# Don't Throw Away Your Beams: Improving Consistency-based Uncertainties in LLMs via Beam Search

# **Anonymous authors**

Paper under double-blind review

# **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 separated into several distinct categories. 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);

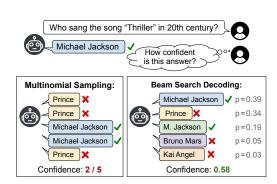


Figure 1: Beam Search vs Multinomial Sampling. Sampling produces multiple identical generations resulting in noisy confidence estimate, while beam search covers top answers from LLM distribution resulting in a better confidence score.

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). 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:

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

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

**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 Consistency-Based UQ Methods and Their Challenges

The objective of uncertainty quantification for LLMs is to measure the level of uncertainty in a particular generation  $y_*$  conditioned on the input sequence x, which we denote as  $U(y_* \mid x)$ .

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:

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

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  $U_D(\mathbf{y}_* \mid \mathbf{x})$  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  $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(M)} \sim p(\cdot \mid \mathbf{x})$  and computes uncertainty in the following way:

$$\widehat{U}_D^{MC}(\mathbf{y}_* \mid \mathbf{x}) = \frac{1}{M} \sum_{i=1}^{M} \left( 1 - s(\mathbf{y}^{(i)}, \mathbf{y}_*) \right). \tag{3}$$

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.

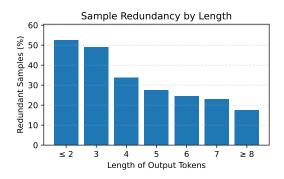


Figure 2: Mean percentage of redundant samples (i.e., outputs already seen among earlier generations) as a function of greedy output length. Results were obtained from 2,000 questions from the TriviaQA dataset using the Gemma 3 4B base model and 10 candidate generations. Redundancy is especially high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

# 3 UNCERTAINTY QUANTIFICATION BASED ON CONSISTENCY OF BEAM SEARCH CANDIDATES

To address the problems outlined in Section 2.2, 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  $\{p(\mathbf{b}^{(i)} \mid \mathbf{x})\}_{i=1}^M$ . To perform an estimation

of  $U_D(\mathbf{y}_* \mid \mathbf{x})$  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:

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

$$(4)$$

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

$$\widehat{U}_{D}^{b}(\mathbf{y}_{*} \mid \mathbf{x}) = \sum_{i=1}^{M} w_{i} (1 - s(\mathbf{b}^{(i)}, \mathbf{y}_{*})).$$
 (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  $\hat{U}_D^b(\mathbf{y}_* \mid \mathbf{x})$  from a theoretical perspective.

# 3.2 THEORETICAL ANALYSIS

We compare the multinomial Monte Carlo estimator (3) with the beam-weighted estimator (5) for the dissimilarity  $U_D(\mathbf{y}_* \mid \mathbf{x})$  defined in equation (2). Let  $\mathcal{B}_M(\mathbf{x}) = \{\mathbf{b}^{(1)}, \dots, \mathbf{b}^{(M)}\}$  be the beam set and  $m_{\mathcal{B}} = \sum_{i=1}^M p(\mathbf{b}^{(i)} \mid \mathbf{x})$  its total probability mass. Define

$$\mu_{\mathcal{B}} = \mathbb{E}_{\mathbf{y} \sim p(\cdot | \mathbf{x})} \left[ 1 - s(\mathbf{y}, \mathbf{y}_*) \mid \mathbf{y} \in \mathcal{B}_M(\mathbf{x}) \right], \qquad \mu_{\overline{\mathcal{B}}} = \mathbb{E}_{\mathbf{y} \sim p(\cdot | \mathbf{x})} \left[ 1 - s(\mathbf{y}, \mathbf{y}_*) \mid \mathbf{y} \notin \mathcal{B}_M(\mathbf{x}) \right].$$

**Proposition 1** (Monte Carlo). Let  $\widehat{U}_D^{MC}$  be the average of M i.i.d. multinomial samples. Then

$$\mathbb{E}[\widehat{U}_D^{MC}] = U_D(\mathbf{y}_* \mid \mathbf{x}), \quad \text{MSE}(\widehat{U}_D^{MC}) = \text{Var}(\widehat{U}_D^{MC}) = \sigma^2/M,$$
where  $\sigma^2 = \text{Var}_{\mathbf{y} \sim p(\cdot \mid \mathbf{x})} (1 - s(\mathbf{y}, \mathbf{y}_*)) \leq \frac{1}{4} \text{ since } s(\mathbf{y}, \mathbf{y}_*) \in [0, 1].$ 

**Proposition 2** (Beam-weighted). Let  $\widehat{U}_D^b$  be the beam-weighted estimator using  $\mathcal{B}_M(\mathbf{x})$ . Then

$$U_D(\mathbf{y}_* \mid \mathbf{x}) = m_{\mathcal{B}}\mu_{\mathcal{B}} + (1 - m_{\mathcal{B}})\mu_{\overline{\mathcal{B}}}, \qquad \widehat{U}_D^b = \mu_{\mathcal{B}},$$

so the estimator is deterministic with squared error

$$SE(\widehat{U}_D^b) = (\widehat{U}_D^b - U_D)^2 = (1 - m_{\mathcal{B}})^2 (\mu_{\mathcal{B}} - \mu_{\overline{\mathcal{B}}})^2.$$

**Theorem 1** (Comparison condition). The squared error of the beam-weighted estimator is smaller than the mean-squared error of the Monte Carlo estimator whenever

$$(1 - m_{\mathcal{B}}) \left| \mu_{\mathcal{B}} - \mu_{\overline{\mathcal{B}}} \right| < \sigma / \sqrt{M}.$$

A distribution-free sufficient condition follows from  $\left|\mu_{\mathcal{B}} - \mu_{\overline{\mathcal{B}}}\right| \leq 1$  and  $\sigma^2 \leq \frac{1}{4}$ :

$$m_{\mathcal{B}} > 1 - \frac{1}{2\sqrt{M}}$$
.

**Discussion.** For M=10, the sufficient condition becomes  $m_{\mathcal{B}}>0.842$ . Short-form QA often concentrates most probability mass on the top few beams, so the condition is frequently met. Figure 3 shows, on the TriviaQA dataset, that the percentage of texts satisfying this condition is 22.7% and especially high for short generations (30-40%) for  $\leq 3$  output tokens).

