

DON'T THROW AWAY YOUR BEAMS: IMPROVING CONSISTENCY-BASED UNCERTAINTIES IN LLMs VIA BEAM SEARCH

Ekaterina Fadeeva¹ Maiya Goloburda² Aleksandr Rubashevskii² Roman Vashurin²

Artem Shelmanov² Preslav Nakov² Mrinmaya Sachan¹ Maxim Panov²

¹ ETH Zurich

² MBZUAI

ABSTRACT

Consistency-based methods have emerged as an effective approach to uncertainty quantification (UQ) in large language models. These methods typically rely on several generations obtained via multinomial sampling, measuring their agreement level. However, in short-form QA, multinomial sampling is prone to producing duplicates due to peaked distributions, and its stochasticity introduces considerable variance in uncertainty estimates across runs. We introduce a new family of methods that employ beam search to generate candidates for consistency-based UQ, yielding improved performance and reduced variance compared to multinomial sampling. We also provide a theoretical lower bound on the beam set probability mass under which beam search achieves a smaller error than multinomial sampling. We empirically evaluate our approach on six QA datasets and find that its consistent improvements over multinomial sampling lead to state-of-the-art UQ performance.

1 INTRODUCTION

Today, large language models (LLMs) are increasingly being adapted in various safety-critical domains, including medicine (Busch et al., 2025), education (Xing et al., 2025), and law (Shu et al., 2024). This rapid adoption has led to a growing body of work focused on the assessment of the quality and reliability of LLM outputs. An important research direction in this field is Uncertainty Quantification (UQ; Xiao & Wang, 2019; Baan et al., 2023; Xia et al., 2025), which measures the LLM’s confidence in their responses.

UQ methods can be categorized into several distinct groups. These include information-based methods that rely on token likelihoods produced by the LLM (Fomicheva et al., 2020); verbalization approaches that prompt models to provide a confidence score (Tian et al., 2023); density-based methods that utilize embeddings (Yoo et al., 2022); and last but not least, consistency-based measures that evaluate agreement between sampled outputs (Lin et al., 2024).

Consistency-based UQ methods are of particular interest due to not only their strong performance but also their applicability to black-box settings (Vashurin et al., 2025a).




Figure 1: Beam Search vs Multinomial Sampling. Sampling produces multiple identical generations resulting in noisy confidence estimate, while beam search covers top answers from LLM distribution resulting in a better confidence score.

Correspondence to: efadeeva@ethz.ch, maxim.panov@mbzuai.ac.ae

Code available at: <https://github.com/IINemo/lm-polygraph/tree/beam-uncertainty>

Moreover, in white-box settings too, it was shown that combining information-based and consistency-based methods yields state-of-the-art performance for a variety of tasks (Kuhn et al., 2023; Duan et al., 2024). A key component of these methods is sampling, which serves as a practical means of approximating the full probability space of all potential model outputs.

Most existing UQ approaches rely on multinomial sampling from the model’s output distribution. However, in short-form QA, multinomial sampling is prone to producing similar or even identical generations, due to its bias towards higher-probability tokens during decoding; see Figure 1. Furthermore, since each run produces a different set of candidate outputs, sample-based uncertainty estimates exhibit high variance, undermining their robustness. This limits their effectiveness as a representation of the full output space, especially since, for computational efficiency, studies typically rely on a small number of samples.

To address this problem, we propose computing output consistency based on samples generated using beam search. Beam search facilitates the exploration of alternative decoding paths, which in turn allows one to generate distinct candidate outputs that better capture the model’s output space in short-form QA. Our approach includes weighting beam search outputs by their probabilities rather than uniformly, thereby preventing the overrepresentation of low-probability outputs. Particularly, when beam search is employed for decoding, uncertainty estimates are obtained at essentially no additional cost. We show that replacing multinomial sampled outputs with those generated via beam search improves the robustness and accuracy of existing consistency-based methods, as well as hybrid methods relying on both output consistency and token likelihoods.

Our main **contributions** are as follows.

- We identify key limitations of existing consistency-based uncertainty quantification methods based on multinomial sampling; see Section 2.
- We propose a new family of UQ methods that employ an importance-weighted estimator of consistency-based uncertainty with beam search output candidates; see Section 3.
- We provide a distribution-free sufficient condition ensuring that the beam-weighted estimator achieves a lower error than the expected error of the multinomial sampler; see Section 3.2.
- We show that applying a beam search-based estimator to existing consistency-based UQ approaches improves their performance on short-form QA tasks, achieving state-of-the-art results; see Section 4.

2 BACKGROUND AND MOTIVATION

2.1 LANGUAGE MODEL DECODING

Autoregressive LLMs produce text sequentially, generating one token at a time. At each step i , the model samples a token $y_i \sim p(\cdot \mid \mathbf{y}_{<i}, \mathbf{x})$, where $\mathbf{y}_{<i}$ denotes the sequence of previously generated tokens. The probability of generating an output sequence $\mathbf{y}$ is:

$$p(\mathbf{y} \mid \mathbf{x}) = \prod_{i=1}^{|\mathbf{y}|} p(y_i \mid \mathbf{y}_{<i}, \mathbf{x}). \quad (1)$$

At each step, the model outputs a probability distribution over the entire vocabulary $\mathcal{V}$ conditioned on the prompt $\mathbf{x}$ and the partial sequence $\mathbf{y}_{<i}$.

Decoding strategies. Since the model defines a probability distribution, a concrete output must be obtained at inference time by applying a decoding strategy. Common decoding strategies include: (i) greedy decoding that selects maximum probability tokens at each step; (ii) multinomial sampling where tokens are drawn according to $p(y_i \mid \mathbf{y}_{<i}, \mathbf{x})$; and (iii) beam search, which maintains the top- k most likely partial sequences at each step. Several other variants of decoding approaches have been proposed, such as top- p nucleus sampling or temperature scaling (Holtzman et al., 2020; Vijayakumar et al., 2018). Each decoding strategy offers different trade-offs between output quality and diversity.

2.2 UNCERTAINTY QUANTIFICATION FOR LLMs

The objective of uncertainty quantification is to measure the level of uncertainty introduced by LLM when generating output sequence $\mathbf{y}_*$ conditioned on input sequence $\mathbf{x}$, denoted by $U(\mathbf{y}_* | \mathbf{x})$. Existing approaches to UQ can be broadly categorized into three main groups.

Information-based methods rely on a single forward pass of the model and compute statistics over the token-level probability distributions to quantify uncertainty. Examples include Sequence Probability, Mean Token Entropy, Perplexity (Fomicheva et al., 2020), and CCP (Fadeeva et al., 2024).

Reflexive methods query the model directly about its confidence in a generated answer using specially designed prompts. A representative example is $P(\text{True})$ (Kadavath et al., 2022), which measures the probability that the model outputs “True” when asked whether its generated answer $\mathbf{y}_*$ is correct.

Sampling-based methods draw multiple samples from the model’s output distribution and evaluate their semantic or lexical similarity to assess uncertainty. Lexical Similarity (Fomicheva et al., 2020) computes mean pairwise similarity between generated texts; other examples include Semantic Entropy (Kuhn et al., 2023), SAR (Duan et al., 2024), and black-box uncertainty measures from (Lin et al., 2024).




Figure 2: Mean percentage of redundant samples (i.e., outputs already seen among earlier generations) as a function of greedy output length. Results were obtained from 2,000 questions from the TriviaQA dataset using the Gemma 3 4B base model and 10 candidate generations. Redundancy is especially high for short answers, leading to wasted computation.

Consistency-based UQ methods. A notable subset of sampling-based methods is *consistency-based UQ* (Vashurin et al., 2025b). These methods estimate uncertainty *with respect to a particular generated output* $\mathbf{y}_* \sim p(\cdot | \mathbf{x})$, rather than the overall uncertainty of the model’s predictive distribution for the input $\mathbf{x}$. This distinction makes consistency-based UQ particularly suited for evaluating confidence in a specific prediction rather than overall model uncertainty, and Vashurin et al. (2025b) empirically demonstrate that such methods outperform other sampling-based approaches in practice.

Let us consider the most straightforward consistency-based method for predictive uncertainty quantification: measuring how semantically different alternative generations are from the produced answer $\mathbf{y}_*$. We refer to this score as *Dissimilarity* and formalize it as the expected semantic dissimilarity between the produced answer $\mathbf{y}_*$ and *all* potential alternatives drawn from the model:

$$U_D(\mathbf{y}_* | \mathbf{x}) = \mathbb{E}_{\mathbf{y} \sim p(\cdot | \mathbf{x})} [1 - s(\mathbf{y}, \mathbf{y}_*)]. \quad (2)$$

Here, $s(\mathbf{y}', \mathbf{y}'') \in [0, 1]$ is a function that measures semantic similarity between two generations $\mathbf{y}'$ and $\mathbf{y}''$. A higher value of $U_D(\mathbf{y}_* | \mathbf{x})$ indicates lower consistency between the chosen answer and alternative candidate outputs, and thus reflects greater predictive uncertainty.

The corresponding Monte Carlo estimator introduced by (Lin et al., 2024) draws M i.i.d. samples $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(M)} \sim p(\cdot | \mathbf{x})$ and computes uncertainty in the following way:

$$\hat{U}_D^{MC}(\mathbf{y}_* | \mathbf{x}) = \frac{1}{M} \sum_{i=1}^M (1 - s(\mathbf{y}^{(i)}, \mathbf{y}_*)). \quad (3)$$

Challenges of consistency-based UQ methods. A natural intuition is that, for consistency-based methods, samples should be generated in a distinct, high-probability, and stable manner. Most existing methods use multinomial sampling, which, especially for shorter generations and small sample sizes, does not satisfy these criteria.

Figure 2 shows the effect of multinomial sampling on the percentage of duplicates depending on the length of generations. The resulting samples contain many duplicates, with the issue being particularly pronounced for shorter generations, where 30–50% of the outputs are duplicates.

This not only contributes to wasted computation, but also leads to high variance estimates. Moreover, drawing M full generations solely for uncertainty estimation can be costly.

Thus, while multinomial sampling is widely used, it does not best serve the goals of consistency-based uncertainty estimation.

3 UNCERTAINTY QUANTIFICATION BASED ON CONSISTENCY OF BEAM SEARCH CANDIDATES

To address the problems outlined above, we propose to utilize an alternative decoding strategy for generating candidate outputs: beam search. Beam search (i) guarantees distinct candidate outputs, (ii) reduces variance (see Section 3.2) and (iii) provides uncertainty estimates essentially “for free” as the beam already provides a distribution over candidate outputs.

3.1 REPLACING MULTINOMIAL SAMPLING

A simple way to approximate dissimilarity from beam-generated candidates would be to reuse equation (3), treating the beam outputs as if they were drawn uniformly. While this offers a plausible alternative, treating the candidates produced by beam search in a uniform manner would overemphasize lower-probability outputs. To better reflect the model distribution while avoiding repeated multinomial draws, we form a probability-weighted estimator over the beam set.

For this purpose, we use beam search with width M to obtain distinct candidates $\mathcal{B}_M(\mathbf{x}) = \{\mathbf{b}^{(1)}, \dots, \mathbf{b}^{(M)}\}$ and their sequence probabilities $\{p(\mathbf{b}^{(i)} | \mathbf{x})\}_{i=1}^M$. To perform an estimation of $U_D(\mathbf{y}_* | \mathbf{x})$ in equation (2) with the help of samples $b^{(i)}$, one needs to perform importance weighting. Thus, we define the restricted (top- M) normalized masses w_i as:

$$w_i = \frac{p(\mathbf{b}^{(i)} | \mathbf{x})}{\sum_{j=1}^M p(\mathbf{b}^{(j)} | \mathbf{x})}, \quad i = 1, \dots, M. \quad (4)$$

The resulting importance-weighted estimator of equation (2) is

$$\hat{U}_D^b(\mathbf{y}_* | \mathbf{x}) = \sum_{i=1}^M w_i (1 - s(\mathbf{b}^{(i)}, \mathbf{y}_*)). \quad (5)$$

This top- M truncation introduces a small bias relative to full multinomial sampling but typically reduces variance and duplication on peaked distributions, yielding more stable estimates per unit budget. In the next section we are going to explore the benefits of beam search-based estimator $\hat{U}_D^b(\mathbf{y}_* | \mathbf{x})$ from a theoretical perspective.

3.2 THEORETICAL ANALYSIS

We compare the multinomial Monte Carlo estimator $\hat{U}_D^{MC}$ (3) with the beam-weighted estimator $\hat{U}_D^b$ (5) for the dissimilarity U_D defined in equation (2).

Theorem 1 (Comparison condition for beam-weighted and Monte Carlo estimators).

Let $\mathcal{B}_M(\mathbf{x}) = \{\mathbf{b}^{(1)}, \dots, \mathbf{b}^{(M)}\}$ be the beam set, $m_B = \sum_{i=1}^M p(\mathbf{b}^{(i)} | \mathbf{x})$ be its total probability mass, and define μ_B and $\mu_{\bar{B}}$ as dissimilarity inside and outside the beam set $\mathcal{B}_M$ correspondingly:

$$\mu_B = \mathbb{E}_{\mathbf{y} \sim p(\cdot | \mathbf{x})} [1 - s(\mathbf{y}, \mathbf{y}_*) | \mathbf{y} \in \mathcal{B}_M(\mathbf{x})], \quad \mu_{\bar{B}} = \mathbb{E}_{\mathbf{y} \sim p(\cdot | \mathbf{x})} [1 - s(\mathbf{y}, \mathbf{y}_*) | \mathbf{y} \notin \mathcal{B}_M(\mathbf{x})].$$

Then the beam-weighted estimator $\hat{U}_D^b$ achieves smaller mean-squared error than the Monte Carlo estimator $\hat{U}_D^{MC}$ whenever

$$(1 - m_B) |\mu_B - \mu_{\bar{B}}| < \sigma / \sqrt{M}, \quad (6)$$

where $\sigma^2 = \text{Var}_{\mathbf{y} \sim p(\cdot | \mathbf{x})} (1 - s(\mathbf{y}, \mathbf{y}_*))$. The corresponding distribution-free sufficient condition is

$$m_B > 1 - \frac{1}{2\sqrt{M}}. \quad (7)$$

Proof. The Monte Carlo estimator averages M i.i.d. samples $\mathbf{y}^{(i)} \sim p(\cdot | \mathbf{x})$, so it is unbiased with $\mathbb{E}[\hat{U}_D^{MC}] = U_D(\mathbf{y}_* | \mathbf{x})$ and $\text{MSE}(\hat{U}_D^{MC}) = \text{Var}(\hat{U}_D^{MC}) = \sigma^2/M$. By Popoviciu’s inequality, any random variable supported on $[0, 1]$ has variance at most $1/4$, hence $\sigma^2 \leq 1/4$.

