

# Mind the Gap: Benchmarking LLM Uncertainty and Calibration with Specialty-Aware Clinical QA and Reasoning-Based Behavioural Features

Alberto Testoni<sup>1,2</sup>, Iacer Calixto<sup>1,2</sup>

<sup>1</sup>Department of Medical Informatics, Amsterdam University Medical Center,  
University of Amsterdam, Amsterdam, The Netherlands,

<sup>2</sup>Amsterdam Public Health, Methodology, Amsterdam, The Netherlands,

Correspondence: a.testoni@amsterdamumc.nl; i.coimbra@amsterdamumc.nl

## Abstract

Reliable uncertainty quantification (UQ) is essential when employing large language models (LLMs) in high-risk domains such as clinical question answering (QA). In this work, we evaluate uncertainty estimation methods for clinical QA focusing, for the first time, on eleven clinical specialties and six question types, and across ten open-source LLMs (general-purpose, biomedical, and reasoning models), alongside representative proprietary models. We analyze score-based UQ methods, present a case study introducing a novel lightweight method based on behavioral features derived from reasoning-oriented models, and examine conformal prediction as a complementary set-based approach. Our findings reveal that uncertainty reliability is not a monolithic property, but one that depends on clinical specialty and question type due to shifts in calibration and discrimination. Our results highlight the need to select or ensemble models based on their distinct, complementary strengths and clinical use.

## 1 Introduction

In recent years, Large Language Model (LLM) performance has consistently increased across medical and clinical tasks (Chen et al., 2023; Wu et al., 2024a; Acikgoz et al., 2024; Sellergren et al., 2025). This includes classification tasks with finite output spaces, such as multiple-choice question answering (MCQA; Jin et al., 2019, 2021; Pal et al., 2022), as well as natural language generation (NLG) tasks with open-ended output spaces (Baan et al., 2023), such as open-ended question answering (Open-QA; Thirunavukarasu et al., 2023; Singhal et al., 2023) and clinical text summarization (Schoonbeek et al., 2024; Van Veen et al., 2024). However, LLMs are known to be overconfident (Groot and Valdengro Toro, 2024), raising concerns about their reliability in high-stakes clinical settings. In such settings, reliability depends not only on *discrimination* (how well confidence scores separate correct

vs. incorrect answers) but also on *calibration* (how closely predicted probabilities match actual accuracy); both are critical for safe deployment (Atf et al., 2025). By quantifying uncertainty, LLMs can spot unreliable predictions and prompt expert oversight if needed (Zhou et al., 2025).

Let  $x \in \mathcal{X}$  be a question and  $y \in \mathcal{Y}$  its answer. Typically, in a classification setting such as MCQA,  $\mathcal{Y}$  is a finite set and quantifying the entropy of an LLM’s predictive uncertainty  $p_\theta(y|x)$  for this problem can be done directly by enumerating all the possible answers. However, uncertainty quantification (UQ) in NLG is a significant challenge (Giulianelli et al., 2023; Geng et al., 2024). In Open-QA, the output space  $\mathcal{Y}$  is the set of all possible strings, which grows exponentially with the length of the answer. Directly calculating the predictive uncertainty  $p_\theta(y|x)$  for every possible output becomes unfeasible. In this work, we benchmark how representative UQ approaches for NLG—token-based confidence (Huang et al., 2025; Aichberger et al., 2024), sample consistency (Kuhn et al., 2023), confidence elicitation (Xiong et al., 2024), and conformal prediction (Campos et al., 2024)—fare specifically in clinical QA. We benchmark uncertainty across 10 LLMs (general-purpose, biomedical, and reasoning-focused) and 2 clinical QA datasets, including 11 different clinical specialties (among which 5 can be considered less frequent; see Section 3 for details). We also introduce a novel annotation scheme that maps QA pairs onto 6 different question types (e.g., diagnosis, treatment, definition questions, etc).

We also present a new case study on behavioral signals from LLM reasoning traces, where we offer a lightweight, interpretable, and single-pass alternative to sampling-based methods. We find that reasoning LLMs often yield superior accuracy and calibration, but this advantage varies significantly across specialties and question types.

We offer a cautionary yet constructive perspec-

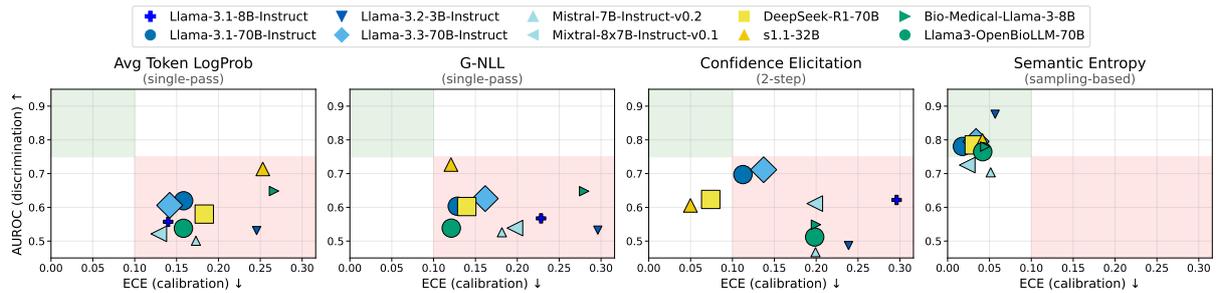


Figure 1: Uncertainty estimation for clinical QA across LLMs featuring AUROC (discrimination) vs. ECE (calibration). Marker size indicates model size; color and shape indicate model type (general, reasoning, biomedical). The green-shaded area (top left) highlights the desirable region: low calibration error and high discrimination. Most single-pass methods fail on both; even G-NLL (Aichberger et al., 2024) underperforms despite good performance reported on non-clinical QA. Semantic Entropy achieves strong performance, though it requires multiple generations.

tive on uncertainty in clinical NLP, pointing toward more context-aware evaluation. Our results show that aggregate metrics obscure important reliability differences and highlight the importance of uncertainty evaluations that account for clinical specialty and question type. We see this work as a first step towards supporting the development of LLMs that are both accurate and trustworthy in clinical use. Our code is available at <https://github.com/albertotestoni/mind-the-gap-clinical-qa>.

## 2 Related Work

Starting from the seminal work of Guo et al. (2017), uncertainty quantification has been studied in NLP tasks such as sentiment analysis, machine translation, and named entity recognition (Xiao and Wang, 2019; Zerva et al., 2022; Hashimoto et al., 2024; Shorinwa et al., 2025). While uncertainty estimation has a long tradition in classification tasks, the rise of LLMs has renewed interest in extending these methods to language generation (Geng et al., 2024; Ulmer et al., 2024; Liu et al., 2025; Huang et al., 2025). Ensuring well-calibrated predictions is essential in clinical NLP, where overconfidence or miscalibrated uncertainty estimates can translate into patient risk or diagnostic error (Ghoshal et al., 2021; Kompa et al., 2021; Dawood et al., 2023).

Few approaches have systematically evaluated these aspects in the clinical domain. Gao et al. (2025) argue that LLM uncertainty signals (e.g., token-level entropy) are poorly suited for diagnosis tasks, but their analysis remains limited to a small set of outdated models and focuses solely on classification tasks. Similarly, studies on open-ended generation (Wu et al., 2024b; Savage et al.,

2025) evaluate only a narrow range of methods and models, limiting the generalizability of their findings. Moreover, existing evaluations typically report aggregated metrics, offering limited insight into finer-grained dimensions such as medical specialties or question types. Another still underexplored aspect is the comparison across different LLM types. Recent studies have evaluated general-purpose and biomedical LLMs on clinical tasks, reporting mixed evidence for the latter’s advantage (Dorfner et al., 2024). However, these comparisons largely overlook uncertainty calibration.

A recent related study to ours is Savage et al. (2025), which investigates uncertainty estimation methods for clinical question answering. The authors found that sample consistency (assessing agreement across multiple answers generated via stochastic decoding) outperformed other approaches in terms of both discrimination and calibration, though the latter remained suboptimal. While their findings provide valuable initial insights into UQ for medical QA, their evaluation is limited in several key ways. First, it focuses on a small set of outdated models, excluding recent advances such as reasoning-tuned and domain-specific LLMs (which we highlight in Figure 1, along with the calibration–discrimination trade-off). Second, their analysis is restricted to diagnosis and treatment questions, ignoring other clinically relevant question types. Third, they do not examine how performance varies across medical specialties. A complementary line of work is Ye et al. (2024), which applies conformal prediction to evaluate uncertainty in LLMs across five (non-clinical) NLP tasks, such as question answering and text summarization. Notably, they find that (i) higher accuracy does not necessarily correlate

with higher certainty, and (ii) larger LLMs can exhibit greater uncertainty than smaller ones. In our work, we investigate whether these patterns extend to clinical question answering.

### 3 Datasets and Question Annotation

We focus on two dimensions of clinical question answering that have been largely overlooked in prior work: medical specialty and question type, which induce systematic variation in task difficulty and uncertainty profiles that is obscured by aggregate evaluation.

**By Medical Specialty** We leverage two English multiple-choice question answering datasets. <sup>1</sup> **S-MedQA** (Yan et al., 2026) combines questions from MedQA (Jin et al., 2021) and MedMCQA (Pal et al., 2022). The dataset includes a validated annotation pipeline that maps each question to one of 55 medical specialties recognized in the European Union. We consider the six most frequent specialties: *Gastroenterology*, *Cardiology*, *Obstetrics and Gynecology*, *Neurology*, *Infectious Diseases*, and *Pediatrics*, which we refer to in the following as ‘common specialities’. To ensure balanced evaluation across specialties while keeping the total number of questions manageable for computational and analytical purposes, we fix the sample size to 145 questions per specialty, resulting in a total of 870 questions. Each question in S-MedQA includes four answer options and one correct answer. **MedExQA** (Kim et al., 2024) focuses on underrepresented specialties: *Biomedical Engineering* (143 questions in the MedExQA test set), *Clinical Laboratory Science* (372 questions), *Clinical Psychology* (106 questions), *Occupational Therapy* (189 questions), and *Speech Language Pathology* (130 questions), summing up to 940 questions. Questions in MedExQA were manually gathered from a variety of freely accessible online sources, such as mock tests and practice exams designed for specific specialties. Each question in MedExQA includes four answer options and one correct answer.

