How Uncertain Is the Grade? A Benchmark of Uncertainty Quantification for LLM-Based Automatic Assessment

Anonymous ACL submission

Abstract

Large language models (LLMs) are rapidly transforming automatic assessment in education by enabling flexible grading across diverse question types and response formats. However, their probabilistic nature introduces output uncertainty, which can undermine the reliability of downstream pedagogical decisions such as feedback generation and instructional support. To better understand this challenge, we benchmark a broad range of uncertainty quantification (UQ) methods for LLM-based automatic assessment. Although these methods have demonstrated effectiveness in domains such as question answering, their reliability in educational grading remains underexplored, particularly due to the ordinal nature of grading scores. To address this challenge, we propose an evaluation framework that considers not only grading correctness but also the magnitude of deviation from ground-truth scores. Through extensive analyses across multiple assessment datasets, LLM families, and decoding settings, we characterize uncertainty behaviors in grading scenarios and investigate how factors such as model family, task type, and decoding strategy influence uncertainty estimation. Our findings highlight both the strengths and limitations of existing UQ methods and provide practical insights for developing more reliable and uncertainty-aware grading systems. Code is available at `https://anonymous.4open.science/r/UQBench-1FF4/`.

1 Introduction

With the surge of Large language models (LLMs), LLM-powered automatic assessment has become increasingly popular in many domains (Gu et al., 2024; Li et al., 2024a). Automated assessment, also called automated grading, is a system that takes questions and candidate solutions, optionally with rubrics and other context information, and generates text to present the correctness of the evaluated solution (Ala-Mutka, 2005; Emirtekin, 2025). Its inputs include questions, answers, rubrics, and optionally demonstrations of solution-grade examples, while its outputs may include grades, aspect-level assessment (Low et al., 2019; Song et al., 2024a), and reasons for assessments (Cohn et al., 2024; Li et al., 2024b). Existing automated assessments are supported by rule-based (Braun et al., 2006; Cline et al., 2010) or supervised machine learning models (Ahmed et al., 2022; Lu and Cutumisu, 2021), which require intensive labor from domain experts for rule and feature design, dataset annotation, validation, and handling anomalies or out-of-distribution cases (Ala-Mutka, 2005; Paiva et al., 2022). Owing to LLMs' broad prior knowledge and instruction-following ability, LLM-based grading can handle diverse question domains, understand rubrics, and produce structured evaluations with interpretable rationales in natural language (Li et al., 2025a; Kortemeyer, 2023; Hashemi et al., 2024). Compared with traditional task-specific grading systems, LLM-based systems offer a unified pipeline capable of handling diverse inputs without training or adaptation, enabling their usage across domains and problem formulations in practical educational workflows.

Although automated assessment is efficient, the uncertainty issue has been a central challenge, and this is even more severe in LLM-based assessment. Uncertainty is the metric denoting how sure the models are when making judgments, with the causes of randomness, absence of precise boundaries of concepts, inaccurate base knowledge, and ambiguity between different objects (Mena et al., 2021). It leads to decreased performance and misleading outputs, especially when information or model abilities are inadequate. This issue becomes more severe in LLM-based frameworks because LLMs hide the internal mechanisms in large-scale parameters, rely on token probability and sampling, and tend to produce seemingly correct responses even without enough information and capabilities.

In conclusion, certainty is highly prioritized by automated assessment tasks, and responses of low certainty are harmful to both the learning procedures and the equality between students.

Uncertainty quantification (UQ) has been studied in various tasks, including classification (Kurz et al., 2022; Chlaily et al., 2023), regression (Levi et al., 2022; Abadie et al., 2020), and Bayesian inference settings (Mena et al., 2021), covering data in text and image modalities. However, many existing methods rely heavily on access to model structure and hidden states, including prediction probabilities represented by neural networks (Sensoy et al., 2018) and the feasible space of possibility distributions (Chlaily et al., 2023). These methods are inapplicable for LLM-based auto-assessment systems as the internal states of LLMs are computationally expensive, hard to handle, or inaccessible when using proprietary LLMs. Existing UQ methods for LLMs mainly include white-box methods that require logits, internal layer outputs, and model parameters (Vashurin et al., 2025); consistency-based methods that compare multiple generations (Zhang et al., 2024); and reflection-based methods that elicit LLMs to self-report or analyze certainty (Tian et al., 2023). Because their model access and settings differ, these works are not systematically compared, and few insights exist into how well they perform. As a prosperous research field, LLM-based auto-assessment has developed many works on modeling and implementing uncertainty, while a comprehensive investigation to validate and compare them is still ignored.

To bridge this gap, in this work, we propose a benchmark on uncertainty quantification for the auto-assessment tasks. The challenges of applying UQ methods to auto-assessment derive from the specificity in the wide range of domain knowledge, the finite set of ordinal output features, and the influences of input prompts. First, auto-assessment deploys general-use LLMs as a judge in a specialized domain, requiring benchmarking work to consider heterogeneous subjects rather than domains of LLMs' expertise only. Second, the expected output, i.e., the grade, is an ordinal variable whose labels have a relative order, but the distances between categories are not comparable. This prevents arithmetic operations on the grade labels and makes numeric metrics such as differences and mean squared error invalid. Similarly, metrics designed for semantic distance may not adequately capture output differences because both correct and wrong gradings are centered on the semantic representations of question contexts. Third, LLMs are sensitive to the input text, making many UQ methods incomparable when they use different templates along with the question and rubrics.

Taking these points in mind, we design the benchmark with multiple learning subjects, different forms of output grades, and unified templates to build the input prompts from questions and rubrics. We collect datasets that are both open-sourced and private, covering essay writing, natural science, chemistry, and math education. For the output grades, the grade sets include binary, ordered, multi-class, etc. In addition, we elicit the LLMs to judge with the text of reasons for the grades, incorporating the semantic features into the analysis. Finally, we use a set of curated-designed templates to build the prompts, forcing the input text similar to prior works or other datasets in this benchmark. To satisfy the need for broad usage and flexibility, we focus on repetition-based (or ensemble-based) methods in this work, which take several runs of generation to estimate uncertainty, without using internal states of LLMs in any way.

2 Related Work

2.1 LLM-based Automatic Assessment

Automatic assessment has long been studied before the emergence of LLMs. Earlier systems were typically built on task-specific machine learning models trained for particular grading formats, such as automated essay scoring, short-answer evaluation, or rubric-based classification (Ala-Mutka, 2005; Paiva et al., 2022; Higgins et al., 2005). These approaches generally relied on engineered linguistic features and supervised neural architectures for fixed domains or tasks. While effective in narrow settings, they require annotated datasets for each task, making it costly to scale assessment systems to new contexts. Besides, rigorous prompt tuning is needed to align the assessment with fine-grained evaluation criteria, limiting transferability across subjects, question types, and scoring criteria.

The advent of large language models has shifted this paradigm toward general-purpose evaluators that can assess with instructions and rubrics. A representative framework is LLM-as-a-Judge (Gu et al., 2024), where a prompted model assigns correctness labels (Zhu et al., 2023b; Liu et al., 2023; Gao et al., 2023) or reviews written works (Brake and Schaaf, 2024; Hosseini et al., 2024). The LLM-

as-a-Judge framework has been used for diverse tasks, including mathematics (Luo et al., 2023), natural and medical science (Zhao et al., 2024; Krolik et al., 2024), programming (Kumar et al., 2024), and language (Song et al., 2024b). Recent studies show that LLM evaluators can produce stable grading judgments across prompts and reasoning settings (Raina et al., 2024); rubric-guided prompting supports fine-grained scoring aligned with evaluation criteria (Deng et al., 2025); and aggregation frameworks such as SURE (Korthals et al., 2025) use self-consistency and majority voting to improve grading reliability. Along with prompt engineering, studies also refine rubrics to achieve fine-grained grades (Singh et al.; Chu et al., 2024). However, these studies reveal limitations: evaluation outcomes may change under minor input perturbations, including synonyms and meaning-invariant rephrasing (Zhu et al., 2023a; Raina et al., 2024; He et al., 2024a), symbols indicating grade labels (Ye et al., 2024), and the order of provided grade labels (Li et al., 2025b; Shi et al., 2024). These findings suggest that although LLM-based assessment improves coverage and flexibility, existing methods remain vulnerable to undesired variation in generated judgments, highlighting the need for uncertainty estimation to complement assessment decisions and mitigate unreliable outputs.

2.2 Uncertainty Estimation in LLMs

With the surge of LLM applications, inaccurate or non-factual generations have motivated uncertainty estimation as a practical signal of when outputs should or should not be trusted. A common starting point is to derive confidence from observable generation behavior, especially via repeated sampling. (Cole et al., 2023) shows that repetition-based metrics from multiple samples provide a reliable criterion for abstention decisions against uncertain responses, outperforming likelihood-based and self-verification alternatives. (Lyu et al., 2024) formalizes sample consistency into confidence measures based on agreement or entropy, demonstrating effectiveness for confidence calibration across models and reasoning tasks, while also exploring the influence of sampling size and explanation prompting. In addition, (Xiong et al., 2023) proposes self-probing, which elicits verbalized confidence from LLMs and helps with calibration and failure prediction, while also showing overconfident statements like humans. Complementing these studies, (Huang et al., 2023) conducts a large-scale empirical comparison across uncertainty estimators, LLMs, and domain tasks, showing that uncertainty signals can surface risky generations while exposing practical limitations. Beyond surface features, more sophisticated approaches refine the notion of equivalence used to aggregate samples. For open-ended generation, (Kuhn et al., 2023) introduces semantic entropy, which clusters responses by meaning invariance and quantifies uncertainty over semantic modes rather than token sequences. Recently, (Xia et al., 2025) synthesizes inference-time uncertainty estimation for LLMs by defining key uncertainty sources and organizing methods into a taxonomy including verbalized confidence, latent-information, consistency-based, and semantic clustering approaches.

