

BENCHMARKING UNCERTAINTY ESTIMATION IN LARGE LANGUAGE MODEL REPLIES FOR NATURAL SCIENCE QUESTION ANSWERING

Anonymous authors

Paper under double-blind review

ABSTRACT

Large Language Models (LLMs) are commonly used in Question Answering (QA) settings, increasingly in the natural sciences if not science at large. Reliable Uncertainty Quantification (UQ) is critical for the trustworthy uptake of generated answers. Existing UQ approaches remain weakly validated in scientific QA, a domain relying on fact-retrieval and reasoning capabilities. We introduce the first large-scale benchmark for evaluating UQ metrics in scientific QA studying calibration of UQ methods, providing an extensible open-source framework to reproducibly assess calibration. Our study spans up to 20 large language models of base, instruction-tuned and reasoning variants. Our analysis covers seven scientific QA datasets, including both multiple-choice and arithmetic question answering tasks. We evaluate and compare methods representative of prominent approaches on a total of 685,000 long-form responses, spanning different reasoning complexities representative of domain-specific tasks. At the token level, we find that instruction tuning induces strong probability mass polarization, reducing the reliability of token-level confidences as estimates of uncertainty. Models further fine-tuned for reasoning are exposed to the same effect, but the reasoning process appears to mitigate it depending on the provider. At the sequence level, we show that verbalized approaches are systematically biased and poorly correlated with correctness, while answer frequency (consistency across samples) yields the most reliable calibration. In the wake of our analysis, we study and report the misleading effect of relying exclusively on ECE as a sole measure for judging performance of UQ methods on benchmark datasets. Our findings expose critical limitations of current UQ methods for LLMs and standard practices in benchmarking thereof.

1 INTRODUCTION

Large Language Models (LLMs) have rapidly emerged as powerful tools for natural language processing, understanding and generation. Among their diverse applications, they are increasingly deployed in chat-based assistants for Question Answering (QA), automated agent-based systems, and serving as surrogates for traditional search engines (Jin et al., 2025; Xiong et al., 2024a; Kelly et al., 2023). Within this landscape, scientific QA constitutes a particularly critical and challenging task. It spans a wide range of use cases, from public science communication (Schäfer, 2023) and education across different levels of expertise (Welbl et al., 2017), to supporting research and knowledge discovery (Gu & Krenn, 2025). In these contexts, accuracy and trustworthiness are essential, as errors may misinform the public, impair learning, or distort scientific practice. A central obstacle to reliability is the phenomenon of hallucinations, where LLMs generate fluent and seemingly confident answers that are factually incorrect or misleading (Ji et al., 2022). While hallucinations are now widely recognized, effective methods to detect and mitigate them remain underdeveloped, particularly in high-stakes domains such as science. Evidence for this can be found in UQ methods being absent from popular products while being subject to research since several years.

Research into UQ aims to develop reliable, automated methods for quantifying how certain a model is in its own predictions (Guo et al., 2017). In the context of LLMs, UQ serves as a critical tool to mitigate hallucinations: by identifying outputs with high uncertainty, systems can flag poten-

054 tially erroneous responses, abstain from answering, or route queries to alternative mechanisms – i.e.
055 larger models, retrieval-augmented systems, or human reviewers. Beyond error mitigation, reliable
056 UQ enhances transparency and user trustworthiness by providing interpretable indicators of answer
057 reliability and soundness (Dhuliawala et al., 2023; Devic et al., 2025; Reyes et al., 2025).

059 2 KEY CONTRIBUTIONS

061 To address open questions regarding the applicability of UQ in scientific QA, this work provides
062 a comprehensive benchmark of uncertainty estimation methods with a qualitative and quantitative
063 focus on calibration. We identify and evaluate a set of core uncertainty estimation approaches that
064 are widely used and form the basis of many derivative methods, covering three principal approaches:
065 token-level confidences, verbalized approaches and semantic consistency.

066 We first select UQ methods based on applicability, computational feasibility, and relevance to current
067 practice. We identify core challenges for LLMs and subsequent UQ on responses in the scientific
068 domain. We evaluate the selected methods on a diverse collection of models and datasets, incor-
069 porating different difficulty levels we identified to make scientific QA especially challenging. For
070 verifiable ground truth we are using multiple choice QA and arithmetic QA.

071 In total, we analyze 685,000 responses across datasets and models, making this one of the largest and
072 most realistic evaluations of uncertainty quantification in long-form scientific QA to date. We report
073 detailed summary statistics, distributional analyses, and probability plots, addressing limitations of
074 prior evaluations that were predominantly conducted on short, saturated datasets and often only in
075 the context of newly proposed methods.

076 By evaluating how the selected UQ methods capture the reliability of reasoning-heavy, long-form
077 predictions, we characterize systematic model behaviour and fundamental limitations in reasoning-
078 demanding settings where reliable uncertainty estimates are particularly valuable. By identifying
079 systematic differences across settings, we uncover potential drivers of divergent UQ behaviour and
080 point to mechanisms that warrant further investigation. These insights motivate future work to un-
081 derstand and mitigate the factors shaping uncertainty estimates in complex reasoning tasks.

082 We empirically evaluate the selected methods and analyze the resulting data to address two guiding
083 research questions:

085 **(RQ1)** To what extent are the token probabilities a calibrated measure of confidence and what
086 effect does the instruction-tuning or reasoning process have on the calibration?

087 **(RQ2)** To what extent do sequence-level uncertainty-quantification methods provide reliable un-
088 certainty estimates for long-form answers in scientific question answering, across tasks
089 spanning retrieval-based to reasoning-intensive questions?

091 Finally, to facilitate ongoing research and rapid iteration as LLMs and UQ methods continue to
092 evolve, we introduce a flexible and extensible framework for LLM benchmarking. We release this
093 framework together with the implementation of benchmarks for calibration assessment presented in
094 this paper, detailed reproduction instructions, raw uncertainty scores from our benchmark runs, and
095 additional visualizations in an open-access repository accompanying this paper¹.

097 3 BACKGROUND

099 3.1 UNCERTAINTY QUANTIFICATION IN PREDICTIVE MODELS

101 Neural networks, including LLMs, face predictive uncertainty as their training data provide only
102 a discrete and incomplete mapping of real-world artifacts (Hüllermeier & Waegeman, 2021). The
103 paradigm of negative log likelihood training for next token prediction (Radford et al., 2018) in LLM
104 pre-training enforces uncertainty of generated tokens. Besides obtaining this predictive uncertainty
105 with dedicated methods (Gal et al., 2016), the act of validating the quantitative soundness of uncer-
106 tainties comes as a challenge to many practitioners (Guo et al., 2017; Chung et al., 2021).

107 ¹<https://anonymous.4open.science/r/llm-uncertainty-bench-9B2B/>

108 The field of UQ methods in LLMs is still nascent compared to methods for classification or regres-
109 sion tasks (Kendall & Gal, 2017; Kuleshov et al., 2018; Papamarkou et al., 2024) in general. In the
110 broader UQ literature, two core tasks are typically distinguished: calibration, which evaluates the
111 alignment between confidence scores and predictive accuracy, and selective prediction, which con-
112 cerns deciding when a model should abstain (Liu et al., 2025a). In this work we focus on calibration,
113 as it forms the statistical backdrop of almost all UQ methods and directly determines the reliability
114 of predicted confidences. Since calibration evaluates predicted probabilities for correctness, only
115 methods that produce scores in the $[0, 1]$ range are applicable. Arbitrary or unbounded scores are
116 not directly usable for calibration metrics.

117 The standard evaluation technique in calibration are calibration plots, which visualise how well
118 predicted confidence scores correlate to the true likelihood of correctness, and summary statis-
119 tics thereof (Guo et al., 2017; Hendrycks & Gimpel, 2018). Many quantitative approaches rely on
120 information-theoretic metrics, such as entropy or perplexity, but these often fail to capture language-
121 specific nuances. Recently, LLM-specific uncertainty estimation methods have emerged, such as
122 Verbalized Uncertainty, P(True), and Claim Conditioned Probability (CCP) (see details in Sec-
123 tion 3.2 and Section 7.1). However, they are studied in narrow domains on simple tasks, such as
124 factual QA on biographies or encyclopedic content (Fadeeva et al., 2024). Consequently, it remains
125 unclear how well these approaches generalize to scientific QA, which poses the unique challenges
126 outlined in Section 3.3. Compounding the issue, current benchmarks are often tightly coupled to
127 specific models, datasets, and UQ methods, limiting adaptability to rapidly evolving LLM architec-
128 tures and use cases.

129 3.2 MULTIPLE-CHOICE QUESTION ANSWERING AND UQ CALIBRATION

130
131 A straight forward way of estimating uncertainty in LLMs is to reformulate Multiple-Choice Ques-
132 tion Answering (MCQA) items as classification tasks, prompting the model to output a single label
133 token (e.g., A/B/C/D) representing the chosen answer. The confidence scores are derived from the
134 probabilities assigned to each label. These probabilities can be used as confidence scores directly,
135 i.e. ignoring probability mass on non-label tokens. Alternatively, a normalization with respect to
136 the total probability mass assigned to all possible answer labels can be applied. By normalizing
137 over the label set, the resulting confidence scores represent relative preferences among the labels –
138 not (un-)certainty in the options. This may obscure uncertainty that would otherwise be expressed
139 through low absolute probabilities, making the normalized scores less reliable as measures of uncer-
140 tainty (Wang et al., 2024a). The GPT-4 Technical Report (OpenAI, 2023) compared the calibration
141 of responses from a base model and instruction-tuned model using this approach. Their results sug-
142 gested good calibration in the base model, but significantly worse calibration in the instruction-tuned
143 model. This sparked a so far understudied controversy about the influence of instruction-tuning on
144 the calibration of models which is often encountered in colloquial discussions.

145 3.3 CHALLENGES IN SCIENTIFIC QUESTION ANSWERING

146 Scientific QA presents challenges that extend far beyond factual recall of general MCQA. Questions
147 often combine domain-specific terminology, symbolic expressions, and quantitative relationships,
148 requiring models to interpret structured scientific information and integrate it into coherent reason-
149 ing processes. Many tasks demand the recall of appropriate formulae, the identification of relevant
150 variables, and the manipulation of non-linear or multi-equation systems. These steps exceed sim-
151 ple pattern matching and due to the lack of intrinsic arithmetic ability, push LLMs toward brittle
152 heuristics. Thus, scientific problem solving frequently involves complex, dependency-rich reason-
153 ing chains. Errors in intermediate steps propagate, rendering downstream inferences unreliable.
154 This complicates UQ, as standard methods assume uncertainty is localized to individual predic-
155 tions. Effective UQ in scientific QA requires accounting for cumulative uncertainty across multiple
156 reasoning steps, each susceptible to distinct failure modes.

157 This dependency structure is precisely what makes scientific QA both difficult for current LLMs
158 and uniquely important for evaluating UQ methods. Whereas short-form benchmarks hide these
159 challenges, scientific QA exposes the complex, compounding nature of uncertainty in realistic,
160 reasoning-intensive settings, where reliable UQ stands to provide the greatest practical benefit.

161

4 RELATED WORK

UQ is particularly critical for LLMs because their token-level probabilistic generation can produce fluent yet confidently incorrect or misleading outputs known as hallucinations (Maynez et al., 2020). Hallucinations, classified as intrinsic contradictions or extrinsic fabrications, diminish trustworthiness and require robust uncertainty detection to uncover them (Sui et al., 2024; Zhang et al., 2023b; Banerjee et al., 2024). Many methods have been established to obtain predictive uncertainties for LLM generated text (Geng et al., 2024).

Token-level uncertainty estimation faces challenges due to varying semantic importance of tokens. Many methods nevertheless assume equal token importance, leading to misrepresentations (Ullrich et al., 2025; Kuhn et al., 2023a). Linguistic calibration by epistemic markers (e.g., “might,” “potentially”) provides interpretable uncertainty cues, but models tend to be overconfident in these expressions, risking user over-reliance (Band et al., 2024; Zhou et al., 2024). Empirical work shows that base LLMs are generally better calibrated than instruction-tuned models, which often become overconfident. (OpenAI, 2023; Tian et al., 2023; Wang et al., 2025). While researching risk scores on a binary classification task in a narrow domain, (Cruz et al., 2024) reported a polarization of binary label probability scores as a consequence of instruction tuning. Whereas overconfidence refers to a systematic tendency for confidence scores to exceed actual accuracy, polarization denotes a collapse of the token-level probability distribution. In polarization, the model allocates nearly all probability mass to a single label. Given the constrained setting in (Cruz et al., 2024), the generalizability of this finding on token-level calibration remains unclear.

Prompt design significantly affects uncertainty and calibration: small prompt variations, the use of epistemic phrasing, and simulating knowledge profiles influence both accuracy and confidence. Strategies have been proposed to mitigate overconfidence in self-evaluations (Cao et al., 2024; He et al., 2024; Sclar et al., 2024; Zhou et al., 2023; Lu & Wang, 2024; Xiong et al., 2024b). This body of work underscores the complexity of uncertainty estimation in LLMs and the need for context-aware, robust and scalable analyses of UQ methods.

First comprehensive studies of UQ effectiveness as well as their calibration were undertaken only recently by Fadeeva et al. (2023a) and Huang et al. (2025). In Fadeeva et al. (2023a), the authors share open-source UQ benchmark tooling (LM-Polygraph) publicly, underlining the importance of UQ analyses as LLM architectures progress rapidly. This study focuses on selective prediction (Geng et al., 2024). Both studies rate the effectiveness of UQ methods by virtue of one (Fadeeva et al., 2023a) and two (Huang et al., 2025) summary statistics respectively. This stands in contrast to complementary best practices (calibration plots) of closely related ML fields to evaluate predictive uncertainties (Guo et al., 2017) both quantitatively and qualitatively. Multiple reports (Zhang et al., 2024a; Vashurin et al., 2025) around the LM-Polygraph project offer a structured comparison of UQ methods considering computational cost, logit access, and granularity of uncertainty of a wide range of uncertainty methods. But these reports target the task of selective prediction exclusively, enabling the use of methods producing unbound scores, and focus on one single summary statistic (Prediction Rejection Ratio, PRR) to compare UQ methods. Here, we identified a gap in existing academic literature which understudied the calibration of UQ methods in favor of selective prediction.

We do observe early studies of calibration as a measure of quality control for UQ in LLMs: Tao et al. (2025) presented a vast analysis with respect to explored models, but restrict themselves to one single dataset only. Liu et al. (2025b) surveyed a variety of published UQ methods and presented available datasets, but did not perform experiments for an empirical comparison.

Our work aims to complement these findings by introducing a benchmark that systematically evaluates methods on the second core task of UQ: calibration. We assess methods representative of the most prominent approaches on scientific QA, stress-testing them on long-form responses that require complex syntactic and multi-step reasoning. By incorporating a diverse range of models and datasets, our benchmark provides reproducible, transparent qualitative and quantitative evidence, offering a solid foundation for future research on calibration in LLMs.

216 5 GENERAL EXPERIMENTAL SETUP

217 5.1 DATASET SELECTION

218 Due to the necessity of verifiable ground truth for computing calibration metrics such as ECE, AU-
 219 ROC and calibration plots, we target structured tasks. We use four multiple-choice QA datasets in
 220 combination with counterfactual prompting using the APriCoT prompting strategy to better approx-
 221 imate open-ended QA behaviour. To complement this, we are using three arithmetic QA datasets
 222 for better generalizability. Our evaluation includes datasets with a strong emphasis on scientific QA,
 223 containing content from physics, chemistry and biology. For low complexity arithmetic reasoning,
 224 math datasets are employed.

