

Fine-Grained Uncertainty Decomposition in Large Language Models: A Spectral Approach

Nassim Walha^{1,2,4}, Sebastian G. Gruber⁵, Thomas Decker^{6,7,8}, Yinchong Yang⁶,
Alireza Javanmardi^{7,8}, Eyke Hüllermeier^{7,8,9}, Florian Buettner^{1,2,3,4}

¹German Cancer Research Center (DKFZ)

²German Cancer Consortium (DKTK)

³Frankfurt Cancer Institute, Germany

⁴Goethe University Frankfurt, Germany

⁵ESAT-PSI, KU Leuven, Belgium

⁶Siemens AG

⁷LMU Munich

⁸Munich Center for Machine Learning (MCML)

⁹German Center for Artificial Intelligence (DFKI), Kaiserslautern, Germany

nassim.walha@dkfz.de, sebgruber@gmail.com, eyke@lmu.de, florian.buettner@dkfz.de

Abstract

As Large Language Models (LLMs) are increasingly integrated in diverse applications, obtaining reliable measures of their predictive uncertainty has become critically important. A precise distinction between aleatoric uncertainty, arising from inherent ambiguities within input data, and epistemic uncertainty, originating exclusively from model limitations, is essential to effectively address each uncertainty source. In this paper, we introduce Spectral Uncertainty, a novel approach to quantifying and decomposing uncertainties in LLMs. Leveraging the Von Neumann entropy from quantum information theory, Spectral Uncertainty provides a rigorous theoretical foundation for separating total uncertainty into distinct aleatoric and epistemic components. Unlike existing baseline methods, our approach incorporates a fine-grained representation of semantic similarity, enabling nuanced differentiation among various semantic interpretations in model responses. Empirical evaluations demonstrate that Spectral Uncertainty outperforms state-of-the-art methods in estimating both aleatoric and total uncertainty across diverse models and benchmark datasets.

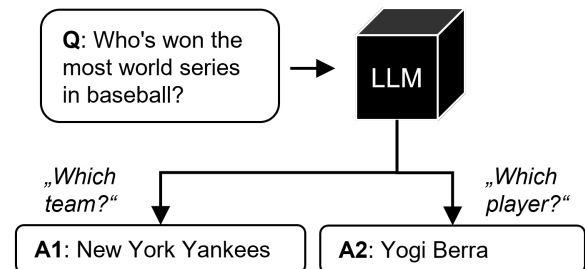
Code — <https://smplu.link/spectralUncertainty>

Extended version — <https://arxiv.org/abs/2509.22272>

1 Introduction

Since the public release of ChatGPT (Leiter et al. 2024), Large Language Models (LLMs) have exhibited exponential improvements in capabilities across numerous benchmarks and have demonstrated increased algorithmic efficiency (Ho et al. 2024). Consequently, LLMs are increasingly employed in critical domains such as scientific research (Feyzollahi and Rafizadeh 2025; Zhuang et al. 2025), politics (Aoki 2024), and medicine (Wang and Zhang 2024; Riedemann, Labonne, and Gilbert 2024). This widespread adoption has underscored the need to not only generate better predictions but also reliably quantify their uncertainty.

Copyright © 2026, Association for the Advancement of Artificial



Spectral Uncertainty Decomposition




Figure 1: Illustration of our Spectral Uncertainty Decomposition. Given an ambiguous query like “Who’s won the most World Series in baseball?”, an LLM may interpret it in multiple valid ways (e.g., by team or by player), leading to high predictive uncertainty.

While several recent approaches have been developed to quantify uncertainty in LLMs, they exhibit key limitations: many rely on token-level representations, treat semantic similarity as a binary relation, or fail to decompose predictive uncertainty into its epistemic and aleatoric components. These limitations hinder the interpretability and practical utility of uncertainty estimates in real-world applications.

To address this gap, we propose the Spectral Uncertainty framework, which introduces a fine-grained, theo-

Intelligence (www.aaai.org). All rights reserved.

retically grounded decomposition of uncertainty in LLMs. Our approach leverages Von Neumann entropy (Von Neumann 2018) and functional Bregman information (Gruber and Buettner 2023) to derive a kernel-based decomposition into aleatoric and epistemic components. A two-stage sampling and embedding process, followed by spectral analysis via a kernel function, enables practical estimation of these uncertainty measures in continuous semantic space.

The main contributions of our work are as follows:

- (a) We introduce Spectral Uncertainty, a novel uncertainty quantification framework for LLMs that enables fine-grained estimation of both aleatoric and epistemic uncertainty. We provide a rigorous theoretical derivation of this framework from a novel and general uncertainty decomposition based on functional Bregman information, applicable to any concave uncertainty measure.
- (b) We instantiate this decomposition using von Neumann entropy and propose practical, kernel-based estimators for each component: aleatoric, epistemic, and total uncertainty. These estimators allow the theoretical framework to be applied in real-world LLM scenarios using embedding-based representations of model outputs.
- (c) We demonstrate through extensive empirical evaluation that Spectral Uncertainty achieves state-of-the-art performance in both ambiguity detection and correctness prediction tasks, outperforming strong semantic and decomposition-based baselines. This highlights its potential for improving the reliability and interpretability of LLM predictions in practice.

2 Related Work

Early approaches to estimating model uncertainty (Malinin and Gales 2021; Jiang et al. 2021) relied exclusively on output token probabilities, which rendered them infeasible in black-box scenarios. Moreover, such methods primarily measured lexical and syntactic confidence, failing to account for the semantic correctness of model responses. For instance, a model generating both “France’s capital is Paris” and “Paris is France’s capital” may appear uncertain under token-level measures despite conveying the same meaning.