The bound is conservative: when the inside-outside gap  $\delta = |\mu_{\mathcal{B}} - \mu_{\overline{\mathcal{B}}}| < 1$ , the break-even requirement relaxes to  $(1-m_{\mathcal{B}})\delta < \sigma/\sqrt{M}$ , allowing the beam-weighted estimator to outperform Monte Carlo even when  $m_{\mathcal{B}} < 0.842$ . In practice,  $\mu_{\overline{\mathcal{B}}}$  is not directly computable due

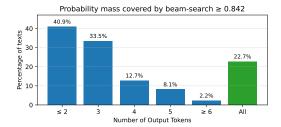


Figure 3: Percentage of texts meeting the distribution-free condition (Theorem 1). Results are based on 2,000 TriviaQA questions, Gemma 3 4B base and M=10. The green "All" bar shows the overall percentage across all lengths.

to the combinatorial output space, so  $\delta$  cannot be measured on real data; nevertheless, our experiments consistently show beam search outperforming multinomial sampling, suggesting that  $\delta$  is modest in practice and thus further loosens the conservative threshold.

# 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  $y_*$  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  $\mathbf{y}^{(M+1)} = \mathbf{y}_*$ :

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

$$(6)$$

Then we compute the degree matrix D:

$$D_{ij} = \begin{cases} \sum_{k=1}^{M+1} W_{ik}, & i = j, \\ 0, & i \neq j, \end{cases}$$
 (7)

and obtain the eigendecomposition of the Graph Laplacian  $L=I-D^{-1/2}WD^{-1/2}$ , 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:

$$\widehat{U}_{Ecc}(\mathbf{y}_* \mid \mathbf{x}) = \left\| \mathbf{v}_* - \frac{1}{M} \sum_{i=1}^{M} \mathbf{v}_i \right\|_2^2, \tag{8}$$

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:

$$\widehat{U}_{Ecc}^{b}(\mathbf{y}_* \mid \mathbf{x}) = \left\| \mathbf{v}_*^b - \sum_{i=1}^M w_i \mathbf{v}_i^b \right\|_2^2.$$
(9)

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

$$\widehat{U}_{CoCoA}(\mathbf{y}_* \mid \mathbf{x}) = u(\mathbf{y}_* \mid \mathbf{x}) \cdot \widehat{U}_D^{MC}(\mathbf{y}_* \mid \mathbf{x}) = u(\mathbf{y}_* \mid \mathbf{x}) \cdot \frac{1}{M} \sum_{i=1}^{M} (1 - s(\mathbf{y}^{(i)}, \mathbf{y}_*)), \quad (10)$$

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

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

$$\widehat{U}_{CoCoA}^{b}(\mathbf{y}_{*} \mid \mathbf{x}) = u(\mathbf{y}_{*} \mid \mathbf{x}) \cdot \widehat{U}_{D}^{b}(\mathbf{y}_{*} \mid \mathbf{x})$$
(11)

**Eigenvectors Dissimilarity.** Both Dissimilarity and Eccentricity produce confidence scores for the generated answer  $y_*$ . Dissimilarity compares  $y_*$  to each sample using the base similarity function s, while Eccentricity measures the distance from  $y_*$  to the centroid in the Laplacian embedding space; see equation (8). To bridge these views, we measure dissimilarity within the embedding space itself, averaging the distances from the embedding of  $y_*$  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

$$\widehat{U}_{EigVecD}(\mathbf{y}_* \mid \mathbf{x}) = \frac{1}{M} \sum_{i=1}^{M} \left\| \mathbf{v}_* - \mathbf{v}_i \right\|_2^2, \tag{12}$$

Table 1: Test dataset settings and statistics.

	Closed-l	Closed-Book QA		ook QA	Multiple Choice		
	TriviaQA	Web Questions	CoQA	HotpotQA	Common senceQA	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	

and the beam-guided, probability-weighted version is

$$\widehat{U}_{EigVecD}^{b}(\mathbf{y}_{*} \mid \mathbf{x}) = \sum_{i=1}^{M} w_{i} \|\mathbf{v}_{*}^{b} - \mathbf{v}_{i}^{b}\|_{2}^{2},$$
(13)

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: *CommonsenceQA* (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.

**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  $\mathcal{D} = \{(\mathbf{x}_j, \mathbf{y}_j^*)\}$ , where  $\mathbf{y}_j$  denotes the output generated by an LLM for input  $\mathbf{x}_j$ , and  $u_j = U(\mathbf{x}_j)$  is the associated uncertainty score. The rejection curve captures how the average quality  $Q(\mathbf{y}_j, \mathbf{y}_j^*)$  of predictions with uncertainty  $u_j < a$  varies with the rejection parameter a. PRR is then defined as the normalized area under the rejection curve, computed as the ratio of the excess AUC over a random baseline to that of the oracle uncertainty score (which ranks instances perfectly by quality):

$$PRR = \frac{AUC_{unc} - AUC_{rnd}}{AUC_{oracle} - AUC_{rnd}}.$$
 (14)

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)

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.

324 325

Table 3: PRR (↑ 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 ↑ when the variant improves over its original multinomial-sampling counterpart.