In addition, the ordinal characteristic of grade labels makes the assessment task different from other tasks. Grades are usually presented as a label or value from a finite set; for example, in SemEval 2013 - Task 7 (Dzikovska et al., 2013), the grade labels indicate whether and to what degree a student answer aligns with the given reference: `3` for fully correct, `2` for partially correct, `1` for contradictory with reference, and `0` for irrelevant. Although the labels are ordered, with higher scores indicating a stronger tie with the reference, the distance between labels is incomparable or even nonsensical (Tellamekala et al., 2023; Xie et al., 2024). This body of work motivates our benchmark design. When evaluating these uncertainty signals in automated assessment, we focus on grading tasks whose predictions lie in a constrained label space.

3 Benchmark Design

3.1 Overview

As described in § 1, we focus on uncertainty estimation approaches based on repeated generation, since such methods can be uniformly applied to both open-source and proprietary LLMs without requiring access to internal model states. Following prior work (Kuhn et al., 2023; Vashurin et al., 2025), we define uncertainty as the variability in a model's predictions. In LLM-based automatic assessment, this task aims to measure how stable or inconsistent the model is when judging the same question, where outputs may include final assessment scores, intermediate reasoning steps, and supporting evidence. Depending on the uncertainty quantification method, our benchmark parses LLM-generated responses in different ways

and computes corresponding uncertainty values. We then evaluate these uncertainty estimates with seven metrics from three perspectives related to their downstream usage. By comparing the ranking results of different uncertainty quantification methods across multiple automatic assessment datasets and evaluation metrics, our benchmark provides a comprehensive analysis of uncertainty quantification methods for LLM-based automatic assessment tasks. An overview of the benchmark workflow is illustrated in Fig. 1.

3.2 Uncertainty Quantification Methods

Existing uncertainty quantification methods for LLMs are often designed for open-ended question answering or reasoning tasks (Xia et al., 2025), where uncertainty is estimated from the variability or equivalence of free-form generations. In this work, we study uncertainty quantification in LLM-based automatic assessment, where repeated generations for the same input usually include both final grading labels and intermediate grading rationales. Accordingly, we organize the evaluated methods into two groups: categorical-based methods, which operate on the distribution of predicted score labels, and relation-based methods, which estimate uncertainty from the relationships among generated responses. Detailed definitions and formulas of all methods are provided in Appx. C.

3.2.1 Categorical Based Methods

Categorical-based methods use the predefined and discrete output label space of automatic assessment. Given multiple grading outputs for the same question, these methods first extract the final score label from each output and then compute uncertainty from the empirical frequency distribution of the labels. We evaluate four methods in this category. **Numset** (Lyu et al., 2024) measures the number of unique labels produced across repeated generations, where more distinct labels indicate larger variability. **Max-Agree-Rate (MAR)** (Xiong et al., 2023) measures the dominance of the most frequent label and converts it into an uncertainty score, so lower agreement corresponds to higher uncertainty. **Categorical-Entropy (CE)** (Lyu et al., 2024; De Boer et al., 2005) computes entropy over the empirical label distribution and therefore considers both majority and minority responses. **First–Second Distance (FSD)** (Lyu et al., 2024) focuses on the gap between the most frequent and second most frequent labels, making it sensitive to cases where the model oscillates between two plausible scores.

3.2.2 Relation Based Methods

Categorical-based methods are efficient, but they only use the extracted final labels and may be unstable when the number of repeated generations is limited. Relation-based methods additionally exploit the textual content of generated outputs, especially the reasoning traces produced by CoT prompting (Wei et al., 2022) and reasoning-oriented models. These methods follow a two-step paradigm. First, they construct a relation graph over repeated outputs, where each node is a generated grading response, and each edge represents the pairwise relationship between two responses. We consider three relation functions: **Jaccard Similarity**, which captures lexical overlap; **Embedding Cosine Similarity** (Reimers and Gurevych, 2019), which captures continuous semantic relatedness; and **Entailment Score** (Kuhn et al., 2023; Lewis et al., 2020; Zhang et al., 2019), which captures logical consistency through NLI-based sentence-pair entailment.

Second, relation-based methods compute graph-level uncertainty from the constructed relation graph. We evaluate four property-driven uncertainty measures. **Normalized Average Degree (NAD)**, also known as variational ratio (VR) (Huang et al., 2023), measures the overall connectivity density of the graph, where weaker connectivity indicates higher uncertainty. **Graph Eccentricity (GE)** measures the global spread of responses through shortest-path distances. **Graph Eigenvalues (Eigen)** use spectral connectivity of the graph Laplacian to quantify how tightly the response graph is connected. **Discrete Semantic Entropy (DSE)** (Kuhn et al., 2023) clusters semantically equivalent responses and computes entropy over the induced cluster distribution. Together, these relation-based methods provide continuous uncertainty estimates by measuring whether repeated grading outputs form a compact and coherent response set or a dispersed and inconsistent one.

3.3 Evaluation Metrics

To comprehensively assess UQ methods for automatic assessment, we design our evaluation from three key perspectives: effectiveness, stability and correlation. For each perspective, we develop various evaluation metrics.

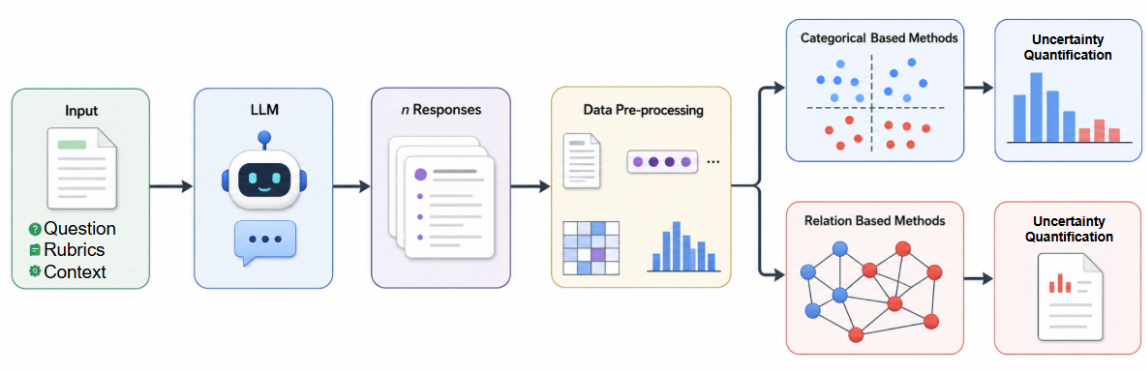

Figure 1: An overview of the benchmark workflow where n denotes the number of generation rollouts in this work.

3.3.1 Effectiveness

Uncertainty values are expected to indicate the reliability of responses produced by LLMs: responses with lower uncertainty should exhibit higher accuracy than those with higher uncertainty. In automatic assessment, this allows uncertainty to serve as a signal for human intervention and prevent low-reliability grading results from being released for downstream use. Based on this premise, we evaluate the effectiveness of each UQ method as an accuracy indicator using four metrics.

AUROC (Xia et al., 2025) evaluates whether uncertainty scores can distinguish correct and incorrect predictions. By treating incorrect ones as positive cases and uncertainty scores as ranking signals, AUROC measures the probability that an incorrect response receives higher uncertainty than a correct response. Higher AUROC indicates better discrimination between correct and incorrect predictions.

C-Index (Steck et al., 2007) is a rank-based metric for assessing consistency between predicted risk scores and observed outcomes. Since assessment scores are ordinal rather than purely binary, we use the C-index to evaluate whether responses with larger true errors receive higher uncertainty scores. Higher C-index indicates better alignment between uncertainty and error severity.

AUARC (Xia et al., 2025) measures the trade-off between prediction accuracy and rejection rate when selectively abstaining from low-confidence responses. The accuracy–rejection curve reports the accuracy of retained predictions as high-uncertainty responses are rejected. Higher AUARC indicates that rejecting a small fraction of uncertain responses yields larger accuracy gains, reflecting the utility of uncertainty for selective prediction and human-in-the-loop workflows.

AUERC (Simeone et al., 2011) is the error-based counterpart of AUARC, where the error-rejection curve plots the mean absolute error of retained predictions against the rejection rate. It evaluates whether high-uncertainty responses correspond to larger errors. Lower AUERC values are better, indicating that rejecting uncertain responses effectively reduces the remaining prediction error.

3.3.2 Stability

All uncertainty quantification methods in this work are derived from stochastic responses generated by LLMs, so the estimated uncertainty values are inherently noisy. Increasing the number of sampled responses can reduce this variance, but repeatedly generating many responses for the same grading request is computationally expensive and impractical in real-world deployment. To study this trade-off, we analyze how uncertainty estimates evolve as the response set size n increases and measure how quickly the values stabilize. Since uncertainty is mainly used for ranking responses, such as flagging high-uncertainty cases for human review, we also evaluate the stability of uncertainty-induced rankings using Spearman's rank correlation coefficient (Bishara and Hittner, 2012) between estimates computed from different response set sizes.

3.3.3 Correlation

Although different uncertainty quantification methods assess response uncertainty from diverse perspectives, they may still be correlated because of shared underlying evidence sources, such as categorical frequency or relational similarity. To summarize relationships among methods and help future work avoid redundant measures, we analyze pairwise correlations between methods using the Pearson correlation coefficient.

