EMBEDDING TRUST: SEMANTIC ISOTROPY PREDICTS NONFACTUALITY IN LONG-FORM TEXT GENERATION

Anonymous authors

000

001

002003004

010 011

012

013

014

015

016

017

018

019

021

025

026

027 028 029

031

033

034

037

038

040

041

042 043

044

046

047

048

049

050 051

052

Paper under double-blind review

ABSTRACT

To deploy large language models (LLMs) in high-stakes application domains that require substantively accurate responses to open-ended prompts, we need reliable, computationally inexpensive methods that assess the trustworthiness of long-form responses generated by LLMs. However, existing approaches often rely on claim-by-claim fact-checking, which is computationally expensive and brittle in long-form responses to open-ended prompts. In this work, we introduce semantic isotropy—the degree of uniformity across normalized text embeddings on the unit sphere—and use it to assess the trustworthiness of long-form responses generated by LLMs. To do so, we generate several long-form responses, embed them, and estimate the level of semantic isotropy of these responses as the angular dispersion of the embeddings on the unit sphere. We find that higher semantic isotropy—that is, greater embedding dispersion—reliably signals lower factual consistency across samples. Our approach requires no labeled data, no fine-tuning, and no hyperparameter selection, and can be used with open- or closed-weight embedding models. Across multiple domains, our method consistently outperforms existing approaches in predicting nonfactuality in long-form responses using only a handful of samples—offering a practical, low-cost approach for integrating trust assessment into real-world LLM workflows.

1 Introduction

Large language models (LLMs) increasingly serve as front-line knowledge workers (Mayer et al., 2025). Yet the application of LLMs to high-stakes settings that require long-form responses to openended prompts is hamstrung by the fact that reliably ascertaining the trustworthiness of long-form text generated by LLMs remains challenging.

Without scalable, computationally-inexpensive methods that are able to reliably indicate the level of trustworthiness of long-form responses generated by a language model, deployment to high-stakes application domains that require substantively accurate responses to open-ended prompts will remain risky. Manual approaches are often infeasible at scale, and the standard stop-gaps—prompt engineering, system messages, and ensemble voting—remain brittle and expensive (Shorinwa et al., 2024). What is needed is a lightweight, model-agnostic, and data-agnostic method that is able to flag potentially untrustworthy text generations without needing to resort to slow, claim-by-claim fact-checking.

Prior work has attempted to gauge factuality in long-form natural language generation (NLG) by aligning atomic claims with a knowledge graph or by training auxiliary classifiers on annotated corpora (Wang et al., 2023). However, these approaches fall short in three important ways. First, they require structured references or costly ground-truth labels that may not exist in specialized domains. Second, they struggle with open-ended, multi-sentence answers where relevant facts are implicit rather than explicit (Nikitin et al., 2024). Third, their computational footprint grows rapidly with the number of sentences or claims, making them ill-suited to real-time applications (Jiang et al., 2024; Liu et al., 2024; Farquhar et al., 2024; Zhang et al., 2024; Manakul et al., 2023).

We address these limitations by introducing *semantic isotropy*—the degree of uniformity across normalized text embeddings on the unit sphere—and use it as a proxy for nonfactuality in long-form text generation. Intuitively, if a prompt admits a single, factually grounded explanation, independently sampled responses should cluster tightly in embedding space (Qiu and Miikkulainen, 2024; Wang

Figure 1: Overview of the Semantic Isotropy Scoring Pipeline.

and Holmes, 2024). Conversely, when an LLM hallucinates, subtle changes in semantics and content pull the embeddings apart, inflating angular dispersion. To measure the semantic isotropy of a set of embeddings—and gauge the trustworthiness of the corresponding long-form text generations—we compute a *semantic isotropy score* according to the following simple and computationally inexpensive steps: We first generate a handful of responses, we then embed them with any off-the-shelf text encoder, and finally, we compute a semantic isotropy score by estimating the von Neumann entropy of the cosine kernel under the set of embeddings. No labels, fine-tuning, or hyperparameter search are required.

To enable broad evaluation of semantic isotropy as a proxy for nonfactuality—particularly across varying text lengths and at scale—we develop *Segment-Score*, a new factuality scoring method. *Segment-Score* is designed to address key limitations of existing approaches (Min et al., 2023): It is more efficient in terms of token usage, scales effectively to longer responses, and offers clearer, more consistent criteria for labeling statements as true or false. Using *Segment-Score*, we assess the reliability of *semantic isotropy scoring* in predicting nonfactuality long-form text generation. Computing semantic isotropy scores is model-agnostic and can be done with off-the-shelf embedding models, making this approach inexpensive and easy to use in practice.

In our empirical evaluation, we find that *semantic isotropy scoring* achieves state-of-the-art performance across relevant benchmarks. It does so robustly across a wide range of models, response lengths, and evaluation settings, demonstrating its effectiveness and generalizability as a lightweight and scalable proxy for trustworthiness.

To summarize, our key contributions are:

- We introduce semantic isotropy and describe a simple and computationally inexpensive method for semantic isotropy scoring as a means to assess nonfactuality in long-form natural language generation.
- 2. We develop the *Segment-Score* protocol to generate and score datasets of long-form LLM generations for open-ended prompts and create a dataset of 1,182 unique entities along with $\approx \! \! 65,450$ scored responses by three different models, each containing roughly between 25 and 60 distinct claims.
- 3. We demonstrate the efficacy of using *semantic isotropy scores* as a proxy for factual inconsistency in long-form natural language generation across a range of generative models, embeddings, and relevant experimental settings.

2 RELATED WORK

2.1 Uncertainty Quantification in Large Language Models

Early work on uncertainty quantification in LLMs largely involved analyzing token-level logits and derived distributions. OpenAI (2023) demonstrated the differences in token-level calibration dynamics of pre-trained and RLHF post-trained LLMs with Kirk et al. (2024) inferring that the latter process led to model miscalibration.

On the other hand, Monte Carlo sampling-based approaches (Zhang, 2020; Malinin and Gales, 2021; Lin et al., 2024; Lakshminarayanan et al., 2017) have shown promise in estimating uncertainty in LLMs. However, a key challenge highlighted by these approaches is the need to standardize LLM responses to effectively compare them. In multiple-choice style discrete finite class output settings,

studies have used a combination of Natural Language Inference (NLI) (Williams et al., 2018), Oracle LLMs (Band et al., 2024), and Embedding Models (Qiu and Miikkulainen, 2024; Grewal et al., 2025) to standardize model responses for precise class attribution. However, in NLI methods specifically, models have historically had limited context windows and several of the comparison algorithms grow quadratically in the number of inputs (Farquhar et al., 2024; Zhang et al., 2024; Manakul et al., 2023; Wang et al., 2024c; Chen et al., 2025). As a result, while the techniques used are in principle capable of addressing the semantic variations introduced by short-form responses, several limitations remain that hinder widespread adoption in practical settings.

A direction that addresses some of the pain points of NLI methods is the introspection of the model's internal state to estimate uncertainty. Kadavath et al. (2022) use the activations in the model's last hidden layer to train a linear classifier probe that indicates whether the model knows the answer to the prompt. Kossen et al. (2024) adopt a similar approach by developing a cheap predictive model on top of the LLM's internal state that is capable of estimating semantic entropy (Farquhar et al., 2024) without repeated sampling; Nikitin et al. (2024) extend this approach to analyze a graph representation and its Laplacian. Ji et al. (2024) use the state information to specifically predict the likelihood of hallucination while Ao et al. (2024) use latent representations in a multi-modal context to develop an uncertainty quantification method for both text and images simultaneously. While the success of these approaches justify the value of leveraging latent representations for uncertainty quantification, these existing methods generally do not scale well with increases in response size and model complexity.

Recently, several studies have made meaningful advances in addressing uncertainty quantification, specifically for longer-format settings. Long Uncertainty Quantification (LUQ) (Zhang et al., 2024) and INSIDE (EigenScore) (Chen et al., 2024) are most relevant to our method and serve as key benchmarks in our analysis. While LUQ uses NLI methods to calculate a similarity matrix between samples, the LUQ-Atomic variant that we benchmark against averages normalized entailment at a sentence / statement level, providing a much more fine-grained and robust measure of similarity. EigenScore also samples a batch of responses and computes the Eigen-distribtion over the batch of latent representations in the generating model.

2.2 FACTUALITY AND LONG-FORM NLG

While several NLP datasets are designed for multiple choice-style tasks (Joshi et al., 2017; Rajpurkar et al., 2016; Lin et al., 2022; Reddy et al., 2019), the options for open-ended long-form responses are much more limited. A key challenge is defining a suitable metric that is a meaningful measure of quality but also tractable to compute. Factuality, or the degree to which a response is supported by a ground truth, has been a go-to metric for several studies (Guo et al., 2022). Min et al. (2023) developed FactScore, a popular scoring algorithm for generating factuality scores using a ground truth text. FactScore operates by decomposing a text into a set of atomic facts by using an oracle model (LLM) and scoring each of these facts with the help of the ground truth reference document. This method has been applied to create datasets such as FactScore-Bio (Zhang et al., 2024), that operates on biographies taken from Wikipedia.

FactScore has given rise to several derivative methods such as LongFact & SAFE (Wei et al., 2024) (which uses dynamic web searches to find supporting references), FELM (Chen et al., 2023) (applied to Math and Reasoning problems), Multi-FAct (Shafayat et al., 2024) (multi-lingual setting), FactAlign (Huang and Chen, 2024) and LoGU (Yang et al., 2024b) (post-training / RLHF to incorporate uncertainty phrasing into model generations), and VeriScore (Song et al., 2024) (better precision by scoring verifiable facts only) to name a few. But most of these methods share key limitations of the FactScore framework, scaling poorly with response length and requiring several LLM calls per sample. We address some of these drawbacks through our method, Segment-Score (see Algorithm 1), that utilizes capabilities of more modern LLMs such as longer context windows and the ability to respond with structured outputs.