To overcome these limitations, semantic entropy (Kuhn, Gal, and Farquhar 2023; Farquhar et al. 2024) defines entropy in a semantic rather than token-based space. This technique samples multiple responses and groups them into clusters of semantically equivalent responses, using a Natural Language Inference (NLI) model (Bowman et al. 2015). Subsequently, entropy is computed from the resulting categorical distribution of clusters. While this method significantly improves over lexical measures, it treats semantic equivalence as binary, thus failing to capture finer semantic nuances—such as gradations between “extremely high,” “somewhat high,” and “moderate” temperatures.

Kernel Language Entropy (Nikitin et al. 2024) provides a finer semantic representation by computing discrete similarity scores between generated responses using weighted NLI predictions. Although this improves granularity, it still discretizes the semantic space and does not fully leverage continuous embeddings.

A further advancement, predictive kernel entropy (Gruber and Buettner 2024), represents model outputs using sentence embeddings and computes similarity via kernel functions in a continuous semantic space. This method currently achieves state-of-the-art performance and offers the most refined representation of uncertainty.

Despite these advancements, existing methods focus solely on predictive (or total) uncertainty and do not disentangle its underlying components. Predictive uncertainty captures the overall confidence in a model’s response but offers no insight into the source of that uncertainty. In particular, two major types of uncertainty are well-recognized (Hüllermeier and Waegeman 2021):

- *Aleatoric uncertainty* arises from inherent ambiguity or noise in the input (e.g., unclear queries or underspecified instructions) and cannot be reduced by improving the model.
- *Epistemic uncertainty*, by contrast, reflects the model’s lack of knowledge, often due to gaps in training data, and can potentially be reduced through additional learning or data collection.

Uncertainty decomposition aims to separate total predictive uncertainty into its constituent components: aleatoric and epistemic uncertainty. While this has been studied extensively in the context of classification tasks (Depeweg et al. 2018; Gruber and Buettner 2023), its application to LLMs remains relatively underexplored. A recent effort by Hou et al. (2024) extends uncertainty decomposition to LLMs using a clustering-based method akin to semantic entropy.

Conceptually, their approach draws on the standard information-theoretic decomposition of uncertainty (Depeweg et al. 2018):

$$\mathcal{H}(q(Y | X)) = \mathbb{E}_{q(\theta|\mathcal{D})} [\mathcal{H}(q(Y | X, \theta))] + \mathcal{I}(Y; \theta | X),$$

where \mathcal{H} denotes Shannon entropy (Shannon 1948) and \mathcal{I} the mutual information. Here, $q(Y | X)$ represents the model’s predictive distribution for output Y given input X , and θ is a latent variable representing different model configurations—typically instantiated via ensembling. In this decomposition, the mutual information term $\mathcal{I}(Y; \theta | X)$ captures the disagreement among ensemble members and is thus interpreted as epistemic uncertainty. The expected conditional entropy $\mathbb{E}_{q(\theta|\mathcal{D})} [\mathcal{H}(q(Y | X, \theta))]$ quantifies the remaining irreducible uncertainty, attributed to aleatoric uncertainty.

However, Hou et al. (2024) diverge from this standard decomposition in two key ways to adapt it to LLMs. First, although uncertainty in classification or regression tasks can be estimated using Bayesian Neural Networks (BNNs) (Graves 2011; Blundell et al. 2015) or Deep Ensembles (Lakshminarayanan, Pritzel, and Blundell 2017), these approaches are computationally infeasible for LLMs. Even limiting such methods to the fine-tuning stage requires white-box access to the model, which is often impractical or unavailable for proprietary LLMs.

To circumvent this limitation, Hou et al. (2024) propose substituting model variability (θ) with input context

variability. Specifically, they generate multiple clarifications C_1, C_2, \dots, C_n of the user’s input question, each representing an interpretation or reformulation of the question. The model is then conditioned on these clarifications, effectively creating an ensemble over input contexts rather than over models. This leads to a reformulated decomposition:

$$\mathcal{H}(q(Y | X)) = \mathbb{E}_{q(C|D)} [\mathcal{H}(q(Y | X, C))] + \mathcal{I}(Y; C | X),$$

where the mutual information term $\mathcal{I}(Y; C | X)$ reflects disagreement between interpretations and is therefore attributed to aleatoric uncertainty. Conversely, the expected conditional entropy $\mathbb{E}_{q(C|D)} [\mathcal{H}(q(Y | X, C))]$ is interpreted as epistemic uncertainty, as it captures residual uncertainty after conditioning on a particular interpretation.

While this method achieves state-of-the-art performance, it inherits a key limitation from earlier clustering-based techniques: the reliance on discrete clusters to compute entropy reduces semantic similarity to a binary notion. As a result, the decomposition remains coarse and unable to fully capture fine-grained distinctions in meaning, which limits the accuracy and expressiveness of its uncertainty estimates.

3 A Novel Uncertainty Decomposition

In this section, we introduce a novel and general uncertainty decomposition of aleatoric and epistemic uncertainty given any total uncertainty represented by a concave function. This is followed by a discussion of the special case of von Neumann entropy, and an introduction of estimators useable in practice. A core definition, which we require for our contribution, is Bregman Information given as follows.

Definition 3.1 ((Gruber and Buettner 2023)). For a random variable X with outcomes in an appropriate space \mathcal{X} , and a convex function $g: \mathcal{X} \rightarrow \mathbb{R}$, the (functional) **Bregman Information** of X generated by g is defined by

$$\mathbb{B}_g(X) := \mathbb{E}[g(X)] - g(\mathbb{E}[X]). \quad (1)$$

The Bregman Information is a generalisation of the variance of a random variable, i.e., if $\mathcal{X} = \mathbb{R}$ and $g_{\text{sq}}(x) = x^2$, then $\mathbb{B}_{g_{\text{sq}}}(X) = \text{Var}(X)$. For any g it holds that $\mathbb{B}_g(X) = 0$ if X has only one outcome with non-zero probability. Further, if g is differentiable, then the Bregman Information is the expected Bregman divergence between X and its expectation $\mathbb{E}[X]$ (Banerjee et al. 2005).