By the law of total expectation, the true dissimilarity U_D decomposes as:

$$U_D(\mathbf{y}_* | \mathbf{x}) = m_B \mu_B + (1 - m_B) \mu_{\bar{B}}, \quad \hat{U}_D^b = \mu_B,$$

so squared error of the beam-weighted estimator $\hat{U}_D^b$ is deterministic:

$$\text{SE}(\hat{U}_D^b) = (\hat{U}_D^b - U_D)^2 = (1 - m_B)^2 (\mu_B - \mu_{\bar{B}})^2.$$

Beam-weighted estimation is therefore more accurate whenever

$$(1 - m_B)^2 (\mu_B - \mu_{\bar{B}})^2 < \sigma^2/M,$$

which yields the stated condition (6). A distribution-free bound (7) follows from $|\mu_B - \mu_{\bar{B}}| \leq 1$ and $\sigma^2 \leq 1/4$. $\square$

From Theorem 1, beam-weighted estimator is more accurate than Monte Carlo estimator whenever total beam probability mass m_B exceeds $1 - \frac{1}{2\sqrt{M}}$. For $M = 10$, the threshold is $m_B > 0.842$. Thus, when the top-10 beam hypotheses capture at least $\sim 84\%$ of the model’s probability mass, beam search provides a lower-error estimator than multinomial sampling with the same sample budget.

In practice, *this condition is frequently satisfied*. On the TriviaQA dataset, Figure 3 shows that 22.7% of examples meet the sufficient condition overall, and up to 30-40% for very short generations (≤ 3 output tokens), where probability mass is highly concentrated on the top beams. When the inside-outside gap $\delta = |\mu_B - \mu_{\bar{B}}| < 1$, the break-even requirement (6) relaxes to $(1 - m_B)\delta < \sigma/\sqrt{M}$, allowing beam search to outperform even when $m_B < 0.842$. Although $\mu_{\bar{B}}$ is not directly computable due to the combinatorial output space, our experiments consistently show beam search outperforming multinomial sampling, suggesting that δ is modest in practice and that the effective threshold is often lower than 0.842.




Figure 3: Percentage of texts meeting the sufficient condition (Theorem 1). Results are based on 2,000 TriviaQA questions, Gemma 3 4B base and $M = 10$. The green “All” bar shows the overall percentage across all lengths.

3.3 ADAPTING OTHER UQ METHODS TO BEAM SEARCH

In a similar manner, other consistency-based methods can be adapted to utilize beam search-based samples in their formulation.

Eccentricity. Eccentricity is a method introduced by Lin et al. (2024). Unlike dissimilarity, which uses only the similarities between the produced answer $\mathbf{y}_*$ and each alternative sample, Eccentricity aggregates the *joint* pairwise relationships among all samples.

In this method, we first construct a similarity matrix of size $(M + 1) \times (M + 1)$ for the M samples and the produced answer $\mathbf{y}^{(M+1)} = \mathbf{y}_*$:

$$W_{ij} = s(\mathbf{y}^{(i)}, \mathbf{y}^{(j)}), \quad 1 \leq i, j \leq M + 1. \quad (8)$$

Then we compute the degree matrix D :

$$D_{ij} = \begin{cases} \sum_{k=1}^{M+1} W_{ik}, & i = j, \\ 0, & i \neq j, \end{cases} \quad (9)$$

and obtain the eigendecomposition of the Graph Laplacian $L = I - D^{-1/2} W D^{-1/2}$, yielding eigenpairs $\{\lambda_i, \mathbf{u}_i\}_{i=1}^{M+1}$.

Smaller eigenvalues (close to zero) capture meaningful semantic structure, whereas larger eigenvalues tend to reflect noise. We therefore retain the eigenvectors whose eigenvalues satisfy $\lambda_i < \alpha$, yielding K vectors in total; K is thus determined by the threshold $\alpha > 0$.

Semantic embeddings are formed as $\mathbf{v}_j = [\mathbf{u}_{1j}, \mathbf{u}_{2j}, \dots, \mathbf{u}_{Kj}]$. For $1 \leq j \leq M$, $\mathbf{v}_j$ represents the embedding of $\mathbf{y}^{(j)}$, and $\mathbf{v}_* = \mathbf{v}_{M+1}$ corresponds to $\mathbf{y}_*$. The confidence score is the distance between the embedding of the produced answer and the mean embedding of the samples:

$$\hat{U}_{Ecc}(\mathbf{y}_* | \mathbf{x}) = \left\| \mathbf{v}_* - \frac{1}{M} \sum_{i=1}^M \mathbf{v}_i \right\|_2^2, \quad (10)$$

where higher values indicate higher uncertainty.

With beam-generated candidates, we weight embeddings by the normalized masses w_i from equation (4) to better reflect the model distribution while avoiding duplicate generations:

$$\hat{U}_{Ecc}^b(\mathbf{y}_* | \mathbf{x}) = \left\| \mathbf{v}_* - \sum_{i=1}^M w_i \mathbf{v}_i^b \right\|_2^2. \quad (11)$$

CoCoA. A white-box approach CoCoA (Vashurin et al., 2025b) combines a model probabilities-based uncertainty with the sample-consistency signal:

$$\hat{U}_{CoCoA}(\mathbf{y}_* | \mathbf{x}) = u(\mathbf{y}_* | \mathbf{x}) \cdot \hat{U}_D^M(\mathbf{y}_* | \mathbf{x}) = u(\mathbf{y}_* | \mathbf{x}) \cdot \frac{1}{M} \sum_{i=1}^M (1 - s(\mathbf{y}^{(i)}, \mathbf{y}_*)), \quad (12)$$

where $u(\mathbf{y} | \mathbf{x})$ is a model-based uncertainty measure for the sequence (e.g., $-\log p(\mathbf{y} | \mathbf{x})$).

For a beam-weighted estimator, we utilize (5) as sample-consistency signal:

$$\hat{U}_{CoCoA}^b(\mathbf{y}_* | \mathbf{x}) = u(\mathbf{y}_* | \mathbf{x}) \cdot \hat{U}_D^b(\mathbf{y}_* | \mathbf{x}). \quad (13)$$

Eigenvectors Dissimilarity. Both Dissimilarity and Eccentricity produce confidence scores for the generated answer $\mathbf{y}_*$. Dissimilarity compares $\mathbf{y}_*$ to each sample using the base similarity function s , while Eccentricity measures the distance from $\mathbf{y}_*$ to the centroid in the Laplacian embedding space; see equation (10). To bridge these views, we measure dissimilarity within the embedding space itself, averaging the distances from the embedding of $\mathbf{y}_*$ to the embeddings of individual samples. This retains the joint-pairwise smoothing of Eccentricity and also reflects the variance among samples, rather than only the centroid. The sampling-based estimate is

$$\hat{U}_{EigVecD}(\mathbf{y}_* | \mathbf{x}) = \frac{1}{M} \sum_{i=1}^M \left\| \mathbf{v}_* - \mathbf{v}_i \right\|_2^2, \quad (14)$$

and the beam-guided, probability-weighted version is

$$\hat{U}_{EigVecD}^b(\mathbf{y}_* | \mathbf{x}) = \sum_{i=1}^M w_i \left\| \mathbf{v}_*^b - \mathbf{v}_i^b \right\|_2^2, \quad (15)$$

where the embeddings $\mathbf{v}_i$ (and $\mathbf{v}_i^b$) are obtained from the Graph Laplacian as in Eccentricity, and w_i are the normalized masses from equation (4). This estimator increases both when $\mathbf{y}_*$ moves away from the bulk and when the samples themselves are more dispersed; by contrast, Eccentricity focuses on the single distance to the weighted centroid.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

Datasets. We evaluate our approach on six QA datasets in total. Those include two closed-book datasets: *TriviaQA* (Joshi et al., 2017) and *Web Questions* (Berant et al., 2013), two open-book datasets: *CoQA* (Reddy et al., 2019) and *HotpotQA* (Yang et al., 2018) and two multiple-choice datasets: *CommonsenseQA* (Talmor et al., 2019) and *ARC-Challenge* (Clark et al., 2018). For each dataset, we randomly sampled several questions from the test set. The statistics for those datasets are available in Table 1. Prompt details and examples of questions are provided in Appendix C.

Table 1: Test dataset settings and statistics.

	Closed-Book QA		Open-Book QA		Multiple Choice	
	TriviaQA	Web Questions	CoQA	HotpotQA	Common senseQA	ARC-Challenge
# Questions	2000	1490	2000	2000	1221	447
# few-shot examples	5	5	all preceding	0	2	2
Max new tokens	20	20	20	20	10	20

Table 2: Summary of baseline UQ methods.

Category	Uncertainty Quantification Method
Information-based	Sequence Probability (Prob)
	Mean Token Entropy (MTE)
	Perplexity
	CCP (Fadeeva et al., 2024)
Reflexive	P(True) (Kadavath et al., 2022)
Sampling-based	Semantic Entropy (Kuhn et al., 2023)
	Shifting Attention to Relevance (SAR) (Duan et al., 2024)
	Lexical Similarity (Fomicheva et al., 2020)
	Sum of Eigenvalues of Laplacian (EigValLaplacian) (Lin et al., 2024)
	Number of Semantic Sets (NumSemSets) (Lin et al., 2024)

Models. We use base and instruct versions of 3 models: Gemma 3 4B (Team, 2025a), Llama 3.1 8B (Dubey et al., 2024), and Qwen 3 8B (Team, 2025b).

Metrics. Following best uncertainty benchmarking practices (Vashurin et al., 2025a), we adopt the Prediction-Rejection Ratio (PRR) as our primary evaluation metric. Consider a test dataset $\mathcal{D} = \{(\mathbf{x}_j, \mathbf{t}_j)\}$, where $\mathbf{t}_j$ denotes target output. Then, we can obtain an output $\mathbf{y}_j^*$ generated by an LLM for input $\mathbf{x}_j$ and the associated uncertainty score $u_j = U(\mathbf{y}_j^* | \mathbf{x}_j)$.

Based on these we can build rejection curve that captures how the average quality $Q(\mathbf{y}_j^*, \mathbf{t}_j)$ over all $\{(\mathbf{y}_j^*, \mathbf{t}_j) : u_j < \tau\}$ changes with the rejection threshold τ . An oracle rejection curve can be defined by substituting $u_j = -Q(\mathbf{y}_j^*, \mathbf{t}_j)$, giving best possible rejection order where lowest-quality outputs are rejected first. A baseline for rejection can be obtained by rejecting outputs uniformly at random. PRR is then defined as the ratio of the area between UQ rejection curve and a random rejection baseline to the area between oracle rejection and the same random baseline:

$$PRR = \frac{AUC_{\text{unc}} - AUC_{\text{rnd}}}{AUC_{\text{oracle}} - AUC_{\text{rnd}}}. \quad (16)$$

A higher PRR indicates a more effective uncertainty score. Following Vashurin et al. (2025a), we use AlignScore (Zha et al., 2023) as the quality metric Q . While PRR serves as our main evaluation measure, we additionally report ROC-AUC and PR-AUC in Appendix D.2.

Baselines. We evaluate four main methods, Dissimilarity, Eccentricity, Eigenvectors Dissimilarity, and CoCoA, under multinomial sampling and their beam-guided, probability-weighted variants. For CoCoA, we consider both *CocoaMSP* based on unnormalized log-probability:

$$u(\mathbf{y}_* | \mathbf{x}) = -\log p(\mathbf{y}_* | \mathbf{x}), \quad (17)$$

and *CocoaPPL* based on perplexity:

$$u(\mathbf{y}_* | \mathbf{x}) = -\frac{1}{|\mathbf{y}_*|} \log p(\mathbf{y}_* | \mathbf{x}). \quad (18)$$

In addition, we compare against several state-of-the-art UQ baselines summarized in Table 2, using implementations from LM-Polygraph (Fadeeva et al., 2023). The simplest baseline, *Sequence Probability*, calculates $-\log p(\mathbf{y}_* | \mathbf{x})$. For detailed descriptions of other methods see Appendix E.

Table 3: PRR ($\uparrow$ is better) averaged over 6 datasets. For each model, the top-1 method is **bold** and the second-best is underlined. For beam-guided variants, we mark $\uparrow$ when the variant improves over its original multinomial-sampling counterpart.

Method	Llama 3.1 8B base	Llama 3.1 8B instruct	Gemma 3 4B base	Gemma 3 4B instruct	Qwen 3 8B base	Qwen 3 8B instruct
<i>Baseline UQ Methods</i>						
Prob	.410 $\pm$.019	.344 $\pm$.031	.471 $\pm$.023	.292 $\pm$.022	.376 $\pm$.03	.289 $\pm$.067
MTE	.422 $\pm$.016	.364 $\pm$.026	.476 $\pm$.022	.317 $\pm$.028	.407 $\pm$.032	.297 $\pm$.064
Perplexity	.452 $\pm$.02	.323 $\pm$.027	.525 $\pm$.024	.288 $\pm$.025	.372 $\pm$.03	.276 $\pm$.058
CCP	.401 $\pm$.02	.364 $\pm$.029	.492 $\pm$.022	.331 $\pm$.026	.355 $\pm$.034	.291 $\pm$.06
SAR	.352 $\pm$.02	.385 $\pm$.029	.386 $\pm$.026	.239 $\pm$.024	.363 $\pm$.033	.292 $\pm$.052
P(True)	.015 $\pm$.023	.072 $\pm$.03	.093 $\pm$.026	-.096 $\pm$.024	.110 $\pm$.03	-.114 $\pm$.055
Semantic Entropy	.414 $\pm$.019	.376 $\pm$.025	.401 $\pm$.023	.293 $\pm$.024	.319 $\pm$.031	.299 $\pm$.058
Lexical Similarity	.411 $\pm$.02	.366 $\pm$.029	.426 $\pm$.025	.247 $\pm$.023	.425 $\pm$.034	.237 $\pm$.055
EigValLaplacian	.426 $\pm$.016	.371 $\pm$.028	.437 $\pm$.03	.233 $\pm$.025	.406 $\pm$.03	.265 $\pm$.056
NumSemSets	.396 $\pm$.018	.319 $\pm$.031	.418 $\pm$.024	.238 $\pm$.023	.365 $\pm$.033	.253 $\pm$.052
<i>Consistency-based UQ: multinomial vs. beamsearch versions</i>						
Dissimilarity	.505 $\pm$.018	.379 $\pm$.028	.630 $\pm$.021	.206 $\pm$.019	.477 $\pm$.037	.327 $\pm$.066
Dissimilarity + beamsearch	.543 $\uparrow$ $\pm$.019	<u>.417</u> $\uparrow$ $\pm$.026	.650 $\uparrow$ $\pm$.022	.252 $\uparrow$ $\pm$.022	.478 $\uparrow$ $\pm$.031	<u>.355</u> $\uparrow$ $\pm$.062
Eccentricity	.453 $\pm$.016	.368 $\pm$.029	.563 $\pm$.021	.231 $\pm$.025	.396 $\pm$.035	.251 $\pm$.058
Eccentricity + beamsearch	.505 $\uparrow$ $\pm$.017	.397 $\uparrow$ $\pm$.029	.603 $\uparrow$ $\pm$.023	.285 $\uparrow$ $\pm$.024	.410 $\uparrow$ $\pm$.03	.376 $\uparrow$ $\pm$.061
EigVecDissimilarity	.463 $\pm$.019	.370 $\pm$.028	.561 $\pm$.026	.236 $\pm$.025	.425 $\pm$.035	.256 $\pm$.051
EigVecDissimilarity + beamsearch	.510 $\uparrow$ $\pm$.021	.414 $\uparrow$ $\pm$.028	.598 $\uparrow$ $\pm$.022	.301 $\uparrow$ $\pm$.019	.450 $\uparrow$ $\pm$.033	.376 $\uparrow$ $\pm$.057
CocoaMSP	.505 $\pm$.018	.404 $\pm$.025	.587 $\pm$.023	.314 $\pm$.024	.461 $\pm$.031	.334 $\pm$.054
CocoaMSP + beamsearch	.521 $\uparrow$ $\pm$.019	.426 $\uparrow$ $\pm$.024	.615 $\uparrow$ $\pm$.021	.345 $\uparrow$ $\pm$.026	.473 $\pm$.03	.347 $\pm$.061
CocoaPPL	.523 $\pm$.017	.397 $\pm$.026	.628 $\pm$.024	.312 $\pm$.023	.461 $\pm$.034	.327 $\pm$.055
CocoaPPL + beamsearch	<u>.536</u> $\uparrow$ $\pm$.02	.412 $\uparrow$ $\pm$.027	<u>.649</u> $\uparrow$ $\pm$.026	<u>.339</u> $\uparrow$ $\pm$.021	.461 $\uparrow$ $\pm$.035	.337 $\uparrow$ $\pm$.057

All experiments use $M = 10$ candidates for both multinomial sampling and beam search. We adopt the entailment probability from the DeBERTa-large model fine-tuned on the MNLI task (He et al., 2021) for similarity function s , following Lin et al. (2024).