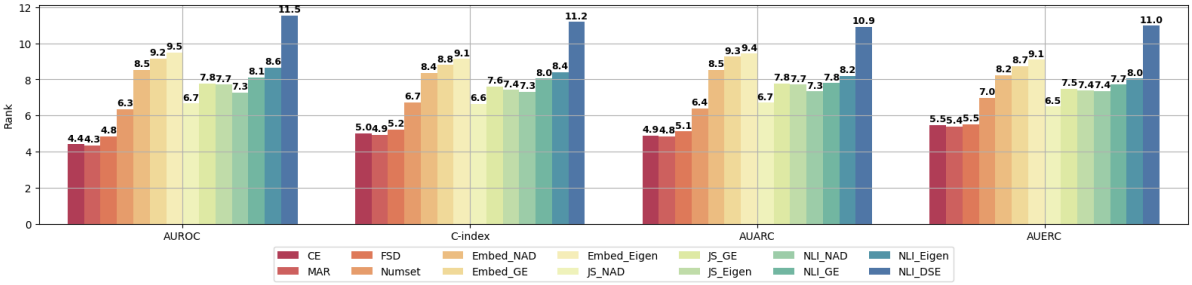

Figure 2: Average ranks of methods across questions, datasets, and generation strategies on effectiveness metrics.

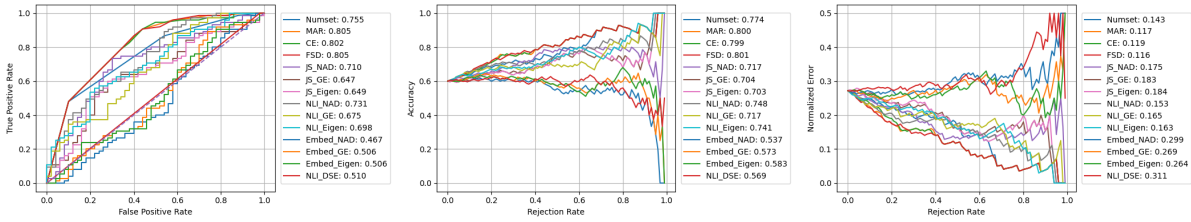

(a) Ministral-3-14B (AUROC) (b) Ministral-3-14B (AUARC) (c) Ministral-3-14B (AUERC)

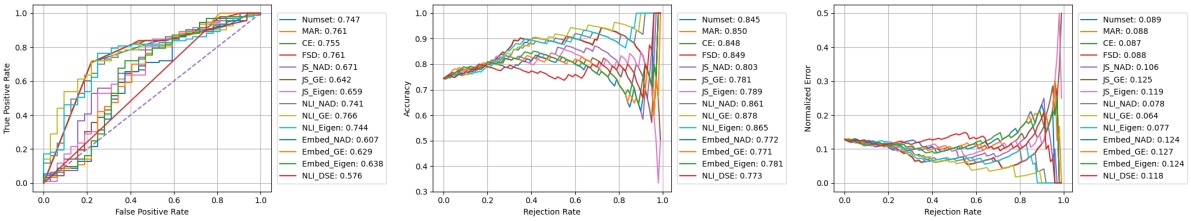

(d) gemini-3-flash (AUROC) (e) gemini-3-flash (AUARC) (f) gemini-3-flash (AUERC)

Figure 3: UQ Method Comparison between Ministral and Gemini over different metrics on the T1-PCK question.

4 Experiment

4.1 Datasets

To ensure the robustness of our conclusions, we benchmark the uncertainty quantification (UQ) methods on 14 grading questions sampled from five different sources, covering topics including essay writing, classroom Q&A, chemistry, earth science, and pedagogical knowledge. For each question, we collect 125 student responses spanning all score categories for subsequent experiments, resulting in a total of 1,750 samples. Detailed information about each question is provided in Appx. A.

4.2 LLM Settings

We benchmark a comprehensive set of mainstream LLMs, including three proprietary models: Claude-Haiku-4.5, Gemini-3-Flash, and GPT-5-Nano, and four open-source models: Qwen3-4B, Llama-3.1-8B, Ministral-3-14B, and DeepSeek-R1-Distill-Qwen-32B. Additional details about each model are provided in Appx. B. To ensure broad coverage of uncertainty quantification methods under realistic automatic assessment settings, we implement three commonly used generation strategies (Chu et al., 2024): (1) Zero-shot, (2) Zero-shot + Chain-of-Thought (CoT), and (3) Few-shot + Chain-of-Thought (CoT). During grading, each model and generation strategy is used to grade every response 15 times under the model's default generation settings. The resulting solution sets are then used to estimate uncertainty values with different UQ methods. More details are provided in Appx. D.

4.3 Effectiveness

Due to differences in score levels and response distributions across questions, the scale of uncertainty values produced by the same UQ method can vary substantially across datasets and tasks. To avoid large-magnitude uncertainty values dominating the evaluation, we rank each UQ method within each experimental configuration (model, dataset, and generation strategy) and aggregate results by averaging rankings across settings. Fig. 2 summarizes the aggregated rankings across the four evaluation metrics, while detailed results are provided in Appx. E. Overall, the categorical-based method CE achieves the best average ranking across all

metrics. This contrasts with prior UQ studies on open-ended Q&A tasks (Xia et al., 2025), where relation-based methods often perform better. Other categorical-based methods, including MAR, FSD, and Numset, also consistently outperform relation-based approaches, suggesting that the discrete output space of automatic assessment is particularly compatible with categorical uncertainty measures. To further investigate the influence of different factors, we conduct case studies on model choice, question characteristics, and generation settings. Fig. 3 compares Gemini and Ministral on the same dataset and generation condition. Although rankings remain relatively consistent across evaluation metrics within the same model, the preferred UQ methods differ substantially between models: Gemini favors relation-based methods, while Ministral shows stronger performance for category-based methods. Similar patterns are observed across datasets, whereas generation strategies have relatively limited influence. These findings suggest that UQ methods should be selected according to the underlying model and task conditions.

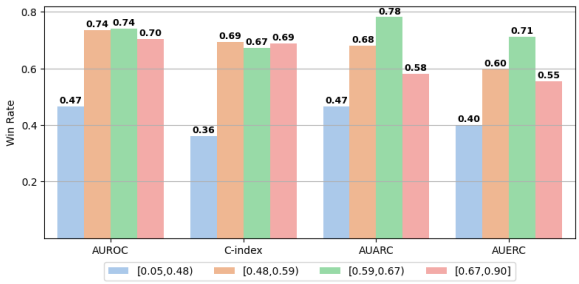

Figure 4: Win rates of category-based UQ over relation-based UQ across different task accuracy groups.

Finally, we explore the relationship between task accuracy and the preference for category- versus relation-based UQ methods. Based on Fig.2, we select MAR and JS_MAD as the best-performing category- and relation-based methods, respectively, and compare their win rates across different accuracy-level groups. As shown in Fig.4, relation-based methods achieve win rates above 50% when task accuracy is low, suggesting that relational information provides more reliable signals under difficult conditions. However, as accuracy increases, category-based methods begin to dominate, reaching win rates above 70%. Under high-accuracy conditions, the performance gap narrows again, with win rates around 55% for metrics such as AUARC and AUERC, indicating that both approaches become similarly effective.

4.4 Stability

In addition to effectiveness, we also evaluate different UQ methods from the perspective of stability, since obtaining many repeated generations for uncertainty estimation can be computationally expensive, especially for recent studies relying on high-cost proprietary LLMs. In Fig.5, we report the average number of repeats required to achieve 95% stability, defined as the change in the UQ value from n to n+1 repeats being less than 5%. More details about the stability metrics across different repeat numbers are provided in Appx.F. From the figure, we observe that relation-based methods using Jaccard similarity and embedding distance require fewer than five repeats on average to achieve stable uncertainty estimates, demonstrating the advantage of leveraging additional relational information compared to score-only UQ methods. In contrast, although category-based UQ methods exhibit relatively larger fluctuations in uncertainty values, their ranking differences across student responses stabilize quickly, requiring fewer than seven repeats on average. Considering practical deployment scenarios, if uncertainty decisions are made based on the relative ranking of responses within the same batch, category-based UQ methods still maintain strong overall performance. However, when uncertainty is determined using a fixed threshold, the higher variance in category-based uncertainty values may reduce consistency, making relation-based methods more stable and cost-efficient in practice.

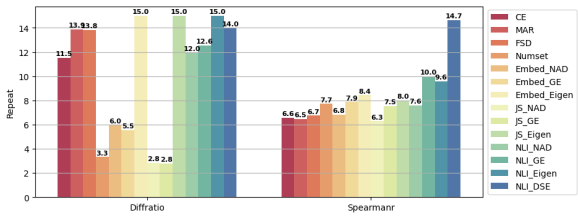

Figure 5: Stability comparison of uncertainty estimates produced by different UQ methods across repeated runs.

Similar to the effectiveness analysis, we further investigate how factors such as model choice, dataset, and generation strategy influence these conclusions. Specifically, Fig. 6 presents box plots illustrating the variance introduced by each factor. From the results, we draw conclusions consistent with the effectiveness study: model choice has the largest influence on UQ behavior, followed by the dataset, while the three generation-control strategies considered in this work have relatively limited

impact. Based on these observations, we suggest that UQ methods should be re-evaluated whenever the underlying model or dataset changes.

