

# SIMBA UQ: Similarity-Based Aggregation for Uncertainty Quantification in Large Language Models

Debarun Bhattacharjya<sup>1</sup>, Balaji Ganesan<sup>1</sup>, Junkyu Lee<sup>1</sup>, Radu Marinescu<sup>1</sup>,  
Katsiaryna Mirylenka<sup>2</sup>, Michael Glass<sup>1</sup>, Xiao Shou<sup>3</sup>

<sup>1</sup>IBM Research, <sup>2</sup>Zalando, <sup>3</sup>Baylor University

{debarunb@us, bganesa1@in, Junkyu.Lee@, radu.marinescu@ie}.ibm.com,  
katya.mirylenka@zalando.ch, mrglass@us.ibm.com, Xiao\_Shou@baylor.edu

## Abstract

When does a large language model (LLM) know what it does not know? Uncertainty quantification (UQ) provides measures of uncertainty, such as an estimate of the *confidence* in an LLM’s generated output, and is therefore increasingly recognized as a crucial component of trusted AI systems. *Black-box* UQ methods do not require access to internal model information from the generating LLM and therefore have numerous real-world advantages, such as robustness to system changes, adaptability to choice of LLM, reduced costs, and computational tractability. In this paper, we investigate the effectiveness of UQ techniques that are primarily but not necessarily entirely black-box, where the consistency between a generated output and other sampled generations is used as a proxy for confidence in its correctness. We propose a high-level non-verbalized *similarity-based aggregation* framework that subsumes a broad swath of UQ approaches suitable for complex generative tasks, as well as introduce specific novel techniques from the framework that train confidence estimation models using small training sets. Through an empirical study with datasets spanning the diverse tasks of question answering, summarization, and text-to-SQL, we demonstrate that our proposed similarity-based methods can yield better calibrated confidences than baselines.

## 1 Introduction

*Uncertainty quantification* (UQ) approaches help inform reliability of model predictions and are therefore critical for deploying large language models (LLMs). UQ refers to a broad suite of techniques that yield measures of uncertainty; here we are interested in assessing the *confidence* of an LLM’s generations for a user-specified task. While confidence is not always a well-defined quantity in the UQ literature, it typically captures some measure of faith in an LLM’s response to a query as

represented by a number between 0 and 1. In this work, we consider tasks where there is a notion of whether an LLM’s response to a user query is *correct* or not, and interpret confidence of a response as the probability that it is correct. This is in contrast to approaches that estimate the *uncertainty* of LLM generations in response to a query, which measure the variability of the output.

An important category of UQ techniques are *black-box* methods, which only assume access to the model being used without requiring information such as the model parameters or even the token log probabilities. Such techniques have numerous practical advantages as they are robust to the constantly evolving landscape of LLMs, computationally lightweight, and quickly deployable at inference time. As a result, black-box UQ has become increasingly popular for tasks such as question answering (Lin et al., 2024b; Cole et al., 2023; Manakul et al., 2023).

Much of the research on black-box UQ can be viewed as *consistency-based*, where the idea is to use the consistency between a generation and other sampled generations as a proxy for its confidence. The implicit underlying assumption behind such approaches is that when a generated response is more different from others, it is more likely to be incorrect, implying that responses that are consistently similar are more likely to be correct; this has been explored for various use cases involving self-consistency (Mitchell et al., 2022; Wang et al., 2023; Chen et al., 2024). Recent work has referred to this assumption as the *consistency hypothesis*, formalizing it through statistical tests and evaluation metrics, and empirically validating its prevalence across a suite of tasks and datasets (Xiao et al., 2025).

In this paper, we provide a fresh perspective on consistency-based approaches through the lens of aggregating pairwise similarities. Figure 1 outlines the procedure for a high-level framework that aims

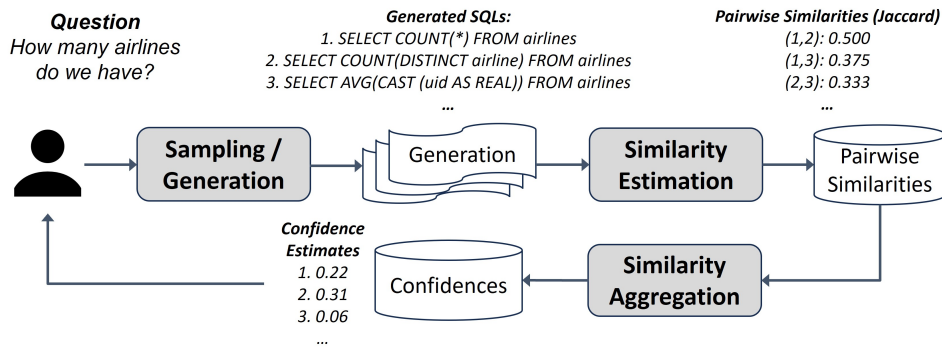


Figure 1: Framework for UQ yielding confidence estimates for generations from an LLM in response to a natural language query. An illustrative natural language query from the Spider dataset for the text-to-SQL task is shown, along with example outputs at each step of the pipeline. Jaccard is chosen as the similarity metric, and confidences are estimated using a proposed ‘aggregation by classification’ method with random forests (described later).

at exploiting the afore-mentioned consistency assumption for UQ. First, multiple outputs/samples are generated by the LLM through some sampling procedure. Pairwise similarities between samples can then be computed using any similarity metric of choice. Finally, these similarities are leveraged to provide confidence estimates for each generation of interest. Our methodological contributions are primarily in the third phase of Figure 1. In particular, rather than viewing the third phase as clustering outputs like in closely related work (Kuhn et al., 2023; Lin et al., 2024b), we propose *similarity-based aggregation* as a framework for estimating confidences. In contrast to aggregating verbalized confidences (Xiong et al., 2024), we aggregate pairwise similarities between generations, making our approach *non-verbalized* and therefore avoiding some empirically observed concerns around potential overconfidence when asking LLMs for probabilities (Hu and Levy, 2023; Xiong et al., 2024). Verbalized confidence aggregation is known to struggle for complex generative tasks.

Confidence estimates from UQ can be evaluated in various ways depending on how they will be used by the system or end user. We are primarily interested in approaches yielding confidences that are well *calibrated*, as gauged by how closely they align with the empirical accuracy of the predictions (Murphy and Epstein, 1967; Dawid, 1982). We also evaluate our proposed approaches on metrics that gauge how confidence estimates benefit when used to select from a set of generations or predict whether a generation is correct or not.

Our **contributions** are summarized as follows:

- We introduce a high-level similarity-based aggregation (SIMBA) framework, unifying and

generalizing various consistency-based UQ approaches by positioning them as different ways to represent (and possibly learn) confidence as a function of similarities with other generations.

- We propose specific novel approaches from the framework, including 1) Bayesian aggregation that uses pairwise similarities as evidence, and 2) viewing confidence estimation as classification and obtaining the probability of a generation being correct using similarities as features.
- We conduct experiments using 9 datasets – 3 each for question answering, summarization, and text-to-SQL tasks – including ablations. Similarity-based aggregation methods are shown to perform well on all chosen evaluation metrics, particularly those measuring calibration error. Importantly, results indicate that the methods perform well across short and long form generations, including when output is structured, like in SQL queries.

## 2 Related Work

Procedures for UQ estimate measures like the variability or confidence of LLM outputs, and can be categorized as either white-box or black-box. White-box methods assume access to the LLM’s internal components, such as model weights, logits, or embeddings. In contrast, black-box methods rely only on outputs, inferring confidence through alternative means. An orthogonal distinction is between verbalized and non-verbalized methods, where the former type prompts an LLM to express uncertainty in natural language, such as using phrases (“I don’t know” or “most probably”) or quantitative indicators (“low” or “50%” or “90%”).

