al., 2022; Chung et al., 2024; Guo et al., 2025), yet they often struggle to reliably assess the quality of their own responses (Xiong et al., 2023; Tian et al., 2023; Kapoor et al., 2024; Liu et al., 2024; Zhang & Zhang, 2025; Da et al., 2025; Liu et al., 2025). This limitation becomes particularly evident in complex reasoning scenarios where models may generate seemingly convincing but incorrect solutions without indicating uncertainty.

Beyond the dominant body of uncertainty estimation methods that largely focus on short-form question answering (Zhang et al., 2023; Yadkori et al., 2024) and classification tasks (Yang et al., 2023; Wang et al., 2024; Shi et al., 2024), two main approaches have been explored for the more challenging setting of sequence uncertainty estimation: (i) Query-level methods (Gal et al., 2016; Osband et al., 2023; Hou et al., 2023), despite their solid theoretical foundation, estimate uncertainties $\mathcal{U}(y|x)$ with respect to input prompts x alone, without evaluating the quality of specific generated responses y conditioned on those inputs (see Sec. 2.1). Besides, these methods require marginalization over the entire output space y; this becomes intractable as sequence length grows. (ii) Response-level methods (Murray & Chiang, 2018; Malinin & Gales, 2021; Kadavath et al., 2022), typically variants of log-probabilities, have shown empirical success but lack strong theoretical grounding (Kuhn et al., 2023). As a result, the limitations of the aforementioned methods in capturing response-specific uncertainty hinder the deployment of LLMs in high-stakes reasoning tasks that demand reliable self-assessment.

To address this challenge, we propose a principled framework, dubbed **Token-level Uncertainty** estimation for **Re**asoning (TokUR), for estimating the uncertainty of generated sequences by aggregating token-level uncertainties based on random low-rank weight perturbation. TokUR introduces carefully calibrated perturbations to the weights of attention layers, creating an ensemble of model variants that enables principled uncertainty estimation without requiring costly retraining or extensive parameter updates. Building on this, we decompose the *total uncertainty* of each generated token into *aleatoric uncertainty* (inherent randomness in the data) and *epistemic uncertainty* (model uncertainty about its parameters), providing a theoretically grounded assessment of confidence across the generation

process. We then aggregate these token-level uncertainties to evaluate entire reasoning responses, demonstrating both theoretical consistency with established uncertainty principles and practical utility in downstream applications.

Empirically, TokUR enhances LLM reasoning in three key aspects: (i) token-level *epistemic uncertainty* effectively identifies incorrect reasoning paths, outperforming baselines across three mathematical reasoning benchmarks, (ii) TokUR excels at selecting high-quality solutions from multiple candidates, and (iii) it functions as an implicit reward to guide reasoning, improving accuracy when combined with off-the-shelf test-time-scaling algorithms (Puri et al., 2025). In summary, our contributions are:

- We introduce TokUR, a training-free token-level uncertainty estimation approach for LLM reasoning through low-rank weight perturbation, providing a principled decomposition of uncertainties with proven theoretical properties.
- We demonstrate that epistemic uncertainty can serve as a good metric to measure the quality
 of generated reasoning paths, consistently outperforming conventional confidence metrics
 across diverse mathematical reasoning tasks.
- We demonstrate practical applications of our uncertainty estimation framework: it improves reasoning performance through incorrect path detection, high-quality solution selection, and uncertainty-guided generation.

2 Preliminaries

In this section, we first introduce the notation used in the remaining sections, and then review the key concepts of uncertainties (Sec. 2.1) and existing Bayesian LLMs for downstream adaptation (Sec. 2.2).

Notation. In this paper, scalars are denoted by lowercase letters (x), vectors by lowercase bold-math letters (x), random vectors by lowercase boldface letters (x), and matrices by uppercase boldface letters (X). We use $[m] = \{1, 2, \cdots, m\}$ to denote the set of consecutive integer numbers from 1 to m. Following convention, we use p for probability, \mathbb{E} for expectation, \mathcal{H} for entropy, and \mathcal{I} for mutual information. Specifically, $\mathcal{H}[\mathbf{y}|\mathbf{x}]$ denotes the conditional entropy between $random\ variables$ \mathbf{y} and \mathbf{x} . We use $\mathcal{H}[p(\mathbf{y}|\mathbf{x}=\mathbf{x})]$ to denote the $predictive\ entropy$ of the output variable conditioned on input \mathbf{x} , with $\mathcal{H}[p(\mathbf{y}|\mathbf{x})]$ as a shorthand notation when context is clear.

2.1 Uncertainty Estimation of Long-Form Generation

Prediction with Bayesian Neural Networks. Bayesian Neural Networks (BNNs) (Neal, 2012; Hernández-Lobato & Adams, 2015; Gal & Ghahramani, 2016; Blundell et al., 2015; Wang & Yeung, 2016; Wang et al., 2016; Lakshminarayanan et al., 2017; Wang & Yeung, 2020) predict responses and estimate their uncertainties using the variational distribution $q(\theta|\mathcal{D})$ that approximates the true weight posterior $p(\theta|\mathcal{D})$. Given an input sequence $\boldsymbol{x}=(x_1,\cdots,x_L)\in\mathcal{X}$, the probability of the output sequence $\boldsymbol{y}=(y_1,\cdots,y_T)\in\mathcal{Y}$ is defined as marginalization over the parameters and estimated by Bayesian Model Averaging (BMA) of size M:

$$p(\boldsymbol{y}|\boldsymbol{x}) = \int p(\boldsymbol{y}|\boldsymbol{x};\boldsymbol{\theta}) q(\boldsymbol{\theta}|\mathcal{D}) d\boldsymbol{\theta} \approx \frac{1}{M} \sum_{m=1}^{M} p(\boldsymbol{y}|\boldsymbol{x};\boldsymbol{\theta}^{(m)}), \quad \boldsymbol{\theta}^{(m)} \sim q(\boldsymbol{\theta}|\mathcal{D}). \tag{1}$$

Query-Level Uncertainty Estimation. Established techniques of uncertainty estimation (Gal et al., 2016) mainly quantify the uncertainty of input x (query-level uncertainty) by

$$\mathcal{H}[p(\mathbf{y}|\mathbf{x})] = \mathbb{E}_{\mathbf{y} \sim p(\mathbf{y}|\mathbf{x})}[-\log p(\mathbf{y}|\mathbf{x})]. \tag{2}$$

In the context of BNNs (Eqn. 1), the predictive distribution of y is the marginalized predictive distribution over the model parameters, and hence Eqn. 2 is defined as "total uncertainty" (Gal et al., 2016; Depeweg et al., 2017).

A model's uncertainty about a specific input cannot be solely attributed to the randomness of the approximate posterior $q(\theta|\mathcal{D})$, which is input-agnostic. For instance, when faced with a query "Name

¹We provide an implementation of our framework that is compatible with vLLM (Kwon et al., 2023) for efficient deployment.

a city in the UK?" (Yadkori et al., 2024), even if an infinite amount of data is observed (eliminating the randomness of the model parameters), the uncertainty of this question remains high, as there are many correct candidate answers. Hence to distinguish different sources of uncertainty, total uncertainty is decomposed into epistemic uncertainty and aleatoric uncertainty (Gal et al., 2016):

$$\underbrace{\mathcal{H}[p(\mathbf{y}|\mathbf{x})]}_{\text{Total Uncertainty}} = \underbrace{\mathbb{E}_{q(\boldsymbol{\theta}|\mathcal{D})}[\mathcal{H}[p(\mathbf{y}|\mathbf{x};\boldsymbol{\theta})]]}_{\text{Aleatoric Uncertainty}} + \underbrace{\mathcal{I}(\mathbf{y};\boldsymbol{\theta}|\mathbf{x})}_{\text{Epistemic Uncertainty}}.$$
(3)

Here, aleatoric uncertainty captures the intrinsic randomness in data and cannot be reduced even with more data observed. In contrast, epistemic uncertainty, defined as the mutual information $\mathcal{I}(\mathbf{y}; \boldsymbol{\theta}|\mathbf{x})$ between \mathbf{y} and $\boldsymbol{\theta}$, reflects the model's uncertainty about its own parameters, which can in principle be reduced by collecting more evidence. We use $\mathcal{U}(\mathbf{y}|\mathbf{x})$ defined in Definition 2.1 to denote any of the three uncertainties.

Definition 2.1 (Query-Level Uncertainty). Query-level uncertainty $\mathcal{U}(\mathbf{y}|\mathbf{x})$ is the uncertainty of the predictive distribution $p(\mathbf{y}|\mathbf{x})$ given an input query \mathbf{x} . Total Uncertainty (TU), Aleatoric Uncertainty (AU), and Epistemic Uncertainty (EU) in Eqn. 3 are all instances of query-level uncertainty.

Limitations of Query-Level Uncertainty. Using the chain rule for conditional entropy (Cover, 1999), the *query-level uncertainty* estimation can be decomposed token-by-token as

$$\mathcal{U}(\mathbf{y}|\mathbf{x}) = \sum_{t=1}^{T} \mathcal{U}(\mathbf{y}_t|\mathbf{y}_{< t}, \mathbf{x}). \tag{4}$$

However, the uncertainty term $\mathcal{U}(\mathbf{y}_t|\mathbf{y}_{< t},x)$ in Eqn. 4 requires marginalization over the random variable $\mathbf{y}_{< t}$, which is (i) computationally intractable, and (ii) only reflecting the quality of the input query. Hence, these query-level uncertainties are not proper indicators for evaluating a concrete output response \mathbf{y} .

2.2 BAYESIAN LARGE LANGUAGE MODELS

Bayesian Low-Rank Adaptation. For a pre-trained network layer with weight matrix W_0 , Low-Rank Adaptation (LoRA) (Hu et al., 2022) optimizes the parameters within a constrained low-rank subspace. Specifically, the weight update matrix is modeled by $\Delta W = BA$, where $\Delta W \in \mathbb{R}^{m \times n}$, $B \in \mathbb{R}^{m \times r}$, $A \in \mathbb{R}^{r \times n}$, and $r \ll \min(m, n)$. The output $z \in \mathbb{R}^{m \times 1}$ of forwarding the input vector $h \in \mathbb{R}^{n \times 1}$ is then

$$z = W_0 h + \Delta W h = W_0 h + BAh. \tag{5}$$

Leveraging LoRA's parameter efficiency, Bayesian LoRAs (Yang et al., 2023; Wang et al., 2024; Shi et al., 2024) aim to further integrate BNN's uncertainty estimation capabilities into LLMs without significant increasing memory complexity. The key idea is to model A and/or B as approximate distributions of the true weight posterior. The asymmetric Bayesianization, exemplified by BLoB (Wang et al., 2024) and TFB (Shi et al., 2024), models the elements of A with independent Gaussian distributions while keeping B deterministic. Specifically, we have

$$q(\mathbf{A}|\{\mathbf{M},\mathbf{\Omega}\}) = \prod_{ij} q(A_{ij}|M_{ij},\Omega_{ij}) = \prod_{ij} \mathcal{N}(A_{ij}|M_{ij},\Omega_{ij}^2), \tag{6}$$

where M and Ω share the same shape as A and denote the mean and standard deviation of the random variable A, respectively. To estimate this distribution, BLoB jointly trains the mean and covariance through the re-parameterization trick (Wang et al., 2024), while TFB uses a simple training-free maximal variance searching technique by fixing the approximate distribution to the family of low-rank isotropic Gaussian distributions (Shi et al., 2024).

Limited Scope of Existing Bayesian LLMs. Existing Bayesian LLMs have been primarily validated in downstream classification tasks of simple single- or multiple-choice problems, where uncertainty estimation is quantitatively assessed via the alignment of prediction confidence and accuracy (Yang et al., 2023; Balabanov & Linander, 2024; Wang et al., 2023; 2024; Shi et al., 2024). However, these methods have not yet demonstrated effective generalization to long-form generation tasks, i.e., LLM reasoning. Therefore, our Tokur, which estimates token-level uncertainties via weight perturbations, represents an initial step toward *extending Bayesian LLMs to long-form generation*, an area where uncertainty estimation remains largely unexplored and technically challenging.

3 TOKUR: TOKEN-LEVEL UNCERTAINTY ESTIMATION VIA LOW-RANK WEIGHT PERTURBATION

Sec. 3.1 introduces the key techniques of token-level uncertainty estimation. Sec. 3.2 then details how token-level uncertainties can be aggregated for response-level uncertainty estimation, and describes the underlying theoretical foundation. Finally, Sec. 3.3 presents our low-rank weight perturbation as posterior approximation. All proofs of propositions can be found in Appendix C.

3.1 TOKEN-LEVEL UNCERTAINTIES IN GENERAL

Given an approximate posterior $q(\boldsymbol{\theta}|\mathcal{D})$, a fixed input query $\boldsymbol{x} \in \mathcal{X}$ and a specific output response $\boldsymbol{y} = (y_1, y_2, \dots, y_T) \in \mathcal{Y}$ sampled from the base policy $p(\boldsymbol{y}|\boldsymbol{x})$, we denote the predictive distribution of the next token y_t produced by marginalization over weights as

$$\bar{p}(y_t|\boldsymbol{y}_{< t}, \boldsymbol{x}) \triangleq \mathbb{E}_{\boldsymbol{\theta} \sim q(\cdot|\mathcal{D})}[p(y_t|\boldsymbol{y}_{< t}, \boldsymbol{x}; \boldsymbol{\theta})]. \tag{7}$$

Assumption 3.1 (Stepwise Posterior Sampling). We assume that the weights θ sampled from the approximate posterior $q(\cdot|\mathcal{D})$ are not shared across decoding steps. Formally, the probability of a sequence is factorized as

$$\bar{p}(\boldsymbol{y}|\boldsymbol{x}) \triangleq \prod_{t=1}^{T} \bar{p}(y_t|\boldsymbol{x}, \boldsymbol{y}_{< t}) = \prod_{t=1}^{T} \Big\{ \mathbb{E}_{\boldsymbol{\theta}_t \sim q(\cdot|\mathcal{D})}[p(y_t|\boldsymbol{x}, \boldsymbol{y}_{< t}, \boldsymbol{\theta}_t)] \Big\}, \tag{8}$$

instead of adopting the joint formulation

$$\bar{p}(\boldsymbol{y}|\boldsymbol{x}) \triangleq \mathbb{E}_{\boldsymbol{\theta} \sim q(\cdot|\mathcal{D})}[p(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{\theta})].$$
 (9)

While both are valid probabilistic models, the joint formulation is incompatible with the autoregressive decoding mechanism of LLMs. Hence, we adopt the stepwise formulation in Assumption 3.1. To validate this assumption, we further conduct an ablation study comparing the stepwise formulation in Assumption 3.1 with the joint formulation, and report the results in Appendix E.4.4.