225 To study the effect of task complexity, the selected datasets span a range of difficulty levels and
 226 cognitive demands, from fact retrieval to syntactic and multi-step reasoning required in scientific
 227 QA. We treat scientific QA as a representative domain with high relevance for safety-critical and
 228 reasoning-intensive applications and with strong potential for generalisation beyond individual sub-
 229 fields. Through this selection, our benchmark is explicitly designed to evaluate uncertainty methods
 230 against the challenges identified in Section 3.1, ensuring that the tasks reflect the reasoning, domain
 231 knowledge and syntactic complexities that make scientific QA particularly demanding.

232 The following provides an overview over selected datasets (more details provided in Section A.3).

233 **MMLU** (Hendrycks et al., 2021) is a multiple-choice benchmark with 15,908 questions across 57
 234 academic subjects, including physics at varying levels. It is primarily testing comprehension, factual
 235 knowledge and single step reasoning.

236 **ARC** (Clark et al., 2018) tests scientific reasoning using grade-school science exam questions split
 237 into **ARC-Easy** and **ARC-Challenge** subsets. The latter emphasizes multi-step reasoning, making
 238 it a strong benchmark for evaluating UQ methods under complex conditions.

239 **SciQ** (Welbl et al., 2017) contains 13,679 multiple-choice questions in physics, chemistry, and
 240 biology. It emphasizes conceptual understanding.

241 **GPQA** (Rein et al., 2023) is a graduate-level science benchmark with 448 expert-written questions
 242 across physics, chemistry and biology, designed to resist simple lookup. Its high difficulty and
 243 reasoning demands make it particularly useful for stress-testing UQ methods.

244 **GSM8K** (Cobbe et al., 2021) is a standard math word problem dataset with 8,500 arithmetic ques-
 245 tions requiring step-by-step symbolic reasoning. **GSM-MC** (Zhang et al., 2024b) is a multiple-
 246 choice variant of GSM8K, using model-generated distractors to reduce evaluation ambiguity.

247 **SVAMP** (Patel et al., 2021) is an arithmetic reasoning dataset containing 1,000 questions that in-
 248 troduce distracting information in the problem text. While featuring low to moderate computational
 249 complexity, these distractors are specifically designed to induce ambiguity and decision uncertainty,
 250 challenging models can recognize and quantify uncertainty in the presence of misleading or irrele-
 251 vant information.

252 **SciBench** (Wang et al., 2024b) comprises 692 college-level questions from math, chemistry, and
 253 physics textbooks. It targets advanced symbolic reasoning involving formulas and physical units.

254 5.2 MODEL SELECTION

255 Models were chosen to cover a broad spectrum of LLMs while ensuring reproducibility through the
 256 use of open-weight, publicly available models. To capture diversity in design, the selection spans
 257 five major providers (OpenAI, Mistral, Meta, Qwen, and Google) and includes models of varying
 258 sizes (from 7B to 70B parameters), variants (base, instruction-tuned and reasoning models) and
 259 architectural designs, such as Mixture-of-Experts. To research the effect of instruction tuning on the
 260 calibration of label probabilities, instruction-tuned and reasoning models are complemented by their
 261 base model counterparts to enable controlled comparison. A detailed list of selected models can be
 262 found in Section A.2.

6 BENCHMARKING LABEL PROBABILITY CALIBRATION

6.1 METHODOLOGY / EXPERIMENTAL SETUP

While confidence scores represent a model’s self-assessed probability of correctness, effective UQ requires calibration methods that align these scores with empirical accuracy (Guo et al., 2017). Expected Calibration Error (ECE) is a widely used metric for this purpose, comparing binned confidence estimates with actual correctness.

To ensure comparability while extending prior work, the experimental setup was designed to remain close to the GPT-4 Technical Report (OpenAI, 2023), while expanding in model coverage, dataset diversity and calibration and score distribution analysis. We evaluated all selected base, instruction-tuned and reasoning models across four MCQA datasets, chosen to represent different reasoning complexities: factual and single-step reasoning (MMLU), symbolic reasoning (GSM8K), and multi-step reasoning (ARC-Reasoning, GPQA). Initial tests revealed substantial differences in task comprehension between base and instruction-tuned models. To address this, we designed four alternative prompts (see Section A.6.1), each using three-shot prompting. We then selected the prompt that maximized the average probability mass assigned to label tokens, thereby ensuring task comprehension across models. For the final results, we employed structured decoding during generation to reliably obtain relative label probabilities and to prevent invalid answers, as detailed in Section A.6.4.

Calibration performance was then assessed by comparing the model variants using label probabilities as confidence scores with calibration plots and summary metrics such as ECE.

6.2 KEY FINDINGS

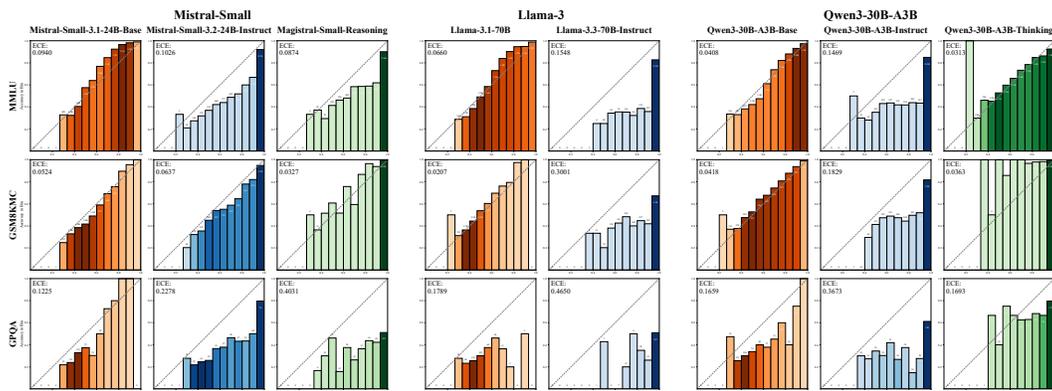


Figure 1: **Calibration Plots Using Label Probabilities as Uncertainty Scores for Only the Most Probable Label Per Prompt.** Columns correspond to three selected model families: base variants are shown in orange, instruction-tuned variants in blue, and reasoning variants in green. Rows refer to three QA datasets on MMLU, GSM8KMC and GPQA. Darker shading indicates a higher number of items within each confidence bin. Each bin is labelled with the sample count contained. The plot shows the results using Prompt Design 1 with structured decoding enabled.

Figure 1 reports a subset of representative calibration plots. All results shown use Prompt Design 1 (see Section A.6.1), which was found to optimize task comprehension across model types (see Section A.6.2). Structured decoding has been employed to generate the final result, after normalization has been validated (see Section A.7.2). Comprehensive results for all models, datasets and prompt designs are included in Section A.7.1.

We find that ECE increases with reasoning complexity (see Figure 1). Tasks requiring symbolic or multi-step reasoning, such as GSM8K and GPQA, exhibit substantially higher ECE compared to factual retrieval tasks, such as MMLU. This distinction highlights that token-level probabilities can reliably capture aleatoric uncertainty in fact retrieval, but become overconfident and unreliable when reasoning is required as a result of epistemic uncertainty.

We reproduce the degradation of ECE observed by (OpenAI, 2023) as a result of instruction-tuning, although not uniformly across models to the same degree. However, visual analysis of calibration plots reveals systematic polarization of token probability distributions (see Figure A.2). Instruction-tuned models tend to concentrate nearly all probability mass on a single label compared to their base model counterparts, thereby degrading the expressiveness of token-level confidence scores. This effect is evident in Figure 1, which shows a pronounced accumulation of items in the highest-confidence bucket. The degree of polarization is consistent across most instruction-tuned model families, with few notable exceptions among Mistral models.

This effect is also evident in the *gpt-oss* and *Magistral* reasoning models, where the reasoning process appears to commit to a single option. In contrast, the *DeepSeek-R1* and *Qwen3* reasoning model families seem to consider multiple options and actively retrieve relevant facts rather than committing prematurely. This allows the model’s uncertainty estimates to better reflect uncertainty arising from the different possible answers, resulting in improved calibration and reduced polarization. Because these differences cluster by model provider rather than by scale or architecture, we hypothesize that provider-specific training pipelines are a key underlying factor. Our benchmark thus provides a foundation for further research into how individual steps in training influence model calibration and polarization.

6.3 STRUCTURAL LIMITATIONS OF ECE FOR CALIBRATION ASSESSMENT

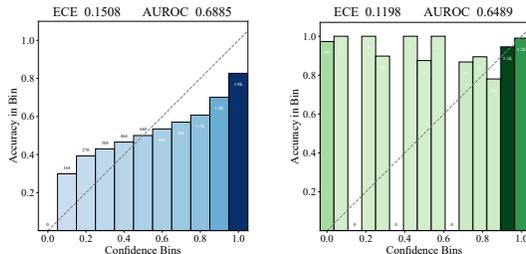


Figure 2: **ECE as a Misleading Proxy for Calibration: An Illustrative Case.** Two calibration plots from our results illustrate that similar or even lower ECE values do not necessarily indicate better calibration. In the left plot, uncertainty scores exhibit a clear relationship with correctness. In the right plot, there is no meaningful relationship between scores and correctness, yet the model still obtains a lower ECE and only a slightly lower AUROC.

The mixed levels of ECE degradation observed in the Label Probability Experiment, see Figure 1, highlights a fundamental limitation of ECE as a calibration metric. Although ECE measures the correlation between predicted confidence and the likelihood of correctness, it is not independent of the model’s overall accuracy. As a result, when confidence scores collapse into a narrow region of the distribution, ECE becomes driven primarily by overall accuracy rather than by the reliability of individual uncertainty estimates. Consequently, as model accuracy increases and consistently high confidences are reported, ECE may appear deceptively low even when the model remains poorly calibrated at the instance level, see Figure 2. This undermines the utility of ECE for UQ, where the primary goal is to assess whether a method can accurately quantify uncertainty on a per-instance basis, independent of its overall accuracy. This limitation is particularly relevant for LLMs, due to our findings of confidence polarization, where models produce consistently very high confidences. This issue has serious implications for the evaluation of novel UQ methods and benchmarks, as prior studies that rely solely on summary metrics like ECE may report misleadingly optimistic results. We advocate for using ECE only in combination with visual inspection of calibration plots and complementary summary metrics.

7 BENCHMARKING SEQUENCE-LEVEL CALIBRATION

7.1 SELECTION OF UNCERTAINTY MEASURES

Sequence level methods often use token-level confidences (examined in Section 6) via different aggregation strategies. Beyond these methods, we identify two more approaches of sequence-level

UQ methods that meet our criteria of theoretical grounding, computational feasibility, and empirical promise: verbalized approaches and semantic consistency. We further consider their practical relevance and the extent to which they underpin recently proposed UQ methods. From these families, we select four prominent and representative methods for evaluation. A justification of our selection criteria, as well as the rationale for excluding other approaches is provided in Section A.5.

Verbalized Uncertainty (Tian et al., 2023) prompts the model subsequent to the answer generation to provide a self-assessed probability estimate of correctness in token space. While straightforward, it ignores token probability distributions and may be susceptible to training data bias.

P(True) (Kadavath et al., 2022) prompts the model subsequent to the answer generation to classify the answer as “(A) True” or “(B) False”. The underlying token probabilities assigned to the corresponding labels “(A)” and “(B)” are then used as confidence scores.

Frequency of Answer estimates certainty by the proportion of semantically equivalent answers among multiple sampled generations for the same prompt. This is computationally expensive and semantic equivalence detection of answers is non-trivial in open ended questions answering. However, it captures semantic consistency and proxies other approaches like self-consistency prompting (Wang et al., 2023).

Claim-Conditioned Probability (CCP) (Fadeeva et al., 2024) evaluates uncertainty at the token level by determining the semantic consistency among the top probable token alternatives. This is done by clustering the token alternatives in tokens that entail and contradict the original meaning by comparing the chosen token with its alternatives using an NLI model. The token-level confidence score is the ratio of probability mass assigned to entailing tokens to the sum of entailing and contradicting tokens. Sequence-level confidence is calculated from product of token confidences.

7.2 METHODOLOGY / EXPERIMENTAL SETUP

To evaluate the calibration of different uncertainty methods in long-form QA, we focus exclusively on instruction-tuned and reasoning models, as base models exhibit poor task comprehension and are unsuitable for QA tasks. For this experiment, all seven previously selected MCQA and arithmetic QA datasets were used. MCQA introduces specific challenges: (1) **Selection bias**. Models may favor certain answer choices due to token frequency or formatting learned during training, regardless of semantic content (Myrzakhan et al., 2024). (2) **Positional bias**. Models can exhibit systematic preference for certain label positions (e.g., always selecting “A”) (Zheng et al., 2024).

We address these biases using the APriCoT prompting strategy (Moore et al., 2025), which combines Chain-of-Thought reasoning with counterfactual prompting. Each answer choice is evaluated independently, and the model classifies it as correct or incorrect, producing a verifiable judgment. This isolates answer evaluation from ordering and formatting effects, reducing both selection and positional bias and improving calibration. Rephrasing MC questions as open-ended queries could avoid format-related biases, but many items depend on predefined choices (e.g., fill-in-the-blank or elimination logic). APriCoT also helps approximate open-ended QA by eliciting reasoning over individual answers while preserving a verifiable format.

For arithmetic QA datasets, Chain-of-Thought (CoT) prompting is used to facilitate multi-step reasoning. To balance coverage and computational cost, each dataset is subsampled to 250 items, and 10 generations are sampled per prompt, resulting in a total of 685,000 long form responses across models and datasets evaluated against each of the 4 selected methods (see Section A.4).

7.3 KEY FINDINGS

Verbalized Uncertainty Our results show that models overwhelmingly defaulted to a small set of uncertainty scores (see Figure 4a), with higher confidences dominating the responses. This leads to a biased distribution of confidence scores, potentially driven by training data and instruction tuning. Calibration plots reveal no meaningful correlation between these verbalized scores and answer accuracy in all instruction tuned models, indicating that Verbalized Uncertainty is not a reliable proxy for true model confidence. This is slightly improved in reasoning models, most notably in the *gpt-oss* model family, providing scores that better correlate with answer accuracy on challenging datasets such as GPQA and SciBench (see Figure A.10). As no architectural or scale-related factors besides

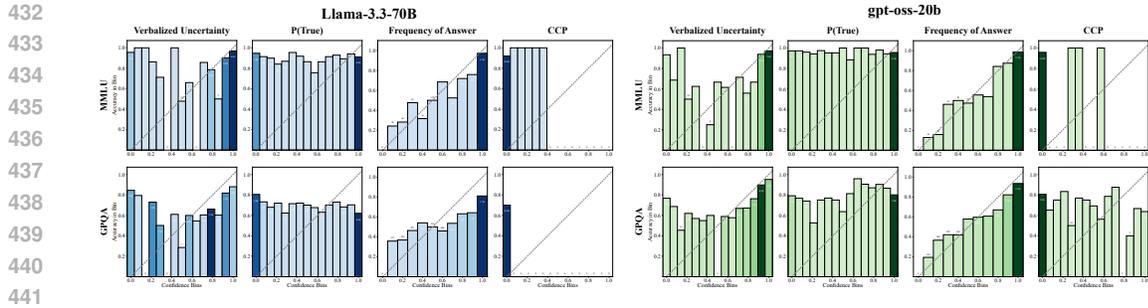
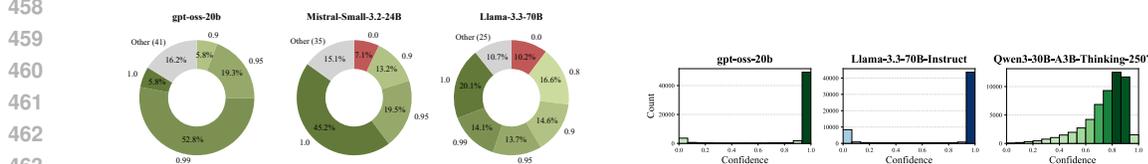


Figure 3: **Selected Calibration Plots For The Four Selected Methods.** Results for *Llama-3.3-70B* (left) and *gpt-oss-20b* (right) are shown, each for MMLU and GPQA for the four computed sequence level uncertainty methods. The full plots for all models and datasets per UQ method can be found in Section A.8.3.

the reasoning process clearly distinguish these models from others, the reason for the reliability is likely driven by training pipelines, which we cannot investigate due to them being undisclosed.