**By Question Type** After manually inspecting a random sample of 50 questions from each dataset, we identify six clinically meaningful question types: *Diagnosis*, *Treatment*, *Diagnostic Test*, *Definition*, *Procedure/Operation*, and *Other*. We annotate all 1,810 questions using

Llama-3.1-70B-Instruct (Meta AI, 2024b), and validate its outputs against manual annotations from one of the authors on a sample of 100 questions. The agreement is high, with 94% accuracy; the few disagreements involve borderline cases of the broad *Other* category. Full details on the prompt, question type definitions, and distribution are provided in Appendix F.

## 4 Methodology

### 4.1 Extracting Uncertainty Estimates

Uncertainty in LLMs can be quantified by exploiting the probability distributions induced by the model over its outputs. In the context of question answering, these distributions can be summarized as scalar scores or prediction sets that reflect how confident a model is in its chosen answer. We consider multiple strategies for extracting such uncertainty estimates, ranging from single-pass methods based on token probabilities to sampling-based approaches and set-valued predictions.

**Score-based methods:** Following the taxonomy of Geng et al. (2024), we consider *logit-based* approaches, such as average token-level probability under greedy decoding (Huang et al., 2025). Consistent with Savage et al. (2025), we found negligible differences when using minimum or maximum token probability. We also include G-NLL (Aichberger et al., 2024), which computes the negative log-likelihood of the greedily decoded output and serves as an efficient uncertainty proxy.<sup>2</sup> For *linguistic confidence* methods, we adopt a two-step confidence elicitation approach (Rivera et al., 2024; Xiong et al., 2024), where the model first produces an answer and is then prompted to reprocess the full question–answer pair and assign a scalar certainty score. We use the same elicitation prompt as Savage et al. (2025).

**Consistency-based methods:** We evaluate Semantic Entropy, inspired by Kuhn et al. (2023); Farquhar et al. (2024) and related to Sample Consistency (Savage et al., 2025). While Sample Consistency measures agreement across sampled outputs, Semantic Entropy quantifies the entropy of the predicted answer distribution. Given the multiple-choice answers available in the datasets we used, we group generations by selected option and compute entropy over the resulting distribution. Higher

<sup>1</sup>In our setup, models generate open-ended responses, and correctness is evaluated against multiple-choice labels (§4.3)

<sup>2</sup>In G-NLL (Aichberger et al., 2024), calibration was not included in the authors’ original analysis.

entropy reflects greater uncertainty. We follow [Manakul et al. \(2023\)](#) and generate 10 responses per question.

**Set-based methods:** As for a representative set-based method, we evaluate a conformal prediction approach inspired by [Ke et al. \(2025\)](#). The method constructs prediction sets by sampling multiple responses per input and estimating empirical coverage over answer options; in our setup, we target 90% marginal coverage of the correct answer. Appendix H provides additional details and an illustrative example.

**Linguistic behavioral features:** We also introduce a novel case study based on *linguistic behavioral features* extracted from reasoning traces, an understudied category in current taxonomies of uncertainty estimation methods for LLMs. We provide additional methodological details and analysis of this approach in Section 5.2.

Finally, following the classification in [Liu et al. \(2025\)](#), we can further distinguish between *single-pass* methods (logit-based scores, confidence elicitation, and linguistic behavioral features) and *sampling-based* methods, which require multiple generations (consistency-based approaches). The linguistic confidence method specifically involves two passes per input.

## 4.2 Models

We evaluate 10 open-source instruction-tuned LLMs to comprehensively estimate uncertainty across model architecture, scale, and domain specialization. We include six **general-purpose** models of varying sizes and architectures: LLaMA-3.1-8B, LLaMA-3.1-70B, LLaMA-3.2-3B, LLaMA-3.3-70B ([Meta AI, 2024a](#)), Mistral-7B-v0.2 and Mistral-8x7B-v0.1 ([Mistral AI, 2023](#)); two **reasoning-oriented** models: DeepSeek-R1-70B ([Guo et al., 2025](#)) and s1.1-32B ([Muennighoff et al., 2025](#)); and two **domain-specific biomedical** models: Bio-Medical-LLaMA-3-8B ([ContactDoctor, 2024](#)) and LLaMA3-OpenBioLLM-70B ([Ankit Pal, 2024](#)), both fine-tuned on biomedical and clinical corpora. In Appendix A, B, and C we provide details about the models, prompts, and decoding hyperparameters, respectively. Appendix G presents a focused analysis on three **OpenAI ChatGPT models** (GPT-4.1, GPT-4.1-mini, and GPT-4o-mini).

## 4.3 Evaluation

**Answer Extraction and Accuracy** Models are prompted to output a long-form answer that includes the letter of the selected option in square brackets (e.g., [C]), which we extract via regular expressions. Only when the format deviates (on average in 5–6% of responses, e.g., full answer text within brackets instead of the letter alone), we use LLaMA-3.1-8B-Instruct to map the response to one of the four options (see Appendix B for additional details). We validate its performance on this task against human annotations over 100 samples, observing 92% accuracy. *QA accuracy is computed by comparing the extracted answer to the ground-truth label*, whereas *uncertainty estimates are computed over the full (long-form) answer generated by the LLM*. To control for position bias, answer options are shuffled four times per question with greedy decoding. Semantic Entropy and Conformal Prediction require sampling (via nucleus sampling) 10 and 20 responses per question, respectively, with the options shuffled at each generation to mitigate position bias.

**Discrimination and Calibration** To evaluate discrimination, i.e., the ability of uncertainty estimates to distinguish correct from incorrect predictions, we compute the Area Under the Receiver Operating Characteristic Curve (AUROC). AUROC measures how well uncertainty scores separate correct from incorrect predictions; a value of 1.0 indicates perfect discrimination, while 0.5 corresponds to random guessing. We assess calibration, i.e., the alignment between predicted confidence and empirical accuracy, using Expected Calibration Error (ECE) and the Brier Score ([Brier, 1950](#)). ECE measures the discrepancy between predicted confidence and actual accuracy. Following [Rivera et al. \(2024\)](#), we use 10 bins to compute ECE. The Brier Score captures the squared error between predicted confidence and correctness, providing a proper scoring rule that reflects both calibration and sharpness. For both scores, a value of 0 indicates perfect calibration. Given that ECE and Brier Score rank models similarly, we report their average for some of the analyses below.

**Conformal Prediction** For Conformal Prediction, we evaluate uncertainty in terms of both empirical coverage (the proportion of test instances for which the correct answer is included in the prediction set) and average set size (how many op-

Model	Gastro.	Cardiology	Obs. and gyn.	Neurology	Infectious dis.	Pediatrics	Biomed. Eng.	Clinical Lab.	Clinical Psy.	Occ. Therapy	Speech-Lang. Path.
Llama-3.1-8B	0.623	0.633	0.651	0.617	0.601	0.616	0.714	0.742	0.785	0.784	0.286
Llama-3.1-70B	<b>0.769</b>	0.781	<b>0.828</b>	0.774	<b>0.750</b>	<b>0.848</b>	<b>0.868</b>	0.887	<b>0.925</b>	<b>0.870</b>	<b>0.342</b>
Llama-3.2-3B	0.538	0.484	0.512	0.543	0.534	0.512	0.572	0.661	0.665	0.713	0.285
Llama-3.3-70B	<b>0.769</b>	<b>0.812</b>	0.809	0.719	0.748	0.769	0.834	0.881	0.851	0.848	<b>0.335</b>
Mistral-7B	0.463	0.529	0.505	0.452	0.401	0.440	0.787	0.700	0.764	0.762	0.287
Mixtral-8x7B	0.590	0.607	0.641	0.573	0.512	0.584	0.811	0.794	0.835	0.792	0.285
DeepSeek-R1-70B	<b>0.828</b>	<b>0.833</b>	<b>0.841</b>	<b>0.831</b>	<b>0.862</b>	<b>0.886</b>	<b>0.877</b>	<b>0.904</b>	<b>0.932</b>	0.856	0.310
s1.1-32B	<b>0.847</b>	<b>0.845</b>	<b>0.850</b>	<b>0.819</b>	<b>0.867</b>	<b>0.867</b>	0.844	<b>0.896</b>	<b>0.925</b>	<b>0.861</b>	0.318
Llama3-OpenBioLLM-70B	0.719	0.757	0.759	<b>0.777</b>	0.723	0.779	<b>0.880</b>	<b>0.897</b>	0.897	<b>0.865</b>	<b>0.345</b>
Bio-Medical-Llama-3-8B	0.653	0.647	0.705	0.648	0.659	0.643	0.743	0.722	0.818	0.737	0.273

Table 1: Multiple-choice QA accuracy by medical specialty, revealing substantial heterogeneity in model performance across domains. For each specialty, the three best-performing models are highlighted in bold.

tions are included on average in the prediction set). We target a nominal coverage level of 90%, and report the realized coverage. Intuitively, a lower average set size indicates more confident (i.e., less uncertain) predictions, while higher empirical coverage reflects greater reliability. However, the two must be balanced: low set size with low coverage suggests overconfidence, while high coverage with large sets indicates conservativeness.

## 5 Score-based UQ Methods Results

We begin by evaluating score-based uncertainty methods by analyzing how their scores relate to prediction correctness, focusing on calibration and discrimination across all models and questions.

### 5.1 From Logits to Semantic Entropy

We compare logit-based, confidence elicitation, and consistency-based methods to assess (a) which yields the most calibrated and discriminative estimates, and (b) whether LLM type or size systematically affect them. Figure 1 shows AUROC vs. ECE for all LLMs and UQ methods. In line with the findings in Savage et al. (2025), *Semantic Entropy* shows the strongest overall performance while requiring multiple generations per input. Reasoning-oriented and biomedical LLMs do not consistently outperform general-purpose models in terms of uncertainty estimates. G-NLL does not perform well in either dimension, although it performs better on the recent s1.1-32B reasoning model. In general, larger models tend to yield more reliable uncertainty estimates (particularly for Confidence Elicitation), though this trend is less pronounced for *Semantic Entropy*, which performs robustly across nearly all models, including the smallest (LLaMA-3.2-3B). By contrast, *Mistral-7B* and

*Mixtral-8x7B* underperform relative to the rest.