Given an input x and a partial output $y_{< t}$, for the time step t, we have the following three uncertainties:

• Total Uncertainty (TU) is the entropy of random variable y_t conditioned on x and $y_{< t}$:

$$TU(\mathbf{y}_t|\mathbf{y}_{< t}, \mathbf{x}) \triangleq \mathcal{H}[\bar{p}(\mathbf{y}_t|\mathbf{y}_{< t}, \mathbf{x})] = -\sum_{y_t \in \mathcal{V}} \bar{p}(y_t|\mathbf{y}_{< t}, \mathbf{x}) \log \bar{p}(y_t|\mathbf{y}_{< t}, \mathbf{x}), \quad (10)$$

• Aleatoric Uncertainty (AU) is the expectation of entropy of random variable y_t over the weights θ sampled from the approximate posterior $q(\cdot|\mathcal{D})$ as in Eqn. 3:

$$AU(\mathbf{y}_t|\mathbf{y}_{< t}, \mathbf{x}) \triangleq \mathbb{E}_{\boldsymbol{\theta} \sim q(\cdot|\mathcal{D})} [\mathcal{H}[p(\mathbf{y}_t|\mathbf{y}_{< t}, \mathbf{x}; \boldsymbol{\theta})]], \tag{11}$$

• Epistemic Uncertainty (EU) is the difference between TU and AU:

$$\mathrm{EU}(\mathbf{y}_t|\mathbf{y}_{\leq t},\mathbf{x}) \triangleq \mathrm{TU}(\mathbf{y}_t|\mathbf{y}_{\leq t},\mathbf{x}) - \mathrm{AU}(\mathbf{y}_t|\mathbf{y}_{\leq t},\mathbf{x}) = \mathcal{I}(\mathbf{y}_t;\boldsymbol{\theta}|\mathbf{y}_{\leq t},\mathbf{x}), \tag{12}$$

where V is the vocabulary and all the expectations are estimated with BMA.

3.2 TOKEN-LEVEL UNCERTAINTY FOR RESPONSE-LEVEL UNCERTAINTY ESTIMATION

Definition 3.1 (Response-Level Uncertainty). Given the token-level uncertainties $U(\mathbf{y}_t|\mathbf{y}_{< t},\mathbf{x})$ defined in Eqn. 10-12, we define response-level uncertainty as their cumulative sum across all tokens in the output sequence:

$$\widetilde{\mathcal{U}}(\boldsymbol{y}|\boldsymbol{x}) \triangleq \sum_{t=1}^{T} \mathcal{U}(\mathbf{y}_{t}|\boldsymbol{y}_{< t}, \boldsymbol{x}),$$
 (13)

where U can denote any of the considered uncertainty measures (TU, AU, or EU in Eqn. 10-12).

Proposition 3.1 (Response-Level Uncertainty as an Unbiased Estimator of Query-Level Uncertainty). Given an input query x, let $y \sim p(y|x)$ be a generated sample of length T. Then the response-level uncertainty $\widetilde{\mathcal{U}}$ (Definition 3.1) is an unbiased estimator of the query-level uncertainty \mathcal{U} (Definition 2.1), i.e.,

$$\mathbb{E}_{\boldsymbol{y} \sim p(\mathbf{y}|\boldsymbol{x})}[\widetilde{\mathcal{U}}(\boldsymbol{y}|\boldsymbol{x})] = \mathcal{U}(\mathbf{y}|\boldsymbol{x}). \tag{14}$$

Proposition 3.2 (Token-Level and Response-Level Uncertainty). Given an input query x, let $y \sim p(y|x)$ be a generated sample of length T. Let $\mathcal{U}(y_t|y_{< t},x)$ denote the token-level uncertainty as defined in Eqn. 10-12, with $\widetilde{\mathcal{U}}(y|x)$ as the corresponding response-level uncertainty (Definition 3.1). Our token-level uncertainty is equivalent to the response-level uncertainty when T = 1:

$$\widetilde{\mathcal{U}}(\mathbf{y}_1|\mathbf{x}) = \mathcal{U}(\mathbf{y}_1|\mathbf{x}). \tag{15}$$

The two propositions above provide key connections between Eqn. 13 and existing uncertainty estimation theory (Malinin & Gales, 2021; Ling et al., 2024). Proposition 3.1 shows that $\widetilde{\mathcal{U}}(\mathbf{y}|\mathbf{x})$ is an unbiased estimator of the true query-level uncertainty $\mathcal{U}(\mathbf{y}|\mathbf{x})$, ensuring its statistical consistency with the ideal formulation. Proposition 3.2 confirms that when the sequence length T=1, e.g., single-token prediction tasks such as multiple-choice QA (Yang et al., 2023; Wang et al., 2024), the estimator exactly recovers the token-level uncertainty, demonstrating structural consistency. These results support the validity and reliability of our approximation.

Advantages of Token-Level Uncertainty. Compared to Query-Level Uncertainty (Definition 2.1), token- and response-level uncertainties (i) avoid expensive marginalization over sequences (note the difference between $\mathbf{y}_{< t}$ in Eqn. 4 and $\mathbf{y}_{< t}$ in Eqn. 14) while still (ii) capturing the expected uncertainty conditioned on the generated output response. Moreover, since $\mathcal{U}(\mathbf{y}_t|\mathbf{y}_{< t},\mathbf{x})$ depends on the quality of the prefix $\mathbf{y}_{< t}$, (iii) the estimate retains rich semantic information, making it well-suited for entropy-based sequential decision-making (Kuhn et al., 2023; Ye et al., 2025) or hallucination detection (Farquhar et al., 2024; Kossen et al., 2024; Ye et al., 2025) in downstream tasks.

3.3 LOW-RANK WEIGHT PERTURBATION AS APPROXIMATION OF WEIGHT POSTERIOR

Suppose that we have an LLM policy p(y|x). To estimate the uncertainty of its output, we cast this model into a Bayesian framework by introducing weight perturbations. Due to the established advantages of *efficiency*, *performance preservation* of pre-perturbation model, and *effectiveness* of uncertainty estimation (Shi et al., 2024), we adopt a low-rank structure for the noise added to the model weights. Given a rank-r weight matrix $W_0 \in \mathbb{R}^{m \times n}$ of a neural network layer, we first perform compact Singular Value Decomposition (SVD) (Klema & Laub, 1980):

$$\mathbf{W}_0 = \mathbf{U} \operatorname{diag}(\mathbf{d}) \mathbf{V}^{\top}, \tag{16}$$

where $d \succ \mathbf{0} \in \mathbb{R}^{r \times 1}$ is the vector of singular values, and $U \in \mathbb{R}^{m \times r}$ and $V \in \mathbb{R}^{n \times r}$ both contain orthonormal columns, i.e., $U^{\top}U = V^{\top}V = I_r$. To ensure computational efficiency, we introduce a low-rank noise matrix $\epsilon \in \mathbb{R}^{n \times r'}$ whose rank $r' \ll r$ is significantly smaller than the rank of weight matrix, and whose entries are sampled i.i.d. from a Gaussian distribution of standard deviation of σ_q , which we refer to as *perturbation strength*, i.e., $\epsilon_{ij} \sim \mathcal{N}(0, \sigma_q^2), \forall i \in [n], j \in [r']$. The perturbed weight matrix is then constructed as

$$\mathbf{W} = \mathbf{W}_0 + \mathbf{U}' \boldsymbol{\epsilon}^{\top}, \tag{17}$$

where the matrix U' contains the top-r' columns of U. This perturbation transforms the deterministic W_0 to a variational low-rank isotropic Gaussian distribution W (Wang et al., 2024; Shi et al., 2024):

$$q(\text{vec}(\boldsymbol{W})|\sigma_q) = \mathcal{N}(\text{vec}(\boldsymbol{W})|\boldsymbol{\mu}_q, \boldsymbol{\Sigma}_q),$$
where $\boldsymbol{\mu}_q = \text{vec}(\boldsymbol{W}_0),$

$$\boldsymbol{\Sigma}_q = \sigma_q^2 \boldsymbol{I}_n \otimes \begin{bmatrix} \boldsymbol{I}_{r'} \\ \boldsymbol{0}_{m-r'} \end{bmatrix}.$$
(18)

Let θ denote the collection of all perturbed weight matrices across the model. By assuming the statistical independence among layers, the overall approximate posterior becomes

$$q(\boldsymbol{\theta}|\sigma_q) = \prod_i q(\text{vec}(\boldsymbol{W}^i)|\sigma_q). \tag{19}$$

Utilizing the Approximate Weight Posterior $q(\theta|\sigma_q)$. Notably, while we leverage the variational posterior formulation of Eqn. 19 to quantify uncertainty (detailed in Sec. 3.1), we use only the mean weights W_0 for decoding of each step rather than BMA as in Eqn. 1. This approach allows for a controlled study of the effects of uncertainty estimation itself, separate from the effects of BNNs. For the complete algorithmic description and overview, please refer to Appendix B.

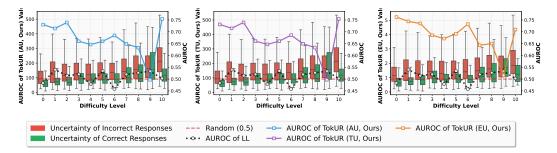


Figure 1: Distribution of TokUR's Uncertainty Scores and AUROC across Different Difficulty Levels, applied to Llama-3.2-1B-Instruct. Left: TokUR (AU, Ours); Middle: TokUR (TU, Ours); Right: TokUR (EU, Ours).

4 EXPERIMENTS

This section presents practical applications of our Tokur for LLM reasoning. For additional experimental results, please refer to Appendix E.

Datasets. We run our experiments on three mathematical reasoning benchmarks of varying difficulty levels: **GSM8K** (Cobbe et al., 2021) (grade-school arithmetic problems), **MATH500** (Lightman et al., 2023) (challenging high school/college mathematics competition problems), and 5,000-example subset of **DeepScaleR** (Luo et al., 2025) (high-difficulty problems from diverse sources). For these complex math problems, LLMs often need to perform multi-step reasoning (Wei et al., 2022b; Yao et al., 2023; Zhou et al., 2023) to reach the final answer. These tasks inherently involve long-form generation, therefore well-suited for evaluation of uncertainty estimation methods.

Models. We use two open-source LLMs in our experiments: Llama-3.2-1B-Instruct and Llama-3.1-8B-Instruct (Grattafiori et al., 2024). These models represent recent advances in downstream tasks and offer a good balance between performance and efficiency. Besides, their different model scales enable comparisons of uncertainty estimation across varying model sizes.

Implementation of our TokUR. We estimate token-level uncertainties by applying random perturbations as in Eqn. 17 to the query and key weight matrices ((W^Q, W^K)) (Vaswani et al., 2017) in all the attention layers of LLMs (Hu et al., 2022; Yang et al., 2023; Wang et al., 2024; Shi et al., 2024). For more details, please refer to Appendix D.1.

4.1 DO TOKUR'S UNCERTAINTIES ACCURATELY REFLECT RESPONSE QUALITY?

This section assesses if our TokUR's uncertainties reflect response quality in math reasoning tasks.

4.1.1 Tokur's Uncertainties and Question Difficulty

Experimental Setting. To better understand the relationships among uncertainty estimates, question difficulty, and their ability to distinguish correct from incorrect responses, we sample a subset of math questions from math-orz (Hu et al., 2025). A question's difficulty level is determined by the number of failed attempts out of 10 when using the <code>Qwen2.5-3B-Instruct</code> model. A difficulty level of 0 means the model solved the question every time, while a level of 10 indicates it failed on every attempt. We sample 500 questions per difficulty level, yielding a 5,500-question dataset. We then prompt <code>Llama-3.2-1B-Instruct</code> to solve each question with greedy decoding and apply <code>TokUR</code> to compute uncertainties for both correct and incorrect responses across difficulty levels. Fig. 1 summarizes the results.

Results. TokUR's uncertainty estimates remain positively correlated with question difficulty: for all three types of uncertainty (AU, TU, and EU), incorrect responses consistently exhibit higher uncertainty than correct ones across difficulty levels. In terms of discriminative power, AUROC values are consistently above random (0.5), confirming that TokUR provides useful signals for distinguishing correct from incorrect reasoning. Yet, AUROC tends to decrease as difficulty increases, especially in the mid-to-high range (levels 7–9), showing that uncertainty estimates become less reliable at separating outcomes on challenging tasks. Interestingly, the AUROC score shows a slight increase at

Table 1: Performance of Uncertainty Estimation Methods for Incorrect Reasoning Path Detection. AUROC, AUPRC, and ACC* are all reported as percentage (%). "ISO?" indicates whether the method utilizes Internal Signal Only for uncertainty estimation. We include the accuracy of CoT (i.e., greedy decoding with Chain-of-Thought prompting) in the first row for reference. Boldface and underlining denote the best and the second-best performance, respectively.