32	1
32	8
32	Ç
33	(
33	1
33	2

331	
332	
333	
334	
335	
336	
337	



338



350 351 352

349

354 355 356

357

353

362 364

365

366

376

377

UQ Method	Llama	a 3.1 8B	Gemr	na 3 4B	Qwe	n 3 8B
- CQ Wellou	base	instruct	base	instruct	base	instruct
	Baseline	UQ metho	ds			
Prob	.410	.344	.471	.292	.376	.289
MTE	.422	.364	.476	.317	.407	.297
Perplexity	.452	.323	.525	.288	.372	.276
CCP	.401	.364	.492	.331	.355	.291
SAR	.352	.385	.386	.239	.363	.292
P(True)	.015	.072	.093	096	.110	114
Semantic Entropy	.414	.376	.401	.293	.319	.299
Lexical Similarity	.411	.366	.426	.247	.425	.237
EigValLaplacian	.426	.371	.437	.233	.406	.265
NumSemSets	.396	.319	.418	.238	.365	.253
Consistency-based	UQ: mult	inomial vs.	beamsed	ırch versio	ns	
Consistency-based Dissimilarity	UQ: mult	inomial vs.	beamsed	ırch versio		.327
					ns .477 .478↑	.327 . <u>355</u> ↑
Dissimilarity	.505	.379	.630	.206	.477	
Dissimilarity Dissimilarity + beamsearch	.505 . <b>543</b> ↑	.379 .417	.630 . <b>650</b> ↑	.206 .252↑	.477 .478↑	<u>.355</u> ↑
Dissimilarity Dissimilarity + beamsearch Eccentricity	.505   <b>.543</b> ↑   .453	.379 <u>.417</u> ↑ .368	.630 . <b>650</b> ↑	.206 .252↑ .231	.477 .478↑ .396	<u>.355</u> ↑ .251
Dissimilarity Dissimilarity + beamsearch Eccentricity Eccentricity + beamsearch	.505   <b>.543</b> ↑   .453   .505↑	.379 .417↑ .368 .397↑	.630 .650↑ .563 .603↑	.206 .252† .231 .285†	<u>.477</u> . <b>478</b> ↑   .396   .410↑	.355↑ .251 .345↑
Dissimilarity Dissimilarity + beamsearch Eccentricity Eccentricity + beamsearch EigVecDissimilarity	.505   <b>.543</b> ↑   .453   .505↑   .463	.379 .417↑ .368 .397↑	.630 <b>.650</b> ↑   .563 .603↑	.206 .252† .231 .285†	.477 .478↑   .396   .410↑   .425	.355↑ .251 .345↑ .256
Dissimilarity Dissimilarity + beamsearch  Eccentricity Eccentricity + beamsearch  EigVecDissimilarity EigVecDissimilarity + beamsearch	.505   <b>.543</b> ↑   .453   .505↑   .463   .510↑	.379 .417↑ .368 .397↑ .370 .414↑	.630   .650↑   .563   .603↑   .561   .598↑	.206 .252↑ .231 .285↑ .236 .301↑	.477   .478↑   .396   .410↑   .425   .450↑	.355↑ .251 .345↑ .256 .376↑
Dissimilarity Dissimilarity + beamsearch  Eccentricity Eccentricity + beamsearch  EigVecDissimilarity EigVecDissimilarity + beamsearch  CocoaMSP	.505   .543↑   .453   .505↑   .463   .510↑	.379 _417↑ .368 .397↑ .370 _414↑ .404	.630   .650↑   .563   .603↑   .561   .598↑	.206 .252↑ .231 .285↑ .236 .301↑	.477   .478↑   .396   .410↑   .425   .450↑	.355↑ .251 .345↑ .256 .376↑ .334

**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 the log-probability form  $u(\mathbf{y}_* \mid \mathbf{x}) = -\log p(\mathbf{y}_* \mid \mathbf{x})$  (CocoaMSP) and the perplexity form  $u(\mathbf{y}_* \mid \mathbf{x}) = -\frac{1}{|\mathbf{y}_*|} \log p(\mathbf{y}_* \mid \mathbf{x})$  (CocoaPPL).

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 simpliest baseline, Sequence *Probability*, calculates  $-\log p(\mathbf{y}_* \mid \mathbf{x})$ . For detailed descriptions of other methods see Appendix E.

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) PRR curves.

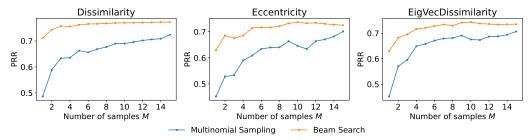


Figure 4: PRR ( $\uparrow$  is better) as a function of the number of candidates M on TriviaQA with Gemma 3 4B base. Each panel reports one estimator (Dissimilarity, Eccentricity, EigVecDissimilarity). Curves compare multinomial sampling and beam search (with probability weights from equation (4)).

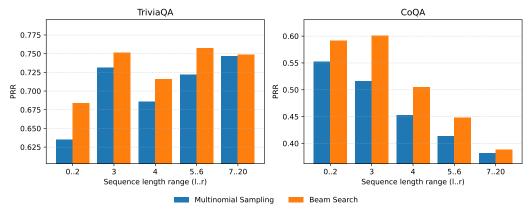


Figure 5: PRR († is better) for Dissimilarity under beam search (with probability weights) vs. multinomial sampling, for different output lengths. Each dataset (TriviaQA, CoQA) with Gemma 3 4B base is partitioned into five approximately equal-size bins token length of greedy output.

### 4.3.1 EFFECT OF SAMPLE COUNT

To assess sensitivity to the sample count, we vary  $M \in \{1, \dots, 15\}$  for Dissimilarity, Eccentricity, and EigVecDissimilarity under multinomial sampling and beam search. Results in Figure 4 show that beam search attains higher PRR at small budgets (notably M = 2–4) and saturates quickly, while multinomial sampling improves more gradually yet remains below beam across the range, the gap narrows only mildly at larger M.

#### 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 5. 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 EigVecDissimilarity on TriviaQA with Llama 3.1 8B base. Across all estimators, beam search consistently dominates multinomial sampling for nearly the entire rejection range. The improvement 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 the most critical.

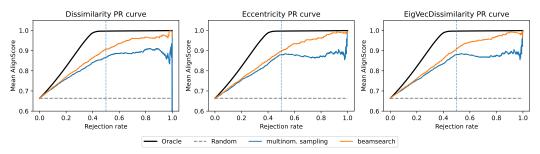


Figure 6: Prediction-Rejection curves for *Dissimilarity*, *Eccentricity*, and *EigVecDissimilarity* on TriviaQA with Llama 3.1 8B base, comparing multinomial sampling (blue) and beam search with weights (orange). Oracle (black) and random (gray dashed) baselines are shown. The vertical dashed line marks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

# 4.3.4 Additional Ablations

 Additional ablations are deferred to the appendix: 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  $\epsilon$  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 using the top-1 beam decode as the produced answer  $\mathbf{y}_*$  (instead of greedy), a natural choice when beam search is already run to obtain a higher-quality output.

# 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.

**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. Conversely, other research focused on making the decoding itself uncertainty-aware (Daheim et al., 2025; Garces Arias et al., 2024). However, these works do not explore consistency-based uncertainty quantification, and more broadly, the interplay between decoding and uncertainty is a relatively underexplored direction.