P(True) We find that P(True) suffers from pronounced response bias. In several models P(True) overwhelmingly assigns near 1.0 confidence with little use of intermediate confidence scores, resulting in a polarized distribution of certainty scores (see Figure 4b). This polarization may stem from the model’s commitment to a single reasoning path before classification, and varies by model. Calibration plots further reveal no meaningful correlation between P(True) scores and actual correctness, indicating that P(True) cannot reliably quantify uncertainty in this setting.



(a) **Verbalized Uncertainty** Confidence scores seen in less than 5% of the total responses have been grouped into “Other”, with the number of distinct confidence scores shown in brackets.

(b) **P(True)** Confidence scores are derived by probability mass assigned to labels (A) and (B), representing the model’s confidence that the answer is true or false respectively.

Figure 4: **Distribution of Confidence Scores Across Selected Representative Models.** Confidence scores have been aggregated across all datasets (57, 500 prompts in total). See Section A.8.4 and Section A.8.5 for extensive plots for all models.

Frequency of Answer Our evaluation (Figure 3) shows that higher answer frequencies strongly correlate with correctness across both multiple-choice and arithmetic tasks. More challenging datasets (e.g., GPQA, SciBench) exhibiting greater answer diversity, which reflects higher model uncertainty. Calibration plots confirm well-aligned confidence estimates based on Frequency of Answer, demonstrating the reliability of this approach. However, due to its computational cost, scaling with the number of samples per prompt, and its dependence on semantic clustering of outputs, it remains challenging or unapplicable to open-ended QA.

Claim Conditioned Probability In practice, we find that CCP suffers from vanishing sequence-level scores as generation length grows, i.e. multiplying a large number of token confidences drives overall confidence near 0. As a result, calibration plots (Figure 3) show no meaningful alignment with correctness. Furthermore, high impact of single NLI misclassifications and the inclusion of stop words in the aggregation further destabilize scores. These issues make CCP unreliable for sequence-level uncertainty estimation. But, its token-level insights could still inform targeted analyses once aggregation and domain-specific entailment are improved upon in future work.

8 LIMITATIONS

This study focuses on scientific QA in structured formats, which constrains the generalizability of the findings to other domains and task types such as open-ended generation or summarization. All datasets are in English, and the potential impact of cross-linguistic variation, input formatting, and prompt phrasing on calibration was not examined. Despite mitigation efforts using prompting strategies like APriCoT, the multiple-choice setting itself may introduce systematic biases, such as steering effects. Furthermore, the benchmark employs controlled inference conditions (e.g. fixed temperature, fixed decoding parameters) that may not capture the variability of real-world deployments. Finally, the evaluation of sequence-level uncertainty was limited to a subset of UQ methods with normalized outputs, selected to enable calibration analysis via calibration plots and summary metrics. This leaves unnormalized or claim-level metrics for future investigation into their potential value for detecting incorrect answers through high uncertainty estimates.

9 CONCLUSION

We presented a systematic evaluation of four UQ methods for LLMs across seven natural science datasets and up to 20 open-weight models, including base, instruction-tuned, and reasoning variants. We acquired our results by developing an open-source framework for benchmarking LLMs with a focus on efficiency and reproducibility. Expanding on existing literature and software frameworks, our contribution targets calibration instead of selective prediction to assess UQ method performance.

Our results reveal a pronounced polarization effect in token-level confidence distributions induced by instruction-tuning, which diminishes their utility as reliable uncertainty signals. Reasoning models are exposed to the same effect, but the reasoning process can mitigate this depending on its structure, as differences between providers are showing.

At the sequence level, we present results based on the evaluation of 685,000 long form responses across models and datasets on 4 selected UQ methods representative for prominent approaches. We report that verbalized UQ methods exhibit consistently poor performance across most models. In particular, **P(True)** is adversely affected by the same confidence polarization previously observed at the token level. **Verbalized Uncertainty** scores are biased toward a narrow range of high-confidence scores and generally fail to correlate with answer correctness. Reasoning can improve on this, as seen in the *gpt-oss* family, underscoring the need for continued benchmarking to disentangle the impact of architectural design choices, fine-tuning data, and training paradigms on UQ behavior. The **Frequency of Answer** method based in semantic consistency showed strong reliability, albeit at high computational cost and with the non-trivial challenge of robust semantic equivalence detection. **Claim-Conditioned Probability (CCP)** suffers from vanishing sequence-level scores as generation length increases, compounded by NLI misclassifications and instability from stop-word aggregation.

We advocate that our results have substantial relevance outside of scientific QA tasks given the variety and number of datasets and models in use. As such, our findings highlight a critical need for the development of more robust, efficient, and theoretically grounded UQ methods for LLMs. Our results demonstrate that algorithmic advancements must always go alongside sustained empirical benchmarking to isolate the contributions of model architecture, fine-tuning strategy, and training data to uncertainty behavior in real world scenarios.

10 FUTURE WORK

Future research on uncertainty estimation in LLMs should explore several directions. Short-term efforts include extending benchmarks with new UQ methods, systematically studying inference parameters (e.g., temperature, sampling) and prompting strategies, and tracking UQ performance across emerging models to relate gains to architecture or training paradigms. Advancing UQ evaluation will also require new datasets that explicitly induce uncertainty or reformulate multiple-choice tasks into open-ended formats to isolate task-specific sources of uncertainty. Longer-term directions include claim-level uncertainty estimation, assessing reliability of individual statements or reasoning steps, and linguistic calibration, studying alignment between epistemic markers and model confidence.

540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593

REPRODUCIBILITY STATEMENT

Reproducibility in UQ for LLMs poses unique challenges due to the multistage evaluation pipelines and inherent randomness at each step, which can compound and introduce measurement noise. This stochasticity affects comparisons across UQ methods and model outputs, particularly when methods are evaluated on different generations, potentially contaminating results. To address these issues, we developed a modular benchmarking framework designed to ensure reproducibility and resource efficiency in large-scale LLM experiments. The framework supports extensible and replaceable computation nodes, caches probabilistic outputs such as model generations to ensure consistent re-evaluation, preserves intermediate outputs for qualitative inspection and allows incremental updates so that only affected steps need recomputation. It also enables sharing of intermediate results to reduce compute costs and allow other researchers to rerun the benchmark with their own methods. This is particularly valuable for ongoing research in UQ on models gated by proprietary access or costly hardware, supporting broader collaboration within the scientific community. These features not only enhance the reliability of the present study but also establish a foundation for future research in the field.

To facilitate reproducibility, the repository² includes the full framework `async-graph-bench` in its current form, final leaf-node outputs containing confidence scores required to reproduce the plots, a container definition file to build the execution environment, exact pip and driver versions, and instructions to install the framework and run the benchmarks presented in this work. While the framework is in its early stages, it is fully functional for reproducing the experiments reported here, and we plan to continue improving documentation and testing before its broader release in the near future.

LLM DISCLOSURE

LLMs were used in a limited manner to assist with writing, coding, and literature search. Specifically, LLMs provided minor support for language polishing, code-related assistance and identifying relevant references. All outputs were carefully reviewed, tested, and verified by the authors. LLMs did not contribute to the conceptualization of the research, the design of experiments, or the interpretation of results.

REFERENCES

- Neil Band, Xuechen Li, Tengyu Ma, and Tatsunori Hashimoto. Linguistic calibration of long-form generations, 2024. URL <https://arxiv.org/abs/2404.00474>.
- Sourav Banerjee, Ayushi Agarwal, and Saloni Singla. Llms will always hallucinate, and we need to live with this, 2024. URL <https://arxiv.org/abs/2409.05746>.
- Bowen Cao, Deng Cai, Zhisong Zhang, Yuexian Zou, and Wai Lam. On the worst prompt performance of large language models, 2024. URL <https://arxiv.org/abs/2406.10248>.
- 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. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024. URL <https://openreview.net/forum?id=Zj12nz1Qbz>.
- Youngseog Chung, Ian Char, Han Guo, Jeff Schneider, and Willie Neiswanger. Uncertainty toolbox: an open-source library for assessing, visualizing, and improving uncertainty quantification. *arXiv preprint arXiv:2109.10254*, 2021.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge, 2018. URL <https://arxiv.org/abs/1803.05457>.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John

²<https://anonymous.4open.science/r/llm-uncertainty-bench-9B2B/>

- 594 Schulman. Training verifiers to solve math word problems, 2021. URL <https://arxiv.org/abs/2110.14168>.
- 595
596
- 597 André F. Cruz, Moritz Hardt, and Celestine Mender-Dünner. Evaluating language models as risk
598 scores, 2024. URL <https://arxiv.org/abs/2407.14614>.
- 599
600 Maxime Darrin, Pablo Piantanida, and Pierre Colombo. RainProof: An umbrella to shield
601 text generator from out-of-distribution data. In Houda Bouamor, Juan Pino, and Kalika Bali
602 (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Pro-*
603 *cessing*, pp. 5831–5857, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.357. URL [https://aclanthology.org/2023.](https://aclanthology.org/2023.emnlp-main.357/)
604 [emnlp-main.357/](https://aclanthology.org/2023.emnlp-main.357/).
- 605
606 Siddhartha Devic, Tejas Srinivasan, Jesse Thomason, Willie Neiswanger, and Vatsal Sharan. From
607 calibration to collaboration: Llm uncertainty quantification should be more human-centered,
608 2025. URL <https://arxiv.org/abs/2506.07461>.
- 609
610 Shehzaad Dhuliawala, Vilém Zouhar, Mennatallah El-Assady, and Mrinmaya Sachan. A diachronic
611 perspective on user trust in ai under uncertainty, 2023. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2310.13544)
612 [2310.13544](https://arxiv.org/abs/2310.13544).
- 613
614 Jinhao Duan, Hao Cheng, Shiqi Wang, Alex Zavalny, Chenan Wang, Renjing Xu, Bhavya Kailkhura,
615 and Kaidi Xu. Shifting attention to relevance: Towards the predictive uncertainty quantification
616 of free-form large language models, 2024. URL <https://arxiv.org/abs/2307.01379>.
- 617
618 Ekaterina Fadeeva, Roman Vashurin, Akim Tsvigun, Artem Vazhentsev, Sergey Petrakov, Kirill
619 Fedyanin, Daniil Vasilev, Elizaveta Goncharova, Alexander Panchenko, Maxim Panov, Timothy
620 Baldwin, and Artem Shelmanov. Lm-polygraph: Uncertainty estimation for language models,
621 2023a. URL <https://arxiv.org/abs/2311.07383>.
- 622
623 Ekaterina Fadeeva, Roman Vashurin, Akim Tsvigun, Artem Vazhentsev, Sergey Petrakov, Kirill
624 Fedyanin, Daniil Vasilev, Elizaveta Goncharova, Alexander Panchenko, Maxim Panov, Timothy
625 Baldwin, and Artem Shelmanov. LM-polygraph: Uncertainty estimation for language models. In
626 Yansong Feng and Els Lefever (eds.), *Proceedings of the 2023 Conference on Empirical Methods*
627 *in Natural Language Processing: System Demonstrations*, pp. 446–461, Singapore, December
628 2023b. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-demo.41. URL
629 <https://aclanthology.org/2023.emnlp-demo.41>.
- 630
631 Ekaterina Fadeeva, Aleksandr Rubashevskii, Artem Shelmanov, Sergey Petrakov, Haonan Li,
632 Hamdy Mubarak, Evgenii Tsymbalov, Gleb Kuzmin, Alexander Panchenko, Timothy Baldwin,
633 Preslav Nakov, and Maxim Panov. Fact-checking the output of large language models via token-
634 level uncertainty quantification, 2024. URL <https://arxiv.org/abs/2403.04696>.
- 635
636 Marina Fomicheva, Shuo Sun, Lisa Yankovskaya, Frédéric Blain, Rico Sennrich, and Lu-
637 cia Specia. Unsupervised quality estimation for neural machine translation. *Transactions*
638 *of the Association for Computational Linguistics*, 8:539–555, 2020. doi: 10.1162/tacl.a_00330. URL https://direct.mit.edu/tacl/article/doi/10.1162/tacl.a_00330/96475/Unsupervised-Quality-Estimation-for-Neural-Machine.
- 639
640 Yarin Gal et al. Uncertainty in deep learning, 2016.
- 641
642 Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koeppel, Preslav Nakov, and Iryna Gurevych. A survey
643 of confidence estimation and calibration in large language models. In Kevin Duh, Helena Gomez,
644 and Steven Bethard (eds.), *Proceedings of the 2024 Conference of the North American Chapter of*
645 *the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long*
646 *Papers)*, pp. 6577–6595. Association for Computational Linguistics, 2024. doi: 10.18653/v1/2024.naacl-long.366. URL <https://aclanthology.org/2024.naacl-long.366/>.
- 647
648 Xuemei Gu and Mario Krenn. Forecasting high-impact research topics via machine learning
649 on evolving knowledge graphs. *Machine Learning: Science and Technology*, 6(2):025041,
650 may 2025. doi: 10.1088/2632-2153/add6ef. URL [https://dx.doi.org/10.1088/](https://dx.doi.org/10.1088/2632-2153/add6ef)
651 [2632-2153/add6ef](https://dx.doi.org/10.1088/2632-2153/add6ef).