Method	ISO?	MATH500				GSM8K		DeepScaleR		
		AUROC	AUPRC	ACC*	AUROC	AUPRC	ACC*	AUROC	AUPRC	ACC*
				Llama-3.	2-1B-Inst	ruct				
CoT (Lower-Bound)	-	-	-	25.60±0.00	-	-	44.43±0.00	-	-	14.25±0.00
SE	X	47.29 ± 3.81	25.71 ± 2.33	24.13 ± 4.42	50.64 ± 4.44	45.09 ± 0.72	42.62 ± 0.16	46.30 ± 0.21	12.94 ± 0.23	12.58 ± 0.49
SAR	X	44.57 ± 2.04	24.03 ± 2.53	21.07 ± 1.62	50.28 ± 0.97	43.24 ± 0.89	43.95 ± 0.77	43.14 ± 1.42	12.34 ± 0.35	11.14 ± 0.47
U_{Ecc}	X	48.75 ± 1.05	25.79 ± 1.83	25.20 ± 0.33	49.05 ± 0.46	60.02 ± 0.44	59.62 ± 0.22	48.68 ± 0.24	13.77 ± 0.29	14.23 ± 0.45
U_{Deq}	X	60.57 ± 2.31	36.32 ± 2.59	30.93 ± 0.94	66.60 ± 0.36	75.72 ± 0.36	71.99 ± 0.39	56.88 ± 0.54	18.04 ± 0.63	16.50 ± 0.39
P(True)	/	54.38 ± 1.20	26.39 ± 1.26	27.60 ± 1.18	56.64 ± 0.04	48.22 ± 0.03	48.92 ± 0.00	59.58 ± 0.43	17.48 ± 0.25	17.52 ± 0.50
LLM-Check	/	56.41 ± 0.96	27.01 ± 1.22	31.33 ± 1.29	71.01 ± 0.02	61.29 ± 0.08	59.54 ± 0.00	55.76 ± 0.48	14.55 ± 0.26	17.30 ± 0.51
INSIDE	/	55.71 ± 4.69	28.82 ± 4.05	29.20 ± 4.33	53.66 ± 0.92	46.03 ± 0.23	45.79 ± 1.25	54.73 ± 0.82	15.50 ± 0.48	16.30 ± 0.35
PE	/	57.08 ± 0.89	26.88 ± 1.05	31.33 ± 0.82	71.21 ± 0.03	61.61 ± 0.08	59.85 ± 0.00	56.09 ± 0.46	14.74 ± 0.23	17.33 ± 0.92
LL	/	55.41 ± 0.54	25.88 ± 0.87	29.87 ± 0.82	69.01 ± 0.03	58.51 ± 0.09	57.38 ± 0.00	53.84 ± 0.47	13.93 ± 0.23	16.83 ± 0.48
Self-Certainty	/	71.17 ± 0.30	48.37 ± 0.50	38.13 ± 0.61	73.41 ± 0.00	68.38 ± 0.00	61.38 ± 0.00	71.93 ± 0.04	33.81 ± 0.08	21.76 ± 0.04
DeepConf	/	71.77 ± 0.12	46.00 ± 0.42	39.87 ± 0.46	75.70 ± 0.00	69.72 ± 0.00	62.77 ± 0.00	71.65 ± 0.04	29.99 ± 0.05	22.00 ± 0.04
TokUR (TU, Ours)	/	80.64 ± 0.29	56.79 ± 0.74	44.67 ± 0.46	75.07 ± 0.05	70.29 ± 0.07	62.31 ± 0.00	83.55 ± 0.02	47.56 ± 0.04	25.71 ± 0.02
TokUR (AU, Ours)	/	80.61 ± 0.27	56.73 ± 0.75	44.67 ± 0.46	75.03 ± 0.06	70.22 ± 0.05	$\overline{62.21\pm0.18}$	83.52 ± 0.02	47.48 ± 0.05	25.71 ± 0.02
Tokur (EU, Ours)	1	79.74 ± 0.21	$\overline{56.64\pm0.41}$	44.13 ± 0.83	$71.79{\scriptstyle\pm0.80}$	$\overline{66.40\pm1.02}$	59.74 ± 1.00	82.87 ± 0.32	46.76 ± 0.38	25.52±0.1
				Llama-3.	1-8B-Inst	ruct				
CoT (Lower-Bound)	-	-	-	48.60±0.00	-	-	85.69±0.00	-	-	24.86±0.00
SE	X	62.93 ± 0.90	55.21 ± 1.04	55.73 ± 0.83	55.61 ± 3.36	87.16±1.14	86.77 ± 1.01	67.68 ± 0.94	35.18 ± 1.00	35.55 ± 0.3
SAR	X	69.42 ± 2.19	63.74 ± 3.03	59.20 ± 1.06	60.16 ± 2.22	89.24 ± 0.74	87.99 ± 0.81	73.01 ± 0.28	42.89 ± 0.65	37.51 ± 0.12
U_{Ecc}	X	50.23 ± 2.23	49.48 ± 2.44	49.60 ± 2.04	47.47 ± 2.15	84.69 ± 0.89	84.87 ± 1.17	50.16 ± 0.66	25.08 ± 0.18	25.48 ± 0.5
U_{Deq}	X	58.62 ± 0.36	57.69 ± 0.90	53.47 ± 1.64	67.22 ± 1.06	92.24 ± 0.53	92.62 ± 0.88	59.14 ± 0.37	32.64 ± 0.43	29.75±0.36
P(True)	1	33.41 ± 0.25	36.05 ± 0.55	35.33 ± 0.19	41.94 ± 0.01	82.19 ± 0.00	82.77 ± 0.00	33.64 ± 0.20	18.06 ± 0.06	16.23 ± 0.03
LLM-Check	1	57.41 ± 0.44	49.69 ± 1.07	52.80 ± 1.38	73.98 ± 0.01	93.37 ± 0.01	93.23 ± 0.00	55.42 ± 0.27	26.46 ± 0.19	28.37 ± 0.40
INSIDE	1	62.94 ± 1.72	55.06 ± 3.19	57.33 ± 1.01	58.86 ± 2.11	87.44 ± 0.94	88.21 ± 0.90	67.05 ± 0.49	33.83 ± 0.42	34.13 ± 0.10
PE	1	57.98 ± 0.49	49.72 ± 0.84	53.07 ± 0.94	74.03 ± 0.01	93.37 ± 0.00	93.23 ± 0.00	55.90 ± 0.23	26.80 ± 0.16	28.65 ± 0.22
LL	1	55.36 ± 0.49	47.24 ± 0.90	51.07 ± 0.94	72.21 ± 0.02	92.64 ± 0.00	92.46 ± 0.00	52.82 ± 0.32	24.48 ± 0.13	26.85±0.19
Self-Certainty	1	76.41 ± 0.61	76.22 ± 0.87	69.07 ± 0.83	80.60 ± 0.11	95.65 ± 0.03	96.26 ± 0.09	76.72 ± 0.09	56.15 ± 0.30	39.03 ± 0.23
DeepConf	1	71.86 ± 0.70	69.57 ± 0.94	66.27 ± 1.15	$83.30{\scriptstyle\pm0.07}$	96.23 ± 0.02	96.56±0.09	73.05 ± 0.08	48.76 ± 0.10	37.48±0.14
TokUR (TU, Ours)	1	82.47 ± 0.47	79.62 ± 0.33	$74.00{\scriptstyle\pm0.69}$	81.01 ± 0.04	95.53 ± 0.05	95.54 ± 0.00	$85.33{\scriptstyle\pm0.07}$	65.25 ± 0.01	43.91 ± 0.09
TokUR (AU, Ours)	1	82.43±0.48	79.56 ± 0.35	$74.00{\scriptstyle\pm0.69}$	80.97 ± 0.02	95.52 ± 0.03	95.49 ± 0.09	85.31 ± 0.07	$\overline{65.20\pm0.02}$	43.89 ± 0.08
TokUR (EU, Ours)	/	$82.86 \!\pm\! 0.42$	81.35 ± 0.66	72.40 ± 1.20	78.31 ± 1.58	94.91±0.59	94.67 ± 0.77	84.92±0.28	65.57 ± 0.43	43.89±0.2

the highest difficulty level (10). This is likely a result of the imbalanced data distribution (mostly incorrect), where the model consistently produces high uncertainty, which causes a misleading high metric value.

4.1.2 TOKUR FOR INCORRECT REASONING PATH DETECTION

Experimental Setting. The preliminary study demonstrates that our Tokur's uncertainty estimation can reflect the quality of generated responses, with lower uncertainty generally associated with better outputs. One important application of uncertainty estimation is hallucination detection in LLMs (Farquhar et al., 2024; Kossen et al., 2024; Ye et al., 2025). In this context, we treat uncertainty as a scoring function to identify hallucinated (incorrect) responses for long-form reasoning tasks. We adopt three metrics: Area Under the Receiver Operating Characteristic Curve (AUROC), Area Under the Precision-Recall Curve (AUPRC), and Top-50% ACC (ACC*) (Farquhar et al., 2024; Ye et al., 2025; Hanley & McNeil, 1982; Boyd et al., 2013). AUROC and AUPRC measure the overall power of uncertainty scores in distinguishing correct from incorrect responses. In addition, we report Top-50% ACC, defined as the accuracy of the top 50% samples ranked by the corresponding score. This metric reflects the model's ability to prioritize higher-quality generations under a fixed budget. We repeat the experiments with three different random seeds to obtain the mean and standard deviation across runs.

Baselines. We systematically categorize our baselines into two distinct types: (i) those relying solely on the LLM's internal signals, including the most recent **Self-Certainty** (Kang et al., 2025), Deep Think With Confidence (**DeepConf**) (Fu et al., 2025), **LLM-Check** (Sriramanan et al., 2024), Degree Matrix (U_{Ecc} and U_{Deg}) (Lin et al., 2023), and INternal States for hallucInation DEtection (**INSIDE**) (Chen et al., 2024a), as well as classic **P(True)** (Kadavath et al., 2022), Predictive Entropy (**PE**) (Malinin & Gales, 2021), Log-Likelihood (**LL**) (Murray & Chiang, 2018), and (ii) those leveraging external signals, such as an auxiliary Natural Language Inference model (He et al., 2020): Semantic Entropy (**SE**) (Kuhn et al., 2023), and Shifting Attention to Relevance (**SAR**) (Duan et al., 2024). Note that, apart from the five baselines with underlines (Kang et al., 2025; Fu et al.,

Table 2: **Performance of Uncertainty Estimation Methods for Test-Time Scaling. Boldface** and underlining denote the best and the second-best performance, respectively.

Dataset	Score	Method	Llama-3.2-1B-Instruct			Llama-3.1-8B-Instruct		
Dutuset	Score	Method	N=16	N=64	N=256	N=16	N=64	N=256
	LL	Maj@N WBoN	47.10±0.85 47.10±0.85	54.11±0.52 54.15±0.55	58.89 ± 0.36 58.92 ± 0.37	86.74±0.62 86.74±0.62	90.48±0.48 90.48±0.49	91.01±0. 91.00±0.
	Self-Certainty	Maj@N WBoN	45.02±0.92 45.02±0.92	52.61±0.72 52.65±0.70	57.18±0.53 57.22±0.54	80.02±0.70 80.02±0.70	87.25±0.49 87.25±0.50	90.05±0 90.05±0
GSM8K (Pass@1: 44.43 / 85.69)	DeepConf	Maj@N WBoN	46.72±0.89 46.72±0.89	53.50±0.66 53.47±0.65	58.05±0.44 58.10±0.45	$86.24{\scriptstyle\pm0.66}\atop86.24{\scriptstyle\pm0.66}$	90.34±0.46 90.32±0.46	90.92±0 91.02±0
(1 855 @ 1. 44.457 65.67)	Tokur (TU, Ours)	Maj@N WBoN	$\frac{50.29{\pm}1.03}{50.29{\pm}1.03}$	57.18±0.45 57.22±0.45	$\substack{60.68 \pm 0.49 \\ 60.71 \pm 0.49}$	$\frac{87.68 \pm 0.57}{87.68 \pm 0.57}$	$\frac{90.67 \pm 0.45}{90.65 \pm 0.46}$	90.96±0 90.98±0
	Tokur (AU, Ours)	Maj@N WBoN	50.20±0.98 50.20±0.98	$\frac{57.21{\pm0.46}}{57.19{\pm0.44}}$	60.70±0.41 60.78±0.42	87.42±0.66 87.42±0.66	90.60±0.44 90.57±0.43	90.99±0 91.02±0
	Tokur (EU, Ours)	Maj@N WBoN	$50.38 {\scriptstyle \pm 0.92} \atop 50.38 {\scriptstyle \pm 0.92}$	56.92 ± 0.60 56.89 ± 0.54	59.88±0.52 59.91±0.58	88.06±0.57 88.06±0.57	90.69±0.47 90.67±0.48	91.07±0
	LL	Maj@N WBoN	26.42±0.84 26.42±0.84	33.28±0.97 33.30±1.10	38.56±0.75 38.58±0.73	50.92±1.77 50.92±1.77	59.36±0.74 59.46±0.78	64.10±0
	Self-Certainty	Maj@N WBoN	20.14±1.14 20.14±1.14	29.12±1.11 29.16±0.99	$\substack{36.68 \pm 0.83 \\ 36.80 \pm 0.80}$	44.00±1.82 44.00±1.82	55.56±1.08 55.58±1.06	62.66±0 62.52±0
MATH500 (Pass@1: 25.60 / 48.60)	DeepConf	Maj@N WBoN	$\substack{25.68 \pm 1.38 \\ 25.68 \pm 1.38}$	33.30±1.10 32.44±1.20	38.52 ± 0.43 37.08 ± 0.78	49.88±1.29 49.88±1.29	59.74±1.17 58.22±1.00	64.30±0
(1 100 0 11 20 100 1 10 100)	Tokur (TU, Ours)	Maj@N WBoN	$\frac{27.06 \pm 0.94}{27.06 \pm 0.94}$	$\frac{33.76 \pm 0.84}{33.60 \pm 0.82}$	39.18±0.70 39.20±0.65	$\frac{51.26{\pm}1.36}{51.26{\pm}1.36}$	59.44±1.31 59.44±1.30	63.86±0
	Tokur (AU, Ours)	Maj@N WBoN	$\frac{27.06 \pm 0.91}{27.06 \pm 0.91}$	33.64±0.76 33.48±0.73	39.12±0.72 39.10±0.69	51.16±1.45 51.16±1.45	59.42±1.16 59.44±1.19	64.00±0 63.92±0
	Tokur (EU, Ours)	Maj@N WBoN	$\substack{ 28.28 \pm 1.32 \\ 28.28 \pm 1.32 }$	35.44±0.79 35.44±0.78	39.44±0.88 39.38±0.87	52.40±1.39 52.40±1.39	$\frac{60.90 \pm 0.93}{61.04 \pm 0.88}$	65.32±0

2025; Kadavath et al., 2022; Malinin & Gales, 2021; Murray & Chiang, 2018), the others were originally designed for query-level de-hallucination in short-form QA tasks and are therefore not directly comparable to Tokur; we include them for completeness (see Appendix D.4 for details).

Results. As shown in Table 1, our proposed Tokur consistently outperforms baselines across AUROC, AUPRC, and ACC*. For example, on Llama-3.2-1B-Instruct, Tokur (TU) achieves an AUROC of 80.64% and an AUPRC of 56.79% on MATH500, clearly surpassing all baselines. On the larger Llama-3.1-8B-Instruct, the improvements are also substantial: Tokur (EU) attains 82.86% AUROC and 81.35% AUPRC on MATH500, establishing new state-of-the-art performance. These results highlight an important insight: Tokur provides a reliable and scalable uncertainty estimation framework, achieving strong performance without relying on external signals.

4.2 CAN TOKUR'S UNCERTAINTIES IMPROVE GENERATION QUALITY?

In this section, we explore the direct application of TokUR to reasoning tasks to enhance generation quality. Following previous works (Fu et al., 2025), we apply TokUR to measure the confidence of reasoning traces generated from a question and aggregate them via voting to obtain a final solution. In addition, we investigate the possibility of utilizing TokUR in an *online* manner to dynamically guide the generation process itself. Further details of online method are provided in E.3.

Baselines. We adopt Log-Likelihood (LL) as a baseline, given its widespread use as a proxy for generation quality (Manakul et al., 2023; Rafailov et al., 2023; Chen et al., 2024b). In addition, we compare against **Self-Certainty** (Kang et al., 2025) and **DeepConf** (Fu et al., 2025), two recent uncertainty-driven approaches for test-time scaling. As our study emphasizes model self-awareness of the boundaries of its knowledge, we do not include baselines that rely on external reward models (Guan et al., 2025; Puri et al., 2025; Beeching et al.; Uesato et al., 2022; Lightman et al., 2023).

Response Aggregation with Uncertainties. We first rank all N candidate responses using one of the scoring methods (**LL**, **Self-Certainty**, **DeepConf**, or TokUR) and retain the top-P% candidates. We then employ two common aggregation strategies: Weighted Best-of-N (**WBoN**) and Majority Voting (**Maj@N**) (Brown et al., 2024). WBoN performs weighted voting by assigning weights to the retained candidates according to their scores, whereas Maj@N simply selects the most frequent response among them, regardless of scoring.

Experimental Setting. We randomly sample 512 responses for each question in MATH500 and GSM8K with a decoding temperature of $\tau = 0.8$. For each N, we first retain the **top-10%** of samples

ranked by their scores. From this subset, the final prediction is determined using either Maj@N or WBoN. Each experiment is repeated 10 times (sample w/o replacement using offline records).

Results. As shown in Table 2, accuracy consistently improves with larger N across both GSM8K and MATH500. Our TokUR-based selection methods achieve clear gains over all baselines, particularly in the low-sample regime (N=16), where they deliver up to 3–4 points of improvement. Notably, TokUR (EU) attains the best overall performance on both datasets, with strong advantages in the challenging MATH500 benchmark. In addition, results for Maj@N and WBoN are similar, indicating that both aggregation strategies are similarly effective once the top candidates are identified.