The Bregman Information arises naturally in our decomposition as follows.

Decompositions: General and Special Cases

Based on the above definition, we provide the following general uncertainty decomposition of a marginal distribution.

Theorem 3.2. Let \mathcal{P} be a set of probability distributions over a set \mathcal{Y} and $H: \mathcal{P} \rightarrow \mathbb{R}$ a concave function. Let Y be a random variable with outcomes in \mathcal{Y} , and marginal distribution \mathbb{P}_Y . Further, let $\mathbb{P}_{Y|W}$ be a conditional distribution of Y given another random variable W . Then,

$$H(\mathbb{P}_Y) = \mathbb{E}_W [H(\mathbb{P}_{Y|W})] + \mathbb{B}_{-H}(\mathbb{P}_{Y|W}). \quad (2)$$

The proof of this general result is remarkably short:

$$\begin{aligned} H(\mathbb{P}_Y) &= H(\mathbb{E}_W [\mathbb{P}_{Y|W}]) + \mathbb{E}_W [H(\mathbb{P}_{Y|W})] \\ &\quad - \mathbb{E}_W [H(\mathbb{P}_{Y|W})] \\ &= \mathbb{E}_W [H(\mathbb{P}_{Y|W})] + \mathbb{B}_{-H}(\mathbb{P}_{Y|W}). \end{aligned}$$

A very common example of a concave function H of distributions is the Shannon entropy (Shannon 1948), defined as $H(p) = -\sum_{i=1}^n p_i \log p_i$ for a discrete probability distribution $p = (p_1, \dots, p_n)$. Substituting H with the Shannon entropy in Theorem 3.2 recovers the classical information-theoretical decomposition of total uncertainty into aleatoric and epistemic uncertainty (Depeweg et al. 2018).

Besides the classical Shannon entropy, the kernel-based von Neumann entropy (Von Neumann 2018; Bach 2022) is another case of a concave function, which is used in recent advances for detecting hallucinations of large language models (Nikitin et al. 2024). Informally, the von Neumann entropy receives a covariance operator as its argument and is equal to the Shannon entropy of the eigenvalues of the respective covariance operator. For a rigorous definition, we require some fundamental concepts related to reproducing kernel Hilbert spaces (RKHS) (Bach 2022). Let \mathcal{X} be a compact set and $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be continuous positive semidefinite (p.s.d.) kernel function. Let \mathcal{H} be the corresponding RKHS. The kernel is normalized if $k(x, x) = 1$ for all $x \in \mathcal{X}$. Further, let $\Phi: \mathcal{X} \rightarrow \mathcal{H}$ be the corresponding feature map, with

$$k(x, y) = \langle \Phi(x), \Phi(y) \rangle_{\mathcal{H}}, \quad \text{for all } x, y \in \mathcal{X}. \quad (3)$$

The respective tensor product $\otimes_{\mathcal{H}}$ is defined for every $f, g, h \in \mathcal{H}$ by an operator $(f \otimes_{\mathcal{H}} g): \mathcal{H} \rightarrow \mathcal{H}$ via

$$(f \otimes_{\mathcal{H}} g)(h) = \langle g, h \rangle_{\mathcal{H}} f. \quad (4)$$

We can now introduce covariance operators, which act as arguments for the von Neumann entropy.

Definition 3.3 ((Bach 2022)). Given a compact set \mathcal{X} , a probability distribution \mathbb{P} over \mathcal{X} , and a continuous, p.s.d., and normalised kernel k with RKHS \mathcal{H} , and feature map $\Phi: \mathcal{X} \rightarrow \mathcal{H}$. The non-central **covariance operator** of the distribution \mathbb{P} w.r.t the kernel k is defined by:

$$\Sigma_{\mathbb{P}} := \mathbb{E}_{X \sim \mathbb{P}} [\Phi(X) \otimes_{\mathcal{H}} \Phi(X)], \quad (5)$$

for a \mathbb{P} -distributed random variable X with values in \mathcal{X} .

Note that $\Sigma_{\mathbb{P}}$ is a self-adjoint and p.s.d. operator, and has unit trace (Bach 2022).

Definition 3.4 ((Bach 2022)). For a self-adjoint, p.s.d. operator A with a unit trace on the Hilbert space \mathcal{H} , the **von Neumann entropy (VNE)** of A is defined as

$$H_{VN}(A) := -\text{Tr}[A \log A]. \quad (6)$$

It holds that H_{VN} is concave. Further, $H_{VN}(A) = -\sum_{\lambda \in \Lambda(A)} \lambda \log \lambda$, with $\Lambda(A)$ being the (possibly infinite) sequence of eigenvalues of A ¹. In that sense, the VNE of the operator A is the Shannon entropy of its eigenvalues. Based on its properties, we can use the covariance operator $\Sigma_{\mathbb{P}}$ as the argument for H_{VN} , yielding the following.