# 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.

# 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 *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1533–1544. Association for Computational Linguistics, 2013. URL https://www.aclweb.org/anthology/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. Bressem. 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*, 2025. URL https://openreview.net/forum?id=hPpyUv1XyQ.
- 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 *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 5050–5063. Association for Computational Linguistics, 2024. 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 *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 446–461. Association for Computational Linguistics, 2023. 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 *Findings of the Association for Computational Linguistics:* ACL 2024, pp. 9367–9385. Association for Computational Linguistics, 2024. 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, 2025. URL https://arxiv.org/abs/2509.16696.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. DeBERTa: Decoding-enhanced BERT with disentangled attention. In *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 *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1601–1611. Association for Computational Linguistics, 2017. 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, et al. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*, 2022. URL https://arxiv.org/abs/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*, 2023. URL https://openreview.net/forum?id=VD-AYtPOdve.
- 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 pretraining 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. Association for Computing Machinery, 2024. 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 *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. Association for Computational Linguistics, 2019. 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 *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 5433–5442. Association for Computational Linguistics, 2023. 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. Uncertainty quantification for LLMs through Minimum Bayes Risk: Bridging confidence and consistency. *arXiv preprint arXiv:2502.04964*, 2025b. URL https://arxiv.org/abs/2502.04964.
- 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), 2018. doi: 10.1609/aaai. v32i1.12340. URL http://dx.doi.org/10.1609/aaai.v32i1.12340.
- Ashwin K Vijayakumar, Michael Cogswell, Ramprasath R Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. Diverse beam search: Decoding diverse solutions from neural sequence models. arXiv preprint arXiv:1610.02424, 2016. URL https://arxiv.org/abs/1610.02424.
- Zhiqiu Xia, Jinxuan Xu, Yuqian Zhang, and Hang Liu. A survey of uncertainty estimation methods on large language models. In *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 21381–21396. Association for Computational Linguistics, 2025. 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)*, 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 *Findings of the Association for Computational Linguistics: ACL 2022*, pp. 3656–3672. Association for Computational Linguistics, 2022. 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 *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 11328–11348. Association for Computational Linguistics, 2023. 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 (Vijayakumar et al., 2016) 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

$$w_i^T = \frac{p\left(\mathbf{y}_T^{(i)} \mid \mathbf{x}\right)^{1-1/T}}{\sum_{j=1}^M p\left(\mathbf{y}_T^{(j)} \mid \mathbf{x}\right)^{1-1/T}}.$$
(15)

**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

$$w_i^H = p\left(\mathbf{b}^{(i)} \mid \mathbf{x}\right), \quad i = 1, \dots, B, \qquad w_j^H = \frac{1 - \sum_{i=1}^B p\left(\mathbf{b}^{(i)} \mid \mathbf{x}\right)}{M - B}, \quad j = B + 1, \dots, M,$$
 so that  $\sum_{i=1}^M w_i^H = 1$ . We test  $B \in \{1, \dots, 9\}$ .

Table 4: PRR († 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 <u>underlined</u>.

				Gemma 3	3 4B base					Llama 3.1	8B base		
		1	ΓriviaQ			CoQA		·	TriviaQ <i>A</i>			CoQA	
		Dissim	Ecc	EigVec Dissim	Dissim	Ecc	EigVec Dissim	Dissim	Ecc	EigVec Dissim	Dissim	Ecc	EigVec Dissim
Bean	n search	.771	.732	.751	.561	.483	.488	.623	.581	.598	.502	.438	.454
	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
Multinomial	T = 1.0	.742	.689	.715	.465	.424	.432	.599	.516	.537	.431	.416	.424
Sampling	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>
	$G=2, \lambda=0.5$	.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
Diverse	$G = 2, \lambda = 1.5$	.763	.522	.714	.537	.420	.441	.618	101	.542	.462	.245	.317
Beam	$G = 2, \lambda = 2.0$	.759	.515	.702	.547	.377	.391	.630	130	.546	.452	.215	.287
Search	$G = 5, \lambda = 0.5$	.758	.546	.736	.493	.401	.453	.591	026	.569	.395	.262	.353
Search	$G = 5, \lambda = 1.0$	.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
	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	.781	.736	.754	.503	.461	.439	.620	.553	.598	.519	.436	.424
Hybrid	B = 4	.784	.750	.769	.516	.467	.428	.622	.572	<u>.617</u>	.436	.388	.382
Multinomial-	B = 5	.757	.733	.749	.538	.500	.466	.655	.578	.609	.470	.412	.419
Beam	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>	.597	.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

 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.

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  $\epsilon$  on the per-candidate mass:

$$w_i^{\epsilon} = \frac{\max(\epsilon, p(\mathbf{b}^{(i)} \mid \mathbf{x}))}{\sum_{j=1}^{M} \max(\epsilon, p(\mathbf{b}^{(j)} \mid \mathbf{x}))}.$$
 (17)

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

We evaluate beam-guided probability-weighted methods for different  $\epsilon$  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  $\epsilon$  parameter. Determining the optimal  $\epsilon$  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  $\epsilon$  in equation (17):  $\epsilon = 0$  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 <u>underlined</u>.

	Gemma 3 4B base			Llama 3.1 8B base								
	1	FriviaQA	4		CoQA		1	TriviaQ/	A		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	.667	.612	.627	.502	.447	.446
$\epsilon = 0.05$	.764	.720	.744	.561	.490	.487	.657	.606	.626	.509	.435	.451
$\epsilon = 0.01$	.766	.718	.749	.559	.478	.489	.630	.584	.602	.496	.437	.452
$\epsilon = 0.001$	.771	.731	.751	.562	.484	.488	.624	.581	.598	.501	.438	.453
$\epsilon = 0.00001$	.771	.732	.751	.561	.483	.488	.623	.581	.598	.502	.438	.454
$\epsilon = 0$	.771	.732	.751	.561	.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:

$$\widehat{U}_{DegMat}(\mathbf{x}) = \frac{1}{M^2} \sum_{i=1}^{M} \sum_{j=1}^{M} (1 - s(\mathbf{y}^{(i)}, \mathbf{y}^{(j)})).$$
(18)

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

$$\widehat{U}_{DegMat}^{b}(\mathbf{x}) = \sum_{i=1}^{M} w_i \sum_{j=1}^{M} w_j (1 - s(\mathbf{b}^{(i)}, \mathbf{b}^{(j)})).$$
(19)

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

$$\hat{p}(c) = \frac{1}{M} \sum_{i=1}^{M} \mathbf{1} \{ \mathbf{y}^{(i)} \in c \} \quad \text{for } c \in C.$$
 (20)

Then Semantic Entropy calculates

$$\widehat{U}_{SemEnt}(\mathbf{x}) = -\frac{1}{|C|} \sum_{c \in C} \log \widehat{p}(c).$$
(21)

For beam candidates, use cluster masses aggregated by  $w_i$ :

$$\hat{p}^{b}(c) = \sum_{i=1}^{M} w_{i} \mathbf{1}\{\mathbf{b}^{(i)} \in c\}, \qquad \hat{U}_{SemEnt}^{b}(\mathbf{x}) = -\sum_{c \in C} \hat{p}^{b}(c) \log \hat{p}^{b}(c).$$
 (22)

Note that these objectives score LLM uncertainty about the *input* x as they are independent of a particular  $y_*$ .