5 RELATED WORK

Uncertainty Estimation of LLMs. Uncertainty estimation in LLMs is gaining traction for improving model calibration in data-scarce adaptation tasks and for reducing hallucinations in text generation (Liu et al., 2025; Vashurin et al., 2025). One prominent approach is Bayesian Adaptation, which combines Bayesian inference with low-rank adaptation (LoRA) (Hu et al., 2022) to approximate weight posterior distributions efficiently, avoiding the high computational cost of full Bayesian modeling (Yang et al., 2023; Wang et al., 2024; Shi et al., 2024). To estimate uncertainty in generation, two main lines of work have emerged. The first focuses on verbalized uncertainty, where models are prompted to express confidence in natural language (Lin et al., 2022; Kadavath et al., 2022; Tian et al., 2023; Kapoor et al., 2024). The second line includes logits-based methods, which estimate uncertainty directly from the model's output distributions (Van Der Poel et al., 2022; Ren et al., 2023; Duan et al., 2024; Darrin et al., 2023). In parallel, other approaches aim to refine these estimation strategies. For instance, (Malinin & Gales, 2021) investigates techniques for estimating epistemic uncertainty in structured prediction tasks, while semantic entropy (Kuhn et al., 2023) captures uncertainty by leveraging invariance in meaning across paraphrases. More recently, (Zhang & Zhang, 2025) introduces a method that leverages the reasoning capabilities of LLMs to enhance uncertainty quantification, using chain-of-thought prompting to better reflect model confidence in multi-step tasks. These works complement verbalized and logits-based methods by offering orthogonal perspectives on how uncertainty can be interpreted and measured.

Uncertainty for Improving LLM Generation. Uncertainty estimation for improving LLM generation, while not entirely novel, has been predominantly limited to approaches based on log-probability or its variants, Self-Certainty (Kang et al., 2025) estimates confidence via KL divergence from a uniform distribution, DeepConf (Fu et al., 2025) aggregates top-K log-probabilities as scores. Beam search (Lowerre, 1976; Sutskever et al., 2014; Freitag & Al-Onaizan, 2017; Xie et al., 2023) selects higher-confidence sequences by retaining candidates with the largest cumulative log-probability. UAG (Yin et al., 2024) leverages abrupt log-probability changes to select appropriate demonstrations for in-context learning (Brown et al., 2020). UnCert-CoT (Zhu et al., 2025) alternates between greedy and Chain-of-Thought decoding based on log-probability scores. Our work differs fundamentally by estimating token-level uncertainties with rigorous theoretical foundations, representing a significant step toward extending Bayesian LLMs to long-form generation scenarios.

6 Conclusion

In this paper, we introduce a novel framework Tokur to quantify uncertainty in LLM reasoning generations. By incorporating low-rank random weight perturbation during the LLM decoding procedure, Tokur provides a new perspective for uncertainty estimation in auto-regressive long-form generation with sound theoretical grounding. Through comprehensive empirical evaluation, we demonstrate that Tokur's uncertainty estimations effectively reflect the quality of generated reasoning paths and can thereby improve reasoning performance in LLMs. These contributions extend Bayesian uncertainty estimation to long-form reasoning, providing both theoretical foundations and practical tools for more reliable, self-aware LLMs.

Limitations. Our work is subject to several limitations. First, although compatible with efficient deployment frameworks such as vLLM (Kwon et al., 2023), repeated weight perturbation sampling during inference still poses efficiency challenges for real-time use. Second, our token-level uncertainty aggregation may miss higher-level semantic or logical inconsistencies across multiple tokens or reasoning steps, limiting its utility in complex generation tasks. Finally, the problem of high-variance estimation in our Tokur remains unresolved, constraining reliability in real-world scenarios.

REFERENCES

- Oleksandr Balabanov and Hampus Linander. Uncertainty quantification in fine-tuned llms using lora ensembles. *arXiv preprint arXiv:2402.12264*, 2024.
- Edward Beeching, Lewis Tunstall, and Sasha Rush. Scaling test-time compute with open models. URL https://huggingface.co/spaces/HuggingFaceH4/blogpost-scaling-test-time-compute.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International conference on machine learning*, pp. 1613–1622. PMLR, 2015.
- Kendrick Boyd, Kevin H Eng, and C David Page. Area under the precision-recall curve: point estimates and confidence intervals. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2013, Prague, Czech Republic, September 23-27, 2013, Proceedings, Part III 13*, pp. 451–466. Springer, 2013.
- Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V Le, Christopher Ré, and Azalia Mirhoseini. Large language monkeys: Scaling inference compute with repeated sampling. *arXiv* preprint arXiv:2407.21787, 2024.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Chao Chen, Kai Liu, Ze Chen, Yi Gu, Yue Wu, Mingyuan Tao, Zhihang Fu, and Jieping Ye. Inside: Llms' internal states retain the power of hallucination detection. *arXiv preprint arXiv:2402.03744*, 2024a.
- Haoxian Chen, Hanyang Zhao, Henry Lam, David Yao, and Wenpin Tang. Mallowspo: Fine-tune your llm with preference dispersions. *arXiv* preprint arXiv:2405.14953, 2024b.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53, 2024.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. arXiv preprint arXiv:2110.14168, 2021.
- Thomas M Cover. *Elements of information theory*. John Wiley & Sons, 1999.
- Longchao Da, Xiaoou Liu, Jiaxin Dai, Lu Cheng, Yaqing Wang, and Hua Wei. Understanding the uncertainty of llm explanations: A perspective based on reasoning topology. *arXiv preprint arXiv:2502.17026*, 2025.
- Maxime Darrin, Pablo Piantanida, and Pierre Colombo. Rainproof: An umbrella to shield text generator from out-of-distribution data. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 5831–5857, 2023.
- Stefan Depeweg, José Miguel Hernández-Lobato, Finale Doshi-Velez, and Steffen Udluft. Decomposition of uncertainty for active learning and reliable reinforcement learning in stochastic systems. *stat*, 1050(1):11, 2017.
- Jinhao Duan, Hao Cheng, Shiqi Wang, Alex Zavalny, Chenan Wang, Renjing Xu, Bhavya Kailkhura, and Kaidi Xu. Shifting attention to relevance: Towards the predictive uncertainty quantification of free-form large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 5050–5063, 2024.
- Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017):625–630, 2024.

- Markus Freitag and Yaser Al-Onaizan. Beam search strategies for neural machine translation. *arXiv* preprint arXiv:1702.01806, 2017.
- Yichao Fu, Xuewei Wang, Yuandong Tian, and Jiawei Zhao. Deep think with confidence. *arXiv* preprint arXiv:2508.15260, 2025.
 - Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059. PMLR, 2016.
 - Yarin Gal et al. Uncertainty in deep learning. 2016.

- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Xinyu Guan, Li Lyna Zhang, Yifei Liu, Ning Shang, Youran Sun, Yi Zhu, Fan Yang, and Mao Yang. rstar-math: Small llms can master math reasoning with self-evolved deep thinking. *arXiv* preprint *arXiv*:2501.04519, 2025.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- James A Hanley and Barbara J McNeil. The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology*, 143(1):29–36, 1982.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*, 2020.
- José Miguel Hernández-Lobato and Ryan Adams. Probabilistic backpropagation for scalable learning of bayesian neural networks. In *International conference on machine learning*, pp. 1861–1869. PMLR, 2015.
- Bairu Hou, Yujian Liu, Kaizhi Qian, Jacob Andreas, Shiyu Chang, and Yang Zhang. Decomposing uncertainty for large language models through input clarification ensembling. *arXiv* preprint *arXiv*:2311.08718, 2023.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=nZeVKeeFYf9.
- Jingcheng Hu, Yinmin Zhang, Qi Han, Daxin Jiang, Xiangyu Zhang, and Heung-Yeung Shum. Open-reasoner-zero: An open source approach to scaling up reinforcement learning on the base model. *arXiv preprint arXiv:2503.24290*, 2025.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El-Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, Deep Ganguli, Danny Hernandez, Josh Jacobson, Jackson Kernion, Shauna Kravec, Liane Lovitt, Kamal Ndousse, Catherine Olsson, Sam Ringer, Dario Amodei, Tom Brown, Jack Clark, Nicholas Joseph, Ben Mann, Sam McCandlish, Chris Olah, and Jared Kaplan. Language models (mostly) know what they know, 2022.
- Zhewei Kang, Xuandong Zhao, and Dawn Song. Scalable best-of-n selection for large language models via self-certainty. *arXiv preprint arXiv:2502.18581*, 2025.
- Sanyam Kapoor, Nate Gruver, Manley Roberts, Katherine Collins, Arka Pal, Umang Bhatt, Adrian Weller, Samuel Dooley, Micah Goldblum, and Andrew Gordon Wilson. Large language models must be taught to know what they don't know. *arXiv preprint arXiv:2406.08391*, 2024.
- Virginia Klema and Alan Laub. The singular value decomposition: Its computation and some applications. *IEEE Transactions on automatic control*, 25(2):164–176, 1980.

- Jannik Kossen, Jiatong Han, Muhammed Razzak, Lisa Schut, Shreshth Malik, and Yarin Gal. Semantic entropy probes: Robust and cheap hallucination detection in llms. *arXiv preprint arXiv:2406.15927*, 2024.
 - Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. In *The Eleventh International Conference on Learning Representations*, 2023.
 - Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023.
 - Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
 - Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *The Twelfth International Conference on Learning Representations*, 2023.
 - Stephanie Lin, Jacob Hilton, and Owain Evans. Teaching models to express their uncertainty in words. *Transactions on Machine Learning Research*, 2022.
 - Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. Generating with confidence: Uncertainty quantification for black-box large language models. *arXiv preprint arXiv:2305.19187*, 2023.
 - Chen Ling, Xujiang Zhao, Xuchao Zhang, Wei Cheng, Yanchi Liu, Yiyou Sun, Mika Oishi, Takao Osaki, Katsushi Matsuda, Jie Ji, et al. Uncertainty quantification for in-context learning of large language models. *arXiv preprint arXiv:2402.10189*, 2024.
 - Ollie Liu, Deqing Fu, Dani Yogatama, and Willie Neiswanger. Dellma: Decision making under uncertainty with large language models. *arXiv preprint arXiv:2402.02392*, 2024.
 - Xiaoou Liu, Tiejin Chen, Longchao Da, Chacha Chen, Zhen Lin, and Hua Wei. Uncertainty quantification and confidence calibration in large language models: A survey. *arXiv preprint arXiv:2503.15850*, 2025.
 - Bruce T Lowerre. The harpy speech recognition system. Carnegie Mellon University, 1976.
 - Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y Tang, Manan Roongta, Colin Cai, Jeffrey Luo, Tianjun Zhang, Li Erran Li, et al. Deepscaler: Surpassing o1-preview with a 1.5 b model by scaling rl. *Notion Blog*, 2025.
 - Andrey Malinin and Mark Gales. Uncertainty estimation in autoregressive structured prediction. In *International Conference on Learning Representations*, 2021.
 - Potsawee Manakul, Adian Liusie, and Mark JF Gales. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models. *arXiv preprint arXiv:2303.08896*, 2023.
 - Kenton Murray and David Chiang. Correcting length bias in neural machine translation. *arXiv* preprint arXiv:1808.10006, 2018.
 - Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
 - Ian Osband, Zheng Wen, Seyed Mohammad Asghari, Vikranth Dwaracherla, Morteza Ibrahimi, Xiuyuan Lu, and Benjamin Van Roy. Epistemic neural networks. *Advances in Neural Information Processing Systems*, 36:2795–2823, 2023.
 - Isha Puri, Shivchander Sudalairaj, Guangxuan Xu, Kai Xu, and Akash Srivastava. A probabilistic inference approach to inference-time scaling of llms using particle-based monte carlo methods. *arXiv preprint arXiv:2502.01618*, 2025.

- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36:53728–53741, 2023.
- Jie Ren, Jiaming Luo, Yao Zhao, Kundan Krishna, Mohammad Saleh, Balaji Lakshminarayanan, and Peter J Liu. Out-of-distribution detection and selective generation for conditional language models. In *The Eleventh International Conference on Learning Representations*, 2023.
- Haizhou Shi, Yibin Wang, Ligong Han, Huan Zhang, and Hao Wang. Training-free bayesianization for low-rank adapters of large language models. *arXiv preprint arXiv:2412.05723*, 2024.
- Gaurang Sriramanan, Siddhant Bharti, Vinu Sankar Sadasivan, Shoumik Saha, Priyatham Kattakinda, and Soheil Feizi. Llm-check: Investigating detection of hallucinations in large language models. *Advances in Neural Information Processing Systems*, 37:34188–34216, 2024.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27, 2014.
- Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher Manning. Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 5433–5442, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.330. URL https://aclanthology.org/2023.emnlp-main.330.
- Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. Solving math word problems with process-and outcome-based feedback. arXiv preprint arXiv:2211.14275, 2022.
- Liam Van Der Poel, Ryan Cotterell, and Clara Meister. Mutual information alleviates hallucinations in abstractive summarization. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 5956–5965, 2022.
- Roman Vashurin, Ekaterina Fadeeva, Artem Vazhentsev, Lyudmila Rvanova, Daniil Vasilev, Akim Tsvigun, Sergey Petrakov, Rui Xing, Abdelrahman Sadallah, Kirill Grishchenkov, et al. Benchmarking uncertainty quantification methods for large language models with lm-polygraph. *Transactions of the Association for Computational Linguistics*, 13:220–248, 2025.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Hao Wang and Dit-Yan Yeung. Towards bayesian deep learning: A framework and some existing methods. *IEEE Transactions on Knowledge and Data Engineering*, 28(12):3395–3408, 2016.
- Hao Wang and Dit-Yan Yeung. A survey on bayesian deep learning. *ACM computing surveys (csur)*, 53(5):1–37, 2020.
- Hao Wang, SHI Xingjian, and Dit-Yan Yeung. Natural-parameter networks: A class of probabilistic neural networks. In *NIPS*, pp. 118–126, 2016.
- Xi Wang, Laurence Aitchison, and Maja Rudolph. Lora ensembles for large language model fine-tuning, 2023.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv* preprint arXiv:2203.11171, 2022.
 - Yibin Wang, Haizhou Shi, Ligong Han, Dimitris Metaxas, and Hao Wang. Blob: Bayesian low-rank adaptation by backpropagation for large language models. *arXiv preprint arXiv:2406.11675*, 2024.

- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *arXiv preprint arXiv*:2206.07682, 2022a.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022b.
- Yuxi Xie, Kenji Kawaguchi, Yiran Zhao, James Xu Zhao, Min-Yen Kan, Junxian He, and Michael Xie. Self-evaluation guided beam search for reasoning. *Advances in Neural Information Processing Systems*, 36:41618–41650, 2023.
- Miao Xiong, Zhiyuan Hu, Xinyang Lu, Yifei Li, Jie Fu, Junxian He, and Bryan Hooi. Can llms express their uncertainty? an empirical evaluation of confidence elicitation in llms. *arXiv preprint arXiv:2306.13063*, 2023.
- Yasin Abbasi Yadkori, Ilja Kuzborskij, András György, and Csaba Szepesvári. To believe or not to believe your llm. *arXiv preprint arXiv:2406.02543*, 2024.
- Adam X Yang, Maxime Robeyns, Xi Wang, and Laurence Aitchison. Bayesian low-rank adaptation for large language models. *arXiv preprint arXiv:2308.13111*, 2023.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36:11809–11822, 2023.
- Zihuiwen Ye, Luckeciano Carvalho Melo, Younesse Kaddar, Phil Blunsom, Sam Staton, and Yarin Gal. Uncertainty-aware step-wise verification with generative reward models. *arXiv preprint arXiv:2502.11250*, 2025.
- Zhangyue Yin, Qiushi Sun, Qipeng Guo, Zhiyuan Zeng, Xiaonan Li, Junqi Dai, Qinyuan Cheng, Xuan-Jing Huang, and Xipeng Qiu. Reasoning in flux: Enhancing large language models reasoning through uncertainty-aware adaptive guidance. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2401–2416, 2024.
- Boxuan Zhang and Ruqi Zhang. Cot-uq: Improving response-wise uncertainty quantification in llms with chain-of-thought. *arXiv* preprint arXiv:2502.17214, 2025.
- Tianhang Zhang, Lin Qiu, Qipeng Guo, Cheng Deng, Yue Zhang, Zheng Zhang, Chenghu Zhou, Xinbing Wang, and Luoyi Fu. Enhancing uncertainty-based hallucination detection with stronger focus. *arXiv preprint arXiv:2311.13230*, 2023.
- Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, and Yu-Xiong Wang. Language agent tree search unifies reasoning acting and planning in language models. *arXiv preprint arXiv:2310.04406*, 2023.
- Yuqi Zhu, Ge Li, Xue Jiang, Jia Li, Hong Mei, Zhi Jin, and Yihong Dong. Uncertainty-guided chain-of-thought for code generation with llms. *arXiv preprint arXiv:2503.15341*, 2025.