4.2 RESULTS AND DISCUSSION

Table 3 presents PRR results for six models, averaged over six datasets. Across all models, incorporating beam search consistently improves the performance of consistency-based uncertainty scores. Moreover, in almost every case, beam search–based methods achieve either the best or second-best PRR compared to both baselines and the original consistency-based approaches. In particular, Dissimilarity + Beam Search achieves the best PRR scores for all base models and the second-best scores for Llama 3.1 8B instruct and Qwen 3 8B instruct. Similarly, CocoaMSP + Beam Search achieves the best results for Llama 3.1 8B instruct and Gemma 3 4B instruct, while CocoaPPL + Beam Search ranks second-best for Llama 3.1 8B base, Gemma 3 4B base, and Gemma 3 4B instruct. We further provide separate results for each dataset in Appendix D.3.

4.3 ABLATIONS

In this section, we study sensitivity to (i) the number of candidates M , (ii) output length, and (iii) rejection rate in PRR curves.

4.3.1 EFFECT OF SAMPLE COUNT

We vary the sample count $M \in \{1, \dots, 15\}$ for Dissimilarity, Eccentricity, and EigVecDissimilarity under multinomial sampling and beam search. Figure 5 shows that beam search generally achieves higher PRR across all budgets $M \geq 2$. Notably, beam search reaches high PRR at small budgets (3-5 samples) and saturates quickly, while multinomial sampling improves more gradually and remains below beam search throughout.

For $M = 1$, beam search reduces to greedy decoding, causing Dissimilarity to be nearly zero because it compares two identical greedy outputs. In contrast, the sampling variant compares greedy decoding to a stochastic sample, yielding a more informative value.




Figure 4: PRR ($\uparrow$ is better) for Dissimilarity under beam search (with probability weights) vs. multinomial sampling, for different output lengths. Each dataset (TriviaQA, CoQA) with Gemma 3 4B base is partitioned into five approximately equal-size bins token length of greedy output.




Figure 5: PRR ($\uparrow$ is better) as a function of the number of candidates M on TriviaQA with Gemma 3 4B base. Each panel reports one estimator (Dissimilarity, Eccentricity, EigVecDissimilarity). Curves compare multinomial sampling and beam search (with probability weights from equation (4)).




Figure 6: Prediction-Rejection curves for *Dissimilarity*, *Eccentricity*, and *EigVecDissimilarity* on TriviaQA with Llama 3.1 8B base, comparing multinomial sampling (blue) and beam search with weights (orange). Oracle (black) and random (gray dashed) baselines are shown. The vertical dashed line marks the maximum rejection rate used in AUC calculations.

4.3.2 EFFECT OF OUTPUT LENGTH

Beam-guided estimators outperform sampling-based ones most clearly when generations are short. As shown earlier in Figure 2, duplicate rates under multinomial sampling are high for 2-4 tokens ($\sim 30\text{--}50\%$) and drop to $\sim 17\%$ for outputs of 8+ tokens. To quantify the impact, we compute PRR for Dissimilarity using beam search (with weights from equation (4)) and multinomial sampling (no weights) across five length bins of approximately equal size on TriviaQA and CoQA with Gemma 3 4B base; see Figure 4. Within each bin, beam search consistently beats multinomial sampling for short outputs; the gap narrows and becomes negligible for lengths of about 7 tokens and above, where duplication is less pronounced.

4.3.3 PREDICTION-REJECTION CURVES

Figure 6 compares full Prediction-Rejection curves for Dissimilarity, Eccentricity, and EigVecDis-similarity on TriviaQA with Llama 3.1 8B base. Across all estimators, beam search consistently dominates multinomial sampling for nearly the entire rejection range. The becomes increasingly pronounced as the rejection rate grows, where beam-guided estimates remain stable while multinomial ones flatten or even degrade. This indicates that beam search is especially beneficial in the high-rejection regime, where distinguishing between stronger and weaker candidates is most critical.

4.3.4 ADDITIONAL ABLATIONS

Additional ablations are deferred to the appendices: Appendix A.1 compares candidate-generation strategies including Diverse Beam Search, temperature sampling, and a hybrid multinomial-beam sampling. Appendix A.2 investigates restricted-mass normalization and shows that introducing a small probability floor ϵ can stabilize the weighting of low-mass beams. Appendix A.3 evaluates other sampling-based objectives (Semantic Entropy, Degree Matrix) under beam generation with probability-weighted formulations. Appendix D.1 examines top-1 beam decode as the produced answer y_* (instead of greedy), a natural choice when beam search has already been run.

5 RELATED WORK

Consistency-based uncertainty estimation. In a black-box setting, consistency-based methods are especially relevant, as they do not require access to the model internals. Lin et al. (2024) introduce several methods that estimate confidence based on a similarity matrix, where each entry represents the similarity between a pair of sampled generations. Fomicheva et al. (2020) present Lexical Similarity, a metric that evaluates the average similarity of words or phrases between each pair of responses. In a white-box setting, consistency signals can be combined with model token-probabilities-based confidence. These hybrid methods, such as Semantic Entropy (Kuhn et al., 2023), CoCoA (Vashurin et al., 2025b) and SAR (Duan et al., 2024) explore different ways of combining these signals and achieve state-of-the-art performance. However, these works are primarily concerned with the introduction of new methods for uncertainty quantification and use multinomial sampling as a way to approximate a variety of consistency-based measures.

Uncertainty and decoding. There were some efforts focused on examining the interaction between decoding strategies and uncertainty quantification. In particular, Hashimoto et al. (2025) explores the impact of decoding strategies on the performance of token probabilities-based UQ methods, namely Sequence Probability and Mean Token Entropy. The authors find that these scores produced with beam search can sometimes under perform compared to greedy or contrastive search. While this work offers interesting insights, no experiments with stochastic decoding strategies or non-likelihood based methods were conducted. Conversely, other research focused on making the decoding itself uncertainty-aware. For example, Daheim et al. (2025) propose Minimum Bayes Risk (MBR) decoding, which incorporates model uncertainty into the MBR objective for improved generation quality. Garces Arias et al. (2024) and Lee et al. (2025) incorporate uncertainty into contrastive search decoding. Lastly, Ding et al. (2025) combines global entropy trends and local deviations to guide a self-adaptive decoding. These works integrate uncertainty into the decoding process to improve the quality of the generation, rather than improving the performance of the uncertainty itself. Although some uncertainty-aware decoding methods have also demonstrated improved uncertainty quantification performance, they are generally not evaluated with consistency-based metrics.

6 CONCLUSION

We present a new family of uncertainty quantification methods for LLMs that employ a beam-weighted estimator of consistency-based uncertainty. Compared to multinomial sampling, commonly used in existing approaches, our method yields lower variance in dissimilarity and greater diversity of candidate answers. We also derive a theoretical lower bound on the beam set probability mass under which the error of the multinomial Monte Carlo estimator is guaranteed to be larger. Finally, we evaluate our approach on six QA datasets and six different models, demonstrating state-of-the-art performance.

LIMITATIONS

Although our method provides an improvement over existing consistency-based estimators, several important considerations remain. First, we evaluated our methods in white-box settings, as they require access to the model’s probability distributions. Nonetheless, we argue that developing methods tailored for white-box settings continues to be of great importance given their continued relevance and usage. Moreover, the methods could be extended to the black-box settings using empirical probability estimates.

Second, our experiments are limited to short-form QA datasets, and the generalizability of our findings to longer-form generation remains an open question.

Lastly, our implementation and evaluation relies on existing neural metrics: AlignScore is used to score the quality of the generation, and pre-trained NLI model is utilized as a measure of consistency. Although widely used in previous work, certain more specialized tasks might require different sample similarity measures and quality metrics.

REFERENCES

- Joris Baan, Nico Daheim, Evgenia Ilia, Dennis Ulmer, Haau-Sing Li, R. Fernández, Barbara Plank, Rico Sennrich, Chrysoula Zerva, and Wilker Aziz. Uncertainty in natural language generation: From theory to applications. *ArXiv*, abs/2307.15703, 2023. URL <https://api.semanticscholar.org/CorpusID:260316110>.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. Semantic parsing on Freebase from question-answer pairs. In David Yarowsky, Timothy Baldwin, Anna Korhonen, Karen Livescu, and Steven Bethard (eds.), *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1533–1544, Seattle, Washington, USA, October 2013. Association for Computational Linguistics. URL <https://aclanthology.org/D13-1160/>.
- Felix Busch, Lena Hoffmann, Christopher Rueger, Elon HC van Dijk, Rawen Kader, Esteban Ortiz-Prado, Marcus R. Makowski, Luca Saba, Martin Hadamitzky, Jakob Nikolas Kather, Daniel Truhn, Renato Cuocolo, Lisa C. Adams, and Keno K. Bressemer. Current applications and challenges in large language models for patient care: A systematic review. *Communications Medicine*, 5(1), 2025. ISSN 2730-664X. doi: 10.1038/s43856-024-00717-2. URL <http://dx.doi.org/10.1038/s43856-024-00717-2>.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? Try ARC, the AI2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018. URL <https://arxiv.org/abs/1803.05457>.
- Nico Daheim, Clara Meister, Thomas Möllenhoff, and Iryna Gurevych. Uncertainty-aware decoding with minimum Bayes risk. In *The Thirteenth International Conference on Learning Representations*, Singapore, 2025. URL <https://openreview.net/forum?id=hPpyUv1XyQ>.
- Yuanhao Ding, Esteban Garces Arias, Meimingwei Li, Julian Rodemann, Matthias Aßenmacher, Danlu Chen, Gaojuan Fan, Christian Heumann, and Chongsheng Zhang. GUARD: Glocal uncertainty-aware robust decoding for effective and efficient open-ended text generation. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2025*, pp. 7202–7226, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-335-7. doi: 10.18653/v1/2025.findings-emnlp.380. URL <https://aclanthology.org/2025.findings-emnlp.380/>.
- Jinhao Duan, Hao Cheng, Shiqi Wang, Alex Zavalny, Chenan Wang, Renjing Xu, Bhavya Kailkhura, and Kaidi Xu. Shifting attention to relevance: Towards the predictive uncertainty quantification of free-form large language models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 5050–5063, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.276. URL <https://aclanthology.org/2024.acl-long.276/>.

- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, and Angela Fan. The Llama 3 herd of models. *CoRR*, abs/2407.21783, 2024. URL <https://doi.org/10.48550/arXiv.2407.21783>.
- 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. LM-polygraph: Uncertainty estimation for language models. In Yansong Feng and Els Lefever (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 446–461, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-demo.41. URL <https://aclanthology.org/2023.emnlp-demo.41/>.
- 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. Fact-checking the output of large language models via token-level uncertainty quantification. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 9367–9385, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.558. URL <https://aclanthology.org/2024.findings-acl.558/>.
- Marina Fomicheva, Shuo Sun, Lisa Yankovskaya, Frédéric Blain, Francisco Guzmán, Mark Fishel, Nikolaos Aletras, Vishrav Chaudhary, and Lucia Specia. Unsupervised quality estimation for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8: 539–555, 2020. doi: 10.1162/tacl.a.00330. URL <https://aclanthology.org/2020.tacl-1.35/>.
- Esteban Garces Arias, Julian Rodemann, Meimingwei Li, Christian Heumann, and Matthias Aßenmacher. Adaptive contrastive search: Uncertainty-guided decoding for open-ended text generation. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 15060–15080, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-emnlp.885. URL <https://aclanthology.org/2024.findings-emnlp.885/>.
- Wataru Hashimoto, Hidetaka Kamigaito, and Taro Watanabe. Decoding uncertainty: The impact of decoding strategies for uncertainty estimation in large language models. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2025*, pp. 14601–14613, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-335-7. doi: 10.18653/v1/2025.findings-emnlp.788. URL <https://aclanthology.org/2025.findings-emnlp.788/>.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. DeBERTa: Decoding-enhanced BERT with disentangled attention. In *Proceedings of the International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=XPZiaotutsD>.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rygGQyrFvH>.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension. In Regina Barzilay and Min-Yen Kan (eds.), *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1601–1611, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1147. URL <https://aclanthology.org/P17-1147/>.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, Deep Ganguli, Danny Hernandez, Josh Jacobson, Jackson Kernion, Shauna Kravec, Liane Lovitt, Kamal Ndousse, Catherine Olsson, Sam Ringer, Dario Amodei, Tom

- Brown, Jack Clark, Nicholas Joseph, Ben Mann, Sam McCandlish, Chris Olah, and Jared Kaplan. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*, 2022. doi: 10.48550/arXiv.2207.05221. URL <https://doi.org/10.48550/arXiv.2207.05221>.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*, 2023. URL <https://openreview.net/pdf?id=VD-AYtP0dve>.
- Hakyung Lee, Subeen Park, Joowang Kim, Sungjun Lim, and Kyungwoo Song. Uncertainty-aware contrastive decoding. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 26376–26391, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-256-5. doi: 10.18653/v1/2025.findings-acl.1352. URL <https://aclanthology.org/2025.findings-acl.1352/>.
- Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. Generating with confidence: Uncertainty quantification for black-box large language models. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=DWkJCSxKU5>.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. RoBERTa: A robustly optimized BERT pre-training approach. *arXiv preprint arXiv:1907.11692*, 2019. URL <https://arxiv.org/abs/1907.11692>.
- Siva Reddy, Danqi Chen, and Christopher D. Manning. CoQA: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266, 2019. doi: 10.1162/tacl.a.00266. URL <https://aclanthology.org/Q19-1016>.
- Dong Shu, Haoran Zhao, Xukun Liu, David Demeter, Mengnan Du, and Yongfeng Zhang. LawLLM: Law large language model for the US legal system. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pp. 4882–4889, Boise, Idaho, USA, 2024. Association for Computing Machinery. doi: 10.1145/3627673.3680020. URL <https://doi.org/10.1145/3627673.3680020>.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. CommonsenseQA: A question answering challenge targeting commonsense knowledge. In Jill Burstein, Christy Doran, and Thamar Solorio (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4149–4158, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1421. URL <https://aclanthology.org/N19-1421/>.
- Gemma Team. Gemma 3 technical report, 2025a. URL <https://arxiv.org/abs/2503.19786>.
- Qwen Team. Qwen3 technical report, 2025b. URL <https://arxiv.org/abs/2505.09388>.
- Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher Manning. Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 5433–5442, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.330. URL <https://aclanthology.org/2023.emnlp-main.330/>.
- Roman Vashurin, Ekaterina Fadeeva, Artem Vazhentsev, Lyudmila Rvanova, Daniil Vasilev, Akim Tsvigun, Sergey Petrakov, Rui Xing, Abdelrahman Sadallah, Kirill Grishchenkov, Alexander Panchenko, Timothy Baldwin, Preslav Nakov, Maxim Panov, and Artem Shelmanov. Benchmarking uncertainty quantification methods for large language models with LM-Polygraph. *Transactions of the Association for Computational Linguistics*, 13:220–248, 2025a. doi: 10.1162/tacl.a.00737. URL <https://aclanthology.org/2025.tacl-1.11/>.

