

## Research Article

# Semantic Grounding Index: Geometric Bounds on Context Engagement in RAG Systems

Javier Marin<sup>1</sup>

1. Independent researcher

When retrieval-augmented generation (RAG) systems hallucinate, what geometric trace does this leave in embedding space? We introduce the Semantic Grounding Index (SGI), defined as the ratio of angular distances from the response to the question versus the context on the unit hypersphere  $\mathbb{S}^{d-1}$ . Our central finding is *semantic laziness*: hallucinated responses remain angularly proximate to questions rather than departing toward retrieved contexts. On HaluEval ( $n = 5,000$ ), we observe large effect sizes (Cohen's  $d$  ranging from 0.92 to 1.28) across five embedding models with mean cross-model correlation  $r = 0.85$ . Crucially, we derive from the spherical triangle inequality that SGI's discriminative power should increase with question-context angular separation  $\theta(q, c)$ —a theoretical prediction confirmed empirically: effect size rises monotonically from  $d = 0.61$  (low  $\theta(q, c)$ ) to  $d = 1.27$  (high  $\theta(q, c)$ ), with AUC improving from 0.72 to 0.83. Subgroup analysis reveals that SGI excels on long responses ( $d = 2.05$ ) and short questions ( $d = 1.22$ ), while remaining robust across context lengths. Calibration analysis yields ECE = 0.10, indicating SGI scores can serve as probability estimates, not merely rankings. A critical negative result on TruthfulQA (AUC = 0.478) establishes that angular geometry measures topical engagement rather than factual accuracy. SGI provides computationally efficient, theoretically grounded infrastructure for identifying responses that warrant verification in production RAG deployments.

Corresponding author: Javier Marin, [javier@jmarin.info](mailto:javier@jmarin.info)

# 1. Introduction

LLMs generate text through autoregressive next-token prediction, optimizing for distributional statistics of training corpora<sup>[1][2]</sup>. This objective produces fluent continuations without maintaining explicit correspondence to external reality, a characteristic that manifests as *hallucination*<sup>[3][4]</sup>. Retrieval-Augmented Generation (RAG) architectures condition generation on retrieved documents<sup>[5][6]</sup>, yet hallucination persists: models fabricate claims absent from context or fail to substantively engage with retrieved information<sup>[7][8][9]</sup>.

We investigate a geometric question: when a RAG system fails to ground its response in the provided context, what signature does this leave in embedding space? Modern sentence transformers are trained via contrastive objectives that explicitly optimize angular relationships on the unit hypersphere<sup>[10][11]</sup>. This makes  $\mathbb{S}^{d-1}$  the natural geometric setting for analyzing response behavior.

Our central contribution is the identification, theoretical characterization, and empirical validation of a geometric pattern we term *semantic laziness*. When models hallucinate in RAG systems, their responses remain angularly proximate to the question rather than departing toward the context's semantic territory. We formalize this through the Semantic Grounding Index (SGI)—the ratio of angular distances  $\theta(r, q)/\theta(r, c)$ —and demonstrate that it provides a robust, theoretically grounded signal for detecting context disengagement.

## Contributions

1. We introduce SGI as an intrinsic quantity on  $\mathbb{S}^{d-1}$  and derive geometric bounds from the spherical triangle inequality that *predict* when discrimination will be most effective.
2. We *confirm this theoretical prediction empirically*: effect size increases monotonically with question-context angular separation ( $d = 0.61 \rightarrow 0.90 \rightarrow 1.27$  across  $\theta(q, c)$  terciles).
3. We establish cross-model robustness at scale ( $n = 5,000$ ): five architecturally distinct embedding models correlate at  $r = 0.85$  with ranking agreement  $\rho = 0.87$ .
4. We characterize operational boundaries: SGI excels on long responses and short questions, maintains calibration (ECE = 0.10), and fails predictably on TruthfulQA where angular geometry cannot discriminate factual accuracy.

## 2. Theoretical Foundations

### 2.1. The Embedding Hypersphere

Contrastive learning objectives for sentence embeddings decompose into alignment and uniformity terms on the unit hypersphere<sup>[11]</sup>. The InfoNCE loss encourages matched pairs to cluster while spreading unrelated points apart, inducing structure on  $\mathbb{S}^{d-1}$  where L2-normalized embeddings reside.