APPENDIX

757 758

759

760

761

762

763

764

765

766

767

768 769

770

771

772

773

774 775

776 777

778

779

781 782

783 784 785

786

787

788

789

790

791

792

793

794

796

798

803 804

805

806

808

In Appendix A, we describe the role of large language models (LLMs) in our work. In Appendix B, we present the full algorithmic description of our method with low-rank weight perturbation. In Appendix C, we provide detailed proofs for all propositions presented in the main paper. In Appendix D, we provide our implementation details of the experiments, including:

- implementation of our TokUR (Appendix D.1),
- dataset details (Appendix D.2),
- prompt templates used in LLM reasoning (Appendix D.3),
- baseline details (Appendix D.4),
- and evaluation metrics (Appendix D.5).

Finally, in Appendix E, we present additional empirical results, including:

- preliminary study on the uncertainty distributions produced by TokUR (Appendix E.1),
- **detailed numerical results** of the test-time scaling (Appendix E.2),
- online test-time scaling of TokUR (Appendix E.3),
- an ablation study on different components of our token-level uncertainties (Appendix E.4),
- and a case study of our token-level uncertainties (Appendix E.5).

A LLM USAGE DISCLOSURE

We used large language models (LLMs) solely to assist with polishing the writing of this paper, including improving grammar, clarity, and readability. The LLMs did not contribute to research ideation, experimental design, analysis, or the generation of scientific content. All technical contributions, claims, and conclusions are the authors' own.

B ALGORITHM DETAILS

Algorithm 1 Low-Rank Weight Perturbation as Approximation of Weight Posterior.

```
1: Input
```

- 2: The base model policy $p(\mathbf{y}|\mathbf{x})$;
- 3: The set of weight matrices to be Bayesianized $\{W_0^k\}_{k=1}^N$;
- 4: rank of noise matrix r';
- 5: The perturbation strength σ_a .

6: for i = 1 to N do

- 7: $U, \operatorname{diag}(d), V^{\top} \leftarrow \operatorname{SVD}(W_0^k)$.
- 8: $U' \leftarrow$ the first r' columns of matrix U.
- 9: Sample noise matrix $\epsilon \in \mathbb{R}^{n \times r'}$: $\epsilon_{ij} \sim N(0, \sigma_q)$.
- 10: Perturb the weight matrix: $\mathbf{W}^k \leftarrow \mathbf{W}_0^k + \mathbf{U}' \boldsymbol{\epsilon}^{\top}$. \triangleright Eqn. 17
- 11: Get the weight posterior: $q(\text{vec}(\boldsymbol{W}^k)|\sigma_a)$. \triangleright Eqn. 18
- 12: **end for**
- 13: **Output:** The overall approximate posterior: $q(\theta|\sigma_q) \leftarrow \prod_k q(\text{vec}(\mathbf{W}^k)|\sigma_q)$

C PROOF OF PROPOSITIONS

Lemma C.1 (Definition of Conditional Entropy (Cover, 1999)). Give $(y, x) \sim p(y, x)$, the conditional entropy $\mathcal{H}(y|x)$ is defined as

$$\mathcal{H}(\mathbf{y}|\mathbf{x}) = \sum_{\boldsymbol{x} \in \mathcal{X}} p(\boldsymbol{x}) \mathcal{H}(\mathbf{y}|\boldsymbol{x})$$

$$= \mathbb{E}_{\boldsymbol{x} \sim p(\mathbf{x})} [\mathcal{H}(\mathbf{y}|\boldsymbol{x})].$$
(20)

⊳ Eqn. 16

Algorithm 2 Particle Filtering for Inference-Time Scaling (Puri et al., 2025)

```
1: Input
```

810

811

812

813

814

815

816 817

818

819

820

821

822

823

824

825

826 827 828

829

830

831 832

833

835

836

837

838

839

840

841 842

843

844

845

846

847

848

849

850

852

853

854 855

856

858

859

861

862 863

- 2: The number of particles N;
- 3: A reward model \hat{r} ;
 - 4: A LLM p_M and a prompt c.
 - 5: Initialize N particles $\{x_1^i \sim p_M(\cdot|c)\}_{i=1}^N$.

 - 7: while not all particles stop do
 - Update rewards $\mathbf{w} = \{\widehat{r}(x_{1:t}^{(1)}), \widehat{r}(x_{1:t}^{(2)}), \dots, \widehat{r}(x_{1:t}^{(N)})\}.$ Compute softmax distribution $\boldsymbol{\theta} = \operatorname{softmax}(\mathbf{w}).$ 8:
 - 9:
 - Sample indices $\{j_t^{(i)}\}_{i=1}^N \sim \mathbb{P}_t(j=i) = \boldsymbol{\theta}_i$. 10:
 - Update the set of particles as $\{x_{1:t}^{(j_t^{(I)})}\}_{i=1}^N$. 11:
- Transition $\{x_{t+1}^i \sim p_M(\cdot|c, x_{1:t}^{(i)})\}_{i=1}^N$. 12:
- $t \leftarrow t + 1$. 13:
- 14: end while
 - 15: **Output:** The set of particles in the end.

Lemma C.2 (Chain rule of Conditional Entropy (Cover, 1999)). Let X and Y be two random variables, then the conditional entropy of the joint distribution $\mathcal{H}(X,Y)$ can be decomposed as:

$$\mathcal{H}(X,Y) = \mathcal{H}(X) + \mathcal{H}(Y|X) \tag{21}$$

Lemma C.1 (Cover, 1999) reveals the relationship between conditional entropy $\mathcal{H}(\mathbf{y}|\mathbf{x})$ and the entropy derived from conditional probability distributions. Lemma C.2 lays the foundation for estimating the uncertainties of sequences. The two lemmas together give us the following proposition.

Proposition C.1 (Decomposition of Query-Level Uncertainty, Eqn. 4). Suppose that we have an input sequence x and a model policy p(y|x). The sequence-level uncertainty $\mathcal{U}(y|x)$ can be decomposed token-by-token as:

$$\mathcal{U}(\mathbf{y}|\mathbf{x}) = \sum_{t=1}^{T} \mathcal{U}(\mathbf{y}_{t}|\mathbf{y}_{< t}, \mathbf{x}), \tag{22}$$

where $\mathcal{U}(\mathbf{y}_t|\mathbf{y}_{< t},\mathbf{x})$ is token-level uncertainty metric as defined in Eqn. 10 ~ Eqn. 12.

Proof. For Aleatoric Uncertainty (AU) and Total Uncertainty (TU) defined in Eqn. 10 and Eqn. 11, both are expressed in terms of entropy. Therefore, the decomposition of sequence-level uncertainty can be directly derived using the chain rule stated in the Lemma C.2.

For Epistemic Uncertainty (EU), also called *mutual information* defined in Eqn. 12, we proceed with the following derivation:

$$\mathcal{H}(p(\mathbf{y}|\mathbf{x})) = \mathcal{H}\left(\mathbb{E}_{p(\boldsymbol{\theta}|\mathcal{D})}[p(\mathbf{y}_1|\mathbf{x};\boldsymbol{\theta})] \cdot \dots \cdot \mathbb{E}_{p(\boldsymbol{\theta}|\mathcal{D})}[p(\mathbf{y}_T|\mathbf{y}_{< T}, \mathbf{x};\boldsymbol{\theta})]\right)$$
(23)

$$= \sum_{t}^{T} \mathcal{H}(\mathbb{E}_{p(\boldsymbol{\theta}|\mathcal{D})}[p(\mathbf{y}_{t}|\mathbf{y}_{< t}, \boldsymbol{x}; \boldsymbol{\theta})])$$
 (24)

$$= \sum_{t}^{T} \mathcal{I}(\mathbf{y}_{t}; \boldsymbol{\theta} | \mathbf{y}_{< t}, \boldsymbol{x}) + \sum_{t}^{T} \mathbb{E}_{p(\boldsymbol{\theta} | \mathcal{D})} [\mathcal{H}(p(\mathbf{y}_{t} | \mathbf{y}_{< t}, \boldsymbol{x}; \boldsymbol{\theta}))]$$
(25)

$$= \sum_{t}^{T} \mathcal{I}(\mathbf{y}_{t}; \boldsymbol{\theta} | \mathbf{y}_{< t}, \boldsymbol{x}) + \mathbb{E}_{p(\boldsymbol{\theta} | \mathcal{D})} \mathcal{H}(p(\mathbf{y} | \boldsymbol{x}; \boldsymbol{\theta}))$$
(26)

Finally, based on the definition of *mutual information*, we obtain:

$$\mathcal{I}(\mathbf{y}; \boldsymbol{\theta} | \boldsymbol{x}) = \mathcal{H}(p(\mathbf{y} | \boldsymbol{x})) - \mathbb{E}_{p(\boldsymbol{\theta} | \mathcal{D})} \mathcal{H}(p(\mathbf{y} | \boldsymbol{x}; \boldsymbol{\theta}))$$

$$= \sum_{t}^{T} \mathcal{I}(\mathbf{y}_{t}; \boldsymbol{\theta} | \mathbf{y}_{< t}, \boldsymbol{x})$$
(27)

Proposition 3.1 (Response-Level Uncertainty as an Unbiased Estimator of Query-Level Uncertainty). Given an input query x, let $y \sim p(y|x)$ be a generated sample of length T. Then the response-level uncertainty $\widetilde{\mathcal{U}}$ (Definition 3.1) is an unbiased estimator of the query-level uncertainty \mathcal{U} (Definition 2.1), i.e.,

$$\mathbb{E}_{\boldsymbol{y} \sim p(\mathbf{y}|\boldsymbol{x})}[\widetilde{\mathcal{U}}(\boldsymbol{y}|\boldsymbol{x})] = \mathcal{U}(\mathbf{y}|\boldsymbol{x}). \tag{28}$$

Proof. Based on Lemma C.1, for the token-level uncertainty $\mathcal{U}(\mathbf{y}_t|\mathbf{y}_{< t},\mathbf{x})$ defined in Eqn. 10~Eqn. 12, we have

$$\mathbb{E}_{\boldsymbol{y}_{< t} \sim p(\cdot|\boldsymbol{x})}[\mathcal{U}(\mathbf{y}_t|\boldsymbol{y}_{< t}, \boldsymbol{x})] = \sum_{\boldsymbol{y}_{< t} \in \mathcal{Y}} p(\boldsymbol{y}_{< t}|\boldsymbol{x})\mathcal{U}(\mathbf{y}_t|\boldsymbol{y}_{< t}, \boldsymbol{x})$$
(29)

$$= \mathcal{U}(\mathbf{y}_t | \mathbf{y}_{< t}, \boldsymbol{x}). \tag{30}$$

Therefore, the uncertainty of the sequence defined in Eqn. 13:

$$\mathbb{E}_{\boldsymbol{y} \sim p(\mathbf{y}|\boldsymbol{x})}[\widetilde{\mathcal{U}}(\boldsymbol{y}|\boldsymbol{x})] = \mathbb{E}_{p(\mathbf{y}|\mathbf{x})}[\sum_{t=1}^{T} \mathcal{U}(\mathbf{y}_{t}|\boldsymbol{y}_{< t}, \boldsymbol{x})]$$
(31)

$$= \sum_{t=1}^{T} \mathbb{E}_{p(\mathbf{y}|\mathbf{x})} [\mathcal{U}(\mathbf{y}_t|\mathbf{y}_{< t}, \mathbf{x})]$$
(32)

$$= \sum_{t=1}^{T} \mathbb{E}_{\boldsymbol{y}_{(33)$$

$$= \sum_{t=1}^{T} \mathcal{U}(\mathbf{y}_t | \mathbf{y}_{< t}, \boldsymbol{x})$$
 (34)

$$= \mathcal{U}(\mathbf{y}|\mathbf{x}),\tag{35}$$

where the final step follows from the chain rule of entropy (Proposition C.1).

Proposition 3.2 (Token-Level and Response-Level Uncertainty). Given an input query x, let $y \sim p(y|x)$ be a generated sample of length T. Let $\mathcal{U}(y_t|y_{< t},x)$ denote the token-level uncertainty as defined in Eqn. 10-12, with $\widetilde{\mathcal{U}}(y|x)$ as the corresponding response-level uncertainty (Definition 3.1). Our token-level uncertainty is equivalent to the response-level uncertainty when T=1:

$$\widetilde{\mathcal{U}}(\mathbf{y}|\mathbf{x}) = \mathcal{U}(\mathbf{y}_1|\mathbf{x}). \tag{36}$$

Proof. When the sequence length T=1, based on the definition of uncertainty of sequence in Eqn. 13, we have

$$\widetilde{\mathcal{U}}(\boldsymbol{y}|\mathbf{x}) = \sum_{t=1}^T \mathcal{U}(\mathbf{y}_t|\boldsymbol{y}_{< t}, \boldsymbol{x}) = \mathcal{U}(\mathbf{y}_1|\boldsymbol{x}).$$

This proposition implies that the sequence uncertainty collapses to token-level uncertainty when the output sequence length is 1, reflecting the structural consistency of the estimator.

Proposition C.2 (Approximate Distribution of the Weight W Perturbed by Low-Rank Noise, Eqn. 18). Given the weight matrix $W_0 \in \mathbb{R}^{m \times n}$, the low-rank noise matrix $\epsilon \in \mathbb{R}^{n \times r'}$ whose rank $r' \ll r$ is significantly smaller than the rank r of W_0 , and whose entries are sampled i.i.d. from a Gaussian distribution of standard deviation of σ_q : $\epsilon_{ij} \sim \mathcal{N}(0, \sigma_q^2), \forall i \in [n], j \in [r']$, we have the perturbed weighted matrix W as defined in Eqn. 17. The variational distribution $q(\text{vec}(W)|\sigma_q)$ defined on the weight matrix W is

$$q(\text{vec}(\boldsymbol{W})|\sigma_q) = \mathcal{N}(\text{vec}(\boldsymbol{W})|\boldsymbol{\mu}_q, \boldsymbol{\Sigma}_q),$$
where $\boldsymbol{\mu}_q = \text{vec}(\boldsymbol{W}_0),$

$$\boldsymbol{\Sigma}_q = \sigma_q^2 \boldsymbol{I}_n \otimes \begin{bmatrix} \boldsymbol{I}_{r'} \\ \boldsymbol{0}_{m-r'} \end{bmatrix}.$$
(37)

Proof. We begin with compact SVD decomposition of the weight matrix W_0 as described in Eqn. 16:

$$\mathbf{W}_0 = \mathbf{U} \operatorname{diag}(\mathbf{d}) \mathbf{V}^{\top}, \tag{38}$$

where $d \succ 0 \in \mathbb{R}^{r \times 1}$ is the vector of singular values, and $U \in \mathbb{R}^{m \times r}$, $V \in \mathbb{R}^{n \times r}$ are orthogonal matrices. We denote the first r' columns of U as $U' \in \mathbb{R}^{m \times r'}$ to analyze the updated matrix $U' \epsilon^{\top}$ in Eqn. 17.