- Roman Vashurin, Maiya Goloburda, Albina Ilina, Aleksandr Rubashevskii, Preslav Nakov, Artem Shelmanov, and Maxim Panov. CoCoA: A minimum Bayes risk framework bridging confidence and consistency for uncertainty quantification in llms. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025b. URL <https://openreview.net/forum?id=H1NGlLNVC>.
- Ashwin Vijayakumar, Michael Cogswell, Ramprasaath Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. Diverse beam search for improved description of complex scenes. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), Apr. 2018. doi: 10.1609/aaai.v32i1.12340. URL <https://ojs.aaai.org/index.php/AAAI/article/view/12340>.
- Zhiqiu Xia, Jinxuan Xu, Yuqian Zhang, and Hang Liu. A survey of uncertainty estimation methods on large language models. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 21381–21396, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-256-5. doi: 10.18653/v1/2025.findings-acl.1101. URL <https://aclanthology.org/2025.findings-acl.1101/>.
- Yijun Xiao and William Yang Wang. Quantifying uncertainties in natural language processing tasks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):7322–7329, 2019. ISSN 2159-5399. doi: 10.1609/aaai.v33i01.33017322. URL <http://dx.doi.org/10.1609/aaai.v33i01.33017322>.
- Wanli Xing, Nia Nixon, Scott Crossley, Paul Denny, Andrew Lan, John Stamper, and Zhou Yu. The use of large language models in education. *International Journal of Artificial Intelligence in Education*, 35(2):439–443, February 2025. ISSN 1560-4306. doi: 10.1007/s40593-025-00457-x. URL <http://dx.doi.org/10.1007/s40593-025-00457-x>.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Brussels, Belgium, 2018. URL <https://aclanthology.org/D18-1259/>.
- KiYoon Yoo, Jangho Kim, Jiho Jang, and Nojun Kwak. Detection of adversarial examples in text classification: Benchmark and baseline via robust density estimation. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Findings of the Association for Computational Linguistics: ACL 2022*, pp. 3656–3672, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.findings-acl.289. URL <https://aclanthology.org/2022.findings-acl.289/>.
- Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. AlignScore: Evaluating factual consistency with a unified alignment function. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 11328–11348, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.634. URL <https://aclanthology.org/2023.acl-long.634/>.

A ABLATION STUDIES

A.1 DIFFERENT SAMPLING STRATEGIES

This section studies how the proposed estimators behave under different sample generation strategies. In addition to multinomial sampling and beam search settings, we evaluate three additional families.

Diverse beam search. We generate $M = 10$ candidates using a diverse beam search (Vijayarumar et al., 2018) with group penalties $\lambda \in \{0.5, 1.0, 1.5, 2.0\}$ and group counts that split the ten candidates into $G \in \{2, 5\}$ groups. As in the main beam setup, we apply the same self-normalized probability weights w_i from equation (4).

Temperature sampling with importance weights. For different temperatures T , we draw $M = 10$ samples with temperature sampling $\{\mathbf{y}_T^{(i)}\}_{i=1}^M$ and re-weight them via self-normalized importance weights

$$w_i^T = \frac{p(\mathbf{y}_T^{(i)} | \mathbf{x})^{1-1/T}}{\sum_{j=1}^M p(\mathbf{y}_T^{(j)} | \mathbf{x})^{1-1/T}}. \quad (19)$$

Hybrid multinomial-beam search. We also consider a joint strategy: first draw B beam candidates, then draw the remaining $M - B$ candidates via multinomial sampling while excluding the beam results. Beam candidates use autoregressive probability weights, and the residual probability mass is distributed uniformly over the multinomial samples. Let $\{\mathbf{b}^{(i)}\}_{i=1}^B$ be the beam outputs and $\{\mathbf{y}^{(j)}\}_{j=B+1}^M$ the multinomial samples (with beam sequences masked out). We assign weights

$$w_i^H = p(\mathbf{b}^{(i)} | \mathbf{x}), \quad i = 1, \dots, B, \quad w_j^H = \frac{1 - \sum_{i=1}^B p(\mathbf{b}^{(i)} | \mathbf{x})}{M - B}, \quad j = B + 1, \dots, M, \quad (20)$$

so that $\sum_{i=1}^M w_i^H = 1$. We test $B \in \{1, \dots, 9\}$.

Evaluations use a subset of 500 examples from TriviaQA and CoQA with two base models, Gemma 3 4B base and Llama 3.1 8B base. Results are summarized in Table 4.

Table 4: PRR ($\uparrow$ is better) under different sampling strategies. Columns list methods (Dissimilarity, Eccentricity, EigVecDissimilarity) and four different model-dataset pairs; rows list strategies with their hyperparameters. Per column, top-1 is **bold**, second-best is underlined.

		Gemma 3 4B base						Llama 3.1 8B base					
		TriviaQA			CoQA			TriviaQA			CoQA		
		Dissim	Ecc	EigVec Dissim	Dissim	Ecc	EigVec Dissim	Dissim	Ecc	EigVec Dissim	Dissim	Ecc	EigVec Dissim
Beam search		.771	.732	.751	.561	.483	.488	.623	.581	.598	<u>.502</u>	.438	.454
Multinomial Sampling	$T = 0.7$	.703	.619	.687	.455	.368	.397	.561	.521	.538	.451	.371	.382
	$T = 0.9$	.734	.664	.710	.468	.417	.425	.588	.521	.533	.418	.407	.410
	$T = 1.0$	.742	.689	.715	.465	.424	.432	.599	.516	.537	.431	.416	.424
	$T = 1.2$	.733	.679	.717	.435	.424	.437	.610	.517	.535	.391	.427	.433
	$T = 1.5$	.718	.685	.717	.406	.426	.436	.555	.515	.532	.397	.431	.440
	$T = 1.7$	.680	.690	.717	.372	.426	.433	.569	.514	.530	.304	.432	<u>.441</u>
Diverse Beam Search	$G = 2, \lambda = 0.5$	.753	.528	.693	.498	.405	.452	.623	-.123	.546	.458	.310	.363
	$G = 2, \lambda = 1.0$	.753	.566	.705	.518	.384	.432	.594	-.138	.539	.466	.285	.355
	$G = 2, \lambda = 1.5$	.763	.522	.714	.537	.420	.441	.618	-.101	.542	.462	.245	.317
	$G = 2, \lambda = 2.0$	.759	.515	.702	.547	.377	.391	.630	-.130	.546	.452	.215	.287
	$G = 5, \lambda = 0.5$	.758	.546	.736	.493	.401	.453	.591	-.026	.569	.395	.262	.353
	$G = 5, \lambda = 1.0$	.768	.523	.746	.515	.369	.423	.615	-.086	.563	.447	.274	.376
	$G = 5, \lambda = 1.5$	.761	.453	.723	.513	.391	.427	.623	-.153	.533	.461	.199	.324
	$G = 5, \lambda = 2.0$	.770	.476	.690	.513	.355	.415	.631	-.093	.548	.453	.132	.254
Hybrid Multinomial- Beam	$B = 1$	.759	.715	.746	.512	.451	.433	.597	.519	.549	.386	.314	.325
	$B = 2$	.765	.731	.745	.519	.470	.435	.617	.564	.586	.445	.338	.345
	$B = 3$	<u>.781</u>	.736	.754	.503	.461	.439	.620	.553	.598	.519	<u>.436</u>	.424
	$B = 4$	.784	.750	.769	.516	.467	.428	.622	.572	<u>.617</u>	.436	.388	.382
	$B = 5$	.757	.733	.749	.538	.500	.466	.655	.578	.609	.470	.412	.419
	$B = 6$	.773	.733	.756	.528	.512	.488	.635	.586	.613	.486	.421	.415
	$B = 7$	.771	.737	.754	.543	.468	.471	.640	.596	.617	.483	.399	.427
	$B = 8$	.764	.733	.755	.548	.504	.507	<u>.648</u>	<u>.597</u>	.610	.491	.434	.425
	$B = 9$	.772	<u>.747</u>	<u>.765</u>	<u>.551</u>	<u>.509</u>	<u>.497</u>	.646	.597	.618	.501	.427	.439

No single strategy dominates across datasets, models, or estimators. Temperature sampling (with importance weights) and diverse beam search systematically yield to beam search and hybrid multinomial-beam. Hybrid multinomial-beam strategy can reach top-1 for specific hyperparameter B , but gains are not systematic and are sensitive to tuning. Given this variability and tuning cost, plain beam search with probability weighting is a reasonable default.

A.2 RESTRICTED-MASS NORMALIZATION

Equation (4) normalizes autoregressive sequence probabilities over the M beam candidates. This choice can be sensitive to tail candidates whose probabilities are tiny and length-dependent. To test robustness, we introduce a floor ϵ on the per-candidate mass:

$$w_i^\epsilon = \frac{\max(\epsilon, p(\mathbf{b}^{(i)} | \mathbf{x}))}{\sum_{j=1}^M \max(\epsilon, p(\mathbf{b}^{(j)} | \mathbf{x}))}. \quad (21)$$

The setting $\epsilon = 0$ recovers equation (4); $\epsilon = 1$ yields uniform weights $w_i^1 = 1/M$. Intermediate ϵ values trade off fidelity to the model distribution against robustness to noisy, length-biased tails.

We evaluate beam-guided probability-weighted methods for different ϵ on a subset of 500 examples from TriviaQA and CoQA with two base models, Gemma 3 4B base and Llama 3.1 8B base. Results are summarized in Table 5.

The results do not indicate a clear best choice of method and corresponding ϵ parameter. Determining the optimal ϵ is a case-dependent task.

Table 5: PRR ($\uparrow$ is better) under restricted-mass normalization ablation. Columns group dataset-model pairs with methods (Dissim, Ecc, EigVecDissim). Rows vary the mass floor ϵ in equation (21): $\epsilon = 0$ recovers equation (4); $\epsilon = 1$ yields uniform weights $w_i = 1/M$. For each dataset-method, the top-1 score is **bold** and the second-best is underlined.

	Gemma 3 4B base						Llama 3.1 8B base					
	TriviaQA			CoQA			TriviaQA			CoQA		
	Dissim	Ecc	EigVec Dissim	Dissim	Ecc	EigVec Dissim	Dissim	Ecc	EigVec Dissim	Dissim	Ecc	EigVec Dissim
$\epsilon = 1.0$	.765	.741	.744	.536	.487	.461	.668	.596	.607	.470	.428	.410
$\epsilon = 0.1$	.765	.727	.745	.556	.497	.483	<u>.667</u>	.612	.627	<u>.502</u>	.447	.446
$\epsilon = 0.05$	.764	.720	.744	.561	<u>.490</u>	.487	.657	<u>.606</u>	<u>.626</u>	.509	.435	.451
$\epsilon = 0.01$	.766	.718	.749	.559	.478	.489	.630	.584	.602	.496	.437	.452
$\epsilon = 0.001$	<u>.771</u>	.731	.751	.562	.484	.488	.624	.581	.598	.501	<u>.438</u>	.453
$\epsilon = 0.00001$	.771	<u>.732</u>	<u>.751</u>	<u>.561</u>	.483	.488	.623	.581	.598	.502	.438	.454
$\epsilon = 0$	.771	<u>.732</u>	<u>.751</u>	<u>.561</u>	.483	.488	.623	.581	.598	.502	.438	<u>.454</u>

A.3 OTHER SAMPLING-BASED METHODS UNDER BEAM SEARCH

Beyond Dissimilarity, Eccentricity, and EigVecDissimilarity, this ablation evaluates two other sampling-based methods under beam-generated candidates: *Degree Matrix* (Lin et al., 2024) and *Semantic Entropy* (Kuhn et al., 2023). We also provide probability-weighted beam formulations using the weights w_i from equation (4).

Degree Matrix. Given M multinomial samples $\{\mathbf{y}^{(i)}\}_{i=1}^M$, Degree Matrix estimates the average pairwise dissimilarity:

$$\hat{U}_{DegMat}(\mathbf{x}) = \frac{1}{M^2} \sum_{i=1}^M \sum_{j=1}^M (1 - s(\mathbf{y}^{(i)}, \mathbf{y}^{(j)})). \quad (22)$$

For beam candidates $\{\mathbf{b}^{(i)}\}_{i=1}^M$, our mass-aware variant averages with weights:

$$\hat{U}_{DegMat}^b(\mathbf{x}) = \sum_{i=1}^M w_i \sum_{j=1}^M w_j (1 - s(\mathbf{b}^{(i)}, \mathbf{b}^{(j)})). \quad (23)$$

Table 6: PR-AUC ($\uparrow$ is better) on 6 datasets with Gemma 3 4B base. Each method is shown as a pair: its multinomial-sampling variant and its beam-search variant; $\uparrow$ denotes an improvement of the beam variant over its multinomial counterpart. Along main methods, the table includes input-uncertainty methods (Semantic Entropy, Lexical Similarity). For each dataset, the top-1 score is **bold** and the second-best is underlined. The rightmost column reports the mean PR-AUC across datasets.