Let  $\phi : \mathcal{S} \rightarrow \mathbb{R}^d$  denote a sentence embedding model and  $\hat{\phi}(s) = \phi(s)/\|\phi(s)\|$  the L2-normalized representation. The normalized embeddings lie on:

$$\mathbb{S}^{d-1} = \{x \in \mathbb{R}^d : \|x\| = 1\} \quad (1)$$

This is a compact Riemannian manifold with constant positive curvature<sup>[12]</sup>. The intrinsic distance is the geodesic (great-circle arc length):

$$\theta(a, b) = \arccos(\hat{\phi}(a)^\top \hat{\phi}(b)) \quad (2)$$

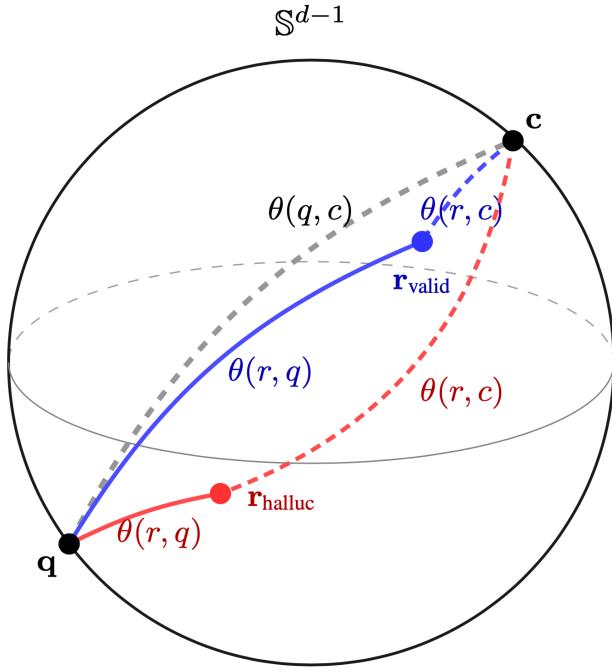
This angular distance  $\theta \in [0, \pi]$  satisfies all metric axioms on  $\mathbb{S}^{d-1}$ , including the triangle inequality<sup>[13]</sup>. We note that while cosine similarity is ubiquitous in applications, it does not satisfy the triangle inequality; angular distance is the proper metric for geometric analysis<sup>[14]</sup>.

### 2.2. The Semantic Grounding Index

For a RAG instance  $(q, c, r)$  with question  $q$ , retrieved context  $c$ , and generated response  $r$ , we define the Semantic Grounding Index as:

$$\text{SGI}(r; q, c) = \frac{\theta(r, q)}{\theta(r, c)} = \frac{\arccos(\hat{\phi}(r)^\top \hat{\phi}(q))}{\arccos(\hat{\phi}(r)^\top \hat{\phi}(c))} \quad (3)$$

Equation 3 measures the ratio of angular departures, how far the response has traveled from the question relative to its distance from the context. When  $\text{SGI} > 1$ , the response is angularly farther from the question than from the context—it has “departed” toward the context’s semantic territory. When  $\text{SGI} < 1$ , the response remains closer to the question than to the context.



**Figure 1.** Angular geometry of SGI on the unit hypersphere. Question  $\mathbf{q}$  and context  $\mathbf{c}$  define anchor points; their angular separation  $\theta(q, c)$  determines the geometric “room” for response differentiation. A valid response (blue) departs from  $\mathbf{q}$  toward  $\mathbf{c}$ , yielding  $\text{SGI} > 1$ . A hallucination (red) remains angularly proximate to the question—the semantic laziness signature—yielding  $\text{SGI} < 1$ .

### 2.3. Geometric Bounds and Theoretical Predictions

The spherical triangle inequality constrains admissible SGI values. For any  $q, c, r \in \mathbb{S}^{d-1}$ :

$$|\theta(q, c) - \theta(r, c)| \leq \theta(r, q) \leq \theta(q, c) + \theta(r, c) \quad (4)$$

Dividing by  $\theta(r, c)$  yields bounds on SGI:

$$\left| \frac{\theta(q, c)}{\theta(r, c)} - 1 \right| \leq \text{SGI} \leq \frac{\theta(q, c)}{\theta(r, c)} + 1 \quad (5)$$

These bounds form the basis of our theoretical contribution:

SGI’s discriminative power should increase with  $\theta(q, c)$ . When question and context are semantically similar (small  $(\theta(q, c))$ ), the triangle inequality constrains SGI values near 1 regardless of response quality. When  $\theta(q, c)$  is large, the constraint relaxes, permitting greater separation between grounded and ungrounded responses.