**White-Box Methods.** Early work on calibration in deep learning and transformers laid the groundwork for using token-level probabilities – derived from model output logits – as a signal for estimating the reliability of model predictions (Kuleshov et al., 2018; Ott et al., 2018; Desai and Durrett, 2020). More recently, Kuhn et al. (2023) propose semantic entropy based clustering on multiple samples generated from the model and then estimating confidence estimates by summing the token-level probabilities in each cluster. The use of token-level probabilities for UQ has become a rapidly burgeoning area of research, driven by growing interest in the predictive confidence of LLMs (Fadeeva et al., 2023; Aichberger et al., 2024; Lin et al., 2024a; Vazhentsev et al., 2025). Kadavath et al. (2022) suggest a verbalized method where the LLM generates responses and then evaluates them as either True/False; the probability that the model assigns to the generated token determines confidence.

Other approaches consider the LLM’s internal state such as embeddings and activation spaces. For instance, Ren et al. (2023) fit the embeddings for both inputs and outputs in the training data to a Gaussian distribution, and estimate the model’s confidence by computing the distance of the evaluated data pair from this Gaussian distribution. Some methods probe the model’s attention layers to discriminate between correct and incorrect answers (Kadavath et al., 2022; Burns et al., 2023; Li et al., 2023; Azaria and Mitchell, 2023). Although these methods provide insights into the model’s linguistic understanding, they require supervised training on specially annotated data.

**Black-Box Methods.** One strand of research considers verbalized black-box methods, such as using an LLM to evaluate the correctness of its own generated answers in a conversational agent scenario (Mielke et al., 2022). Xiong et al. (2024) conduct an empirical study on UQ for reasoning tasks, showing that LLMs tend to be overconfident when verbalizing their own confidence in the correctness of the generated answers. Other related work includes fine-tuning GPT-3 to verbalize uncertainty associated with responses (Lin et al., 2022).

Many black-box methods use similarity between multiple generations given an input question, where common choices of metrics are natural language inference scores (Kuhn et al., 2023) or embedding-based similarity (Gao et al., 2024; Farquhar et al., 2024). Such similarity metrics can be used to ex-

tend clustering algorithms for uncertainty quantification of LLMs (Kuhn et al., 2023; Ao et al., 2024; Da et al., 2024; Jiang et al., 2024). In this line of work, one assumes that inconsistency in responses correlates with incorrect or hallucinatory generations. For instance, Manakul et al. (2023) propose a simple sampling-based approach that uses consistency among generations to find potential hallucinations. Lin et al. (2024b) estimate uncertainty based on analysis of a similarity matrix between generations, such as through the sum of the eigenvalues of the graph Laplacian, the degree matrix, and the eccentricity. Farquhar et al. (2024) and Nikitin et al. (2024) quantify uncertainty in LLM outputs using semantic similarity kernels to capture fine-grained variation among responses. Recent methods have also explored combining white- and black-box UQ (Chen and Mueller, 2024; Shrivastava et al., 2023).

**Comparison with Prior Work.** Our proposed framework and corresponding approaches lie within the (consistency-based) non-verbalized UQ category and are primarily black-box (although they can be made white-box if needed, by choosing features arising from token probabilities). While there is a growing body of work on UQ methods leveraging similarity, our proposed framework generalizes beyond semantic similarity – which has been the main focus in the closest related work – by emphasizing the functional relation between similarities and confidence of generations. This is important because not all generative tasks need to be concerned about semantic similarity of output, e.g., code generation. Also, we note that our proposed supervised methods in the framework differ from concurrent efforts (Liu et al., 2024; Ulmer et al., 2024; Yaldiz et al., 2025) in that: 1) they are also applicable when only similarities (and not token probabilities) are used as features and therefore can be entirely black-box, and 2) they are light-weight and thus easily deployable in real-world systems. We compare variations of the proposed aggregation methods for our extensive experimental study.

### 3 Methodology

We propose an overarching framework as well as specific techniques that provide confidence estimates, possibly using a limited amount of training data. In this section, we first present some basic notation and assumptions, and then describe the components of the workflow depicted in Figure 1.

### 3.1 Basic Notation and Assumptions

Consider an LLM generating output  $y$  for some input query  $x$ . We assume there is an associated ground truth output  $y^*$  for input  $x$  as well as a binary reward  $r \in \{0, 1\}$  from a reward function  $r(x, y, y^*)$ . Importantly, we assume there is a way to gauge whether any particular generation (with corresponding ground truth) is correct or not, i.e. whether the reward is 1 or 0;  $Y^*(x)$  denotes the set of correct responses for  $x$ . This may not be possible in some situations, such as when humans disagree about class labels (Baan et al., 2022). For tasks such as open-ended question answering and summarization, we deem the reward to be 1 if a text similarity metric (e.g. Rouge-L) between the ground truth and generated output exceeds a pre-determined threshold; this has also been assumed by related prior work (Kuhn et al., 2023; Lin et al., 2024b). For text-to-SQL, reward is 1 if the generated and ground truth queries return the same result upon query execution on the underlying database.

### 3.2 Sampling

Consistency-based approaches begin with obtaining multiple generations  $y_1, \dots, y_m$  for an input  $x$ . In this work, we adopt a sampling approach where multiple samples are generated from multiple temperatures using the next-token probability distribution of the LLM. While there are other means of generating diverse samples (Gao et al., 2024), our focus is primarily on leveraging temperature. Varying temperature has been empirically shown to provide ample opportunity for assessing whether response variability is present, which is needed by consistency-based methods to help distinguish correct from incorrect responses in complex generations (Zhu et al., 2024; Xiao et al., 2025).

In practice, one may be interested in confidence estimates for only a subset of the generations, such as the ones most likely to be correct. Generations at higher temperatures could therefore be used merely for better estimating those at lower temperatures.

### 3.3 Computing Pairwise Similarities

After generating samples, consistency-based approaches rely on access to a similarity metric with which one can compute pairwise similarities  $s(y_i, y_j)$  for all sample pairs, assumed to lie in the interval  $[0, 1]$ . As shorthand, we denote the similarities between the  $i^{\text{th}}$  generation and other generations as  $\mathbf{s}_i = s_{i,1}, \dots, s_{i,i-1}, s_{i,i+1}, \dots, s_{i,m}$ , where

$s_{i,k} = s(y_i, y_k)$  is the similarity between samples  $y_i$  and  $y_k$ . For our experiments, we consider metrics that treat samples as general text or sets of tokens, such as the Jaccard coefficient and variations of ROUGE such as Rouge-1 and Rouge-L. We note that any similarity metric can be used, but in our experience, metrics that treat generations as sets of tokens are most suitable for consistency-based UQ. Future work could explore adapting learnable similarity functions over higher level concepts, such as those using graph neural networks for complex entity matching (Krivosheev et al., 2021).

### 3.4 Similarity Aggregation

The final phase of the pipeline relies on leveraging pairwise similarities for UQ. Recall that the underlying assumption behind consistency-based approaches is that correct generations are more similar to other generations than incorrect ones. An aggregated similarity between a generation and others therefore acts as a proxy for correctness.

We present a simple yet broad perspective on consistency-based approaches that is applicable to any generative task. Rather than clustering generations such as around semantic equivalence (Kuhn et al., 2023), the confidence for sample  $y_i$  can be estimated using a suitable aggregation:  $c_i = f(\mathbf{s}_i)$ , where  $\mathbf{s}_i$  denotes pairwise similarities between  $y_i$  and other samples. A deterministic function  $f(\cdot)$  implies that identical generations yield identical confidences for the same sample set, which is a desirable property. Furthermore,  $f(\cdot)$  should reflect the underlying hypothesis around consistency-based methods, which is that more consistency is expected for correct answers. We propose the following 3 categories, each differing in how aggregation function  $f(\cdot)$  is selected and/or learned.