Based on these findings, we adopt *Semantic Entropy* as our primary uncertainty estimation method for the remainder of this section. We revisit single-generation alternatives in Section 5.2. While *Semantic Entropy* provides well-calibrated uncertainty estimates across a broad range of models, two important questions remain: *Does this performance hold uniformly across medical specialties and question types?* And *how are these uncertainty estimates aligned with actual QA accuracy?*

#### 5.1.1 Impact of Medical Specialty

**Model Accuracy:** Table 1 reports model accuracy across medical specialties, with the top three models per specialty highlighted in bold. We observe a consistent effect of model size: larger models generally achieve higher accuracy than their smaller counterparts, particularly within the LLaMA family. Reasoning models (s1.1-32B and DeepSeek-R1-70B), trained on reasoning traces, frequently rank among the top performers, suggesting that prior exposure to reasoning data contributes to robust performance across domains. Biomedical models tend to outperform general-purpose models at smaller scales (e.g., 8B), but this advantage does not hold at larger scales, where models such as LLaMA-3.1-70B surpass their domain-specific counterparts (e.g., Llama3-OpenBioLLM-70B). Notably, all models perform only marginally above chance on *Speech-Language Pathology*, echoing findings from the original dataset authors (Kim et al., 2024). Overall, these observations raise the question of whether more accurate models are also more calibrated or reliable in their predictions, and how this aligns with findings by Ye et al. (2024), who report that

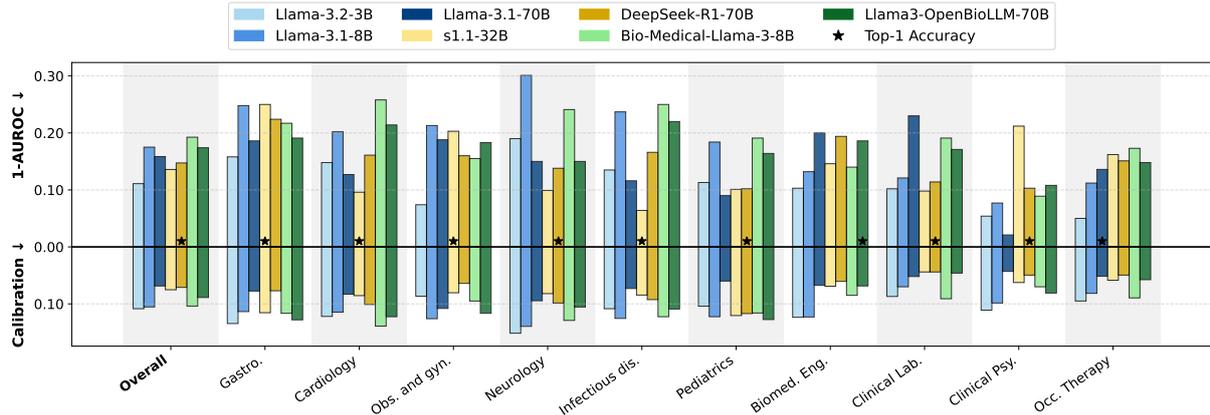


Figure 2: Discrimination ( $1 - \text{AUROC}$ ) and calibration (average ECE and Brier score) of semantic entropy-based uncertainty estimates across specialties.  $\star$  indicates the most accurate model per specialty.

larger-scale LLMs may exhibit greater uncertainty.

**Discrimination vs Calibration:** Figure 2 shows discrimination ( $1 - \text{AUROC}$ ) and calibration (avg. ECE and Brier score) for a representative set of models. While larger models generally achieve higher accuracy (no model  $< 32\text{B}$  ranked top-3 in any specialty), this does not guarantee superior uncertainty estimates. For instance, Llama-3.2-3B achieves the highest AUROC, significantly outperforming the second-best model, s1.1-32B ( $p = 0.03$ , bootstrap test). However, it simultaneously ranks among the worst in calibration, highlighting a clear decoupling between discrimination and calibration. Interesting trade-offs emerge between accuracy and calibration, e.g., in *Pediatrics*, DeepSeek-R1-70B achieves the highest accuracy ( $\sim 89\%$ ) but has worse calibration than Llama-3.1-70B ( $p = 0.05$ , Brier score), whose accuracy is slightly lower ( $\sim 85\%$ ) but much better calibrated. This suggests that gains in accuracy can come at the cost of overconfidence.

**Model Specialization:** Smaller models occasionally demonstrate strong discrimination despite low accuracy. In *Obstetrics and Gynecology*, Llama-3.2-3B answers just  $\sim 51\%$  correctly but achieves the best AUROC, suggesting it is reliably uncertain when unsure. Conversely, some larger models show worse AUROC despite higher accuracy, indicating overconfidence on their few errors. DeepSeek-R1-70B consistently exhibits strong calibration, achieving the lowest average ECE and Brier scores in about half the specialties. Meanwhile, domain-specialized models such as Llama3-OpenBioLLM-70B excel in some focus areas (e.g., *Biomedical Engineering*) but underper-

Model	Diagnosis	Treatment	Diagnostic Test	Definition	Procedure/Operation	Other
Llama-3.1-8B	0.620	0.646	0.659	0.730	0.618	0.558
Llama-3.1-70B	<b>0.790</b>	0.772	0.818	<b>0.853</b>	<b>0.686</b>	0.733
Llama-3.2-3B	0.538	0.543	0.570	0.645	0.443	0.448
Llama-3.3-70B	0.766	0.761	0.805	0.835	<b>0.633</b>	0.688
Mistral-7B-v0.2	0.482	0.503	0.618	0.720	0.549	0.530
Mixtral-8x7B-v0.1	0.588	0.598	0.719	0.776	0.611	0.616
DeepSeek-R1-70B	<b>0.839</b>	<b>0.806</b>	<b>0.863</b>	<b>0.853</b>	0.631	<b>0.772</b>
s1.1-32B	<b>0.848</b>	<b>0.795</b>	<b>0.833</b>	0.845	0.617	<b>0.819</b>
Llama3-OpenBioLLM-70B	0.739	<b>0.775</b>	<b>0.825</b>	<b>0.869</b>	<b>0.715</b>	<b>0.765</b>
Bio-Medical-Llama-3-8B	0.657	0.639	0.627	0.737	0.591	0.597

Table 2: Accuracy per question type. Top-3 models per question type are shown in bold.

form on general domains like *Cardiology* or *Infectious Disease*, highlighting the limited generalizability of domain tuning. Together, these results underscore that high performance in one evaluation dimension (accuracy, discrimination, calibration) does not imply strength in the others, reinforcing the need to evaluate all three dimensions for safe deployment.

### 5.1.2 Impact of Question Type

Table 2 reports accuracy by question type, leveraging the annotation pipeline introduced in Section 3. Several trends observed in the per-specialty analysis persist: larger LLaMA models consistently outperform their smaller counterparts, and reasoning-oriented models frequently rank among the top in accuracy. Our new question-type annotations also reveal *Procedure/Operation* as a particularly challenging category, with all models underperforming and only OpenBioLLM-70B exceeding 70% accuracy. By contrast, *Definition* questions emerge as the easiest category, with nearly all models achiev-

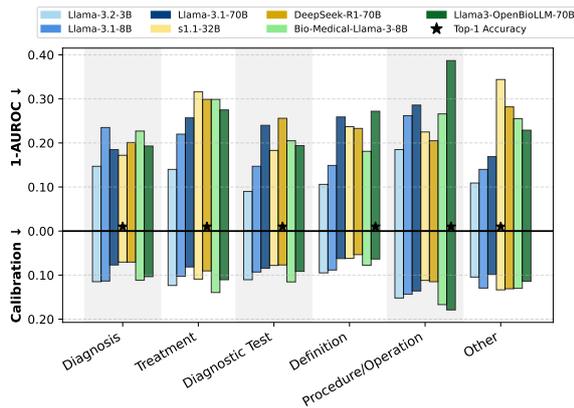


Figure 3: Calibration/discrimination per question type using Semantic Entropy estimates.

ing high accuracy. This may reflect the fact that definition-style queries are closer in form and content to pretraining data, often resembling glossary entries or encyclopedic descriptions.

The quality of uncertainty estimates varies substantially across question types (Figure 3). On harder question types (in terms of accuracy), models not only make more errors but also exhibit weaker confidence separation between correct and incorrect answers. Interestingly, the smallest model (3B) achieves the best discrimination across all question types despite the lowest accuracy and poor calibration, confirming the pattern observed for medical specialties. No single model dominates across all types: Llama3-OpenBioLLM-70B achieves the highest accuracy on the challenging *Procedure/Operation* questions, yet also exhibits the worst calibration and discrimination on the same category, suggesting that its confident predictions are often misaligned with actual correctness.

## 5.2 Behavioural Signals of Uncertainty in Reasoning Models: A Case Study

Although performance varies across medical specialties and question types, Semantic Entropy substantially outperforms sentence-level logprob methods in both calibration and discrimination. While effective, the method requires repeated generations: this feature may hinder its deployment in real-world applications, particularly in latency-sensitive or resource-constrained settings, such as clinical decision support. In contrast, single-pass methods like G-NLL and average token probability require only one forward pass per input, albeit with reduced discrimination and calibration performance.

To address this trade-off, we explore whether be-

havioral signals in LLM reasoning traces can serve as lightweight proxies for uncertainty. Our motivation stems from empirical observations: models like DeepSeek-R1-70B and s1.1-32B often produce intermediate reasoning steps with linguistic cues that appear to reflect uncertainty. For example, we frequently observe self-directed questions and tokens such as “Wait...”, which may signal self-verification or reconsideration. We also hypothesize that longer responses may indicate greater difficulty in reaching a final answer. Appendix D provides illustrative examples of these patterns. Notably, the authors of s1.1-32B show that “Wait” tokens can be leveraged at test time to improve prediction calibration (Muennighoff et al., 2025).