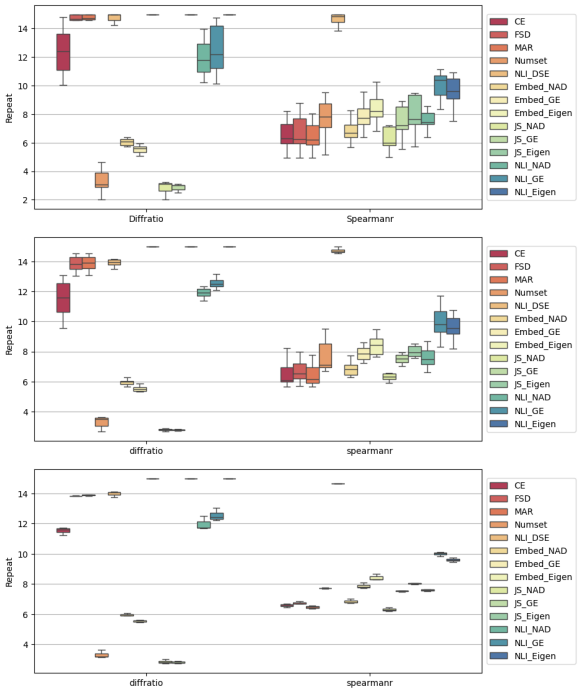

Figure 6: Distribution of the minimum number of repeats required to achieve stable UQ values across different models, questions, and generation strategies.

4.5 Correlation

To analyze correlations among methods, we compute the Pearson correlation coefficient for each pair of uncertainty quantification methods and aggregate these correlations across questions, models, and generation strategies by averaging. The resulting aggregated correlation matrix is shown in Fig. 7a. From the figure, we observe strong correlations among categorical-based methods, reflecting their reliance on similar frequency-based evidence. Among relation-based methods, those that share the same graph construction strategy exhibit particularly high correlations. In addition, methods derived from embedding-based and Jaccard similarity graphs show moderate correlation, as both capture semantic relatedness between responses, albeit at different representational levels. Similar to our analyses of effectiveness and stability, we examine the robustness of our correlation findings under variations in questions, models, and generation strategies. As shown in Fig. 7c, Fig. 7b and Fig. 7d, although variability exists across different models, questions, and strategies, the magnitude of this variance is relatively small (typically < 0.2) compared to the average correlation values (generally > 0.5). This indicates that the overall correlation patterns reported in our general results are robust to these sources of variation.

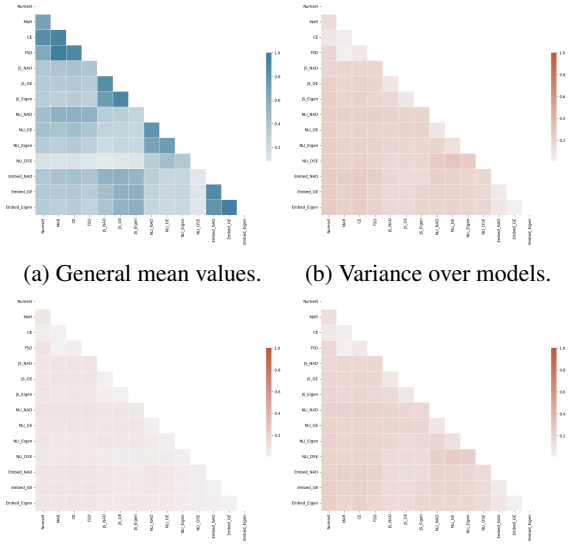

(a) General mean values. (b) Variance over models.

(c) Variance over questions. (d) Variance over strategics.

Figure 7: Mean and variance of aggregated pairwise correlation values among methods across factors.

5 Conclusion

In this work, we present a comprehensive benchmark of UQ methods for LLM-based automatic assessment. Although UQ has been widely studied in domains such as question answering, its application to automated assessment remains underexplored, particularly given the unique characteristics of assessment tasks, including categorical and ordinal outputs as well as rubric-driven evaluation settings. To address this gap, we establish a novel benchmark framework that categorizes existing UQ methods into two groups and systematically compares them from the perspectives of effectiveness, stability, and interrelationships. Through experiments across multiple assessment datasets spanning different subjects and LLMs of varying scales, our study reveals several important findings. Category-based methods demonstrate strong effectiveness in capturing uncertainty in assessment judgments, while relation-based methods exhibit greater robustness to sampling variation and provide more stable uncertainty estimates. Overall, our results suggest that no single UQ method is universally optimal, and that selecting an appropriate UQ metric should depend on the underlying LLM capability, assessment task, and evaluation setting.

Limitations

This work has three main limitations. First, our evaluation is limited to proprietary models and small-to-medium-scale open-source LLMs. While this choice reflects current computational constraints and realistic K–12 deployment conditions where large GPU resources are scarce, benchmarking larger open-source models will be essential as inference costs decrease. We argue that auto assessment by propriety and smaller-scale open-sourced LLMs is more practical, which is the reason we focus on LLMs out of large open-sourced models. Second, our assessment tasks are limited to questions with no more than 5 score levels and clear grading boundaries. Real-world assessment often involves more complex rubric settings (e.g., percentage scores, deduction-based, additive, or rubric items with dependency) that introduce different uncertainty patterns. However, we find the selected datasets are representative enough and of high quality, and these tasks are challenging enough for current LLMs. Thus, benchmarking the LLMs selected by this work is sufficient to showcase the auto-assessment advances in educational scenarios, and we leave these harder tasks for future work. Third, we focus on fixed grading criteria rather than personalized assessment based on student characteristics or teacher-specific rubrics. Although personalization is crucial for realistic scenarios, it requires substantial human-designed rubrics and gold labels, which increases grading inconsistency and raises privacy and ethical concerns. As the auto-assessment is still primarily deployed with fixed rubrics, and personalized assessment inherently shares the same capabilities of LLMs, we believe the results of evaluation with fixed rubrics are insightful enough, and will leave these personalized requirements for future exploration.

References

Alberto Abadie, Susan Athey, Guido W Imbens, and Jeffrey M Wooldridge. 2020. Sampling-based versus design-based uncertainty in regression analysis. *Econometrica*, 88(1):265–296.

Abbirah Ahmed, Arash Joorabchi, and Martin Hayes. 2022. On deep learning approaches to automated assessment: Strategies for short answer grading.

Kirsti M Ala-Mutka. 2005. A survey of automated assessment approaches for programming assignments. *Computer science education*, 15(2):83–102.

Anthropic. 2025. Introducing claude haiku 4.5. https://www.anthropic.com/news/claude-haiku-4-5. Large language model.

Anthony J Bishara and James B Hittner. 2012. Testing the significance of a correlation with nonnormal data: comparison of pearson, spearman, transformation, and resampling approaches. *Psychological methods*, 17(3):399.

Nathan Brake and Thomas Schaaf. 2024. Comparing two model designs for clinical note generation; is an llm a useful evaluator of consistency? In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 352–363.

Henry Braun, Isaac I Bejar, and David M Williamson. 2006. Rule based methods for automated scoring: Application in a licensing context. *Automated scoring of complex tasks in computer-based testing*, pages 83–122.

Saloua Chlaily, Debanshu Ratha, Pigi Lozou, and Andrea Marinoni. 2023. On measures of uncertainty in classification. *IEEE Transactions on Signal Processing*, 71:3710–3725.

Yucheng Chu, Hang Li, Kaiqi Yang, Harry Shomer, Hui Liu, Yasemin Copur-Gencturk, and Jiliang Tang. 2024. A llm-powered automatic grading framework with human-level guidelines optimization. *arXiv preprint arXiv:2410.02165*.

Ben E Cline, Carlyle C Brewster, and Richard D Fell. 2010. A rule-based system for automatically evaluating student concept maps. *Expert systems with applications*, 37(3):2282–2291.

Clayton Cohn, Nicole M. Hutchins, Tuan Le, and Gautam Biswas. 2024. A chain-of-thought prompting approach with llms for evaluating students' formative assessment responses in science. In *AAAI Conference on Artificial Intelligence*.

Jeremy Cole, Michael Zhang, Dan Gillick, Julian Eisenschlos, Bhuwan Dhingra, and Jacob Eisenstein. 2023. Selectively answering ambiguous questions. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 530–543.

Yasemin Copur-Gencturk and Tammy Tolar. 2022. Mathematics teaching expertise: A study of the dimensionality of content knowledge, pedagogical content knowledge, and content-specific noticing skills. *Teaching and Teacher Education*, 114:103696.

Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. 2019. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9268–9277.

Pieter-Tjerk De Boer, Dirk P Kroese, Shie Mannor, and Reuven Y Rubinstein. 2005. A tutorial on the cross-entropy method. *Annals of operations research*, 134(1):19–67.

DeepSeek-AI. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *Preprint*, arXiv:2501.12948.

Haotian Deng, Chris Farber, Jiyoon Lee, and David Tang. 2025. Rubric-conditioned llm grading: Alignment, uncertainty, and robustness. *arXiv preprint arXiv:2601.08843*.

Myroslava O Dzikovska, Rodney Nielsen, Chris Brew, Claudia Leacock, Danilo Giampiccolo, Luisa Bentivogli, Peter Clark, Ido Dagan, and Hoa Trang Dang. 2013. Semeval-2013 task 7: The joint student response analysis and 8th recognizing textual entailment challenge. In *Second Joint Conference on Lexical and Computational Semantics (* SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 263–274.

Emrah Emirtekin. 2025. Large language model-powered automated assessment: A systematic review. *Applied Sciences*.

Mingqi Gao, Jie Ruan, Renliang Sun, Xunjian Yin, Shiping Yang, and Xiaojun Wan. 2023. Human-like summarization evaluation with chatgpt. *arXiv preprint arXiv:2304.02554*.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, and 1 others. 2024. A survey on llm-as-a-judge. *The Innovation*.