¹We use the convention $0 \log 0 = 0$

Definition 3.5 (Kernel-based von Neumann entropy (Bach 2022)). Let \mathbb{P} be a probability distribution over \mathcal{X} and $\Sigma_{\mathbb{P}}$ be the respective covariance operator. The **kernel-based VNE** of \mathbb{P} is defined as

$$H_{VN}(\mathbb{P}) := H_{VN}(\Sigma_{\mathbb{P}}) = -\text{Tr}[\Sigma_{\mathbb{P}} \log \Sigma_{\mathbb{P}}]. \quad (7)$$

Since H_{VN} is also concave with the distribution as argument, we can use it to generate a Bregman Information for the conditional distribution $\mathbb{P}_{Y|W}$. This recovers the Holevo Information (Nielsen and Chuang 2010) given by

$$\mathbb{H}(\mathbb{P}_{Y|W}) := \mathbb{B}_{-H_{VN}}(\mathbb{P}_{Y|W}). \quad (8)$$

Now, we apply Theorem 3.2 to obtain the following important special case.

Corollary 3.6 (Spectral uncertainty decomposition). *The following holds for given random variables Y and W :*

$$\underbrace{H_{VN}(\mathbb{P}_Y)}_{\text{I}} = \underbrace{\mathbb{E}_W[H_{VN}(\mathbb{P}_{Y|W})]}_{\text{II}} + \underbrace{\mathbb{H}(\mathbb{P}_{Y|W})}_{\text{III}}. \quad (9)$$

The corollary follows directly from combining Theorem 3.2 with the definition of the Holevo Information.

For Spectral Uncertainty, we follow (Hou et al. 2024)’s approach and condition on the input clarifications (interpretations), represented here by the random variable W . In this decomposition:

- Term I denotes the total predictive uncertainty, expressed as the von Neumann entropy (VNE) of the marginal predictive distribution \mathbb{P}_Y .
- Term II corresponds to the expected conditional entropy, computed by marginalizing out the clarification variable W . This captures the model’s intrinsic uncertainty after conditioning on a specific interpretation of the input. As such, we interpret this term as measuring epistemic uncertainty, which reflects the model’s limitations in knowledge or training data, independent of input ambiguity.
- Term III represents the functional Bregman information (here using VNE as the uncertainty functional), quantifying the variability in the conditional distributions $\mathbb{P}_{Y|W}$ as W varies. In our framework, following the intuition of (Hou et al. 2024), the variable W captures different plausible clarifications or interpretations of the user’s input.

When aleatoric uncertainty is low (i.e., the input is unambiguous) there is little variation in W , and the conditional distribution $\mathbb{P}_{Y|W}$ remains stable. In the limiting case where W is almost surely constant, we have $\mathbb{B}_{-H_{VN}}(\mathbb{P}_{Y|W}) = 0$, indicating the absence of aleatoric uncertainty. In contrast, high aleatoric uncertainty manifests as greater variability in how the input can be interpreted, leading to greater variation in $\mathbb{P}_{Y|W}$, and therefore, a larger Bregman information term. Based on this behavior, we attribute term III to aleatoric uncertainty.

The decomposition provides a principled way to disentangle total predictive uncertainty into its epistemic and aleatoric components. The term spectral refers to the fact that all three quantities—total, aleatoric, and epistemic uncertainty—can be computed directly from the eigenvalues of the covariance operator in a Reproducing Kernel Hilbert Space (RKHS) via spectral decomposition, as follows.

Finite-Sum Spectral Estimators

Having established a novel uncertainty decomposition with a relevant special case, we now describe how to estimate the individual terms of the spectral uncertainty decomposition in Corollary 3.6.

The estimation of $H_{VN}(\mathbb{P}_X)$ is based on (Bach 2022). Let $\mathbf{X} = X_1, \dots, X_n \sim \mathbb{P}_X$ be i.i.d. random variables with values in \mathcal{X} . We can estimate $\Sigma_{\mathbb{P}_X}$ via

$$\widehat{\Sigma}_{\mathbb{P}_X} := \frac{1}{n} \sum_{i=1}^n \Phi(X_i) \otimes \Phi(X_i). \quad (10)$$

This yields a plug-in estimator for $H_{VN}(\mathbb{P}_X)$:

$$\widehat{H}_{VN}(\mathbf{X}) := -\text{Tr}[\widehat{\Sigma}_{\mathbb{P}_X} \log \widehat{\Sigma}_{\mathbb{P}_X}]. \quad (11)$$

Since \mathcal{H} and $\widehat{\Sigma}_{\mathbb{P}_X}$ are possibly infinite-dimensional, computing $\widehat{H}_{VN}(\mathbf{X})$ based on the above formula may be practically infeasible. In consequence, we use the following property.

Proposition 3.7 ((Bach 2022)). *Let $K \in \mathbb{R}^{n \times n}$ be the empirical kernel matrix defined by $[K]_{ij} := k(X_i, X_j)$ with $i, j \in [n] := \{1, \dots, n\}$, and denote with $\hat{\lambda}_1, \dots, \hat{\lambda}_n$ the eigenvalues of $\frac{1}{n}K$. Then*

$$\text{Tr}[\widehat{\Sigma}_{\mathbb{P}_X} \log \widehat{\Sigma}_{\mathbb{P}_X}] = \sum_{i=1}^n \hat{\lambda}_i \log \hat{\lambda}_i. \quad (12)$$

We restate the proof in the Appendix. Thus, we can express $\widehat{H}_{VN}(\mathbf{X})$ as a finite sum and compute it in practice.

Building upon that, we propose an analogous plug-in estimator of the kernel-based Holevo Information $\mathbb{H}(\mathbb{P}_{Y|W})$. Similarly to Theorem 3.2, we consider Y to be a random variable with values in \mathcal{X} and $\mathbb{P}_{Y|W}$ its conditional distribution, conditioned on another random variable W . To estimate this quantity, we need a two-stage sampling procedure:

- First, an outer sample $W_1, \dots, W_n \stackrel{\text{i.i.d.}}{\sim} \mathbb{P}_W$.
- Second, for each $i \in [n]$, an inner sample $Y_{i1}, \dots, Y_{im} \stackrel{\text{i.i.d.}}{\sim} \mathbb{P}_{Y|W_i}$, yielding a sample matrix $\mathbf{Y} := (Y_{ij})_{i \in [n], j \in [m]}$.

