

# A Survey of Uncertainty Estimation Methods on Large Language Models

Zhiqiu Xia, Jinxuan Xu, Yuqian Zhang, Hang Liu

Rutgers, The State University of New Jersey

{zhiqiu.xia, jinxuan.xu, yqz.zhang, hang.liu}@rutgers.edu

## Abstract

Large language models (LLMs) have demonstrated remarkable capabilities across various tasks. However, these models could offer biased, hallucinated, or non-factual responses camouflaged by their fluency and realistic appearance. Uncertainty estimation is the key method to address this challenge. While research efforts in uncertainty estimation are ramping up, there is a lack of *comprehensive* and *dedicated* surveys on LLM uncertainty estimation. This survey presents four major avenues of LLM uncertainty estimation. Furthermore, we perform extensive experimental evaluations across multiple methods and datasets. At last, we provide critical and promising future directions for LLM uncertainty estimation.

## 1 Introduction

Large Language Models (LLMs) have emerged as state-of-the-art solutions for a wide range of problems, mainly due to their unparalleled ability to generate coherent and contextually appropriate responses to diverse user prompts (Ouyang et al., 2022; Zhao et al., 2024). However, with the increasing adoption of LLMs, concerns have grown regarding their tendency to produce biased, hallucinated, non-factual, and misaligned outputs (Zhang et al., 2023b; Huang et al., 2024b). These issues are further exacerbated by the fact that such flawed responses often appear highly fluent and convincingly realistic, making them difficult to detect. A promising approach to addressing the challenge of misleading yet plausible responses is uncertainty estimation, which assigns an uncertainty or confidence score to the model's output. Figure 1 provides an overview of this process. First, the LLM generates an initial response based on the input. Next, a confidence score is computed for this response. The score is then evaluated against a pre-defined threshold to determine the final output. If

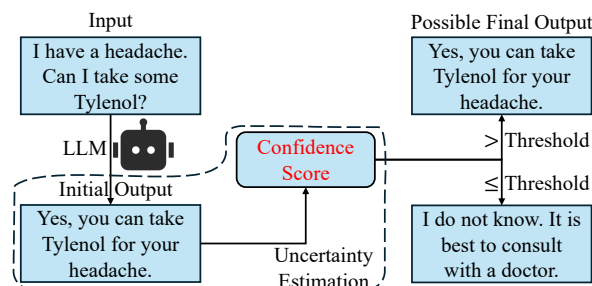


Figure 1: Illustration of uncertainty estimation.

the confidence score meets or exceeds the threshold, the initial response is accepted; otherwise, the model outputs "I do not know," thereby reducing the risk of providing incorrect but convincingly realistic information to users.

There is an urgent need for a comprehensive survey on LLM uncertainty estimation. Below, we highlight three of them: (i) Although uncertainty estimation has been extensively studied in traditional deep neural networks (DNNs)—with Bayesian and ensemble methods being notable examples (Gawlikowski et al., 2023)—these techniques are not easily transferable to LLMs, due to the large number of parameters in LLMs. (ii) LLMs significantly transform society, creating a strong demand for a thorough study of uncertainty estimation tailored to LLMs. While numerous uncertainty estimation methods exist, recent benchmark studies (Fadeeva et al., 2023; Vashurin et al., 2025) have focused on their empirical evaluation rather than in-depth methodological discussions. Consequently, a survey of recent advances in LLM uncertainty estimation that synthesizes recent progress is crucial, offering a solid foundation for future development in the field. (iii) Besides surveys focusing on natural language processing (Fomicheva et al., 2020; Baan et al., 2023), there are three existing surveys concentrating on LLM uncertainty estimation. However, each has notable limitations. Specifically, Huang et al. (2024a) dedicates a substantial portion

of its content to traditional DNN uncertainty estimation rather than focusing on LLMs. Geng et al. (2024) shifts its attention to uncertainty calibration and the applications of LLM uncertainty estimation, rather than providing a deep exploration of the core techniques. Similarly, Shorinwa et al. (2024) devotes much of its content to benchmarks and applications while lacking a complete view of the uncertainty estimation methods on LLMs.

This work focuses on studying the uncertainty estimation methods within the context of LLMs, introducing a new taxonomy from the perspective of LLMs. We center our scope around techniques applicable during the inference stage. We focus on the uncertainty estimation methods, excluding confidence calibration methods (Zhou et al., 2023; Detommaso et al., 2024) from our scope. Besides, we emphasize the methods that do not require additional data (Ren et al., 2023; Kumar et al., 2023; Tonolini et al., 2024) or model modifications (Huang et al., 2023a; Liu et al., 2024), ensuring the broad applicability of this survey. Furthermore, this survey conducts a thorough evaluation of representative uncertainty estimation approaches across various datasets and domains. Built on the insights from our evaluations, we postulate two interesting future directions for LLM uncertainty estimation.

## 2 Uncertainty Sources in LLM

There are two primary sources of uncertainty: aleatoric and epistemic uncertainties (Kendall and Gal, 2017; Hüllermeier and Waegeman, 2021). In the context of LLMs (Gao et al., 2024; Ahdriz et al., 2024; Hou et al., 2024), these sources manifest in the following ways:

- **Aleatoric uncertainty** refers to the uncertainty inherent in the data. For LLMs, this arises from ambiguous or incomplete information and inherent properties of natural language itself. Examples include vague or contextually dependent prompts, as well as linguistic phenomena where multiple valid interpretations or responses naturally coexist.
- **Epistemic uncertainty** reflects the model’s lack of knowledge or understanding. In LLMs, this occurs when the model encounters unfamiliar concepts or data that are underrepresented in its training set. This type of uncertainty can potentially be reduced by improving the training datasets and models.

## 3 Uncertainty Estimation in LLMs

### 3.1 Problem Definition and Overview

**Token generation in LLMs.** LLMs output responses in an auto-regressive manner, predicting the probability distribution of the next token given the prompt and the previously generated tokens. We denote the model as  $f$ , the prompt as  $x$ , and the generated response (or the answer) as  $r$ , which consists of  $N$  tokens, denoted as  $\{z_1, z_2, z_3, \dots, z_N\}$ . The tokens can be either words, subwords, or characters from a predefined vocabulary  $\mathcal{Z}$ . At each step of token generation, the model computes the conditional probability distribution over the vocabulary for the next token, based on the prompt  $x$  and all previously generated tokens  $r_{<i} = \{z_1, z_2, \dots, z_{i-1}\}$ . The probability distribution for the  $i$ -th token is given by  $p_i = \text{Softmax}(f(x, r_{<i}))$ . Here,  $p_i$  is a vector of length  $|\mathcal{Z}|$ , with each entry representing the probability of a specific token in  $\mathcal{Z}$  being chosen as the next token. It allows strategies such as sampling or beam search to choose from these token candidates according to their probabilities. Such an auto-regressive process ends when # of generated tokens reaches a preset number or LLM generates the end-of-sequence (EOS) token.




Figure 2: Illustration of uncertainty versus confidence.

It is important to note that uncertainty is the innate nature of LLMs, regardless of whether we estimate it. Now, we provide an intuitive understanding of uncertainty and how to estimate it.