3 Semantic Isotropy in Long-Form Language Generation

In this section, we introduce *semantic isotropy*, a measure of semantic variation in long-form text generations, and use it to motivate *semantic isotropy scoring*, a method that expresses the level of semantic isotropy in long-form responses and can serve as a proxy for nonfactuality.

Algorithm 1 SEGMENT-SCORE(t)

```
 \begin{array}{lll} \textbf{Require:} & \textbf{Topic } t, \textbf{LLM response } Y(t), \textbf{ reference document } \mathcal{D}(t), \textbf{ Oracle LLM } \mathcal{O} \\ \textbf{1:} & S(t) \leftarrow \textbf{Segmenter}_{\mathcal{O}}\big(Y(t)\big) = \langle s_1, s_2, \ldots, s_m \rangle, & \textbf{where } m \leftarrow |S(t)| \triangleright \textbf{ Partition } Y(t) \textbf{ into } m \\ \textbf{atomic segments} \\ \textbf{2:} & \textbf{ Initialize } \mathbf{v} \leftarrow [] \\ \textbf{3:} & \textbf{ for } i = 1 \textbf{ to } m \textbf{ do} \\ \textbf{4:} & C_i \leftarrow \langle s_1, s_2, \ldots, s_{i-1} \rangle & \triangleright \textbf{ Context before segment } s_i \\ \textbf{5:} & v_i \leftarrow \textbf{Score}_{\mathcal{O}}\big(C_i, s_i, \mathcal{D}(t)\big) & \triangleright \textbf{ Verify } s_i \textbf{ against } \mathcal{D}(t) \textbf{ in context} \\ \textbf{6:} & \textbf{ Append } v_i \textbf{ to } \mathbf{v} \\ \textbf{7:} & \textbf{ end for} \\ \textbf{8:} & \phi \leftarrow \frac{1}{m} \sum_{i=1}^m \mathbf{1}(v_i = \textbf{True}) & \triangleright \textbf{ Fraction of segments verified as True} \\ \end{array}
```

3.1 SEMANTIC ISOTROPY

We start with the conventional definition of isotropy. Intuitively, isotropy describes a state where a set of vectors or a distribution has no preferred direction—its properties are the same in all orientations. More formally, for a set of unit vectors, we can express this property as follows:

Definition (Isotropy). The column vectors
$$x_1,...,x_N \in \mathbb{R}^D$$
 with $||x_i||_2 = 1$ for all $i \in \{1,...,N\}$ are said to be isotropic if for the matrix $X = [x_1,...,x_N]$: $X^\top X = I_N$.

This condition means that, on average, the transformation represented by X preserves directions and spreads uniformly across all directions in \mathbb{R}^D . Geometrically, if the unit vectors $x_1, ..., x_N$ are isotropic, this implies that the unit vectors they are uniformly dispersed across the unit sphere, giving us an intuitively meaningful measure of variation within $x_1, ..., x_N$.

We will now build on this intuition to define *semantic isotropy*. Let $\mathcal{R} = \{r_1,...,r_N\} \in \mathcal{T}^N$ be a set of N long-form NLG responses given a prompt p and let $\Gamma_{\theta}(\cdot): \mathcal{T} \to \mathbb{R}^D$ be an embedding model parameterized by neural network parameters θ . We define $E = \{e_1,...,e_N\}$ to be the collection of embedding vectors, where $e_i = \Gamma_{\theta}(r_i) \in \mathbb{R}^D$ is the embedding vector of response $r_i, \bar{e}_i \in \mathbb{R}^D$ is the corresponding normalized embedding vector, $\bar{e}_i = e_i/||e_i||$, and \bar{E} is the matrix of normalized embedding vectors, $\bar{E} = [\bar{e}_1,...,\bar{e}_N]^{\top} \in \mathbb{R}^{N \times D}$. The cosine kernel function for two embedding vectors, $e_i, e_j \in \mathbb{R}^D$, is then defined as

$$k^{\cos}(e_i, e_j) = [\bar{E}\bar{E}^\top]_{i,j} = \bar{e}_i^\top \bar{e}_j. \tag{1}$$

We denote the cosine kernel matrix for E by $K_E^{\cos} \in \mathbb{R}^{N \times N}$.

With these preliminaries in place, we are now ready to formally define *semantic isotropy* to represent uniform dispersion of the normalized embedding vectors:

Definition (Semantic Isotropy). A set $\mathcal{R} = \{r_1, ..., r_N\}$ of long-form responses with corresponding embeddings $E = \{e_1, ..., e_N\}$ is semantically isotropic if $K_E^{\cos} = I_N$.

3.2 Measuring Semantic Isotropy in Long-Form Language Generation

As defined, isotropy represents a strict condition that is either met or not met. Therefore, to use semantic isotropy in practice as a means of ascertaining nonfactuality in long-form natural language generation, we need to be able to gauge the *level* of isotropy of a given cosine kernel. We want to formalize the intuition that factuality corresponds to a high level of semantic alignment within a set of long-form generations (highly anisotropic), whereas uncertainty and non-facuality correspond to a high level of dispersion (very isotropic).

To estimate the level of isotropy of a set of long-form text generations, we use the von Neumann entropy (vNE) (von Neumann and Beyer, 2018) of the cosine kernel K_E^{\cos} , normalized to have

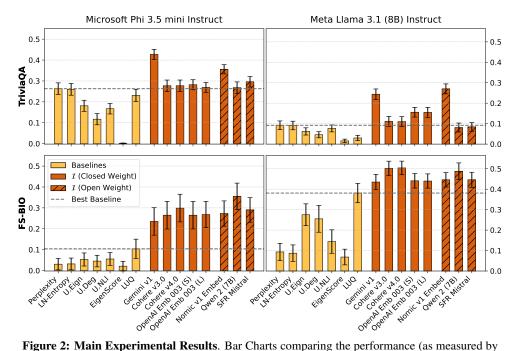


Figure 2: Main Experimental Results. Bar Charts comparing the performance (as measured by the R^2 of a linear model of Factuality \sim Semantic Isotropy); implemented using various embedding models and benchmark uncertainty metrics on two datasets: TriviaQA Entities [TriviaQA] (Top Row) and FactScore-Biographies [FS-BIO] (Bottom Row); scored using the Segment-Score (SS) algorithm. Responses are \approx 500 words. 1500 Bootstrapped samples are used to generate 1-SD error bars. Semantic isotropy \mathcal{I} (ours) outperforms all other baselines for all embedding models.

trace = 1, which allows us to view the eigenvalues as defining a discrete probability distribution. Let $\bar{K}_E^{\cos} = K_E^{\cos}/\mathrm{trace}(K_E^{\cos})$ and define

$$vNE(K_E^{\cos}) = -trace(\bar{K}_E^{\cos}\log\bar{K}_E^{\cos}). \tag{2}$$

To be precise, we view the above definition in the eigenspace of the positive semidefinite matrix K_E^{\cos} , that is, $\text{vNE}(K_E^{\cos}) = -\sum_i^N \lambda_i \log \lambda_i$, where $\lambda_1,...,\lambda_N$ are the eigenvalues of \bar{K}_E^{\cos} . Note that $\text{vNE}(K_E^{\cos})$ attains its maximum value of $\log N$ for semantically (perfectly) isotropic responses (i.e., isotropic normalized embedding vectors) and is minimized at 0 when all embedding vectors are aligned (parallel). It thus measures the degree of dispersion of $\{\bar{e}_i\}_{i=1}^N$ on the unit sphere and serves as a measure of the level of isotropy within the set of text embeddings.

We are now ready to define the *semantic isotropy score* $\mathcal{I}(\cdot)$ that will serve as our proxy for factuality. Consider N sampled long-form text generations in response to a prompt and define

$$\mathcal{I}(K_E^{\cos}) = \text{vNE}(K_E^{\cos})/\log N, \tag{3}$$

with $0 \le \mathcal{I}(K_E^{\cos}) \le 1$. Under this definition, (i) a lower degree of semantic isotropy (a semantic isotropy score closer to 0) is associated with stronger alignment/lower dispersion and consequently higher certainty and factuality, whereas (ii) a higher degree of semantic isotropy (a semantic isotropy score closer to 1) is associated with weaker alignment/higher dispersion and consequently lower certainty and factuality. In a nutshell: *The lower* the semantic isotropy score, *the more trustworthy* the generating model's long-form responses to the prompt are.

We also considered alternative matrix measures—like the Frobenius norm, $\|K_E^{\cos}\|_F$, and the log determinant, $\log(\det(K_E^{\cos}))$ —for estimating the level of embedding isotropy and found that using the von Neumann entropy works best (see Table 1). This is consistent with the findings of Nikitin et al. (2024). When assessing the efficacy of other measures of isotropy, we found that several other measures performed well, suggesting that even under different measurement approaches, isotropy is a robust proxy for nonfactuality.