Similar to Proposition 3.7, we require the following eigenvalues. Define the inner kernel matrices $K_i \in \mathbb{R}^{m \times m}$ with $[K_i]_{j_1 j_2} := k(Y_{ij_1}, Y_{ij_2})$ for each $i \in [n]$ and $j_1, j_2 \in [m]$. Further, define the outer kernel matrix $K^{\text{out}} \in \mathbb{R}^{nm \times nm}$ with

$$[K^{\text{out}}]_{(i_1-1)m+j_1, (i_2-1)m+j_2} := k(Y_{i_1 j_1}, Y_{i_2 j_2})$$

for $i_1, i_2 \in [n]$ and $j_1, j_2 \in [m]$. Denote with $\hat{\lambda}_{i1}, \dots, \hat{\lambda}_{im}$ the eigenvalues of $\frac{1}{m}K_i$ for every $i \in [n]$ and with $\hat{\lambda}_1^{\text{out}}, \dots, \hat{\lambda}_{nm}^{\text{out}}$ the nm eigenvalues of $\frac{1}{nm}K^{\text{out}}$.

Now, we propose for the aleatoric uncertainty $\mathbb{H}(\mathbb{P}_{Y|W})$ the novel estimator

$$\widehat{\mathbb{H}}(\mathbf{Y}) := \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^m \hat{\lambda}_{ij} \log \hat{\lambda}_{ij} - \sum_{i=1}^{nm} \hat{\lambda}_i^{\text{out}} \log \hat{\lambda}_i^{\text{out}}. \quad (13)$$

In the Appendix, we show how this is a plug-in estimator based on Proposition 3.7. Using Proposition 3.7, we also derive the estimator for the epistemic uncertainty $\mathbb{E}_W [H_{VN}(\mathbb{P}_{Y|W})]$ as

$$-\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^m \hat{\lambda}_{ij} \log \hat{\lambda}_{ij}. \quad (14)$$

Finally, adding up both estimators yields the total uncertainty estimator.

4 Methodology

We consider a scenario in which a user provides a question or an instruction (e.g., “Who has won the most World Series championships in baseball?”) to a *target LLM*. Our method computes various uncertainty measures based on the Spectral Uncertainty decomposition presented in Equation 9. An overview of the proposed method is illustrated in Figure 1.

To compute estimators corresponding to each component of the decomposition, we employ a two-stage sampling procedure for generating answers from the *target LLM*. First, following the methodology described by Hou et al. (2024), we utilize a *clarification LLM* to generate n clarifications W_1, \dots, W_n of the user’s original input. For example, these clarifications could be “Which team has won the most World Series championships in baseball?” and “Which player has won the most World Series championships in baseball?”. The *clarification LLM* may either coincide with or differ from the *target LLM*. A more capable model, such as *GPT-4o*, naturally provides superior clarifications, thereby yielding improved uncertainty estimates, albeit with higher inference costs. Nonetheless, all n clarifications can be efficiently generated within a single prompt, significantly reducing the computational cost. Moreover, Hou et al. (2024) demonstrate that supervised fine-tuning of a smaller model (e.g., *Llama-3-8B*) on clarification generation substantially improves performance, resulting in uncertainty estimation capabilities comparable to those of larger proprietary models.

Second, for each clarification W_i , we generate m answers X_{i1}, \dots, X_{im} from the *target LLM* using multinomial sampling at a temperature $t > 0$.

Next, we employ a pretrained *sentence embedding model* on all generated answers to obtain corresponding embeddings Y_{ij} for each sampled answer, indexed by $i \in [n]$ and $j \in [m]$.

Finally, we apply the estimators introduced in Section 3 to compute the respective uncertainty measures. In our example, the generated answers to the “which team” clarification could be identically “The New York Yankees”, while answers to the “which player” clarification would be identically “Yogi Barra”. Spectral Uncertainty correctly attributes zero epistemic uncertainty and high aleatoric uncertainty in this case (which is equal to total uncertainty).

The complete procedure is summarized in Algorithm 1.

5 Experiments

We validate our proposed uncertainty decomposition framework through a comprehensive experimental analysis.

Algorithm 1: Spectral Uncertainty

Input: Target LLM $\mathcal{M}_{\text{target}}$, clarification LLM $\mathcal{M}_{\text{clarification}}$, sentence embedding model f_{emb} , kernel k , user task t

Output: Total, aleatoric, and epistemic uncertainty estimates

```

1:  $W_1, \dots, W_n \leftarrow \mathcal{M}_{\text{clarification}}(t)$ 
    $\triangleright$  Generate clarifications
2: for  $i \leftarrow 1$  to  $n$  do
3:    $X_{i1}, \dots, X_{im} \leftarrow \mathcal{M}_{\text{target}}(W_i)$ 
    $\triangleright$  Sample model answers
4:   for  $j \leftarrow 1$  to  $m$  do
5:      $Y_{ij} \leftarrow f_{\text{emb}}(X_{ij})$ 
    $\triangleright$  Compute answer embeddings
6:   end for
7:    $K_i \leftarrow \text{pairwiseCompute}(k, Y_{i,1:m})$ 
    $\triangleright$  Compute pairwise kernel values
8:    $\hat{\lambda}_{i1}, \dots, \hat{\lambda}_{im} \leftarrow \text{computeEigenvalues}(\frac{1}{m} K_i)$ 
9:   end for
10:  $K^{\text{out}} \leftarrow \text{pairwiseCompute}(k, \text{flatten}(Y_{1:n,1:m}))$ 
11:  $\hat{\lambda}_1^{\text{out}}, \dots, \hat{\lambda}_{nm}^{\text{out}} \leftarrow \text{computeEigenvalues}(\frac{1}{nm} K^{\text{out}})$ 
12: aleatoric  $\leftarrow \text{computeAleatEstimator}(\hat{\lambda}_{1:n,1:m}, \hat{\lambda}_{1:nm}^{\text{out}})$ 
    $\triangleright$  Apply Eq. 13
13: epistemic  $\leftarrow \text{computeEpistEstimator}(\hat{\lambda}_{1:n,1:m})$ 
    $\triangleright$  Apply Eq. 14
14: total  $\leftarrow$  aleatoric + epistemic
15: return total, aleatoric, epistemic