To operationalize these hypotheses, we extract three surface-level behavioral features from each model response: (1) the number of tokens in the reasoning trace, (2) the number of self-posed questions, and (3) the number of occurrences of the token “Wait”. We explore different strategies for turning these features into uncertainty estimates, detailed in Appendix D. The most successful method treats the three features as input to a Calibration Regression supervised model (which fits a linear function with L2 regularization as *Ridge Regression*). To train this model, we use ground-truth correctness labels (1 for correct, 0 for incorrect) for each response. Although the regression model is trained to predict correctness, we hypothesize that its outputs can serve as proxies for confidence. We define uncertainty as its complement,  $u = 1 - f(\mathbf{x})$ , enabling direct comparison with other scalar-based uncertainty estimation methods. To ensure fair comparison and avoid overfitting, we train the regression models on 40% of the dataset and evaluate on the remaining 60%, preserving the original distribution of medical specialties. We assess the quality of the behavioral uncertainty signals using AUROC, ECE, and Brier Score. To further disentangle model behavior, we report results separately for common and uncommon medical specialties (see Appendix D for the latter). We include results with Semantic Entropy for comparison, noting that it relies on repeated sampling.

The results, shown in Figure 4, reveal that our lightweight behavioral feature method outperforms sentence-level methods and closely matches Semantic Entropy, while requiring only a single forward pass per input. This pattern holds across both reasoning-oriented LLMs, DeepSeek-R1-70B and s1.1-32B, and applies to both calibration and dis-

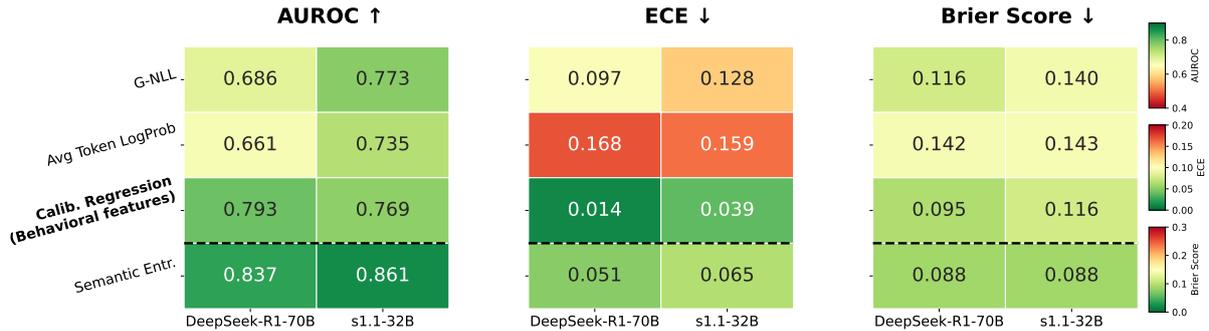


Figure 4: Our proposed method, based on behavioral features in reasoning-oriented LLMs (shown in bold), performs strongly across AUROC, ECE, and Brier Score, approaching Semantic Entropy, which instead relies on multiple generations via sampling.

Model	Overall	Cardio	Gastro	Inf. Dis.	Neuro	Ob/Gyn	Peds	Def.	Diag.	Dx Test	Treat.	Other
<b>Llama-3.1-70B</b>	85.5% (1.08)	83.4% (1.06)	83.4% (1.11)	83.3% (1.08)	85.4% (1.10)	89.7% (1.10)	87.6% (1.04)	80.0% (1.00)	85.2% (1.09)	86.8% (1.06)	86.8% (1.09)	82.6% (1.07)
<b>DeepSeek-R1-70B</b>	93.8% (1.18)	94.5% (1.21)	91.0% (1.14)	92.4% (1.16)	95.2% (1.21)	93.1% (1.16)	96.5% (1.21)	93.3% (1.13)	95.3% (1.19)	90.6% (1.09)	92.2% (1.17)	91.3% (1.24)
<b>Llama3-OpenBioLLM-70B</b>	90.1% (1.40)	92.4% (1.43)	89.0% (1.38)	86.2% (1.40)	91.0% (1.43)	89.7% (1.32)	92.4% (1.42)	86.7% (1.27)	89.7% (1.44)	94.3% (1.34)	91.1% (1.34)	87.0% (1.37)

Table 3: Conformal Prediction. Empirical coverage (top) and average set size (bottom, in parentheses) for each model. Coverage is colored green ( $\geq 90\%$ ) or red ( $< 90\%$ ). Set size transitions from blue to purple based on value.

crimination metrics. In Appendix D.3, we provide further details on the regression model coefficients. Although restricted to two models, these results support our hypothesis that behavioral cues in reasoning traces encode relevant information about model uncertainty and highlight a promising direction for efficient, lightweight, and interpretable uncertainty estimation without requiring sampling or token-level probabilities. This has important implications for real-world applications, where uncertainty needs to be estimated reliably and efficiently. Our case study offers a starting point for future work to validate this approach across broader range of tasks and models. We view our contribution as an initial step toward establishing a more general framework for LLM behavioral uncertainty estimation leveraging structured reasoning traces.

## 6 Conformal Prediction Results

Analyses based on score-based uncertainty estimates (Section 5) reveal notable trends: reasoning LLMs often achieve higher accuracy and better calibration, while biomedical models often appear overconfident (likely due to domain-specific tuning) and underperform outside a narrow range of specialties, leading to less reliable estimates.

We next examine conformal prediction (follow-

ing the method in Ke et al. 2025) to assess whether set-based uncertainty estimates offer complementary insights. We focus on three representative 70B LLMs (one per model type) and restrict the analysis to common medical specialties with sufficient held-out data. Table 3 reports empirical coverage and average set size across models, specialties, and question types. The three models achieve strong overall coverage, though not all of them reach the 90% coverage target. DeepSeek-R1-70B performs most reliably, achieving 93.8% overall coverage with compact sets, and generalizes well across specialties and question types. Llama3-OpenBioLLM-70B also reaches high coverage (90.1%) but with larger sets compared to DeepSeek-R1-70B, and it fails to meet the target in several categories, particularly Infectious Diseases and Definition questions.

In contrast, Llama-3.1-70B falls short of the coverage target (85.5%) with highly restrictive predictions, indicating overconfidence and poor generalization from calibration to test data. For this model, no specialty or question type meets the 90% threshold, suggesting that its strong score-based UQ performance does not extend to set-based methods. These findings underscore that accuracy alone does not suffice. Robust uncertainty estimation depends on careful calibration and generalization

across specialties, question types, and uncertainty paradigms, particularly in light of the misalignment observed between score- and set-based methods.

## 7 Discussion and Conclusion

We investigate uncertainty estimation for LLMs in clinical QA. By jointly assessing calibration and discrimination spanning diverse models and across clinical specialties and question types, we address key limitations of prior work. Our findings show that model performance varies not only in accuracy but also in uncertainty quality, with substantial differences across clinical contexts. Reasoning-oriented LLMs often perform well, but no method succeeds consistently across settings; evaluation and deployment must be context-specific.

In conclusion, our work raises a crucial cautionary note for the deployment of LLMs in clinical NLP, while also offering a constructive path forward. Our analysis demonstrates that a model’s ability to signal uncertainty is not a monolithic capability but one that varies significantly with the clinical context. To support more nuanced evaluation, we propose a lightweight and interpretable UQ case study based on behavioral cues from reasoning traces, which approximates the performance of more resource-intensive sampling-based methods. The primary lesson for NLP practitioners is that for high-stakes domains, ‘one-size-fits-all’ evaluations are insufficient and potentially misleading. The true measure of an LLM’s reliability lies in its performance on the specific tasks and in the precise domains where it will be used. Future work should involve clinicians in assessing questions flagged as uncertain, to determine which signals best indicate when additional human review is warranted.

### Limitations

In the following, we discuss some limitations of our experimental setup.

**Open-ended generation mapped to MCQA.** While we benchmark on multiple-choice question answering for comparability and ground truth clarity, our setup involves open-ended generation: models first produce free-form answers and reasoning traces, which we then map to the multiple-choice options for evaluation. We believe that fully open-ended QA is crucial for clinical realism, but it comes with significant challenges in defining gold standards. We believe this focus is justified as a first

step: to the best of our knowledge, ours is the first study to systematically evaluate UQ across model types, medical specialties, and question types at this scale. Nonetheless, future work should extend this analysis to more open-ended clinical scenarios.

**Automated prediction and annotation pipelines.** We rely on automated extraction of model predictions and question-type annotations. Although our heuristics and model-assisted annotations are validated against manual labels with high agreement, some residual noise may persist, especially for ambiguous questions or borderline categories (e.g., “Other”).

**Behavioral features as exploratory proxies.** We use behavioral features as proxies for uncertainty in an exploratory case study, as highlighted throughout the main text (Section 5.2). Validation with domain experts or clinicians is planned as an important next step to confirm their practical value. While we show that behavioral features offer a promising, interpretable alternative to multi-sample methods like Semantic Entropy, our feature space is deliberately simple and task-specific. Generalizing these findings to other tasks, languages, or reasoning paradigms will require broader validation and potentially more sophisticated feature design.

**Sources of uncertainty.** We do not investigate in detail sources of uncertainty. We believe future work should look into promising approaches in this direction, such as [Yang et al. \(2025\)](#).

**Post-hoc calibration.** We do not apply post-hoc calibration techniques such as temperature scaling. Our focus is on comparing raw uncertainty estimates without additional tuning. Applying temperature scaling consistently is challenging in our setup, as models differ in tokenization schemes and output structure – e.g., in how they format answers (direct label vs. more complex sentences), affecting how likelihoods are computed and making calibration non-trivial. We leave the integration of model-specific calibration as an important direction for future work. As for additional considerations on the temperature setup and decoding details, see [Appendix C](#).

**Uncertainty communication.** We benchmark uncertainty estimates (including confidence elicitation) but do not study how uncertainty should be communicated to end users ([Ulmer et al., 2025](#)); verbalized uncertainty is shaped by contextual and

sociopragmatic factors and can be biased by alignment for helpfulness, raising additional challenges for clinician-facing deployment.

**Coverage for ChatGPT models.** Our analysis of OpenAI ChatGPT models (see Appendix G) is limited to Semantic Entropy, which proved to be the most robust UQ method. Future work should evaluate other methods, such as Confidence Elicitation and Conformal Prediction, also for these models.