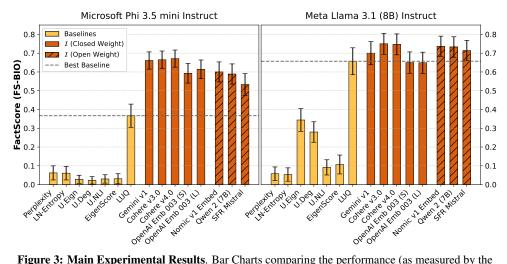


Figure 3: Main Experimental Results. Bar Charts comparing the performance (as measured by the \mathbb{R}^2 of a linear model of factuality with semantic isotropy as the explanatory variable; implemented using various embedding models and benchmark uncertainty metrics on the FS-BIO dataset using the FactScore (FS) scoring algorithm. Left hand side: Benchmark UQ metrics. Right Hand Side: semantic isotropy (\mathcal{I}) implemented using various embedding models. Same experimental setting as in Figure 2. We observe that our method's performance is robust to the scoring scheme.

4 EMPIRICAL EVALUATION

In this section, we conduct an empirical evaluation to assess whether semantic isotropy scores can serve as effective a proxy for nonfactuality in long-form natural language generation with LLMs. First, we will describe the creation of a ground truth evaluation dataset using *Segment-Score* and present our evaluation metrics and benchmarks. We will then present our main findings about the efficacy of semantic isotropy scoring as a means to predict nonfactuality in long-form text generation.

4.1 CREATION OF A DATASET OF LONG-FORM RESPONSES TO OPEN-ENDED PROMPTS

To construct a dataset of long-form response to open-ended prompts, we must select a corpus of entities such that each is associated with an underlying ground truth reference document. These entities are used to create a set of open-ended prompts that can be fed into an LLM to generate long-form responses. We use two sources—FactScore-Bio [FS-BIO] (182 unique entities) (Min et al., 2023; Zhang et al., 2024) and TriviaQA (Joshi et al., 2017) (1,000 unique entities). Both datasets use Wikipedia as their ground truth to source reference documents. For TriviaQA, we exclude all entities that correspond to days, dates or numerical values and only select those that match the title of the underlying Wikipedia page. Of the 5,245 qualifying TriviaQA entities across the training and validation set, we randomly select a subset of 1,000 entities to create open-ended prompts for long-form text generation.

To generate long-form responses, we use Meta Llama 3.1 8B Instruct (Touvron et al., 2023), Microsoft Phi 3.5 Mini Instruct (Abdin et al., 2024), and OpenAI GPT 4.1 Mini (OpenAI, 2023) to generate responses for **TriviaQA** and **FS-BIO**. For **TriviaQA** we sample up to k=20 responses while for **FS-BIO** we sample k=10 responses, targeting approximately 500 words for each response. For **TriviaQA** using Llama 3.1 8B Instruct, we also generate a longer dataset of approximately 1,000 words and derived datasets of intermediate word counts of approximately 125, 250, 375, 500, and 750 words by truncating the 1,000 word variant to the nearest sentence. Model inference is run with temperature $\tau=0.7$ and FP16 quantization using the vLLM framework (Kwon et al., 2023). Finally, we compute factuality scores for the datasets using the *Segment-Score* (**SS**) procedure (Algorithm 1). Implementation details and prompts are presented in Appendix B.3. The dataset and Segment-Score algorithm are released to allow future benchmarking by the research community.

We also generate *FactScore* (**FS**; Min et al., 2023) scores for the **FS-BIO** datasets from each model. This allow us to benchmark our scoring algorithm (**SS**) to the existing state-of-the-art method, and we can show that our results hold robustly. See Appendix C.2 for a comparison of the scoring methods. We use GPT 4.1 Mini (OpenAI, 2023) as the oracle LLM in both algorithms.

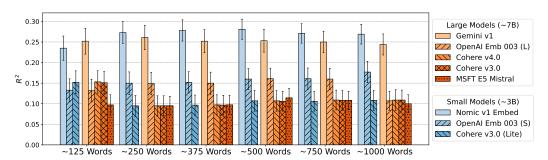


Figure 4: Performance by Response Length: Comparing *Semantic Isotropy* in terms of explained variance (R^2) of Factuality (SS) across various response lengths, contrasted by model size: Small (up to 3B parameters) and Large ($\approx 7B$ or more parameters). 1500 Bootstrapped samples are used to generate 1-SD error bars.

4.2 EVALUATION METRICS

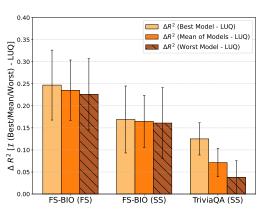
To assess the performance of our methods and compare them to appropriate baselines from the literature, we measure the degree of the relationship between each individual measure and the average factuality of a given topic, aggregated across all responses. First, we selected a set of embedding models that rank highly on the MTEB Benchmark (Muennighoff et al., 2023) to use in computing our isotropy scores. We measure predictiveness of these scores using R^2 , the explained variance of a linear model with factuality scores as the dependent variable and isotropy scores as the independent variable. To estimate error bounds, we perform bootstrap sampling and generate 1- σ error bars for each measure. We compare Semantic Isotropy to the following baseline measures: **Perplexity** -Exponential of the average negative log-likelihood of the generation (Ren et al., 2023); LN-Entropy - Length Normalized Entropy (Malinin and Gales, 2021); **U.Eign** - Sum of Eigenvalues of the Graph Laplacian, U.Deg - Average pairwise distance measured using an NLI model (Lin et al., 2024); U.NLI - NLI uncertainty derived using SelfCheckNLI (Manakul et al., 2023); Semantic Entropy -Entropy of the probability distribution of semantic clusters over the responses (Farquhar et al., 2024); LUQ-Atomic - Long Uncertainty Quantification using atomic facts via LLM based fact extraction (Zhang et al., 2024); **EigenScore** - the log-determinant of the centered embedding matrix of the generating model's internal state activations over a set of sampled responses Chen et al. (2024).

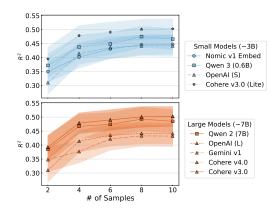
While relevant to our study, we do not include comparisons with **Kernel Language Entropy (KLE)** (Nikitin et al., 2024) and **Graph Longform Uncertainty** (Jiang et al., 2024). We found that the KLE computation was highly numerically unstable given the degree of semantic entailment overlap among the sampled long-form responses and did not produce meaningful results. For Graph Uncertainty, the distinct claim union algorithm proved computationally difficult to implement at the response lengths considered in our study. For context, a typical set of responses had over 600 distinct claims, which implied 6,000 LLM calls to measure claim to response entailment for one topic, assuming N=10.

Sampling the appropriate embedding. Unlike closed-weight models that return just an embedding vector, for open-weight models we select the activations corresponding to the last token in the final hidden state, except for Nomic v1 (Nussbaum et al., 2025) where we take the average of the activations over the token dimension instead. For further details, see Appendix B.2.

4.3 Main Result: Factuality in Long-Form Responses to Open-Ended Prompts

We present our main empirical results in Figure 2 (see also Table 2 in Appendix A). We find that semantic isotropy scores are superior to existing methods in predicting nonfactuality as per the (SS) algorithm on the FS-BIO and TriviaQA datasets. Using \mathbb{R}^2 (i.e., the OLS-explained variance) as our performance measure, \mathcal{I} outperforms all existing metrics, in most cases by a wide margin. We see some variance in the performance under different embedding models, and discuss this finding in more detail in Section 5. While certain embedding models are better suited to specific tasks, semantic isotropy scoring outperforms existing state-of-the-art methods across the board. Notably, the Nomic V1 embedding (Nussbaum et al., 2025) demonstrates exceptional performance on all datasets, occasionally surpassing even larger and closed-source models such as Gemini (Team, 2024) and OpenAI Embeddings (OpenAI, 2023).





- (a) Model Agnostic Performance Comparison: Performance difference between current state-of-the-art LUQ-Atomic (Zhang et al., 2024) and semantic isotropy based on different embedding models across each dataset using Phi 3.5 Mini Instruct as the generator model. In each case, all versions of semantic isotropy are more predictive than LUQ. 1500 bootstrapped samples are used to generate 1-SD error bars.
- (b) Performance by # of Samples: Line Charts comparing the performance of semantic isotropy by number of sampled responses on FS-BIO (SS) dataset of ≈500 words, contrasted by model size: Small (up to 3B parameters) and Large (≈7B or more parameters). 1,500 Bootstrapped samples are used to generate 1-SD error regions.

Figure 5: Ablation Studies: Comparing the performance of semantic isotropy when varying (a) Embedding Model Used [Left] and (b) Number of samples used to measure Isotropy [Right].

4.4 FACTSCORE EVALUATION

To ensure that our results are not merely an artifact of the *Segment-Score* method, we also evaluate semantic isotropy scoring under the *FactScore* (**FS**) algorithm applied to the **FS-BIO** dataset (see Figure 3). We find that the results mirror those in Figure 2 even improving the predictive performance of semantic isotropy scoring. This consistency across scoring schemes highlights the robustness and generalizability of semantic isotropy as a measure of nonfactuality. A more detailed comparison of the two scoring methods can be found in Appendix C.2.

5 SENSITIVITY STUDIES

Effect of Response Length. One important study we undertake is to understand the relationship between measure performance and response length. While our main study focuses on ≈ 500 word responses, we explore the relationship between semantic isotropy and factuality for shorter and longer texts as well. We outline these results in Figure 4. While certain models perform better (Gemini and Nomic v1), performance is generally consistent across the board. OpenAI's Text Embedding 3 is a notable exception, with performance consistently increasing with response length, independent of model size.