```

Specifically, we assess the effectiveness of our estimates of aleatoric and total uncertainty in tasks where each type of uncertainty is relevant.

Metrics and Tasks

To evaluate aleatoric uncertainty estimates, we follow prior work (Kirchhof, Kasneci, and Oh 2023; Mucsányi, Kirchhof, and Oh 2024) and treat label disagreement as ground truth. We use datasets in which samples with high annotator disagreement are labeled as “ambiguous” and measure how well an uncertainty estimator can discriminate between ambiguous and unambiguous inputs.

To evaluate total predictive uncertainty, we adopt the correctness prediction task (Mucsányi, Kirchhof, and Oh 2024; Kuhn, Gal, and Farquhar 2023), which measures an estimator’s ability to predict whether a model’s output is correct.

In both tasks, we quantify performance using the Area Under the Receiver Operating Characteristic curve (AUROC), which reflects the quality of the uncertainty ranking. Additionally, we report the Area Under the Precision-Recall curve (AUPR) as a complementary metric.

Baselines

We focus on baselines that leverage semantic representations of model outputs, as methods based solely on token-level probabilities have shown limited performance (Kuhn, Gal, and Farquhar 2023; Farquhar et al. 2024). Our evaluation includes Semantic Entropy (Kuhn, Gal, and Farquhar

2023), Kernel Language Entropy (Nikitin et al. 2024), Predictive Kernel Entropy (Gruber and Buettner 2024), and Input Clarification Ensembling (Hou et al. 2024). The latter is the state-of-the-art decomposition method specifically targeting aleatoric uncertainty by generating multiple clarifications of the input.

Datasets and Models

For ambiguity detection, we use the *AmbigQA* dataset (Min et al. 2020), which provides ambiguity annotations for questions. We conduct our evaluation on a randomly selected subset of 200 samples. Additionally, we include the synthetic *AmbigInst* dataset (Hou et al. 2024), which focuses on instruction-based tasks rather than general knowledge questions.

For correctness prediction, we evaluate on two widely-used question answering benchmarks: *TriviaQA* (Joshi et al. 2017) and *Natural Questions* (Kwiatkowski et al. 2019), randomly sampling 300 questions from the development set of each.

To evaluate model performance across different scales, we utilize two large language models: the 109B-parameter *LLaMA 4 Maverick* and the 14B-parameter *Phi-4*. Both models serve as the *target LLM* for generating responses across all methods. For clarification-based approaches — namely Input Clarification Ensembling as well as Spectral Uncertainty — we employ *GPT-4o* as the *clarification LLM* to generate high-quality input clarifications. All prompts used for generating model responses and clarifications are detailed in the Appendix. To compute semantic similarity, we use normalized sentence embeddings from the *all-mpnet-base-v2* model.

Ambiguity Detection Task (Aleatoric Uncertainty)

Table 1 presents AUROC and AUPR scores for aleatoric uncertainty estimation on the *AmbigQA* and *AmbigInst* datasets, evaluated using both the *Phi-4 14B* and *LLaMA 4 Maverick* models.

Across both datasets and model scales, Spectral Uncertainty consistently achieves the best performance among all baselines. On *AmbigQA*, which contains real-world ambiguous questions, Spectral Uncertainty yields the highest AUROC (69.15% and 60.39%) and AUPR (67.98% and 60.48%) on *Phi-4* and *LLaMA 4*, respectively. Notably, it provides almost a 9% higher AUROC for *Phi-4*, compared to Input Clarification Ensembling, the best-performing baseline.

The performance gap becomes even more pronounced on *AmbigInst*. Here, Spectral Uncertainty reaches AUROC scores of 86.37% (*Phi-4*) and 85.95% (*LLaMA 4*), significantly outperforming all baselines. Compared to the next-best method—Predictive Kernel Entropy for *Phi-4* and Input Clarification Ensembling for *LLaMA 4*—our method improves AUROC by over 13% and 23%, respectively. Similar trends are observed in AUPR scores. These results indicate that our decomposition-based method is particularly effective at capturing aleatoric uncertainty. Further results on *Qwen-3* family models are included in the Appendix.

These findings are supported by kernel density plots (cf. Appendix) showing the probability distributions of uncertainty values for ambiguous vs. unambiguous samples across all considered methods. In particular, they visually highlight how Spectral Uncertainty provides substantially better separation of ambiguous samples from unambiguous ones compared to baselines.

Correctness Prediction Task (Total Uncertainty)

In the correctness prediction task, we aim to quantify the ability of an LLM to give a correct answer to a given question. To establish ground truth for correctness prediction, we follow the protocol of Farquhar et al. (2024): First, we sample the most likely answer from the model at temperature $t = 0.1$, treating this as the model’s best effort. Then, we prompt GPT-4.1, to compare this answer to the ground truth and determine whether it is correct (see prompt in the Appendix).