- 648 Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural
649 networks. In *International conference on machine learning*, pp. 1321–1330. PMLR, 2017.
- 650
- 651 Jia He, Mukund Rungta, David Koleczek, Arshdeep Sekhon, Franklin X Wang, and Sadid Hasan.
652 Does prompt formatting have any impact on llm performance?, 2024. URL <https://arxiv.org/abs/2411.10541>.
- 653
- 654 Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution
655 examples in neural networks, 2018. URL <https://arxiv.org/abs/1610.02136>.
- 656
- 657 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Ja-
658 cob Steinhardt. Measuring massive multitask language understanding, 2021. URL <https://arxiv.org/abs/2009.03300>.
- 659
- 660 Yuheng Huang, Jiayang Song, Zhijie Wang, Shengming Zhao, Huaming Chen, Felix Juefei-Xu,
661 and Lei Ma. Look before you leap: An exploratory study of uncertainty measurement for large
662 language models. *IEEE Transactions on Software Engineering*, 51(2):413–429, 2025. ISSN
663 0098-5589, 1939-3520, 2326-3881. doi: 10.1109/TSE.2024.3519464. URL <http://arxiv.org/abs/2307.10236>.
- 664
- 665 Eyke Hüllermeier and Willem Waegeman. Aleatoric and epistemic uncertainty in machine learning:
666 An introduction to concepts and methods. *Machine learning*, 110(3):457–506, 2021.
- 667
- 668 Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Yejin Bang, Andrea
669 Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *CoRR*,
670 abs/2202.03629, 2022. URL <https://arxiv.org/abs/2202.03629>.
- 671
- 672 Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Serkan Arik, Dong Wang, Hamed Zamani, and
673 Jiawei Han. Search-rl: Training llms to reason and leverage search engines with reinforcement
674 learning, 2025. URL <https://arxiv.org/abs/2503.09516>.
- 675
- 676 Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez,
677 Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer
678 El-Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bow-
679 man, Stanislav Fort, Deep Ganguli, Danny Hernandez, Josh Jacobson, Jackson Kernion, Shauna
680 Kravec, Liane Lovitt, Kamal Ndousse, Catherine Olsson, Sam Ringer, Dario Amodei, Tom
681 Brown, Jack Clark, Nicholas Joseph, Ben Mann, Sam McCandlish, Chris Olah, and Jared Ka-
682 plan. Language models (mostly) know what they know, 2022. URL <https://arxiv.org/abs/2207.05221>.
- 683
- 684 Dominique Kelly, Yimin Chen, Sarah E Cornwell, Nicole S Delellis, Alex Mayhew, Sodiq Onao-
685 lapo, and Victoria L Rubin. Bing chat: The future of search engines? *Proceedings of the Association for Information Science and Technology*, 60(1):1007–1009, 2023.
- 686
- 687 Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer
688 vision?, 2017. URL <https://arxiv.org/abs/1703.04977>.
- 689
- 690 Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. Semantic uncertainty: Linguistic invariances for
691 uncertainty estimation in natural language generation, 2023a. URL <https://arxiv.org/abs/2302.09664>.
- 692
- 693 Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. Semantic uncertainty: Linguistic invariances
694 for uncertainty estimation in natural language generation. In *Proceedings of the International
695 Conference on Learning Representations (ICLR)*, 2023b. URL <https://openreview.net/forum?id=VD-AYtP0dve>.
- 696
- 697 Volodymyr Kuleshov, Nathan Fenner, and Stefano Ermon. Accurate uncertainties for deep learning
698 using calibrated regression, 2018. URL <https://arxiv.org/abs/1807.00263>.
- 699
- 700 Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detect-
701 ing out-of-distribution samples and adversarial attacks. In *Advances in Neural Information Pro-
cessing Systems*, volume 31, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/abdeb6f575ac5c6676b747bca8d09cc2-Abstract.html>.

- 702 Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. Generating with confidence: Uncertainty quan-
703 tification for black-box large language models. *arXiv preprint arXiv:2305.19187*, 2023. URL
704 <https://arxiv.org/abs/2305.19187>.
705
- 706 Xiaoou Liu, Tiejun Chen, Longchao Da, Chacha Chen, Zhen Lin, and Hua Wei. Uncertainty
707 quantification and confidence calibration in large language models: A survey, 2025a. URL
708 <https://arxiv.org/abs/2503.15850>.
- 709 Xiaoou Liu, Tiejun Chen, Longchao Da, Chacha Chen, Zhen Lin, and Hua Wei. Uncertainty
710 quantification and confidence calibration in large language models: A survey, 2025b. URL
711 <https://arxiv.org/abs/2503.15850>.
- 712 Xinyi Lu and Xu Wang. Generative students: Using llm-simulated student profiles to support ques-
713 tion item evaluation. In *Proceedings of the Eleventh ACM Conference on Learning @ Scale,*
714 *L@S '24*, pp. 16–27, New York, NY, USA, 2024. Association for Computing Machinery. ISBN
715 9798400706332. doi: 10.1145/3657604.3662031. URL [https://doi.org/10.1145/](https://doi.org/10.1145/3657604.3662031)
716 [3657604.3662031](https://doi.org/10.1145/3657604.3662031).
- 717 Andrey Malinin and Mark Gales. Uncertainty estimation in autoregressive structured prediction.
718 *arXiv preprint arXiv:2002.07650*, 2020. URL <https://arxiv.org/abs/2002.07650>.
719
- 720 Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. On faithfulness and factuality
721 in abstractive summarization, 2020. URL <https://arxiv.org/abs/2005.00661>.
- 722 Kyle Moore, Jesse Roberts, Thao Pham, and Douglas Fisher. Chain of thought still thinks fast:
723 Apricot helps with thinking slow, 2025. URL <https://arxiv.org/abs/2408.08651>.
724
- 725 Aidar Myrzakhan, Sondos Mahmoud Bsharat, and Zhiqiang Shen. Open-llm-leaderboard: From
726 multi-choice to open-style questions for llms evaluation, benchmark, and arena, 2024. URL
727 <https://arxiv.org/abs/2406.07545>.
- 728 Alexander Nikitin, Jannik Kossen, Yarin Gal, and Pekka Marttinen. Kernel language entropy:
729 Fine-grained uncertainty quantification for llms from semantic similarities. *arXiv preprint*
730 *arXiv:2405.20003*, 2024. URL <https://arxiv.org/pdf/2405.20003>.
- 731 OpenAI. Gpt-4 technical report, 2023. URL <https://arxiv.org/abs/2303.08774>.
732 [arXiv:2303.08774 \[cs.CL\]](https://arxiv.org/abs/2303.08774).
733
- 734 Theodore Papamarkou, Maria Skoularidou, Konstantina Palla, Laurence Aitchison, Julyan Arbel,
735 David Dunson, Maurizio Filippone, Vincent Fortuin, Philipp Hennig, José Miguel Hernández-
736 Lobato, et al. Position: Bayesian deep learning is needed in the age of large-scale ai. *arXiv*
737 *preprint arXiv:2402.00809*, 2024.
- 738 Arkil Patel, Satwik Bhattamishra, and Navin Goyal. Are nlp models really able to solve simple math
739 word problems?, 2021. URL <https://arxiv.org/abs/2103.07191>.
- 740 Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language
741 understanding by generative pre-training, 2018. URL [https://cdn.openai.com/](https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf)
742 [research-covers/language-unsupervised/language_understanding_](https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf)
743 [paper.pdf](https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf).
744
- 745 David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien
746 Dirani, Julian Michael, and Samuel R. Bowman. Gpqa: A graduate-level google-proof q&a
747 benchmark, 2023. URL <https://arxiv.org/abs/2311.12022>.
- 748 Jie Ren, Jiaming Luo, Yao Zhao, Kundan Krishna, Mohammad Saleh, Balaji Lakshminarayanan,
749 and Peter J. Liu. Out-of-distribution detection and selective generation for conditional language
750 models. In *International Conference on Learning Representations (ICLR)*, 2023. URL [https://](https://openreview.net/forum?id=kJUS5nD0vPB)
751 openreview.net/forum?id=kJUS5nD0vPB.
- 752 Jonatan Reyes, Anil Ufuk Batmaz, and Marta Kersten-Oertel. Trusting AI: does un-
753 certainty visualization affect decision-making? *Frontiers in Computer Science*,
754 7, February 2025. ISSN 2624-9898. doi: 10.3389/fcomp.2025.1464348. URL
755 [https://www.frontiersin.org/journals/computer-science/articles/](https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2025.1464348/full)
[10.3389/fcomp.2025.1464348/full](https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2025.1464348/full). Publisher: Frontiers.

- 756 Mike S Schäfer. The notorious gpt: science communication in the age of artificial intelligence.
757 *JCOM: Journal of Science Communication*, 22(2):Y02, 2023.
758
- 759 Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. Quantifying language models’ sen-
760 sitivity to spurious features in prompt design or: How i learned to start worrying about prompt
761 formatting, 2024. URL <https://arxiv.org/abs/2310.11324>.
- 762 Peiqi Sui, Eamon Duede, Sophie Wu, and Richard Jean So. Confabulation: The surprising value
763 of large language model hallucinations, 2024. URL [https://arxiv.org/abs/2406.](https://arxiv.org/abs/2406.04175)
764 04175.
- 765 Junya Takayama and Yuki Arase. Relevant and informative response generation using pointwise
766 mutual information. In Yun-Nung Chen, Tania Bedrax-Weiss, Dilek Hakkani-Tur, Anuj Ku-
767 mar, Mike Lewis, Thang-Minh Luong, Pei-Hao Su, and Tsung-Hsien Wen (eds.), *Proceed-*
768 *ings of the First Workshop on NLP for Conversational AI*, pp. 133–138, Florence, Italy, Au-
769 gust 2019. Association for Computational Linguistics. doi: 10.18653/v1/W19-4115. URL
770 <https://aclanthology.org/W19-4115/>.
771
- 772 Linwei Tao, Yi-Fan Yeh, Minjing Dong, Tao Huang, Philip Torr, and Chang Xu. Revisiting un-
773 certainty estimation and calibration of large language models, 2025. URL [https://arxiv.](https://arxiv.org/abs/2505.23854)
774 [org/abs/2505.23854](https://arxiv.org/abs/2505.23854).
- 775 Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea
776 Finn, and Christopher D. Manning. Just ask for calibration: Strategies for eliciting calibrated
777 confidence scores from language models fine-tuned with human feedback, 2023. URL [https://](https://arxiv.org/abs/2305.14975)
778 arxiv.org/abs/2305.14975.
- 779 Herbert Ullrich, Tomáš Mlynář, and Jan Drchal. Claim extraction for fact-checking: Data, models,
780 and automated metrics, 2025. URL <https://arxiv.org/abs/2502.04955>.
781
- 782 Liam van der Poel, Ryan Cotterell, and Clara Meister. Mutual information alleviates hallucina-
783 tions in abstractive summarization. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.),
784 *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp.
785 5956–5965, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational
786 Linguistics. doi: 10.18653/v1/2022.emnlp-main.399. URL [https://aclanthology.org/](https://aclanthology.org/2022.emnlp-main.399/)
787 [2022.emnlp-main.399/](https://aclanthology.org/2022.emnlp-main.399/).
- 788 Roman Vashurin, Ekaterina Fadeeva, Artem Vazhentsev, Lyudmila Rvanova, Akim Tsvigun, Daniil
789 Vasilev, Rui Xing, Abdelrahman Boda Sadallah, Kirill Grishchenkov, Sergey Petrakov, Alexan-
790 der Panchenko, Timothy Baldwin, Preslav Nakov, Maxim Panov, and Artem Shelmanov. Bench-
791 marking uncertainty quantification methods for large language models with LM-polygraph. 13:
792 220–248, 2025. ISSN 2307-387X. doi: 10.1162/tacl.a.00737. URL [http://arxiv.org/](http://arxiv.org/abs/2406.15627)
793 [abs/2406.15627](http://arxiv.org/abs/2406.15627).
- 794 Artem Vazhentsev, Gleb Kuzmin, Akim Tsvigun, Alexander Panchenko, Maxim Panov, Mikhail
795 Burtsev, and Artem Shelmanov. Hybrid uncertainty quantification for selective text classifica-
796 tion in ambiguous tasks. In *Proceedings of the 61st Annual Meeting of the Association for*
797 *Computational Linguistics (Volume 1: Long Papers)*, pp. 11659–11681, Toronto, Canada, July
798 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.652. URL
799 <https://aclanthology.org/2023.acl-long.652/>.
- 800 Haochun Wang, Sendong Zhao, Zewen Qiang, Nuwa Xi, Bing Qin, and Ting Liu. Llms may per-
801 form mcqa by selecting the least incorrect option, 2024a. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2402.01349)
802 [2402.01349](https://arxiv.org/abs/2402.01349).
- 803 Xiaoxuan Wang, Ziniu Hu, Pan Lu, Yanqiao Zhu, Jieyu Zhang, Satyen Subramaniam, Arjun R.
804 Loomba, Shichang Zhang, Yizhou Sun, and Wei Wang. Scibench: Evaluating college-level scien-
805 tific problem-solving abilities of large language models, 2024b. URL [https://arxiv.org/](https://arxiv.org/abs/2307.10635)
806 [abs/2307.10635](https://arxiv.org/abs/2307.10635).
807
- 808 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-
809 ury, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models,
2023. URL <https://arxiv.org/abs/2203.11171>.

- 810 Ziming Wang, Zeyu Shi, Haoyi Zhou, Shiqi Gao, Qingyun Sun, and Jianxin Li. Towards objective
811 fine-tuning: How llms’ prior knowledge causes potential poor calibration?, 2025. URL <https://arxiv.org/abs/2505.20903>.
812
813
- 814 Johannes Welbl, Nelson F. Liu, and Matt Gardner. Crowdsourcing multiple choice science questions,
815 2017. URL <https://arxiv.org/abs/1707.06209>.
- 816 Haoyi Xiong, Jiang Bian, Yuchen Li, Xuhong Li, Mengnan Du, Shuaiqiang Wang, Dawei Yin, and
817 Sumi Helal. When search engine services meet large language models: visions and challenges.
818 *IEEE Transactions on Services Computing*, 2024a.
- 819 Miao Xiong, Zhiyuan Hu, Xinyang Lu, Yifei Li, Jie Fu, Junxian He, and Bryan Hooi. Can llms
820 express their uncertainty? an empirical evaluation of confidence elicitation in llms, 2024b. URL
821 <https://arxiv.org/abs/2306.13063>.
822
- 823 KiYoon Yoo, Jangho Kim, Jiho Jang, and Nojun Kwak. Detection of adversarial examples in text
824 classification: Benchmark and baseline via robust density estimation. In Smaranda Muresan,
825 Preslav Nakov, and Aline Villavicencio (eds.), *Findings of the Association for Computational Lin-*
826 *guistics: ACL 2022*, pp. 3656–3672, Dublin, Ireland, May 2022. Association for Computational
827 Linguistics. doi: 10.18653/v1/2022.findings-acl.289. URL [https://aclanthology.org/](https://aclanthology.org/2022.findings-acl.289/)
828 [2022.findings-acl.289/](https://aclanthology.org/2022.findings-acl.289/).
- 829 Caiqi Zhang, Fangyu Liu, Marco Basaldella, and Nigel Collier. Luq: Long-text uncertainty quan-
830 tification for llms. In *Proceedings of the 2024 Conference on Empirical Methods in Natu-*
831 *ral Language Processing*, pp. 5244–5262, Miami, Florida, USA, November 2024a. Associa-
832 tion for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.299. URL <https://aclanthology.org/2024.emnlp-main.299/>.
833
- 834 Tianhang Zhang, Lin Qiu, Qipeng Guo, Cheng Deng, Yue Zhang, Zheng Zhang, Chenghu Zhou,
835 Xinbing Wang, and Luoyi Fu. Enhancing uncertainty-based hallucination detection with stronger
836 focus. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language*
837 *Processing*, pp. 915–932, Singapore, December 2023a. Association for Computational Linguis-
838 tics. doi: 10.18653/v1/2023.emnlp-main.58. URL [https://aclanthology.org/2023.](https://aclanthology.org/2023.emnlp-main.58/)
839 [emnlp-main.58/](https://aclanthology.org/2023.emnlp-main.58/).
- 840 Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemaoy Liu, Tingchen Fu, Xinting Huang, Enbo Zhao,
841 Yu Zhang, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi.
842 Siren’s song in the ai ocean: A survey on hallucination in large language models, 2023b. URL
843 <https://arxiv.org/abs/2309.01219>.
844
- 845 Ziyin Zhang, Zhaokun Jiang, Lizhen Xu, Hongkun Hao, and Rui Wang. Multiple-choice ques-
846 tions are efficient and robust llm evaluators, 2024b. URL [https://arxiv.org/abs/2405.](https://arxiv.org/abs/2405.11966)
847 [11966](https://arxiv.org/abs/2405.11966).
- 848 Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. Large language models are
849 not robust multiple choice selectors, 2024. URL <https://arxiv.org/abs/2309.03882>.
850
- 851 Kaitlyn Zhou, Dan Jurafsky, and Tatsunori Hashimoto. Navigating the grey area: How expressions
852 of uncertainty and overconfidence affect language models, 02 2023. URL [https://arxiv.](https://arxiv.org/abs/2302.13439)
853 [org/abs/2302.13439](https://arxiv.org/abs/2302.13439).
- 854 Kaitlyn Zhou, Jena D. Hwang, Xiang Ren, and Maarten Sap. Relying on the unreliable: The impact
855 of language models’ reluctance to express uncertainty, 2024. URL [https://arxiv.org/](https://arxiv.org/abs/2401.06730)
856 [abs/2401.06730](https://arxiv.org/abs/2401.06730).
857
858
859
860
861
862
863

A APPENDIX

A.1 REPOSITORY

The repository accompanying this paper is available online at <https://anonymous.4open.science/r/llm-uncertainty-bench-9B2B/>.