Effect of Choice of Embedding Model. We analyze how the size and type of embedding model impact performance. As shown in Figure 5a, *Semantic Isotropy* consistently outperforms LUQ across all three datasets/scoring methods and embedding models. Notably, even the least effective embedding model yields improvements over LUQ, while the best performs substantially better. These results highlight the robustness of *Semantic Isotropy*, demonstrating its effectiveness regardless of the embedding model used. We compare small and large variants of closed source models such as those provided by OpenAI and Cohere. In both cases, the performance difference is negligible and likely attributable to sampling noise. This finding demonstrates that semantic isotropy is robust even at small model sizes.

Effect of Number of Sampled Responses. Figure 5b illustrates how performance scales with the number of sampled responses. Encouragingly, semantic isotropy scoring does not require many samples (\approx 6-8) to achieve comparable performance to its implied asymptotic level (at N=10 samples), across both small and large models. While the method is generally scalable, this result confirms that it can deliver consistent truthfulness indicators with minimal overhead.

 Effect of Generator Capability Level. In addition to the results discussed in Section 4.3, where we used long-form responses generated by Phi 3.5 Mini Instruct and Llama 3.1 8B Instruct, we also considered long-form responses generated by a more capable, closed-weight model, GPT 4.1 Mini. Applying the Segment-Score method to long-form responses for FS-BIO and TriviaQA, we find that the absolute performance of semantic isotropy scoring as well as its relative performance compared to the baselines significantly increases on FS-BIO but decreases on TriviaQA. On TriviaQA, semantic isotropy scoring achieves a higher R^2 than all baselines using Nomic v1 embeddings, but the best absolute R^2 value is barely above 10%, implying that almost none of the variation in the factuality scores can be explained by any of the methods used. Since, for FS-BIO, semantic isotropy scoring performs better on long-form responses generated by GPT 4.1 Mini than on long-form responses generated by Phi 3.5 Mini Instruct and Llama 3.1 8B Instruct, the poor performance (across methods) on TriviaQA is unlikely due to increased model capability per se and more likely due to dataset idiosyncrasies, such as inherent entity ambiguities, or model pre-training artifacts. We present a detailed analysis and discussion in Appendix C.1.

Effect of Choice of Isotropy Measure. In Table 1 we ablate a variety of alternative measures of isotropy, among them the Frobenius norm $\|K_E^{\cos}\|$, $\log(\det(K_E^{\cos}))$ and $\operatorname{trace}(K_E^{\cos})^{-1}$. We find that von Neumann entropy-based isotropy scoring performs best overall, outperforming other measures on average across different embedding model choices and dataset settings. However, most of these measures demonstrate nearly equivalent performance, further underscoring the robustness of semantic isotropy as a proxy for nonfactuality.

Table 1: Comparison of different isotropy measures: Analysis of various isotropy measures, each of which can represent the isotropy condition. Comparing explained variance (R^2) of Factuality on ≈ 500 word length responses generated using Phi 3.5 Mini Instruct on the **TriviaQA** (SS) dataset, across various embedding models. 1,500 Bootstrapped samples are used to generate 1-SD error bars.

Model	Frobenius	Inv. Trace	LogDet	vNE(OURS)	
Gemini	0.412 ± 0.02	0.376 ± 0.02	$\textbf{0.431} \pm \textbf{0.02}$	0.43 ± 0.02	
Nomic V1	0.326 ± 0.02	0.329 ± 0.02	$\textbf{0.393} \pm \textbf{0.02}$	0.354 ± 0.02	
OpenAI (L)	0.257 ± 0.02	0.208 ± 0.02	0.255 ± 0.02	$\textbf{0.266} \pm \textbf{0.02}$	
OpenAI (S)	0.274 ± 0.02	0.167 ± 0.02	0.241 ± 0.02	$\textbf{0.279} \pm \textbf{0.02}$	
Cohere v4.0	0.272 ± 0.02	0.222 ± 0.02	0.259 ± 0.02	$\textbf{0.276} \pm \textbf{0.02}$	
Cohere v3.0	0.272 ± 0.02	0.221 ± 0.02	0.261 ± 0.02	$\textbf{0.276} \pm \textbf{0.02}$	
Qwen 2 (7B)	$\textbf{0.27} \pm \textbf{0.02}$	0.169 ± 0.02	0.232 ± 0.02	0.268 ± 0.02	
SFR Mistral (7B)	0.291 ± 0.02	0.186 ± 0.02	0.268 ± 0.02	$\textbf{0.297} \pm \textbf{0.02}$	

Computational Considerations. One of the key drawbacks of existing methods such as LUQ and Graph Uncertainty is the $\mathbf{O}(MN^2)$ complexity, where M is the average number of facts or segments per response and N is the number of sampled responses. Even for Semantic Entropy and derived methods, the computation of entailment scores is $\mathbf{O}(N^2)$. Entailment scores are also computationally expensive as premises and hypothesis cannot be naively entailed in a vectorized fashion. In contrast, while semantic isotropy is naively $\mathbf{O}(N^2)$, where N is the number of samples, the ability to compute it in a parallelized vectorized way yields significant performance improvements. We generally find that for scoring one batch of N=20 responses for a given topic of $\approx\!500$ words using $\mathcal I$ requires 1.8 ± 0.05 seconds on a V100 GPU (amortizing the overhead over 20 trials), compared to 302 ± 48 seconds for LUQ-Atomic, using models of comparable parameter counts—that is, Nomic v1 for $\mathcal I$ and Deberta V3 Large (MNLI) (Manakul et al., 2023) for LUQ.

6 Conclusion

In this work, we introduced *Semantic Isotropy*, proposed a computationally inexpensive method for semantic isotropy scoring, presented *Segment-Score* for cheap long-form ground truth data generation, and demonstrated that semantic isotropy scores yield state-of-the-art factuality prediction in long-form natural language generation. We found that semantic isotropy is highly predictive of nonfactuality in long-form LLM responses to open-ended prompts and that semantic-isotropy scores outperform alternative approaches—in most cases by a large margin. We also found that, in most of the cases considered in our analysis, semantic isotropy scoring is robust to the choice of embedding model, the choice of isotropy measure, the response length, and the number of samples used. We release code to reproduce the empirical evaluation and implement the *Segment-Score* method, and we make the *Segment-Score*-annotated dataset publicly available to facilitate further research into assessing factuality in long-form text generation. Our findings suggest that semantic isotropy is a simple and yet effective predictor of nonfactuality and open the door to more reliable and cost-effective evaluation of long-form LLM text generation at scale.

REPRODUCIBILITY STATEMENT

We have made a significant effort to ensure the reproducibility of our results. An anonymized implementation of our method is provided at anonymous.4open.science/r/semantic_isotropy-C927, which includes training, evaluation, and analysis scripts. The experimental setup—including hyperparameters, model configurations, and sampling parameters are described in Section 4.1. All datasets used in our experiments are publicly available, and we additionally provide scripts for data preparation. Our experiments rely on open-weight language models and certain API-only models; instructions for reproducing results with both included in the README.md file in the released code. Infrastructure and compute requirements are documented in Appendix B.1. Finally, we intend to release the dataset constructed for this study upon completion of the review process, enabling direct benchmarking of our results.

REFERENCES

486

487 488

489

490

491

492

493

494

495

496

497 498

499 500

501

504

505

506

507

509

510

511

512

513

514

515

516

517

519

520 521

522 523

524

525

526

527 528

529

530

531

532

534

535

536

538

Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Martin Cai, Qin Cai, Vishrav Chaudhary, Dong Chen, Dongdong Chen, Weizhu Chen, Yen-Chun Chen, Yi-Ling Chen, Hao Cheng, Parul Chopra, Xiyang Dai, Matthew Dixon, Ronen Eldan, Victor Fragoso, Jianfeng Gao, Mei Gao, Min Gao, Amit Garg, Allie Del Giorno, Abhishek Goswami, Suriya Gunasekar, Emman Haider, Junheng Hao, Russell J. Hewett, Wenxiang Hu, Jamie Huynh, Dan Iter, Sam Ade Jacobs, Mojan Javaheripi, Xin Jin, Nikos Karampatziakis, Piero Kauffmann, Mahoud Khademi, Dongwoo Kim, Young Jin Kim, Lev Kurilenko, James R. Lee, Yin Tat Lee, Yuanzhi Li, Yunsheng Li, Chen Liang, Lars Liden, Xihui Lin, Zeqi Lin, Ce Liu, Liyuan Liu, Mengchen Liu, Weishung Liu, Xiaodong Liu, Chong Luo, Piyush Madan, Ali Mahmoudzadeh, David Majercak, Matt Mazzola, Caio César Teodoro Mendes, Arindam Mitra, Hardik Modi, Anh Nguyen, Brandon Norick, Barun Patra, Daniel Perez-Becker, Thomas Portet, Reid Pryzant, Heyang Qin, Marko Radmilac, Liliang Ren, Gustavo de Rosa, Corby Rosset, Sambudha Roy, Olatunji Ruwase, Olli Saarikivi, Amin Saied, Adil Salim, Michael Santacroce, Shital Shah, Ning Shang, Hiteshi Sharma, Yelong Shen, Swadheen Shukla, Xia Song, Masahiro Tanaka, Andrea Tupini, Praneetha Vaddamanu, Chunyu Wang, Guanhua Wang, Lijuan Wang, Shuohang Wang, Xin Wang, Yu Wang, Rachel Ward, Wen Wen, Philipp Witte, Haiping Wu, Xiaoxia Wu, Michael Wyatt, Bin Xiao, Can Xu, Jiahang Xu, Weijian Xu, Jilong Xue, Sonali Yaday, Fan Yang, Jianwei Yang, Yifan Yang, Ziyi Yang, Donghan Yu, Lu Yuan, Chenruidong Zhang, Cyril Zhang, Jianwen Zhang, Li Lyna Zhang, Yi Zhang, Yue Zhang, Yunan Zhang, and Xiren Zhou. Phi-3 technical report: A highly capable language model locally on your phone, 2024. URL https://arxiv.org/abs/2404.14219. 6, 17