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.

Table 6: PR-AUC († 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; † 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 <u>underlined</u>. The rightmost column reports the mean PR-AUC across datasets.

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

#### A.4 GRAPH LAPLACIAN EMBEDDING PARAMETERS

Both multinomial and beam-guided versions of Eccentricity and EigVecDissimilarity depend on the threshold parameter  $\alpha$ , 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  $\alpha = 0.9$ .

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  $\alpha$  (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.

			1 AD 1		Llama 3.1 8B base				
		Gemma 3							
	Triv	/iaQA	C	οQA	Triv	/iaQA	C	οQA	
	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	.622	.643	.438	.441	
$\alpha = 0.8$	<u>.751</u>	.757	.508	.475	.616	.630	.444	.450	
$\alpha = 0.9$	.732	.751	.483	.488	.581	.598	.438	.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	

Here we vary  $\alpha$  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  $\alpha$  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 <u>underlined</u>;  $\uparrow$  marks an improvement of a beam variant over its multinomial counterpart.

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

# B ANALYSIS AND EXAMPLES

#### B.1 PROBABILITY MASS COVERAGE

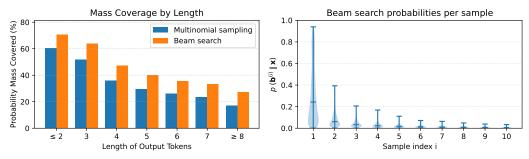


Figure 7: Left: average probability mass covered by the candidate set (M=10) across output-length bins (averaged over examples in the bin) on TriviaQA with Gemma 3 4B base. Right: for beam search, distribution of sequence probabilities  $p(\mathbf{b}^{(i)} \mid \mathbf{x})$  by beam rank i (1 = highest-probability text).

Figure 7 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  $w_i$  (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.

Kathleen Ferrier? <b>Greedy:</b> breasts of	eancer	
Multinomial samples	Beam-sea sample	
cancer	cancer	p=0.228
breast cancer	tuberculosis	p=0.154
pulmonary	breast cancer	p=0.089
breast cancer	lung cancer	p=0.041
cancer	pneumonia	p=0.039
breast cancer	leukaemia	p=0.034
myx	myel	p=0.023
cancer	leuk	p=0.011
cancer	pulmonary	p=0.01
pneumonia	lymphoma	p=0.010

Multinomial samples	Beam- sam	
six	sixth	p=0.31
seventh	6	p=0.169
sixth	6th	p=0.10
sixth	ninth	p=0.06
sixth	seventh	p=0.03
6	9	$p=0.02^{\circ}$
seventh	six	p=0.02
no	9th	p=0.02
sixteenth	no.	p=0.01
n6	7	p=0.003

Figure 8: 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

Multinomial samples	Beam-search samples					
Euro	Euro	p=0.439				
Euro	euro	p=0.201				
Euro	Cyprus pound	p=0.091				
Euro	Cypriot	p=0.072				
euro	Cyprus Pound	p=0.016				
euro	Euros	p=0.014				
euro	EURO	p=0.007				
euros	euros	p=0.007				
Euro	cyprus	p=0.007				
Euro	Cyprus	p=0.006				

Dissimilarity: 0.976

 **Dissimilarity + beamsearch:** 0.800

**Question:** who plays charlie in the santa clause movies?

Greedy: Tim Allen

Multinomial samples	Beam-search samples				
Tim Allen	Tim Allen Jeff Daniels	p=0.318 p=0.017			
Scott Calvin	Timothy Oly	p=0.017 p=0.012			
Tim Allen	Ed Asner	p=0.010			
Tim Allen Edward Arnold	Scott Calvin Edward Asner	p=0.008 p=0.008			
Tim Allen	Tony Cox	p=0.007			
Jeremy nault	Tim Allen	p=0.007			
Tim Allen Tim Allen	Tim allen Eric Lloyd	p=0.005 p=0.004			

**Dissimilarity:** 0.427

**Dissimilarity + beamsearch:** 0.313

Figure 9: 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-searc samples	
Boundaries.	Boundaries.	p=0.185
Set limits.	Limits.	p=0.079
Boundaries.	They don't	p=0.040
limits.	Rules.	p=0.032
Boundaries in their	Boundaries.	p=0.029
Charging station.	Boundaries.	p=0.028
Similar limitations.	A charging station.	p=0.016
Boundaries.	Boundaries that protect	p=0.010
Boundaries.	Limits in their own	p=0.007
Set up similar limits	Boundaries	p=0.007