**Simple Aggregation.** A simple approach is to find an aggregate distance between  $y_i$  and other generations,  $\bar{d}_i = g(d_{i,1}, \dots, d_{i,m})$  where distance  $d_{i,k} = 1 - s_{i,k}$ , and compute  $c_i = 1 - \bar{d}_i$  since the aggregate distance lies in  $[0, 1]$ . The rationale is that the consistency assumption suggests that a generation further removed from others is more likely to be incorrect. While any form of aggregation  $g(\cdot)$  is possible, we use the arithmetic mean for experiments, which simplifies as  $c_i = \bar{s}_i$ . This form of aggregation is mathematically equivalent to the spectral clustering by degree approach in Lin et al. (2024b) and therefore treated as a baseline for experiments. We show later that this performs reasonably well on some (but not all) UQ metrics.

**Bayesian Aggregation.** We propose a Bayesian form of aggregation that updates beliefs about confidence using similarities as evidence. Specifically, we compute the posterior probability of  $y_i$  being correct given similarities with other generations:

$$P(y_i \in Y^* | \mathbf{s}_i) = \frac{p_0 \prod_{k \neq i} \alpha_i}{p_0 \prod_{k \neq i} \alpha_i + (1 - p_0) \prod_{k \neq i} \beta_i},$$

where  $p_0$  denotes the prior  $P(y_i \in Y^*)$ ,  $\mathbf{s}_i$  are pairwise similarities between  $y_i$  and other generations,  $\alpha_i = P(s_{i,k} | y_i \in Y^*)$ ,  $\beta_i = P(s_{i,k} | y_i \notin Y^*)$ . The formula makes two important assumptions: 1) similarities in  $\mathbf{s}_i$  depend only on whether  $y_i$  is correct, and 2) similarities in  $\mathbf{s}_i$  are conditionally independent. The first reflects the underlying assumption about the relation between consistency and correctness, allowing for a less variable distribution if  $y_i$  is correct, while the second is made for tractability. Bayesian approaches are popular in related but different areas, such as modeling uncertainty in parameters of neural networks (Gal and Ghahramani, 2016; Maddox et al., 2019).

Note that this approach requires a small training set to learn the parameters of the probabilistic model. For experiments, we assume Beta distributions for the conditional similarity distributions; this requires 5 parameters to be learned – prior  $p_0$  and 2 parameters each for the 2 Beta distributions.

**Aggregation by Classification.** We also propose treating similarity aggregation as a classification task; specifically, we train a probabilistic classifier for whether a response is correct using supervised learning with features based on similarities. Denoting  $\mathbf{s}_i^f$  as the set of similarity features for classifying correctness, confidence for generation  $y_i$  is computed as:  $c_i = P(y_i \in Y^*) = f(\mathbf{s}_i^f)$ . This method can be generalized by also including other non-similarity features  $\mathbf{o}_i^f$ , such as the generative score from the LLM, in which case  $c_i = f(\mathbf{s}_i^f, \mathbf{o}_i^f)$ .

For our experiments, we learn the function  $f(\cdot)$  using a random forest as the probabilistic classifier, since it was observed to perform better than other methods like logistic regression. We compare variations of our proposed classification approach with different feature sets. (Details are provided later.)

This is a simple yet powerful extension of simple aggregation where the function is learned using a small training set. Both this approach and the Bayesian one are more likely to be effective when the sampling procedure for training is identical to that during test time. In practical applications, this

is reasonably straightforward to control, assuming the training and test sets are not too dissimilar, and the data with similarity features that is needed for training a classifier can be easily compiled using a small labeled dataset with ground truth responses.

## 4 Empirical Investigation

### 4.1 Experimental Setup

We summarize our experimental setup in this subsection. Note that we restrict ourselves to using representative open-source LLMs for generation.

**Datasets.** We consider the following datasets:

- **QA:** We consider the open-book dataset CoQA (Reddy et al., 2019), the closed-book dataset TriviaQA (Joshi et al., 2017), as well as the closed-book dataset Natural Questions (Kwiatkowski et al., 2019). QA is widely studied in the literature on UQ for LLMs.
- **Summarization:** For this task, we experiment with the following datasets: XSum (Narayan et al., 2018), SamSum (Gliwa et al., 2019), and CNN Dailymail (Nallapati et al., 2016). Note that summarization typically results in longer form generations as compared to QA.
- **Text-to-SQL:** We consider the popular Spider benchmark (Yu et al., 2018), Spider-Realistic (Deng et al., 2021), which is a more challenging version of Spider, and BIRD (Li et al., 2024), a recent cross-domain benchmark covering many professional domains. Text-to-SQL with LLMs is an increasingly popular area of research for LLMs, with ongoing efforts to improve generation robustness and reliability through techniques like schema linking and advanced grounding (Dragusin et al., 2025).

**Models.** For QA and summarization, we generate responses using two open-source models: **LLaMA 3.3 70B** instruct (Touvron et al., 2023) and **Granite 3.1 8B** instruct (Mishra et al., 2024) models. For text-to-SQL, we use few-shot prompting on a **Codellama 34B** instruct model (Rozière et al., 2024), a code-specialized version of Llama 2, and **Granite 34B Code** instruct (Mishra et al., 2024).

For generation, we use the following default parameters: maximum number of new tokens = 200, and for sampling-based methods, we set top-k = 20 and top-p = 0.7 when applicable. Input sequences are truncated to a maximum length of 700 tokens with padding. For consistency-based methods, we generate samples across multiple tem-

peratures ranging from 0.25 to 1.5 in increments of 0.25, and compute average log probabilities across generated tokens for use in classification-based UQ approaches.

**Evaluation Metrics.** We consider the following 3 evaluation metrics, each capturing a different facet of how confidence estimates may be utilized:

- As a performance metric, we propose *accuracy from top selection (ATS)*, which measures accuracy (fraction of correct instances) in a test set when confidences are used to select one generation from a set of generations for each query. This represents the situation where the system must provide a single response for every query and confidences are used as scores for selection among generations.
- As a calibration metric, we choose *adaptive calibrated error (ACE)*, which bins confidence estimates into probability ranges such that each bin contains the same number of data points (Nixon et al., 2019). Formally,  $ACE = \frac{1}{KB} \sum_{k=1}^K \sum_{b=1}^B |acc(b, k) - c(b, k)|$ , where  $acc(b, k)$  and  $c(b, k)$  are the accuracy and confidence of adaptive calibration bin  $b$  for class label  $k$ . We prefer using adaptive bin sizes instead of fixed bin sizes, as the latter often results in unbalanced datapoints across bins. We set the # of bins  $B = 5$  for all experiments.
- As a prediction metric, we consider the *area under the receiver operating characteristic (AUROC)*, which computes the area under the curve of the false positive rate vs. true positive rate when confidences are used as a probabilistic classifier for the correctness of generations. This is a standard metric that is widely used for evaluating confidence estimation.

**Baselines.** We consider the following baselines, many of which are state-of-the-art approaches spanning all categories of UQ in LLMs:

- *avg. log prob* computes a probability by exponentiating the avg. logit over generated tokens; the generative score for sample  $y_i$  is denoted  $p_i^g$ .
- *spec-ecc* is a spectral clustering approach for UQ that leverages a graph Laplacian matrix computed from pairwise similarities and uses eccentricity (Lin et al., 2024b).
- *p(true)* is when an LLM is asked whether a generation is either True or False and the probability of the generated token (True/False) determines confidence (Kadavath et al., 2022).
- *arith-agg* estimates a generation’s confidence as

the arithmetic mean of pairwise similarities; it is mathematically equivalent to a spectral clustering approach that uses degree (Lin et al., 2024b).

- *clf-gen* is a classification approach with generative score  $p_i^g$  (as described above) as only feature. Note that we do not consider approaches that use natural language inference for similarity (Kuhn et al., 2023; Chen and Mueller, 2024) or those requiring fine-tuning LLMs as baselines, since they are either unsuitable for structured output such as SQL or require substantial training data.