Since each entry in ϵ has zero mean, it is evident that the updated matrix also has zero mean. Consequently, we have $\mu_q = \text{vec}(W_0) + 0 = \text{vec}(W_0)$.

Next, we focus on the proof of the variance Σ_q . Gien $U' = (u_1, u_2, \dots, u_{r'}) \in \mathbb{R}^{m \times r'}$, and $\epsilon = (\epsilon_1, \epsilon_2, \dots, \epsilon_{r'}) \in \mathbb{R}^{n \times r'}$ as defined above, we have the following properties:

$$U'U'^{\top} = \sum_{i=1}^{r'} \boldsymbol{u}_i \boldsymbol{u}_i^{\top} = \begin{bmatrix} \boldsymbol{I}_{r'} & \\ & \boldsymbol{0}_{m-r'} \end{bmatrix}, \tag{39}$$

$$\operatorname{vec}(\boldsymbol{U}'\boldsymbol{\epsilon}^{\top}) = \operatorname{vec}(\sum_{i=1}^{r'} \boldsymbol{u}_i \boldsymbol{\epsilon}_i^{\top}) = \sum_{i=1}^{r'} (\boldsymbol{\epsilon}_i \otimes \boldsymbol{u}_i). \tag{40}$$

We can now derive the covariance matrix as:

$$\Sigma_q = \text{Var}[\text{vec}(\boldsymbol{W})] = \text{Var}[\text{vec}(\boldsymbol{W}_0 + \boldsymbol{U}'\boldsymbol{\epsilon}^\top)] = \text{Var}[\text{vec}(\boldsymbol{U}'\boldsymbol{\epsilon}^\top)]$$
(41)

$$= \operatorname{Var}\left[\sum_{i=1}^{r'} \boldsymbol{\epsilon}_i \otimes \boldsymbol{u}_i\right] = \sum_{i=1}^{r'} \operatorname{Var}\left[\boldsymbol{\epsilon}_i \otimes \boldsymbol{u}_i\right]$$
(42)

$$= \sum_{i=1}^{r'} \left\{ \mathbb{E}_{\boldsymbol{\epsilon}_i} [(\boldsymbol{\epsilon}_i \otimes \boldsymbol{u}_i) (\boldsymbol{\epsilon}_i \otimes \boldsymbol{u}_i)^\top] - \mathbb{E}_{\boldsymbol{\epsilon}_i} [(\boldsymbol{\epsilon}_i \otimes \boldsymbol{u}_i)] \mathbb{E}_{\boldsymbol{\epsilon}_i} [(\boldsymbol{\epsilon}_i \otimes \boldsymbol{u}_i)^\top] \right\}$$
(43)

$$= \sum_{i=1}^{r'} \left\{ \mathbb{E}_{\boldsymbol{\epsilon}_i} [\boldsymbol{\epsilon}_i \boldsymbol{\epsilon}_i^{\top}] \otimes (\boldsymbol{u}_i \boldsymbol{u}_i^{\top}) - (\mathbb{E}_{\boldsymbol{\epsilon}_i} [\boldsymbol{\epsilon}_i] \mathbb{E}_{\boldsymbol{\epsilon}_i} [\boldsymbol{\epsilon}_i]^{\top}) \otimes (\boldsymbol{u}_i \boldsymbol{u}_i^{\top}) \right\}$$
(44)

$$= \sum_{i=1}^{r'} \sigma_q^2 \mathbf{I}_n \otimes (\mathbf{u}_i \mathbf{u}_i^{\top}) = \sigma_q^2 \mathbf{I}_n \otimes \sum_{i=1}^{r'} \mathbf{u}_i \mathbf{u}_i^{\top} = \sigma_q^2 \mathbf{I}_n \otimes \begin{bmatrix} \mathbf{I}_{r'} \\ \mathbf{0}_{m-r'} \end{bmatrix}. \tag{45}$$

D IMPLEMENTATION DETAILS

D.1 IMPLEMENTATION OF TOKUR'S TOKEN-LEVEL UNCERTAINTIES

Unless otherwise specified, we set the rank of low-rank noise to r'=8, the perturbation strength $\sigma_q=0.1$, and the number of samples per uncertainty estimation to M=2. For the test-time scaling experiments in Sec. 4.2, we apply length normalization to TokuR to mitigate the bias introduced by varying sequence lengths. In contrast, the effect of length normalization may differ in hallucination detection tasks. To investigate this, we conduct additional ablation studies in Appendix E.4.3, examining the impact of length normalization in that setting. To ensure practical applicability in real-world scenarios, we implement our method as a seamless integration with vLLM (Kwon et al., 2023).

D.2 DATASETS

Table 3 shows the statistics of datasets in our experiments. These datasets collectively span a wide range of difficulty levels, from moderate to highly challenging, covering both elementary-level numerical reasoning and advanced symbolic mathematical tasks. In addition, the problem domains are diverse, including: algebra, geometry, and number theory. Such a design ensures that our experiments are comprehensive and representative, facilitating a thorough assessment of the model's capability across varied reasoning scenarios.

Table 3: Statistics of the datasets used in our experiments.

Dataset	Samples Used	Split	Task Type	Language	Level	
GSM8K	1,300	Training split	Mathematical Reasoning	English	Moderate	
MATH500	500	Full set	Mathematical Reasoning	English	Difficult	
DeepScaleR	5,000	First 5,000 samples	Mathematical Reasoning	English	Highly Challenging	

D.3 PROMPT TEMPLATES

In this work, we use the following prompts published by Meta ².

Prompt Example

972

973974975976977978979980

981 982

983

984 985

986 987

989 990

991

992

993

994 995

996

997 998 999

1000

1001

1002 1003 1004

1005

1007

1008

1009

1010

1011

1012

1013

1014

1015

1016

1017

1020

1023

1024 1025 Solve the following math problem efficiently and clearly:

-For simple problems (2 steps or fewer):

Provide a concise solution with minimal explanation.

-For complex problems (3 steps or more): Use this step-by-step format:

.....

Step 1: [Concise description]
[Brief explanation and calculations]

Step 2: [Concise description]
[Brief explanation and calculations]

•••

Regardless of the approach, always conclude with:

Therefore, the final answer is: answer. I hope it is correct.

Where [answer] is just the final number or expression that solves the problem.

D.4 BASELINES

We compare our uncertainty estimation approach against several baseline methods:

- Log-Likelihood (LL) (Murray & Chiang, 2018): Mean of token-wise log-probabilities of the output sequence, representing the model's overall confidence in its generation.
- Predictive Entropy (PE) (Malinin & Gales, 2021): Mean entropy of the predicted distribution of each token.
- **P(True)** (Kadavath et al., 2022): Directly queries the model about the correctness of its own output and uses the predicted probability of the token "True", normalized by the sum of probabilities of token "True" and "False", as a confidence score.
- **Self-Certainty** (Kang et al., 2025): Quantifies confidence using the KL divergence between the predicted token distribution and a uniform distribution over the vocabulary at each decoding step.
- DeepConf (Fu et al., 2025): Computes confidence scores by aggregating the log-probabilities
 of the top-K candidate tokens at each decoding step.
- The Degree Matrix (Lin et al., 2023): Utilizes the degree matrix of the graph Laplacian of the similarities matrix of responses.
- LLM-Check (Sriramanan et al., 2024): We faithfully reproduced the official implementation for comparison.
- **INSIDE** (Chen et al., 2024a): INSIDE is a method to estimate query-level uncertainty. We tailored INSIDE to our setting by asking the LLM to verify the same response multiple times and then calculating the semantic entropy across these verification attempts.

²https://huggingface.co/datasets/meta-llama/Llama-3.2-1B-Instruct-evals

- Semantic Entropy (SE) (Kuhn et al., 2023): We adapted SE to our setting by prompting the LLM to verify the same response multiple times and computing the semantic entropy of these verification attempts. While this provides a signal of response quality, we note that SE requires an external NLI or embedding model, giving it an inherent advantage compared to our method.
- SAR (Duan et al., 2024): We adapted SAR to our setting by computing sentence-level SAR scores over multiple verification attempts, following a similar procedure to SE. While this method provides a meaningful proxy for uncertainty, it requires an external semantic similarity model, which raises fairness concerns compared to our approach, which operates solely with the base LLM.

D.5 EVALUATION PARSING AND METRICS

Parsing. To automate the evaluation of outputs generated by large language models, we design specific prompts (see Appendix D.3) that constrain the model to follow a fixed structure and require it to place the final answer within a \box{}. Considering that in mathematical reasoning tasks, the same answer can be expressed in various forms, we standardize all answers into a canonical form before comparison (Beeching et al.). During the evaluation, we assess the correctness from two perspectives: numerical equality and symbolic equality, to label each generation as "True" or "False".

Metrics. To comprehensively assess model performance in binary classification tasks, we adopt the following metrics: Area Under the Receiver Operating Characteristic Curve (AUROC), Area Under the Precision-Recall Curve (AUPRC), and Top 50% Accuracy (Farquhar et al., 2024; Ye et al., 2025; Hanley & McNeil, 1982; Boyd et al., 2013).

• AUROC measures the trade-off between true positive rate (TPR) and false positive rate (FPR) at various threshold settings. Formally, for a set of predictions with associated confidence scores, AUROC is computed as:

$$AUROC = \int_0^1 TPR(FPR^{-1}(x)) dx, \tag{46}$$

where TPR and FPR are defined as:

$$TPR = \frac{TP}{TP+FN}, \quad FPR = \frac{FP}{FP+TN}.$$

• AUPRC evaluates the trade-off between precision and recall, which is particularly useful in imbalanced datasets. It is calculated as:

$$AUPRC = \int_0^1 Precision(Recall^{-1}(x)) dx, \tag{47}$$

where precision and recall are defined as:

Precision =
$$\frac{TP}{TP+FP}$$
, Recall = $\frac{TP}{TP+FN}$.

• Top 50% Accuracy evaluates the correctness of the top half predictions ranked by confidence. Let N be the total number of predictions and S be the set of indices corresponding to the top $\lceil N/2 \rceil$ predictions with highest confidence. The metric is defined as:

Top 50% Accuracy =
$$\frac{1}{|S|} \sum_{i \in S} \delta(\widehat{y}_i = y_i),$$
 (48)

where \hat{y}_i is the predicted label and y_i is the ground-truth label.

E ADDITIONAL EXPERIMENTAL RESULTS

E.1 Preliminary Study: Distribution of Uncertainties

We conduct a preliminary study to examine the relationship between responses' token-level uncertainties and their correctness. We generate responses on the GSM8K dataset using a greedy decoding strategy with Llama-3.2-1B-Instruct, and label each response as correct or incorrect based

on an exact match with the ground-truth answer. Considering the class imbalance in the model responses, we construct a balanced subset for visualization. Specifically, we retain all incorrect responses and randomly sample an equal number of correct responses. We compute the TokUR (EU) and TokUR (AU) with our proposed token-level uncertainties in Eqn. 13, and plot the results in the Normalized EU-AU space (Fig. 3). We observe that both TokUR (EU) and TokUR (AU) show a better-than-chance separation between correct and incorrect outputs. Although some overlap exists, their distribution peaks differ significantly, indicating that our uncertainty estimates meaningfully correlate with generation quality.

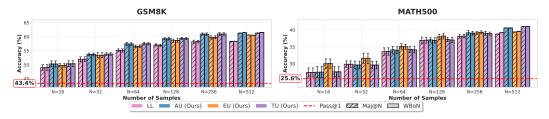


Figure 2: Performance on GSM8K (Left) and MATH500 (Right) when scaling up sample size N at test time of Llama-3.2-1B-Instruct. Our Tokur (AU, EU, and TU) consistently outperforms the LL baseline, particularly when N is small. Please refer to Table 4 for detailed numerical results.

E.2 TEST-TIME SCALING VIA UNCERTAINTY ESTIMATION

We provide an additional visualization of the test-time scaling results in Fig. 2. While the complete numerical results are reported in Table 2, this new figure offers an intuitive view of how accuracy improves with increasing numbers of test-time samples ($N \in 16, 32, 64, 128, 256, 512$). All experiments use Llama-3.2-1B-Instruct as the base model. For reference, the Pass@1 baseline accuracy (GSM8K: 44.43%; MATH500: 25.60%) is also shown as red dashed lines, highlighting the gains achieved through test-time scaling. Moreover, we provide an extended version of the results in Table 4, which builds on Table 2 to include additional test-time sample configurations.

E.3 TOKUR FOR TEST-TIME SCALING (ONLINE)

One popular approach to improving model performance uses a Process Reward Model (PRM) to score each intermediate step during multi-step generation (Guan et al., 2025; Puri et al., 2025; Beeching et al.; Uesato et al., 2022; Lightman et al., 2023), thereby guiding the model's reasoning path. In this section, we explore an alternative: guiding the generation process using uncertainty as an intrinsic reward, without relying on an explicit reward model.

Experimental Setting. Particle Filtering (PF) (Puri et al., 2025) is an inference-time scaling method for LLM reasoning (details in Appendix B). Building upon this algorithm, we use uncertainty as the score for each particle at each step to guide the model's generation process. We set the number of particles to N=16 and the decoding temperature to $\tau=0.8$. We repeat the experiments with three different random seeds to obtain the mean and standard deviation across runs.

Results. Table 5 shows the results. Compared to LL, our TokUR, especially TokUR (EU), yields a slight performance gain. Given that guiding generation through stepwise scoring is inherently challenging, we consider the lack of a significant performance gain from uncertainty estimation to be acceptable. Nevertheless, we believe this experiment offers valuable insights that may inform the future design of process reward models.

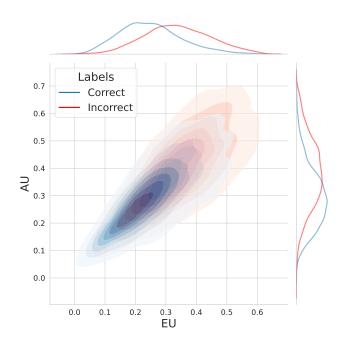


Figure 3: Distribution of responses from GSM8K (Cobbe et al., 2021) plotted in the Length Normalized EU-AU uncertainty space, as quantified by our token-level uncertainty metrics (Eqn. 13).

Table 5: TokUR as Implicit Reward for Test-Time Scaling (Online), on MATH500.

Intrinsic Reward	BoN	WBoN		
LL	26.27±0.25	26.27±0.41		
Tokur (TU, Ours) Tokur (AU, Ours) Tokur (EU, Ours)	$\begin{array}{c} 27.93 {\pm} 0.25 \\ 25.20 {\pm} 0.99 \\ \textbf{28.93} {\pm} \textbf{0.08} \end{array}$	28.13±0.38 25.13±0.74 29.20 ±0.98		

E.4 ABLATION STUDY

This section presents an ablation study on our token-level uncertainty estimation method using low-rank perturbations. Appendix E.4.1 examines the effect of varying perturbation strength σ_q , while Appendix E.4.2 analyzes the impact of different decoding temperatures. In Appendix E.4.3, we investigate the effect of length normalization in the context of hallucination detection tasks. Finally, in Appendix E.4.4, we assess the validity of Assumption 3.1 as it pertains to TokUR.