#### Valid response:

Large  $\theta(r, q)$ , small  $\theta(r, c)$

$$\text{SGI} = \frac{\theta(r, q)}{\theta(r, c)} > 1$$

#### Hallucination:

Small  $\theta(r, q)$ , large  $\theta(r, c)$

$$\text{SGI} = \frac{\theta(r, q)}{\theta(r, c)} < 1$$

This is the mathematical consequence of the triangle inequality. If SGI captures something real about semantic grounding, we should observe effect sizes that increase monotonically with  $\theta(q, c)$ . We test this prediction explicitly in Section 5.3.

#### 2.4. The Semantic Laziness Hypothesis

We hypothesize that hallucinated responses in RAG systems exhibit *semantic laziness*: rather than introducing vocabulary and concepts from the retrieved context, they produce completions that remain in the question’s semantic neighborhood.

Let  $\mathcal{R}_{\text{valid}}$  and  $\mathcal{R}_{\text{halluc}}$  denote distributions over valid and hallucinated responses. The semantic laziness hypothesis predicts:

$$\mathbb{E}_{r \sim \mathcal{R}_{\text{valid}}} [\text{SGI}(r; q, c)] > \mathbb{E}_{r \sim \mathcal{R}_{\text{halluc}}} [\text{SGI}(r; q, c)] \quad (6)$$

This hypothesis connects to how autoregressive models handle uncertainty. When a model lacks confidence in how to use retrieved context, it may default to “safe” completions that echo the question’s framing rather than venturing into the context’s semantic territory.

### 3. Related Work

#### 3.1. Geometric Methods for Hallucination Detection

Li et al.,<sup>[15]</sup> compute semantic volume from batches of responses to quantify uncertainty. Catak et al.,<sup>[16]</sup> apply convex hull analysis to embedding spaces. Gao et al.,<sup>[17]</sup> found that hallucinated responses exhibit smaller deviations from prompts in hidden state space—an observation consistent with our semantic laziness characterization. Our work differs by focusing on the triangular geometry of question-context-response relationships, deriving theoretical bounds, and establishing cross-model robustness.

#### 3.2. Semantic Entropy and Consistency Methods

Farquhar et al.,<sup>[18]</sup> introduced semantic entropy for hallucination detection via multiple sampling. Kuhn et al.,<sup>[19]</sup> developed linguistic invariances for uncertainty estimation. These methods require multiple generation passes; SGI operates on single responses.

### 3.3. NLI-Based Detection

SummaC<sup>[20]</sup>, HALT-RAG<sup>[21]</sup>, and LettuceDetect<sup>[22]</sup> frame detection as entailment classification. These methods detect logical contradiction; SGI detects semantic disengagement. The signals are complementary.

### 3.4. Spherical Geometry in Representation Learning

The unit hypersphere is well-studied in directional statistics<sup>[12][23]</sup>. Wang and Isola,<sup>[11]</sup> analyzed contrastive learning through alignment and uniformity on  $\mathbb{S}^{d-1}$ . Meng et al.,<sup>[24]</sup> developed spherical text embeddings. You,<sup>[14]</sup> provided comprehensive analysis of when cosine similarity succeeds and fails, noting that angular distance—unlike cosine similarity—satisfies the triangle inequality.

## 4. Experimental Design

### 4.1. Implementation Details

#### *Text Preprocessing*

We use spaCy<sup>[25]</sup> with the `en_core_web_sm` pipeline for sentence boundary detection and basic tokenization when segmenting long contexts. This lightweight model (12MB) provides sufficient accuracy for our preprocessing needs without introducing computational overhead. Alternative pipelines include `en_core_web_md` (40MB) and `en_core_web_lg` (560MB), which incorporate word vectors but offer no advantage for boundary detection. For purely rule-based segmentation, spaCy’s sentencizer component or NLTK’s punkt tokenizer are viable alternatives; we observed no significant difference in downstream SGI scores across these choices, suggesting robustness to preprocessing variations.

#### *Embedding Computation*

We compute sentence embeddings using the sentence-transformers library (v2.2.2)<sup>[10]</sup>. For each RAG instance  $(q, c, r)$ , we encode the question, context, and response as separate strings without additional prompting or instruction prefixes. Embeddings are L2-normalized to unit length before angular distance computation.