Table 2 reports results for correctness prediction, measuring the effectiveness of total uncertainty estimates across the *TriviaQA* and *Natural Questions* datasets.

On *TriviaQA*, Spectral Uncertainty achieves the best performance across both models. For *Phi-4*, it reaches 91.92% AUROC and 80.79% AUPR, outperforming Input Clarification Ensembling by 2.5 and 6.2 percentage points, respectively. For *LLaMA 4*, Spectral Uncertainty again leads with an AUROC of 84.82% and AUPR of 60.84%.

On *Natural Questions*, the differences among methods are more nuanced. While Kernel Language Entropy achieves the highest AUROC on *Phi-4* (82.77%), Spectral Uncertainty attains the highest AUPR (81.98%), indicating stronger precision-recall performance. On *LLaMA 4*, Spectral Uncertainty yields the best scores on both metrics.

We also validate our approach via kernel density plots of correct vs. incorrect predictions across different methods (see Appendix). Once again, Spectral Uncertainty provides visibly better separation of uncertainty values than the baselines.

Overall, these results demonstrate that Spectral Uncertainty outperforms state-of-the-art baselines in most scenarios, achieving robust performance regardless of model scale or dataset.

Implementation Details

To generate multiple model outputs required by our method and the selected baselines, we use multinomial sampling from the LLM with a temperature setting of $t = 0.5$, following prior work by Kuhn, Gal, and Farquhar (2023); Gruber and Buettner (2024). In line with the recommendations of Farquhar et al. (2024), we sample $m = 10$ model answers per question for the Semantic Entropy, Kernel Language Entropy, and Predictive Kernel Entropy baselines. Similarly, for clarification-based methods — including our Spectral Uncertainty approach and Input Clarification Ensembling — we sample $m = 10$ model answers for each generated clarification.

For both clarification-based approaches, the number of clarifications per input is determined dynamically by the

Uncertainty Method	Phi-4 14B		LLaMA 4 Maverick	
	AUROC (%)	AUPR (%)	AUROC (%)	AUPR (%)
AmbigQA				
Semantic Entropy	53.29	51.85	46.14	49.36
Kernel Language Entropy	49.88	48.11	45.59	48.84
Predictive Kernel Entropy	48.37	48.94	45.10	49.12
Input Clarification Ensembling (aleatoric)	63.46	62.23	59.51	60.12
Spectral Uncertainty (aleatoric)	69.15	67.98	60.39	60.48
AmbigInst				
Semantic Entropy	60.58	69.18	55.88	64.37
Kernel Language Entropy	60.60	69.35	55.80	61.83
Predictive Kernel Entropy	75.93	79.90	66.83	71.31
Input Clarification Ensembling (aleatoric)	71.70	80.62	69.66	79.04
Spectral Uncertainty (aleatoric)	86.37	90.10	85.95	89.46

Table 1: Comparison of aleatoric uncertainty estimation for different methods on the AmbigQA and AmbigInst datasets using Phi-4 14B and LLaMA 4 Maverick. Metrics are reported as percentages.

Uncertainty Method	Phi-4 14B		LLaMA 4 Maverick	
	AUROC (%)	AUPR (%)	AUROC (%)	AUPR (%)
TriviaQA				
Semantic Entropy	84.70	71.10	71.64	43.75
Kernel Language Entropy	86.20	76.64	71.95	45.72
Predictive Kernel Entropy	85.88	74.66	73.85	45.86
Input Clarification Ensembling (total)	89.45	74.54	82.76	55.95
Spectral Uncertainty (total)	91.92	80.79	84.82	60.84
Natural Questions (NQ)				
Semantic Entropy	76.24	74.36	70.24	52.35
Kernel Language Entropy	82.77	81.84	71.60	60.79
Predictive Kernel Entropy	77.67	76.57	70.56	58.73
Input Clarification Ensembling (total)	81.91	81.04	74.93	60.72
Spectral Uncertainty (total)	81.63	81.98	75.02	62.87

Table 2: Comparison of predictive uncertainty estimation for different methods on the TriviaQA and Natural Questions datasets using Phi-4 14B and LLaMA 4 Maverick. Metrics are reported as percentages.

clarification LLM. To ensure computational tractability, we impose an upper bound of 10 clarifications per input.

For kernel choice, we follow common practice (Gruber and Buettner 2024) and employ the Radial Basis Function (RBF) kernel in our experiments (cf. Appendix).

As compute infrastructure, we use NVIDIA Quadro RTX 5000 GPUs to compute sentence embeddings and run *Phi-4* experiments. *GPT* models and *LLaMA 4* are accessed via OpenAI and Groq API calls, respectively.

6 Discussion and Conclusion

We introduced Spectral Uncertainty, a novel framework for decomposing predictive uncertainty in LLMs into aleatoric and epistemic components. Our approach is theoretically grounded in a general uncertainty decomposition based on functional Bregman information, and instantiated using von Neumann entropy in a kernel-induced semantic space. This

yields fine-grained, theoretically motivated uncertainty estimates that outperform existing baselines across standard benchmarks.

While effective, the method involves computational cost due to the number of generated responses ($n \times m$), even though each clarification uses only a small number of samples. We benchmark the effect of this limitation on compute time in the Appendix. Reducing this cost via more efficient or adaptive sampling strategies is a promising direction for future work. Moreover, although the decomposition and estimators are derived from first principles, our evaluation remains empirical—consistent with the broader trend in LLM uncertainty research.

Overall, Spectral Uncertainty offers a principled and practical decomposition framework for modeling uncertainty in language models, with potential applications in safety-critical and interactive AI systems.