E.4.1 The Effect of Perturbation Strength σ_q on Uncertainty Estimation

To investigate the impact of perturbation strength on uncertainty estimation, we conducted a series of experiments under varying σ_q settings, as shown in Fig. 4. First, we computed the average uncertainty estimates (TU, AU, and EU) on samples generated from the GSM8K test dataset using Llama-3.2-1B-Instruct. As illustrated in Fig. 4 Left, the model's uncertainty increases steadily with higher perturbation strength. However, once σ_q exceeds a critical threshold (e.g., 0.2), a sharp rise in uncertainty is observed. This rise illustrates that the current approximate distribution of the weights has deviated too far from the pre-trained point estimation of the parameters, leading to unreliable uncertainty estimates.

We further evaluate the effect of perturbation strength on downstream task performance. Specifically, we assess how effectively the uncertainty estimates, obtained under different σ_q values, can be used as scoring signals to distinguish between correct and incorrect samples, as described in Sec. 4.1.2. As shown in Fig. 4 **Right**, for Tokur (EU), too small σ_q does not lead to meaningful

Table 4: **Test-Time Scaling for GSM8K and MATH500.** Performance comparison of different methods with varying numbers of test-time samples (N = 16 to 512) using Llama-3.2-1B-Instruct as the base model. Methods evaluated include log-likelihood (LL) and three variants of TokUR (TU, AU and EU) with both Maj@N and WBoN strategies. **Boldface** and <u>underlining</u> denote the best and the second-best performance, respectively.

Dotoset	C	Math - 2	Number of Samples N							
Dataset	Score	Method	N=16	N=32	N=64	N=128	N=256	N=512		
		Ll	ama-3.2-1	B-Instruct						
GSM8K (Pass@1: 44.43)	LL	Maj@N WBoN	47.10±0.85 47.10±0.85	50.45±0.64 50.45±0.64	54.11±0.52 54.15±0.55	$\begin{array}{c} 56.77{\scriptstyle \pm 0.40} \\ 56.72{\scriptstyle \pm 0.42} \end{array}$	58.89 ± 0.36 58.92 ± 0.37	59.72±0.0 59.81±0.0		
	Self-Certainty	Maj@N WBoN	$^{45.02\pm 0.92}_{45.02\pm 0.92}$	$^{48.97 \pm 0.81}_{48.97 \pm 0.81}$	$\begin{array}{c} 52.61{\scriptstyle \pm 0.72} \\ 52.65{\scriptstyle \pm 0.70} \end{array}$	$\begin{array}{c} 55.22{\pm}0.67 \\ 55.30{\pm}0.64 \end{array}$	$57.18{\scriptstyle\pm0.53}\atop57.22{\scriptstyle\pm0.54}$	58.03±0.0 58.10±0.0		
	DeepConf	Maj@N WBoN	$^{46.72 \pm 0.89}_{46.72 \pm 0.89}$	$50.12{\scriptstyle\pm0.71}\atop50.12{\scriptstyle\pm0.71}$	$53.50{\scriptstyle \pm 0.66}\atop 53.47{\scriptstyle \pm 0.65}$	$56.10{\scriptstyle \pm 0.52}\atop 56.08{\scriptstyle \pm 0.49}$	$\begin{array}{c} 58.05{\scriptstyle \pm 0.44} \\ 58.10{\scriptstyle \pm 0.45} \end{array}$	58.97±0.0 59.05±0.0		
	Tokur (TU, Ours)	Maj@N WBoN	$\frac{50.29{\pm}1.03}{50.29{\pm}1.03}$	$\tfrac{53.72 \pm 0.77}{53.72 \pm 0.77}$	$57.18 \scriptstyle{\pm 0.45} \\ 57.22 \scriptstyle{\pm 0.45}$	$\begin{array}{c} 59.10{\scriptstyle \pm 0.60} \\ \textbf{59.21} {\scriptstyle \pm \textbf{0.66}} \end{array}$	$\frac{60.68 \pm 0.49}{60.71 \pm 0.49}$	61.23±0. 61.31±0.		
	Tokur (AU, Ours)	Maj@N WBoN	$50.20{\scriptstyle \pm 0.98}\atop 50.20{\scriptstyle \pm 0.98}$	$53.77{\scriptstyle \pm 0.90}\atop 53.77{\scriptstyle \pm 0.90}$	$\frac{57.21{\pm}0.46}{57.19{\pm}0.44}$	$\begin{array}{c} 58.99{\scriptstyle \pm 0.61} \\ \underline{59.13}{\scriptstyle \pm 0.67} \end{array}$	$\substack{60.70 \pm 0.41 \\ \textbf{60.78} \pm \textbf{0.42}}$	61.38±0.		
	Tokur (EU, Ours)	Maj@N WBoN	$50.38 {\scriptstyle \pm 0.92}\atop 50.38 {\scriptstyle \pm 0.92}$	$52.98{\scriptstyle\pm0.67}\atop52.98{\scriptstyle\pm0.67}$	$\begin{array}{c} 56.92{\scriptstyle \pm 0.60} \\ 56.89{\scriptstyle \pm 0.54} \end{array}$	$\begin{array}{c} 58.77{\scriptstyle \pm 0.38} \\ 58.70{\scriptstyle \pm 0.40} \end{array}$	$\begin{array}{c} 59.88{\scriptstyle \pm 0.52} \\ 59.91{\scriptstyle \pm 0.58} \end{array}$	60.69±0. 60.85±0.		
	LL	Maj@N WBoN	$\substack{26.42 \pm 0.84 \\ 26.42 \pm 0.84}$	$\tfrac{29.28 \pm 0.89}{29.28 \pm 0.89}$	$\substack{33.28 \pm 0.97 \\ 33.30 \pm 1.10}$	$\begin{array}{c} 37.04{\scriptstyle \pm 0.74} \\ 37.02{\scriptstyle \pm 0.84} \end{array}$	$38.56 \pm 0.75 \ 38.58 \pm 0.73$	39.00 ± 0 39.00 ± 0		
	Self-Certainty	Maj@N WBoN	$\substack{20.14 \pm 1.14 \\ 20.14 \pm 1.14}$	$\substack{23.24 \pm 1.44 \\ 23.24 \pm 1.44}$	$\substack{29.12 \pm 1.11 \\ 29.16 \pm 0.99}$	$\begin{array}{c} 33.80 {\pm} 0.89 \\ 33.82 {\pm} 0.82 \end{array}$	$\substack{36.68 \pm 0.83 \\ 36.80 \pm 0.80}$	38.60 ± 0 38.60 ± 0		
MATH500	DeepConf	Maj@N WBoN	$\substack{25.68 \pm 1.38 \\ 25.68 \pm 1.38}$	$\substack{28.36 \pm 0.91 \\ 28.08 \pm 1.08}$	$\begin{array}{c} 33.30{\scriptstyle \pm 1.10} \\ 32.44{\scriptstyle \pm 1.20} \end{array}$	$\begin{array}{c} 37.34{\pm}1.31 \\ 36.00{\pm}1.22 \end{array}$	$\substack{38.52 \pm 0.43 \\ 37.08 \pm 0.78}$	40.40±0 38.60±0		
(Pass@1: 25.60)	Tokur (TU, Ours)	Maj@N WBoN	$\frac{27.06 \pm 0.94}{27.06 \pm 0.94}$	29.18±1.06 29.18±1.06	$\frac{33.76 \pm 0.84}{33.60 \pm 0.82}$	37.62±0.70 37.60±0.79	39.18±0.70 39.20±0.65	39.00±0 39.40±0		
	Tokur (AU, Ours)	Maj@N WBoN	$\tfrac{27.06 \pm 0.91}{27.06 \pm 0.91}$	29.08±1.14 29.08±1.14	33.64 ± 0.76 33.48 ± 0.73	37.52 ± 0.82 37.64 ± 0.73	39.12±0.72 39.10±0.69	39.20±0 39.60±0		
	Tokur (EU, Ours)	Maj@N WBoN	$\substack{ 28.28 \pm 1.32 \\ 28.28 \pm 1.32 }$	31.36±1.05 31.36±1.05	35.44±0.79 35.44±0.78	38.00±0.77 37.86±0.84	39.44±0.88 39.38±0.87	39.60±0 40.00±0		
		Ll	ama-3.1-81	3-Instruct						
	LL	Maj@N WBoN	$86.74 \scriptstyle{\pm 0.62} \\ 86.74 \scriptstyle{\pm 0.62}$	$\substack{89.16 \pm 0.53 \\ 89.16 \pm 0.53}$	$90.48 \scriptstyle{\pm 0.48} \\ 90.48 \scriptstyle{\pm 0.49}$	$90.99{\scriptstyle\pm0.35}\atop90.99{\scriptstyle\pm0.36}$	$91.01{\scriptstyle\pm0.28}\atop91.00{\scriptstyle\pm0.29}$	91.00±0 91.00±0		
	Self-Certainty	Maj@N WBoN	$80.02{\scriptstyle \pm 0.70}\atop80.02{\scriptstyle \pm 0.70}$	$\substack{84.13 \pm 0.66 \\ 84.13 \pm 0.66}$	$87.25{\scriptstyle\pm0.49\atop87.25{\scriptstyle\pm0.50}}$	$89.22{\scriptstyle\pm0.40}\atop89.21{\scriptstyle\pm0.39}$	$90.05{\scriptstyle \pm 0.40}\atop90.05{\scriptstyle \pm 0.41}$	90.77±0 90.77±0		
GSM8K	DeepConf	Maj@N WBoN	86.24±0.66 86.24±0.66	88.74±0.64 88.74±0.64	90.34±0.46 90.32±0.46	90.88±0.46 90.90±0.45	90.92±0.28 90.94±0.28	91.01±0 91.02±0		
(Pass@1: 85.69)	Tokur (TU, Ours)	Maj@N WBoN	87.68±0.57 87.68±0.57	89.72±0.55 89.72±0.55	$\frac{90.67 \pm 0.45}{90.65 \pm 0.46}$	$\frac{91.06 \pm 0.38}{91.06 \pm 0.37}$	90.96±0.36 90.98±0.37	91.02±0		
	Tokur (AU, Ours)	Maj@N WBoN	87.42±0.66 87.42±0.66	$\begin{array}{c} 89.59{\pm}0.55 \\ 89.59{\pm}0.55 \end{array}$	90.60±0.44 90.57±0.43	91.01±0.31 91.04±0.35	$90.99{\scriptstyle\pm0.32}\atop90.98{\scriptstyle\pm0.30}$	90.93±0 91.02±0		
	Tokur (EU, Ours)	Maj@N WBoN	$88.06{\scriptstyle \pm 0.57}\atop88.06{\scriptstyle \pm 0.57}$	$89.88{\scriptstyle\pm0.39}\atop89.88{\scriptstyle\pm0.39}$	90.69±0.47 90.67±0.48	$91.19{\scriptstyle\pm0.40}\atop91.19{\scriptstyle\pm0.39}$	$\frac{91.07{\pm}0.33}{91.09{\pm}0.36}$	91.02±0 91.05±0		
	LL	Maj@N WBoN	50.92±1.77 50.92±1.77	55.24±0.51 55.24±0.51	59.36±0.74 59.46±0.78	$\substack{62.86 \pm 0.70 \\ 62.80 \pm 0.78}$	64.10±0.61 64.02±0.71	65.00±0 65.00±0		
	Self-Certainty	Maj@N WBoN	44.00±1.82 44.00±1.82	48.48±1.06 48.48±1.06	55.56±1.08 55.58±1.06	60.04±0.56 59.90±0.51	62.66±0.75 62.52±0.53	65.40±0		
MATH500 (Pass@1: 48.60)	DeepConf	Maj@N WBoN	49.88±1.29 49.88±1.29	55.04±1.44 54.42±1.46	59.74±1.17 58.22±1.00	62.40±0.63 60.90±1.02	64.30±0.63 63.14±0.55	65.20±0		
	Tokur (TU, Ours)	Maj@N WBoN	51.26±1.36 51.26±1.36	$\frac{55.54{\pm}0.70}{55.54{\pm}0.70}$	59.44±1.31 59.44±1.30	62.28±0.95 62.32±1.08	63.86±0.44 63.84±0.51	65.20±0		
	Tokur (AU, Ours)	Maj@N WBoN	51.16±1.45 51.16±1.45	55.52±0.66 55.52±0.66	59.42±1.16 59.44±1.19	62.32±1.07 62.34±1.16	$64.00{\pm0.44} \atop 63.92{\pm0.47}$	65.60±0		
	Tokur (EU, Ours)	Maj@N WBoN	52.40±1.39 52.40±1.39	57.02±0.61 57.02±0.61	$\frac{60.90{\scriptstyle \pm 0.93}}{61.04 {\scriptstyle \pm 0.88}}$	64.24±0.83 64.20±0.76	$\frac{65.32{\pm}0.80}{\textbf{65.48}{\pm}\textbf{0.75}}$	67.00±0		

improvements in log-likelihood, whereas an excessively large σ_q harms performance of all three TokUR variants (AU, TU and EU) by distorting the original semantic content. Based on these findings, we set $\sigma_q = 0.1$ for the experiments reported in Sec. 4.

Table 6: **Uncertainties for Incorrect Reasoning Path Detection.** AUROC, AUPRC, and ACC* are all reported as percentage (%), where ACC* (%) denotes the accuracy of the Top 50% generations identified by different uncertainty measures. Rows with shading indicate methods *without* Length Normalization (LN) for uncertainty estimation.