UQ Method	TriviaQA	Web Questions	CoQA	HotpotQA	Common senseQA	ARC-Challenge	Mean
Semantic Entropy	.622	.505	.301	.140	.407	.431	.401
Semantic Entropy + beamsearch	.685 $\uparrow$	.614 $\uparrow$	.365 $\uparrow$	.278 $\uparrow$	.436 $\uparrow$	.454 $\uparrow$	.472 $\uparrow$
Degree Matrix	.682	.605	.385	.311	.409	.419	.469
Degree Matrix + beamsearch	.673	.642 $\uparrow$	.328	.244	.444 $\uparrow$	.473 $\uparrow$	.467
SAR	.656	.571	.347	.296	.183	.264	.386
SAR + beamsearch	.671 $\uparrow$	.589 $\uparrow$	.329	.266	.209 $\uparrow$	.269 $\uparrow$	.372
Dissimilarity	<u>.755</u>	<u>.715</u>	<u>.578</u>	<u>.626</u>	.561	.545	.630
Dissimilarity + beamsearch	.766$\uparrow$	.722$\uparrow$	.600$\uparrow$	.611	.595$\uparrow$	.604 $\uparrow$	.650$\uparrow$
Eccentricity	.714	.653	.459	.453	.549	.549	.563
Eccentricity + beamsearch	.739 $\uparrow$	.633	.505 $\uparrow$	.514 $\uparrow$	<u>.590$\uparrow$</u>	.636$\uparrow$	.603 $\uparrow$
EigVecDissimilarity	.738	.661	.443	.448	.512	.562	.561
EigVecDissimilarity + beamsearch	.753 $\uparrow$	.668 $\uparrow$	.497 $\uparrow$	.487 $\uparrow$	.562 $\uparrow$	<u>.621$\uparrow$</u>	.598 $\uparrow$
CocoaMSP	.738	.666	.509	.430	.583	.595	.587
CocoaMSP + beamsearch	.747 $\uparrow$	.679 $\uparrow$	.548 $\uparrow$	.523 $\uparrow$	.586 $\uparrow$	.606 $\uparrow$	.615 $\uparrow$
CocoaPPL	.739	.678	.548	.625	.580	.595	.628
CocoaPPL + beamsearch	.748 $\uparrow$	.694 $\uparrow$	.577 $\uparrow$	.681$\uparrow$	.582 $\uparrow$	.610 $\uparrow$	<u>.649$\uparrow$</u>

Semantic Entropy. Multinomial samples are clustered into semantic equivalence classes C . For each class, we calculate its probability

$$\hat{p}(c) = \frac{1}{M} \sum_{i=1}^M \mathbf{1}\{\mathbf{y}^{(i)} \in c\} \quad \text{for } c \in C. \quad (24)$$

Then Semantic Entropy calculates

$$\hat{U}_{SemEnt}(\mathbf{x}) = -\frac{1}{|C|} \sum_{c \in C} \log \hat{p}(c). \quad (25)$$

For beam candidates, use cluster masses aggregated by w_i :

$$\hat{p}^b(c) = \sum_{i=1}^M w_i \mathbf{1}\{\mathbf{b}^{(i)} \in c\}, \quad \hat{U}_{SemEnt}^b(\mathbf{x}) = -\sum_{c \in C} \hat{p}^b(c) \log \hat{p}^b(c). \quad (26)$$

Note that these objectives score LLM uncertainty about the *input* $\mathbf{x}$ as they are independent of a particular $\mathbf{y}_*$.

The results are summarized in Table 6. Beam search yields significant gains for Semantic Entropy and little to no improvement for Degree Matrix. Even with the beam-adapted formulations above, both objectives show worse results in terms of absolute PR-AUC compared to other methods. The primary reason is the target mismatch: as noted, these scores quantify uncertainty of the input $\mathbf{x}$ and are independent of the produced answer $\mathbf{y}_*$, whereas our main methods, Dissimilarity, Eccentricity and EigVecDissimilarity, focuses on ranking the correctness of $\mathbf{y}_*$ itself.

To further assess performance under different numbers of samples M used for UQ, we plot PRR as a function of M for one selected baseline, Semantic Entropy, as well as for Dissimilarity (both sampling and beam-search variants) for reference. Figure 7 presents the results, showing that for all $M > 1$, Semantic Entropy underperforms both Dissimilarity variants. This occurs because Dissimilarity measures the targeted uncertainty of the specific generation $\mathbf{y}_*$ rather than the overall uncertainty associated with $\mathbf{x}$, measured by Semantic Entropy.




Figure 7: PRR ($\uparrow$ is better) as a function of the number of candidates M on TriviaQA with Gemma 3 4B base for 3 UQ methods: Semantic Entropy, and sampling and beam search versions of Dissimilarity.

Table 7: PRR ($\uparrow$ is better) for Eccentricity and EigVecDissimilarity under different Graph Laplacian embedding choices on four dataset-model pairs. Top block varies the eigenvalue threshold α (retaining all $\lambda_i < \alpha$); bottom block fixes the embedding dimension K . For each pair, the best score is **bold** and the second-best is underlined.

	Gemma 3 4B base				Llama 3.1 8B base			
	TriviaQA		CoQA		TriviaQA		CoQA	
	Ecc	EigVec Dissim	Ecc	EigVec Dissim	Ecc	EigVec Dissim	Ecc	EigVec Dissim
$\alpha = 0.3$	.717	.710	.434	.355	.601	.599	.408	.364
$\alpha = 0.5$	.752	.740	.497	.460	.627	.626	.431	.420
$\alpha = 0.7$	.750	.749	.539	.498	<u>.622</u>	.643	.438	.441
$\alpha = 0.8$	<u>.751</u>	.757	<u>.508</u>	.475	.616	<u>.630</u>	.444	<u>.450</u>
$\alpha = 0.9$	.732	.751	.483	<u>.488</u>	.581	.598	<u>.438</u>	.454
$\alpha = 0.99$	.725	<u>.755</u>	.432	.454	.535	.561	.397	.409
$K = 1$	.454	.358	.346	.332	.444	.357	.304	.280
$K = 2$	.510	.532	.383	.361	.474	.469	.318	.338
$K = 3$	.619	.639	.434	.429	.538	.534	.351	.348
$K = 4$	.645	.643	.418	.412	.519	.514	.359	.352
$K = 5$	.638	.655	.366	.352	.487	.492	.314	.316
$K = 6$	.529	.545	.244	.236	.363	.368	.210	.211
$K = 7$	.210	.242	-.101	-.076	.050	.062	-.208	-.211
$K = 8$	-.265	-.209	-.210	-.171	-.358	-.349	-.327	-.308
$K = 9$	-.566	-.414	-.268	-.186	-.462	-.410	-.339	-.317
$K = 10$	-.659	-.484	-.283	-.189	-.467	-.397	-.330	-.249

A.4 GRAPH LAPLACIAN EMBEDDING PARAMETERS

Both multinomial and beam-guided versions of Eccentricity and EigVecDissimilarity depend on the threshold parameter α , which selects eigenvectors of the Graph Laplacian $L = I - D^{-1/2}WD^{-1/2}$ used to form semantic embeddings. Specifically, after computing the eigenpairs $\{\lambda_i, \mathbf{u}_i\}_{i=1}^{M+1}$, we retain those with $\lambda_i < \alpha$, yielding K eigenvectors in total and embeddings $\mathbf{v}_j = [\mathbf{u}_{1j}, \mathbf{u}_{2j}, \dots, \mathbf{u}_{Kj}]$ of dimension K . All eigenvalues lie in $[0, 1]$; smaller values capture stronger semantic structure, whereas values closer to 1 tend to reflect noise (Lin et al., 2024). In the main experiments we follow the original Eccentricity setting and use $\alpha = 0.9$.

Here we vary α and also test a fixed- K strategy (i.e., keeping exactly K leading low-spectrum eigenvectors irrespective of the threshold).

Table 7 reports the performance for Eccentricity and EigVecDissimilarity across α and K on four dataset-model pairs. A fixed embedding size performs poorly: the optimal number of informative directions varies between candidate sets, so fixing K either underfits or includes noisy directions. Thresholding is more robust: $\alpha \in [0.7, 0.9]$ consistently yields strong results across methods and pairs, supporting our default choice $\alpha = 0.9$.

A.5 CROSS-ENCODER SIMILARITY

In the main text, we instantiate the similarity function s using an NLI score: the entailment probability from a DeBERTa model. CoCoA, however, originally used a RoBERTa-large *cross-encoder* fine-tuned on the Semantic Textual Similarity benchmark (Liu et al., 2019). Table 8 reports PRR for Gemma 3 4B base when replacing the NLI-based s with this cross-encoder; all other settings are unchanged.

Table 8: PRR ($\uparrow$ is better) on 6 datasets with Gemma 3 4B base using a RoBERTa-large cross-encoder (STS) as the similarity function s in place of NLI. For each dataset, the top-1 is **bold** and the second-best is underlined; $\uparrow$ marks an improvement of a beam variant over its multinomial counterpart.

UQ Method	TriviaQA	Web Questions	CoQA	HotpotQA	Common senseQA	ARC-Challenge	Mean
Dissimilarity	.725	<u>.683</u>	.497	.597	.481	.421	.567
Dissimilarity + beamsearch	.746$\uparrow$	.693$\uparrow$	.513$\uparrow$	.654$\uparrow$	.505 $\uparrow$	.479 $\uparrow$	<u>.598$\uparrow$</u>
Eccentricity	.722	.647	.489	.544	.455	.500	.560
Eccentricity + beamsearch	.734 $\uparrow$	.647 $\uparrow$	.483	.604 $\uparrow$	.362	.421	.542
EigVecDissimilarity	.737	.649	.453	.523	.489	.529	.563
EigVecDissimilarity + beamsearch	<u>.744$\uparrow$</u>	.675 $\uparrow$	.484 $\uparrow$	.582 $\uparrow$	.439	.496	.570 $\uparrow$
CocoaMSP	.731	.642	.438	.397	<u>.553</u>	.577	.556
CocoaMSP + beamsearch	.740 $\uparrow$	.648 $\uparrow$	.462 $\uparrow$	.479 $\uparrow$	.558$\uparrow$	.593$\uparrow$	.580 $\uparrow$
CocoaPPL	.728	.653	.488	.607	.546	.567	.598
CocoaPPL + beamsearch	.737 $\uparrow$	.658 $\uparrow$	<u>.498$\uparrow$</u>	<u>.650$\uparrow$</u>	.548 $\uparrow$	<u>.586$\uparrow$</u>	.613$\uparrow$

A.6 NUMBER OF SAMPLES ACROSS DIFFERENT TASKS




Figure 8: PRR ($\uparrow$ is better) as a function of the number of candidates M across different datasets and models.

To evaluate the performance of the proposed beam-search variations under different numbers of samples M across models, we computed PRR for both the sampling and beam-search versions of Dissimilarity on 200 random subsamples of TriviaQA for three LLMs: Gemma 3 4B base, Llama 3.1 8B base, and Qwen 3 8B base. To further assess performance across datasets, we additionally evaluated PRR for Gemma 3 4B base on the 200 subsamples of ARC-Challenge dataset. All four resulting plots are shown in Figure 8.

The results show that for all budgets $M > 1$, beam search consistently outperforms sampling, yielding higher PRR. On the open-ended TriviaQA dataset, PRR increases steadily with M , with beam search reaching a plateau around $M = 5$ for all 3 models tested. On the multiple-choice ARC-Challenge dataset, PRR plateaus at a considerably smaller budget ($M = 2$), likely due to the small output space (i.e., a limited set of answer choices).

Overall, these results indicate that the beam-search variant of Dissimilarity remains effective even at relatively small sample budgets: $M \approx 5$ for open-ended short-form generation tasks, and $M = 2$ for multiple-choice settings, where the constrained output space enables faster saturation.

B ANALYSIS AND EXAMPLES

B.1 PROBABILITY MASS COVERAGE




Figure 9: *Left*: average probability mass covered by the candidate set ($M=10$) across output-length bins (averaged over examples in the bin) on TriviaQA with Gemma 3 4B base. *Right*: for beam search, distribution of sequence probabilities $p(\mathbf{b}^{(i)} | \mathbf{x})$ by beam rank i (1 = highest-probability text).

Figure 9 summarizes two observations. First (left), beam search covers a larger share of the model’s probability mass than multinomial sampling across length bins. Second (right), beam probabilities decay sharply with rank: the first few beams capture most of the mass, while lower-ranked beams contribute little. This motivates mass-aware weighting w_i (see equation (4)) and helps explain why probability-weighted beam variants are effective, especially at small candidate budgets.

B.2 EXAMPLES

We include qualitative examples for Gemma 3 4B base: two from TriviaQA, two from WebQuestions, and one from CoQA. Each panel shows the question, the greedy answer, ten multinomial samples, and ten beam-search samples with autoregressive probabilities, together with the corresponding uncertainty scores (e.g., Dissimilarity and its beam-guided variant). The cases illustrate how beam search reduces duplication and enhances uncertainty.

Question: What claimed the life of singer Kathleen Ferrier? Greedy: breasts cancer			Question: Which number Beethoven symphony is known as ‘The Pastoral’? Greedy: 6		
Multinomial samples	Beam-search samples		Multinomial samples	Beam-search samples	
cancer	cancer	p=0.228	six	sixth	p=0.314
breast cancer	tuberculosis	p=0.154	seventh	6	p=0.169
pulmonary	breast cancer	p=0.089	sixth	6th	p=0.104
breast cancer	lung cancer	p=0.041	sixth	ninth	p=0.061
cancer	pneumonia	p=0.039	sixth	seventh	p=0.037
breast cancer	leukaemia	p=0.034	6	9	p=0.027
myx	myel	p=0.023	seventh	six	p=0.023
cancer	leuk	p=0.011	no	9th	p=0.021
cancer	pulmonary	p=0.011	sixteenth	no.	p=0.013
pneumonia	lymphoma	p=0.010	n6	7	p=0.008
Dissimilarity: 0.330 Dissimilarity + beamsearch: 0.533			Dissimilarity: 0.634 Dissimilarity + beamsearch: 0.561		

Figure 10: Two examples from Gemma 3 4B base on TriviaQA. Each panel shows the question, greedy answer, multinomial and beam-search samples with autoregressive probabilities, plus dissimilarity and beamsearch-guided dissimilarity.

Question: what currency does cyprus use? Greedy: Cyprus pound			Question: who plays charlie in the santa clause movies? Greedy: Tim Allen		
Multinomial samples	Beam-search samples		Multinomial samples	Beam-search samples	
Euro	Euro	p=0.439	Tim Allen	Tim Allen	p=0.318
Euro	euro	p=0.201	Tim Allen	Jeff Daniels	p=0.017
Euro	Cyprus pound	p=0.091	Scott Calvin	Timothy Oly	p=0.012
Euro	Cypriot	p=0.072	Tim Allen	Ed Asner	p=0.010
euro	Cyprus Pound	p=0.016	Tim Allen	Scott Calvin	p=0.008
euro	Euros	p=0.014	Edward Arnold	Edward Asner	p=0.008
euro	EURO	p=0.007	Tim Allen	Tony Cox	p=0.007
euros	euros	p=0.007	Jeremy nault	Tim Allen	p=0.007
Euro	cyprus	p=0.007	Tim Allen	Tim allen	p=0.005
Euro	Cyprus	p=0.006	Tim Allen	Eric Lloyd	p=0.004
Dissimilarity: 0.976 Dissimilarity + beamsearch: 0.800			Dissimilarity: 0.427 Dissimilarity + beamsearch: 0.313		

Figure 11: Two examples from Gemma 3 4B base on WebQ. Each panel shows the question, greedy answer, multinomial and beam-search samples with autoregressive probabilities, plus dissimilarity and beamsearch-guided dissimilarity.

Story: A couple of weeks ago, my 12-year-old daughter, Ella threatened to take my phone and break it. "At night you'll always have your phone out and break you'll just type," Ella says. "I'm ready to go to bed, and try to get you to read stories for me and you're just standing there reading your texts and texting other people," she adds. I came to realize that I was ignoring her as a father...			Question: She mentions a lot of grown ups don't make what in their lifetime? Greedy: Limits.		
Multinomial samples	Beam-search samples		Multinomial samples	Beam-search samples	
Boundaries.	Boundaries.	p=0.185	Boundaries.	Boundaries.	p=0.185
Set limits.	Limits.	p=0.079	Set limits.	Limits.	p=0.079
Boundaries.	They don't	p=0.040	Boundaries.	They don't	p=0.040
limits.	Rules.	p=0.032	limits.	Rules.	p=0.032
Boundaries in their	Boundaries.	p=0.029	Boundaries in their	Boundaries.	p=0.029
Charging station.	Boundaries.	p=0.028	Charging station.	Boundaries.	p=0.028
Similar limitations.	A charging station.	p=0.016	Similar limitations.	A charging station.	p=0.016
Boundaries.	Boundaries that protect	p=0.010	Boundaries.	Boundaries that protect	p=0.010
Boundaries.	Limits in their own	p=0.007	Boundaries.	Limits in their own	p=0.007
Set up similar limits	Boundaries	p=0.007	Set up similar limits	Boundaries	p=0.007
Dissimilarity: 0.310 Dissimilarity + beamsearch: 0.190					

Figure 12: One example from Gemma 3 4B base on CoQA. Shown are the question, greedy answer, multinomial and beam-search samples with autoregressive probabilities, plus dissimilarity and beamsearch-guided dissimilarity.