## Acknowledgments

This publication is part of the project CaRe-NLP with file number NGF.1607.22.014 of the research programme AiNed Fellowship Grants which is (partly) financed by the Dutch Research Council (NWO). We thank the members of the NLP4Health Lab and KIK AI at Amsterdam UMC, as well as the DMG group at the ILLC, University of Amsterdam, for their insightful feedback and helpful suggestions at different stages of this work.

## References

- Emre Can Acikgoz, Osman Batur İnce, Rayene Bench, Arda Anil Boz, İlker Kesen, Aykut Erdem, and Erkut Erdem. 2024. *Hippocrates: An open-source framework for advancing large language models in healthcare*. Preprint, arXiv:2404.16621.
- Lukas Aichberger, Kajetan Schweighofer, and Sepp Hochreiter. 2024. Rethinking uncertainty estimation in natural language generation. *arXiv preprint arXiv:2412.15176*.
- Malaikannan Sankarasubbu Ankit Pal. 2024. Openbiollms: Advancing open-source large language models for healthcare and life sciences. <https://huggingface.co/aaditya/OpenBioLLM-Llama3-70B>.
- Zahra Atf, Seyed Amir Ahmad Safavi-Naini, Peter R Lewis, Aref Mahjoubfar, Nariman Naderi, Thomas R Savage, and Ali Soroush. 2025. The challenge of uncertainty quantification of large language models in medicine. *arXiv preprint arXiv:2504.05278*.
- Joris Baan, Nico Daheim, Evgenia Iliia, Dennis Ulmer, Haau-Sing Li, Raquel Fernández, Barbara Plank, Rico Sennrich, Chrysoula Zerva, and Wilker Aziz. 2023. Uncertainty in natural language generation: From theory to applications. *arXiv preprint arXiv:2307.15703*.
- Glenn W. Brier. 1950. Verification of forecasts expressed in terms of probability. *Monthly weather review*, 78(1):1–3.
- Margarida Campos, António Farinhas, Chrysoula Zerva, Mário AT Figueiredo, and André FT Martins. 2024. Conformal prediction for natural language processing: A survey. *Transactions of the Association for Computational Linguistics*, 12:1497–1516.
- Zeming Chen, Alejandro Hernández-Cano, Angelika Romanou, Antoine Bonnet, Kyle Matoba, Francesco Salvi, Matteo Pagliardini, Simin Fan, Andreas Köpf, Amirkeivan Mohtashami, Alexandre Sallinen, Alireza Sakhaeirad, Vinitra Swamy, Igor Krawczuk, Deniz Bayazit, Axel Marmet, Syrielle Montariol, Mary-Anne Hartley, Martin Jaggi, and Antoine Bosselut. 2023. *Meditron-70b: Scaling medical pretraining for large language models*. Preprint, arXiv:2311.16079.
- ContactDoctor. 2024. Bio-medical: A high-performance biomedical language model. <https://huggingface.co/ContactDoctor/Bio-Medical-Llama-3-8B>.
- Tareen Dawood, Chen Chen, Baldeep S Sidhu, Bram Ruijsink, Justin Gould, Bradley Porter, Mark K Elliott, Vishal Mehta, Christopher A Rinaldi, Esther Puyol-Antón, and 1 others. 2023. Uncertainty aware training to improve deep learning model calibration for classification of cardiac mr images. *Medical Image Analysis*, 88:102861.
- Felix J Dorfner, Amin Dada, Felix Busch, Marcus R Makowski, Tianyu Han, Daniel Truhn, Jens Kleesiek, Madhumita Sushil, Jacqueline Lammert, Lisa C Adams, and 1 others. 2024. Biomedical large language models seem not to be superior to generalist models on unseen medical data. *arXiv preprint arXiv:2408.13833*.
- Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. 2024. Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017):625–630.
- Yanjun Gao, Skatje Myers, Shan Chen, Dmitriy Dligach, Timothy Miller, Danielle S Bitterman, Guanhua Chen, Anoop Mayampurath, Matthew M Churpek, and Majid Afshar. 2025. Uncertainty estimation in diagnosis generation from large language models: next-word probability is not pre-test probability. *JAMIA open*, 8(1):ooae154.
- Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koeppel, Preslav Nakov, and Iryna Gurevych. 2024. *A survey of confidence estimation and calibration in large language models*. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6577–6595, Mexico City, Mexico. Association for Computational Linguistics.
- Biraja Ghoshal, Allan Tucker, Bal Sanghera, and Wai Lup Wong. 2021. Estimating uncertainty in deep learning for reporting confidence to clinicians in medical image segmentation and diseases detection. *Computational Intelligence*, 37(2):701–734.

- Mario Giulianelli, Joris Baan, Wilker Aziz, Raquel Fernández, and Barbara Plank. 2023. [What comes next? evaluating uncertainty in neural text generators against human production variability](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14349–14371, Singapore. Association for Computational Linguistics.
- Tobias Groot and Matias Valdenegro Toro. 2024. [Overconfidence is key: Verbalized uncertainty evaluation in large language and vision-language models](#). In *Proceedings of the 4th Workshop on Trustworthy Natural Language Processing (TrustNLP 2024)*, pages 145–171, Mexico City, Mexico. Association for Computational Linguistics.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shitong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Wataru Hashimoto, Hidetaka Kamigaito, and Taro Watanabe. 2024. Are data augmentation methods in named entity recognition applicable for uncertainty estimation? In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 18852–18867.
- Yuheng Huang, Jiayang Song, Zhijie Wang, Shengming Zhao, Huaming Chen, Felix Juefei-Xu, and Lei Ma. 2025. Look before you leap: An exploratory study of uncertainty analysis for large language models. *IEEE Transactions on Software Engineering*.
- Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2021. What disease does this patient have? a large-scale open domain question answering dataset from medical exams. *Applied Sciences*, 11(14):6421.
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. [PubMedQA: A dataset for biomedical research question answering](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2567–2577, Hong Kong, China. Association for Computational Linguistics.
- Yusong Ke, Hongru Lin, Yuting Ruan, Junya Tang, and Li Li. 2025. Correctness coverage evaluation for medical multiple-choice question answering based on the enhanced conformal prediction framework. *Mathematics*, 13(9):1538.
- Yunsoo Kim, Jinge Wu, Yusuf Abdulle, and Honghan Wu. 2024. Medexqa: Medical question answering benchmark with multiple explanations. In *Proceedings of the 23rd Workshop on Biomedical Natural Language Processing*, pages 167–181.
- Benjamin Kompa, Jasper Snoek, and Andrew L Beam. 2021. Second opinion needed: communicating uncertainty in medical machine learning. *NPJ Digital Medicine*, 4(1):4.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. In *The Eleventh International Conference on Learning Representations*.
- Xiaou Liu, Tiejun Chen, Longchao Da, Chacha Chen, Zhen Lin, and Hua Wei. 2025. Uncertainty quantification and confidence calibration in large language models: A survey. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pages 6107–6117.
- Potsawee Manakul, Adian Liusie, and Mark Gales. 2023. [SelfCheckGPT: Zero-resource black-box hallucination detection for generative large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9004–9017, Singapore. Association for Computational Linguistics.
- Meta AI. 2024a. Llama 3 models. <https://ai.meta.com/blog/meta-llama-3>.
- Meta AI. 2024b. Meta llama 3.1 70b instruct. <https://huggingface.co/meta-llama/Llama-3.1-70B-Instruct>.
- Mistral AI. 2023. Mistral-7b and mixtral-8x7b models. <https://mistral.ai/news/mixtral-of-experts>.
- Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candes, and Tatsunori Hashimoto. 2025. [s1: Simple test-time scaling](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 20275–20321, Suzhou, China. Association for Computational Linguistics.
- Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. 2022. Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering. In *Conference on health, inference, and learning*, pages 248–260. PMLR.
- Mauricio Rivera, Jean-François Godbout, Reihaneh Rabbany, and Kellin Pelrine. 2024. [Combining confidence elicitation and sample-based methods for uncertainty quantification in misinformation mitigation](#). In *Proceedings of the 1st Workshop on Uncertainty-Aware NLP (UncertainNLP 2024)*, pages 114–126, St Julians, Malta. Association for Computational Linguistics.