### How to estimate uncertainty and confidence?

Following the conceptualization by Lin et al. (2024), we illustrate the process as shown in Figure 2. For each input  $x$ , an LLM model has an underlying response distribution for it (1). For ease of illustration, we assume the distribution is a normal distribution  $N(\mu, \sigma^2)$ . Uncertainty estimation is to estimate the underlying variance  $\sigma^2$ . For example, the sample variance of  $M$  different responses  $r_1, \dots, r_M$  (2) can be an estimator for

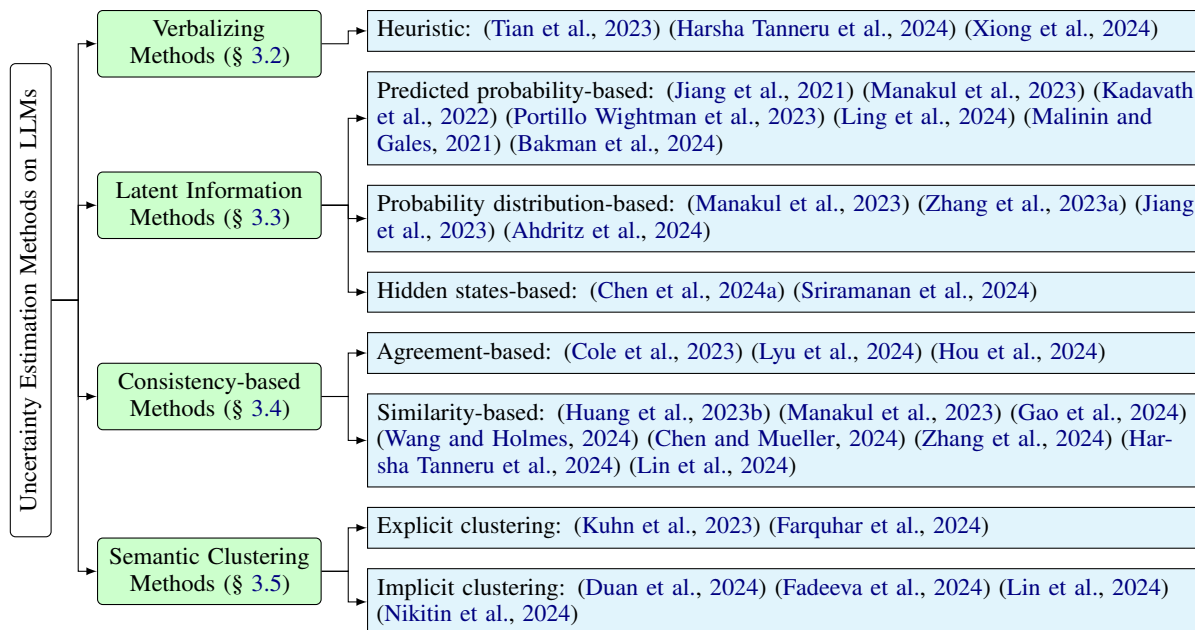


Figure 3: Taxonomy of uncertainty estimation methods on LLMs.

the variance, which indicates the variations of responses ( $U$  in Figure 2).

There are generally two types of confidence, i.e., overall confidence  $C$  and the confidence  $C_i$  associated with each response candidate  $r_i$ . The overall confidence  $C$  is complementary to  $U$ , i.e., the precision  $1/\sigma^2$  of the distribution is a confidence  $C$  to the input. The associated confidence is related to  $x$  and the tokens in a specific response  $r_i$ . To provide the final response to answer the input  $x$  given sampled responses, some literature resort to majority voting to select the most-voted response  $r_*$  (3) (Wang et al., 2023), while others choose to generate one extra response  $r_{M+1}$  with low-temperature settings (4) (Farquhar et al., 2024).

**Survey papers overview.** Figure 3 categorizes all the uncertainty estimation papers for LLM into four classes: verbalizing methods, latent information methods, consistency-based methods, and semantic clustering methods. We review each of them through Sections 3.2 - 3.5.

### 3.2 Verbalizing Methods




Figure 4: Illustration on verbalizing methods.

Figure 4 demonstrates the main workflow of ver-

balizing methods. Firstly, the input is transformed into an engineered prompt that explicitly asks the model to provide both an answer and its confidence level (1). Secondly, the LLM processes this prompt and generates an output that includes the answer and a verbalized confidence score (2), representing its self-assessed certainty about the correctness of its response.

Lin et al. (2022a) pioneers this cohort of efforts. As the capabilities of LLMs continue to develop, they can provide reasonable confidence under proper guidance, even without fine-tuning. Subsequently, Tian et al. (2023) proposes three verbalizing variants: (i) Generate multiple response candidates with confidence scores and select the highest-rated one as the final response, (ii) derive the response and confidence through two rounds of prompt-and-answer interactions, and (iii) use words instead of numerical values to indicate the confidence. Recently, Harsha Tanneru et al. (2024) introduces two methods inspired by Chain-of-Thought (CoT) prompting. The first method requests the LLM to assign an importance score to each word in the input, while the second one prompts the LLM to provide confidence for each reasoning step in the response. Finally, LLM will offer a final confidence score for the overall response. Beyond that, Xiong et al. (2024) presents a systematic framework for verbalizing methods with three parts: prompting, sampling, and aggregation. It employs specific confidence-eliciting prompts

and generates diverse response samples containing confidence scores. After that, the final confidence score is derived through inter-sample agreement or response ranking information.

While verbalizing methods offer intuitive and straightforward uncertainty estimation, they face significant limitations. Kadavath et al. (2022) shows that LLMs tend to be over-confident in their answers as the reinforcement learning from human feedback (RLHF) nature pushes LLMs to do so.

### 3.3 Latent Information Methods

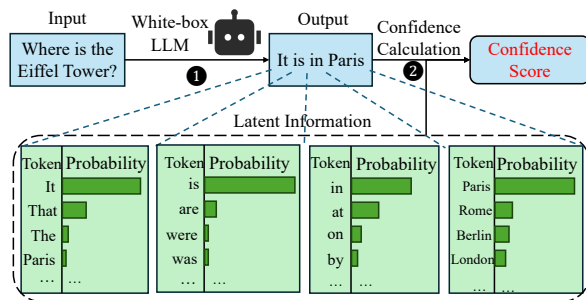


Figure 5: Illustration on latent information methods.

Figure 5 illustrates the concept of latent information methods. First of all, the LLM is prompted to provide an output to the input (❶). Of note, latent information methods require a white-box LLM, which offers latent information in the output, such as the full probability distribution over each generated token. Subsequently, this method leverages the generated information to estimate the uncertainty/confidence score via specific metrics or measures (❷). We refer the readers to Section A.2 for the formula of different latent information methods.

Jiang et al. (2021) directly uses the predicted probability of the response tokens to measure the confidence score. Manakul et al. (2023) proposes to use the negative log-likelihood of the response tokens, either average or maximum across tokens, to serve as an uncertainty measure. The averaged negative log-likelihood across tokens is also known as perplexity (Ren et al., 2023). In contrast, Kadavath et al. (2022) proposes a method that prompts the model to evaluate its answers by answering true or false, using the latent probability associated with “True” as the confidence score.

The analysis of token probabilities can be extended beyond a single response for more robust uncertainty estimation. Portillo Wightman et al. (2023) proposes to average the predicted probabilities across multiple responses. Ling et al. (2024)

picks the key token from the responses and aggregates them into a distribution, and the uncertainty is from the entropy of the distribution. Kadavath et al. (2022) considers all the tokens in the responses, calculates the probability for each response using token probabilities, and measures uncertainty through the entropy of the response distribution, called predictive entropy. However, varying response lengths can introduce undesirable noise to the estimation. To address this limitation, Malinin and Gales (2021) proposes the length-normalized entropy, incorporating the response length based on predictive entropy. Furthermore, Bakman et al. (2024) proposes to replace the length normalization by assigning a weight to each token with a BERT model to consider both the sequence length and the semantic contribution of tokens.