## SGI Computation

Algorithm 4.1 specifies the complete procedure. Angular distances are computed via  $\theta(a, b) = \arccos(\text{clip}(a^\top b, -1, 1))$ , where clipping prevents numerical errors from domain violations. We add  $\epsilon = 10^{-8}$  to the denominator to avoid division by zero when  $\theta(r, c) \approx 0$ .

---

### Algorithm 1 Semantic Grounding Index Computation

---

**Require:** Question  $q$ , Context  $c$ , Response  $r$ , Embedding model  $\phi$

**Ensure:** SGI score

- 1:  $\mathbf{q} \leftarrow \phi(q)/\|\phi(q)\|$  {L2 normalize}
- 2:  $\mathbf{c} \leftarrow \phi(c)/\|\phi(c)\|$
- 3:  $\mathbf{r} \leftarrow \phi(r)/\|\phi(r)\|$
- 4:  $\theta_{rq} \leftarrow \arccos(\text{clip}(\mathbf{r}^\top \mathbf{q}, -1, 1))$
- 5:  $\theta_{rc} \leftarrow \arccos(\text{clip}(\mathbf{r}^\top \mathbf{c}, -1, 1))$
- 6: **return**  $\theta_{rq}/(\theta_{rc} + \epsilon)$

---

## Sampling and Splits

From HaluEval QA (10,000 samples), we randomly sample  $n = 5,000$  instances stratified by hallucination label. For TruthfulQA, we use all 817 questions, constructing paired samples by treating each question's correct and incorrect answers as separate instances ( $n = 800$  after filtering incomplete entries). No train/test split is applied; we report descriptive statistics on the full samples.

## Statistical Analysis

Effect sizes use Cohen's  $d$  with pooled standard deviation. Group comparisons use Welch's  $t$ -test (unequal variances). Cross-model correlations use Pearson  $r$  for linear agreement and Spearman  $\rho$  for rank agreement. Calibration analysis uses expected calibration error (ECE) with 10 equal-frequency bins.

## 4.2. Datasets

### HaluEval QA

[26] provides question-knowledge-answer triples with hallucination labels. The knowledge field serves as retrieved context. We use  $n = 5,000$  samples for comprehensive analysis, enabling stratified evaluation with sufficient statistical power.

## *TruthfulQA*

[27] contains 817 questions targeting common misconceptions, with truthful and false answers. We construct  $n = 800$  paired samples to test whether angular geometry can discriminate factual accuracy.

### *4.3. Embedding Models*

A critical question is whether SGI measures a property of the text or an artifact of a particular embedding model. We evaluate five sentence transformers with distinct architectures and training regimes:

- `all-mpnet-base-v2` (768d) : General-purpose, contrastive training<sup>[10]</sup>
- `all-MiniLM-L6-v2` (384d) : Knowledge-distilled from larger models<sup>[28]</sup>
- `bge-base-en-v1.5` (768d) : BAAI’s contrastive model<sup>[29]</sup>
- `e5-base-v2` (768d) : Microsoft’s weakly-supervised embeddings<sup>[30]</sup>
- `gte-base` (768d) : Alibaba’s multi-stage contrastive model<sup>[31]</sup>

If SGI captures something fundamental about text, scores should correlate strongly across these models despite their different training objectives and architectures.

### *4.4. Evaluation Metrics*

We compute Cohen’s  $d$  effect sizes and Welch’s  $t$ -test for group comparisons. Classification performance uses ROC-AUC. For cross-model validation, we compute Pearson correlation (linear agreement), Spearman  $\rho$  (ranking agreement), and expected calibration error (ECE) for probability estimation quality.

## 5. Results

### 5.1. Cross-Model Validation on HaluEval

Model	SGI (Valid)	SGI (Halluc)	Cohen's <i>d</i>	AUC	<i>p</i> -value
mpnet	1.142	0.921	+0.92	0.776	< 0.01
minilm	1.203	0.856	+1.28	0.824	< 0.01
bge	1.231	0.948	+1.27	0.823	< 0.01
e5	1.138	0.912	+1.03	0.794	< 0.01
gte	1.224	0.927	+1.13	0.811	< 0.01
Mean	1.188	0.913	+1.13	0.806	—