- Thomas Savage, John Wang, Robert Gallo, Abdessalem Boukil, Vishwesh Patel, Seyed Amir Ahmad Safavi-Naini, Ali Soroush, and Jonathan H Chen. 2025. Large language model uncertainty proxies: discrimination and calibration for medical diagnosis and treatment. *Journal of the American Medical Informatics Association*, 32(1):139–149.
- Rosanne Schoonbeek, Jessica Workum, Stephanie CE Schuit, Job Doornberg, Tom P Van Der Laan, and Charlotte MHHT Bootsma-Robroeks. 2024. Completeness, correctness and conciseness of physician-written versus large language model generated patient summaries integrated in electronic health records. Available at SSRN 4835935.
- Andrew Sellergren, Sahar Kazemzadeh, Tiam Jaroensri, Atilla Kiraly, Madeleine Traverse, Timo Kohlberger, Shawn Xu, Fayaz Jamil, Cían Hughes, Charles Lau, Justin Chen, Fereshteh Mahvar, Liron Yatziv, Tiffany Chen, Bram Sterling, Stefanie Anna Baby, Susanna Maria Baby, Jeremy Lai, Samuel Schmidgall, and 62 others. 2025. [Medgemma technical report](#). Preprint, arXiv:2507.05201.
- Ola Shorinwa, Zhiting Mei, Justin Lidard, Allen Z Ren, and Anirudha Majumdar. 2025. A survey on uncertainty quantification of large language models: Taxonomy, open research challenges, and future directions. *ACM Computing Surveys*.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, and 1 others. 2023. Large language models encode clinical knowledge. *Nature*, 620(7972):172–180.
- Arun James Thirunavukarasu, Darren Shu Jeng Ting, Kabilan Elangovan, Laura Gutierrez, Ting Fang Tan, and Daniel Shu Wei Ting. 2023. Large language models in medicine. *Nature medicine*, 29(8):1930–1940.
- Dennis Ulmer, Martin Gubri, Hwaran Lee, Sangdoon Yun, and Seong Oh. 2024. [Calibrating large language models using their generations only](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15440–15459, Bangkok, Thailand. Association for Computational Linguistics.
- Dennis Ulmer, Alexandra Lorson, Ivan Titov, and Christian Hardmeier. 2025. Anthropomorphic uncertainty: What verbalized uncertainty in language models is missing. *arXiv preprint arXiv:2507.10587*.
- Dave Van Veen, Cara Van Uden, Louis Blanke-meier, Jean-Benoit Delbrouck, Asad Aali, Christian Bluethgen, Anuj Pareek, Malgorzata Polacin, Eduardo Pontes Reis, Anna Seehofnerová, and 1 others. 2024. Adapted large language models can outperform medical experts in clinical text summarization. *Nature medicine*, 30(4):1134–1142.
- Chaoyi Wu, Weixiong Lin, Xiaoman Zhang, Ya Zhang, Weidi Xie, and Yanfeng Wang. 2024a. PMC-LLaMA: toward building open-source language models for medicine. *J Am Med Inform Assoc*, 31(9):1833–1843.
- Jiixin Wu, Yizhou Yu, and Hong-Yu Zhou. 2024b. Uncertainty estimation of large language models in medical question answering. *arXiv preprint arXiv:2407.08662*.
- Yijun Xiao and William Yang Wang. 2019. Quantifying uncertainties in natural language processing tasks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 7322–7329.
- Miao Xiong, Zhiyuan Hu, Xinyang Lu, YIFEI LI, Jie Fu, Junxian He, and Bryan Hooi. 2024. Can llms express their uncertainty? an empirical evaluation of confidence elicitation in llms. In *The Twelfth International Conference on Learning Representations*.
- Xinlan Yan, Di Wu, Yibin Lei, Christof Monz, and Iacer Calixto. 2026. What does neuro mean to cardio? investigating the role of clinical specialty data in medical llms. In *The 19th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2026)*.
- Yongjin Yang, Haneul Yoo, and Hwaran Lee. 2025. Maqa: Evaluating uncertainty quantification in llms regarding data uncertainty. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 5846–5863.
- Fanghua Ye, Mingming Yang, Jianhui Pang, Longyue Wang, Derek Wong, Emine Yilmaz, Shuming Shi, and Zhaopeng Tu. 2024. Benchmarking llms via uncertainty quantification. *Advances in Neural Information Processing Systems*, 37:15356–15385.
- Chrysoula Zerva, Taisiya Glushkova, Ricardo Rei, and André FT Martins. 2022. Disentangling uncertainty in machine translation evaluation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8622–8641.
- Shuang Zhou, Jiashuo Wang, Zidu Xu, Song Wang, David Brauer, Lindsay Welton, Jacob Cogan, Yuen-Hei Chung, Lei Tian, Zaifu Zhan, and 1 others. 2025. Uncertainty-aware large language models for explainable disease diagnosis. *npj Digital Medicine*, 8(1):690.

## Appendix

### A Model Details and Licenses

The LLaMA family of models and LLaMA3-OpenBioLLM-70B are distributed under the Meta AI community license agreement. Mistral-7B-v0.2, Mixtral-8x7B-v0.1, s1.1-32B are released under Apache license 2.0. DeepSeek-R1-Distill-Llama-70B is distributed

under MIT license. Bio-Medical-Llama-3-8B is released under Bio-Medical-Llama-3-8B LLM License (Non-Commercial Use Only). All models are accessed through Hugging Face (<https://huggingface.co/>). We run inference on a cluster with 4 A100 GPUs and a batch size of 8. We run small and medium models (3B-32B parameters) on 2 GPUs and larger models on 4 GPUs. We use the vLLM library. Running inference takes on average a couple of hours for small models and 24 hours for larger models and reasoning models, which produce longer model responses. We used all datasets in accordance with their licenses and consented use policies.

## B Prompts

**Model Responses** In order to get the model responses to the input question, we employ the following prompt (example of an instantiation in parenthesis):

Question: (A 59-year-old overweight woman presents to the urgent care clinic with the complaint of severe abdominal pain for the past 2 hours. She also complains of a dull pain in her back with nausea and vomiting several times. Her pain has no relation with food. Her past medical history is significant for recurrent abdominal pain due to cholelithiasis. Her father died at the age of 60 with some form of abdominal cancer. Her temperature is 37°C (98.6°F), respirations are 15/min, pulse is 67/min, and blood pressure is 122/98 mm Hg. Physical exam is unremarkable. However, a CT scan of the abdomen shows a calcified mass near her gallbladder. Which of the following diagnoses should be excluded first in this patient?)

Provide a concise answer and start your response with the letter of the selected option in square brackets. Options:

(  
[A] Acute cholecystitis  
[B] Gallbladder cancer  
[C] Choledocholithiasis  
[D] Pancreatitis"  
)

Your Answer:

For models that come with a recommended system prompt, we use it as provided by the model developers/authors.

**Mapping Responses to the Corresponding Option** We use the following prompt to map model responses that do not follow the required format (i.e., the letter of the selected option in square brackets) to the corresponding answer choice. Before

running the annotation, we inspected a random sample of 100 ill-formatted responses and found no need to include an “other”/“invalid” class for this task. Placeholders in curly brackets.

You will be given a multiple-choice question, four answer options labeled A, B, C, and D, and a free-form response.

Question: question

Options:

A. {option a}

B. {option b}

C. {option c}

D. {option d}

Response: {response}

Your task is to determine which of the four options the response refers to. Output only the corresponding letter in square brackets, like [X]. Do not include any explanation or additional text.

Your annotation must be exactly one of the following: [A], [B], [C], or [D].

If the response includes a different letter, you must still map it to one of the four valid options above, based on the content of the response.

Keep in mind that the response may be poorly formatted or contain irrelevant letters—focus only on identifying the most likely intended option.

If multiple options seem plausible, choose the one that is most strongly implied by the response alone, without relying on external knowledge or context.

Your annotation:

**Question Type Annotation** We used the following prompt to assign each question to a specific question type (placeholders in curly brackets):

You are a medical expert assistant. Your task is to classify the type of each medical question according to the following categories:

1. Diagnosis – Identifying a disease or condition.
2. Treatment – Interventions like medication, surgery, or therapy.
3. Diagnostic Test – Lab tests, imaging, or diagnostic procedures.
4. Definition – Asking for an explanation of a medical concept.
5. Procedure/Operation – Questions about medical or surgical procedures.
6. Other

Assign the most appropriate type based on the question intent.

Classify the following question. Provide only the corresponding label number (1–6) in square brackets as your response.

Q: { }

Question type:

Model	Accuracy	AUROC	ECE	Brier Score
Llama-3.1-8B-Instruct	0.583	0.806	0.048	0.176
Llama-3.1-70B-Instruct	0.784	0.873	0.040	0.104
Llama-3.2-3B-Instruct	0.422	0.868	0.101	0.153
Llama-3.3-70B-Instruct	0.774	0.858	0.036	0.112
Mistral-7B-Instruct-v0.2	0.474	0.665	0.078	0.293
Mixtral-8x7B-Instruct-v0.1	0.602	0.708	0.040	0.225
DeepSeek-R1-70B	0.767	0.898	0.119	0.113
s1.1-32B	0.837	0.852	0.081	0.090
Llama3-OpenBioLLM-70B	0.699	0.856	0.093	0.140
Bio-Medical-Llama-3-8B	0.666	0.781	0.030	0.169

Table 4: Evaluation results with  $T = 1$  (uncertainty estimates extracted as Semantic Entropy).

## C Decoding Details

For experiments requiring multiple samples, we reviewed commonly suggested decoding configurations for nucleus sampling to approximate how language models are typically used in practice. Based on this review, we set the top- $p$  value to 0.9 and the temperature to 0.6. As an exploratory analysis, we also repeated our experiments using a higher temperature of 1.0 (Table 4). We found that this change had little effect on the overall ranking of models in terms of accuracy or on the calibration and discrimination performance of uncertainty estimates derived from Semantic Entropy. However, for DeepSeek-R1-70B, we observed a decline in both accuracy and calibration. This effect warrants further investigation, which would require a held-out validation set and a systematic analysis across a broader range of temperature values. Such an analysis might ultimately support model-specific or even task-specific decoding configurations, varying across models, medical specialties, or question types. In this work, however, we adhere to decoding settings—temperature and top- $p$ —that reflect typical real-world usage, especially for users with limited technical knowledge. As for annotating question types and mapping the responses to the corresponding option, we employ greedy decoding.

## D Behavioral Features in Reasoning Traces

In this section, we provide additional details on the experiments using behavioral features for reasoning models (Section 5.2), including regression-based results. In earlier stages of the project, we also tested alternative approaches such as Random Forests, a simple MLP, and sentence embeddings from SentenceTransformer, all of which performed worse than the models reported here and were therefore discarded. We further experimented

with subsets of the three behavioral features, but observed consistent performance drops.

### D.1 Additional details on the method

Given the features described in Section 5.2, the first is a simple normalized sum of the three features, referred to as the *Linear Sum* method. Let  $x_1$ ,  $x_2$ , and  $x_3$  denote the normalized values of the three features, and compute the uncertainty estimate as  $u = x_1 + x_2 + x_3$ . The second approach treats the three features as input to a supervised model: either a *Logistic Regression* or a *Calibration Regression* (the former models a sigmoid likelihood, the latter fits a linear function with L2 regularization as *Ridge Regression*). *Calibration Regression* is the method discussed in Section 5.2, while here we provide details on all methods. To train these models, we use ground-truth correctness labels (1 for correct, 0 for incorrect) associated with each response. Formally, the model learns a function  $f : \mathbb{R}^3 \rightarrow [0, 1]$ . Although the regression models are trained to predict correctness, we hypothesize that their outputs can serve as proxies for confidence. We define uncertainty as their complement,  $u = 1 - f(\mathbf{x})$ , enabling direct comparison with other uncertainty estimation methods. To ensure a fair comparison and avoid overfitting, we train the regression models on 40% of the dataset while evaluating on the remaining 60%, preserving the original distribution of medical specialties. All responses are generated via greedy decoding to enable direct comparison with single-pass baselines such as G-NLL and average token probability. We assess the quality of the behavioral uncertainty signals using AUROC, ECE, and Brier Score. To further disentangle model behavior, we report results separately for common and uncommon medical specialties. We include results with Semantic Entropy for comparison, noting that it relies on repeated sampling to estimate uncertainty.