Ben Hamner, Jaison Morgan, lynnvandev, Mark Shermis, and Tom Vander Ark. 2012. The hewlett foundation: Automated essay scoring. https://kaggle.com/competitions/asap-aes. Kaggle Competition.

Helia Hashemi, Jason Eisner, Corby Rosset, Benjamin Van Durme, and Chris Kedzie. 2024. Llm-rubric: A multidimensional, calibrated approach to automated evaluation of natural language texts. In *Annual Meeting of the Association for Computational Linguistics*.

Jia He, Mukund Rungta, David Koleczek, Arshdeep Sekhon, Franklin X Wang, and Sadid A. Hasan. 2024a. Does prompt formatting have any impact on llm performance? *ArXiv*, abs/2411.10541.

Peng He, Namsoo Shin, and Joseph Krajcik. 2024b. School level. *Handbook of Research on Science Learning Progressions*.

Colin A Higgins, Geoffrey Gray, Pavlos Symeonidis, and Athanasios Tsintsifas. 2005. Automated assessment and experiences of teaching programming. *Journal on Educational Resources in Computing (JERIC)*, 5(3):5–es.

Pedram Hosseini, Jessica M Sin, Bing Ren, Bryceton G Thomas, Elnaz Nouri, Ali Farahanchi, and Saeed Hassanpour. 2024. A benchmark for long-form medical question answering. *arXiv preprint arXiv:2411.09834*.

Yuheng Huang, Jiayang Song, Zhijie Wang, Shengming Zhao, Huaming Chen, Felix Juefei-Xu, and Lei Ma. 2023. Look before you leap: An exploratory study of uncertainty measurement for large language models. *arXiv preprint arXiv:2307.10236*.

Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, and 1 others. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.

Gerd Kortemeyer. 2023. Performance of the pre-trained large language model gpt-4 on automated short answer grading. *Discover Artificial Intelligence*, 4.

Luke Korthals, Emma Akrong, Gali Geller, Hannes Rosenbusch, Raoul Grasman, and Ingmar Visser. 2025. Towards reliable llm grading through self-consistency and selective human review: Higher accuracy, less work.

Jack Krolik, Herprit Mahal, Feroz Ahmad, Gaurav Trivedi, and Bahador Saket. 2024. Towards leveraging large language models for automated medical q&a evaluation. *arXiv preprint arXiv:2409.01941*.

Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. *arXiv preprint arXiv:2302.09664*.

Abhishek Kumar, Sonia Haiduc, Partha Pratim Das, and Partha Pratim Chakrabarti. 2024. Llms as evaluators: A novel approach to evaluate bug report summarization. *arXiv preprint arXiv:2409.00630*.

Alexander Kurz, Katja Hauser, Hendrik Alexander Mehrtens, Eva Krieghoff-Henning, Achim Hekler, Jakob Nikolas Kather, Stefan Fröhling, Christof Von Kalle, and Titus Josef Brinker. 2022. Uncertainty estimation in medical image classification: systematic review. *JMIR Medical Informatics*, 10(8):e36427.

Dan Levi, Liran Gispan, Niv Giladi, and Ethan Fetaya. 2022. Evaluating and calibrating uncertainty prediction in regression tasks. *Sensors*, 22(15):5540.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 7871–7880.

Dawei Li, Bohan Jiang, Liangjie Huang, Alimohammad Beigi, Chengshuai Zhao, Zhen Tan, Amrita Bhattacharjee, Yuxuan Jiang, Canyu Chen, Tianhao Wu, Kai Shu, Lu Cheng, and Huan Liu. 2024a. From generation to judgment: Opportunities and challenges of llm-as-a-judge. In *Conference on Empirical Methods in Natural Language Processing*.

Hai Li, Chenglu Li, Wanli Xing, Sami Baral, and Neil T. Heffernan. 2024b. Automated feedback for student math responses based on multi-modality and fine-tuning. *Proceedings of the 14th Learning Analytics and Knowledge Conference*.

Hang Li, Yucheng Chu, Kaiqi Yang, Yasemin Copur-Gencturk, and Jiliang Tang. 2025a. Llm-based automated grading with human-in-the-loop. *2025 IEEE International Conference on Teaching, Assessment, and Learning for Engineering (TALE)*, pages 1–8.

Qingquan Li, Shaoyu Dou, Kailai Shao, Chao Chen, and Haixiang Hu. 2025b. Evaluating scoring bias in llm-as-a-judge. *ArXiv*, abs/2506.22316.

Alexander H Liu, Kartik Khandelwal, Sandeep Subramanian, Victor Jouault, Abhinav Rastogi, Adrien Sadé, Alan Jeffares, Albert Jiang, Alexandre Cahill, Alexandre Gavaudan, and 1 others. 2026. Ministral 3. *arXiv preprint arXiv:2601.08584*.

Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. G-eval: Nlg evaluation using gpt-4 with better human alignment. In *Proceedings of the 2023 conference on empirical methods in natural language processing*, pages 2511–2522.

Daniel Mark Low, Kate H Bentley, and Satrajit S. Ghosh. 2019. Automated assessment of psychiatric disorders using speech: A systematic review. *Laryngoscope Investigative Otolaryngology*, 5:96 – 116.

Chang Lu and Maria Cutumisu. 2021. Integrating deep learning into an automated feedback generation system for automated essay scoring. *International Educational Data Mining Society*.

Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, and Dongmei Zhang. 2023. Wizardmath: Empowering mathematical reasoning for large language models via reinforced evol-instruct. *arXiv preprint arXiv:2308.09583*.

Qing Lyu, Kumar Shridhar, Chaitanya Malaviya, Li Zhang, Yanai Elazar, Niket Tandon, Marianna Apidianaki, Mrinmaya Sachan, and Chris Callison-Burch. 2024. Calibrating large language models with sample consistency. *ArXiv*, abs/2402.13904.

José Mena, Oriol Pujol, and Jordi Vitrià. 2021. A survey on uncertainty estimation in deep learning classification systems from a bayesian perspective. *ACM Computing Surveys (CSUR)*, 54(9):1–35.

José Carlos Paiva, José Paulo Leal, and Álvaro Figueira. 2022. Automated assessment in computer science education: A state-of-the-art review. *ACM Transactions on Computing Education (TOCE)*, 22(3):1–40.

Vyas Raina, Adian Liusie, and Mark Gales. 2024. Is llm-as-a-judge robust? investigating universal adversarial attacks on zero-shot llm assessment. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 7499–7517.

Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. 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 3982–3992.

Murat Sensoy, Lance Kaplan, and Melih Kandemir. 2018. Evidential deep learning to quantify classification uncertainty. *Advances in neural information processing systems*, 31.

Lin Shi, Chiyu Ma, Wenhua Liang, Xingjian Diao, Weicheng Ma, and Soroush Vosoughi. 2024. Judging the judges: A systematic study of position bias in llm-as-a-judge.

Paolo Simeone, Claudio Marrocco, and Francesco Tortorella. 2011. Shaping the error-reject curve of error correcting output coding systems. In *International Conference on Image Analysis and Processing*, pages 118–127. Springer.

Aaditya Singh, Adam Fry, Adam Perelman, Adam Tart, Adi Ganesh, Ahmed El-Kishky, Aidan McLaughlin, Aiden Low, AJ Ostrow, Akhila Ananthram, and 1 others. 2025. Openai gpt-5 system card. *arXiv preprint arXiv:2601.03267*.

Juhi Singh, Ziqiao Ao, and Sebastian Antinome. Optimizing prompt refinement: Algorithmic strategies for large language model-based text classification. *Available at SSRN 5525975*.

Hwanjun Song, Hang Su, Igor Shalyminov, Jason (Jinglun) Cai, and Saab Mansour. 2024a. Finesure: Fine-grained summarization evaluation using llms. *ArXiv*, abs/2407.00908.

Yishen Song, Qianta Zhu, Huaibo Wang, and Qinhua Zheng. 2024b. Automated essay scoring and revising based on open-source large language models. *IEEE Transactions on Learning Technologies*, 17:1880–1890.

Harald Steck, Balaji Krishnapuram, Cary Dehing-Oberije, Philippe Lambin, and Vikas C Raykar. 2007. On ranking in survival analysis: Bounds on the concordance index. *Advances in neural information processing systems*, 20.

Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan

Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, and 1 others. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.

Qwen Team. 2025. Qwen3 technical report. *Preprint*, arXiv:2505.09388.

Mani Kumar Tellamekala, Shahin Amiriparian, Björn W Schuller, Elisabeth André, Timo Giesbrecht, and Michel Valstar. 2023. Cold fusion: Calibrated and ordinal latent distribution fusion for uncertainty-aware multimodal emotion recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(2):805–822.

Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher D Manning. 2023. Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5433–5442.

Roman Vashurin, Ekaterina Fadeeva, Artem Vazhentsev, Lyudmila Rvanova, Daniil Vasilev, Akim Tsvigun, Sergey Petrakov, Rui Xing, Abdelrahman Sadallah, Kirill Grishchenkov, and 1 others. 2025. Benchmarking uncertainty quantification methods for large language models with lm-polygraph. *Transactions of the Association for Computational Linguistics*, 13:220–248.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.

Zhiqiu Xia, Jinxuan Xu, Yuqian Zhang, and Hang Liu. 2025. A survey of uncertainty estimation methods on large language models. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 21381–21396.

Zhuyang Xie, Yan Yang, Jie Wang, Xiaorong Liu, and Xiaofan Li. 2024. Trustworthy multimodal fusion for sentiment analysis in ordinal sentiment space. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(8):7657–7670.