A.2 MODEL SELECTION

Table A.1: Overview of LLMs used in the experiments of this paper by provider, type, size, and release date. Model families of corresponding base, instruct and reasoning variants have been grouped together.

Provider	Model Name	Type	Size	Release Date
OpenAI	gpt-oss-20b	Reasoning	20B	Aug 2025
	gpt-oss-120b	Reasoning	120B	Aug 2025
Mistral AI	Mistral-Nemo-Base-2407	Base	8B	Jul 2024
	Mistral-Nemo-Instruct-2407	Instruct	8B	Jul 2024
	Ministral-8B-Instruct-2410	Instruct	8B	Oct 2024
	Mistral-Small-3.1-24B-Base-2503	Base	24B	Mar 2025
	Mistral-Small-3.2-24B-Instruct-2506	Instruct	24B	Jun 2025
	Magistral-Small-2507	Reasoning	24B	Jul 2025
Meta LLaMA	Llama-3.1-70B	Base	70B	Jul 2024
	Llama-3.3-70B-Instruct	Instruct	70B	Dec 2024
	Llama-4-Scout-17B-16E	Base	109B	Apr 2025
	Llama-4-Scout-17B-16E-Instruct	Instruct	109B	Apr 2025
Qwen	Qwen3-30B-A3B-Base	Base	30B	Jul 2025
	Qwen3-30B-A3B-Instruct-2507	Instruct	30B	Jul 2025
	Qwen3-30B-A3B-Thinking-2507	Reasoning	30B	Jul 2025
DeepSeek AI	DeepSeek-R1-Distill-Llama-70B	Reasoning	70B	Jan 2025
	DeepSeek-R1-Distill-Qwen-32B	Reasoning	32B	Jun 2024
Google	gemma-3-27b-pt	Base	27B	Mar 2025
	gemma-3-27b-it	Instruct	27B	Mar 2025

Please note that the model *Magistral-Small-2507* requires using a specific system prompt to enable reasoning³. We have tested the model without a system prompt (acting as an instruction tuned model) and with the system prompt provided, acting as a reasoning model. To highlight whether the reasoning behaviour was enabled through system prompt usage, we added a "Reasoning-Enabled" postfix to the model name in tables and plots.

A.2.1 MODEL CONFIGURATION AND EXCEPTIONS

All models were evaluated in their default configurations. No system prompts were employed, with the sole exception of *Magistral-Small-2507*, which requires a system prompt to enable reasoning prior to generating a final answer⁴. Without the system prompt, the model will behave like an instruction-tuned model. For the label-probability experiments, this model was tested in two variants: with reasoning enabled (*Magistral-Small-2507-Reasoning-Enabled*) and without reasoning (*Magistral-Small-2507*).

The base (pre-trained) model *gemma-3-27b-pt*, corresponding to the instruction-tuned *gemma-3-27b-it*, was excluded from the label-probability calibration experiments. In preliminary evaluations,

³<https://huggingface.co/mistralai/Magistral-Small-2507>

⁴see <https://huggingface.co/mistralai/Magistral-Small-2506>

the base model exhibited insufficient task comprehension, resulting in negligible probability mass assigned to label tokens within the top-20 most probable tokens. Consequently, label probabilities could not be retrieved through `vllm`, preventing the computation of confidence scores.

All models were used in their base configuration. With the exception of *Magistral-Small-2507*, no system prompts were used. *Magistral-Small-2507* uses the system prompt to elicit reasoning steps before providing a final answer. It was included in the label probability experiment both with (*Magistral-Small-2507-Reasoning-Enabled*) and without (*Magistral-Small-2507*) the system prompt. *gemma-3-27b-pt*, the pre-trained or base variant of *gemma-3-27b-it* was excluded from the experiment researching the label probability calibration due to lack in task comprehension. This resulted in no probability mass being assigned to label tokens within the 20 most probable tokens, rendering probabilities unaccessible via `vllm`. Therefore no certainties could be retrieved.

For the exact configuration parameters supplied to `vllm`, please refer to the `models.py` files located in the experiment subdirectories of the repository accompanying this paper.

A.3 DATASET SELECTION

Table A.2: **Scientific QA datasets surveyed.** Selection criteria emphasized natural-science domains (especially physics) and inclusion of different reasoning requirements for answering, while providing a verifiable ground-truth.

Dataset	Size	Task Format	Domain
MMLU	15,908	Multiple Choice	General (57 topics, including Physics)
ARC	7,787	Multiple Choice	Science (Physics, Chemistry, Biology, Earth Science)
SciQ	13,679	Multiple Choice	Science (Physics, Chemistry, Biology)
GPQA	448	Multiple Choice	Science (Physics, Chemistry, Biology) – Graduate-level
GSM8K	8,792	Arithmetic	Mathematics
GSM-MC	8,787	Multiple Choice	Mathematics
SVAMP	1,000	Arithmetic	Mathematics
SciBench	2,229	Arithmetic	Science (Physics, Chemistry, Biology, Medicine, Earth Science)

For the exact dataset configuration parameters, please refer to the `data_source` subdirectories located in the experiment directories of the repository accompanying this paper.

A.4 PROMPT COMPOSITION

The second experiment evaluates a total of 685,000 long-form responses across models and datasets using sequence-level UQ methods. We include five multiple-choice QA datasets (MMLU, ARC-Easy, ARC-Reasoning, SciQ, GPQA) and three arithmetic QA datasets (GSM8K, SVAMP, SciBench).

For the MCQA datasets, we apply counterfactual prompting using the APriCoT strategy, evaluating each of the four answer options independently. This expands each item into four separate prompts. All datasets are subsampled to 250 items, and we sample 10 responses per prompt. Consequently, each arithmetic dataset yields 2,500 responses per model, and each MCQA dataset yields 10,000 responses per model, totaling 57,500 responses per model.

Due to high inference costs, we reduced the number of generated responses specifically for the *deepseek-ai/DeepSeek-R1-Distill-Llama-70B* model on the GPQA dataset by half, resulting in 5,000 evaluated responses for this model–dataset pair.

Across all 12 evaluated models, this results in a total of 685,000 responses, each assessed individually using the four selected UQ methods in Experiment 2.

972 A.5 UNCERTAINTY QUANTIFICATION METHOD EXCLUSION

973
974 Our selection of UQ methods reflects principled constraints: feasibility, interpretability, and appli-
975 cability to realistic scientific reasoning tasks. An overview over exclusion of methods and short
976 form explanation is provided in Table A.3. Only methods producing normalized sequence-level
977 uncertainty scores are included in our analysis to enable reliability UQ validation by calibration.
978 Subsequently, we focus our work on the following UQ methods: *Verbalized Uncertainty* (Tian et al.,
979 2023), *P(True)* (Kadavath et al., 2022), *Frequency of Answer* (Wang et al., 2023), *Claim-Conditioned*
980 *Probability (CCP)* (Fadeeva et al., 2024). More details of these methods are discussed in Section 7.1.

981 A.5.1 SURVEY OF UQ METHODS

982
983 We used the taxonomy from LM-Polygraph Fadeeva et al. (2023a) as a starting point and adapted
984 the categorization of UQ methods into information-based, ensemble-based, density-based, reflex-
985 ive or verbalized and meaning-diversity approaches. However, it is important to emphasize that
986 LM-Polygraph targets selective generation, whereas our work focuses on calibration, the second
987 core UQ task. Calibration requires normalized confidence scores in $[0,1]$, which immediately ex-
988 cludes many of the 28 methods surveyed in LM-Polygraph. Thus, while we relied on Fadeeva
989 et al. (2023a) for a broad overview of UQ approaches, the spectrum of methods applicable to cali-
990 bration is necessarily narrower.

991 A.5.2 DENSITY-BASED METHODS

992
993 Density-based methods require access to the model’s training data to estimate likelihoods or
994 information-theoretic quantities. In our benchmark, the evaluated models are open-weight but not
995 open-source, and their training corpora are unavailable. Consequently, density-based approaches
996 cannot be applied, making them unsuitable for our calibration-focused evaluation.

997 A.5.3 ENSEMBLE-BASED METHODS

998
999 Ensemble methods combine multiple distinct models to estimate uncertainty. While potentially
1000 effective, they prevent the assessment of per-model calibration, which is central to our study. Addi-
1001 tionally, ensembles introduce substantial design ambiguity regarding which combinations of models
1002 to use and incur significant computational costs. In contrast, the Frequency of Answer method
1003 we benchmark samples a single model multiple times, preserving per-model interpretability while
1004 allowing assessment of semantic consensus as an UQ signal.

1005 A.5.4 CLAIM-LEVEL METHODS

1006
1007 Semantic Entropy as a claim-level method was excluded due to methodological and computational
1008 limitations. Existing evaluations of this method are restricted to short-form factual QA datasets
1009 and do not assess calibration, offering limited evidence for their suitability in reasoning-intensive
1010 scientific QA. Furthermore, scientific answers frequently span hundreds or thousands of tokens. In
1011 our internal tests, claim-extraction prompts often produced tens to hundreds of trivial or irrelevant
1012 claims, each requiring separate LLM evaluation. This leads to computational costs far exceeding
1013 those of the Frequency of Answer method, our computationally most expensive method employed,
1014 making claim-level approaches impractical for large-scale scientific QA.

1015 A.5.5 OVERVIEW OVER EXCLUDED METHODS

1016
1017
1018 Table A.3: **Reasons for Exclusion of Uncertainty Metrics in Benchmark.** List of UQ Meth-
1019 ods was adapted from the comprehensive evaluation and classification of uncertainty methods by
1020 Fadeeva et al. (2023a).

Method	Category	Exclusion Reasons
Perplexity (Fomicheva et al., 2020)	Information-based	Produces unnormalized scores

Method	Category	Exclusion Reasons
Mean/max token entropy (Fomicheva et al., 2020)	Information-based	Produces unnormalized scores
Monte Carlo sequence entropy (Kuhn et al., 2023b)	Information-based	Produces unnormalized scores
Pointwise mutual information (PMI) (Takayama & Arase, 2019)	Information-based	Produces unnormalized scores
Conditional PMI (van der Poel et al., 2022)	Information-based	Produces unnormalized scores
Rényi divergence (Darrin et al., 2023)	Information-based	Produces unnormalized scores
Fisher-Rao distance (Darrin et al., 2023)	Information-based	Produces unnormalized scores
Focus (Zhang et al., 2023a)	Information-based	Produces unnormalized scores
Semantic entropy (Kuhn et al., 2023b)	Meaning diversity	Designed to work at the individual claim level rather than on entire sequences Very high computational cost
TokenSAR (Duan et al., 2024)	Meaning diversity	Alters sentences in a way that violates autoregressive assumptions Relies on NLI models whose performance in scientific contexts is inadequate
SentenceSAR (Duan et al., 2024)	Meaning diversity	Alters sentences in a way that violates autoregressive assumptions Relies on NLI models whose performance in scientific contexts is inadequate
SAR (Duan et al., 2024)	Meaning diversity	Alters sentences in a way that violates autoregressive assumptions Relies on NLI models whose performance in scientific contexts is inadequate
EigenScore (Chen et al., 2024)	Meaning diversity	Produces unnormalized scores
Sentence-level ensemble-based measures (Malinin & Gales, 2020)	Ensembling	Requires running multiple independent models, leading to high computational cost Introduces extra variability that complicates comparison to single-model methods
Token-level ensemble-based measures (Malinin & Gales, 2020)	Ensembling	Requires running multiple independent models, leading to high computational cost Introduces extra variability that complicates comparison to single-model methods
Mahalanobis distance (MD) (Lee et al., 2018)	Density-based	Density-based approach that requires proprietary training data
Robust density estimation (RDE) (Yoo et al., 2022)	Density-based	Density-based approach that requires proprietary training data
Relative Mahalanobis distance (RMD) (Ren et al., 2023)	Density-based	Density-based approach that requires proprietary training data
Hybrid Uncertainty Quantification (HUQ) (Vazhentsev et al., 2023)	Density-based	Density-based approach that requires proprietary training data
Number of semantic sets (Num-Sets) (Lin et al., 2023)	Meaning Diversity	Produces count-based outputs that are not normalized to [0,1]
Sum of eigenvalues of the graph Laplacian (EigV) (Lin et al., 2023)	Meaning Diversity	Produces unnormalized scores
Degree matrix (Deg) (Lin et al., 2023)	Meaning Diversity	Produces unnormalized scores
Eccentricity (Ecc) (Lin et al., 2023)	Meaning Diversity	Produces unnormalized scores
Lexical similarity (LexSim) (Fomicheva et al., 2020)	Meaning Diversity	Relies on surface-level token overlap instead of semantic meaning
Kernel Language Entropy (Nikitin et al., 2024)	Meaning Diversity	Produces unnormalized scores
LUQ ((Zhang et al., 2024a))	Meaning diversity	Produces unnormalized scores

A.6 LABEL PROBABILITY CALIBRATION EXPERIMENT DESIGN

For the Label Probability Calibration Experiment, preliminary tests revealed that base models struggled with task comprehension: in many cases, none of the valid label tokens (A/B/C/D) appeared among the top-20 first-token predictions. To address this, we designed four prompt variants with different structural formats to identify the formulation that maximizes task comprehension. We define task comprehension as the average total probability mass assigned to the label tokens (A/B/C/D) in the model’s first generated token.

The four prompt designs are presented in Section A.6.1. An overview of task comprehension scores for each prompt is provided in Section A.6.2, where Prompt 1 emerges as the most effective. Nevertheless, even with improved task comprehension, a substantial fraction of responses remained invalid (i.e. no extractable probability assigned to any label), as detailed in Section A.6.3.

We analyze the causes of these invalid responses in Section A.6.4 and outline the mitigation strategies applied to reduce them in Section A.6.5. Finally, we report the updated invalid-answer statistics for Prompt 1 after applying these adjustments in Section A.6.3. We report the distribution of token counts for the reasoning chains for reasoning models in Section A.6.7.