- Cohere AI. Introducing embed v3. https://cohere.com/blog/introducing-embed-v3, 2023. 15
- Cohere AI. Introducing embed 4: Multimodal search for business. https://cohere.com/blog/embed-4, 2025. 15
- Shuang Ao, Stefan Rueger, and Advaith Siddharthan. Css: Contrastive semantic similarity for uncertainty quantification of llms, 2024. URL https://arxiv.org/abs/2406.03158. 3
- Neil Band, Xuechen Li, Tengyu Ma, and Tatsunori Hashimoto. Linguistic calibration of long-form generations, 2024. URL https://arxiv.org/abs/2404.00474. 3
- Chao Chen, Kai Liu, Ze Chen, Yi Gu, Yue Wu, Mingyuan Tao, Zhihang Fu, and Jieping Ye. INSIDE: LLMs' internal states retain the power of hallucination detection. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=Zj12nzlQbz. 3, 7, 15
- Shiqi Chen, Yiran Zhao, Jinghan Zhang, I-Chun Chern, Siyang Gao, Pengfei Liu, and Junxian He. Felm: Benchmarking factuality evaluation of large language models, 2023. URL https://arxiv.org/abs/2310.00741. 3
- Tiejin Chen, Xiaoou Liu, Longchao Da, Jia Chen, Vagelis Papalexakis, and Hua Wei. Uncertainty quantification of large language models through multi-dimensional responses, 2025. URL https://arxiv.org/abs/2502.16820. 3

```
Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. Detecting hallucinations in large language models using semantic entropy. Nature, 630(8017):625–630, jun 2024. ISSN 1476-4687. doi: 10.1038/s41586-024-07421-0. URL https://doi.org/10.1038/s41586-024-07421-0. 1, 3, 7, 15
```

- Yashvir Singh Grewal, Edwin V. Bonilla, and Thang D Bui. Improving uncertainty quantification in large language models via semantic embeddings, 2025. URL https://openreview.net/forum?id=N4mb3MBV6J. 3
- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. A survey on automated fact-checking. *Transactions of the Association for Computational Linguistics*, 10:178–206, 02 2022. ISSN 2307-387X. doi: 10.1162/tacl_a_00454. URL https://doi.org/10.1162/tacl_a_00454. 3
- Chao-Wei Huang and Yun-Nung Chen. Factalign: Long-form factuality alignment of large language models, 2024. URL https://arxiv.org/abs/2410.01691. 3
- Ziwei Ji, Delong Chen, Etsuko Ishii, Samuel Cahyawijaya, Yejin Bang, Bryan Wilie, and Pascale Fung. Llm internal states reveal hallucination risk faced with a query, 2024. URL https://arxiv.org/abs/2407.03282. 3
- Mingjian Jiang, Yangjun Ruan, Prasanna Sattigeri, Salim Roukos, and Tatsunori Hashimoto. Graph-based uncertainty metrics for long-form language model generations. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37, pages 32980–33006. Curran Associates, Inc., 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/39c5871aa13be86ab978cba7069cbcec-Paper-Conference.pdf. 1, 7
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension. In Regina Barzilay and Min-Yen Kan, editors, *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics* (*Volume 1: Long Papers*), pages 1601–1611, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1147. URL https://aclanthology.org/P17-1147/. 3. 6
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El-Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, Deep Ganguli, Danny Hernandez, Josh Jacobson, Jackson Kernion, Shauna Kravec, Liane Lovitt, Kamal Ndousse, Catherine Olsson, Sam Ringer, Dario Amodei, Tom Brown, Jack Clark, Nicholas Joseph, Ben Mann, Sam McCandlish, Chris Olah, and Jared Kaplan. Language models (mostly) know what they know, 2022. URL https://arxiv.org/abs/2207.05221. 3
- Robert Kirk, Ishita Mediratta, Christoforos Nalmpantis, Jelena Luketina, Eric Hambro, Edward Grefenstette, and Roberta Raileanu. Understanding the effects of rlhf on llm generalisation and diversity, 2024. URL https://arxiv.org/abs/2310.06452. 2
- Jannik Kossen, Jiatong Han, Muhammed Razzak, Lisa Schut, Shreshth Malik, and Yarin Gal. Semantic entropy probes: Robust and cheap hallucination detection in llms, 2024. URL https://arxiv.org/abs/2406.15927. 3, 19
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention, 2023. URL https://arxiv.org/abs/2309.06180. 6, 19
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles, 2017. URL https://arxiv.org/abs/1612.01474. 2
- Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human falsehoods, 2022. URL https://arxiv.org/abs/2109.07958. 3
- Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. Generating with confidence: Uncertainty quantification for black-box large language models, 2024. URL https://arxiv.org/abs/2305.19187. 2, 7, 15

- Linyu Liu, Yu Pan, Xiaocheng Li, and Guanting Chen. Uncertainty estimation and quantification for llms: A simple supervised approach, 2024. URL https://arxiv.org/abs/2404.15993. 1
 - Andrey Malinin and Mark Gales. Uncertainty estimation in autoregressive structured prediction, 2021. URL https://arxiv.org/abs/2002.07650. 2, 7, 15
 - Potsawee Manakul, Adian Liusie, and Mark J. F. Gales. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models, 2023. URL https://arxiv.org/abs/2303.08896. 1, 3, 7, 9, 15, 19
 - Hannah Mayer, Lareina Yee, Michael Chui, and Roger Roberts. Superagency in the workplace: Empowering people to unlock AI's full potential. Technical report, McKinsey, 2025. 1
 - Rui Meng, Ye Liu, Shafiq Rayhan Joty, Caiming Xiong, Yingbo Zhou, and Semih Yavuz. Sfrembedding-mistral: Enhance text retrieval with transfer learning. Salesforce AI Research Blog, 2024. URL https://www.salesforce.com/blog/sfr-embedding/. 15, 16
 - Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. Factscore: Fine-grained atomic evaluation of factual precision in long form text generation, 2023. URL https://arxiv.org/abs/2305.14251. 2, 3, 6
 - Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. Mteb: Massive text embedding benchmark, 2023. URL https://arxiv.org/abs/2210.07316. 7
 - Alexander Nikitin, Jannik Kossen, Yarin Gal, and Pekka Marttinen. Kernel language entropy: Fine-grained uncertainty quantification for Ilms from semantic similarities. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37, pages 8901–8929. Curran Associates, Inc., 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/10c456d2160517581a234dfde15a7505-Paper-Conference.pdf. 1, 3, 5, 7
 - Zach Nussbaum, John X. Morris, Brandon Duderstadt, and Andriy Mulyar. Nomic embed: Training a reproducible long context text embedder, 2025. URL https://arxiv.org/abs/2402.01613.7, 15, 16
 - OpenAI. Gpt-4 technical report. arXiv preprint arXiv:2303.08774, 2023. 2, 6, 7, 15, 16, 17
 - Xin Qiu and Risto Miikkulainen. Semantic density: Uncertainty quantification for large language models through confidence measurement in semantic space, 2024. URL https://arxiv.org/abs/2405.13845. 1, 3
 - Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions for machine comprehension of text, 2016. URL https://arxiv.org/abs/1606.05250. 3
 - Siva Reddy, Danqi Chen, and Christopher D. Manning. Coqa: A conversational question answering challenge, 2019. URL https://arxiv.org/abs/1808.07042. 3
 - Jie Ren, Jiaming Luo, Yao Zhao, Kundan Krishna, Mohammad Saleh, Balaji Lakshminarayanan, and Peter J. Liu. Out-of-distribution detection and selective generation for conditional language models, 2023. URL https://arxiv.org/abs/2209.15558. 7, 15
 - Emanuele Ricco, Lorenzo Cima, and Roberto Di Pietro. Hallucination detection: A probabilistic framework using embeddings distance analysis, 2025. URL https://arxiv.org/abs/2502.08663. 19
 - Sheikh Shafayat, Eunsu Kim, Juhyun Oh, and Alice Oh. Multi-fact: Assessing factuality of multilingual llms using factscore, 2024. URL https://arxiv.org/abs/2402.18045. 3
 - Ola Shorinwa, Zhiting Mei, Justin Lidard, Allen Z. Ren, and Anirudha Majumdar. A survey on uncertainty quantification of large language models: Taxonomy, open research challenges, and future directions, 2024. URL https://arxiv.org/abs/2412.05563. 1