**Proposed Methods.** We consider the following proposed methods. Recall that  $\mathbf{s}_i^f$  and  $\mathbf{o}_i^f$  refer to similarity and other features respectively, for methods leveraging aggregation by classification:

- *bayes-beta* is the Bayesian agg. approach, where Beta distribution parameters are learned.
- *clf-pairs* is agg. by classification when all pairwise similarities are features:  $\mathbf{s}_i^f = \mathbf{s}_i$ ,  $\mathbf{o}_i^f = \emptyset$ .
- *clf-mean+gen* includes the generative score with mean similarity:  $\mathbf{s}_i^f = \bar{\mathbf{s}}_i$ ,  $\mathbf{o}_i^f = p_i^g$ .
- *clf-pairs+gen* includes the generative score with all pairwise similarities:  $\mathbf{s}_i^f = \mathbf{s}_i$ ,  $\mathbf{o}_i^f = p_i^g$ .

Further experimental details about datasets, hyper-parameter choices, etc. are in Appendix A.

## 4.2 Main Results

We investigate the effectiveness of our proposed classification approach using generations from 2 different models on all datasets. We use a sampling procedure that generates 5 samples each over 6 temperatures, from 0.25 to 1.5 in increments of 0.25. Evaluations are performed only on samples from the lower 3 temperatures since the higher temperatures provide generations with lower performance. This captures the realistic scenario where the user wishes to obtain confidence estimates for only those samples they will even consider. We split the data randomly into half for train/test sets, and repeat the experiment 5 times to understand variability of the results. To gauge the correctness of generations, we use a Rouge-L threshold of 0.5 and 0.3 for QA and summarization datasets respectively. Jaccard is used as a similarity metric for all methods that leverage pairwise similarity.

Table 1 compares various baseline and proposed UQ approaches for generations from 2 models for the CoQA, SamSum, and Spider datasets. All 3 evaluation metrics are considered – lower ACE and higher ATS and AUROC are preferred. Comparing the performance of each UQ method as shown in

Table 1: Comparing different UQ approaches over 3 evaluation metrics on generations from 2 models each on 3 datasets (1 per task) – CoQA, SamSum, and Spider. Each approach is marked as either black-box (BB) or white-box (WB). Proposed approaches are listed in bold, others are baselines. Jaccard is used for all similarity-based methods. Error bars are from max. and min. values over 5 runs, each with a random 50% train / 50% test split.

Eval Metrics:	ATS $\uparrow$	ACE $\downarrow$	AUROC $\uparrow$	ATS $\uparrow$	ACE $\downarrow$	AUROC $\uparrow$
<b>CoQA</b>	Llama3.3 70B			Granite3.1 8B		
<i>avg. log prob (WB)</i>	0.27 $\pm$ 0.01	0.272 $\pm$ 0.019	0.72 $\pm$ 0.02	0.06 $\pm$ 0.01	0.808 $\pm$ 0.002	0.86 $\pm$ 0.03
<i>p(true) (BB)</i>	0.19 $\pm$ 0.02	1.195 $\pm$ 0.012	0.36 $\pm$ 0.01	0.03 $\pm$ 0.00	1.004 $\pm$ 0.006	0.60 $\pm$ 0.03
<i>spec-ecc (BB)</i>	0.12 $\pm$ 0.02	0.331 $\pm$ 0.018	0.21 $\pm$ 0.02	0.06 $\pm$ 0.01	0.685 $\pm$ 0.010	0.19 $\pm$ 0.02
<i>arith-agg (BB)</i>	0.27 $\pm$ 0.01	0.272 $\pm$ 0.019	0.83 $\pm$ 0.01	0.05 $\pm$ 0.01	0.206 $\pm$ 0.003	0.71 $\pm$ 0.02
<i>clf-gen (WB)</i>	0.29 $\pm$ 0.02	0.074 $\pm$ 0.011	0.73 $\pm$ 0.02	0.06 $\pm$ 0.00	0.028 $\pm$ 0.002	0.86 $\pm$ 0.03
<i>bayes-beta (BB)</i>	<b>0.38</b> $\pm$ 0.02	0.093 $\pm$ 0.010	0.90 $\pm$ 0.01	0.17 $\pm$ 0.01	0.028 $\pm$ 0.001	<b>0.95</b> $\pm$ 0.02
<i>clf-pairs (BB)</i>	0.37 $\pm$ 0.02	<b>0.041</b> $\pm$ 0.012	<b>0.91</b> $\pm$ 0.01	0.18 $\pm$ 0.01	0.015 $\pm$ 0.004	<b>0.95</b> $\pm$ 0.02
<i>clf-mean+gen (WB)</i>	0.31 $\pm$ 0.01	0.047 $\pm$ 0.019	0.84 $\pm$ 0.02	0.16 $\pm$ 0.01	0.021 $\pm$ 0.003	0.89 $\pm$ 0.03
<i>clf-pairs+gen (WB)</i>	0.37 $\pm$ 0.02	0.043 $\pm$ 0.011	<b>0.91</b> $\pm$ 0.01	<b>0.19</b> $\pm$ 0.01	<b>0.014</b> $\pm$ 0.002	<b>0.95</b> $\pm$ 0.02
<b>SamSum</b>	Llama3.3 70B			Granite3.1 8B		
<i>avg. log prob (WB)</i>	0.51 $\pm$ 0.03	0.419 $\pm$ 0.025	0.57 $\pm$ 0.02	0.15 $\pm$ 0.02	0.656 $\pm$ 0.016	0.40 $\pm$ 0.01
<i>p(true) (BB)</i>	0.50 $\pm$ 0.03	1.099 $\pm$ 0.015	0.50 $\pm$ 0.00	0.46 $\pm$ 0.02	0.911 $\pm$ 0.013	0.50 $\pm$ 0.01
<i>spec-ecc (BB)</i>	0.47 $\pm$ 0.02	0.314 $\pm$ 0.023	0.41 $\pm$ 0.01	0.01 $\pm$ 0.01	0.358 $\pm$ 0.012	0.44 $\pm$ 0.01
<i>arith-agg (BB)</i>	0.53 $\pm$ 0.02	<b>0.045</b> $\pm$ 0.010	0.63 $\pm$ 0.01	0.32 $\pm$ 0.03	0.086 $\pm$ 0.008	0.72 $\pm$ 0.01
<i>clf-gen (WB)</i>	0.52 $\pm$ 0.03	0.056 $\pm$ 0.015	0.52 $\pm$ 0.03	0.28 $\pm$ 0.02	<b>0.050</b> $\pm$ 0.007	0.62 $\pm$ 0.02
<i>bayes-beta (BB)</i>	0.53 $\pm$ 0.02	0.322 $\pm$ 0.018	0.63 $\pm$ 0.01	0.32 $\pm$ 0.02	0.911 $\pm$ 0.013	0.72 $\pm$ 0.01
<i>clf-pairs (BB)</i>	<b>0.55</b> $\pm$ 0.02	0.046 $\pm$ 0.013	<b>0.65</b> $\pm$ 0.02	0.52 $\pm$ 0.02	0.060 $\pm$ 0.005	<b>0.86</b> $\pm$ 0.01
<i>clf-mean+gen (WB)</i>	0.53 $\pm$ 0.02	0.052 $\pm$ 0.020	0.62 $\pm$ 0.01	0.42 $\pm$ 0.04	0.055 $\pm$ 0.016	0.81 $\pm$ 0.02
<i>clf-pairs+gen (WB)</i>	<b>0.55</b> $\pm$ 0.02	<b>0.045</b> $\pm$ 0.015	<b>0.65</b> $\pm$ 0.02	<b>0.53</b> $\pm$ 0.02	0.061 $\pm$ 0.005	<b>0.86</b> $\pm$ 0.01
<b>Spider</b>	Codellama 34B			Granite 34B Code		
<i>avg. log prob (WB)</i>	0.28 $\pm$ 0.02	0.654 $\pm$ 0.012	0.53 $\pm$ 0.01	0.24 $\pm$ 0.03	0.632 $\pm$ 0.012	0.65 $\pm$ 0.01
<i>p(true) (BB)</i>	0.23 $\pm$ 0.00	0.784 $\pm$ 0.006	0.52 $\pm$ 0.01	0.19 $\pm$ 0.02	0.892 $\pm$ 0.003	0.63 $\pm$ 0.02
<i>spec-ecc (BB)</i>	0.01 $\pm$ 0.00	0.201 $\pm$ 0.006	0.37 $\pm$ 0.01	0.02 $\pm$ 0.00	0.551 $\pm$ 0.009	0.20 $\pm$ 0.01
<i>arith-agg (BB)</i>	0.29 $\pm$ 0.02	0.238 $\pm$ 0.009	0.68 $\pm$ 0.01	<b>0.33</b> $\pm$ 0.01	0.070 $\pm$ 0.005	0.81 $\pm$ 0.01
<i>clf-gen (WB)</i>	0.28 $\pm$ 0.02	0.050 $\pm$ 0.009	0.53 $\pm$ 0.02	0.24 $\pm$ 0.03	0.050 $\pm$ 0.008	0.64 $\pm$ 0.01
<i>bayes-beta (BB)</i>	0.29 $\pm$ 0.02	0.298 $\pm$ 0.018	0.68 $\pm$ 0.01	<b>0.33</b> $\pm$ 0.01	0.317 $\pm$ 0.022	0.80 $\pm$ 0.01
<i>clf-pairs (BB)</i>	<b>0.30</b> $\pm$ 0.01	0.036 $\pm$ 0.005	<b>0.69</b> $\pm$ 0.01	<b>0.33</b> $\pm$ 0.02	<b>0.034</b> $\pm$ 0.004	<b>0.85</b> $\pm$ 0.01
<i>clf-mean+gen (WB)</i>	0.29 $\pm$ 0.02	0.036 $\pm$ 0.005	0.67 $\pm$ 0.01	0.31 $\pm$ 0.02	0.039 $\pm$ 0.002	0.80 $\pm$ 0.01
<i>clf-pairs+gen (WB)</i>	0.29 $\pm$ 0.02	<b>0.034</b> $\pm$ 0.006	<b>0.69</b> $\pm$ 0.01	<b>0.33</b> $\pm$ 0.02	0.036 $\pm$ 0.004	<b>0.85</b> $\pm$ 0.01