Acknowledgements

Co-funded by the European Union (ERC, TAIPO, 101088594 to FB). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

References

- Aoki, G. 2024. Large Language Models in Politics and Democracy: A Comprehensive Survey. *arXiv preprint arXiv:2412.04498*.
- Bach, F. 2022. Information theory with kernel methods. *IEEE Transactions on Information Theory*, 69(2): 752–775.
- Banerjee, A.; Merugu, S.; Dhillon, I. S.; and Ghosh, J. 2005. Clustering with Bregman divergences. *Journal of machine learning research*, 6(Oct): 1705–1749.
- Blundell, C.; Cornebise, J.; Kavukcuoglu, K.; and Wierstra, D. 2015. Weight uncertainty in neural network. In *International conference on machine learning*, 1613–1622. PMLR.
- Bowman, S.; Angeli, G.; Potts, C.; and Manning, C. D. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 632–642.
- Depeweg, S.; Hernandez-Lobato, J.-M.; Doshi-Velez, F.; and Udluft, S. 2018. Decomposition of uncertainty in Bayesian deep learning for efficient and risk-sensitive learning. In *International conference on machine learning*, 1184–1193. PMLR.
- Farquhar, S.; Kossen, J.; Kuhn, L.; and Gal, Y. 2024. Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017): 625–630.
- Feyzollahi, M.; and Rafizadeh, N. 2025. The adoption of Large Language Models in economics research. *Economics Letters*, 250: 112265.
- Graves, A. 2011. Practical variational inference for neural networks. *Advances in neural information processing systems*, 24.
- Gruber, S. G.; and Buettner, F. 2023. Uncertainty Estimates of Predictions via a General Bias-Variance Decomposition. In *International Conference on Artificial Intelligence and Statistics*, 11331–11354.
- Gruber, S. G.; and Buettner, F. 2024. A bias-variance-covariance decomposition of kernel scores for generative models. In *Proceedings of the 41st International Conference on Machine Learning*, 16460–16501.
- Ho, A.; Besiroglu, T.; Erdil, E.; Owen, D.; Rahman, R.; Guo, Z. C.; Atkinson, D.; Thompson, N.; and Sevilla, J. 2024. Algorithmic progress in language models. *Advances in Neural Information Processing Systems*, 37: 58245–58283.
- Hou, B.; Liu, Y.; Qian, K.; Andreas, J.; Chang, S.; and Zhang, Y. 2024. Decomposing uncertainty for large language models through input clarification ensembling. In *Proceedings of the 41st International Conference on Machine Learning*, 19023–19042.
- Hüllermeier, E.; and Waegeman, W. 2021. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine learning*, 110(3): 457–506.
- Jiang, Z.; Araki, J.; Ding, H.; and Neubig, G. 2021. How can we know when language models know? on the calibration of language models for question answering. *Transactions of the Association for Computational Linguistics*, 9: 962–977.
- Joshi, M.; Choi, E.; Weld, D. S.; and Zettlemoyer, L. 2017. TriviaQA: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1601–1611.
- Kirchhof, M.; Kasneci, E.; and Oh, S. J. 2023. Probabilistic contrastive learning recovers the correct aleatoric uncertainty of ambiguous inputs. In *International Conference on Machine Learning*, 17085–17104. PMLR.
- Kuhn, L.; Gal, Y.; and Farquhar, S. 2023. Semantic Uncertainty: Linguistic Invariances for Uncertainty Estimation in Natural Language Generation. In *The Eleventh International Conference on Learning Representations*.
- Kwiatkowski, T.; Palomaki, J.; Redfield, O.; Collins, M.; Parikh, A.; Alberti, C.; Epstein, D.; Polosukhin, I.; Devlin, J.; Lee, K.; et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7: 453–466.
- Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30.
- Leiter, C.; Zhang, R.; Chen, Y.; Belouadi, J.; Larionov, D.; Fresen, V.; and Eger, S. 2024. Chatgpt: A meta-analysis after 2.5 months. *Machine Learning with Applications*, 16: 100541.
- Malinin, A.; and Gales, M. 2021. Uncertainty Estimation in Autoregressive Structured Prediction. In *International Conference on Learning Representations*.
- Min, S.; Michael, J.; Hajishirzi, H.; and Zettlemoyer, L. 2020. AmbigQA: Answering Ambiguous Open-domain Questions. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Mucsányi, B.; Kirchhof, M.; and Oh, S. J. 2024. Benchmarking uncertainty disentanglement: Specialized uncertainties for specialized tasks. *Advances in neural information processing systems*, 37: 50972–51038.
- Nielsen, M. A.; and Chuang, I. L. 2010. *Quantum computation and quantum information*. Cambridge university press.
- Nikitin, A.; Kossen, J.; Gal, Y.; and Marttinen, P. 2024. Kernel language entropy: Fine-grained uncertainty quantification for llms from semantic similarities. *Advances in Neural Information Processing Systems*, 37: 8901–8929.
- Riedemann, L.; Labonne, M.; and Gilbert, S. 2024. The path forward for large language models in medicine is open. *npj Digital Medicine*, 7(1): 339.

- Shannon, C. E. 1948. A mathematical theory of communication. *The Bell system technical journal*, 27(3): 379–423.
- Von Neumann, J. 2018. *Mathematical foundations of quantum mechanics: New edition*. Princeton university press.
- Wang, D.; and Zhang, S. 2024. Large language models in medical and healthcare fields: applications, advances, and challenges. *Artificial intelligence review*, 57(11): 299.
- Zhuang, Z.; Chen, J.; Xu, H.; Jiang, Y.; and Lin, J. 2025. Large language models for automated scholarly paper review: A survey. *Information Fusion*, 103332.