- Yixiao Song, Yekyung Kim, and Mohit Iyyer. Veriscore: Evaluating the factuality of verifiable claims in long-form text generation, 2024. URL https://arxiv.org/abs/2406.19276. 3
- Gemini Team, 2024. URL https://arxiv.org/abs/2312.11805. 7, 15, 18
 - Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models, 2023. URL https://arxiv.org/abs/2302.13971. 6, 17
 - John von Neumann and Robert T. Beyer. *Mathematical Foundations of Quantum Mechanics: New Edition*. Princeton University Press, ned new edition edition, 2018. ISBN 9780691178561. URL http://www.jstor.org/stable/j.ctt1wq8zhp. 4
 - Cunxiang Wang, Xiaoze Liu, Yuanhao Yue, Xiangru Tang, Tianhang Zhang, Cheng Jiayang, Yunzhi Yao, Wenyang Gao, Xuming Hu, Zehan Qi, Yidong Wang, Linyi Yang, Jindong Wang, Xing Xie, Zheng Zhang, and Yue Zhang. Survey on factuality in large language models: Knowledge, retrieval and domain-specificity, 2023. URL https://arxiv.org/abs/2310.07521. 1
 - Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. Improving text embeddings with large language models, 2024a. URL https://arxiv.org/abs/2401.00368. 15
 - Xiyao Wang, Linfeng Song, Ye Tian, Dian Yu, Baolin Peng, Haitao Mi, Furong Huang, and Dong Yu. Towards self-improvement of llms via mcts: Leveraging stepwise knowledge with curriculum preference learning, 2024b. URL https://arxiv.org/abs/2410.06508. 19
 - Yu-Hsiang Wang, Andrew Bai, Che-Ping Tsai, and Cho-Jui Hsieh. Clue: Concept-level uncertainty estimation for large language models, 2024c. URL https://arxiv.org/abs/2409.03021. 3
 - Ziyu Wang and Chris Holmes. On subjective uncertainty quantification and calibration in natural language generation, 2024. URL https://arxiv.org/abs/2406.05213. 1
 - Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu, Nathan Hu, Jie Huang, Dustin Tran, Daiyi Peng, Ruibo Liu, Da Huang, Cosmo Du, and Quoc V. Le. Long-form factuality in large language models, 2024. URL https://arxiv.org/abs/2403.18802. 3
 - Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics, 2018. URL http://aclweb.org/anthology/N18-1101. 3
 - An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. Qwen2 technical report, 2024a. URL https://arxiv.org/abs/2407.10671. 15, 16
 - An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report, 2025. URL https://arxiv.org/abs/2505.09388. 15

Ruihan Yang, Caiqi Zhang, Zhisong Zhang, Xinting Huang, Sen Yang, Nigel Collier, Dong Yu, and Deqing Yang. Logu: Long-form generation with uncertainty expressions, 2024b. URL https://arxiv.org/abs/2410.14309. 3

Caiqi Zhang, Fangyu Liu, Marco Basaldella, and Nigel Collier. Luq: Long-text uncertainty quantification for llms, 2024. URL https://arxiv.org/abs/2403.20279. 1, 3, 6, 7, 8, 15, 17

Jiaxin Zhang. Modern monte carlo methods for efficient uncertainty quantification and propagation: A survey, 2020. URL https://arxiv.org/abs/2011.00680. 2

Yu Zhao, Huifeng Yin, Bo Zeng, Hao Wang, Tianqi Shi, Chenyang Lyu, Longyue Wang, Weihua Luo, and Kaifu Zhang. Marco-o1: Towards open reasoning models for open-ended solutions, 2024. URL https://arxiv.org/abs/2411.14405.

APPENDIX

TABLE OF CONTENTS

A	Full	Full Experimental Results						
В	Experimental Details							
	B.1	Infrastructure Requirements	15					
	B.2	Selecting the Appropriate Embedding	15					
	B.3	Segment-Score Implementation Specifics	16					
C	Abla	ation Studies	16					
	C.1	Response generation using OpenAI GPT 4.1 Mini	16					
	C.2	Comparing the efficiency of the two Scoring Methods	17					
D	Lim	itations	19					

A FULL EXPERIMENTAL RESULTS

	TriviaQA Entities (Segment Score)			FactScore-BIO (Segment Score)		FactScore-BIO (FactScore)			
Measure / Metric	Llama 3.1 8B	Phi 3.5 Mini	GPT 4.1 Mini	Llama 3.1 8B	Phi 3.5 Mini	GPT 4.1 Mini	Llama 3.1 8B	Phi 3.5 Mini	GPT 4.1 Mini
Perplexity (Ren et al., 2023)	0.092 ± 0.02	0.263 ± 0.03	0.00074 ± 0.003	0.0933 ± 0.04	0.0305 ± 0.03	0.0137 ± 0.007	0.0578 ± 0.04	0.0622 ± 0.04	0.0158 ± 0.02
LN Entropy (Malinin and Gales, 2021)	0.091 ± 0.02	0.26 ± 0.03	0.0306 ± 0.01	0.0872 ± 0.04	0.0316 ± 0.03	0.0241 ± 0.02	0.0535 ± 0.04	0.0608 ± 0.04	0.00625 ± 0.009
U.Eign (Lin et al., 2024)	0.062 ± 0.02	0.181 ± 0.03	0.00166 ± 0.002	0.275 ± 0.05	0.0529 ± 0.03	0.0155 ± 0.02	0.344 ± 0.06	0.0279 ± 0.02	0.0185 ± 0.02
U.Deg (Lin et al., 2024)	0.0457 ± 0.01	0.117 ± 0.03	0.00286 ± 0.003	0.255 ± 0.07	0.0452 ± 0.03	0.0254 ± 0.02	0.28 ± 0.05	0.0224 ± 0.02	0.0255 ± 0.02
U.NLI (Manakul et al., 2023)	0.0767 ± 0.02	0.167 ± 0.03	0.00164 ± 0.002	0.144 ± 0.06	0.0552 ± 0.03	0.0128 ± 0.01	0.0916 ± 0.04	0.0299 ± 0.02	0.0263 ± 0.02
EigenScore (Chen et al., 2024)	0.0158 ± 0.009	0.00131 ± 0.002	N/A	0.0683 ± 0.04	0.021 ± 0.02	N/A	0.107 ± 0.05	0.0309 ± 0.03	N/A
LUQ (Zhang et al., 2024)	0.03 ± 0.01	0.231 ± 0.03	0.1 ± 0.02	0.381 ± 0.05	0.104 ± 0.05	0.116 ± 0.04	0.657 ± 0.07	0.367 ± 0.06	0.192 ± 0.06
Semantic Entropy (Farquhar et al., 2024)	0.0291 ± 0.01	0.0817 ± 0.02	0.00242 ± 0.003	0.156 ± 0.07	0.0298 ± 0.02	0.00692 ± 0.009	0.136 ± 0.06	0.00868 ± 0.01	0.0152 ± 0.02
I: Gemini [†] (Team, 2024)	0.244 ± 0.03	0.427 ± 0.02	0.028 ± 0.01	0.435 ± 0.04	0.235 ± 0.07	0.342 ± 0.04	0.7 ± 0.06	0.661 ± 0.05	0.48 ± 0.08
I: OpenAI Small [†] (OpenAI, 2023)	0.156 ± 0.02	0.282 ± 0.02	0.00196 ± 0.003	0.441 ± 0.04	0.265 ± 0.07	0.314 ± 0.04	0.65 ± 0.06	0.593 ± 0.05	0.38 ± 0.08
I: OpenAI Large† (OpenAI, 2023)	0.155 ± 0.02	0.269 ± 0.02	0.00215 ± 0.003	0.44 ± 0.03	0.268 ± 0.06	0.315 ± 0.04	0.649 ± 0.06	0.614 ± 0.05	0.376 ± 0.08
I: Nomic v1 Embed (Nussbaum et al., 2025)	0.27 ± 0.02	0.356 ± 0.02	0.136 ± 0.03	0.446 ± 0.04	0.273 ± 0.06	0.417 ± 0.04	0.737 ± 0.05	0.6 ± 0.05	0.547 ± 0.08
I: Qwen 2 (7B) (Yang et al., 2024a)	0.0811 ± 0.02	0.268 ± 0.03	0.0203 ± 0.01	0.488 ± 0.04	0.355 ± 0.06	0.42 ± 0.04	0.733 ± 0.05	0.589 ± 0.05	0.569 ± 0.08
I: Mistral (SalesForce) (Meng et al., 2024)	0.0842 ± 0.02	0.296 ± 0.03	0.0128 ± 0.01	0.446 ± 0.04	0.29 ± 0.06	0.46 ± 0.04	0.714 ± 0.05	0.533 ± 0.06	$\textbf{0.607} \pm \textbf{0.08}$
I: Mistral (E5) (Wang et al., 2024a)	0.112 ± 0.02	0.318 ± 0.02	0.00276 ± 0.004	0.445 ± 0.04	0.281 ± 0.06	0.471 ± 0.04	0.714 ± 0.06	0.508 ± 0.06	0.604 ± 0.08
I: Cohere v4.0 † (AI, 2025)	0.11 ± 0.02	0.277 ± 0.03	0.0035 ± 0.004	0.504 ± 0.03	0.299 ± 0.07	0.408 ± 0.04	0.747 ± 0.06	0.671 ± 0.05	0.5 ± 0.08
I: Cohere v3.0 † (AI, 2023)	0.111 ± 0.02	0.277 ± 0.03	0.00325 ± 0.004	0.502 ± 0.04	0.265 ± 0.07	0.408 ± 0.04	0.75 ± 0.06	0.665 ± 0.05	0.503 ± 0.08
I: Cohere v3.0 (Lite) † (AI, 2023)	0.11 ± 0.02	0.278 ± 0.03	0.00324 ± 0.004	0.502 ± 0.03	0.236 ± 0.06	0.41 ± 0.04	0.749 ± 0.05	0.661 ± 0.05	0.503 ± 0.08
I: Qwen 3 (0.6B) (Yang et al., 2025)	0.0477 ± 0.02	0.172 ± 0.02	0.0147 ± 0.01	0.466 ± 0.04	0.305 ± 0.06	0.421 ± 0.04	0.719 ± 0.06	0.525 ± 0.06	0.593 ± 0.08
I: Qwen 3 (4B) (Yang et al., 2025)	0.00524 ± 0.006	0.165 ± 0.03	0.135 ± 0.03	0.441 ± 0.04	0.307 ± 0.06	0.368 ± 0.05	0.713 ± 0.06	0.588 ± 0.05	0.554 ± 0.08
I: Qwen 3 (8B) (Yang et al., 2025)	0.00247 ± 0.004	0.116 ± 0.02	0.126 ± 0.03	0.432 ± 0.04	0.279 ± 0.06	0.331 ± 0.05	0.709 ± 0.06	0.561 ± 0.06	0.528 ± 0.07

Table 2: Main Experimental Results. Comparison of semantic isotropy with other uncertainty metrics across TriviaQA Entities (Segment Score) and FactScore-Bio (Segment Score and FactScore). Each dataset column is subdivided by base model: Llama 3.1 8B, Phi 3.5 Mini, and OpenAI 4.1 Mini. Values are R^2 (explained variance) of a simple linear model of Factuality \sim Isotropy score. 1500 Bootstrapped samples are used to generate 1-SD error bars. †denotes API-only models. Note that Eigenscore values are not generated for GPT 4.1 Mini as it is a black-box model and the internal layer activations are not available.