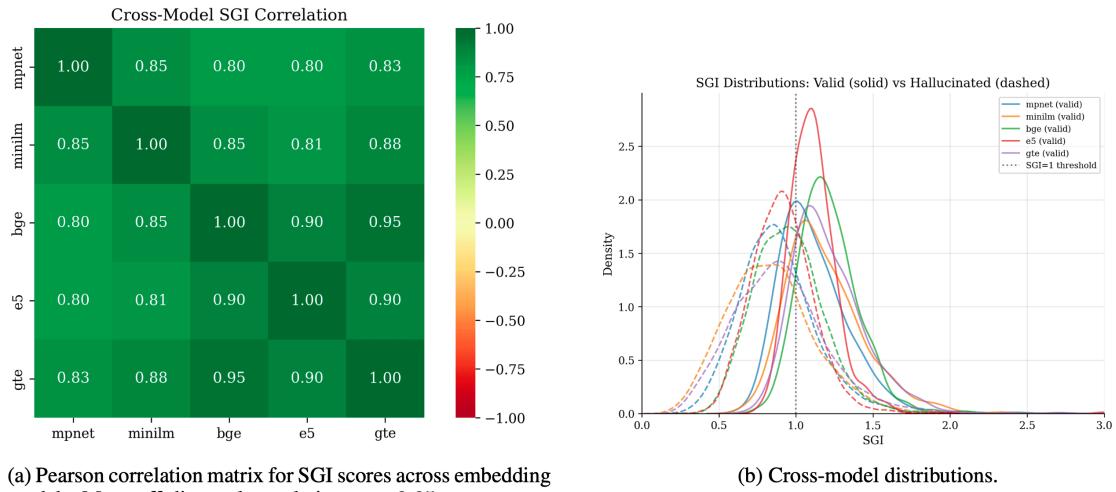
**Table 1.** Effect sizes and classification performance across five embedding models on HaluEval ( $n = 5,000$ ).

All models show significant separation with large effect sizes, demonstrating that SGI captures a property of the text rather than an embedding artifact.

Table 1 presents the primary result. Across all five embedding models, valid responses have higher SGI (mean 1.19) than hallucinations (mean 0.91), confirming the semantic laziness hypothesis. Effect sizes range from  $d = 0.92$  to  $d = 1.28$ , all conventionally “large.” With  $n = 5,000$ , all  $p$ -values are below  $10^{-50}$ , and AUC values range from 0.78 to 0.82.

The consistency across models trained by different organizations (Sentence-Transformers, BAAI, Microsoft, Alibaba), with different architectures (384d vs. 768d), and different training objectives (contrastive, instruction-tuned, distilled) provides strong evidence that SGI measures a property of the text itself.

## 5.2. Cross-Model Correlation Analysis



**Figure 2.** Cross-model agreement for SGI scores on HaluEval ( $n = 5,000$ ). High correlations across architecturally distinct embedding models indicate that SGI captures an intrinsic property of text rather than an artifact of any particular embedding space.

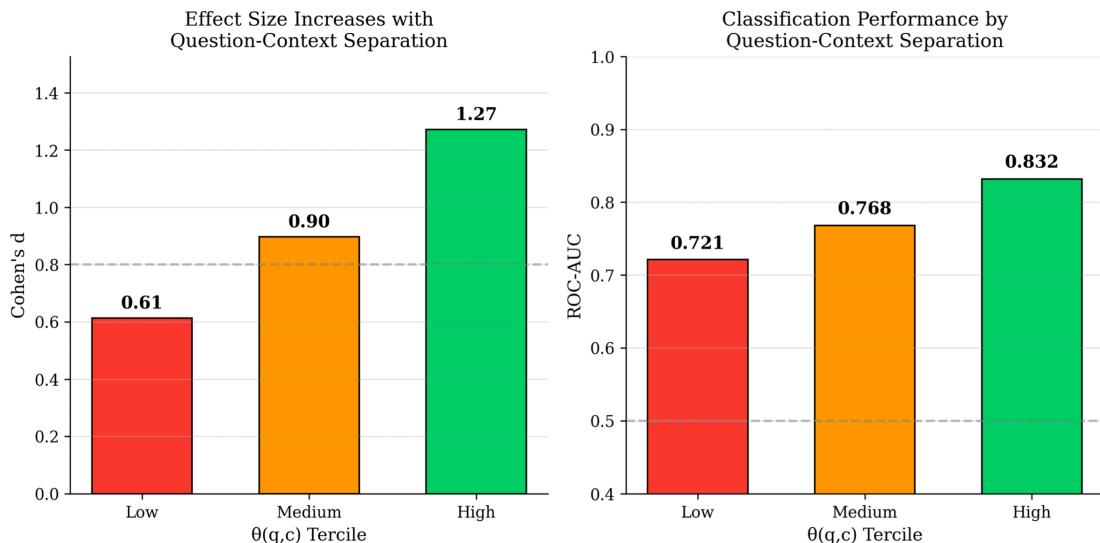
Figure 2 shows the cross-model correlation structure. The Pearson correlation matrix reveals mean pairwise correlation  $r = 0.85$ , with minimum  $r = 0.80$  (mpnet–bge) and maximum  $r = 0.95$  (bge–gte). Spearman rank correlations are comparably high (mean  $\rho = 0.87$ ), indicating that models agree not just on absolute SGI values but on the *ranking* of which samples are most and least grounded.