C DATASETS

Table 9 lists the prompts used to form inputs for each dataset (separately for base and instruct models). Table 10 reports mean accuracy for each model–dataset pair. We measure accuracy as the fraction of predictions whose AlignScore with the gold answer exceeds 0.5.

Table 9: Prompt templates used for each dataset and model type. Few-shot exemplars are shown as placeholders (e.g., <5 few-shot QA pairs>); run-time inputs are denoted by <question>, <context>, <title 1>, etc.

Dataset	Base Prompt	Instruct Prompt
TriviaQA	<5 few-shot QA pairs> Question: <question> Answer:	Answer the following question as briefly as possible. <5 few-shot QA pairs> Now answer the following question: Question: <question> Answer:
Web Questions	<5 few-shot QA pairs> Question: <question> Answer:	Below are questions with short factual answers. Return only the short answer (a name, phrase, number, or year). <5 few-shot QA pairs> Now answer this. Q: <question> A:
CoQA	Story: <context> <all preceding QA pairs> Question: <question> Answer:	Story: <context> <all preceding QA pairs> Answer the following question as briefly as possible. Question: <question> Answer:
HotpotQA	Title: <title 1> <paragraph 1> Title: <title 2> <paragraph 2> Question: <question> Short answer:	Instruction: Read the context and answer with a short factual span (a few words) copied from the context. Reply with the short answer only. Title: <title 1> <paragraph 1> Title: <title 2> <paragraph 2> Question: <question> Short answer:
Common senseQA	<2 few-shot QA pairs> Question: <question> Options: <(A) - (D) options> Answer:	Instruction: Choose the single best answer from the options. Answer with the option text only (not the letter). <2 few-shot QA pairs> Now answer this. Question: <question> Options: <(A) - (D) options> Answer:
ARC-Challenge	<2 few-shot QA pairs> Question: <question> Options: <(A) - (D) options> Answer:	Instruction: Choose the single best answer from the options. Answer with the option text only (not the letter). <2 few-shot QA pairs> Now answer this. Question: <question> Options: <(A) - (D) options> Answer:

Table 10: Mean accuracy (%): proportion of predictions with AlignScore to the gold answer > 0.5.

	Closed-Book QA		Open-Book QA		Multiple Choice	
	TriviaQA	Web Questions	CoQA	HotpotQA	Common senseQA	ARC-Challenge
Llama 3.1 8B base	63%	47%	74%	53%	74%	72%
Llama 3.1 8B instruct	69%	40%	80%	72%	77%	76%
Gemma 3 4B base	47%	33%	69%	41%	65%	70%
Gemma 3 4B instruct	51%	35%	76%	66%	76%	77%
Qwen 3 8B base	52%	48%	81%	47%	89%	91%
Qwen 3 8B instruct	54%	42%	76%	76%	84%	88%

D ADDITIONAL RESULTS

D.1 SCORING TOP-BEAM OUTPUT

In the main text we score the greedy decode as the produced answer y_* . Table 11 complements these results by scoring the *top-1 beam* as y_* , a natural choice when beam search is already used to obtain a higher-quality decode. The beam-weighted family of approaches achieves higher PRR than the original methods and baselines in the majority of cases.

Table 11: PRR ($\uparrow$ is better) averaged over 6 datasets, when scoring the top-1 beam produced answer (instead of greedy). For each dataset, the top-1 score is **bold** and the second-best is underlined. For beam-guided variants, we mark $\uparrow$ when the variant improves over its original multinomial-sampling counterpart.

UQ Method	Llama 3.1 8B		Gemma 3 4B		Qwen 3 8B	
	base	instruct	base	instruct	base	instruct
<i>Baseline UQ methods</i>						
Prob	.399	.174	.400	.213	.390	.090
MTE	.320	.164	.317	.228	.334	.255
Perplexity	.376	.121	.359	.185	.318	.009
CCP	.395	.155	.369	.243	.352	.226
SAR	.333	.221	.336	.348	.342	.246
P(True)	.019	-.075	.031	.090	.012	-.080
SemanticEntropy	.345	.286	.397	.320	.299	.250
LexicalSimilarity	.377	.221	.384	.291	.404	.210
EigValLaplacian	.366	.209	.402	.307	.384	.223
NumSemSets	.349	.215	.365	.262	.344	.208
<i>Consistency-based UQ: multinomial vs. beamsearch versions</i>						
Dissimilarity	.437	.229	.424	.333	.446	.272
Dissimilarity + beamsearch	<u>.455</u> $\uparrow$	.266 $\uparrow$	<u>.466</u> $\uparrow$	.390 $\uparrow$	.440	.346 $\uparrow$
Eccentricity	.405	.238	.395	.310	.375	.208
Eccentricity + beamsearch	<u>.444</u> $\uparrow$	<u>.301</u> $\uparrow$	<u>.450</u> $\uparrow$	<u>.348</u> $\uparrow$	<u>.380</u> $\uparrow$	<u>.308</u> $\uparrow$
EigVecDissimilarity	.402	.243	.412	.316	.403	.213
EigVecDissimilarity + beamsearch	<u>.446</u> $\uparrow$	.316 $\uparrow$	<u>.457</u> $\uparrow$	<u>.366</u> $\uparrow$	<u>.415</u> $\uparrow$	<u>.334</u> $\uparrow$
CocoaMSP	.447	.284	.450	.347	<u>.454</u>	.272
CocoaMSP + beamsearch	.471 $\uparrow$	.290 $\uparrow$	.478 $\uparrow$	.407 $\uparrow$	.459 $\uparrow$	<u>.345</u> $\uparrow$
CocoaPPL	.440	.251	.433	.340	.422	.261
CocoaPPL + beamsearch	<u>.450</u> $\uparrow$	.273 $\uparrow$	<u>.444</u> $\uparrow$	<u>.395</u> $\uparrow$	.410	.318 $\uparrow$

D.2 ROC-AUC AND PR-AUC

In the main text we report PRR. Tables 12 and 13 complements these results with ROC-AUC and PR-AUC on Gemma 3 4B base. We binarize by marking an answer as correct if its AlignScore to the gold answer exceeds 0.5, and incorrect otherwise (the positive class for PR-AUC is the incorrect label). The pattern mirrors PRR: beam-guided variants generally match or outperform multinomial sampling.

D.3 DETAILED RESULTS FOR EACH DATASET

Complementing the main-table results in Table 3, Tables 14–19 report PRR for six datasets separately for Gemma 3 4B base, Gemma 3 4B instruct, Llama 3.1 8B base, Llama 3.1 8B instruct, Qwen 3 8B base, and Qwen 3 8B instruct.

Table 12: ROC-AUC $\uparrow$ for 6 datasets with Gemma 3 4B base. For each dataset, the top-1 method is **bold** and the second-best is underlined. Beam-guided and probability-weighted variants are marked with $\uparrow$ when they improve over their multinomial-sampling baseline. The two rightmost columns report the mean ROC-AUC across datasets.

UQ Method	TriviaQA	Web Questions	CoQA	HotpotQA	Common senceQA	ARC-Challenge	Mean
<i>Baseline UQ methods</i>							
Prob	.863	.768	.698	.632	.796	.821	.763
MTE	.867	.793	.710	.721	.737	.753	.763
Perplexity	.863	.785	.729	.735	.796	.820	.788
CCP	.881	.781	.698	.660	.775	.793	.764
SAR	.867	.776	.701	.713	.653	.696	.748
P(True)	.642	.473	.524	.513	.571	.545	.545
SemanticEntropy	.849	.758	.690	.591	.755	.774	.736
LexicalSimilarity	.842	.766	.713	.656	.739	.756	.745
EigVallaplacian	.867	.766	.701	.633	.739	.775	.747
NumSemSets	.856	.754	.653	.639	.702	.757	.727
<i>Consistency-based UQ: multinomial vs. beamsearch versions</i>							
Dissimilarity	.916	.836	.822	.809	.817	.818	.836
Dissimilarity + beamsearch	.923$\uparrow$	.852$\uparrow$	.826$\uparrow$	<u>.814$\uparrow$</u>	<u>.831$\uparrow$</u>	<u>.841$\uparrow$</u>	.848$\uparrow$
Eccentricity	.897	.808	.768	.737	.809	.821	.806
Eccentricity + beamsearch	.911 $\uparrow$	.816 $\uparrow$	.790 $\uparrow$	.771 $\uparrow$	.833$\uparrow$	.859$\uparrow$	.830 $\uparrow$
EigVecDissimilarity	.902	.813	.761	.728	.798	.825	.805
EigVecDissimilarity + beamsearch	<u>.920$\uparrow$</u>	.827 $\uparrow$	.787 $\uparrow$	.763 $\uparrow$	.820 $\uparrow$	<u>.856$\uparrow$</u>	.829 $\uparrow$
CocoaMSP	.904	.823	.791	.726	.826	.839	.818
CocoaMSP + beamsearch	.910 $\uparrow$	.836 $\uparrow$	.811 $\uparrow$	.779 $\uparrow$	.827 $\uparrow$	.847 $\uparrow$	.835 $\uparrow$
CocoaPPL	.907	.832	.810	.799	.825	.837	.835
CocoaPPL + beamsearch	.912 $\uparrow$	<u>.845$\uparrow$</u>	<u>.823$\uparrow$</u>	.828$\uparrow$	.825 $\uparrow$	.844 $\uparrow$	<u>.846$\uparrow$</u>

Table 13: PR-AUC $\uparrow$ for 6 datasets with Gemma 3 4B base. For each dataset, the top-1 method is **bold** and the second-best is underlined. Beam-guided and probability-weighted variants are marked with $\uparrow$ when they improve over their multinomial-sampling baseline. The two rightmost columns report the mean PR-AUC across datasets.

UQ Method	TriviaQA	Web Questions	CoQA	HotpotQA	Common senceQA	ARC-Challenge	Mean
<i>Baseline UQ methods</i>							
Prob	.855	.838	.477	.678	.623	.628	.683
MTE	.875	.874	.545	.799	.558	.540	.699
Perplexity	.860	.861	.539	.814	.629	.632	.722
CCP	.866	.853	.475	.715	.676	.641	.704
SAR	.865	.861	.484	.753	.437	.422	.646
P(True)	.657	.662	.326	.634	.410	.355	.507
SemanticEntropy	.838	.823	.456	.649	.572	.511	.642
LexicalSimilarity	.833	.848	.514	.711	.545	.509	.660
EigVallaplacian	.865	.855	.481	.682	.565	.532	.663
NumSemSets	.841	.825	.427	.685	.508	.497	.631
<i>Consistency-based UQ: multinomial vs. beamsearch versions</i>							
Dissimilarity	.911	.904	.715	.838	.722	.648	.789
Dissimilarity + beamsearch	.919$\uparrow$	.915$\uparrow$	.660	.822	.754$\uparrow$	.693 $\uparrow$	<u>.794$\uparrow$</u>
Eccentricity	.888	.887	.561	.758	.685	.625	.734
Eccentricity + beamsearch	.906 $\uparrow$	.884	.576 $\uparrow$	.789 $\uparrow$	<u>.744$\uparrow$</u>	.717$\uparrow$	.769 $\uparrow$
EigVecDissimilarity	.902	.889	.573	.766	.677	.651	.743
EigVecDissimilarity + beamsearch	<u>.916$\uparrow$</u>	.900 $\uparrow$	.588 $\uparrow$	.784 $\uparrow$	.717 $\uparrow$	.689 $\uparrow$	.766 $\uparrow$
CocoaMSP	.897	.894	.605	.761	.711	.680	.758
CocoaMSP + beamsearch	.907 $\uparrow$	.904 $\uparrow$	.632 $\uparrow$	.801 $\uparrow$	.715 $\uparrow$	.691 $\uparrow$	.775 $\uparrow$
CocoaPPL	.902	.902	.672	<u>.861</u>	.712	.686	.789
CocoaPPL + beamsearch	.909 $\uparrow$	<u>.910$\uparrow$</u>	<u>.690$\uparrow$</u>	.881$\uparrow$	.718 $\uparrow$	<u>.695$\uparrow$</u>	.801$\uparrow$

Table 14: PRR ($\uparrow$ is better) for 6 datasets with Gemma 3 4B base. For each dataset, the top-1 method is **bold** and the second-best is underlined. Beam-guided and probability-weighted variants are marked with $\uparrow$ when they improve over their multinomial-sampling baseline.