Miao Xiong, Zhiyuan Hu, Xinyang Lu, Yifei Li, Jie Fu, Junxian He, and Bryan Hooi. 2023. Can llms express their uncertainty? an empirical evaluation of confidence elicitation in llms. *arXiv preprint arXiv:2306.13063*.

Jiayi Ye, Yanbo Wang, Yue Huang, Dongping Chen, Qihui Zhang, Nuno Moniz, Tian Gao, Werner Geyer, Chao Huang, Pin-Yu Chen, Nitesh V. Chawla, and Xiangliang Zhang. 2024. Justice or prejudice? quantifying biases in llm-as-a-judge. *ArXiv*, abs/2410.02736.

Caiqi Zhang, Fangyu Liu, Marco Basaldella, and Nigel Collier. 2024. Luq: Long-text uncertainty quantification for llms. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5244–5262.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.

Tianjie Zhao, Sheng Wang, Chaojun Ouyang, Min Chen, Chenying Liu, Jin Zhang, Long Yu, Fei Wang, Yong Xie, Jun Li, and 1 others. 2024. Artificial intelligence for geoscience: Progress, challenges, and perspectives. *The Innovation*, 5(5).

Yanqi Zhou, Tao Lei, Hanxiao Liu, Nan Du, Yanping Huang, Vincent Zhao, Andrew M Dai, Quoc V Le, James Laudon, and 1 others. 2022. Mixture-of-experts with expert choice routing. *Advances in Neural Information Processing Systems*, 35:7103–7114.

Kaijie Zhu, Jindong Wang, Jiaheng Zhou, Zichen Wang, Hao Chen, Yidong Wang, Linyi Yang, Weirong Ye, Neil Zhenqiang Gong, Yue Zhang, and Xingxu Xie. 2023a. Promptrobust: Towards evaluating the robustness of large language models on adversarial prompts. *Proceedings of the 1st ACM Workshop on Large AI Systems and Models with Privacy and Safety Analysis*.

Lianghui Zhu, Xinggang Wang, and Xinlong Wang. 2023b. Judgelm: Fine-tuned large language models are scalable judges. *arXiv preprint arXiv:2310.17631*.

A Dataset Detail

To provide a comprehensive evaluation, we collect 14 grading questions from five different sources for our experiments, including two open-source datasets and three private datasets. Specifically, the two open-source datasets are the Automated Student Assessment Prize (ASAP) dataset (Hamner et al., 2012), which focuses on short-essay grading (150–550 words) written by students in Grades 7–10, and SemEval-2013 Task 7 (SemEval-13-T7) (Dzikovska et al., 2013), which scores student responses to questions collected from in-class exercises, tests, and tutorial dialogues. For the three private datasets, we manually collect two sets of questions from elementary and middle school classes covering chemistry and general science following the 3DLP standard (He et al., 2024b), and one set from adult learners focusing on pedagogical skills (Copur-Gencturk and Tolar, 2022). For the open-source datasets, we use the original ground-truth labels to evaluate the correctness of LLM-based automatic grading. For the private datasets, the ground-truth labels are obtained via agreement between two expert human graders. In cases of disagreement, a third grader adjudicates to determine the final label. Based on these labels, we compute the correctness of the automatic assessment results. In Tab. 1, we summarize the statistics and characteristics of each dataset and the corresponding grading questions. For each question, we randomly sample 125 answers covering all score categories for evaluation.

Table 1: List of questions evaluated by our benchmark.

Question	Score	Subject	Question	Score	Subject
ASAP	0-3	Essay Writing	Semveal	0-3	Mixture
T1-DCI	0-3	Chemistry	T1-SEP	0-3	Chemistry
T2-DCI	0-3	Chemistry	T2-SEP	0-3	Chemistry
T1-CK	0-3	Math Pedagogy	T1-PCK	0-3	Math Pedagogy
T2-CK	0-3	Math Pedagogy	T2-PCK	0-3	Math Pedagogy
U41Q1	0-4	Earth Science	U41Q2	0-4	Earth Science
U42Q1	0-4	Earth Science	U42Q2	0-4	Earth Science

B LLM Detail

To ensure the robustness of our conclusions, we benchmark a comprehensive set of mainstream LLMs, including seven closed-source API-based models from three leading vendors, i.e., OpenAI, Anthropic, and Google, and 4 open-source models released by research teams worldwide. For the closed-source models, we choose the most common used models by prior LLM-based automatic assessment studies (Chu et al., 2024) to ensure the validenss of the conclusion. For the open-source models, we cover recent representative variants of standard LLMs, including mixture-of-experts models (Zhou et al., 2022) and explicitly reasoning-oriented models (Jaech et al., 2024). With this broad model coverage, we aim to ensure that our benchmark results generalize across diverse deployment scenarios of automatic assessment. Tab. 2 summarizes the details of all evaluated models.

Table 2: List of LLMs evaluated by our benchmark.

Proprietary Models		
Model	**Model ID**	**Scale**
Calude Haiku 4.5 (Anthropic, 2025)	claude-haiku-4-5-20251001	/
Gemini 3 Flash (Team et al., 2023)	gemini-3-flash-preview	/
GPT-5 nano (Singh et al., 2025)	gpt-5-nano	/
Open-source Models		
Model	**Model ID**	**Scale**
Qwen3-4B (Team, 2025)	Qwen3-4B-Instruct-2507	4B
Llama-3.1-8B (Grattafiori et al., 2024)	Llama-3.1-8B-Instruct	8B
Ministral-3-14B (Liu et al., 2026)	Ministral-3-14B-Instruct-2512	14B
DeepSeek-R1-Distill-Qwen-32B (DeepSeek-AI, 2025)	DeepSeek-R1-Distill-Qwen-32B	32B

C Details of Uncertainty Quantification Methods

C.1 Categorical Based Methods

Existing studies on uncertainty quantification have primarily focused on open-ended question answering (Q&A) or reasoning tasks (Xia et al., 2025), where model responses are free-form. In contrast, automatic assessment is typically formulated as a classification or scoring task, in which model outputs are constrained to a fixed label or score space. Therefore, we select uncertainty quantification methods (Xia et al., 2025; Lyu et al., 2024; Kuhn et al., 2023) that can capture uncertainty from categorical outputs. Meanwhile, with the recent adoption of chain-of-thought (CoT) prompting (Wei et al., 2022) and the emergence of reasoning-oriented models, model outputs in automatic assessment may include not only final grading scores but also intermediate grading rationales. To capture uncertainty manifested in these intermediate representations, we further incorporate relation-based uncertainty quantification methods.

Let $\mathcal{O} = \{o_i\}_{i=1}^{N}$ denote the set of repeated grading outputs for the same input, where N is the

number of repeated generations. For categorical-based methods, each o_i is mapped to its final assessment score label, and $\mathcal{O}_u$ denotes the set of unique labels observed across repeated generations. Owing to the pre-defined and discrete output label space, methods in this category are relatively simple to compute and computationally efficient. In general, they derive uncertainty from the empirical frequency distribution of score labels observed across repeated queries to the same request.

Numset. Numset (Lyu et al., 2024) measures uncertainty as the number of unique answers observed across multiple samples:

$$U_{\text{Numset}} = |\mathcal{O}_u|. \tag{1}$$

A larger $|\mathcal{O}_u|$ indicates greater variability in the model's outputs and thus higher uncertainty.

Max-Agree-Rate (MAR). MAR (Xiong et al., 2023) captures uncertainty by incorporating the full distribution of answer frequencies. Compared to Numset, which only counts the number of distinct answers, MAR reflects the dominance of the most frequent answer among all sampled responses. A larger agreement rate indicates that the model consistently produces the same answer and is therefore more confident. To express this as an uncertainty measure, we use the complement of the maximum empirical frequency:

$$U_{\text{MAR}} = 1 - \max_{o \in \mathcal{O}_u} \frac{1}{N} \sum_{i=1}^{N} \mathbf{1}[o_i = o]. \tag{2}$$

A larger U_{MAR} indicates lower agreement and higher uncertainty.

Categorical Entropy (CE). Categorical Entropy (Lyu et al., 2024) is a classical distribution chaos quantification method over output class probabilities for classification problems (De Boer et al., 2005). As automatic assessment is also commonly treated as a classification task, we compute entropy over the empirical distribution of predicted grade labels:

$$U_{\text{CE}} = \sum_{o \in \mathcal{O}_u} -p_o \log(p_o), \quad p_o = \frac{1}{N} \sum_{i}^{N} \mathbf{1}[o_i = o]. \tag{3}$$

Here, p_o is the normalized frequency of each unique answer $o \in \mathcal{O}_u$. Compared to MAR, CE not only captures the majority class but also accounts for all categories, so changes in minority responses are reflected in the uncertainty value. A higher U_{CE} indicates a more chaotic answer distribution and thus higher uncertainty.

First-Second Distance (FSD). FSD (Lyu et al., 2024) is a frequency-based uncertainty measure that addresses limitations of commonly used methods such as MAR and Entropy. It quantifies the gap between the proportions of samples supporting the most frequent answer $\bar{o}$ and the second most frequent answer $\bar{\bar{o}}$. To align this measure with our uncertainty formulation, where larger values indicate higher uncertainty, we use the complement of this gap:

$$U_{\text{FSD}} = 1 - \frac{1}{N} \left(\sum_{i=1}^{N} \mathbf{1}[o_i = \bar{o}] - \sum_{i=1}^{N} \mathbf{1}[o_i = \bar{\bar{o}}] \right). \tag{4}$$

Compared with MAR and CE, FSD focuses on the relative dominance between the top two categories. It is particularly informative when the model is uncertain between two plausible answers, since it avoids overconfidence induced by considering only the most-voted answer.