the rows, separately for each model, we observe that the proposed classification approaches using similarity features are generally high performing across all metrics. Bayesian aggregation does well on ATS and AUROC but poorly on ACE, perhaps because the conditional independence assumption that was made for tractability may not actually hold in practice. The contrast with baselines is pronounced for CoQA where the proposed approaches are notably better. Computing the arithmetic mean of similarities (*arith-agg*) is reasonably strong for AUROC. Table 2 compares a smaller set of UQ approaches for generations on the ACE metric using generations from Llama-based models for the 6

other datasets – Natural Questions, TriviaQA, CNN Daily, XSum, BIRD, and Spider-Realistic. We note again that the proposed approaches perform well across datasets.

### 4.3 Ablations

We conduct ablational studies to understand the impact of some of our choices on the main results.

**Similarity Metric.** The Jaccard coefficient was used as our primary choice of similarity metric for the main experiments. Figure 2 compares 3 similarity-based UQ approaches for the ACE metric on CoQA, SamSum and Spider, where we con-

Table 2: Comparing different UQ approaches over the ACE evaluation metric on 6 datasets: Natural Questions (NQ), TriviaQA, CNN, XSum, BIRD, and Spider-Realistic. Each approach is marked as either black-box (BB) or white-box (WB). Proposed approaches are in bold, others are baselines. Jaccard is used for all similarity-based methods. Error bars are from max. and min. values over 5 runs, each with a random 50% train / 50% test split.

Dataset:	NQ	TriviaQA	CNN	XSum	BIRD	Spider-Realistic
<i>avg. log prob</i> (WB)	0.913±0.000	0.914±0.002	0.700±0.011	0.824±0.013	0.552±0.006	0.325±0.013
<i>arith-agg</i> (BB)	0.361±0.002	0.383±0.003	0.319±0.012	0.433±0.016	0.096±0.007	0.075±0.008
<i>clf-gen</i> (WB)	<b>0.001</b> ±0.000	0.004±0.001	0.052±0.013	0.038±0.017	0.078±0.002	0.063±0.003
<i>clf-pairs</i> (BB)	0.005±0.010	<b>0.003</b> ±0.001	<b>0.038</b> ±0.007	0.035±0.016	0.089±0.007	<b>0.054</b> ±0.008
<i>clf-mean+gen</i> (WB)	<b>0.001</b> ±0.002	<b>0.003</b> ±0.001	0.049±0.017	0.035±0.019	<b>0.069</b> ±0.007	0.056±0.007
<i>clf-pairs+gen</i> (WB)	0.005±0.007	<b>0.003</b> ±0.001	0.039±0.007	<b>0.034</b> ±0.016	0.085±0.009	<b>0.054</b> ±0.005




Figure 2: Effect of choice of similarity metric on the ACE metric for the CoQA, SamSum, and Spider datasets.




Figure 3: Three different ablations on the ACE metric for the CoQA dataset.

sider the Rouge-1 and Rouge-L similarity metrics in addition to Jaccard. The plots indicate that the trends generally remain the same regardless of choice of similarity metric – the proposed classification approaches are comparable to each other and perform better than *arith-agg* on ACE.

**Three Different Ablations for CoQA.** Each panel in Figure 3 performs a different ablation for the ACE metric on the CoQA dataset, where we use generations from Llama 3.3 70B and similarity-based methods use Jaccard. Recall that lower ACE means better calibration.

- Figure 3(a) explores the effect of the Rouge-L threshold, which determines whether a generation is correct. We compare 3 UQ approaches (including 2 baselines) for the ACE metric, observing that trends generally remain the same across thresholds in  $\{0.3, 0.5, 0.7\}$ .

- Figure 3(b) compares two types of classifiers – random forests and logistic regression – showing similar results.
- Figure 3(c) tracks the change in ACE for two methods that use consistency as a function of the # of samples, where the supervised method *clf-pairs* is more robust than unsupervised *arith-agg*.

**Similarity Metric and Aggregation Function for BIRD.** In Appendix B, we conduct an extensive comparison of similarity metrics and aggregation functions for the text-to-SQL BIRD benchmark, incorporating a broader set of similarity metrics including those catering specifically to SQL generations. Results indicate that text/token similarity metrics like Jaccard and Rouge-L are suitable for UQ even when the output is SQL, and classification with random forests is a high-performing aggregator regardless of choice of text/token metrics.



## 5 Conclusion

Assessing the confidence of LLM generations in response to a query is a crucial endeavor for enabling trusted AI. Real-world systems can benefit greatly from UQ modules that are flexible in terms of applicability to diverse generative tasks and models, and also reasonably performant for varied downstream applications that use confidence estimates.

We have proposed a general high-level similarity-based aggregation framework for UQ, leveraging pairwise similarities between multiple generated samples, as well as specific novel approaches within that framework. One such approach views confidence estimation as a probabilistic classification task, where the objective is to predict the correctness of a generation using similarities with other generations for the same query as features. Our methods do not rely on asking an LLM for its confidence about a generation and can be categorized as consistency-based, since they rely on using consistency between generations as a signal for confidence. Through an extensive empirical evaluation with 9 datasets spanning the tasks of QA, summarization, and text-to-SQL, we show that using similarity features results in confidence estimates that fare well on various UQ evaluation metrics, particular around minimizing calibration error.

Our data-driven methods are designed under the assumption of no data uncertainty and necessitate a small training set with samples generated in the same way for both training and testing. Since many real-world applications often provide at least a small number of in-domain samples, and our focus in this work is specifically on no-data or low-data scenarios, we have limited our experiments to in-domain settings. Investigating out-of-domain performance with such methods is important for future work as it involves challenges such as distribution shift and domain adaptation that warrant separate study.