A.6.1 PROMPT DESIGNS

Listing 1: Prompt Design 1 for Experiment 1. The prompt design uses 3-shot prompting. <QUESTION> and <ANSWER CHOICE X> are replaced with the individual questions and answer choices from the benchmarked dataset. Prompt Design 1 features markdown construct.

```

1 You are a highly capable language model trained for multiple-choice
2 question answering.
3 Below are three examples of multiple-choice questions with labeled answer
4 choices. Each example includes the correct answer.
5 After the examples, you will be given a new question with four labeled
6 answer choices (A, B, C, D).
7
8 Your task is to select the answer choice you believe is correct by
9 responding with only the corresponding label: A, B, C, or D.
10 Do not include any explanation or additional text.
11
12 ### Example 1:
13 **Question:** What is the capital of France?
14 A) Berlin
15 B) Madrid
16 C) Paris
17 D) Rome
18
19 **Correct Answer:** C
20
21 <Two More Examples Omitted for Readability>
22
23 ---
24
25 **Question:** <QUESTION>
26 A) <ANSWER CHOICE A>
27 B) <ANSWER CHOICE B>
28 C) <ANSWER CHOICE C>
29 D) <ANSWER CHOICE D>
30
31 **Correct Answer:**

```

Listing 2: Prompt Design 2 for Experiment 1. The prompt design uses 3-shot prompting. <QUESTION> and <ANSWER CHOICE X> are replaced with the individual questions and answer choices from the benchmarked dataset. Prompt Design 2 features no introductory text, role or task description. The format is designed to represent natural language without special formatting.

```

1 Question: What is the capital of France?

```

```

1134 2 A) Berlin
1135 3 B) Madrid
1136 4 C) Paris
1137 5 D) Rome
1138 6
1139 7 The correct answer is C
1140 8
1141 9 <Two More Examples Omitted for Readability>
1142 10
1143 11 Question: <QUESTION>
1144 12 A) <ANSWER CHOICE A>
1145 13 B) <ANSWER CHOICE B>
1146 14 C) <ANSWER CHOICE C>
1147 15 D) <ANSWER CHOICE D>
1148 16
1149 17 The correct answer is

```

1149 **Listing 3: Prompt Design 3 for Experiment 1.** The prompt design uses 3-shot prompting.
1150 <QUESTION> and <ANSWER CHOICE X> are replaced with the individual questions and answer
1151 choices from the benchmarked dataset. The format of Prompt 3 is designed to represent natural
1152 language without special formatting. The format of the answer specifically requests the label, not
1153 the answer in general.

```

1154 1 You are a highly capable language model trained for multiple-choice
1155 2 question answering. In the following examples, you will see questions
1156 3 with answer choices. The answer choices are preceded by the phrase "
1157 4 Answer Choices:". Each answer choice is annotated with one of the labels
1158 5 A, B, C or D. The correct answer to the question is given by the sentence
1159 6 "The label of the correct answer choice is" followed by the
1160 7 corresponding label. Your task is to answer the new question in the same
1161 8 format, outputting only the label of the correct answer to the question
1162 9 you are provided. Do not output anything other than one of the labels A,
1163 10 B, C or D.
1164 11
1165 12 Question: What is the capital of France?
1166 13 Answer Choices:
1167 14 A) Berlin
1168 15 B) Madrid
1169 16 C) Paris
1170 17 D) Rome
1171 18
1172 19 The label of the correct answer choice is C
1173 20
1174 21 <Two More Examples Omitted for Readability>
1175 22
1176 23 Question: <QUESTION>
1177 24 Answer Choices:
1178 25 A) <ANSWER CHOICE A>
1179 26 B) <ANSWER CHOICE B>
1180 27 C) <ANSWER CHOICE C>
1181 28 D) <ANSWER CHOICE D>
1182 29
1183 30 The label of the correct answer choice is

```

1181 **Listing 4: Prompt Design 4 for Experiment 1.** The prompt design uses 3-shot prompting.
1182 <QUESTION> and <ANSWER CHOICE X> are replaced with the individual questions and answer
1183 choices from the benchmarked dataset. Prompt Design 2 features special tags to mark the answer
1184 given as a label.

```

1185 1 You are a highly capable multiple-choice question answering model. Below
1186 2 are three examples that show the format you must follow. Each question
1187 3 has four answer choices labeled A, B, C, and D. Your task is to answer a
1188 4 new question by outputting the correct answer in the following format: <

```

```

1188 ANSWER>X<ANSWER>, where X is the label corresponding to the correct
1189 answer, A, B, C or D. Do not add any extra text or explanation.
1190 2
1191 3 Example 1:
1192 4 Question: What is the capital of France?
1193 5 Answer Choices:
1194 6 A) Berlin
1195 7 B) Madrid
1196 8 C) Paris
1197 9 D) Rome
1198 10
1199 11 <ANSWER>C<ANSWER>
1200 12
1201 13 <Two More Examples Omitted for Readability>
1202 14
1203 15 Now, please answer the following question in the same format.
1204 16
1205 17 Question: <QUESTION>
1206 18 Answer Choices:
1207 19 A) <ANSWER CHOICE A>
1208 20 B) <ANSWER CHOICE B>
1209 21 C) <ANSWER CHOICE C>
1210 22 D) <ANSWER CHOICE D>
1211 23
1212 24 <ANSWER>

```

A.6.2 TASK COMPREHENSION PER PROMPT DESIGN

Table A.4: **Task Comprehension Measured by Probability Mass Assigned to Answer Labels.** Mean over the sum of label probabilities per question for base, instruction-tuned, and reasoning models. On average, task comprehension is highest under Prompt 1.

Model Category	Prompt 1	Prompt 2	Prompt 3	Prompt 4
Base Models	0.2368	0.5928	0.4135	0.1632
Instruction-Tuned Models	0.9912	0.3597	0.9561	0.2646
Reasoning Models	0.7002	0.0895	0.4475	0.0022
Average	0.6427	0.3474	0.6057	0.1434

A.6.3 INVALID ANSWER COUNT

Although we selected the prompt that yielded the highest task comprehension for our benchmark, the original experiment exhibited a substantial number of invalid answers, varying with prompt design. These invalid outputs contaminated the results. After performing analysis on the cause of this behaviour and employing mitigation strategies, we were able to significantly decrease the amount of invalid answers, as detailed in Table A.5.

Table A.5: **Number of Invalid Answers Given by Models for Different Prompt Designs Across All Datasets** ($n = 25,316$). Invalid answers assign no probability mass to any answer-choice labels. This phenomenon is discussed above. Using the prompt design yielding the best task comprehension (as discussed in Section A.6.2), we adjusted the experiment to use structured decoding and increased the maximum token generation limit. After this adjustment, invalid answers may only occur when reasoning chains exceed 10240 tokens. The result is shown in column "Prompt 1 (SD)", drastically decreasing the amount of invalid answers.

Model	Prompt 1	Prompt 2	Prompt 3	Prompt 4	Prompt 1 (SD)
gpt-oss-20b	307	3356	498	25304	395
gpt-oss-120b	38	3323	27	24873	0
Ministral-8B-Instruct-2410	0	0	0	0	0
Mistral-Nemo-Base-2407	0	0	12574	0	0
Mistral-Nemo-Instruct-2407	0	0	0	0	0
Mistral-Small-3.1-24B-Base-2503	0	0	0	0	0
Mistral-Small-3.2-24B-Instruct-2506	0	0	0	0	0
Magistral-Small-2507	0	0	0	0	0
Magistral-Small-2507-Reasoning-Enabled	8604	18716	4749	21232	123
Llama-3.1-70B	0	13	0	6636	0
Llama-3.3-70B-Instruct	0	2527	19	125	0
Llama-4-Scout-17B-16E	2	3	0	5048	0
Llama-4-Scout-17B-16E-Instruct	0	15	0	0	0
Qwen3-30B-A3B-Base	2	1	0	55	0
Qwen3-30B-A3B-Instruct-2507	0	6270	0	722	0
Qwen3-30B-A3B-Thinking-2507	6242	11746	8656	21840	58
DeepSeek-R1-Distill-Llama-70B	425	7538	725	7477	148
DeepSeek-R1-Distill-Qwen-32B	914	7465	1170	3337	243
gemma-3-27b-pt	25316	25316	25316	25316	0
gemma-3-27b-it	0	431	0	2	0

A.6.4 ANALYSIS OF INVALID ANSWERS

Our investigation identified two primary causes of invalid responses:

- Missing answer labels within top token alternatives:** Models sometimes did not generate any of the expected labels (A/B/C/D) within the top 20 token alternatives that $vLLM$ provides per token. This occurred in:
 - Base models*, which lacked sufficient task comprehension and failed to include any of the answer labels in the top 20 token alternatives.
 - Reasoning models*, which faced the additional challenge of adhering to the expected output format after generating long-form reasoning, as the attention mechanism favours more recent tokens. Reasoning of several thousand tokens appeared to negatively effect task comprehension, leading to invalid answers in extreme cases.
- Overflowing generations:** Reasoning models produced outputs exceeding the originally set maximum token limit (4096 tokens), appearing to be stuck in prolonged reasoning cycles without converging to a final answer.

A special case is *Magistral-Small-2507*, which has reasoning enabled via the system prompt provided by MistralAI⁵. For this model, we consistently observed non-converging generations that produced a very high number of invalid answers, even on the relatively simple MMLU dataset compared to GPQA. Increasing the maximum token generation did not resolve the issue – the model still consistently hit the upper generation limit. Our solution for handling these invalid outputs is detailed in Section A.6.6.

⁵<https://huggingface.co/mistralai/Magistral-Small-2506>

1296 A.6.5 MITIGATION STRATEGIES

1297

1298 We addressed these issues as follows:

1299

- 1300 • Leveraged v_{LLM} 's structured decoding feature to enforce compliance with the expected
1301 output format after reasoning using the regex "[ABCD]".
- 1302 • Increased the maximum token limit for reasoning from 4096 to 10240 tokens, a $2.5\times$ in-
1303 crease.
- 1304 • Adjusted the generation of *Magistral* to handle infinite generations, as detailed in Sec-
1305 tion A.6.6.

1306

1307 While structured decoding enforces compliance with the expected output format, the prior evaluation
1308 of task comprehension per prompt remains essential. Forcing task compliance without adequate
1309 comprehension would result in outputs that, although syntactically valid, are semantically incorrect
1310 or inconsistent with the question, thereby undermining the reliability of uncertainty quantification
1311 and downstream analyses.

1312 Based on the prior evaluation, we chose Prompt 1 for the updated experiment and discarded the
1313 other Prompt Designs.

1314

1315 A.6.6 MAGISTRAL-SMALL-2507 INVALID ANSWER SOLUTION

1316

1317 To evaluate *Magistral-Small-2507* on label probability calibration and score distribution de-
1318 spite persistent infinite-generation issues, we first set the maximum token limit for rea-
1319 soning chain generation to 4096 tokens. Generated outputs almost always contained un-
1320 finished reasoning chains, indicated by a missing end-of-reasoning token "[/THINK]".
1321 We removed "[/THINK]" if present, trimmed any trailing incomplete sentence, and ap-
1322 pended ". I should now respond with the single label A, B, C, or D
1323 associated with the answer I consider most correct.[/THINK]" to arti-
1324 ficially terminate the chain and prompt a single-label response. The model was then queried again
1325 with this reasoning chain as a prefix, and label probabilities were extracted from the next gener-
1326 ated token. Although structured decoding cannot be applied in this prefix-prompting setup, this
1327 procedure reduced invalid answers from 8604 to 123, as shown in Table A.5.

1327

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349

A.6.7 REASONING TOKEN COUNT DISTRIBUTION FOR REASONING MODELS

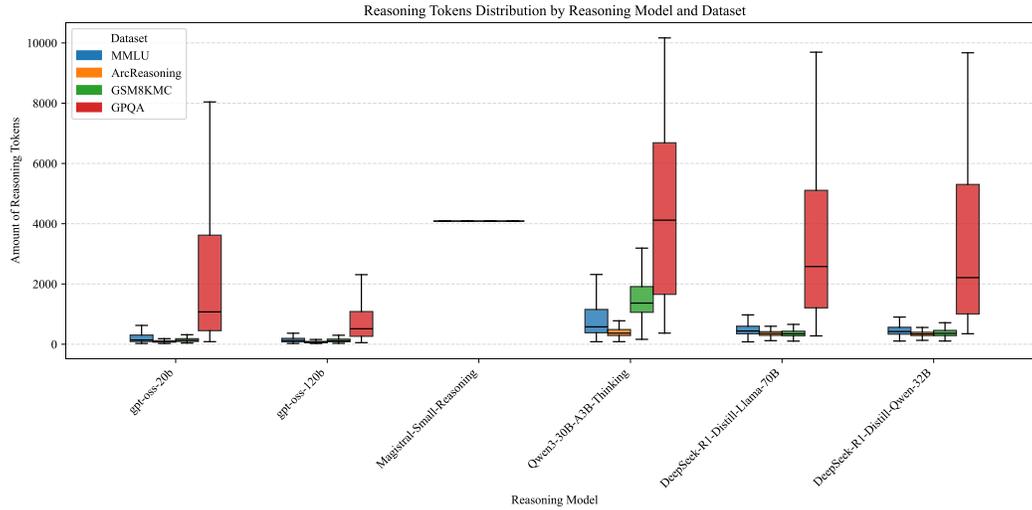


Figure A.1: **Distribution of Reasoning Tokens Produced by Models.** The figure displays a grouped boxplot illustrating the distribution of the Amount of Reasoning Tokens generated by different Reasoning Models across the individual Datasets. Models are grouped along the X-axis, with each boxplot representing the token count distribution for a specific dataset within that model. The central line in each box represents the median, and the box edges represent the first and third quartiles.

A.7 LABEL PROBABILITY EXPERIMENT RESULTS

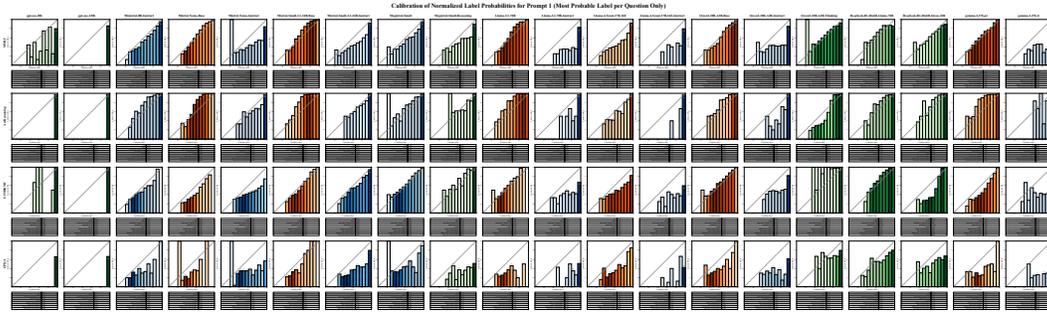
A.7.1 COMPREHENSIVE PLOTS PER PROMPT WITH TABLES

Detailed calibration plots of label probabilities for all prompt designs and configurations are available in the project repository at `label_prob_calibration/resources/figures/full_plots` for the initial prompt designs without structured decoding and at `label_prob_calibration/resources_struct_decoding/figures/full_plots` for the chosen Prompt 1 and structured decoding enabled. Each file follows the naming convention `cal_plot_prompt<id>.table<t>_chosenonly<c>.norm<n>.svg`, where the placeholders encode the following settings:

- `prompt`: Identifier of the prompt design used to generate the plot (see Section A.6.1).
- `table`: Indicator of whether a table summarizing key calibration statistics is included.
- `chosenonly`: Flag specifying whether confidence scores are computed for all candidate labels (0) or restricted to the chosen (most probable) label (1).
- `norm`: Flag denoting whether label probabilities are normalized such that their sum equals one across all candidate labels.