C.2 Relation Based Methods

Although evaluating uncertainty directly in the categorical label space is computationally convenient, it has notable limitations. Due to the discrete nature of label categories in automatic assessment tasks, estimated uncertainty values can exhibit large fluctuations when the number of sampled responses N is small. For example, entropy-based measures such as CE are particularly sensitive to sampling noise before the response distribution stabilizes (Cui et al., 2019). In practical automatic assessment settings, repeatedly querying an LLM for the same input is constrained by efficiency and cost considerations, making such instability especially problematic.

The recent emergence of reasoning-path generation in large language models provides rich intermediate signals that help mitigate the limitations of category-based uncertainty quantification. Specifically, by modeling semantic similarity (Reimers and Gurevych, 2019) or logical entailment (Kuhn et al., 2023) among the reasoning traces associated with grading outputs, relation-based uncertainty quantification methods enable more fine-grained and continuous assessments of model uncertainty. These methods construct a relation graph $\mathcal{G} = (V, E)$ over the output set $\mathcal{O}$, where each ver-

tex $v \in V$ corresponds to one generated grading response and each edge $e_{ij} \in E$ is weighted by the pairwise relationship between responses o_i and o_j. A compact and coherent graph indicates lower uncertainty, whereas a dispersed graph indicates higher uncertainty.

C.2.1 Relation Graph Construction

Inspired by graph-based clustering analysis, recent studies (Huang et al., 2023; Xia et al., 2025) construct a relation graph over generated outputs and measure the tightness of relationships among responses. A key factor influencing the analysis lies in how this relation graph is constructed: different definitions of pairwise relationships yield different notions of answer similarity and capture uncertainty at varying semantic depths and computational costs. In this work, we consider three representative relation functions.

Jaccard Similarity. Jaccard Similarity measures lexical overlap between two responses based on their token sets:

$$s_{ij}^{\text{jac}} = \frac{|T(o_i) \cap T(o_j)|}{|T(o_i) \cup T(o_j)|}, \tag{5}$$

where $T(o)$ denotes the set of tokens in response o. This approach captures surface-level similarity and is computationally efficient, but it is limited in its ability to reflect deeper semantic equivalence between paraphrased or lexically diverse answers.

Embedding Cosine Similarity. Embedding Cosine Similarity (Reimers and Gurevych, 2019) computes the cosine similarity between vector representations of responses obtained from a sentence embedding model:

$$s_{ij}^{\text{emb}} = \frac{\mathbf{h}_i^\top \mathbf{h}_j}{\|\mathbf{h}_i\|_2 \, \|\mathbf{h}_j\|_2}, \quad \mathbf{h}_i = f_{\text{emb}}(o_i). \tag{6}$$

Here, $\mathbf{h}_i \in R^d$ denotes the semantic embedding of response o_i produced by a pre-trained encoder f_{emb}. By operating in a continuous semantic embedding space, this method captures higher-level semantic relatedness beyond exact lexical overlap, at the cost of additional computational overhead for embedding extraction.

Entailment Score. Entailment Score (Kuhn et al., 2023) models pairwise relationships using natural language inference (NLI) by estimating the degree to which one answer entails another. Off-the-shelf NLI models output class probabilities over entailment, neutral, and contradiction (Lewis et al., 2020). However, the limited input context window makes them unsuitable for directly processing long, multi-step reasoning traces produced by LLMs. To address this, we adopt a BERTScore-inspired strategy (Zhang et al., 2019): each response is split into sentences, entailment scores are computed over sentence pairs, and a mean–max aggregation is applied. Following prior work (Kuhn et al., 2023), we use the entailment probability to construct the relation graph. Since entailment is directional, we symmetrize the pairwise relationship by averaging both directions:

$$\begin{aligned} s_{i\to j}^{\text{nli}} &= \frac{1}{M} \sum_{m=1}^{M} \max_k P(\text{entail} \mid o_{i,m}, o_{j,k}), \\ s_{ij}^{\text{nli}} &= \frac{1}{2}\big(s_{i\to j}^{\text{nli}} + s_{j\to i}^{\text{nli}}\big). \end{aligned} \tag{7}$$

Here, $P(\cdot)$ denotes the entailment probability predicted by the NLI model; M and K are the numbers of sentences after splitting answers o_i and o_j, respectively; and $o_{i,m}$ and $o_{j,k}$ denote the m-th and k-th sentences of answers o_i and o_j. Entailment scores capture fine-grained semantic and logical consistency between responses, but they are computationally expensive due to NLI inference over sentence pairs.

C.2.2 Property-driven Uncertainty

Given the relation graph $\mathcal{G}$, different graph-theoretic properties can be leveraged to quantify its tightness. Below, we describe four representative property-driven uncertainty measures that have been widely adopted in prior work on uncertainty quantification for Q&A and reasoning tasks (Xia et al., 2025). Let s_{ij}^* denote the pairwise similarity between responses o_i and o_j obtained from one of the relation graph construction methods.

Normalized Average Degree (NAD). Normalized Average Degree, also known as the variational ratio (VR) (Huang et al., 2023), measures the connectivity density of the relation graph. Intuitively, if each answer is closely connected to many others, the response set forms a dense and compact cluster, indicating low uncertainty. We define the NAD-based uncertainty as:

$$U_{\text{NAD}} = 1 - \frac{1}{N(N-1)} \sum_{i=1}^{N} \sum_{i \neq j} s_{ij}^*. \tag{8}$$

Larger values of U_{NAD} indicate weaker local connectivity and thus higher uncertainty.

Graph Eccentricity (GE). Graph Eccentricity quantifies the maximum shortest-path distance from each node to all other nodes in the graph and summarizes the global spread of responses:

$$U_{\mathrm{ECC}} = \frac{1}{N} \sum_i \max_j [d^{\mathrm{sp}}_{ij}]. \tag{9}$$

Here, d^{sp} denotes the all-pairs shortest-path distance induced by the complement of similarity, $1 - s^*_{ij}$. Smaller values indicate that all answers are mutually close in the relation graph, whereas larger values imply that some answers are distant from the rest, reflecting higher uncertainty.

Graph Eigenvalues (Eigen). Graph Eigenvalues use spectral properties of the graph Laplacian to provide a global measure of connectivity. In particular, the second smallest eigenvalue, $\lambda_2(L)$, known as algebraic connectivity, reflects how tightly connected the graph is. We define the eigenvalue-based uncertainty as:

$$U_{\mathrm{Eigen}} = \frac{1}{\lambda_2(L)}, \quad L = D - A, \quad A_{ij} = s^*_{ij}. \tag{10}$$

Here, L is the graph Laplacian constructed from the degree matrix D and adjacency matrix A. A more tightly connected graph exhibits larger $\lambda_2(L)$ and thus lower uncertainty, whereas smaller $\lambda_2(L)$ corresponds to higher uncertainty.

Discrete Semantic Entropy (DSE). Discrete Semantic Entropy extends entropy-based uncertainty measures to the relation-graph setting by computing entropy over semantically clustered responses (Kuhn et al., 2023). By grouping answers into equivalence classes based on pairwise relations in $\mathcal{G}$ and computing the entropy of the induced discrete distribution, DSE captures both the diversity of semantic modes and their relative frequencies. Specifically, we define:

$$U_{\mathrm{DSE}} = -\sum_{m=1}^{M} p_m \log p_m, \quad p_m = \frac{|\mathcal{C}_m|}{N}. \tag{11}$$

Here, $\{\mathcal{C}_m\}_{m=1}^{M}$ denotes the clusters obtained from the relation graph $\mathcal{G}$, and $|\mathcal{C}_m|$ is the size of cluster $\mathcal{C}_m$. As clustering over $\mathcal{G}$ incurs additional computational cost, we follow prior work (Kuhn et al., 2023) and compute this metric only for the NLI-based relation graph using bidirectional entailment relations. Larger U_{DSE} indicates a more fragmented semantic landscape and thus higher uncertainty.

D Setting Detail

To ensure broad coverage of uncertainty quantification methods under realistic automatic assessment settings, we implement three widely used generation strategies adopted in recent automatic assessment studies (Chu et al., 2024).

Zero-shot is the most straightforward way to leverage LLMs for automatic assessment. Given a question, a student answer, and a grading rubric, the LLM is prompted to directly produce a score without any demonstrations or intermediate reasoning steps.

Zero-shot + Chain of Thought (COT) extends zero-shot by eliciting intermediate reasoning steps from the LLM. By decomposing the grading process into a step-by-step reasoning trajectory, the model is better able to align with the grading criteria specified in the rubric and to analyze student responses more comprehensively, often resulting in improved grading performance.

Few-shot + Chain of Thought (COT) further augments the prompt with exemplar answers corresponding to different score levels. These demonstrations provide concrete interpretations of the abstract criteria in the rubric, helping the model better internalize grading standards and produce more accurate and consistent assessment outcomes. Specifically, in our evaluation, we include two demonstration examples for each score level.

Finally, to ensure the robustness of our conclusions, we set the number of repeated grading samples for each answer to fifteen and use the default generation parameters recommended by each model for every generation. The prompts used for each query setting are presented as follows: Fig. 8 shows the zero-shot prompt, Fig. 10 presents the zero-shot + CoT prompt, and Fig. **??** illustrates the few-shot + CoT prompt.