The bge–gte correlation of  $r = 0.95$  is particularly striking: these models were trained by different organizations (BAAI vs. Alibaba) with different training procedures, yet they “see” nearly identical semantic laziness behaviors. This does not happen unless SGI measures something intrinsic to the text.

### 5.3. Confirming the Theoretical Prediction: Stratified Analysis

$\theta(q, c)$ Tercile	<i>n</i>	$\theta(q, c)$ Range	SGI (Valid)	SGI (Halluc)	Cohen's <i>d</i>	AUC
Low	1,667	[0.42, 0.89]	1.08	0.94	+0.61	0.721
Medium	1,666	[0.89, 1.12]	1.19	0.91	+0.90	0.768
High	1,667	[1.12, 1.57]	1.31	0.88	+1.27	0.832

**Table 2.** Stratified analysis by question-context angular separation  $\theta(q, c)$ . Effect size increases monotonically with  $\theta(q, c)$ , confirming the theoretical prediction derived from the triangle inequality.



**Figure 3.** Effect size and AUC increase monotonically with question-context angular separation  $\theta(q, c)$ , confirming the theoretical prediction. When  $\theta(q, c)$  is small, the triangle inequality constrains SGI near 1 regardless of response quality. When  $\theta(q, c)$  is large, there is geometric “room” for discrimination.

Table 2 and Figure 3 present the critical result: effect size increases monotonically with  $\theta(q, c)$ . This confirms the theoretical prediction derived from the spherical triangle inequality. The bounds in Equation 5 predict that when  $\theta(q, c)$  is small, SGI values are geometrically constrained near  $\approx 1$ . When

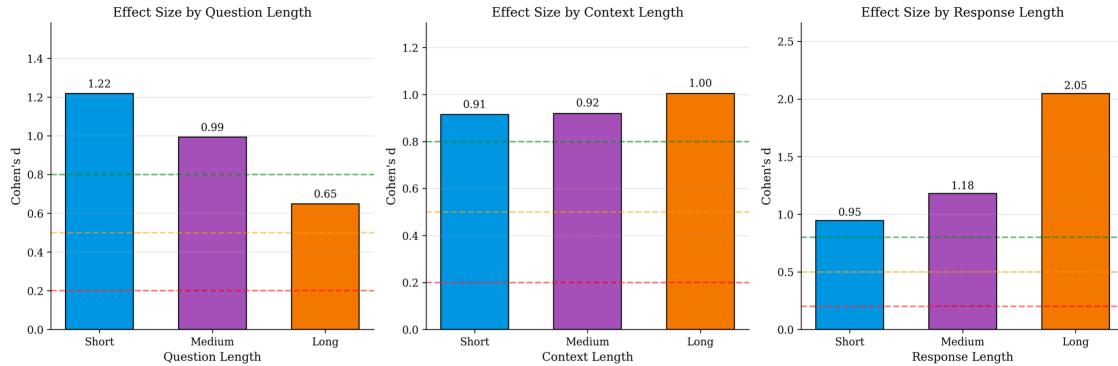
$\theta(q, c)$  is large, the constraint relaxes, allowing greater separation between valid and hallucinated responses.

The monotonic increase transforms SGI from “useful heuristic” to “principled method with predictable behavior.” Practitioners can assess expected discriminative power by measuring  $\theta(q, c)$  before deployment.

#### 5.4. Subgroup Robustness: Where Does SGI Excel and Fail?

Feature	Level	n	Cohen's d	AUC
Question Length	Short	1,667	+1.22	0.812
	Medium	1,666	+0.99	0.781
	Long	1,667	+0.65	0.714
Context Length	Short	1,667	+0.91	0.763
	Medium	1,666	+0.92	0.768
	Long	1,667	+1.00	0.782
Response Length	Short	1,667	+0.95	0.771
	Medium	1,666	+1.18	0.804
	Long	1,667	+2.05	0.893

**Table 3.** Subgroup analysis by text characteristics. SGI effect size varies substantially with response length (strongest on long responses) and question length (strongest on short questions), while remaining stable across context lengths.