While the methods above only require access to the probability value of the response tokens, the following papers would require access to the complete probability distributions: Manakul et al. (2023) computes the entropy of the probability distribution for each generated token, using either the mean or maximum entropy as the uncertainty. Zhang et al. (2023a) proposes an uncertainty metric that combines the sum of negative log probabilities and entropy, while also considering token importance, preceding context, and token properties such as entity type and token frequency. For multiple-choice questions, Jiang et al. (2023) presents a specialized methodology. It computes probability distributions over potential options for each response sample and aggregates these distributions to form an ensemble probability distribution for uncertainty estimation. Ahdriz et al. (2024) introduces a heuristic two-stage method. Initially, the LLM is prompted to generate multiple next-token candidates. Subsequently, through a “repeated prompt” mechanism, the model produces the next token. The final uncertainty score is then computed from the probability distribution of these next tokens.

Beyond the methods using the probability distributions of tokens in the response, some researchers utilize the hidden states of LLMs. Chen et al. (2024a) proposes to use the embeddings in the middle layer of LLMs to construct a covariance matrix for responses, which captures the correlation relationships among them. By manipulating the eigenvalues of the covariance matrix, the degree of divergence among responses can be estimated and considered an uncertainty measure. Different from this approach, Sriramanan et al. (2024)



constructs a covariance matrix within a single response, where each row corresponds to a token in the response. The log-determinant of this covariance matrix is then calculated as an uncertainty metric. Additionally, it proposes another method based on the internal components of the attention mechanism. Specifically, it calculates the sum of the log-determinants of each self-attention head’s kernel similarity map, which serves as an alternative measure of uncertainty.

### 3.4 Consistency-based Methods

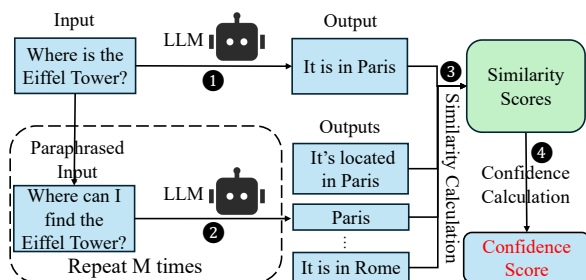


Figure 6: Illustration of consistency-based methods.

Figure 6 illustrates the workflow of consistency-based methods. First, LLM gives an output to the original input (1). Second, the input is paraphrased to maintain the same meaning as the original one but has different contents, where LLM is prompted to answer this changed input. Such process is repeated  $M$  times to generate various sampled outputs (2). Third, the similarities between the original output and each sampled output are computed (3). Finally, the confidence score is calculated based on derived similarities (4). We refer the readers to Section A.3 for detailed mathematical definitions of this consistency-based method.

The fundamental principle of consistency-based methods is that response consistency typically correlates with confidence levels, a.k.a. high response variability suggests higher uncertainty, while consistent responses indicate greater confidence.

Cole et al. (2023) introduces sampling diversity and sampling repetition. Sampling diversity quantifies the ratio of unique answers to the total number of samples, while sampling repetition measures the proportion of samples that align with the most frequent answer. Extending this framework, Lyu et al. (2024) enhances the sampling repetition metric by incorporating the most frequent and second-most frequent responses in its analysis. Hou et al. (2024) presents a more nuanced approach by introducing clarification-based uncertainty estimation. It first

generates multiple clarifications for the input and then produces responses based on these clarified inputs. The estimated uncertainty combines two parts: one from answer frequency distribution and the other from input clarification variance.

While the methods primarily focus on analyzing answer agreement patterns to estimate uncertainty, more methods emphasize evaluating the similarities among responses (3). For domain-specific tasks, targeted metrics like BLEU (Papineni et al., 2002) and CodeBLEU (Ren et al., 2020) have been successfully applied to machine translation and code generation tasks, respectively (Huang et al., 2023b). In general question-answering scenarios, token-level similarity metrics such as BERTScore (Zhang et al., 2020) and RougeL (Lin, 2004) have been widely adopted (Huang et al., 2023b; Manakul et al., 2023; Gao et al., 2024). Moving beyond token-level comparisons, more sophisticated approaches that capture semantic relationships have emerged, including SentenceBERT and NLI-based methods (Wang and Holmes, 2024; Gao et al., 2024; Chen and Mueller, 2024; Zhang et al., 2024). SentenceBERT computes the cosine similarity between two sentences using embeddings generated by the Sentence Transformer model. The NLI-based method leverages natural language inference (NLI) classifiers to categorize sentence relationships as entailment, neutral, or contradiction, regarding the probability the NLI classifier assigns to the “entailment” class as the similarity score. Moreover, Harsha Tanneru et al. (2024) proposes token importance uncertainty and CoT uncertainty. The former quantifies uncertainty through token agreement and token rank metrics, while the latter evaluates inter-step relationships using NLI classification techniques.

The generation of diverse LLM outputs in step 2 represents another critical avenue for enhancing consistency-based methods. Harsha Tanneru et al. (2024) presents two fundamental approaches: sample probing, which employs semantically equivalent prompts, and model probing, which manipulates temperature settings to introduce output stochasticity. Chen and Mueller (2024) introduces a method that modifies CoT steps specifically for prompts employing CoT techniques. Additional approaches have been proposed by Gao et al. (2024), including the strategic insertion of dummy tokens (such as newline characters and tab spaces) and modifications to system messages within prompts.

While most methods estimate confidence by av-

eraging similarities among responses in step (4), Lin et al. (2024) proposes using the number of semantic sets as a measure of uncertainty. These semantic sets are defined as “semantic equivalent” subsets, which are grouped from all responses using an NLI classifier.

Notably, consistency-based methods are computationally expensive due to the need for multiple inferences. While some approaches rely on auxiliary modules, such as an NLI classifier, their computational cost is considerably lower than that of performing an LLM inference. As a result, the computational overhead is primarily dominated by the number of inferences required.

### 3.5 Semantic Clustering Methods




Figure 7: Illustration on semantic clustering methods.

Figure 7 depicts the workflow of semantic clustering methods, which leverages both the latent information and the semantic relationships among responses to offer a more comprehensive estimation of the uncertainty. The first two steps are similar to the consistency-based methods, where the LLM generates responses to the original input and its paraphrased versions (1–2). Next, instead of calculating the similarities, the sampled outputs are partitioned into clusters with a new probability for each cluster (3). Finally, the probability distribution over these clusters calculates a confidence score (4). The motivation for semantic clustering methods stems from the limitations of consistency-based approaches, which can only deliver a perfect uncertainty score when the responses use identical wording. However, in reality, responses may convey the same meaning through varied expressions. Therefore, semantic clustering of the responses is proposed to deal with the limitations. We refer the readers to Section A.4 for the formula of different semantic clustering methods.

Kuhn et al. (2023) introduces semantic entropy for uncertainty estimation. The method comprises three phases: generation (1–2), clustering (3), and entropy estimation (4). In step 3, a bi-

directional entailment algorithm is employed to cluster semantically equivalent responses. It assesses the entailment relationship between each pair of responses, considering them to express the same meaning if they mutually entail each other. The entailment relationship can be determined with the help of an NLI classifier or by simply requesting a general-purpose LLM. The uncertainty is the entropy calculated from the cluster probabilities in step 4. In case there is no access to the token probability, Farquhar et al. (2024) introduces discrete semantic entropy, extending the work of Kuhn et al. (2023), which leverages the frequency of each cluster to calculate an entropy as the uncertainty.