E Effectiveness Result

We present the detailed metric values and ranking results of different UQ methods across datasets, LLMs, and generation controls in the following tables. Due to the large volume of results, we report the results for each LLM separately. Specifically, Tab.3 presents the results for GPT-5-Nano, Tab.4 for Gemini-3-Flash, Tab.5 for Claude-Haiku-4.5, Tab.6 for Qwen3-4B, Tab.7 for Llama-3.1-8B, Tab.8 for Ministral-3-14B, and Tab. 9 for

```
As a virtual teaching assistant, your role
involves assisting in the evaluation of
student assignment submissions. Please be
impartial and only use the grading guidelines,
the student`s submission, and the task in your
evaluation.
Analyze the submission and evaluate how well
the student`s submission answers the task.
Grade this submission of a student in 10th,
with a score in {grade_range}
# COT Command:
  In this task, you should give the response
directly WITHOUT using the chain-of-thought or
solving the problem step by step.
# Question (the students will work on):
  {question}
# Prompt
  {grading_rubrics}
# Submission (to be evaluated):
  {student_answer}
# Output Format:
  `text (reasons for your decisions)...  ##
Final: final_score`.
```

Figure 8: Prompt template for Zero-shot without COT Setting.

DeepSeek-R1-Distill-Qwen-32B. In addition to the detailed results, we further visualize the influence of different factors, including model choice, dataset, and generation control, on the ranking behavior of UQ methods. Specifically, Fig. 11a illustrates the ranking variance introduced by model differences, Fig. 11b shows the variance caused by datasets, and Fig. 11c presents the variance under different generation controls. From these plots, we observe that data choice is the most influential factor affecting UQ method rankings, while model also introduce noticeable variation. In contrast, the influence of generation-control strategies is relatively limited.

F Stability Result

To visualize how stability changes with the number of repetitions, we present the two stability metrics in Fig. 12. From the figure, we observe that for the DiffRatio metric, UQ methods based on embedding relations achieve less than a 5 change ratio at a very early stage, while methods using NLI-based relations exhibit a much slower convergence speed. This finding highlights the importance of selecting appropriate relation metrics in UQ studies. For ranking-based metrics such as Spearman correlation, most UQ methods converge within a similar number of repetitions, and all methods reach the

Table 3: Metric value and ranking results of each uncertainty quantification methods with GPT-5-nano over 14 questions.

| Criteria | Ensemble | AUROC | | |

Table 4: Metric value and ranking results of each uncertainty quantification methods with Gemini-3-flash over 14 questions.

Control	Question	AUROC														C-AUROC														AUARC														AUERC													
		CE	Embed_Eigen	Embed_GE	Embed_NAD	PSD	JS_Eigen	JS_GE	JS_NAD	MAR	NLI_DSE	NLI_Eigen	NLI_GE	NLI_NAD	Numet

Table 6: Metric value and ranking results of each uncertainty quantification methods with Qwen3-4B over 14 questions.

Control	Question	AUROC															C-AUROC																			AUARC															AUERC													
		CE	Embed_Eigen	Embed_GE	Embed_NAD	PSD	JS_Eigen	JS_GE	JS_NAD	MAR	NLI_DSE	NLI_Eigen	NLI_GE	NLI_NAD	Nonest

Table 8: Metric value and ranking results of each uncertainty quantification methods with Ministral-3-14B over 14 questions.

Control	Question	AUROC													C-AUROC													AUARC													AUERC												
		CE	Embed_Eigen	Embed_GE	Embed_NAD	PSD	JS_Eigen	JS_GE	JS_NAD	MAR	NLI_DSE	NLI_Eigen	NLI_GE	NLI_NAD	Nonest

```
As a virtual teaching assistant, your role
involves assisting in the evaluation of
student assignment submissions. Please be
impartial and only use the grading guidelines,
the student`s submission, and the task in your
evaluation.
Analyze the submission and evaluate how well
the student`s submission answers the task.
Grade this submission of a student in 10th,
with a score in {grade_range}
# COT Command:
  In this task, you should use
chain-of-thought to solve the problem
step by step before giving the response.
# Question (the students will work on):
  {question}
# Prompt
  {grading_rubrics}
# Submission (to be evaluated):
  {student_answer}
# Output Format:
  `text (reasons for your decisions)...  ##
Final: final_score`.
```

Figure 9: Prompt template for Zero-shot with COT Setting.

95% stability threshold after approximately 10 repeats.

```
As a virtual teaching assistant, your role
involves assisting in the evaluation of
student assignment submissions. Please be
impartial and only use the grading guidelines,
the student`s submission, and the task in your
evaluation.
Analyze the submission and evaluate how well
the student`s submission answers the task.
Grade this submission of a student in 10th,
with a score in {grade_range}
# COT Command:
  In this task, you should use
chain-of-thought to solve the problem
step by step before giving the response.
# Question (the students will work on):
  {question}
# Prompt
  {grading_rubrics}
# Few-shot Information (examples to follow):
  {example_pairs_of_answer_grading}
# Submission (to be evaluated):
  {student_answer}
# Output Format:
  `text (reasons for your decisions)...  ##
Final: final_score`.
```

Figure 10: Prompt template for Few-shot with COT Setting.

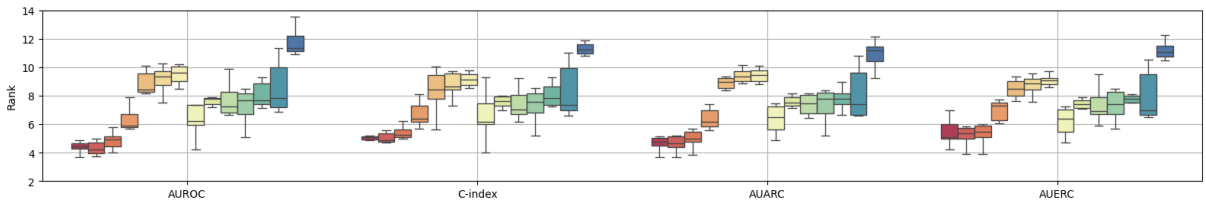

(a) Models.

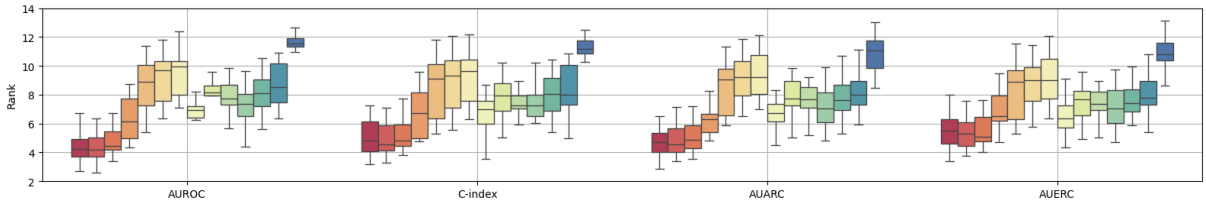

(b) Questions.

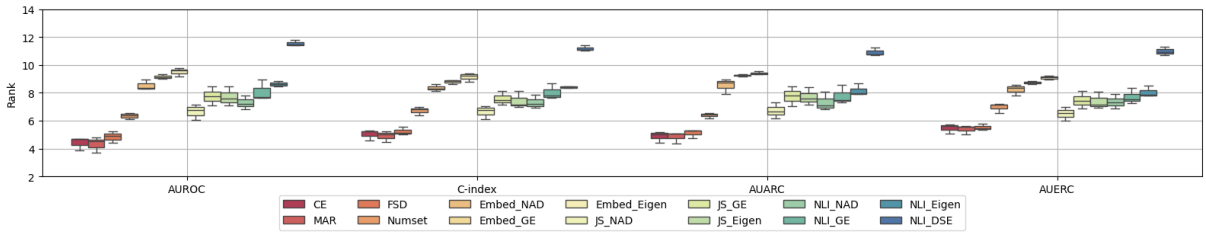

(c) Generation strategies.

Figure 11: Distribution of various in-group average rank for each model, question and generation strategies variants.

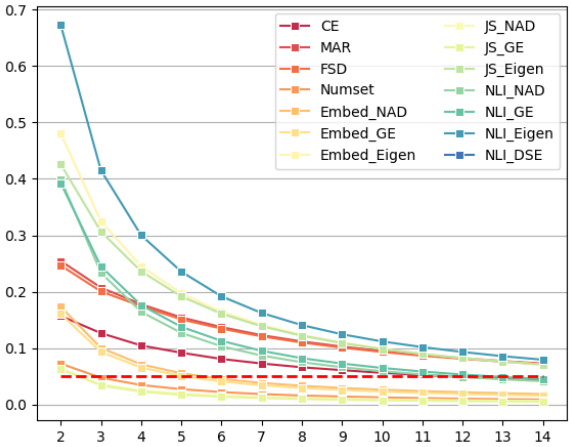

(a) Diffratio.

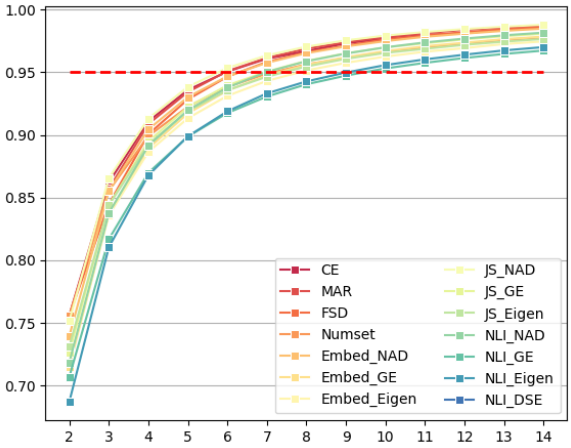

(b) Spearmanr.

Figure 12: Aggregated stability metrics, DiffRatio and Spearmanr, as functions of the number of repeats across different UQ methods, models, questions, and generation strategies.