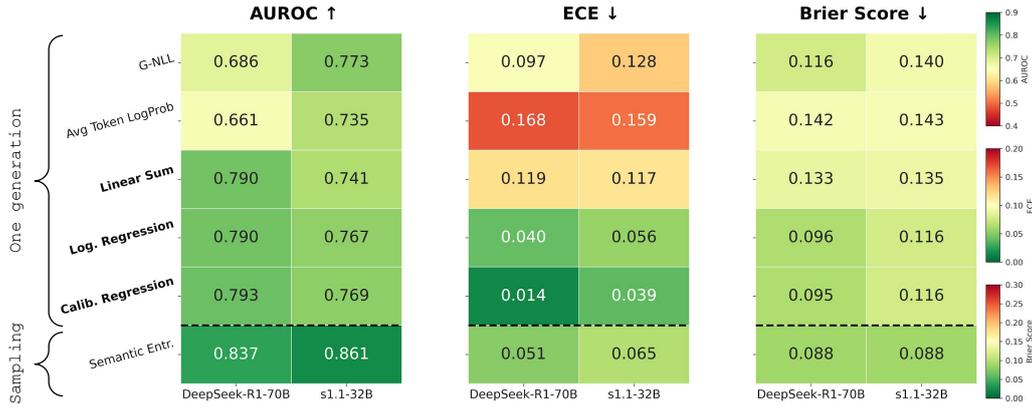


Figure 5: Our proposed methods, based on behavioral features (shown in bold), offer strong performance across AUROC, ECE, and Brier Score, approaching that of Semantic Entropy, which instead relies on multiple generations via sampling.

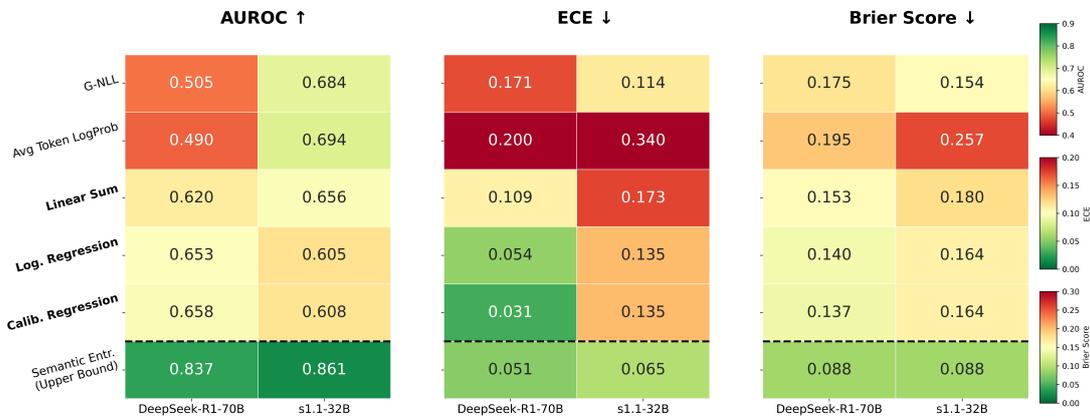


Figure 6: Comparison of UQ methods for DeepSeek-R1-70B and s1.1-32B for uncommon medical specialties.

The results, reported in Figure 5, reveal that our lightweight behavioral feature-based methods—*Linear Sum*, *Logistic Regression*, and *Calibration Regression*—outperform other sentence-level methods and closely match the performance of Semantic Entropy, while requiring only a single forward pass per input. This pattern holds for both reasoning LLMs, DeepSeek-R1-70B and s1.1-32B. Even the simple *Linear Sum* method outperforms token-level baselines like G-NLL and average token log-probability in most cases, particularly in discrimination (AUROC). *Calibration Regression* performs consistently well, especially in terms of calibration. In Section D.3, we provide further details on the regression model coefficients. Results on less-represented specialties are reported in Figure 6. They confirm the effectiveness of our approach for DeepSeek-R1-70B, while showing limited improvement for s1.1-32B, highlighting the need for model-specific tuning or more expressive behavioral features in lower-resource

settings. Our case study offers a starting point for future work that should validate this approach across a broader range of tasks, datasets, and models. We view our contribution as an initial step toward establishing a more general framework for behavioral uncertainty estimation that leverages the output of recently released reasoning LLMs.

## D.2 Example of Model Response

In the sample model response below, generated by DeepSeek-R1-70B, we highlight in bold some of the behavioural features examined in our work, including self-verification questions and “wait” tokens (we highlight relevant sentences below).

**Question:** A 28-year-old male presents to the emergency department with chest pain. He reports that one hour ago he was climbing the stairs in his home when he suddenly developed excruciating chest pain. He describes the pain as ripping and radiating to the back. The patient’s past medical history is significant for a heart murmur that is being followed by his outpatient cardiologist. He takes no medications. The patient works as

a carpenter and reports being monogamous with his wife. He denies any history of intravenous drug use or recent travel. In the emergency department, the patient’s temperature is 99.4°F (37.4°C), blood pressure is 94/56 mmHg, pulse is 121/min, and respirations are 14/min. On physical exam, the patient appears to be in moderate distress and is clutching his chest. He is tall and has joint hypermobility. His breath sounds are equal and full bilaterally and has weak distal pulses.

Which of the following is most likely to be found in this patient?

[A] Late systolic crescendo murmur at the apex with mid-systolic click

[B] Mid-diastolic murmur at the left lower sternal border

[C] Systolic ejection murmur at the right second intercostal space

[D] Holosystolic, high-pitched blowing murmur at the lower left sternal border

**Model Response:**

[C]

<think>

Okay, so I’m trying to figure out this medical question. Let me read through it again and break it down step by step.

The patient is a 28-year-old male presenting with chest pain. He was climbing stairs when he suddenly developed excruciating, ripping chest pain that radiates to the back. [...]

**Wait, but aortic regurgitation typically presents with a diastolic murmur, not systolic. So maybe I’m mixing things up.** Alternatively, if the patient has a bicuspid aortic valve, which is common in Marfan, that could cause a systolic ejection murmur. Also, in aortic dissection, the murmur might not be present unless there’s associated aortic regurgitation.

D. Holosystolic, high-pitched blowing murmur at the lower left sternal border: This is typical of mitral regurgitation. The holosystolic nature suggests blood flowing back through the mitral valve throughout systole. But again, the patient’s presentation points more towards aortic issues.

**Wait, but the patient’s chest pain and physical exam suggest aortic dissection. So, what murmur is associated with that?** If the dissection leads to aortic regurgitation, the murmur would be a diastolic one, like an early diastolic murmur best heard at the left sternal border. But none of the options are diastolic.

Alternatively, if the patient has a bicuspid aortic valve, which is a common congenital heart defect, that could present with a systolic ejection murmur at the right second intercostal space. Bicuspid aortic valve is associated with an increased risk of aortic dissection, especially in the context of Marfan syndrome or other connective tissue disorders. [...]

**Wait, but earlier I thought that aortic regurgitation would present with a diastolic murmur. So, if the dissection caused aortic regurgitation, the murmur would be diastolic, but that’s not**

Feature	Logistic Regr.		Calibration Regr.	
	DeepSeek	s1.1	DeepSeek	s1.1
tokens	-0.6476	-1.4611	-0.0925	-0.1527
questions	-0.1052	0.0755	-0.0196	-0.0181
waits	-0.3456	0.3961	-0.0839	0.0708
bias	1.8559	1.9407	0.8269	0.8556

Table 5: Regression Coefficients.

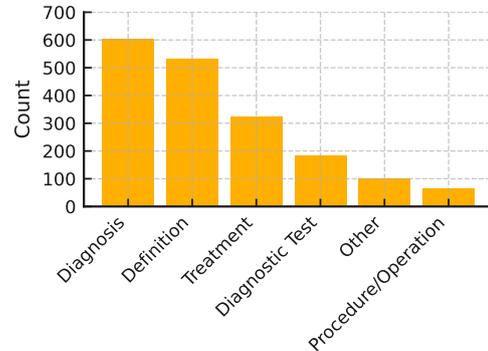


Figure 7: Frequency of question types.

**one of the options. So perhaps the murmur is from the underlying bicuspid valve, which is a systolic ejection murmur.**

Yes, that makes sense. The patient’s heart murmur is likely from the bicuspid aortic valve, which is a systolic ejection murmur heard at the right second intercostal space. So the answer would be option C. </think>

The patient’s presentation of sudden, severe, tearing chest pain radiating to the back, along with tall stature, joint hypermobility, and weak distal pulses, strongly suggests aortic dissection. This condition is often associated with underlying bicuspid aortic valve, which typically presents with a systolic ejection murmur at the right second intercostal space.

[C] Systolic ejection murmur at the right second intercostal space

**D.3 Regression Coefficients**

Table 5 reports the regression coefficients for the methods based on behavioural features of reasoning models. For both models, a higher token count is associated with lower accuracy and worse calibration, suggesting that verbosity reflects uncertainty or confusion. This effect is more pronounced in s1.1-32B, indicating it may overgenerate even when less confident.

**E Additional Details on Uncertainty Quantification Methods**

First, we chose not to include confidence elicitation methods that directly prompt LLMs to report a confidence score. While such approaches have been

explored in prior work, their reliability remains limited. In particular, [Savage et al. \(2025\)](#) report poor performance and calibration when relying on model-generated confidence scores.

Second, for methods based on log-probabilities (logprobs), we carefully considered how to handle outputs from models that include explicit reasoning traces. A notable example is DeepSeek-R1-70B, whose responses are structured with clearly delimited reasoning segments, marked by special tokens such as <think>. In these cases, a key decision is whether to compute uncertainty scores over the entire response (including the reasoning trace) or only over the final answer portion. After empirical inspection and preliminary analyses, we found that restricting the computation of logprobs to the final answer (excluding the reasoning trace) produced more reliable and comparable uncertainty estimates.