While the above method clusters the responses explicitly, some methods propose implicit clustering. Duan et al. (2024) introduces sentence relevance scores between each response pair, which is more effective over long sentences than the bi-directional entail algorithm in the work of (Kuhn et al., 2023). Fadeeva et al. (2024) proposes to further alleviate the impact of claim-type uncertainty by grouping words into several claim types. Besides, it only requires one-time inference by selecting the top-k choices of words using the latent information to generate different responses.

In contrast to these latent-information-dependent techniques, some semantic clustering methods operate without requiring access to such information. Lin et al. (2024) proposes to treat generated responses as nodes and obtain the degree matrix and the graph Laplacian matrix from the pairwise similarities of the responses. Correspondingly, this method defines several uncertainty and confidence measures from the matrices. Nikitin et al. (2024) further considers the distances between the clusters. The method encodes similarities among responses via positive semidefinite unit trace kernels. It offers a more fine-grained uncertainty measure using the von Neumann entropy of these kernels.

## 4 Evaluation

### 4.1 Metrics

We use two primary metrics to evaluate the uncertainty estimation: AUROC (Area Under the Receiver Operating Characteristics curve) (Bradley, 1997) and AUARC (Area Under the Accuracy-Rejection Curve) (Nadeem et al., 2009). Both metrics range from 0 to 1, with higher scores reflecting better uncertainty estimation methods. § B contains more details about AUROC and AUARC. Of note,




Figure 8: **TruthfulQA**: ROC (left), ARC (right) curves, and AUROC and AUARC.

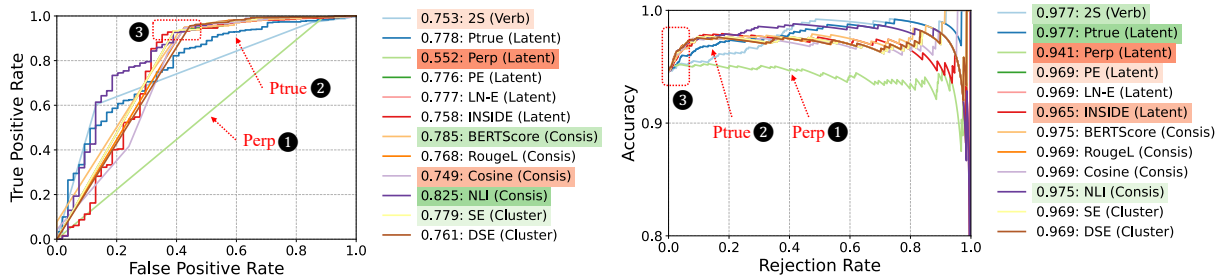


Figure 9: **SciQ**: ROC (left), ARC (right) curves, and AUROC and AUARC.

while some research (Huang et al., 2024c; Chen et al., 2024b) propose novel metrics for uncertainty estimation, they focus more on the calibration ability. In alignment with established practices in the field (Kuhn et al., 2023; Farquhar et al., 2024), our evaluation primarily emphasizes a method’s proficiency in discriminating between correct and incorrect answers based on the estimated uncertainty.

## 4.2 Evaluated Methods

We select several representative methods from each method category as follows:

- *Verbalizing methods (Verb)*: We evaluate the **2S** (Tian et al., 2023) method that asks for confidence in a second-round dialogue.
- *Latent information methods (Latent)*: We select the self-evaluation method (**Ptrue**) (Kadavath et al., 2022), perplexity (**Perp**), predictive entropy (**PE**), length-normalized entropy (**LN-E**), and the method leveraging hidden states of LLMs (**INSIDE**) (Chen et al., 2024a).
- *Consistency-based methods (Consis)*: We adopt four similarity measures: **BERTScore**, **RoughL**, cosine similarity from BERT embeddings (**Cosine**), and the “entailment” probability from an NLI classifier (**NLI**). The confidence score is averaged from similarities.
- *Semantic clustering methods (Cluster)*: We include semantic entropy (**SE**) and discrete semantic entropy (**DSE**).

## 4.3 Model Settings

We use **LLaMA3.1-8B-Instruct** (Llama Team, 2024) in our experiments. Following (Farquhar et al., 2024), we first set the temperature = 0.1 and generate an answer as the final answer. Then, we set the temperature to be 1 and generate 20 answers, which are used for methods that need extra samples. We employ the multinomial sampling as the decoding strategy and set top\_k equal to 50. Due to the varying types of questions and domains, we used the same model to determine the correctness of an answer. The prompts used are in § C.

## 4.4 Illustrative Results

Figures 8 - 12 show the ROC and ARC with the corresponding AUROC and AUARC values in the legend for five different datasets (Details about the datasets are in § E). For the AUROC and AUARC values from the legend, we color-coded the “best”, “2nd best”, “3rd best”, “3rd worst”, “2nd worst”, and “worst”.

**TruthfulQA** (Figure 8) is a benchmark designed to evaluate the truthfulness of language models in answering questions spanning 38 categories (Lin et al., 2022b). The questions in the dataset appear in a multiple-choice form, providing the LLM with clear guidance and ensuring a fixed response format. Therefore, most of the uncertainty is epistemic uncertainty. In the ROC curve, **Perp** and **INSIDE** (1) demonstrate the lowest performance, close to random guessing. The ROC curve of **2S** (2) starts with the steepest rise, indicating most responses assigned with high confidence are cor-




Figure 10: **TriviaQA**: ROC (left), ARC (right) curves, and AUROC and AUARC.




Figure 11: **GSM8K**: ROC (left), ARC (right) curves, and AUROC and AUARC.

rect. In the ARC curve, the worst-performing method (①) shows no improvement in accuracy as the rejection rate increases until the rejection rate is high. Although **2S** (②) shows a slower initial improvement, it enjoys higher improvements afterward, again demonstrating its high accuracy for high-confidence answers. **2S** achieves the best performance on this dataset, showing that LLMs can tell their uncertainty, especially when this is mainly epistemic uncertainty.

**SciQ** (Figure 9) is another multiple-choice Q&A dataset, with a collection of science-focused questions (Welbl et al., 2017). In the ROC curve, **Perp** (①) performs like random guessing (analogous to **TruthfulQA**), whereas all other methods achieve significantly better performance, including **Ptrue** (②). Most of the methods (③) achieve a very high True Positive Rate (TPR) when the False Positive Rate (FPR) approaches 0.4, indicating they assign most of the low confidence scores to negative samples correctly. As for AUARC, most methods exhibit similar performance, as the dataset is considered simple for the LLM, evidenced by a high initial accuracy of about 0.95 (③). However, the accuracy of **Perp** (①) decreases from the very beginning, resulting in the worst AUARC. In contrast, **Ptrue** (②), another variant of the latent information-based method, gains better accuracy with higher rejection rates. The difference between **Perp** and **Ptrue** shows the aggregated predicted probability of tokens is not well-calibrated, but the probability of answering the true/false of the entire response is well-calibrated.