Dissimilarity: 0.310

**Dissimilarity + beamsearch:** 0.190

Figure 10: 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	<pre>&lt;5 few-shot QA pairs&gt; Question: <question> Answer:</question></pre>	Answer the following question as briefly as possible. <5 few-shot QA pairs> Now answer the following question: Question: <question> Answer:</question>
Web Questions	<pre>&lt;5 few-shot QA pairs&gt; Question: <question> Answer:</question></pre>	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:</question>
CoQA	Story: <context> <all pairs="" preceding="" qa=""> Question: <question> Answer:</question></all></context>	Story: <context> <all pairs="" preceding="" qa=""> Answer the following question as briefly as possible. Question: <question> Answer:</question></all></context>
HotpotQA	Title: <title 1=""> &lt;paragraph 1&gt; Title: &lt;title 2&gt; &lt;paragraph 2&gt; Question: &lt;question&gt; Short answer:&lt;/td&gt;&lt;td&gt;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: &lt;title 1&gt; &lt;paragraph 1&gt; Title: &lt;title 2&gt; &lt;paragraph 2&gt; Question: &lt;question&gt; Short answer:&lt;/td&gt;&lt;/tr&gt;&lt;tr&gt;&lt;td&gt;Common senceQA&lt;/td&gt;&lt;td&gt;&lt;pre&gt;&lt;2 few-shot QA pairs&gt; Question: &lt;question&gt; Options: &lt;(A) - (D) options&gt; Answer:&lt;/pre&gt;&lt;/td&gt;&lt;td&gt;&lt;pre&gt;Instruction: Choose the single best answer from the   options. Answer with the option text only (not the   letter).   &lt;2 few-shot QA pairs&gt;   Now answer this.   Question: &lt;question&gt;   Options:   &lt;(A) - (D) options&gt;   Answer:&lt;/pre&gt;&lt;/td&gt;&lt;/tr&gt;&lt;tr&gt;&lt;td&gt;ARC-Challenge&lt;/td&gt;&lt;td&gt;&lt;pre&gt;&lt;2 few-shot QA pairs&gt; Question: &lt;question&gt; Options: &lt;(A) - (D) options&gt; Answer:&lt;/pre&gt;&lt;/td&gt;&lt;td&gt;&lt;pre&gt;Instruction: Choose the single best answer from the   options. Answer with the option text only (not the   letter).   &lt;2 few-shot QA pairs&gt;   Now answer this.   Question: &lt;question&gt;   Options:   &lt;(A) - (D) options&gt;   Answer:&lt;/pre&gt;&lt;/td&gt;&lt;/tr&gt;&lt;/tbody&gt;&lt;/table&gt;</title>	

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 senceQA	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 <u>underlined</u>. For beam-guided variants, we mark  $\uparrow$  when the variant improves over its original multinomial-sampling counterpart.

UQ Method		3.1 8B		na 3 4B	_	n 3 8B
	base	instruct	base	instruct	base	instruct
	Baseline	UQ method	ds			
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
Consistency-based	UQ: mult	inomial vs.	beamse	arch versio	ns	
Dissimilarity	.437	.229	.424	.333	.446	.272
Dissimilarity + beamsearch	<u>.455</u> ↑	.266↑	<u>.466</u> ↑	.390↑	.440	.346↑
Eccentricity	.405	.238	.395	.310	.375	.208
Eccentricity + beamsearch	.444↑	<u>.301</u> ↑	.450↑	.348↑	.380↑	.308↑
EigVecDissimilarity	.402	.243	.412	.316	.403	.213
EigVecDissimilarity + beamsearch	.446↑	.316↑	.457↑	.366↑	.415↑	.334↑
CocoaMSP	.447	.284	.450	.347	.454	.272
CocoaMSP + beamsearch	<b>.471</b> ↑	.290↑	<b>.478</b> ↑	<b>.407</b> ↑	<b>.</b> 459↑	<u>.345</u> ↑
CocoaPPL	.440	.251	.433	.340	.422	.261
CocoaPPL + beamsearch	.450↑	.273↑	.444↑	.395↑	.410	.318↑

# 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 <u>underlined</u>. 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			
Baseline UQ methods										
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			
Consist	tency-based (	UQ: multinom	ial vs. be	amsearch vers	sions					
Dissimilarity	.916	.836	.822	.809	.817	.818	.836			
Dissimilarity + beamsearch	.923↑	.852↑	.826↑	<u>.814</u> ↑	<u>.831</u> ↑	.841↑	.848			
Eccentricity	.897	.808	.768	.737	.809	.821	.806			
Eccentricity + beamsearch	.911↑	.816↑	.790↑	.771↑	.833↑	.859↑	.830			
EigVecDissimilarity	.902	.813	.761	.728	.798	.825	.805			
EigVecDissimilarity + beamsearch	<u>.920</u> ↑	.827↑	.787↑	.763↑	.820↑	<u>.856</u> ↑	.829			
CocoaMSP	.904	.823	.791	.726	.826	.839	.818			
CocoaMSP + beamsearch	.910↑	.836↑	.811↑	.779↑	.827↑	.847↑	.835			
CocoaPPL	.907	.832	.810	.799	.825	.837	.835			
CocoaPPL + beamsearch	.912↑	<u>.845</u> ↑	<u>.823</u> ↑	.828↑	.825↑	.844↑	.846			

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 <u>underlined</u>. 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			
Baseline UQ methods										
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			
Consist	tency-based (	UQ: multinon	nial vs. be	amsearch vers	sions					
Dissimilarity	.911	.904	.715	.838	.722	.648	.789			
Dissimilarity + beamsearch	.919↑	.915↑	.660	.822	.754↑	.693↑	<u>.794</u> ↑			
Eccentricity	.888	.887	.561	.758	.685	.625	.734			
Eccentricity + beamsearch	.906↑	.884	.576↑	.789↑	<u>.744</u> ↑	<b>.717</b> ↑	.769↑			
EigVecDissimilarity	.902	.889	.573	.766	.677	.651	.743			
EigVecDissimilarity + beamsearch	<u>.916</u> ↑	.900↑	.588↑	.784↑	.717↑	.689↑	.766↑			
CocoaMSP	.897	.894	.605	.761	.711	.680	.758			
CocoaMSP + beamsearch	.907↑	.904↑	.632↑	.801↑	.715↑	.691↑	.775↑			
CocoaPPL	.902	.902	.672	.861	.712	.686	.789			
CocoaPPL + beamsearch	.909↑	<u>.910</u> ↑	<u>.690</u> ↑	. <del>881</del> ↑	.718↑	<u>.695</u> ↑	.801↑			