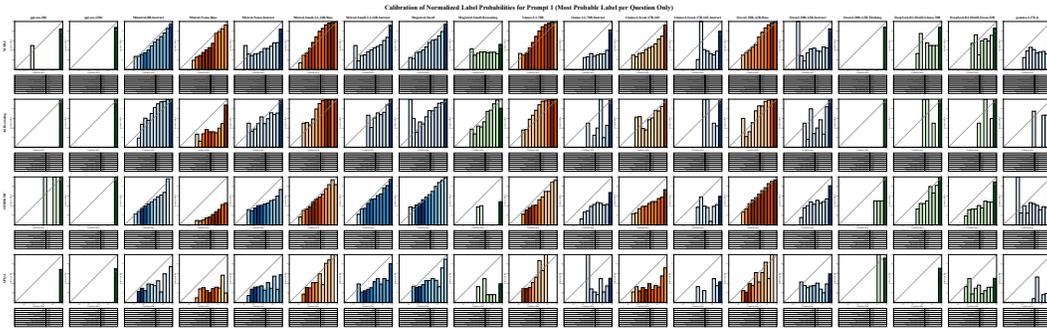
In the following, the plots showing the calibration plots for Prompt 1 (best task comprehension) can be seen.

1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414



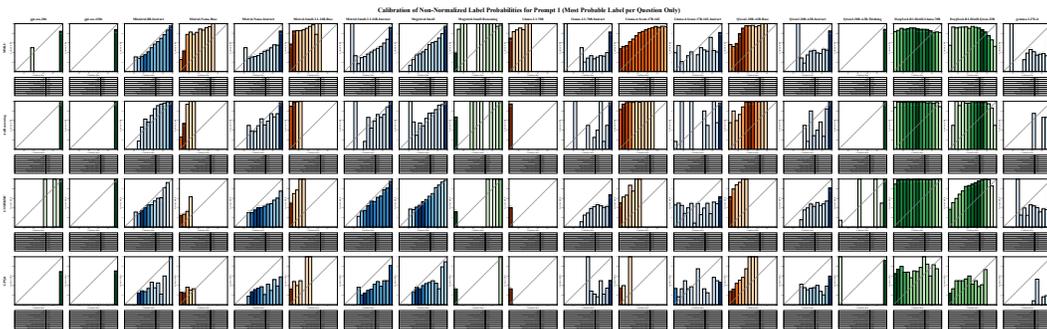
1415 **Figure A.2: Calibration Plots for Prompt 1 with structured decoding enabled, using only the most probable label.** The columns represent different models, while the rows represent datasets. Tables below the plots list summary metrics such as ECE, normalized entropy of bucket counts and AUROC. Base models are shown in orange, instruction-tuned models in blue, and reasoning models in green. Darker colors indicate a higher number of items in the bin.

1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431



1432 **Figure A.3: Calibration Plots for Prompt 1, with structured decoding disabled, using normalized label probabilities and only the most probable label.** The columns represent different models, while the rows represent datasets. Tables below the plots list summary metrics such as ECE, normalized entropy of bucket counts and AUROC. Base models are shown in orange, instruction-tuned models in blue, and reasoning models in green. Darker colors indicate a higher number of items in the bin.

1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450



1451 **Figure A.4: Calibration Plots for Prompt 1, with structured decoding disabled, using unnormalized label probabilities and only the most probable label.** The columns represent different models, while the rows represent datasets. Tables below the plots list summary metrics such as ECE, normalized entropy of bucket counts and AUROC. Base models are shown in orange, instruction-tuned models in blue, and reasoning models in green. Darker colors indicate a higher number of items in the bin.

1452
1453
1454
1455
1456
1457

A.7.2 EFFECT OF NORMALIZATION

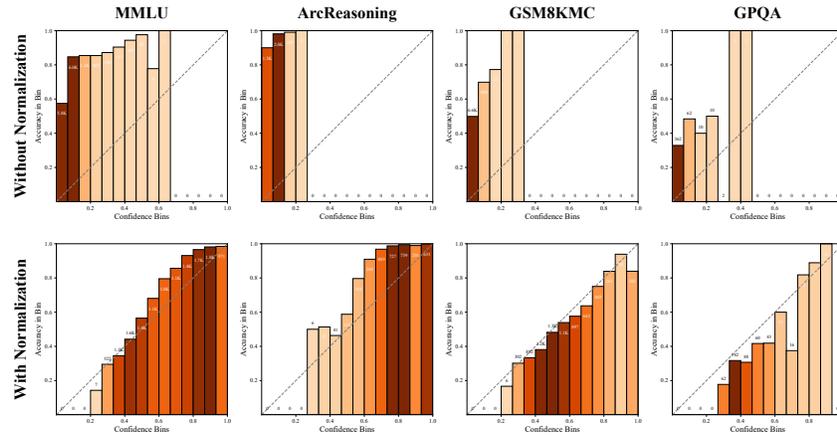


Figure A.5: **Representative Comparison of Calibration Plots for Unnormalized and Normalized Label Probabilities for the Model *Mistral-Small-3.1-24B-Base-2503* and Prompt 1 across all datasets used.** Calibration improves significantly after normalizing label probabilities.

While we had previously hypothesised that using raw label probabilities may enable the model to express general uncertainty, our results show that the task comprehension poses a bigger impact on the total probability mass assigned to label tokens. As a result, we find that raw probabilities are confounded by overall task comprehension and yield misleading calibration, while normalization (excluding non-label tokens) enables meaningful confidence estimates (see Figure A.5).

A.7.3 SELECTION BIAS

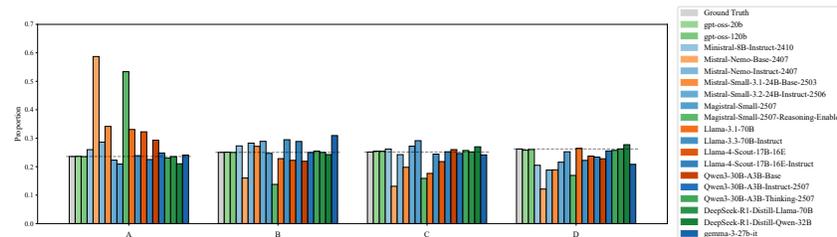


Figure A.6: **Probabilities for the Labels summed across all Datasets for each individual Model using Prompt 1.** The ground truth, represented by the distribution of the labels of the correct answers across all items in the datasets, is visualized by the grey bars and the dashed baselines.

A.7.4 TASK COMPREHENSION PER PROMPT DESIGN

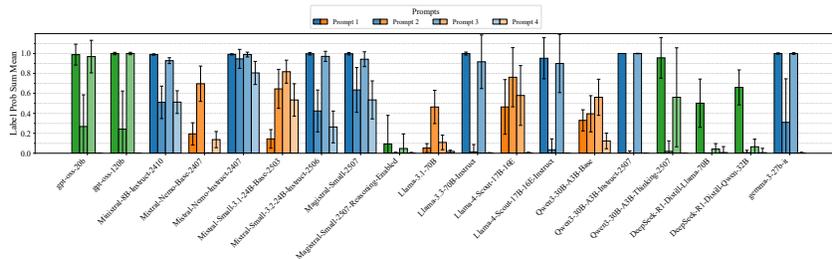


Figure A.7: **Task Comprehension per Prompt Design.** Mean over Sum of Label Probabilities for the Different Prompt Designs and Models. The data has been aggregated across all datasets, spanning 25,316 items per model and prompt design. Base models are shown in orange, instruction-tuned models in blue, and reasoning models in green.

A.8 SEQUENCE-LEVEL CALIBRATION

A.8.1 PROMPTS USED FOR QUESTION ANSWERING

The sequence-level experiments employed the APriCoT prompting strategy for MCQA and a standard Chain-of-Thought (CoT) approach for arithmetic question answering (Arithmetic QA), followed by final answer extraction. The exact prompts for both answer generation and final answer extraction are available in the repository accompanying this paper, specifically in `/llm-uncertainty-bench/seq_ue_calibration/nodes/apricot_mc_calc.py` for MCQA and `/llm-uncertainty-bench/seq_ue_calibration/nodes/arithmetric_calc.py` for Arithmetic QA.

A.8.2 METRIC IMPLEMENTATION

Verbalized Uncertainty The prompt for Verbalized Uncertainty was adopted directly from the original work (Tian et al., 2023). The exact prompt can be found in `/llm-uncertainty-bench/seq_ue_calibration/run.py`, where it is provided to the corresponding computation node as `verbalized.prompt`.

P(True) The prompt used for the P(True) metric follows the formulation of (Kadavath et al., 2022). For APriCoT prompting in MCQA, a minor adaptation was applied, while preserving the core semantics of the original design. The exact prompts are available in `/llm-uncertainty-bench/seq_ue_calibration/nodes/ptrue.py`.

Frequency of Answer For this metric, 10 samples were generated per prompt. The *Frequency of Answer* of a given response is defined as the proportion of semantically equivalent answers within the set of 10 samples. Invalid answers are assigned a confidence of 0.0. Semantic equivalence was determined according to the task type:

- **Multiple-Choice QA:** Using the APriCoT prompting strategy, each option is independently evaluated, and the model classifies each option as correct or incorrect. Semantic equivalence is established when different generations reach the same classification decision for a given option.
- **Simple Arithmetic QA:** For datasets such as GSM8K, which primarily involve integers and rarely require floating-point precision, the final numeric result was extracted using a dedicated prompt and parsed into a numeric representation. Semantic equivalence is then determined by strict numeric equality of the extracted results. For details on the final answer extraction prompt, please refer to `/llm-uncertainty-bench/seq_ue_calibration/nodes/arithmetric_calc.py` in the repository.
- **SciBench:** This dataset presents additional complexity due to intricate computations and the inclusion of physical units, rendering the simple numeric matching used

1566 for other arithmetic datasets insufficient. To address this, a specialized clustering
1567 prompt was developed to group sampled answers into semantically equivalent cate-
1568 gories, with Llama-3.3-70B-Instruct serving as the judging model. Implementa-
1569 tion details are provided in /llm-uncertainty-bench/seq_ue_calibration/
1570 leaf_nodes/answered_correctly_scibench.py.

1571 For evaluation of the calibration of the Frequency of Answer metric, the binning strategy will be
1572 slightly modified. Unlike the other methods, this methods yields only discrete confidence values,
1573 determined by the number of sampled generations. With 10 generations per question, the resulting
1574 confidence scores can only take on values from 0.1 (indicating that all of the other nine sampled
1575 generations resulted in a different answer) to 1.0 (all generations produced the same result), in steps
1576 of 0.1. For responses that fail to yield a numeric outcome in arithmetic datasets, a confidence score
1577 of 0.0 is assigned. To accommodate these discrete confidence levels, 11 bins centered on the possible
1578 certainty scores will be used for generating the calibration plots and summary statistics thereof for
1579 the Frequency of Answer metric.

1580 **Claim-Conditioned Probability (CCP)** The Claim-Conditioned Probability (CCP) metric, pro-
1581 posed by Fadeeva et al. (2024), was originally designed for claim-based uncertainty estimation.
1582 While conceptually valuable, applying CCP to long, complex generations proved challenging: ex-
1583 tracting meaningful claims was computationally expensive, often unreliable, and complicated by
1584 interdependent claims that hindered aggregation. Nevertheless, CCP was included by aggregating
1585 token-level confidence scores through multiplicative composition. The implementation builds upon
1586 the authors’ implementation of the metric in their LM-Polygraph framework (Fadeeva et al.,
1587 2023b) and was optimized for improved computational efficiency.

1588 A.8.3 EXTENSIVE CALIBRATION PLOTS

1589 In the following, extensive calibration plots for the evaluation of sequence level uncertainty methods
1590 are provided. Again, instruction tuned models are highlighted in blue, while reasoning models are
1591 highlighted in green. Darker shading indicates a higher number of items within each confidence bin.
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619

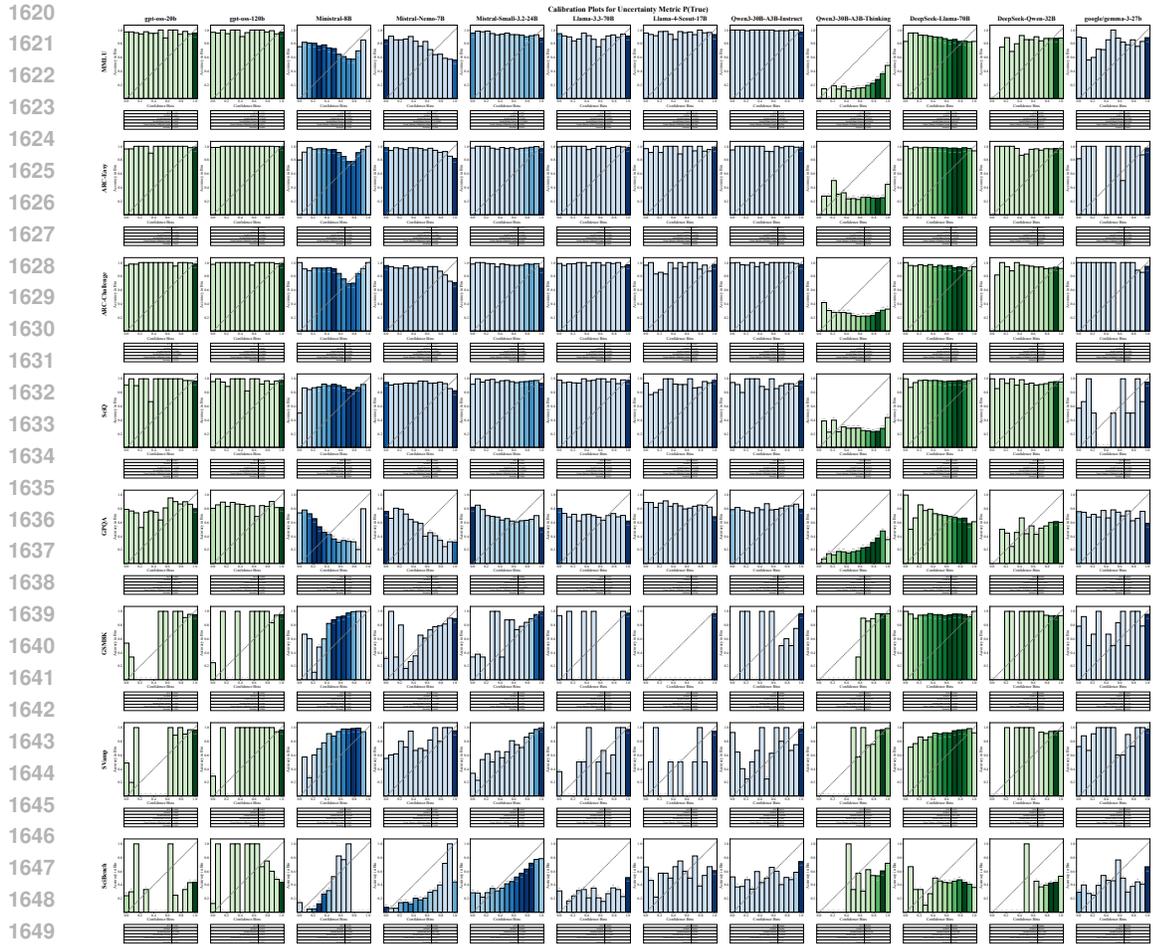


Figure A.8: **Calibration Plots for P(True)**. Columns correspond to the models and rows to datasets.