**TriviaQA** (Figure 10) is a reading comprehension dataset where no context is provided in our settings (Joshi et al., 2017). As a free-form Q&A dataset, it allows responses to a question to vary while still expressing the same meaning. Therefore, the aleatoric uncertainty caused by language ambiguity in questions and responses exists. The ROC curve reveals that **2S** and **Perp** (①) demonstrate relatively poor performance. In contrast, **NLI** (②) achieves the highest performance. In the ARC curve, the accuracy of **2S** and **Perp** (①) deteriorates as the rejection rate increases from 0.5, while **DSE** (②) achieves the highest AUARC score.

**GSM8K** (Figure 11) is comprised of math problems that need reasoning steps to solve (Cobbe et al., 2021). The responses thus can be more diverse than **TriviaQA** due to the variability in reasoning steps. Hence, the aleatoric uncertainty is even higher. The results on AUROC demonstrate that **INSIDE** (①) performs below random guessing. On the contrary, **NLI**, **DSE**, and **SE** (②) maintain more gains on TPR with the increase of FPR. A noteworthy observation is that **NLI** and **SE** (③) achieve positive TPR even when FPR = 0 because they perfectly classify the high-confidence responses. In the ARC curve, this phenomenon is once again reflected that these methods achieve perfect accuracy when considering only the top 20% high-confidence responses (③). From the point where the rejection rate is 0, better methods exhibit faster rates of improvement (②), while the worst one (i.e., **INSIDE**) has a negative rate (①).

Comparing the **TriviaQA** and **GSM8K** datasets,



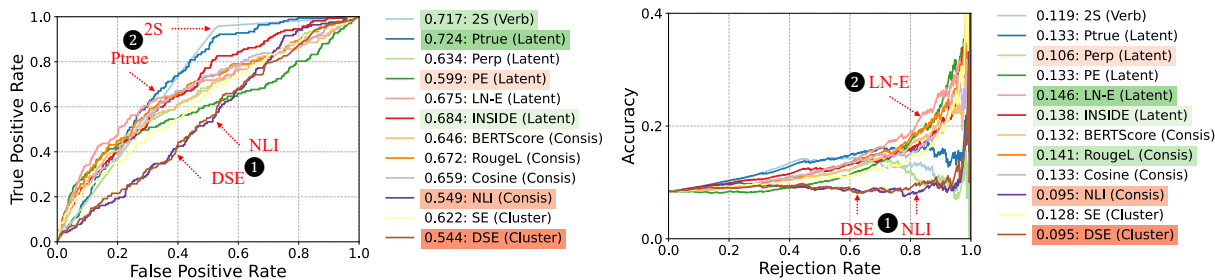


Figure 12: SimpleQA: ROC (left), ARC (right) curves, and AUROC and AUARC.

NLI, SE, and DSE perform the best on the free-form questions. They all consider the entailment relationship among responses, which can tremendously eliminate the aleatoric uncertainty and thus better estimate epistemic uncertainty. By doing so, they obtain better final results.

SimpleQA (Figure 12) is a recent Q&A dataset that presents significant challenges for state-of-the-art LLM models as of 2024 (Wei et al., 2024). Interestingly, in the ROC curve, methods (1) that traditionally demonstrate superior performance on other datasets exhibit notably poor outcomes here. 2S and Ptrue (2) emerge as the top performers, distinguished by their ability to maintain low FPR while TPR approaches 1. In the ARC curve, there is no accuracy improvement for NLI and DSE as the rejection rate increases (1). Notably, LN-E (2) becomes the highest because its accuracy continues to grow after the rejection rate passes 0.8, while others drop. Although SimpleQA is still a free-form dataset, NLI, SE, and DSE do not show their superior performance here. It shows they cannot estimate the epistemic uncertainty well if it is too big, postulating whether current benchmarks adequately evaluate LLM uncertainty estimation.

## 5 Future Directions

**Uncertainty estimation benchmark.** We need a dataset specifically designed for uncertainty estimations on LLMs. Existing datasets are designed to evaluate the capability of LLMs (not their uncertainty). They always have unambiguous questions, resulting in low aleatoric uncertainty. We anticipate three rules for designing this dataset: First, it should incorporate a diverse set of question types, including general Q&A problems, math problems, translation problems, etc. Second, the questions should have varying difficulty levels, from simple to extremely challenging. Finally, the dataset can control the degree of ambiguity for the questions to directly evaluate the uncertainty.

**Uncertainty estimation method enhancement.**

Uncertainty estimation for long responses remains under-explored. While some papers propose to break long responses into shorter segments and process each part individually (Zhang et al., 2024; Farquhar et al., 2024), they ignore the inter-sentence relationships that are critical for capturing the overall uncertainty of the response. Further, the large vocabulary in long responses challenges the effectiveness of consistency-based and semantic clustering methods. Current uncertainty estimation methods, predominantly validated on short-answer scenarios, may not adequately address the complexities inherent in longer, multi-step reasoning processes.

## 6 Conclusion

This survey paints a comprehensive landscape for uncertainty estimation methods on LLMs during the inference stage, classifying them into four classes: verbalizing, latent information, consistency-based, and semantic clustering methods. We further enrich our survey with extensive evaluations and promising future directions.

## 7 Limitations

This survey contains three limitations, mainly due to space constraints. First, we omitted detailed methodological explanations for various methods from the main text. Second, we did not evaluate and report the results of all the introduced methods. Finally, we exclude the literature that does not surround the inference stage of LLMs. We acknowledge these limitations and remain open to academic discussion and collaborative efforts to address them in future work.

## Acknowledgement

We thank the anonymous reviewers and the Area Chair for their helpful suggestions and feedback. This work was in part supported by the NSF CRII Award No. 2331536, CAREER Award No. 2326141, and NSF Awards 2212370, 2319880, 2328948, 2411294, and 2417750.