Method	TriviaQA	Web Questions	CoQA	HotpotQA	Common senceQA	ARC-Challenge
<i>Baseline UQ methods</i>						
Prob	.659 $\pm$ 0.018	.521 $\pm$ 0.031	.312 $\pm$ 0.024	.274 $\pm$ 0.014	.511 $\pm$ 0.025	.548 $\pm$ 0.077
MTE	.670 $\pm$ 0.013	.583 $\pm$ 0.029	.363 $\pm$ 0.02	.494 $\pm$ 0.034	.364 $\pm$ 0.031	.381 $\pm$ 0.052
Perplexity	.647 $\pm$ 0.024	.553 $\pm$ 0.022	.369 $\pm$ 0.02	.527 $\pm$ 0.023	.503 $\pm$ 0.022	.547 $\pm$ 0.062
CCP	.686 $\pm$ 0.021	.569 $\pm$ 0.031	.326 $\pm$ 0.022	.337 $\pm$ 0.025	.506 $\pm$ 0.034	.527 $\pm$ 0.062
SAR	.656 $\pm$ 0.02	.571 $\pm$ 0.028	.347 $\pm$ 0.023	.296 $\pm$ 0.018	.183 $\pm$ 0.037	.264 $\pm$ 0.055
P(True)	.272 $\pm$ 0.026	-.004 $\pm$ 0.034	.031 $\pm$ 0.026	.075 $\pm$ 0.025	.090 $\pm$ 0.028	.090 $\pm$ 0.048
SemanticEntropy	.622 $\pm$ 0.021	.505 $\pm$ 0.022	.301 $\pm$ 0.019	.140 $\pm$ 0.022	.407 $\pm$ 0.028	.431 $\pm$ 0.051
Lexical Similarity	.602 $\pm$ 0.017	.540 $\pm$ 0.032	.349 $\pm$ 0.025	.286 $\pm$ 0.016	.386 $\pm$ 0.032	.392 $\pm$ 0.054
EigValLaplacian	.666 $\pm$ 0.014	.555 $\pm$ 0.028	.320 $\pm$ 0.036	.246 $\pm$ 0.024	.386 $\pm$ 0.027	.452 $\pm$ 0.046
NumSemSets	.656 $\pm$ 0.017	.538 $\pm$ 0.028	.257 $\pm$ 0.027	.268 $\pm$ 0.019	.338 $\pm$ 0.03	.454 $\pm$ 0.042
<i>Consistency-based UQ: multinomial vs. beamsearch versions</i>						
Dissimilarity	.755 $\pm$ 0.019	.715 $\pm$ 0.03	.578 $\pm$ 0.022	.626 $\pm$ 0.016	.561 $\pm$ 0.04	.545 $\pm$ 0.062
Dissimilarity + beamsearch	.766 $\uparrow \pm$ 0.023	.722 $\uparrow \pm$ 0.028	.600 $\uparrow \pm$ 0.016	.611 $\pm$ 0.021	.595 $\uparrow \pm$ 0.028	.604 $\uparrow \pm$ 0.052
Eccentricity	.714 $\pm$ 0.012	.653 $\pm$ 0.029	.459 $\pm$ 0.02	.453 $\pm$ 0.026	.549 $\pm$ 0.034	.549 $\pm$ 0.054
Eccentricity + beamsearch	.739 $\uparrow \pm$ 0.019	.633 $\pm$ 0.035	.505 $\uparrow \pm$ 0.025	.514 $\uparrow \pm$ 0.027	.590 $\uparrow \pm$ 0.024	.636 $\uparrow \pm$ 0.066
EigVecDissimilarity	.738 $\pm$ 0.021	.661 $\pm$ 0.027	.443 $\pm$ 0.031	.448 $\pm$ 0.02	.512 $\pm$ 0.032	.562 $\pm$ 0.035
EigVecDissimilarity + beamsearch	.753 $\uparrow \pm$ 0.028	.668 $\uparrow \pm$ 0.032	.497 $\uparrow \pm$ 0.021	.487 $\uparrow \pm$ 0.016	.562 $\uparrow \pm$ 0.028	.621 $\uparrow \pm$ 0.06
CocoaMSP	.738 $\pm$ 0.023	.666 $\pm$ 0.028	.509 $\pm$ 0.021	.430 $\pm$ 0.028	.583 $\pm$ 0.03	.595 $\pm$ 0.052
CocoaMSP + beamsearch	.747 $\uparrow \pm$ 0.02	.679 $\uparrow \pm$ 0.02	.548 $\uparrow \pm$ 0.02	.523 $\uparrow \pm$ 0.027	.586 $\uparrow \pm$ 0.029	.606 $\uparrow \pm$ 0.072
CocoaPPL	.739 $\pm$ 0.015	.678 $\pm$ 0.025	.548 $\pm$ 0.019	.625 $\pm$ 0.023	.580 $\pm$ 0.031	.595 $\pm$ 0.039
CocoaPPL + beamsearch	.748 $\uparrow \pm$ 0.024	.694 $\uparrow \pm$ 0.024	.577 $\uparrow \pm$ 0.024	.681 $\uparrow \pm$ 0.019	.582 $\uparrow \pm$ 0.035	.610 $\uparrow \pm$ 0.048

Table 15: PRR ($\uparrow$ is better) for 6 datasets with Gemma 3 4B instruct. For each dataset, the top-1 method is **bold** and the second-best is underlined. Beam-guided and probability-weighted variants are marked with $\uparrow$ when they improve over their multinomial-sampling baseline.

UQ Method	TriviaQA	Web Questions	CoQA	HotpotQA	Common senceQA	ARC-Challenge
<i>Baseline UQ methods</i>						
Prob	.442 $\pm$.018	.425 $\pm$.031	.162 $\pm$.024	.220 $\pm$.014	.254 $\pm$.025	.252 $\pm$.077
MTE	.534 $\pm$.013	.465 $\pm$.029	.161 $\pm$.02	.232 $\pm$.034	.253 $\pm$.031	.256 $\pm$.052
Perplexity	.422 $\pm$.024	.419 $\pm$.022	.157 $\pm$.02	.223 $\pm$.023	.252 $\pm$.022	.256 $\pm$.062
CCP	.533 $\pm$.021	.478 $\pm$.031	.117 $\pm$.022	.303 $\pm$.025	.264 $\pm$.034	.290 $\pm$.062
SAR	.533 $\pm$.02	.426 $\pm$.028	.176 $\pm$.023	.214 $\pm$.018	.033 $\pm$.037	.050 $\pm$.055
P(True)	-.076 $\pm$.026	-.155 $\pm$.034	-.161 $\pm$.026	-.090 $\pm$.025	-.046 $\pm$.028	-.047 $\pm$.048
SemanticEntropy	.449 $\pm$.021	.415 $\pm$.022	.166 $\pm$.019	.223 $\pm$.022	.254 $\pm$.028	.252 $\pm$.051
Lexical Similarity	.527 $\pm$.017	.427 $\pm$.032	.176 $\pm$.025	.127 $\pm$.016	.052 $\pm$.032	.172 $\pm$.054
EigValLaplacian	.578 $\pm$.014	.472 $\pm$.028	.190 $\pm$.036	.134 $\pm$.024	.014 $\pm$.027	.010 $\pm$.046
NumSemSets	.556 $\pm$.017	.442 $\pm$.028	.123 $\pm$.027	.106 $\pm$.019	.046 $\pm$.03	.153 $\pm$.042
<i>Consistency-based UQ: multinomial vs. beamsearch versions</i>						
Dissimilarity	.549 $\pm$.019	.415 $\pm$.03	.111 $\pm$.022	.068 $\pm$.016	.024 $\pm$.04	.070 $\pm$.062
Dissimilarity + beamsearch	.413 $\pm$.023	.321 $\pm$.028	.204 $\uparrow \pm$.016	.273 $\uparrow \pm$.021	.218 $\uparrow \pm$.028	.085 $\uparrow \pm$.052
Eccentricity	.540 $\pm$.012	.429 $\pm$.029	.167 $\pm$.02	.175 $\pm$.026	-.020 $\pm$.034	.094 $\pm$.054
Eccentricity + beamsearch	.441 $\pm$.019	.367 $\pm$.035	.235 $\uparrow \pm$.025	.314 $\uparrow \pm$.027	.246 $\uparrow \pm$.024	.108 $\uparrow \pm$.066
EigVecDissimilarity	.561 $\pm$.021	.437 $\pm$.027	.169 $\pm$.031	.173 $\pm$.02	-.017 $\pm$.032	.095 $\pm$.035
EigVecDissimilarity + beamsearch	.478 $\pm$.028	.416 $\pm$.032	.240 $\uparrow \pm$.021	.308 $\uparrow \pm$.016	.253 $\uparrow \pm$.028	.113 $\uparrow \pm$.06
CocoaMSP	.531 $\pm$.023	.456 $\pm$.028	.183 $\pm$.021	.198 $\pm$.028	.252 $\pm$.03	.266 $\pm$.052
CocoaMSP + beamsearch	.535 $\uparrow \pm$.02	.473 $\uparrow \pm$.02	.237 $\uparrow \pm$.02	.287 $\uparrow \pm$.027	.282 $\uparrow \pm$.029	.258 $\pm$.072
CocoaPPL	.523 $\pm$.015	.454 $\pm$.025	.174 $\pm$.019	.201 $\pm$.023	.247 $\pm$.031	.271 $\pm$.039
CocoaPPL + beamsearch	.522 $\pm$.024	.467 $\uparrow \pm$.024	.222 $\uparrow \pm$.024	.285 $\uparrow \pm$.019	.277 $\uparrow \pm$.035	.264 $\pm$.048

Table 16: PRR ($\uparrow$ is better) for 6 datasets with Llama 3.1 8B base. For each dataset, the top-1 method is **bold** and the second-best is underlined. Beam-guided and probability-weighted variants are marked with $\uparrow$ when they improve over their multinomial-sampling baseline.

UQ Method	TriviaQA	Web Questions	CoQA	HotpotQA	Common senceQA	ARC-Challenge
<i>Baseline UQ methods</i>						
Prob	.517 $\pm$.019	.414 $\pm$.029	.310 $\pm$.022	.213 $\pm$.024	.504 $\pm$.029	.505 $\pm$.043
MTE	.544 $\pm$.018	.420 $\pm$.015	.286 $\pm$.022	.327 $\pm$.02	.448 $\pm$.029	.511 $\pm$.055
Perplexity	.507 $\pm$.015	.441 $\pm$.027	.316 $\pm$.031	.375 $\pm$.018	.501 $\pm$.027	.570 $\pm$.047
CCP	.575 $\pm$.016	.420 $\pm$.026	.276 $\pm$.024	.247 $\pm$.029	.442 $\pm$.023	.446 $\pm$.031
SAR	.548 $\pm$.017	.452 $\pm$.028	.331 $\pm$.03	.263 $\pm$.031	.189 $\pm$.021	.330 $\pm$.044
P(True)	-.055 $\pm$.021	.059 $\pm$.023	-.020 $\pm$.018	-.223 $\pm$.026	.034 $\pm$.024	.292 $\pm$.044
SemanticEntropy	.538 $\pm$.019	.409 $\pm$.023	.330 $\pm$.021	.199 $\pm$.024	.492 $\pm$.023	.514 $\pm$.05
Lexical Similarity	.467 $\pm$.018	.396 $\pm$.03	.366 $\pm$.024	.289 $\pm$.026	.437 $\pm$.028	.511 $\pm$.041
EigValLaplacian	.569 $\pm$.019	.418 $\pm$.022	.377 $\pm$.023	.247 $\pm$.025	.449 $\pm$.035	.499 $\pm$.047
NumSemSets	.550 $\pm$.014	.409 $\pm$.033	.319 $\pm$.019	.241 $\pm$.028	.378 $\pm$.025	.477 $\pm$.044
<i>Consistency-based UQ: multinomial vs. beamsearch versions</i>						
Dissimilarity	.576 $\pm$.02	.445 $\pm$.024	.473 $\pm$.023	.446 $\pm$.02	.449 $\pm$.028	.640 $\pm$.056
Dissimilarity + beamsearch	.654 $\uparrow \pm$.017	.504 $\uparrow \pm$.023	.485 $\uparrow \pm$.019	.424 $\pm$.024	.510 $\pm$.023	.683 $\uparrow \pm$.044
Eccentricity	.555 $\pm$.016	.404 $\pm$.025	.405 $\pm$.023	.297 $\pm$.021	.464 $\pm$.028	.591 $\pm$.038
Eccentricity + beamsearch	.613 $\uparrow \pm$.021	.458 $\uparrow \pm$.019	.429 $\uparrow \pm$.017	.361 $\uparrow \pm$.023	.512 $\pm$.025	.657 $\uparrow \pm$.031
EigVecDissimilarity	.570 $\pm$.015	.452 $\pm$.022	.409 $\pm$.019	.289 $\pm$.02	.469 $\pm$.04	.587 $\pm$.038
EigVecDissimilarity + beamsearch	.630 $\uparrow \pm$.019	.492 $\uparrow \pm$.022	.427 $\uparrow \pm$.019	.357 $\uparrow \pm$.02	.506 $\pm$.035	.650 $\uparrow \pm$.032
CocoaMSP	.595 $\pm$.013	.458 $\pm$.021	.463 $\pm$.023	.366 $\pm$.021	.510 $\pm$.028	.641 $\pm$.038
CocoaMSP + beamsearch	<u>.631</u> $\uparrow \pm$.019	.487 $\pm$.023	.465 $\pm$.027	.372 $\pm$.027	.532 $\uparrow \pm$.022	.639 $\pm$.041
CocoaPPL	.587 $\pm$.017	.464 $\pm$.024	.464 $\pm$.023	.465 $\pm$.02	.501 $\pm$.031	.660 $\pm$.034
CocoaPPL + beamsearch	.616 $\uparrow \pm$.016	<u>.498</u> $\uparrow \pm$.029	.459 $\pm$.024	<u>.456</u> $\pm$.018	<u>.525</u> $\uparrow \pm$.028	<u>.661</u> $\uparrow \pm$.046

Table 17: PRR ($\uparrow$ is better) for 6 datasets with Llama 3.1 8B instruct. For each dataset, the top-1 method is **bold** and the second-best is underlined. Beam-guided and probability-weighted variants are marked with $\uparrow$ when they improve over their multinomial-sampling baseline.

UQ Method	TriviaQA	Web Questions	CoQA	HotpotQA	Common senceQA	ARC-Challenge
<i>Baseline UQ methods</i>						
Prob	.524 $\pm$.023	.357 $\pm$.036	.327 $\pm$.021	.213 $\pm$.022	.283 $\pm$.026	.363 $\pm$.044
MTE	.604 $\pm$.015	.424 $\pm$.028	.307 $\pm$.02	.253 $\pm$.031	.260 $\pm$.027	.339 $\pm$.055
Perplexity	.498 $\pm$.018	.367 $\pm$.025	.262 $\pm$.025	.221 $\pm$.03	.255 $\pm$.028	.332 $\pm$.053
CCP	.576 $\pm$.023	.406 $\pm$.028	.291 $\pm$.018	.265 $\pm$.022	.248 $\pm$.034	.402 $\pm$.048
SAR	.599 $\pm$.021	.420 $\pm$.029	.338 $\pm$.024	.236 $\pm$.02	.301 $\pm$.025	.418 $\pm$.04
P(True)	.236 $\pm$.023	.012 $\pm$.031	.018 $\pm$.035	.045 $\pm$.024	-.011 $\pm$.024	.135 $\pm$.051
SemanticEntropy	.591 $\pm$.016	.381 $\pm$.027	.335 $\pm$.032	.231 $\pm$.029	.301 $\pm$.038	.418 $\pm$.061
Lexical Similarity	.566 $\pm$.023	.395 $\pm$.029	.347 $\pm$.024	.232 $\pm$.03	.275 $\pm$.032	.380 $\pm$.045
EigValLaplacian	.615 $\pm$.021	.389 $\pm$.026	.355 $\pm$.029	.238 $\pm$.023	.252 $\pm$.029	.377 $\pm$.051
NumSemSets	.569 $\pm$.021	.363 $\pm$.031	.228 $\pm$.03	.180 $\pm$.023	.208 $\pm$.035	.368 $\pm$.051
<i>Consistency-based UQ: multinomial vs. beamsearch versions</i>						
Dissimilarity	.616 $\pm$.016	.382 $\pm$.031	.349 $\pm$.018	.270 $\pm$.021	.277 $\pm$.037	.378 $\pm$.061
Dissimilarity + beamsearch	.662 $\uparrow \pm$.015	.411 $\uparrow \pm$.029	.358 $\uparrow \pm$.029	.349 $\uparrow \pm$.019	.288 $\pm$.032	.434 $\uparrow \pm$.054
Eccentricity	.598 $\pm$.021	.379 $\pm$.032	.319 $\pm$.016	.248 $\pm$.031	.273 $\pm$.035	.389 $\pm$.058
Eccentricity + beamsearch	.620 $\uparrow \pm$.016	.396 $\uparrow \pm$.027	.330 $\uparrow \pm$.021	.281 $\uparrow \pm$.021	.306 $\uparrow \pm$.031	<u>.451</u> $\uparrow \pm$.047
EigVecDissimilarity	.611 $\pm$.019	.378 $\pm$.033	.325 $\pm$.025	.249 $\pm$.029	.264 $\pm$.037	.390 $\pm$.061
EigVecDissimilarity + beamsearch	.640 $\uparrow \pm$.017	.425 $\pm$.028	.347 $\pm$.027	.291 $\pm$.022	.318 $\uparrow \pm$.034	.461 $\uparrow \pm$.046
CocoaMSP	.629 $\pm$.018	.409 $\pm$.023	.366 $\pm$.03	.278 $\pm$.02	<u>.314</u> $\pm$.029	.426 $\pm$.051
CocoaMSP + beamsearch	.665 $\uparrow \pm$.016	.428 $\uparrow \pm$.029	.378 $\uparrow \pm$.017	<u>.344</u> $\uparrow \pm$.019	.302 $\pm$.036	.439 $\pm$.041
CocoaPPL	.626 $\pm$.022	.410 $\pm$.03	.354 $\pm$.024	.278 $\pm$.024	.299 $\pm$.038	.413 $\pm$.056
CocoaPPL + beamsearch	.653 $\uparrow \pm$.018	<u>.427</u> $\uparrow \pm$.032	.356 $\uparrow \pm$.021	.334 $\uparrow \pm$.018	.285 $\pm$.04	.419 $\uparrow \pm$.056

Table 18: PRR ($\uparrow$ is better) for 6 datasets with Qwen 3 8B base. For each dataset, the top-1 method is **bold** and the second-best is underlined. Beam-guided and probability-weighted variants are marked with $\uparrow$ when they improve over their multinomial-sampling baseline.