B EXPERIMENTAL DETAILS

B.1 Infrastructure Requirements

Our work was performed using Google Cloud Platform compute to run inference on the white-box models using a single node with up to 4 NVIDIA Tesla V100 GPUs. GPT 4.1 Mini was used as the Oracle LLM in Algorithm 1. OpenAI text-embedding-3-small, text-embedding-3-large, Gemini gemini-embedding-001, Cohere Embedding v4.0, v3.0 (AI, 2023; 2025) were also used to compute response embeddings for semantic isotropy computation. Additionally, OpenAI GPT 4.1, Anthropic Claude 3.5 Sonnet, and DeepSeek V3 were also used in verification runs to study robustness of the scoring pipeline. Overall, we utilize approximately 1,440 GPU hours (in addition to CPU compute time) to generate samples and score the responses.

B.2 SELECTING THE APPROPRIATE EMBEDDING

Our experimental results require an embedding vector $z_i \in \mathbb{R}^D$. Closed-weight models return a single vector per response and allow the user to configure embedding side (D) as an input. We use the model default for

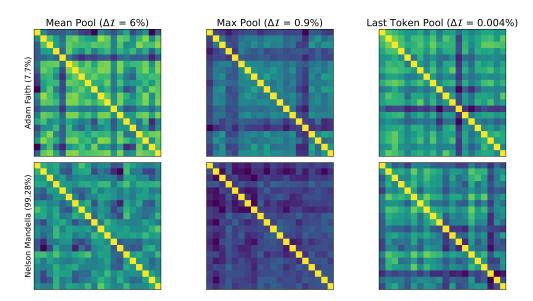


Figure 6: Nomic v1 Embedding Choice. K_E^{cos} for two sample topics: Adam Faith (low factuality, top row) and Nelson Mandella (high factuality, bottom row). The number in the parenthesis next to the topic name is the average factuality across sampled responses. $\Delta \mathcal{I}$ for each pooling method denotes the percentage difference in semantic isotropy scores between the two topics (as a percentage of \mathcal{I} for Nelson Mandella). We infer that Mean Pooling for Nomic V1 yields the most informative matrices for an isotropy analysis.

Gemini (768), OpenAI (1536) and Cohere (1024). However, Open-weight embedding models like Qwen2 (Yang et al., 2024a) and Mistral (Meng et al., 2024) specify using the last token embedding in the final hidden state. This can be found in their reference implementations on HuggingFace. In contrast, for Nomic V1 (Nussbaum et al., 2025), authors do not recommend a pooling method and encourage users to find methods that work well for their use case. We find that the last token embedding method yields a weak semantic isotropy score. Instead, we use the mean pooling over the token dimension. Given model activations of the last hidden layer as a matrix $E \in \mathbb{R}^{L \times D}$, where E is the number of tokens in the response and E0 the embedding dimension, we produce E1 mean (E1, dim=1) such that E1 such that E2 nomic eq. (E3. In Figure 6, we show the difference between a highly trustworthy and a low trustworthy topic by mean, max and last token pooling for Nomic v1.

B.3 SEGMENT-SCORE IMPLEMENTATION SPECIFICS

Unlike FactScore, the Segment-Score algorithm is achieved through a single prompt per response. The basic prompt template for our method can be found in Figure 11. The prompt is structured into three components. First, an instruction block that covers the overall objective of the scoring exercise, expected data, desired output structure, and specific handling for edge cases. Next, we provide two manually curated examples to guide the model's output. One such example can be found in Figure 12. Last, we provide the entity name, contents of the reference document and response to be scored. The model is prompted to respond using a XML tag structure which simplifies the data extraction process. We also force the model to respond with only '0' or '1' for the output classes. For 99.8% of segments, the top 2 tokens are the '0' or '1' class labels, and we can use the log-probabilities from the API to generate a normalized probability score for each labeled segment.

C ABLATION STUDIES

C.1 RESPONSE GENERATION USING OPENAI GPT 4.1 MINI

To understand the impact of the generator model on the performance of the semantic isotropy score, we also generate responses using OpenAI GPT 4.1 Mini (OpenAI, 2023), scored using Segment-Score. Figure 7 shows the results of this experiment in the same format as our main results. For the FS-BIO dataset, we see semantic isotropy remains predictive of overall response factualtiy. However, we see a notable drop in performance in the case of TriviaQA. Notably, this extends to all benchmark metrics as well. We also observe much greater variance in the performance across embedding models, suggesting there are some settings where the choice of embedding model is crucial.

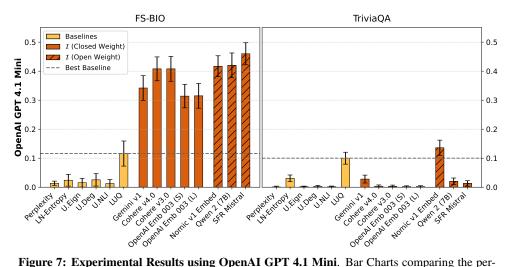


Figure 7: Experimental Results using OpenAI GPT 4.1 Mini. Bar Charts comparing the performance (as measured by the \mathbb{R}^2 of a linear model of factuality with semantic isotropy as the explanatory variable; implemented using various embedding models and benchmark uncertainty metrics on both the FS-BIO ant TriviaQA datasets using the Segment-Score (SS) scoring algorithm. Left hand side: Benchmark UQ metrics. Right Hand Side: semantic isotropy (\mathcal{I}) implemented using various embedding models. Same experimental setting as in Figure 2. We observe that performance on generations using OpenAI GPT 4.1 Mini is significantly lower than other generation models, however this extends to all benchmark metrics as well. Semantic isotropy continues to outperform all other baselines. Note that Eigenscore values are not generated for GPT 4.1 Mini as it is a black-box model and the internal layer activations are not available.

Consequently, it was important to understand the impact of the generation model on the performance of the semantic isotropy score. Zhang et al. (2024) note that the LUQ method does not perform well using GPT 4, a fact we were able to replicate.

"We also observe that LUQ is better suited for models with relatively lower factuality and a lack of self-expressiveness regarding uncertainty. For models with high factuality capabilities, such as GPT-4, LUQ only demonstrates a moderate correlation with factuality scores" - Zhang et al. (2024)

One would expect that GPT 4.1 Mini would perform better in terms of instruction following and maintaining consistent structure. This is anecdotally observed in the token count distribution (see Figure 9b), where GPT 4.1 Mini adhere's much better to the 500 word guidance given in the prompt across different topic areas, as compared to the other generative models being considered.

When studying the distribution of semantic isotropy scores, in Figure 9a, we observe that OpenAI GPT 4.1 Mini's scores are smaller in magnitude and more tightly clustered. This highlighted that there may be some inherrent structural properties of the OpenAI GPT 4.1 Mini generations that manifested in the embeddings in this specific way. To tackle this, we performed a comparitive study between generations from GPT 4.1 Mini (OpenAI, 2023), Phi 3.5 Mini (Abdin et al., 2024) and Llama 3.1 8B (Touvron et al., 2023). Using a LLM as a judge setup, we asked OpenAI GPT 4.1 to compare and contrast 10 generations from GPT 4.1 Mini vs 10 from either Phi 3.5 Mini or Llama 3.1 8B for 50 different topics where the all the models had relatively low factuality (i.e. $\leq 33\%$ accuracy on average). The LLM generated comparisons were summarized into a brief analysis (See Figure 8). This highlighted that properties such as structure, tone and other qualitative aspects of a generation materially impacted the resulting embeddings. We expect that improvements in embedding models or models specifically tuned to disregard qualitative aspects of a generation could improve the performance of the semantic isotropy score, but leave this exploration for future work.

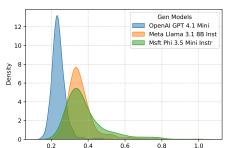
C.2 Comparing the efficiency of the two Scoring Methods

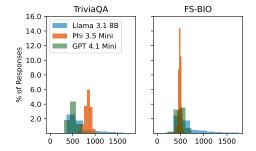
Segment-Score's overall process closely follows the FactScore algorithm. However, we make several key optimizations that improve runtime. First, unlike FactScore, we break the input text into atomic segments rather than extracting atomic facts. As a result, each sentence may be broken into at most 2-3 segments that need to be scored whereas FactScore may generate several atomic facts pertaining to each sentence segment. For FS-BIO,

Across the comparisons, Model A consistently demonstrates formal, polished, and structured writing with high factual accuracy, depth, and analytical rigor. Its responses are typically essay-like, varied in perspective, and nuanced, often employing advanced vocabulary, rhetorical devices, and critical engagement with broader themes such as legacy, cultural context, and impact.

In contrast, Model B tends to be more informal, conversational, and formulaic, frequently exhibiting repetition, superficial coverage, and factual errors. Model B's diversity often arises from inconsistency or error rather than true creativity, and it outputs are generally less coherent, less organized, and less engaging than Model A's.