Finally, we note that results show only moderate gains over baselines on the ATS and AUROC metrics, indicating room for improvement when using the estimated confidences to distinguish between correct and incorrect responses. Further studies are needed to understand fundamental limitations of consistency-based UQ for LLMs.

## Limitations

Confidence estimation is concerned with providing confidences for LLM generations. The consistency-

based approaches proposed in this work rely on the assumption that correct generations are more similar to other generations, compared to incorrect ones. We note that while there is empirical evidence to support this assumption along with widespread practical adoption of similar ideas, further work is required to provide theoretical justification and a more complete understanding of when the assumption may or may not hold. Thus, there are no guarantees associated with confidence estimates of individual instances and such approaches should be deployed with suitable caution.

## Ethical Statement

We recognize both the positive and negative societal impacts of LLMs, including the potential misuse of our work on uncertainty quantification to lend unwarranted credibility to model outputs. While our methods are intended to improve transparency and reliability in LLM use, we acknowledge that they could be misapplied in high-stakes contexts without proper additional safeguards. The datasets we consider are publicly available and peer-reviewed; there are no human subjects involved, and to the best of our knowledge, our work carries no direct harmful consequences. All creators and original owners of assets have been properly credited, and licenses and terms of use have been respected. We have not conducted crowdsourcing experiments or research with human participants. More broadly, we encourage continued reflection on the implications of deploying LLMs.

## Acknowledgments

We thank Nhan Pham, Kavitha Srinivas, Dharmashankar Subramanian, Long Vu, and other colleagues for their valuable comments and feedback while discussing this work.

## References

- Lukas Aichberger, Kajetan Schweighofer, Mykyta Ielanskyi, and Sepp Hochreiter. 2024. Semantically diverse language generation for uncertainty estimation in language models. *arXiv preprint arXiv:2406.04306*.
- Julien Aligon, Matteo Golfarelli, Patrick Marcel, Stefano Rizzi, and Elisa Turricchia. 2014. Similarity measures for OLAP sessions. *Knowledge and Information Systems*, 39:463–489.
- Shuang Ao, Stefan Rueger, and Advait Siddharthan. 2024. CSS: Contrastive semantic similarity for uncer-

- tainty quantification of LLMs. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*.
- Kamel Aouiche, Pierre-Emmanuel Jouve, and Jérôme Darmont. 2006. Clustering-based materialized view selection in data warehouses. In *East European Conference on Advances in Databases and Information Systems*, pages 81–95.
- Amos Azaria and Tom Mitchell. 2023. The internal state of an LLM knows when it’s lying. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 967–976.
- Joris Baan, Wilker Aziz, Barbara Plank, and Raquel Fernandez. 2022. Stop measuring calibration when humans disagree. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1892–1915.
- Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. 2023. Discovering latent knowledge in language models without supervision. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- 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.
- Xinyun Chen, Renat Aksitov, Uri Alon, Jie Ren, Kefan Xiao, Pengcheng Yin, Sushant Prakash, Charles Sutton, Xuezhi Wang, and Denny Zhou. 2024. Universal self-consistency for large language models. In *ICML 2024 Workshop on In-Context Learning*.
- Jeremy Cole, Michael Zhang, Dan Gillick, Julian Eisen-schlos, 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.
- Longchao Da, Tiejun Chen, Lu Cheng, and Hua Wei. 2024. LLM uncertainty quantification through directional entailment graph and claim level response augmentation. *arXiv preprint arXiv:2407.00994*.
- A Philip Dawid. 1982. The well-calibrated Bayesian. *Journal of the American Statistical Association*, 77(379):605–610.
- Xiang Deng, Ahmed Hassan Awadallah, Christopher Meek, Oleksandr Polozov, Huan Sun, and Matthew Richardson. 2021. Structure-grounded pretraining for text-to-sql. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Shrey Desai and Greg Durrett. 2020. Calibration of pre-trained transformers. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 295–302.
- Catalina Dragusin, Katsiaryna Mirylenka, Christoph Miksovic Czasch, Michael Glass, Nahuel Defosse, Paolo Scotton, and Thomas Gschwind. 2025. Grounding LLMs for database exploration: Intent scoping and paraphrasing for robust NL2SQL. *Proceedings of the VLDB Endowment*, page 8097.
- Ekaterina Fadeeva, Roman Vashurin, Akim Tsvigun, Artem Vazhentsev, Sergey Petrakov, Kirill Fedyanin, Daniil Vasilev, Elizaveta Goncharova, Alexander Panchenko, Maxim Panov, et al. 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.
- Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. 2024. Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017):625–630.
- Yarin Gal and Zoubin Ghahramani. 2016. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of The International Conference on Machine Learning*, pages 1050–1059.
- Xiang Gao, Jiaxin Zhang, Lalla Mouatadid, and Kamalika Das. 2024. SPUQ: Perturbation-based uncertainty quantification for large language models. In *Proceedings of the European Chapter of the Association for Computational Linguistics*, pages 2336–2346.
- Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. 2019. SAMSum Corpus: A human-annotated dialogue dataset for abstractive summarization. *CoRR*, abs/1911.12237.
- Jennifer Hu and Roger Levy. 2023. Prompting is not a substitute for probability measurements in large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5040–5060.
- Mingjian Jiang, Yangjun Yangjun, Prasanna Sattigeri, Salim Roukos, and Tatsunori Hashimoto. 2024. Graph-based uncertainty metrics for long-form language model generations. In *Annual Conference on Neural Information Processing Systems*.
- 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)*, page 1601–1611.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield Doddsand Nova DasSarma, Eli Tran-Johnson, et al. 2022. Language models (mostly) know what they know. *preprint arXiv: 2207.05221 [cs.CL]*.