## References

- Gustaf Ahndritz, Tian Qin, Nikhil Vyas, Boaz Barak, and Benjamin L. Edelman. 2024. [Distinguishing the knowable from the unknowable with language models](#). In *Forty-first International Conference on Machine Learning*.
- Joris Baan, Nico Daheim, Evgenia Ilia, 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](#). *Preprint*, arXiv:2307.15703.
- Yavuz Faruk Bakman, Duygu Nur Yaldiz, Baturalp Buyukates, Chenyang Tao, Dimitrios Dimitriadis, and Salman Avestimehr. 2024. [MARS: Meaning-aware response scoring for uncertainty estimation in generative LLMs](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7752–7767, Bangkok, Thailand. Association for Computational Linguistics.
- Ondřej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, Radu Soricut, Lucia Specia, and Aleš Tamchyna. 2014. [Findings of the 2014 workshop on statistical machine translation](#). In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 12–58, Baltimore, Maryland, USA. Association for Computational Linguistics.
- Andrew P. Bradley. 1997. [The use of the area under the roc curve in the evaluation of machine learning algorithms](#). *Pattern Recognition*, 30(7):1145–1159.
- Chao Chen, Kai Liu, Ze Chen, Yi Gu, Yue Wu, Mingyuan Tao, Zhihang Fu, and Jieping Ye. 2024a. [INSIDE: LLMs’ internal states retain the power of hallucination detection](#). In *The Twelfth International Conference on Learning Representations*.
- Jiuhai Chen and Jonas Mueller. 2024. [Quantifying uncertainty in answers from any language model and enhancing their trustworthiness](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5186–5200, Bangkok, Thailand. Association for Computational Linguistics.
- Lihu Chen, Alexandre Perez-Lebel, Fabian M. Suchanek, and Gaël Varoquaux. 2024b. [Reconfiguring LLMs from the grouping loss perspective](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 1567–1581, Miami, Florida, USA. Association for Computational Linguistics.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.
- Jeremy Cole, Michael Zhang, Daniel 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, Singapore. Association for Computational Linguistics.
- Gianluca Detommaso, Martin Bertran, Riccardo Fogliato, and Aaron Roth. 2024. [Multicalibration for confidence scoring in llms](#). In *Proceedings of the 41st International Conference on Machine Learning, ICML’24*. JMLR.org.
- Jinhao Duan, Hao Cheng, Shiqi Wang, Alex Zavalny, Chenan Wang, Renjing Xu, Bhavya Kailkhura, and Kaidi Xu. 2024. [Shifting attention to relevance: Towards the predictive uncertainty quantification of free-form large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5050–5063, Bangkok, Thailand. Association for Computational Linguistics.
- Ekaterina Fadeeva, Aleksandr Rubashevskii, Artem Shelmanov, Sergey Petrakov, Haonan Li, Hamdy Mubarak, Evgenii Tsymbalov, Gleb Kuzmin, Alexander Panchenko, Timothy Baldwin, Preslav Nakov, and Maxim Panov. 2024. [Fact-checking the output of large language models via token-level uncertainty quantification](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 9367–9385, Bangkok, Thailand. Association for Computational Linguistics.
- Ekaterina Fadeeva, Roman Vashurin, Akim Tsvigun, Artem Vazhentsev, Sergey Petrakov, Kirill Fedyanin, Daniil Vasilev, Elizaveta Goncharova, Alexander Panchenko, Maxim Panov, Timothy Baldwin, and Artem Shelmanov. 2023. [LM-polygraph: Uncertainty estimation for language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 446–461, Singapore. Association for Computational Linguistics.
- Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. 2024. [Detecting hallucinations in large language models using semantic entropy](#). *Nature*, 630(8017):625–630.
- Marina Fomicheva, Shuo Sun, Lisa Yankovskaya, Frédéric Blain, Francisco Guzmán, Mark Fishel, Nikolaos Aletras, Vishrav Chaudhary, and Lucia Specia. 2020. [Unsupervised quality estimation for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:539–555.
- Xiang Gao, Jiaxin Zhang, Lalla Mouatadid, and Kamalika Das. 2024. [SPUQ: Perturbation-based uncertainty quantification for large language models](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2336–2346, St. Julian’s, Malta. Association for Computational Linguistics.

- Jakob Gawlikowski, Cedrique Rovile Njieutcheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, et al. 2023. A survey of uncertainty in deep neural networks. *Artificial Intelligence Review*, 56(Suppl 1):1513–1589.
- Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koepl, 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.
- Sree Harsha Tanneru, Chirag Agarwal, and Himabindu Lakkaraju. 2024. Quantifying uncertainty in natural language explanations of large language models. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pages 1072–1080. PMLR.
- Bairu Hou, Yujian Liu, Kaizhi Qian, Jacob Andreas, Shiyu Chang, and Yang Zhang. 2024. Decomposing uncertainty for large language models through input clarification ensembling. In *Forty-first International Conference on Machine Learning*.
- Hsiu-Yuan Huang, Yutong Yang, Zhaoxi Zhang, Sanwoo Lee, and Yunfang Wu. 2024a. A survey of uncertainty estimation in llms: Theory meets practice. *Preprint*, arXiv:2410.15326.
- Jiaxin Huang, Shixiang Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. 2023a. Large language models can self-improve. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1051–1068, Singapore. Association for Computational Linguistics.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2024b. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Trans. Inf. Syst.* Just Accepted.
- Xinmeng Huang, Shuo Li, Mengxin Yu, Matteo Sesia, Hamed Hassani, Insup Lee, Osbert Bastani, and Edgar Dobriban. 2024c. Uncertainty in language models: Assessment through rank-calibration. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 284–312, Miami, Florida, USA. Association for Computational Linguistics.
- Yuheng Huang, Jiayang Song, Zhijie Wang, Shengming Zhao, Huaming Chen, Felix Juefei-Xu, and Lei Ma. 2023b. Look before you leap: An exploratory study of uncertainty measurement for large language models. *Preprint*, arXiv:2307.10236.
- Eyke Hüllermeier and Willem Waegeman. 2021. Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods. *Machine Learning*, 110(3):457–506.
- Mingjian Jiang, Yangjun Ruan, Sicong Huang, Saifei Liao, Silviu Pitis, Roger Baker Grosse, and Jimmy Ba. 2023. Calibrating language models via augmented prompt ensembles. *Workshop on Challenges in Deployable Generative AI at International Conference on Machine Learning*.
- Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. 2021. How can we know when language models know? on the calibration of language models for question answering. *Transactions of the Association for Computational Linguistics*, 9:962–977.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, and et al. 2022. Language models (mostly) know what they know. *Preprint*, arXiv:2207.05221.
- Alex Kendall and Yarin Gal. 2017. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- 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*.
- Bhawesh Kumar, Charlie Lu, Gauri Gupta, Anil Palepu, David Bellamy, Ramesh Raskar, and Andrew Beam. 2023. Conformal prediction with large language models for multi-choice question answering. *Preprint*, arXiv:2305.18404.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022a. Teaching models to express their uncertainty in words. *Transactions on Machine Learning Research*.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022b. TruthfulQA: Measuring how models mimic human falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, Dublin, Ireland. Association for Computational Linguistics.



- Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. 2024. [Generating with confidence: Uncertainty quantification for black-box large language models](#). *Transactions on Machine Learning Research*.
- Chen Ling, Xujiang Zhao, Xuchao Zhang, Wei Cheng, Yanchi Liu, Yiyun Sun, Mika Oishi, Takao Osaki, Katsushi Matsuda, Jie Ji, Guangji Bai, Liang Zhao, and Haifeng Chen. 2024. [Uncertainty quantification for in-context learning of 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 3357–3370, Mexico City, Mexico. Association for Computational Linguistics.
- Linyu Liu, Yu Pan, Xiaocheng Li, and Guanting Chen. 2024. [Uncertainty estimation and quantification for llms: A simple supervised approach](#). *Preprint*, arXiv:2404.15993.
- AI @ Meta Llama Team. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- 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](#). *Preprint*, arXiv:2402.13904.
- Andrey Malinin and Mark Gales. 2021. [Uncertainty estimation in autoregressive structured prediction](#). In *International Conference on Learning Representations*.
- 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.
- Malik Sajjad Ahmed Nadeem, Jean-Daniel Zucker, and Blaise Hanczar. 2009. [Accuracy-rejection curves \(arcs\) for comparing classification methods with a reject option](#). In *Proceedings of the third International Workshop on Machine Learning in Systems Biology*, volume 8 of *Proceedings of Machine Learning Research*, pages 65–81, Ljubljana, Slovenia. PMLR.
- Alexander Nikitin, Jannik Kossen, Yarin Gal, and Pekka Martinen. 2024. [Kernel language entropy: Fine-grained uncertainty quantification for llms from semantic similarities](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 8901–8929. Curran Associates, Inc.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Gwenyth Portillo Wightman, Alexandra Delucia, and Mark Dredze. 2023. [Strength in numbers: Estimating confidence of large language models by prompt agreement](#). In *Proceedings of the 3rd Workshop on Trustworthy Natural Language Processing (TrustNLP 2023)*, pages 326–362, Toronto, Canada. Association for Computational Linguistics.
- Jie Ren, Jiaming Luo, Yao Zhao, Kundan Krishna, Mohammad Saleh, Balaji Lakshminarayanan, and Peter J Liu. 2023. [Out-of-distribution detection and selective generation for conditional language models](#). In *The Eleventh International Conference on Learning Representations*.
- Shuo Ren, Daya Guo, Shuai Lu, Long Zhou, Shujie Liu, Duyu Tang, Neel Sundaresan, Ming Zhou, Ambrosio Blanco, and Shuai Ma. 2020. [Codebleu: a method for automatic evaluation of code synthesis](#). *Preprint*, arXiv:2009.10297.
- Ola Shorinwa, Zhiting Mei, Justin Lidard, Allen Z. Ren, and Anirudha Majumdar. 2024. [A survey on uncertainty quantification of large language models: Taxonomy, open research challenges, and future directions](#). *Preprint*, arXiv:2412.05563.
- Gaurang Sriramanan, Siddhant Bharti, Vinu Sankar Sadasivan, Shoumik Saha, Priyatham Kattakinda, and Soheil Feizi. 2024. [Llm-check: Investigating detection of hallucinations in large language models](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 34188–34216. Curran Associates, Inc.
- Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher 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, Singapore. Association for Computational Linguistics.
- Francesco Tonolini, Nikolaos Aletras, Jordan Massiah, and Gabriella Kazai. 2024. [Bayesian prompt ensembles: Model uncertainty estimation for black-box large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 12229–12272, Bangkok, Thailand. Association for Computational Linguistics.