## F Question Type Distribution

Figure 7 reports the frequency of the question types analyzed in our experiments.

## G GPT Results with Semantic Entropy

Tables 6, 7, and 8 report the performance of three representative OpenAI models (GPT-4.1, GPT-4.1-mini, and GPT-4o-mini) using Semantic Entropy to extract uncertainty estimates across common and uncommon medical specialties. These models show strong accuracy and calibration (average ECE and Brier score) but surprisingly poor discrimination (1-AUROC), sometimes even below smaller open-source models (see Figure 2). Importantly, we once again observe differences across specialties, reinforcing our conclusions. We observe an effect of model size, but less pronounced than expected. GPT-4.1 is generally more accurate and better calibrated than GPT-4.1-mini, its small counterpart. However, this does not hold for all specialties: for instance, for questions in the *Clinical Laboratory Science* domain, the uncertainty estimates extracted from GPT-4.1-mini are better calibrated. In line with the results in Section 5, GPT models show consistently weak performance in the *Speech Language Pathology* specialty. Taken together, these results suggest that the proprietary models evaluated achieve strong accuracy and calibration, yet fall short in producing uncertainty estimates with sufficient discriminative power. Furthermore, the unavailability of log prob-

abilities and reasoning traces limits transparency and constrains deeper analysis.

Specialty	Accuracy ↑	1-AUROC ↓	Avg(ECE,Brier) ↓
Cardiology	0.8945	0.1005	0.0232
Gastroenterology	0.9007	0.2528	0.0223
Infectious diseases	0.8890	0.2430	0.0230
Neurology	0.8752	0.2466	0.0256
Obs. and gyn.	0.8986	0.2492	0.0401
Pediatrics	0.8986	0.1559	0.0290
Biomedical Eng.	0.8853	0.1946	0.0292
Clinical Lab. Science	0.9320	0.2931	0.0135
Clinical Psychology	0.9387	0.0453	0.0189
Occupational Therapy	0.8942	0.1176	0.0176
Speech-Lang. Path.	0.3123	0.5021	0.0159

Table 6: GPT-4.1 results.

Specialty	Accuracy ↑	1-AUROC ↓	Avg(ECE,Brier) ↓
Cardiology	0.8593	0.2784	0.0434
Gastroenterology	0.8290	0.3944	0.0283
Infectious diseases	0.8310	0.2688	0.0358
Neurology	0.8276	0.2802	0.0361
Obs. and gyn.	0.8379	0.2885	0.0252
Pediatrics	0.8572	0.1477	0.0396
Biomedical Eng.	0.8706	0.2255	0.0256
Clinical Lab. Science	0.9153	0.1851	0.0184
Clinical Psychology	0.9377	0.1446	0.0145
Occupational Therapy	0.8974	0.2672	0.0199
Speech-Lang. Path.	0.3100	0.4567	0.0133

Table 7: GPT-4.1-mini results.

Specialty	Accuracy ↑	1-AUROC ↓	Avg(ECE,Brier) ↓
Cardiology	0.7800	0.2923	0.0477
Gastroenterology	0.7552	0.2840	0.0406
Infectious diseases	0.7352	0.2900	0.0537
Neurology	0.7131	0.3006	0.0612
Obs. and gyn.	0.7324	0.3463	0.0448
Pediatrics	0.7807	0.1904	0.0465
Biomedical Eng.	0.8531	0.2442	0.0251
Clinical Lab. Science	0.8680	0.2147	0.0279
Clinical Psychology	0.9151	0.1071	0.0246
Occupational Therapy	0.8852	0.2529	0.0265
Speech-Lang. Path.	0.3085	0.4826	0.0223

Table 8: GPT-4o-mini results.

## H Conformal Prediction for Multiple-Choice QA

### H.1 Method

We adopt a conformal prediction framework for multiple-choice question answering (MCQA), following the methodology proposed by [Ke et al. \(2025\)](#) and summarized below.

#### Step 1: Calibration

**1A. Sample Generations.** For each calibration question with known correct answer, sample  $M$  model outputs and extract the selected answer option.

*Example:*

**Question:** Which organ produces insulin?

**Correct answer:** B

**Sampled answers (M=20):** B, B, C, B, A, B, B, B, B, C, B, B, B, B, B, B, B, B, B, B

**Empirical frequencies:** A: 1, B: 17, C: 2, D: 0

$\Rightarrow \hat{p}_{\text{correct}} = \frac{17}{20} = 0.85$ , so  $S = 1 - 0.85 = 0.15$

**1B. Compute Non-conformity Scores.** Repeat the above process for all calibration questions to compute a set of scores  $S_1, S_2, \dots, S_N$ .

**1C. Estimate the Threshold.** Sort the  $S$  values and set  $\hat{q}$  as the  $(1 - \alpha)$  quantile.

*Example:*

Calibration scores: [0.05, 0.10, 0.15, 0.20, ..., 0.85]

For  $\alpha = 0.1$ , 90th percentile  $\Rightarrow \hat{q} = 0.65$

In our setup, we extract 320 questions per specialty from the S-MedQA validation set to serve as calibration set.

## Step 2: Prediction

**2A. Sample Generations.** For a test-time question, sample  $M$  model outputs and extract the selected options.

*Example:*

**Question:** What is the most common cause of X?

**Sampled answers (M=20):** B, B, B, C, B, B, B, A, B, **Frequencies:** A: 1, B: 18, C: 1, D: 0

**2B. Compute Option-wise Scores.** For each option  $o$ , compute:

$$S(o) = 1 - \hat{p}(o)$$

**Scores:**  $S(A) = 0.95$ ,  $S(B) = 0.10$ ,  $S(C) = 0.95$ ,  $S(D) = 1.00$

**2C. Construct Prediction Set.** Include all options  $o$  such that  $S(o) \leq \hat{q}$ .

*Example:* With  $\hat{q} = 0.65$ :

$$\text{Prediction set} = \{B\}$$

Options A, C, and D are excluded as their non-conformity scores exceed the threshold.

## H.2 Full Results

In Tables 9, 10, 11, we report the full conformal prediction results on the test set for the three representative 70B models, broken down by medical specialty and question type. Each table reports four metrics: **Empirical Coverage:** The proportion of test instances where the correct answer is included in the prediction set, averaged over questions in the group. This is expected to be close to the target coverage (e.g., 90%); **Average Set Size:** The average number of answer options included in the prediction set. Lower values indicate more informative predictions; **Majority Voting Accuracy:** The accuracy of the most frequently selected answer across  $M$  sampled generations (i.e., argmax over vote counts). This serves as a reference for model performance under deterministic decoding; **Hit Rate:** The proportion of test instances where at least one of the  $M$  sampled generations matches the correct answer. This serves as a sort of upper-bound for empirical coverage.

These metrics are reported separately for each medical specialty and question type to highlight performance heterogeneity across clinical subdomains and reasoning tasks.

<b>Group</b>	<b>Empirical Coverage</b>	<b>Avg Set Size</b>	<b>Maj. Voting Accuracy</b>	<b>Hit Rate</b>
<i>Overall</i>	85.50%	1.08	82.72%	95.39%
Cardiology	83.40%	1.06	80.69%	95.17%
Gastroenterology	83.40%	1.11	82.76%	91.72%
Infectious Diseases	83.30%	1.08	78.47%	95.83%
Neurology	85.40%	1.10	80.56%	95.14%
Obstetrics/Gynecology	89.70%	1.10	86.90%	95.17%
Pediatrics	87.60%	1.04	86.90%	99.31%
Definition	80.00%	1.00	80.00%	100.00%
Diagnosis	85.16%	1.09	82.32%	96.14%
Diagnostic Test	86.79%	1.06	83.02%	92.45%
Other	82.61%	1.07	76.09%	97.83%
Treatment	86.77%	1.09	84.83%	93.77%

Table 9: Conformal prediction results for LLaMA-3.1-70B-instruct. Target coverage: 90% ( $\hat{q} = 0.650$ ).

<b>Group</b>	<b>Empirical Coverage</b>	<b>Avg Set Size</b>	<b>Maj. Voting Accuracy</b>	<b>Hit Rate</b>
<i>Overall</i>	93.79%	1.18	89.65%	98.28%
Cardiology	94.48%	1.21	91.03%	97.93%
Gastroenterology	91.03%	1.14	86.90%	96.55%
Infectious Diseases	92.41%	1.16	88.97%	99.31%
Neurology	95.17%	1.21	91.03%	98.62%
Obstetrics/Gynecology	93.10%	1.16	88.28%	99.31%
Pediatrics	96.55%	1.21	91.72%	97.93%
Definition	93.33%	1.13	86.67%	100.00%
Diagnosis	95.34%	1.19	90.87%	98.38%
Diagnostic Test	90.57%	1.09	86.79%	98.11%
Other	91.30%	1.24	89.13%	100.00%
Treatment	92.25%	1.17	88.37%	97.67%

Table 10: Conformal prediction results for DeepSeek-R1-70B. Target coverage: 90% ( $\hat{q} = 0.750$ ).

<b>Group</b>	<b>Empirical Coverage</b>	<b>Avg Set Size</b>	<b>Maj. Voting Accuracy</b>	<b>Hit Rate</b>
<i>Overall</i>	90.12%	1.40	79.89%	96.21%
Cardiology	92.41%	1.43	82.07%	97.24%
Gastroenterology	88.97%	1.38	77.24%	96.55%
Infectious Diseases	86.21%	1.40	75.17%	92.41%
Neurology	91.03%	1.43	79.31%	95.86%
Obstetrics/Gynecology	89.66%	1.32	83.45%	97.24%
Pediatrics	92.41%	1.42	82.07%	97.93%
Definition	86.67%	1.27	80.00%	100.00%
Diagnosis	89.66%	1.44	77.28%	96.55%
Diagnostic Test	94.34%	1.34	84.91%	98.11%
Other	86.96%	1.37	82.61%	93.48%
Treatment	91.09%	1.34	83.33%	95.74%

Table 11: Conformal prediction results for Llama3-OpenBioLLM-70B. Target coverage: 90% ( $\hat{q} = 0.800$ ).