UQ Method	TriviaQA	Web Questions	CoQA	HotpotQA	Common senseQA	ARC-Challenge
<i>Baseline UQ methods</i>						
Prob	.617 $\pm$.017	.449 $\pm$.025	.267 $\pm$.025	.111 $\pm$.033	.337 $\pm$.039	.475 $\pm$.085
MTE	.602 $\pm$.022	.409 $\pm$.027	.267 $\pm$.023	.279 $\pm$.023	.443 $\pm$.044	.444 $\pm$.077
Perplexity	.597 $\pm$.018	.426 $\pm$.028	.278 $\pm$.023	.256 $\pm$.026	.294 $\pm$.045	.381 $\pm$.065
CCP	.640 $\pm$.018	.406 $\pm$.028	.213 $\pm$.028	.153 $\pm$.025	.296 $\pm$.048	.421 $\pm$.09
SAR	.617 $\pm$.023	.457 $\pm$.023	.323 $\pm$.03	.243 $\pm$.028	.220 $\pm$.042	.317 $\pm$.066
P(True)	.322 $\pm$.021	.282 $\pm$.025	.005 $\pm$.031	.168 $\pm$.024	-.043 $\pm$.045	-.074 $\pm$.069
Semantic Entropy	.549 $\pm$.018	.411 $\pm$.02	.247 $\pm$.025	.173 $\pm$.023	.230 $\pm$.026	.305 $\pm$.058
Lexical Similarity	.595 $\pm$.023	.430 $\pm$.024	.338 $\pm$.019	.310 $\pm$.025	.367 $\pm$.042	.508 $\pm$.076
EigValLaplacian	.602 $\pm$.015	.423 $\pm$.027	.301 $\pm$.027	.284 $\pm$.028	.349 $\pm$.032	.475 $\pm$.081
NumSemSets	.593 $\pm$.016	.403 $\pm$.029	.268 $\pm$.024	.250 $\pm$.023	.311 $\pm$.039	.367 $\pm$.069
<i>Consistency-based UQ: multinomial vs. beamsearch versions</i>						
Dissimilarity	.668 $\pm$.014	.462 $\pm$.024	<u>.406</u> $\pm$.023	.531 $\pm$.017	.315 $\pm$.038	.476 $\pm$.086
Dissimilarity + beamsearch	.680 $\uparrow \pm$.019	.484 $\uparrow \pm$.024	.409 $\uparrow \pm$.03	<u>.504</u> $\pm$.019	.335 $\uparrow \pm$.044	.457 $\pm$.088
Eccentricity	.615 $\pm$.016	.416 $\pm$.023	.320 $\pm$.022	.319 $\pm$.024	.266 $\pm$.053	.440 $\pm$.068
Eccentricity + beamsearch	.640 $\uparrow \pm$.013	.437 $\uparrow \pm$.025	.368 $\uparrow \pm$.02	.407 $\uparrow \pm$.026	.243 $\pm$.04	.366 $\pm$.072
EigVecDissimilarity	.628 $\pm$.014	.454 $\pm$.028	.325 $\pm$.026	.314 $\pm$.027	.373 $\pm$.035	.456 $\pm$.071
EigVecDissimilarity + beamsearch	.660 $\uparrow \pm$.016	.460 $\uparrow \pm$.024	.380 $\uparrow \pm$.025	.394 $\uparrow \pm$.025	.353 $\pm$.045	.453 $\pm$.086
CocoaMSP	.667 $\pm$.019	<u>.492</u> $\pm$.019	.385 $\pm$.025	.320 $\pm$.025	.378 $\pm$.028	.523 $\pm$.071
CocoaMSP + beamsearch	<u>.678</u> $\uparrow \pm$.018	.498 $\uparrow \pm$.028	.391 $\uparrow \pm$.02	.378 $\uparrow \pm$.029	<u>.385</u> $\pm$.037	<u>.510</u> $\pm$.079
CocoaPPL	.665 $\pm$.015	.478 $\pm$.019	.388 $\pm$.035	.397 $\pm$.024	.353 $\pm$.038	.484 $\pm$.081
CocoaPPL + beamsearch	.667 $\uparrow \pm$.016	.486 $\pm$.021	.387 $\pm$.036	.437 $\uparrow \pm$.026	.339 $\pm$.044	.450 $\pm$.06

Table 19: PRR ($\uparrow$ is better) for 6 datasets with Qwen 3 8B instruct. For each dataset, the top-1 method is **bold** and the second-best is underlined. Beam-guided and probability-weighted variants are marked with $\uparrow$ when they improve over their multinomial-sampling baseline.

UQ Method	TriviaQA	Web Questions	CoQA	HotpotQA	Common senseQA	ARC-Challenge
<i>Baseline UQ methods</i>						
Prob	.564 $\pm$.017	.353 $\pm$.032	.215 $\pm$.02	.250 $\pm$.026	.174 $\pm$.034	.181 $\pm$.078
MTE	.564 $\pm$.018	.345 $\pm$.028	.164 $\pm$.025	.251 $\pm$.028	.183 $\pm$.03	.272 $\pm$.095
Perplexity	.491 $\pm$.023	.341 $\pm$.036	.169 $\pm$.026	.250 $\pm$.028	.175 $\pm$.037	.229 $\pm$.058
CCP	.563 $\pm$.02	.383 $\pm$.029	.169 $\pm$.018	.258 $\pm$.029	.173 $\pm$.034	.202 $\pm$.068
SAR	.590 $\pm$.016	.425 $\pm$.036	.146 $\pm$.026	.159 $\pm$.029	.201 $\pm$.033	.233 $\pm$.051
P(True)	-.105 $\pm$.023	-.222 $\pm$.035	-.126 $\pm$.017	.018 $\pm$.021	-.083 $\pm$.03	-.164 $\pm$.071
Semantic Entropy	.597 $\pm$.016	.404 $\pm$.034	.214 $\pm$.022	.231 $\pm$.026	.174 $\pm$.041	.176 $\pm$.08
Lexical Similarity	.530 $\pm$.023	.425 $\pm$.029	.193 $\pm$.031	.101 $\pm$.026	.121 $\pm$.039	.053 $\pm$.06
EigValLaplacian	.626 $\pm$.015	.417 $\pm$.04	.196 $\pm$.026	.083 $\pm$.024	.134 $\pm$.031	.134 $\pm$.066
NumSemSets	.608 $\pm$.021	<u>.437</u> $\pm$.036	.110 $\pm$.019	.096 $\pm$.024	.113 $\pm$.041	.154 $\pm$.065
<i>Consistency-based UQ: multinomial vs. beamsearch versions</i>						
Dissimilarity	.588 $\pm$.017	.382 $\pm$.03	.165 $\pm$.02	.187 $\pm$.025	.246 $\pm$.038	.394 $\pm$.072
Dissimilarity + beamsearch	<u>.637</u> $\uparrow \pm$.018	.386 $\pm$.026	.269 $\uparrow \pm$.019	.264 $\uparrow \pm$.026	.213 $\pm$.031	<u>.362</u> $\pm$.083
Eccentricity	.565 $\pm$.019	.367 $\pm$.034	.167 $\pm$.025	.125 $\pm$.023	.150 $\pm$.026	.132 $\pm$.078
Eccentricity + beamsearch	.600 $\uparrow \pm$.016	.392 $\uparrow \pm$.034	<u>.288</u> $\uparrow \pm$.029	<u>.291</u> $\uparrow \pm$.022	.211 $\uparrow \pm$.035	.285 $\uparrow \pm$.084
EigVecDissimilarity	.590 $\pm$.024	.385 $\pm$.031	.169 $\pm$.026	.121 $\pm$.032	.143 $\pm$.033	.131 $\pm$.066
EigVecDissimilarity + beamsearch	.645 $\uparrow \pm$.016	.439 $\uparrow \pm$.032	.328 $\uparrow \pm$.019	.297 $\uparrow \pm$.017	<u>.242</u> $\pm$.029	.306 $\uparrow \pm$.058
CocoaMSP	.607 $\pm$.015	.394 $\pm$.03	.204 $\pm$.016	.272 $\pm$.023	.230 $\pm$.042	.298 $\pm$.061
CocoaMSP + beamsearch	.635 $\pm$.02	.404 $\uparrow \pm$.024	.263 $\pm$.023	.282 $\pm$.025	.206 $\pm$.029	.290 $\pm$.061
CocoaPPL	.581 $\pm$.02	.389 $\pm$.032	.179 $\pm$.024	.272 $\pm$.022	.232 $\pm$.032	.309 $\pm$.082
CocoaPPL + beamsearch	.609 $\uparrow \pm$.02	.395 $\pm$.031	.233 $\pm$.025	.282 $\pm$.026	.207 $\pm$.03	.299 $\pm$.084

E DETAILED DESCRIPTION OF UNCERTAINTY QUANTIFICATION METHODS

In this section, we describe the uncertainty quantification methods used in our experiments.

Sequence Probability (Prob) is the most straightforward approach to uncertainty quantification. We define it formally as the negative log-probability of the generating sequence:

$$U_{SP}(\mathbf{y} \mid \mathbf{x}) = -\log P(\mathbf{y} \mid \mathbf{x}). \quad (27)$$

Mean Token Entropy (MTE) measures an average entropy of tokens in a sequence:

$$U_{MTE}(\mathbf{y} \mid \mathbf{x}) = \frac{1}{L} \sum_{l=1}^L \mathcal{H}(y_l \mid \mathbf{y}_{<l}, \mathbf{x}), \quad (28)$$

where $\mathcal{H}(y_l \mid \mathbf{y}_{<l}, \mathbf{x}) = -\sum_v P(y_l = v \mid \mathbf{y}_{<l}, \mathbf{x}) \log P(y_l = v \mid \mathbf{y}_{<l}, \mathbf{x})$.

Perplexity computes negative average log-likelihood of tokens in a sequence:

$$U_{PPL}(\mathbf{y} \mid \mathbf{x}) = -\frac{1}{L} \log P(\mathbf{y} \mid \mathbf{x}), \quad (29)$$

Claim Conditioned Probability (CCP), introduced in (Fadeeva et al., 2024), measures uncertainty on a claim level by perturbing claim’s tokens with alternative generations:

$$U_{CCP}(C \mid \mathbf{x}) = 1 - \prod_{j \in C} \text{CCP}(y_j \mid \mathbf{y}_{<j}, \mathbf{x}). \quad (30)$$

$$\text{Where } \text{CCP}(y_j \mid \mathbf{y}_{<j}, \mathbf{x}) = \frac{\sum_{k: \text{NLI}(y_j^k, y_j) = \text{true}} P(y_j^k \mid \mathbf{y}_{<j}, \mathbf{x})}{\sum_{k: \text{NLI}(y_j^k, y_j) \in \{\text{true}, \text{false}\}} P(y_j^k \mid \mathbf{y}_{<j}, \mathbf{x})}$$

Shifting Attention to Relevance (SAR) is a method combining TokenSAR and SentenceSAR, as introduced by Duan et al. (2024). SentenceSAR is defined as follows:

$$U_{\text{SentSAR}}(\mathbf{x}) = -\frac{1}{M} \sum_{i=1}^M \log \left(p(\mathbf{y}^{(i)} \mid \mathbf{x}) + \frac{1}{t} R_S(\mathbf{y}^{(i)}, \mathbf{x}) \right), \quad (31)$$

Here, $R_S(\mathbf{y}^{(j)}, \mathbf{x}) = \sum_{k \neq j} s(\mathbf{y}^{(j)}, \mathbf{y}^{(k)}) p(\mathbf{y}^{(k)} \mid \mathbf{x})$. To obtain SAR score, the generative probability $p(\mathbf{y} \mid \mathbf{x})$ is replaced with relevance-reweighted probability on a sequence level. *TokenSAR* is defined as:

$$U_{\text{TokenSAR}}(\mathbf{x}) = -\sum_{l=1}^L \tilde{R}_T(y_l, \mathbf{y}, \mathbf{x}) \log P(y_l \mid \mathbf{y}_{<l}, \mathbf{x}), \quad (32)$$

where $R_T(\cdot)$ denotes some token relevance function and relevance weight for token y_l is given by $\tilde{R}_T(y_k, \mathbf{y}, \mathbf{x}) = \frac{R_T(y_k, \mathbf{y}, \mathbf{x})}{\sum_{l=1}^L R_T(y_l, \mathbf{y}, \mathbf{x})}$.

P(True), introduced in (Kadavath et al., 2022), evaluates the confidence in a generation by asking the model the original question and answer, then asking if it is true or false. We then use the negative log-probability of the token “True” as an uncertainty score.

Lexical Similarity, introduced in (Fomicheva et al., 2020), measures average pairwise similarity between M sampled generations using some similarity function $s(\mathbf{y}, \mathbf{y}')$:

$$U_{\text{LSRL}}(\mathbf{x}) = 1 - \frac{2}{M(M-1)} \sum_{i < j} s(\mathbf{y}^{(i)}, \mathbf{y}^{(j)}). \quad (33)$$

Number of Semantic Sets, introduced in (Lin et al., 2024), estimates how many distinct meanings the model produces by clustering its outputs with an NLI model. Two answers are placed in the same cluster if they mutually entail each other more than they contradict and the final number of distinct clusters serves as an uncertainty score $U_{\text{NumSemSets}}$.

Sum of Eigenvalues of Laplacian, introduced in (Lin et al., 2024), constructs a similarity matrix among the sampled outputs and computes a uncertainty score from the eigenvalues of the Laplacian of that similarity matrix:

$$U_{\text{EigV}}(\mathbf{x}) = \sum_{i=1}^M \max(0, 1 - \lambda_i(\mathbf{x})). \quad (34)$$

F COMPUTATIONAL BUDGET

All experiments were run on $2 \times \text{NVIDIA A100 (80 GB)}$. Evaluating a single model across all six datasets took approximately 2 wall-clock days on this setup (4 GPU-days); with six models, this amounts to 12 wall-clock days (24 GPU-days). Additional ablations (sampling strategies, top-1 beam scoring, and other objectives) required a further 5 wall-clock days on the same hardware (10 GPU-days). In total, the study used about 34 GPU-days.

G THE USAGE OF LLMs

In this study, large language models are examined primarily as the focus of analysis. For practical tasks such as programming and writing, we also make limited use of LLM-based assistants (e.g., ChatGPT) to support grammar correction and code debugging, with all usage carefully monitored by humans.