Roman Vashurin, Ekaterina Fadeeva, Artem Vazhentsev, Lyudmila Rvanova, Akim Tsvigun, Daniil Vasilev, Rui Xing, Abdelrahman Boda Sadallah, Kirill Grishchenkov, Sergey Petrakov, Alexander Panchenko, Timothy Baldwin, Preslav Nakov, Maxim Panov, and Artem Shelmanov. 2025. [Benchmarking uncertainty quantification methods for large language models with lm-polygraph](#). *Preprint*, arXiv:2406.15627.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#). In *The Eleventh International Conference on Learning Representations*.

Ziyu Wang and Chris Holmes. 2024. [On subjective uncertainty quantification and calibration in natural language generation](#). *Preprint*, arXiv:2406.05213.

Jason Wei, Nguyen Karina, Hyung Won Chung, Yunxin Joy Jiao, Spencer Papay, Amelia Glaese, John Schulman, and William Fedus. 2024. [Measuring short-form factuality in large language models](#). *Preprint*, arXiv:2411.04368.

Johannes Welbl, Nelson F. Liu, and Matt Gardner. 2017. [Crowdsourcing multiple choice science questions](#). In *Proceedings of the 3rd Workshop on Noisy User-generated Text*, pages 94–106, Copenhagen, Denmark. Association for Computational Linguistics.

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*.

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, Miami, Florida, USA. Association for Computational Linguistics.

Rui Zhang and Joel Tetreault. 2019. [This email could save your life: Introducing the task of email subject line generation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 446–456, Florence, Italy. Association for Computational Linguistics.

Tianhang Zhang, Lin Qiu, Qipeng Guo, Cheng Deng, Yue Zhang, Zheng Zhang, Chenghu Zhou, Xinbing Wang, and Luoyi Fu. 2023a. [Enhancing uncertainty-based hallucination detection with stronger focus](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 915–932, Singapore. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.

Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. 2023b. [Siren’s song in the ai ocean: A survey on hallucination in large language models](#). *Preprint*, arXiv:2309.01219.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2024. [A survey of large language models](#). *Preprint*, arXiv:2303.18223.

Kaitlyn Zhou, Dan Jurafsky, and Tatsunori Hashimoto. 2023. [Navigating the grey area: How expressions of uncertainty and overconfidence affect language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5506–5524, Singapore. Association for Computational Linguistics.

## A Mathematical Formulation of the Methods

### A.1 Common notations

We list some common notations in Table 1 for mathematical definitions.

Notation	Description
$f$	Large language model
$x$	Input
$r_i$	The $i$ -th sampled response
$r_*$	The most-voted response from samples
$N$	Number of tokens in a response
$M$	Number of sampled responses
$R_i$	The $i$ -th response cluster
$K$	Number of response clusters
$z_i$	the $i$ -th token in a response
$Z$	Vocabulary of the large language model
$r_{<i}$	All tokens before the $i$ -th token
$p_i$	The probability distribution for the $i$ -th token
$p_{z_i}$	The probability for the token $z_i$
$p$	The probability of something
$a(r_i, r_j)$	Similarity score between $r_i$ and $r_j$
$U$	Estimated uncertainty
$C$	Estimated overall confidence score
$C_i$	Estimated confidence score for response $r_i$

Table 1: Common notations and descriptions.

### A.2 Latent Information Methods

**Average over negative logarithm likelihood (Perplexity) (Manakul et al., 2023; Ren et al., 2023):**

$$U = -\frac{1}{|r|} \sum_{i=1}^N \log p_{z_i}$$

**Maximum over negative logarithm likelihood** (Manakul et al., 2023):

$$U = \max_i (-\log p_{z_i}), \quad i \in [1, N]$$

**Ptrue** (Kadavath et al., 2022):

$$C = p(z_{true} | \mathbf{x}'),$$

where  $z_{true}$  is the token for ‘‘true’’, and  $\mathbf{x}'$  is the designed prompt to ask LLM to decide whether the answer is true or false.

**Predictive entropy** (Kadavath et al., 2022):

$$U = -\frac{1}{M} \sum_{j=1}^M \sum_{i=1}^{|r_j|} \log p_{z_i}$$

**Length-normalized entropy** (Malinin and Gales, 2021):

$$U = -\frac{1}{M} \sum_j \frac{1}{|r_j|} \sum_{i=1}^{|r_j|} \log p_{z_i}$$

**Average over tokens’ probability distributions** (Manakul et al., 2023; Ren et al., 2023):

$$U = -\frac{1}{|\mathbf{r}|} \sum_{i=1}^N \sum \mathbf{p}_i \circ \log \mathbf{p}_i,$$

where  $\circ$  is the element-wise multiplication, and the second  $\sum$  means sum over all the elements in a vector.

**Maximum over tokens’ probability distributions** (Manakul et al., 2023):

$$U = \max_i (-\sum \mathbf{p}_i \circ \log \mathbf{p}_i), \quad i \in [1, N],$$

where  $\circ$  is the element-wise multiplication, and  $\sum$  means sum over all the elements in a vector.

**INSIDE** (Chen et al., 2024a):

$$U = \frac{1}{N} \log \det(\Sigma + \alpha \mathbf{I}) = \frac{1}{N} \sum_{i=1}^N \log(\lambda_i),$$

where  $\Sigma$  is the covariance matrix,  $\alpha$  is a small regularization term,  $\mathbf{I}$  is an identity matrix, and  $\lambda_i$  is the  $i$ -th eigenvalue of the matrix  $\Sigma + \alpha \mathbf{I}$ . Specifically,

$$\Sigma = \mathbf{V} \cdot \mathbf{J}_d \cdot \mathbf{V}, \quad \mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N],$$

where  $\mathbf{v}_i$  is the representative embedding for  $r_i$ ,  $\mathbf{J}_d = \mathbf{I}_d - \frac{1}{d} \mathbf{1}_N \mathbf{1}_N^T$  represents the centering matrix, and  $d$  corresponds to the dimension of the embeddings.