Figure 8: Summary Analysis of Low Factuality Generations by GPT 4.1 Mini (Model A) vs Llama 3.1 8B (Model B).





- (a) Kernel Desity Estimation plots of semantic isotropy scores for TriviaQA. Kernel Density Estimation plots of semantic isotropy scores for TriviaQA, comparing the distributions of scores for GPT 4.1 Mini, Phi 3.5 Mini and Llama 3.1 8B. The embedding used in Gemini v001 (Team, 2024). We observe that GPT 4.1 Mini's scores are larger in magnitude and more tightly clustered.
- (b) Word Count by Generation Model and Dataset: Density histograms comparing the response length (in tokens) for each generative model and dataset for a prompt with a guidance length of 500 words. We observe that GPT 4.1 Mini is much better at instruction following that Llama 3.1 and Phi 3.5 Mini, sticking fairly closely to the intended word limit.

Figure 9: Response Characteristics: Comparing the characteristics of generated responses by model type and dataset. Distribution of semantic isotropy scores for TriviaQA using Gemini [Left] and Histogram of response lengths in tokens by dataset and generative model [Right].

this translates to 28 ± 6 distinct segments for each response in *Segment-Score*, while *FactScore* produces 100 ± 20 atomic facts $(2-\sigma \text{ interval})$.

Second, FactScore does not distinguish between verifiable and ambiguous facts which increases noise in the final results. We explicitly prompt the model to classify such examples as False. This generally biases our result lower, which can be seen in the middle plot of Figure 10.

Finally, our entire process is achieved using one single query to the oracle, while FactScore may use an order of magnitude more depending on the number of atomic facts extracted and the length of the ground truth document. Consequently, our method scales more favorably as the length of the response to be scored increases. Despite these differences, we see comparable results for **FS-BIO** in Table 2 for both scoring methods when considereing the various isotropy configurations and response generating models.

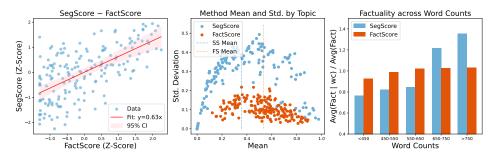


Figure 10: FS vs SS [FS-BIO]. Comparison of *FactScore* and *Segment-Score* on the **FS-BIO** dataset. *Segment-Score* yields a dataset with lower overall factuality and higher variance, however the relationship between the two scores is robust and consistent.

D LIMITATIONS

The main drawback of our method is the need for repeated sampling, which can be expensive to generate. Approaches such as Kossen et al. (2024) have shown that it is possible to train cheap probes that can a priori estimate uncertainty scores without needing repeated sampling. With the advent of reasoning models with long Chain-of-Thought reasoning contexts (Zhao et al., 2024) and Monte-Carlo Tree Search (Wang et al., 2024b), as well as fast inference engines with significantly higher throughput (Kwon et al., 2023), the relative overhead of repeated sampling has been greatly diminished, making methods like ours compelling. Additionally, while studies such as Manakul et al. (2023) have observed that hallucinated responses tend to diverge, and generally are inconsistent with each other, Ricco et al. (2025) find that this effect is not universal and introduce the taxonomy of *extrinsic* (characterized by arbitrary and orthogonal content) vs *intrinsic* (contradictory or anti-parallel content) hallucination. The semantic isotropy framework would likely fail in cases of *intrinsic* hallucinations. However, in practice this intrinsic hallucinations would be exceedingly rare, especially in the long-form response format.

```
1026
        // INSTRUCTION PROMPT FOR SEGMENT SCORE
1027
        You are an NLP segmentation and evaluation engine.
1028
       Examine the scenario below. You are given:
        1. The name of an entity/person/place/thing etc. in <entity> tags.
1029
        2. A reference document regarding the entity in <reference_doc> tags.
1030
       3. A response about the entity to evaluate in <response> tags.
1031
1032
        ### Your Tasks:
1033
        1. **Segmentation Task:**
          Segment the `<response>` into individual statements. Each statement can be a sentence,
1034
                phrase or word and should convey a single, complete, and independent piece of
1035
               information about the `<entity>`. Do not modify, rephrase, or paraphrase the
1036
               original text. Ensure no semantic overlaps exist between statements. Individual
1037
               proper nouns should be part of their own statement however when determining the
               appropriate classification, preceeding context can be used when appropriate.
1039
          Verify that the concatenated content of all statements exactly matches the original
               response.
1040
1041
          Format the segmented response as follows:
1042
1043
          <statements>
          <statement>Statement 1</statement> <class>Class 1</class>
1044
          <statement>Statement 2</statement> <class>Class 2</class>
1045
1046
          </statements>
1047
1048
1049
        2. **Factual Classification Task:**
          For each segmented `<statement>`, classify it as 1 (True) or 0 (False) based solely
1050
               on the information in the `<reference_doc>`. Follow these guidelines:
1051
            If a statement is factually accurate and supported by the `<reference_doc>`,
1052
               classify it as '1'.
1053
          - If a statement is inaccurate, unverifiable, or not supported by the `<reference_doc
1054
               >`, classify it as '0'.
          - If a statement is partially true, but contains incorrect or unsupported information,
1055
                classify it as '0'.
1056
          - Do not rely on any external knowledge or context beyond the `<reference_doc>`.
1057
          - Include only the classification for each statement. Do not provide any explanations
1058
                or additional information.
          - Specify the class in <class> tags.
1059
          - The ONLY valid class values are `1` and `0`. No other values or words should appear
                within the `<class>` tags.
1061
1062
        3. **Error Handling:**
1063
          - If the `<response>` contains unparseable text, incomplete sentences, or conflicting
                information that cannot be resolved using the `<reference_doc>`, include the
1064
               flagged statement as is and classify it as 'False'.
1065
1066
       Examples:
1067
        ###### EXAMPLE 1 #####
1068
       {{example_one}}
        1069
        ###### EXAMPLE 2 #####
1070
        {{example_two}}
1071
        ##########################
1072
1073
       Entity:
        <entity>
1074
       {entity}
1075
       </entity>
1076
1077
       Reference Document:
1078
        <reference_doc>
1079
        {reference_doc}
        </reference_doc>
        Response to Evaluate:
                                                   20
        <response>
        {response}
```

```
1080
1081
        <entity>
1082
       London, UK
1083
        </entity>
1084
       Reference Document:
        <reference_doc>
1086
       London, England's capital, boasts a rich history spanning millennia. Founded by the
1087
            Romans as Londinium around 47 AD, it became a major port and trading center. After
1088
            the Roman withdrawal, Anglo-Saxons established Lundenwic, which later fell to
1089
            Viking raids. The Norman Conquest in 1066 led to the construction of the Tower of
            London, a symbol of royal power. London thrived during the medieval period,
1090
            becoming a major center for trade, finance, and culture. It weathered plagues,
1091
            fires, and civil wars, emerging as a global metropolis and the heart of the British
1092
             Empire. Today, London remains a vibrant hub, blending its historical legacy with
1093
            modern dynamism, home to over 9 million people.
1094
        </reference_doc>
1095
       Response to Evaluate:
1096
        <response>
       London, the capital city of England and the United Kingdom, is a vibrant metropolis
1098
            steeped in history and brimming with modern energy. With a population of over 9
1099
            million people, it stands as one of the world's most influential global cities,
            known for its diverse culture, iconic landmarks, and rich heritage.
1100
1101
        The city's history stretches back over three millennia, founded by the Romans as
1102
            Londinium in 43 AD. Throughout the centuries, London has played a pivotal role in
1103
            world affairs, serving as the heart of the British Empire and surviving tumultuous
1104
            events such as the Great Fire of 1666 and the Blitz during World War I.
1105
        Today, London is a melting pot of cultures, with over 300 languages spoken within its
1106
            boundaries. This diversity is reflected in its neighborhoods, each with its own
1107
            unique character and charm. From the trendy streets of Shoreditch to the upscale
1108
            boutiques of Mayfair, there's something for everyone in this cosmopolitan city.
1109
        </response>
1110
        Segmented and classified response:
1111
        <statements>
1112
        <statement>London, the capital city of England and the United Kingdom</statement> <class
1113
            >1</class>
1114
        <statement>is a vibrant metropolis steeped in history and brimming with modern energy</
            statement> <class>1</class>
1115
        <statement>With a population of over 9 million people</statement> <class>1</class>
1116
        <statement>it stands as one of the world's most influential global cities, known for its
1117
             diverse culture, iconic landmarks, and rich heritage.</statement> <class>1</class>
1118
        <statement>The city's history stretches back over three millennia, founded by the Romans
1119
             as Londinium in 43 AD</statement> <class>0</class>
        <statement>Throughout the centuries, London has played a pivotal role in world affairs,
1120
            serving as the heart of the British Empire</statement> <class>1</class>
1121
        <statement>and surviving tumultuous events such as the Great Fire of 1666</statement> <
1122
            class>1</class>
1123
        <statement>and the Blitz during World War I</statement> <class>0</class>
1124
        <statement>Today, London is a melting pot of cultures, with over 300 languages spoken
            within its boundaries</statement> <class>1</class>
1125
        <statement>This diversity is reflected in its neighborhoods, each with its own unique
1126
            character and charm.</statement> <class>1</class>
1127
        <statement>From the trendy streets of Shoreditch to the upscale boutiques of Mayfair,
1128
            there's something for everyone in this cosmopolitan city</statement> <class>1</
1129
            class>
        </statements>
1130
```

Figure 12: Segment-Score Example used as part of model in-context training for oracle LLM.