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

UQ Method	TriviaQA	Web Questions	CoQA	HotpotQA	Common senceQA	ARC- Challenge			
Baseline UQ methods									
Prob	.659	.521	.312	.274	.511	.548			
MTE	.670	.583	.363	.494	.364	.381			
Perplexity	.647	.553	.369	.527	.503	.547			
CCP	.686	.569	.326	.337	.506	.527			
SAR	.656	.571	.347	.296	.183	.264			
P(True)	.272	004	.031	.075	.090	.090			
SemanticEntropy	.622	.505	.301	.140	.407	.431			
Lexical Similarity	.602	.540	.349	.286	.386	.392			
EigValLaplacian	.666	.555	.320	.246	.386	.452			
NumSemSets	.656	.538	.257	.268	.338	.454			
Consistenc	y-based UQ:	multinomial v	s. beams	earch versions					
Dissimilarity	.755	.715	.578	.626	.561	.545			
Dissimilarity + beamsearch	. <del>766</del> ↑	<b>.722</b> ↑	. <del>600</del> ↑	.611	.595↑	.604↑			
Eccentricity	.714	.653	.459	.453	.549	.549			
Eccentricity + beamsearch	.739↑	.633	.505↑	.514↑	<u>.590</u> ↑	.636↑			
EigVecDissimilarity	.738	.661	.443	.448	.512	.562			
EigVecDissimilarity + beamsearch	.753↑	.668↑	.497↑	.487↑	.562↑	<u>.621</u> ↑			
CocoaMSP	.738	.666	.509	.430	.583	.595			
CocoaMSP + beamsearch	.747↑	.679↑	.548↑	.523↑	.586↑	.606↑			
CocoaPPL	.739	.678	.548	.625	.580	.595			
CocoaPPL + beamsearch	.748↑	.694↑	.577↑	.681↑	.582↑	.610↑			

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 <u>underlined</u>. 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			
Baseline UQ methods									
Prob	.442	.425	.162	.220	.254	.252			
MTE	.534	.465	.161	.232	.253	.256			
Perplexity	.422	.419	.157	.223	.252	.256			
CCP	.533	.478	.117	.303	.264	.290			
SAR	.533	.426	.176	.214	.033	.050			
P(True)	076	155	161	090	046	047			
SemanticEntropy	.449	.415	.166	.223	.254	.252			
Lexical Similarity	.527	.427	.176	.127	.052	.172			
EigValLaplacian	.578	.472	.190	.134	.014	.010			
NumSemSets	.556	.442	.123	.106	.046	.153			
Consistency	v-based UQ:	multinomial v	s. beamse	earch versions					
Dissimilarity	.549	.415	.111	.068	.024	.070			
Dissimilarity + beamsearch	.413	.321	.204↑	.273↑	.218↑	.085↑			
Eccentricity	.540	.429	.167	.175	020	.094			
Eccentricity + beamsearch	.441	.367	.235↑	.314↑	.246↑	.108↑			
EigVecDissimilarity	.561	.437	.169	.173	017	.095			
EigVecDissimilarity + beamsearch	.478	.416	<b>.240</b> ↑	<u>.308</u> ↑	.253↑	.113↑			
CocoaMSP	.531	.456	.183	.198	.252	.266			
CocoaMSP + beamsearch	.535↑	<u>.473</u> ↑	<u>.237</u> ↑	.287↑	.282↑	.258			
CocoaPPL	.523	.454	.174	.201	.247	.271			
CocoaPPL + beamsearch	.522	.467↑	.222↑	.285↑	<u>.277</u> ↑	.264			

Table 16: PRR (↑ 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 <u>underlined</u>. Beam-guided and probability-weighted variants are marked with ↑ when they improve over their multinomial-sampling baseline.

UQ Method	TriviaQA	Web Questions	CoQA	HotpotQA	Common senceQA	ARC- Challenge			
Baseline UQ methods									
Prob	.517	.414	.310	.213	.504	.505			
MTE	.544	.420	.286	.327	.448	.511			
Perplexity	.507	.441	.316	.375	.501	.570			
CCP	.575	.420	.276	.247	.442	.446			
SAR	.548	.452	.331	.263	.189	.330			
P(True)	055	.059	020	223	.034	.292			
SemanticEntropy	.538	.409	.330	.199	.492	.514			
Lexical Similarity	.467	.396	.366	.289	.437	.511			
EigValLaplacian	.569	.418	.377	.247	.449	.499			
NumSemSets	.550	.409	.319	.241	.378	.477			
Consistency	-based UQ:	multinomial v	s. beamse	earch versions					
Dissimilarity	.576	.445	.473	.446	.449	.640			
Dissimilarity + beamsearch	.654↑	.504↑	. <del>485</del> ↑	.424	.510↑	.683↑			
Eccentricity	.555	.404	.405	.297	.464	.591			
Eccentricity + beamsearch	.613↑	.458↑	.429↑	.361↑	.512↑	.657↑			
EigVecDissimilarity	.570	.452	.409	.289	.469	.587			
EigVecDissimilarity + beamsearch	.630↑	.492↑	.427↑	.357↑	.506↑	.650↑			
CocoaMSP	.595	.458	.463	.366	.510	.641			
CocoaMSP + beamsearch	<u>.631</u> ↑	.487↑	.465↑	.372↑	.532↑	.639			
CocoaPPL	.587	.464	.464	.465	.501	.660			
CocoaPPL + beamsearch	.616↑	<u>.498</u> ↑	.459	<u>.456</u>	<u>.525</u> ↑	<u>.661</u> ↑			

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 <u>underlined</u>. 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			
Baseline UQ methods									
Prob	.524	.357	.327	.213	.283	.363			
MTE	.604	.424	.307	.253	.260	.339			
Perplexity	.498	.367	.262	.221	.255	.332			
CCP	.576	.406	.291	.265	.248	.402			
SAR	.599	.420	.338	.236	.301	.418			
P(True)	.236	.012	.018	.045	011	.135			
SemanticEntropy	.591	.381	.335	.231	.301	.418			
Lexical Similarity	.566	.395	.347	.232	.275	.380			
EigValLaplacian	.615	.389	.355	.238	.252	.377			
NumSemSets	.569	.363	.228	.180	.208	.368			
Consistency	y-based UQ:	multinomial v	s. beamse	earch versions					
Dissimilarity	.616	.382	.349	.270	.277	.378			
Dissimilarity + beamsearch	<u>.662</u> ↑	.411↑	.358↑	.349↑	.288↑	.434↑			
Eccentricity	.598	.379	.319	.248	.273	.389			
Eccentricity + beamsearch	.620↑	.396↑	.330↑	.281↑	.306↑	<u>.451</u> ↑			
EigVecDissimilarity	.611	.378	.325	.249	.264	.390			
EigVecDissimilarity + beamsearch	.640↑	.425↑	.347↑	.291↑	.318↑	<b>.461</b> ↑			
CocoaMSP	.629	.409	.366	.278	.314	.426			
CocoaMSP + beamsearch	.665↑	.428↑	<b>.</b> 378↑	<u>.344</u> ↑	.302	.439↑			
CocoaPPL	.626	.410	.354	.278	.299	.413			
CocoaPPL + beamsearch	.653↑	<u>.427</u> †	.356↑	.334↑	.285	.419↑			