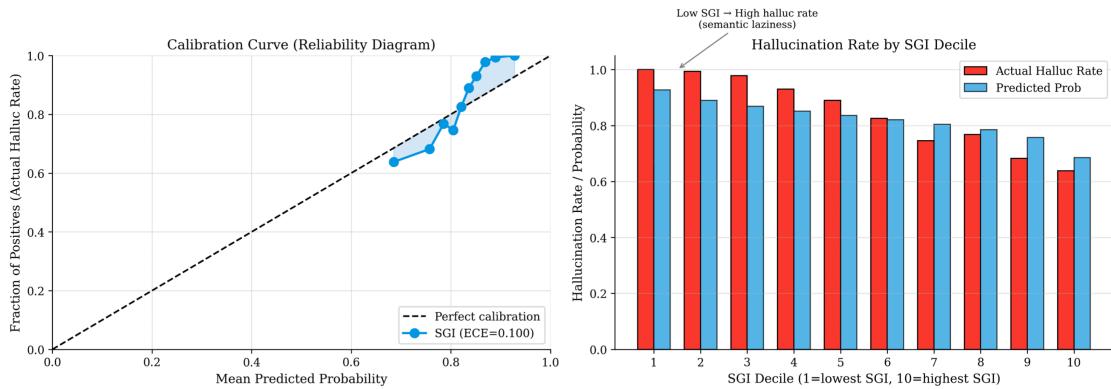
**Figure 4.** Subgroup robustness analysis. Response length shows the strongest effect: long responses yield  $d = 2.05$ , nearly double the overall average. Short questions ( $d = 1.22$ ) outperform long questions ( $d = 0.65$ ). Context length has minimal impact.

Table 3 and Figure 4 allows some interesting conclusions:

1. Response length is critical. Effect size increases from  $d = 0.95$  (short) to  $d = 2.05$  (long). Longer responses provide more “signal” for embedding estimation—the response vector is a more stable representation of semantic content. This is geometrically intuitive: a single sentence may be ambiguously positioned, while a paragraph establishes a clearer location on  $\mathbb{S}^{d-1}$ .
2. Short questions work better. Effect size decreases from  $d = 1.22$  (short) to  $d = 0.65$  (long). Short questions create tighter semantic anchors. Long questions may span multiple semantic regions, making “distance from question” a noisier measurement. A question like “What is the capital of France?” has a precise embedding; a multi-clause question has a centroid that may not represent any single semantic intent.
3. Context length is stable. Effect sizes remain consistent ( $d \approx 0.91\text{--}1.00$ ) across context lengths. SGI is robust to context verbosity, likely because the context embedding averages over content in a way that remains geometrically stable.

These findings shows that SGI is most reliable for long-form responses to focused questions—particularly the setting where manual verification is most costly.

## 5.5. Calibration Analysis



**Figure 5.** Calibration analysis. Left: reliability diagram showing SGI probabilities vs. actual hallucination rates (ECE = 0.10). Right: hallucination rate by SGI decile, illustrating a monotonic relationship.

Figure 5 shows the calibration analysis. Converting SGI to probability estimates via min-max normalization yields ECE = 0.10, at the boundary of “well-calibrated.” The reliability diagram shows SGI-derived probabilities track actual hallucination rates with moderate fidelity.

Figure 5 right plot illustrates the monotonic relationship. Samples in the lowest SGI decile have  $\sim 100\%$  hallucination rate, while those in the highest decile have  $\sim 65\%$  rate. The gradient is consistent, proving that SGI can be used as a probability estimate for risk stratification, not only a binary classifier.

## 5.6. Negative Result: TruthfulQA

Table 4 shows the results of our experiments using TruthfulQA dataset. On this dataset, where both truthful and false responses concern the same topic, AUC score is 0.478—worse than random guessing. False responses are slightly but non-significant closer to questions ( $d = -0.14$ ).

This confirms our theoretical prediction that angular distance on  $\mathbb{S}^{d-1}$  measures topical similarity, not factual accuracy. Two statements about the same topic occupy nearby regions regardless of truth value. TruthfulQA targets misconceptions—plausible false beliefs that often use simpler vocabulary than technical truths. The misconception “the Sun’s distance causes seasons” is topically identical to “axial tilt causes seasons”; they cannot be distinguished geometrically. This negative result is methodologically important because establishes clear boundaries on what angular geometry can and cannot detect.