Method		MATH500		GSM8K			DeepScaleR		
	AUROC	AUPRC	ACC*	AUROC	AUPRC	ACC*	AUROC	AUPRC	ACC*
Llama-3.2-1B-In	struct								
SE	47.29±3.81	25.71±2.33	24.13±4.42	50.64±4.44	45.09±0.72	42.62±0.16	46.30±0.21	12.94±0.23	12.58±0.49
SAR	44.57 ± 2.04	24.03 ± 2.53	21.07 ± 1.62	50.28 ± 0.97	43.24 ± 0.89	$43.95{\scriptstyle\pm0.77}$	$43.14{\scriptstyle\pm1.42}$	12.34 ± 0.35	11.14±0.47
U_{Ecc}	48.75 ± 1.05	25.79 ± 1.83	25.20 ± 0.33	49.05 ± 0.46	60.02 ± 0.44	59.62 ± 0.22	$48.68{\scriptstyle\pm0.24}$	13.77 ± 0.29	14.23±0.45
U_{Deg}	60.57 ± 2.31	36.32 ± 2.59	30.93 ± 0.94	66.60 ± 0.36	75.72 ± 0.36	71.99 ± 0.39	56.88 ± 0.54	18.04 ± 0.63	16.50±0.39
P(True)	54.38 ± 1.20	26.39 ± 1.26	27.60 ± 1.18	56.64 ± 0.04	48.22 ± 0.03	$48.92{\scriptstyle\pm0.00}$	59.58 ± 0.43	17.48 ± 0.25	17.52 ± 0.50
LLM-Check	56.41 ± 0.96	27.01 ± 1.22	31.33±1.29	$71.01 {\pm} 0.02$	$61.29{\scriptstyle\pm0.08}$	59.54 ± 0.00	$55.76{\scriptstyle\pm0.48}$	$14.55 \!\pm\! 0.26$	17.30±0.5
INSIDE	55.71 ± 4.69	28.82 ± 4.05	29.20 ± 4.33	53.66 ± 0.92	46.03 ± 0.23	45.79 ± 1.25	54.73 ± 0.82	15.50 ± 0.48	16.30±0.35
PE	57.08 ± 0.89	26.88 ± 1.05	31.33 ± 0.82	71.21 ± 0.03	61.61 ± 0.08	59.85 ± 0.00	56.09 ± 0.46	14.74 ± 0.23	17.33±0.92
LL	55.41±0.54	25.88 ± 0.87	29.87 ± 0.82	69.01 ± 0.03	58.51±0.09	57.38±0.00	53.84 ± 0.47	13.93 ± 0.23	16.83±0.48
- LN	79.38 ± 0.27	54.64±0.75	43.73 ± 0.61	73.67 ± 0.00	68.88 ± 0.00	60.92 ± 0.00	82.62 ± 0.01	45.76 ± 0.02	25.43±0.02
Self-Certainty		48.37±0.50							
- LN	23.76±0.34	17.04 ± 0.10	11.20 ± 0.00	34.42 ± 0.00	34.20 ± 0.00	31.54±0.00	21.33 ± 0.02	8.57 ± 0.01	4.04 ± 0.00
DeepConf	71.77 ± 0.12	46.00±0.42	39.87 ± 0.46	75.70 ± 0.00	69.72 ± 0.00	62.77 ± 0.00	71.65 ± 0.04	29.99 ± 0.05	22.00±0.04
- LN		17.41 ± 0.10							
Tokur (TU, Ours)									
- LN		56.79±0.74							
Tokur (AU, Ours)									
- LN		56.73±0.75							
TokUR (EU, Ours)									
- LN		56.64±0.41							
Negative Length	76.27±0.34	49.55±1.00	41.87±0.46	65.69±0.00	56.72±0.00	56.87±0.18	78.74±0.02	35.97±0.07	24.48±0.04
Llama-3.1-8B-In	struct								
SE	62.93+0.90	55.21±1.04	55.73+0.83	55 61+3 36	87.16+1.14	86 77+101	67 68+0 94	35.18+1.00	35.55+0.33
SAR		63.74 ± 3.03							
U_{Ecc}		49.48±2.44							
U_{Deg}		57.69±0.90							
P(True)		36.05 ± 0.55							
LLM-Check		49.69 ± 1.07							
INSIDE		55.06±3.19							
PE		49.72±0.84							
LL		47.24±0.90							
- LN		78.80±0.36							
Self-Certainty		76.22 ± 0.87							
- LN		33.57 ± 0.42							
DeepConf		69.57 ± 0.42							
- LN		34.08 ± 0.44							
Tokur (TU, Ours)									
- LN		79.62 ± 0.83							
Tokur (AU, Ours)									
- LN Tokur (EU, Ours)		79.56 ± 0.35							
- LN		81.35 ± 0.66							
Negative Length		73.81±0.28							
riegative Length	10.11 ± 0.54	13.01 ± 0.28	00.00±0.40	13.04±0.01	ラン・3U±0.00	フン・ン4±0.00	01.20±0.20	J1.14±0.31	+1.J/±0.10

E.4.2 THE EFFECT OF TOKEN DECODING TEMPERATURE τ ON UNCERTAINTY ESTIMATION

During text generation with large language models, the decoding temperature introduces uncertainty into the model's output. In general, higher temperatures lead to more diverse responses. In this section, we investigate the relationship between decoding temperature τ and uncertainties estimated by our token-level approach. Specifically, we use Llama-3.2-1B-Instruct to answer questions from the MATH500 dataset under different decoding temperature settings and estimate the average uncertainty of the model's responses.

As shown in Fig. 5 **Left**, increasing the decoding temperature τ results in a notable rise in token-level Aleatoric Uncertainty (AU) of the model, whereas the Epistemic Uncertainty (EU) remains relatively unaffected. Additionally, we report the AUROC scores of various uncertainty estimation approaches across different temperature settings in Fig. 5 **Right**. These results indicate that varying

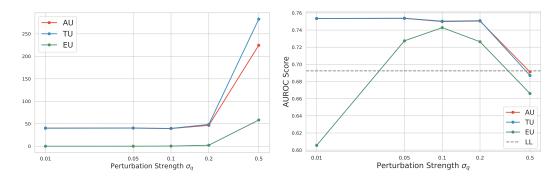


Figure 4: **Left:** Uncertainty estimation with different perturbation strength σ_q . **Right:** Influence of perturbation strength on uncertainty-based AUROC scores.

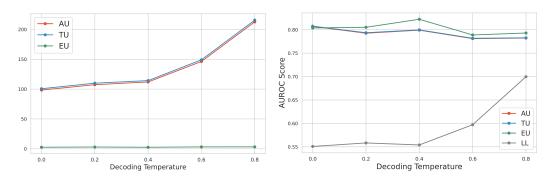


Figure 5: **Left:** Uncertainty estimations in different token decoding temperature τ . **Right:** Influence of token decoding temperature on uncertainty-based AUROC scores.

the temperature τ does not harm the performance of Tokur, highlighting its robustness to changes in decoding temperature.

E.4.3 ABLATION STUDY OF LENGTH NORMALIZATION

Length normalization is a standard technique for aggregating *token-level* uncertainty into *sequence-level* uncertainty (Fu et al., 2025; Kang et al., 2025), as it mitigates the bias introduced by sequence length when evaluating generation confidence. However, as described in Eqn. 13, we do not apply normalization when computing TokUR. To assess the impact of sequence length on uncertainty estimation, we therefore conduct an ablation study on length normalization.

Experimental Setup. We investigate the effect of length normalization on incorrect reasoning path detection across three datasets (MATH500, GSM8K, and DEEPSCALER), following the same settings as in Table 1. We compare Tokur with and without Length Normalization (LN), along with representative baselines. In addition, we introduce a naive baseline, **Negative Length**, which uses sequence length alone as a confidence signal.

Results. As shown in Table 6, the impact of length normalization varies significantly across methods. For both **LL** and Tokur, normalization consistently reduces AUROC and AUPRC, indicating that raw sequence length introduces a favorable bias that benefits uncertainty aggregation in de-hallucination tasks. This observation is further reinforced by the strong performance of the **Negative Length** baseline, which alone achieves competitive results across all datasets. In contrast, Self-Certainty and DeepConf show clear gains with normalization (e.g., Self-Certainty improves from 23.76 to 71.17 AUROC on MATH500), suggesting that normalization is essential for stabilizing their performance. Overall, these findings reveal that the role of length normalization is highly method-dependent.

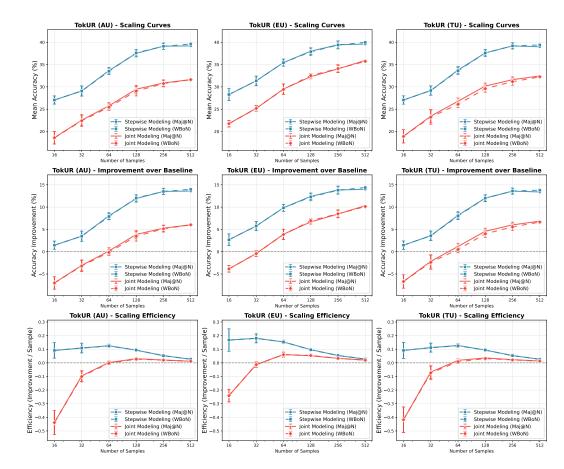


Figure 6: **Ablation of stepwise posterior sampling.** Comparison of *stepwise* vs. *joint* modeling on Llama-3.2-1B-Instruct across accuracy, improvement, and efficiency. Stepwise modeling consistently achieves better scaling performance, validating Assumption 3.1.

E.4.4 ABLATION STUDY OF STEPWISE POSTERIOR SAMPLING

To examine the validity of Assumption 3.1, we perform an ablation study comparing step-wise posterior sampling against the joint posterior formulation. Concretely, we evaluate test-time scaling on MATH500 dataset, using TokUR with both stepwise and joint modeling on Llama-3.2-1B-Instruct, while keeping all other settings consistent with Table 2.

As shown in Fig. 6, stepwise modeling consistently outperforms joint modeling across all uncertainty measures (AU, EU, and TU). Specifically, stepwise sampling achieves higher mean accuracy and larger improvements over the baseline, while also demonstrating superior scaling efficiency with increasing numbers of samples. These results provide strong empirical support for our assumption that posterior samples should not be shared across decoding steps, validating the design choice in Assumption 3.1.

E.5 CASE STUDY

In this section, we present several representative examples from the MATH500 and GSM8K datasets, along with their corresponding solutions generated by Llama-3.2-1B-Instruct. We estimate token-level uncertainty for each output using the definitions provided in Eqn. 10~Eqn. 12. The visualizations are shown in Fig. 7~Fig. 10, where Aleatoric Uncertainty (AU, in RED) and Epistemic Uncertainty (EU, in GREEN) are visualized as text-heatmap. The background shading of each token corresponds to the magnitude of its uncertainty: the darker the shade, the higher the uncertainty, indicating a lower model confidence for that token.

We observe that incorrect solutions consistently exhibit elevated uncertainty at or near the token where the wrong final answer is generated. For instance, as shown in Fig. 7, sharp spikes in uncertainties happens with the arithmetic error of reversing "9600 - 7200" into "7200 - 9600". In contrast, correct solutions tend to show lower uncertainty overall and maintain low uncertainty on key answer tokens.

Furthermore, incorrect outputs tend to contain a higher density of high-uncertainty tokens throughout the solution, whereas correct outputs are generally more consistent and confident. These observations suggest that our token-level uncertainty estimation method can serve as a useful signal for identifying potential reasoning failures or unreliable outputs, offering a valuable diagnostic tool for both model interpretability and downstream error detection.

 Problem: Oula and Tona work for a delivery service company, supplying different parts of their state with milk products. They are paid \$100 for each delivery made in a day. In a particular month, Oula made 96 deliveries while Tona made 3/4 times as many deliveries as Oula. Calculate their difference in pay that month.

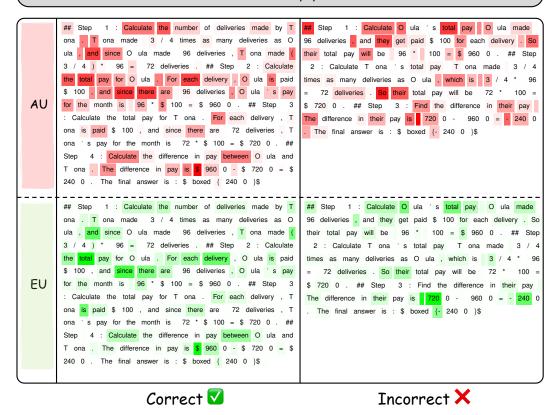


Figure 7: Case Study (1/4): The sample is from GSM8K, whose correct answer is 2400. In the incorrect solution, the model demonstrated significant uncertainty when mistakenly reversing "9600 - 7200" as "7200 - 9600", and also exhibited high uncertainty at the negative sign "—" in the final answer.

Problem: Kiarra is twice as old as Bea. Job is 3 times older than Bea. Figaro is 7 years older than Job. Harry is half as old as Figaro. If Kiarra is 30, how old is Harry?



Figure 8: Case Study (2/4): The sample is from GSM8K. In this example, the incorrect solution ignores the critical condition that "Figaro is 7 years older than Job," leading to the use of 45 instead of 52 in the final calculation. Notably, the model exhibits high uncertainty at the token "45" indicating a lack of confidence in its own response at that point.

158115821583

1584

1585

1586

1587

1588

1589

159015911592

1593

1594

1595

1596

1597

1598

1599 1600 1601

160216031604

1605

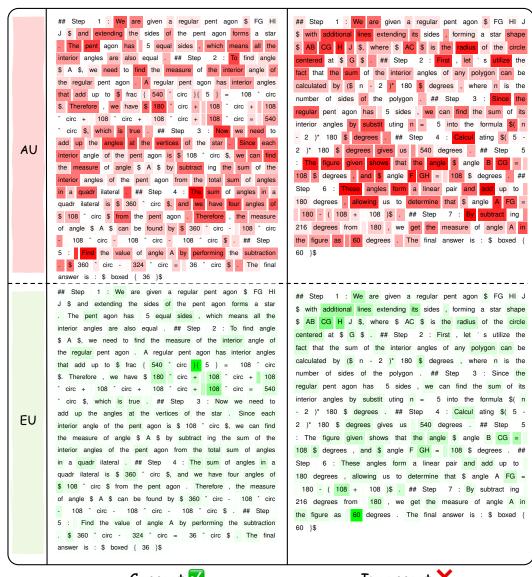
1606

1607

Problem: Simplify the following expression in x: \$19x + 1 - 4x - 81.\$1 : Combine like terms Combine the \$ x \$ 1 : Combine like terms To simplify the given terms . ## Step 2 : Simpl ify S implify the expression by expression , we need to combine the like terms with the performing the additions and sub tr actions from left to right . variable \$ x \$ and the constant terms 19 x + 1 - 4 x - 81 = (19 - 4)x + 1 - 81 Combine constants : \$ - 81 \$ We will = 15 x - 80 The final answer is : \$ boxed { 15 x terms \$ - 4 x \$ and \$ - 81 \$. ## Step 3 : Sub tract AU \$ - 4 x \$ from \$ - 81 \$: **\$** 19 x - 81 **\$** ## Step Combine like terms : \$ 19 x us with \$ 15 x \$ ## Step 5 : Therefore , the simplified expression is \$ boxed { 15 x }\$ 1 : Combine like terms To simplify the given ## Step 1 : Combine like terms Combine the \$ x \$ expression , we need to combine the like terms with the terms . ## Step 2 : Simpl ify S implify the expression by variable \$ x \$ and the constant terms performing the additions and sub tr actions from left to right . 19 x + 1 - 4 x - 81 = (19 - 4)x + 1 - 81 ΕU Sub tract = 15 x - 80 The final answer is : \$ boxed { 15 x terms \$ - 4 x \$ and \$ - 81 \$. ## Step 3 : 80 }\$ \$ - 4 x \$ from \$ - 81 \$: \$ 19 x - 81 \$ ## Step Combine like terms : \$ 19 x - 4 x \$ This leaves us with $$15 \times $#$ Step 5: Therefore , the simplified expression is \$ boxed { 15 x }\$ Incorrect X Correct V

Figure 9: Case Study (3/4): The sample is from MATH500. In this example, the incorrect solution gives its final answer "15x" in step 4. The model exhibits high uncertainty at the token next to "15x" because it overlooks the constant term. Furthermore, it can be observed that tokens associated with high uncertainty occur more frequently in the incorrect solution.

Problem: In regular pentagon \$FGHIJ\$, extending the sides of the pentagon, as shown, forms a star. What is the measure of angle \$A\$ in the figure?\n\n [IMAGE]



Correct ✓ Incorrect ×

Figure 10: Case Study (4/4): The sample is from MATH500. In this example, the model demonstrated notably high uncertainty at the incorrect answer token "60". In the correct solution on the left, the model had low uncertainty for the correct answer "36".