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 <u>underlined</u>. 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				
Baseline UQ methods										
Prob	.617	.449	.267	.111	.337	.475				
MTE	.602	.409	.267	.279	.443	.444				
Perplexity	.597	.426	.278	.256	.294	.381				
CCP	.640	.406	.213	.153	.296	.421				
SAR	.617	.457	.323	.243	.220	.317				
P(True)	.322	.282	.005	.168	043	074				
Semantic Entropy	.549	.411	.247	.173	.230	.305				
Lexical Similarity	.595	.430	.338	.310	.367	.508				
EigValLaplacian	.602	.423	.301	.284	.349	.475				
NumSemSets	.593	.403	.268	.250	.311	.367				
Consistency-based UQ: multinomial vs. beamsearch versions										
Dissimilarity	.668	.462	.406	.531	.315	.476				
Dissimilarity + beamsearch	.680↑	.484↑	<b>.409</b> ↑	<u>.504</u>	.335↑	.457				
Eccentricity	.615	.416	.320	.319	.266	.440				
Eccentricity + beamsearch	.640↑	.437↑	.368↑	.407↑	.243	.366				
EigVecDissimilarity	.628	.454	.325	.314	.373	.456				
EigVecDissimilarity + beamsearch	.660↑	.460↑	.380↑	.394↑	.353	.453				
CocoaMSP	.667	.492	.385	.320	.378	.523				
CocoaMSP + beamsearch	<u>.678</u> ↑	<b>.498</b> ↑	.391↑	.378↑	<u>.385</u> ↑	<u>.510</u>				
CocoaPPL	.665	.478	.388	.397	.353	.484				
CocoaPPL + beamsearch	.667↑	.486↑	.387	.437↑	.339	.450				

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 <u>underlined</u>. 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
	Base	line UQ meth	nods			
Prob	.564	.353	.215	.250	.174	.181
MTE	.564	.345	.164	.251	.183	.272
Perplexity	.491	.341	.169	.250	.175	.229
CCP	.563	.383	.169	.258	.173	.202
SAR	.590	.425	.146	.159	.201	.233
P(True)	105	222	126	.018	083	164
Semantic Entropy	.597	.404	.214	.231	.174	.176
Lexical Similarity	.530	.425	.193	.101	.121	.053
EigValLaplacian	.626	.417	.196	.083	.134	.134
NumSemSets	.608	<u>.437</u>	.110	.096	.113	.154
Consistency	v-based UQ:	multinomial v	s. beamse	earch versions		
Dissimilarity	.588	.382	.165	.187	.246	.394
Dissimilarity + beamsearch	<u>.637</u> †	.386↑	.269↑	.264↑	.213	.362
Eccentricity	.565	.367	.167	.125	.150	.132
Eccentricity + beamsearch	.600↑	.392↑	<u>.288</u> ↑	<u>.291</u> ↑	.211↑	.285↑
EigVecDissimilarity	.590	.385	.169	.121	.143	.131
EigVecDissimilarity + beamsearch	.645↑	.439↑	.328↑	<b>.297</b> ↑	<u>.242</u> ↑	.306↑
CocoaMSP	.607	.394	.204	.272	.230	.298
CocoaMSP + beamsearch	.635↑	.404↑	.263↑	.282↑	.206	.290
CocoaPPL	.581	.389	.179	.272	.232	.309
CocoaPPL + beamsearch	.609↑	.395↑	.233↑	.282↑	.207	.299

# 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:

$$U_{SP}(\mathbf{y} \mid \mathbf{x}) = -\log P(\mathbf{y} \mid \mathbf{x}). \tag{23}$$

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

$$U_{\text{MTE}}(\mathbf{y} \mid \mathbf{x}) = \frac{1}{L} \sum_{l=1}^{L} \mathcal{H}(y_l \mid \mathbf{y}_{< l}, \mathbf{x}), \tag{24}$$

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:

$$U_{\text{PPL}}(\mathbf{y} \mid \mathbf{x}) = -\frac{1}{L} \log P(\mathbf{y} \mid \mathbf{x}), \tag{25}$$

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

$$U_{\text{CCP}}(C \mid \mathbf{x}) = 1 - \prod_{j \in C} \text{CCP}(y_j \mid y_{< j}, \mathbf{x}). \tag{26}$$

Where 
$$\mathrm{CCP}(y_j \mid \mathbf{y}_{< j}, \mathbf{x}) = \frac{\sum_{k: \mathrm{NLI}(y_j^k, y_j) = 'e'} P(y_j^k \mid \mathbf{y}_{< j}, \mathbf{x})}{\sum_{k: \mathrm{NLI}(y_j^k, y_j) \in \{'e', 'c'\}} P(y_j^k \mid \mathbf{y}_{< j}, \mathbf{x})}$$

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

$$U_{\text{SentSAR}}(\mathbf{x}) = -\frac{1}{M} \sum_{i=1}^{M} \log \left( p(\mathbf{y}^{(i)} \mid \mathbf{x}) + \frac{1}{t} R_{S}(\mathbf{y}^{(i)}, \mathbf{x}) \right), \tag{27}$$

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:

$$U_{\text{TokenSAR}}(\mathbf{x}) = -\sum_{l=1}^{L} \tilde{R}_{T}(y_{l}, \mathbf{y}, \mathbf{x}) \log P(y_{l} \mid \mathbf{y}_{< l}, \mathbf{x}),$$
(28)

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.

**Semantic Entropy**, introduced in Kuhn et al. (2023), clusters M sampled generations into K clusters of semantically equivalent responses. The entropy is then computed over these meaning clusters:

$$U_{\text{SE}}(\mathbf{x}) = -\sum_{k=1}^{K} \frac{|\mathcal{C}_k|}{M} \log \hat{p}_k(\mathbf{x}). \tag{29}$$

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

$$U_{LSRL}(\mathbf{x}) = 1 - \frac{2}{M(M-1)} \sum_{i < j} s(\mathbf{y}^{(i)}, \mathbf{y}^{(j)}).$$
(30)

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:

$$U_{\text{EigV}}(\mathbf{x}) = \sum_{i=1}^{M} \max(0, 1 - \lambda_i(\mathbf{x})). \tag{31}$$

# F COMPUTATIONAL BUDGET

All experiments were run on  $2\times NVIDIA$  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.