- Evgeny Krivosheev, Mattia Atzeni, Katsiaryna Mirylenka, Paolo Scotton, Christoph Mikovic, and Anton Zorin. 2021. Business entity matching with siamese graph convolutional networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 16054–16056.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. In *International Conference on Learning Representations*.
- Volodymyr Kuleshov, Peter Fenner, and Stefano Ermon. 2018. Accurate uncertainties for deep learning using calibrated regression. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 2796–2804.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Jinyang Li, Binyuan Hui, Ge Qu, Jiayi Yang, Binhua Li, Bowen Li, Bailin Wang, Bowen Qin, Ruiying Geng, Nan Huo, et al. 2024. Can LLM already serve as a database interface? A big bench for large-scale database grounded text-to-SQLs. *Advances in Neural Information Processing Systems*, 36.
- Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2023. Inference-time intervention: Eliciting truthful answers from a language model. *preprint arXiv: 2306.03341 [cs.CL]*.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. Teaching models to express their uncertainty in words. *Transactions on Machine Learning Research*.
- Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. 2024a. Contextualized sequence likelihood: Enhanced confidence scores for natural language generation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 10351–10368.
- Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. 2024b. Generating with confidence: Uncertainty quantification for black-box large language models. *Transactions on Machine Learning Research*.
- Linyu Liu, Weijia Xiang, Yichi Huang, Haoran Wang, Yisen Li, Sen Zhao, Bo Yang, and Quanquan Gu. 2024. Uncertainty estimation and quantification for llms: A simple supervised approach. *arXiv preprint arXiv:2404.15993*.
- Wesley J. Maddox, Pavel Izmailov, Timur Garipov, Dmitry P. Vetrov, and Andrew Gordon Wilson. 2019. A simple baseline for Bayesian uncertainty in deep learning. In *Advances in Neural Information Processing Systems*.
- Vitor Hirota Makiyama, M Jordan Raddick, and Rafael DC Santos. 2015. Text mining applied to SQL queries: A case study for the SDSS SkyServer. In *Symposium on Information Management and Big Data*, pages 66–72.
- 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.
- Sabrina J. Mielke, Arthur Szlam, Emily Dinan, and Y-Lan Boureau. 2022. Reducing conversational agents’ overconfidence through linguistic calibration. *Transactions of the Association for Computational Linguistics*, 10:857–872.
- Mayank Mishra, Matt Stallone, Gaoyuan Zhang, Yikang Shen, Aditya Prasad, Adriana Meza Soria, Michele Merler, Parameswaran Selvam, Saptha Surendran, Shivdeep Singh, Manish Sethi, Xuan-Hong Dang, Pengyuan Li, Kun-Lung Wu, Syed Zawad, Andrew Coleman, Matthew White, Mark Lewis, Raju Pavuluri, Yan Koyfman, Boris Lublinsky, Maximilien de Bayser, Ibrahim Abdelaziz, Kinjal Basu, Mayank Agarwal, Yi Zhou, Chris Johnson, Aanchal Goyal, Hima Patel, Yousaf Shah, Petros Zerfos, Heiko Ludwig, Asim Munawar, Maxwell Crouse, Pavan Kanipathi, Shweta Salaria, Bob Calio, Sophia Wen, Seetharami Seelam, Brian Belgodere, Carlos Fonseca, Amith Singhee, Nirmal Desai, David D. Cox, Ruchir Puri, and Rameswar Panda. 2024. Granite code models: A family of open foundation models for code intelligence. *preprint arXiv: 2405.04324 [cs.CL]*.
- Eric Mitchell, Joseph Noh, Siyan Li, Will Armstrong, Ananth Agarwal, Patrick Liu, Chelsea Finn, and Christopher D Manning. 2022. Enhancing self-consistency and performance of pre-trained language models through natural language inference. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1754–1768.
- Allan H Murphy and Edward S Epstein. 1967. Verification of probabilistic predictions: A brief review. *Journal of Applied Meteorology and Climatology*, 6(5):748–755.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gülçehre, and Bing Xiang. 2016. Abstractive text summarization using sequence-to-sequence RNNs and beyond. In *Proceedings of the SIGNLL Conference on Computational Natural Language Learning*, pages 280–290.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. Don’t give me the details, just the summary! Topic-aware convolutional neural networks for extreme summarization. *CoRR*, abs/1808.08745.

- Alexander Nikitin, Mark de Jong, Bart van Merriënboer, Marie-Francine Moens, and Ivan Titov. 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.
- Jeremy Nixon, Michael W Dusenberry, Linchuan Zhang, Ghassen Jerfel, and Dustin Tran. 2019. Measuring calibration in deep learning. In *CVPR workshops*, volume 2.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. Analyzing uncertainty in neural machine translation. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 3956–3965.
- Mohammadreza Pourreza and Davood Rafiei. 2024. DIN-SQL: Decomposed in-context learning of text-to-SQL with self-correction. *Advances in Neural Information Processing Systems*, 36.
- Siva Reddy, Danqi Chen, and Christopher D Manning. 2019. CoQA: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266.
- 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*.
- 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 *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. 2024. Code llama: Open foundation models for code. *arXiv: 2308.12950 [cs.CL]*.
- Vaishnavi Shrivastava, Percy Liang, and Ananya Kumar. 2023. Llamas know what GPTs don’t show: Surrogate models for confidence estimation. *preprint arXiv: 2311.08877 [cs.CL]*.
- Xiu Tang, Sai Wu, Mingli Song, Shanshan Ying, Feifei Li, and Gang Chen. 2022. Preqr: Pre-training representation for sql understanding. In *Proceedings of the 2022 International Conference on Management of Data, SIGMOD ’22*, page 204–216.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Dennis Ulmer, Martin Gubri, Hwaran Lee, Sangdoon Yun, and Seong Oh. 2024. Calibrating large language models using their generations only. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15440–15459, Bangkok, Thailand. Association for Computational Linguistics.
- Artem Vazhentsev, Lyudmila Rvanova, Ivan Lazichny, Alexander Panchenko, Maxim Panov, Timothy Baldwin, and Artem Shelmanov. 2025. Token-level density-based uncertainty quantification methods for eliciting truthfulness of large language models. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2246–2262.
- 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 *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Quan Xiao, Debarun Bhattacharjya, Balaji Ganesan, Radu Marinescu, Katsiaryna Mirylenka, Nhan H. Pham, Michael Glass, and Junkyu Lee. 2025. The consistency hypothesis in uncertainty quantification for large language models. In *Proceedings of the Forty-First Conference on Uncertainty in Artificial Intelligence, UAI ’25*. JMLR.org.
- 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 *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Duygu Nur Yaldiz, Yavuz Faruk Bakman, Baturalp Buyukates, Chenyang Tao, Anil Ramakrishna, Dimitrios Dimitriadis, Jieyu Zhao, and Salman Avestimehr. 2025. Do not design, learn: A trainable scoring function for uncertainty estimation in generative LLMs. In *Findings of the Association for Computational Linguistics: NAACL 2025*.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, et al. 2018. Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-SQL task. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 3911–3921.
- Yuqi Zhu, Jia Li, Ge Li, Yunfei Zhao, Zhi Jin, and Hong Mei. 2024. Hot or cold? Adaptive temperature sampling for code generation with large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 437–445.

## A Experimental Details

### Dataset Details

We provide additional information about the datasets considered for experiments. For QA and summarization tasks, we sub-select the first 1000 queries in the dev/validation splits for our experimental study:

- **Question Answering Task**
  - **CoQA** (Reddy et al., 2019) is an open-book conversational QA dataset that measures the ability of machines to understand a text passage and answer a series of interconnected questions that appear in a conversation.
  - **TriviaQA** (Joshi et al., 2017) is a closed-book QA dataset that is more challenging than standard QA benchmarks as the answers for a question may not be directly obtained by span prediction and the context is very long.
  - **Natural Questions** (Kwiatkowski et al., 2019) is also a challenging closed-book QA that contains questions from real users and requires QA systems to read and comprehend an entire Wikipedia article that may or may not contain the answer to the question.
- **Summarization Task**
  - **XSum** (Narayan et al., 2018) is a popular benchmark for abstractive summarization, focusing on generating a single-sentence summary for BBC news articles, essentially capturing their essence in one concise sentence.
  - **SamSum** (Gliwa et al., 2019) is a collection of human-annotated, messenger-style conversations designed for abstractive summarization research, featuring diverse styles and topics to mimic real-life dialogues.
  - **CNN Dailymail** (Nallapati et al., 2016) is a dataset for text summarization built from human generated abstractive summary bullets generated from news stories in CNN and Daily Mail websites.
- **Text-to-SQL Task**
  - **Spider** (Yu et al., 2018) is a popular text-to-SQL benchmark covering 138 domains with 200 databases, such as academic, booking systems, and geography-related databases. The dev set has 1034 queries.
  - **Spider-Realistic** (Deng et al., 2021) is a more challenging version of the Spider dev set as it modifies the natural language queries in Spider in an attempt to reflect realistic scenarios where questions do not make explicit mention

of column names. It includes 508 queries.

- **BIRD** (BIG Bench for LaRge-scale Database Grounded Text-to-SQL Evaluation) (Li et al., 2024) is a recent cross-domain benchmark of 95 databases covering over 37 professional domains. The dev set includes 1533 queries.

### Model Details

- **QA and Summarization Tasks:** We generate responses using a **LLaMA 3.3 70B** instruct model (Touvron et al., 2023) and a **Granite 3.1 8B** instruct model (Mishra et al., 2024).
- **Text-to-SQL Tasks:** We use a few-shot **Codellama 34B** instruct model (Rozière et al., 2024), which is a code-specialized version of Llama 2, trained with 500B tokens of code and code-related data, and a few-shot **Granite 34B Code** instruct model (Mishra et al., 2024), trained on 3-4 trillion tokens sourced from 116 programming languages.