1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673

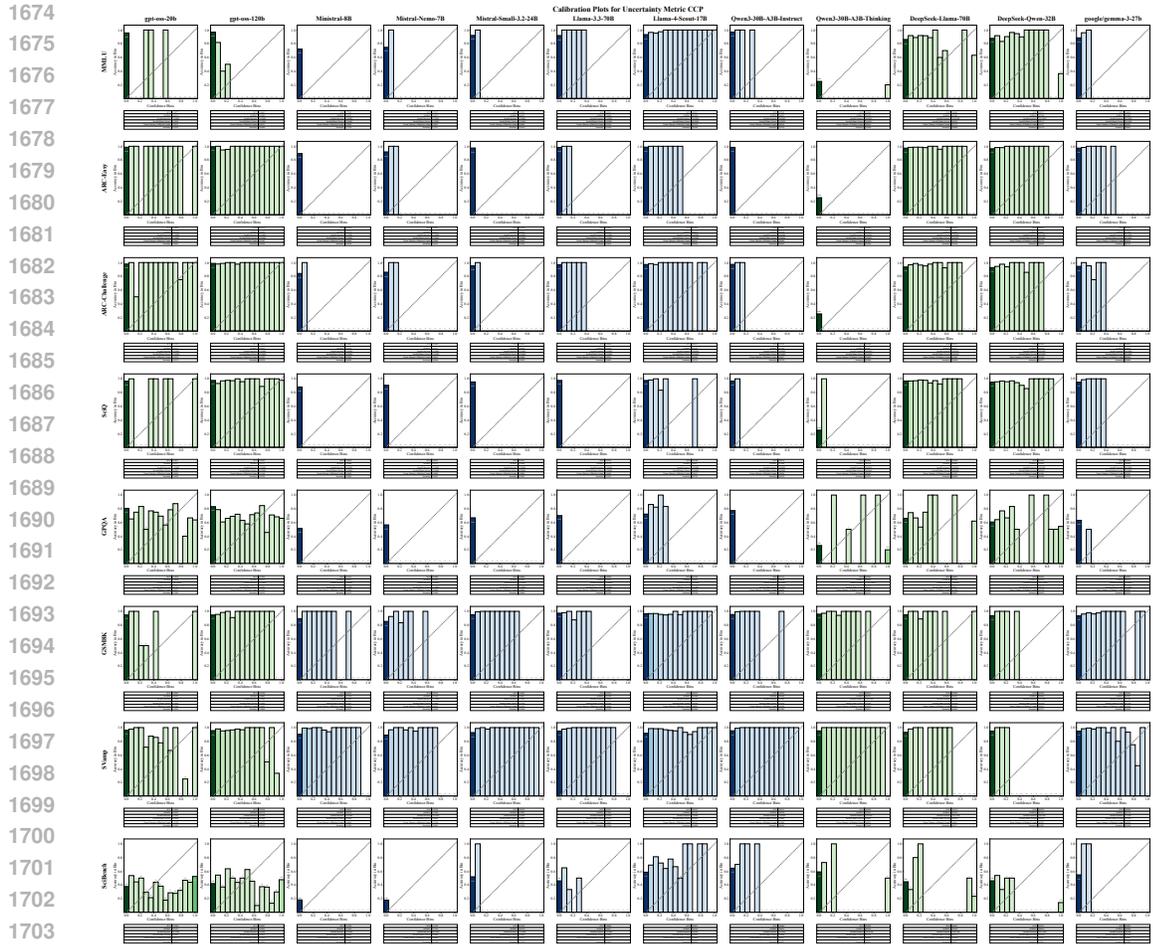


Figure A.9: Calibration Plots for CCP. Columns correspond to the models and rows to datasets.

1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727

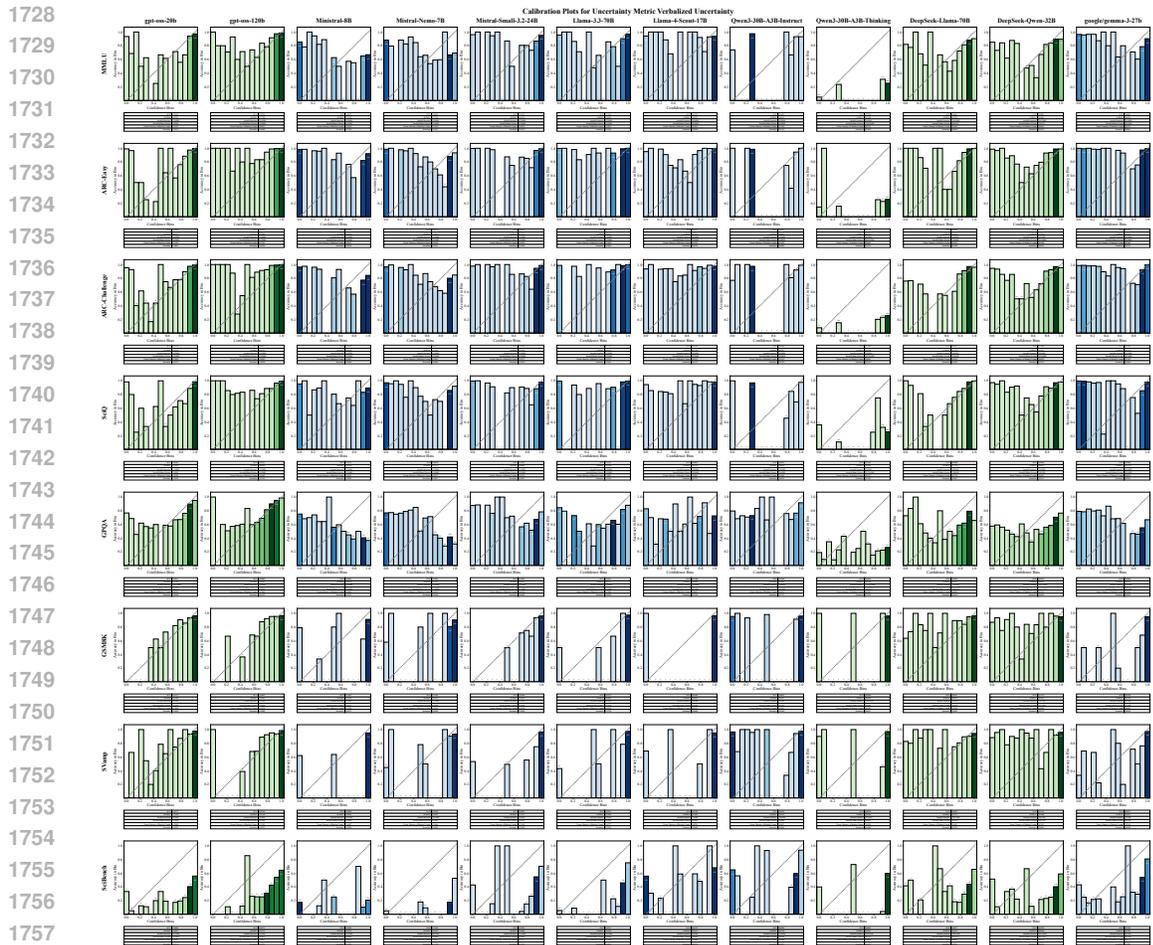


Figure A.10: Calibration Plots for Verbalized Uncertainty. Columns correspond to the models and rows to datasets.

1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781

1782
 1783
 1784
 1785
 1786
 1787
 1788
 1789
 1790
 1791
 1792
 1793
 1794
 1795
 1796
 1797
 1798
 1799
 1800
 1801
 1802
 1803
 1804
 1805
 1806
 1807
 1808
 1809
 1810
 1811
 1812
 1813
 1814
 1815
 1816
 1817
 1818
 1819
 1820
 1821
 1822
 1823
 1824
 1825
 1826
 1827
 1828
 1829
 1830
 1831
 1832
 1833
 1834
 1835

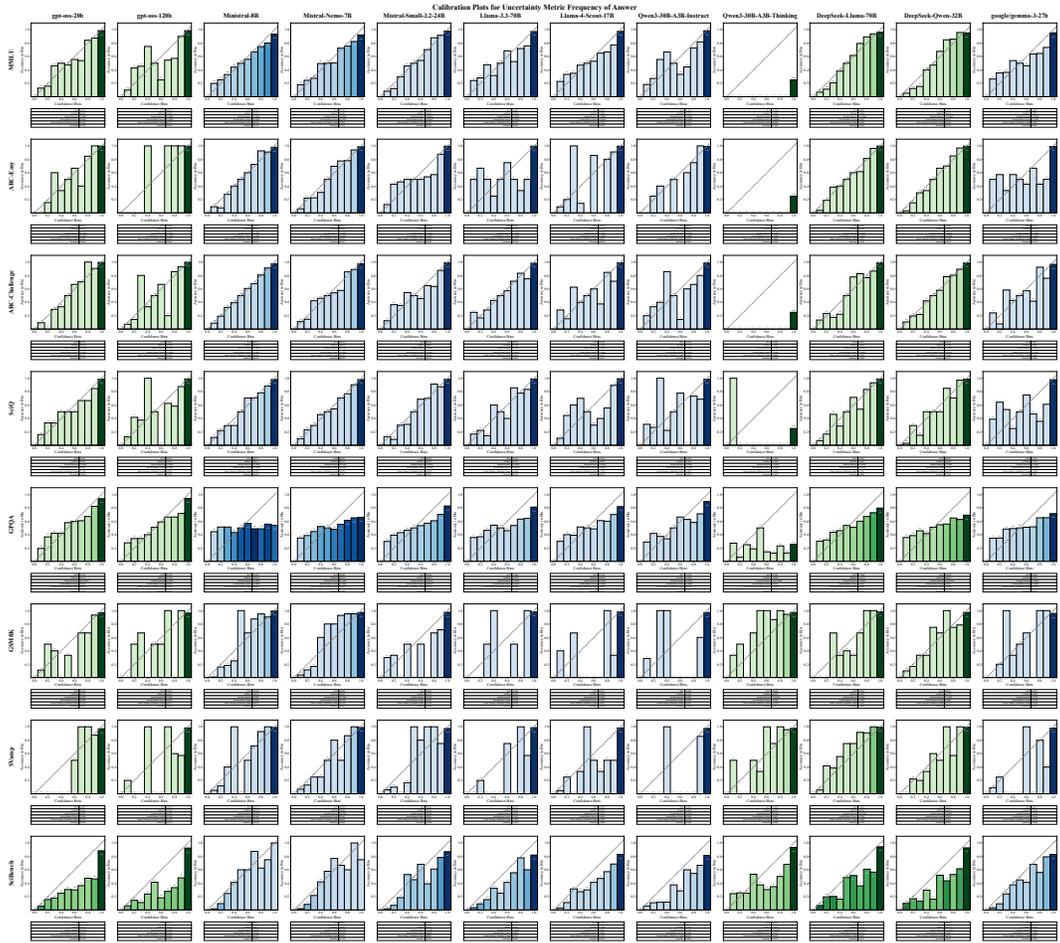


Figure A.11: Calibration Plots for Frequency of Answer. Columns correspond to the models and rows to datasets.

A.8.4 VERBALIZED UNCERTAINTY CONFIDENCE SCORE DISTRIBUTION



Figure A.12: **Distribution of Confidence Scores Provided During Verbalized Uncertainty Prompting Across Models.** Value counts are aggregated over all datasets (57, 500 prompts in total). Confidence scores seen in less than 5% of the total responses have been grouped into “Other”, with the number of distinct confidence scores shown in brackets.

A.8.5 P(TRUE) CONFIDENCE SCORE DISTRIBUTION

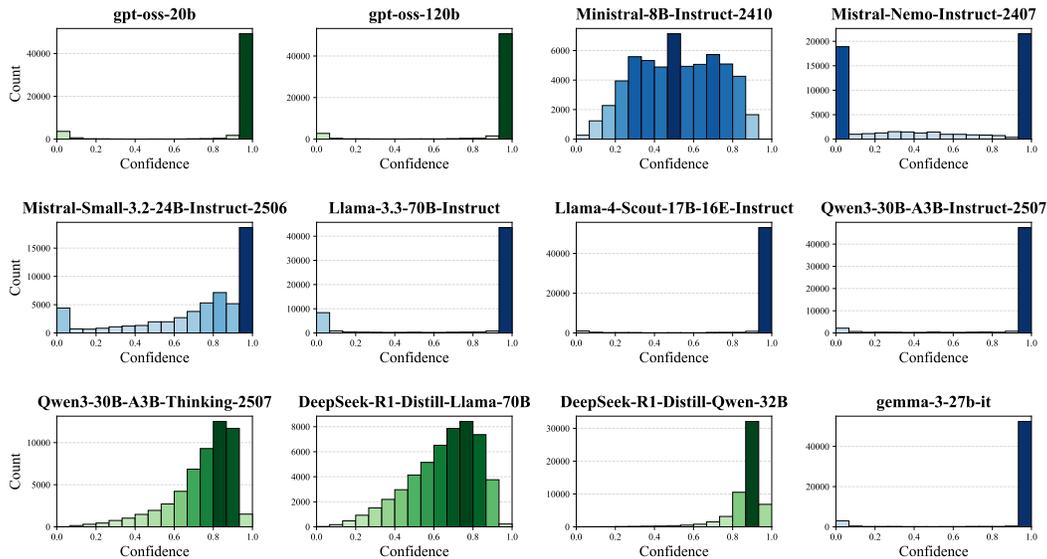


Figure A.13: **Distribution of Confidence Scores Assigned by P(True) Across Models.** Confidence scores have been aggregated across all datasets (57,500 prompts in total). Most models exhibit a clear polarization towards either (A) (representing the model’s confidence that the answer is true) or (B) (representing the model’s confidence that the answer is false) regarding the token probabilities.

A.9 EFFECT OF TEMPERATURE

The following sections clarify why additional temperature-sweep experiments would not meaningfully affect our findings or improve the interpretability of the results.

A.9.1 EXPERIMENT 1: TEMPERATURE IN TOKEN-LEVEL CALIBRATION

Experiment 1 evaluates calibration using the label probability of the first generated token under a fixed temperature of 1.0 (reasoning traces, when used, are decoded greedily at the same temperature). The central finding is a near-complete polarization of probability mass onto a single token, driven by extremely large logit margins. This collapse reflects a structural distributional issue, not a mild miscalibration. In such regimes, temperature scaling is ineffective: once margins grow this large, relative logit differences become dominated by noise and cease to carry calibration-relevant information. Adjusting temperature therefore cannot meaningfully “repair” the distribution or alter the calibration behaviour observed.

A.9.2 EXPERIMENT 2: TEMPERATURE DURING RESPONSE GENERATION

In Experiment 2, temperature affects only the generative content of responses, not the utility of the evaluated UQ methods. Since our goal is not to optimize accuracy or reduce hallucinations, but to identify responses that exhibit uncertain or erroneous reasoning, additional generations at different temperature levels would not strengthen the benchmark. Running a full temperature sweep would require repeating 685,000 generations per temperature value, with negligible methodological benefit: temperature primarily alters response variability, not the structure or comparability of the UQ scores themselves.

A.10 EXPERIMENT 2: TEMPERATURE FOR VERBALIZED METHODS

For verbalized uncertainty methods, we use temperature 1.0 with greedy decoding. This ensures the model provides the reliability judgment it is most confident in. For P(True), the situation mirrors Experiment 1: the method depends on label probabilities, which frequently exhibit polarization

1944 and show no reliable alignment with predictive accuracy. Increasing temperature cannot repair this
1945 mismatch, nor does it meaningfully alter calibration quality when the underlying distribution is
1946 already collapsed.

1947 For these reasons, additional temperature-adjustment experiments would not yield meaningful ben-
1948 efits for either experiment.
1949

1950

1951

1952

1953

1954

1955

1956

1957

1958

1959

1960

1961

1962

1963

1964

1965

1966

1967

1968

1969

1970

1971

1972

1973

1974

1975

1976

1977

1978

1979

1980

1981

1982

1983

1984

1985

1986

1987

1988

1989

1990

1991

1992

1993

1994

1995

1996

1997