### A.3 Consistency-based Methods

**Sampling diversity** (Cole et al., 2023):

$$C = 1 - \frac{K}{M}$$

**Sampling diversity** (Cole et al., 2023):

$$C = \frac{1}{M} \sum_{i=1}^M \mathbb{1}(\mathbf{r}_i = \mathbf{r}_*),$$

where  $\mathbb{1}()$  is the indicator function.

**First-second-distance-based (FSD) method** (Lyu et al., 2024):

$$C = \frac{1}{M} \sum_{i=1}^M \mathbb{1}(\mathbf{r}_i = \mathbf{r}_*) - \frac{1}{M} \sum_{i=1}^M \mathbb{1}(\mathbf{r}_i = \mathbf{r}_{**}),$$

where  $\mathbb{1}()$  is the indicator function, and  $\mathbf{r}_{**}$  denotes the second most-voted answer.

**Variation ratio (VR) (when the final response is  $\mathbf{r}_*$ )** (Huang et al., 2023b):

$$U = 1 - \frac{\sum_{i=1}^M \frac{\sum_{j=1, j \neq i}^M a(\mathbf{r}_i, \mathbf{r}_j)}{M-1}}{M}$$

**Variation ratio (VR) (when the final response is  $\mathbf{r}_{M+1}$ )** (Huang et al., 2023b):

$$U = 1 - \frac{\sum_{i=1}^M a(\mathbf{r}_i, \mathbf{r}_{M+1})}{M}$$

Based on **VR** and **VRO**, using different similarity calculation methods for  $a(\cdot, \cdot)$  can achieve different estimated uncertainty.

### A.4 Semantic Clustering Methods

**Semantic entropy** (Kuhn et al., 2023):

$$U = -\sum_{k=1}^K p(\mathbf{R}_k) \log p(\mathbf{R}_k),$$

where

$$p(\mathbf{R}_k) = \sum_{r_j \in \mathbf{R}_k} \exp\left(\frac{1}{|r_j|} \sum_{i=1}^{|r_j|} \log p_{z_i}\right)$$

**Discrete semantic entropy** (Farquhar et al., 2024):

$$U = -\sum_{k=1}^K p(\mathbf{R}_k) \log p(\mathbf{R}_k),$$

where

$$p(\mathbf{R}_k) = |\mathbf{R}_k|/K$$

Dataset	Year	Question	Answer
TruthfulQA	2021	What are some cities that get more rainfall than Seattle? A) Miami and New Orleans get more rainfall than Seattle. B) No cities get more rainfall than Seattle. C) Seattle gets the most rain. D) It's always raining in Seattle.	A) Miami and New Orleans get more rainfall than Seattle.
SciQ	2017	Who proposed the theory of evolution by natural selection? A) darwin B) Linnaeus C) shaw D) Scopes	A) darwin
TriviaQA	2017	Who was the man behind The Chipmunks?	David Seville
GSM8K	2021	Janet's ducks lay 16 eggs per day. She eats three for breakfast every morning and bakes muffins for her friends every day with four. She sells the remainder at the farmers' market daily for \$2 per fresh duck egg. How much in dollars does she make every day at the farmers' market?	18
SimpleQA	2024	Who received the IEEE Frank Rosenblatt Award in 2010?	Michio Sugeno

Table 2: Samples from each dataset.

## B Detailed Explanation of AUROC and AUARC

**AUROC:** For each response, we consider it as a positive sample (correct) or a negative sample (incorrect) based on whether it matches the ground-truth label. The ROC curve is then created by plotting the true positive rate (TPR) against the false positive rate (FPR). To derive TPRs and FPRs, the accepted confidence threshold is changed to get different Predicted Positives and Negatives (i.e., PP and PN), where a response with confidence higher than the threshold is regarded as PP or PN otherwise. The AUROC is the area under the ROC curve, measuring the discriminability of confidence scores to distinguish between correct and false responses.

**AUARC:** Accuracy-Rejection Curve (RAC) is specifically designed for uncertainty estimation, which plots how the accuracy on the accepted samples changes as more low-confidence answers are rejected. The area under it indicates the uncertainty estimation's ability to maintain high accuracy when low-confidence answers are rejected.

## C Prompts

The prompt for Q&A questions is as follows:

---

### System:

You are a highly knowledgeable assistant. Answer the following question as briefly as possible.

... (several few-shot examples)

### User:

[Question]

---

The prompt for correctness decisions is as follows:

---

### User:

We are assessing the quality of answers to the following question: [Question]

The expected answer is: [Gt\_answer]

The proposed answer is: [Predicted\_answer]

Within the context of the question, does the proposed answer mean the same as the expected answer? Respond only with yes or no.

Response:

---

## D Detailed Explanation of Dataset

We give a sample for each dataset in Table 2.

**TruthfulQA** (Lin et al., 2022b) is a benchmark designed to evaluate the truthfulness of language models in generating answers to questions. It consists of 817 questions spanning 38 diverse categories such as health, law, finance, and politics. The dataset is intentionally crafted with questions that humans may answer falsely due to misconceptions or false beliefs

Dataset	Metric	2S	Ptrue	Perp	PE	LN-E	INSIDE	BERTScore	RougeL	Cosine	NLI	SE	DSE
AESLC	AUROC	0.530	0.526	0.506	0.585	0.593	0.578	0.576	0.582	0.593	0.543	0.584	0.554
AESLC	AUARC	0.383	0.387	0.358	0.435	0.435	0.430	0.421	0.426	0.439	0.406	0.426	0.403
WMT14 De-En	AUROC	0.529	0.688	0.479	0.490	0.636	0.502	0.624	0.621	0.662	0.613	0.639	0.568
WMT14 De-En	AUARC	0.785	0.881	0.789	0.797	0.855	0.800	0.847	0.846	0.866	0.832	0.855	0.806

Table 3: Supplementary Results on AESLC and WMT14 De-En.

**SciQ** (Welbl et al., 2017) is a dataset with 13,7K multiple-choice science questions spanning topics such as biology, chemistry, earth science, and physics. We chose to test our method using its validation set, which contains 1K samples.

**TriviaQA** (Joshi et al., 2017) is a large-scale reading comprehension benchmark containing over 650K question-answer-evidence triples, designed to challenge models with complex, compositional questions and diverse evidence sources. In our experimental setup, we do not provide context to the LLM but directly ask it the questions. We selected 2K samples from the validation set for testing.

**GSM8K** (Cobbe et al., 2021) is a dataset of 8.5K high-quality linguistically diverse grade school math word problems. Each problem requires 2 to 8 steps to solve, using elementary arithmetic operations (+, −, ×, ÷). In our experiments, we included reasoning steps in the examples provided in the prompts, and we used their test dataset, which consists of 1.32K samples.

**SimpleQA** (Wei et al., 2024) is a benchmark consisting of 4,326 short, fact-seeking questions designed to evaluate the factual accuracy of large language models. It covers a diverse range of topics, including science, politics, art, and so on. The Latest LLMs showed poor accuracy and calibration results on this result. We used 2K samples from the dataset for testing.

## E Supplementary Results

We supplement our experiments by including two additional datasets to evaluate summarization (AESLC (Zhang and Tetreault, 2019)) and machine translation (WMT14 De-En dataset (Bojar et al., 2014)) tasks. The results are shown in Table 3. The results on the AESLC dataset are consistent with those observed on the SimpleQA dataset, as both exhibit high epistemic uncertainty. The aleatoric uncertainty for the machine translation task lies between that of the multiple-choice Q&A and free-form Q&A datasets. Consequently, we observe comparable performance between consistency-based methods and semantic clustering methods on this task.