### Prompt Templates

For QA datasets, we use the ‘prompt’ field from the datasets, similar to prior work (Lin et al., 2024b). Table 3 shows the prompt template used for summarization datasets. Table 4 shows prompt templates/examples for our few-shot approach for SQL generation with the Codellama 34B model, as well as those for a baseline (*p-true*). These are illustrative and representative of prompts for other datasets for text-to-SQL.

### Details about Select Methods

We provide some additional details about select baselines and proposed approaches below:

- *spec-ecc* (Lin et al., 2024b): We apply a threshold of 0.9 to keep only the selected eigen vectors for the spectral clustering with eccentricity baseline.
- *p-true* (Kadavath et al., 2022): We prompt the LLM used for generations to provide their belief about whether a generation is True or False. An illustration of the zero-shot prompt template is shown in Table 4.
- *clf-?*: All classification approaches (including the baseline *clf-gen*) use a random forest with a maximum depth of 4 in our experiments.

### Computational Details

Our UQ experiments can be conducted on a standalone CPU machine, but we use GPU machines (typically NVIDIA A100s with more than 40GB memory) for generating samples from the various

Table 3: Prompt template for summarization.

---

**Summary Generation (Zero-Shot)**

[INST]

1. You are given an article or document. Your task is to summarize the input article in one sentence.
2. When generating your response, prioritize correctness, i.e., ensure that your response is correct given the context and user query, and that it is grounded in the context.
3. Furthermore, make sure that the response is supported by the given document or context.

[/INST]

Summarize the following document in one sentence: { }

---

Table 4: Prompt templates for few-shot SQL generation with Codellama, as well as those for the p-true baseline.

---

**SQL Generation (Few-Shot)**

[INST]

Your task is to generate a SQL query for the given question.

«SYS»

You are given the following database schema: { } The SQL query must include one or more of the tables and columns from this schema. If there is only one table, do not use an alias. For multiple tables, assign aliases such as t1, t2, and prefix each column reference with the table alias (e.g., t1.age, t2.phone). The SQL query should not contain more than one table unless required by the question. Aim for efficient queries.

«/SYS»

**Few-Shot Examples:**

Question: Show the ids and names of all documents.  
 SQL query: SELECT document\_id, document\_name FROM Documents

Question: Show the number of documents.  
 SQL query: SELECT count(\*) FROM Documents

Question: Find the name and access counts of all documents, in alphabetical order of document name.  
 SQL query: SELECT document\_name, access\_count FROM documents ORDER BY document\_name

Question: Show all document ids and number of paragraphs in each document. Order by document id.  
 SQL query:

[/INST]

---

**p(True) (Zero-Shot)**

**Instructions:**

1. You are given an input question and a generated SQL query. Determine if the SQL query is correct with respect to the question.
2. Your output must be only True or False, with no extra formatting.
3. You are given the following database schema: { }

**True or False?**

Input: { }

SQL query: { }

Output:

---

LLMs as these are large models. We also access a single machine, either with GPUs or CPUs. models through APIs hosted as a service but this is optional and the experiments can be conducted on

Table 5: Comparing similarity metrics and aggregation approaches using generations from a few-shot Codellama 34B model on the BIRD dev set. We include 6 similarity metrics (across both SQL and token/text categories), 2 evaluation metrics (ACE and AUROC), and 5 UQ techniques – 2 baselines and 3 proposed methods of Bayesian aggregation with conditional Beta distributions, and aggregation by classification using logistic regression and random forests. Error bars are from max. and min. values over 5 runs, each with a random 50%train / 50%test split.

Eval. Metric	ACE ↓					AUROC ↑				
	<i>spec-ecc</i>	<i>arith-agg</i>	<i>bayes-beta</i>	<i>clf-lr</i>	<i>clf-rf</i>	<i>spec-ecc</i>	<i>arith-agg</i>	<i>bayes-beta</i>	<i>clf-lr</i>	<i>clf-rf</i>
Makiyama	0.652 ±0.010	0.112 ±0.006	0.171 ±0.066	0.105 ±0.007	0.109 ±0.007	0.31 ±0.02	0.69 ±0.02	0.70 ±0.02	0.70 ±0.02	0.70 ±0.02
Output type	0.257 ±0.003	0.484 ±0.005	0.188 ±0.009	0.136 ±0.005	0.136 ±0.005	0.43 ±0.00	0.58 ±0.00	0.58 ±0.00	0.57 ±0.02	0.57 ±0.02
Jaccard	0.648 ±0.010	0.226 ±0.005	0.126 ±0.006	0.114 ±0.008	0.093 ±0.008	0.27 ±0.03	0.76 ±0.01	0.74 ±0.01	0.75 ±0.02	<b>0.78</b> ±0.02
Rouge-1	0.460 ±0.011	0.388 ±0.005	0.137 ±0.008	0.097 ±0.007	0.090 ±0.006	0.31 ±0.01	0.73 ±0.01	0.77 ±0.01	0.77 ±0.01	0.77 ±0.01
Rouge-L	0.505 ±0.010	0.360 ±0.006	0.149 ±0.011	0.098 ±0.007	<b>0.089</b> ±0.008	0.28 ±0.01	0.75 ±0.01	0.77 ±0.01	0.77 ±0.02	0.77 ±0.02
Sbert-cos	0.309 ±0.008	0.557 ±0.006	0.117 ±0.006	0.099 ±0.009	0.091 ±0.007	0.40 ±0.01	0.69 ±0.01	<b>0.78</b> ±0.01	0.77 ±0.01	0.77 ±0.01

## B Effect of Similarity Metric and Aggregation Technique on BIRD

We conduct a more in-depth investigation into confidence estimation using similarity-based aggregation for the more complex text-to-SQL task. Specifically, we analyze the choice of similarity metric and similarity aggregation technique using generations from a Codellama 34B model on the BIRD dataset. For this experiment, we generate 5 samples each over 6 temperatures ( $\{0.25, 0.5, \dots, 1.5\}$ ), and evaluations are performed using all 6 samples across all queries. We randomly split the data into half for train/test sets, and repeat the experiment 5 times to study variability of the results.

We expand the set of similarity metrics to also include the following:

- *Embedding-based*: We include the cosine similarity between sentence BERT (sbert) (Reimers and Gurevych, 2019) representations of the generations as an embedding-based metric that compares semantic similarity.
- *SQL-specific*: We also consider similarity metrics specific to SQL queries, such as the binary metric of whether two generations belong to the same SQL output type among 3 categories (simple/join/nested) (Pourreza and Rafiei, 2024), as well as those that rely on parsing the SQL and comparing the contents of various clauses – Aligon (Aligon et al., 2014), Aouiche (Aouiche et al., 2006), and Makiyama (Makiyama et al., 2015). Makiyama is representative and has been shown to perform well among these on a query clustering task (Tang et al., 2022).

We compare a small set of aggregation functions – 2 baselines and 3 proposed methods. For baselines, we include two spectral clustering approaches for UQ that leverage a graph Laplacian matrix computed from pairwise similarities – one that uses eccentricity (*spec-ecc*) and another that uses degree (Lin et al., 2024b); as mentioned previously, the latter is equivalent to simple aggregation using arithmetic mean (*arith-agg*). For the proposed approaches, we include Bayesian aggregation (*bayes-beta*) and classification with logistic regression (*clf-lr*) as well as random forests (*clf-rf*).

The rows in Table 5 correspond to 6 similarity metrics and the columns correspond to 5 UQ techniques with evaluations along 2 metrics – ACE and AUROC. Comparing similarity metrics, we observe that all token/text-based metrics (i.e. Jaccard, Rouge-1, Rouge-L, and Sbert-cos) generally perform well on ACE with a powerful aggregation method such as a random forest classifier. Rouge-L and sbert-cos are high performing metrics for this dataset. We also note that our proposed aggregation methods are better for calibration metrics such as ACE rather than AUROC, as sometimes they only provide marginal improvements over averaging similarities with the *arith-agg* baseline.