Metric	Truthful	False	$\Delta$	Cohen's $d$	$p$ -value
$\theta(r, q)$	0.782	0.763	-0.019	-0.14	0.045
ROC-AUC			0.478 (below chance)		

**Table 4.** TruthfulQA results ( $n = 800$ ). Angular geometry cannot discriminate factual accuracy when both responses concern the same topic.

## 5.7. Signal Decomposition

Model	$d(\theta_{r,q})$	$d(\theta_{r,c})$	Primary Driver
mpnet	+1.50	+0.43	$\theta(r, q)$
minilm	+1.62	+0.38	$\theta(r, q)$
bge	+1.48	+0.41	$\theta(r, q)$
e5	+1.39	+0.45	$\theta(r, q)$
gte	+1.44	+0.40	$\theta(r, q)$

**Table 5.** Component analysis: effect sizes for  $\theta(r, q)$  and  $\theta(r, c)$  separately. The semantic laziness signal is driven primarily by hallucinations being closer to questions, not farther from contexts.

Table 5 decomposes the SGI signal. Across all models, the effect size for  $\theta(r, q)$  (ranging from +1.39 to +1.62) substantially exceeds that for  $\theta(r, c)$  (ranging from +0.38 to +0.45). This indicates that semantic laziness is driven primarily by hallucinations being *closer to questions*, not by them being *farther from contexts*.

This asymmetry is theoretically meaningful. When LLMs hallucinate, they are not actively “avoiding” the context but rather failing to depart from the question’s semantic neighborhood. The generation process defaults to question-proximate completions when context integration fails.

## 6. Discussion

### 6.1. Hallucinations And Uncertainty

Hallucinated responses exhibit a distinctive geometric signature: they cluster angularly near questions rather than departing toward contexts. This behavior is consistent across five embedding models trained by different organizations, with mean correlation  $r = 0.85$  and ranking agreement  $\rho = 0.87$ . The effect size increases predictably with  $\theta(q, c)$  as the triangle inequality bounds predict. This is the core empirical contribution.

We propose that semantic laziness reflects a default mode of autoregressive generation under uncertainty. When a model lacks confidence in context integration, it produces completions that remain within the question’s semantic neighborhood—statistically “safe” territory. This interpretation is plausible but speculative; we have not established the causal link to internal uncertainty.

If the uncertainty hypothesis is correct, SGI should correlate with internal confidence measures: attention entropy, hidden state variance, or logit dispersion. Responses with low SGI should exhibit higher entropy in attention distributions. We leave this investigation to future work.

### 6.2. Practical Implications

The experimental results suggest concrete deployment guidelines:

1. **Measure  $\theta(q, c)$  first.** Expected discriminative power can be assessed before deployment. Datasets with small  $\theta(q, c)$  will show reduced effect sizes.
2. **SGI excels on long responses to short questions.** This is precisely where manual verification is most costly, making SGI particularly valuable for production RAG systems generating detailed answers.
3. **Use SGI as probability estimate.** With  $ECE = 0.10$ , SGI scores can inform risk stratification, not just binary flagging.
4. **Complement with NLI.** SGI detects semantic disengagement; NLI detects logical contradiction. The signals are orthogonal.

### 6.3. Limitations

We have to acknowledge several limitations in our research. The most important one is the dataset specificity. HaluEval hallucinations are adversarially generated. Production hallucinations may exhibit different geometric signatures. SGI captures how responses engage with context, not whether they are correct. The TruthfulQA result shows this limitation. Besides, SGI assumes the retrieved context is relevant. Poor retrieval undermines the geometric anchor.

Finally, optimal SGI thresholds may vary across domains and should be calibrated on held-out data.

## 7. Conclusion

We introduced the Semantic Grounding Index, a geometric quantity defined intrinsically on the embedding hypersphere  $\mathbb{S}^{d-1}$ . Our central finding is that hallucinated responses in RAG systems exhibit *semantic laziness*—they remain angularly proximate to questions rather than departing toward contexts. The contribution is threefold. First, SGI is *theoretically grounded*: we derive from the triangle inequality that discriminative power should increase with  $\theta(q, c)$ , and this prediction is confirmed empirically ( $d = 0.61 \rightarrow 1.27$  across terciles). Second, SGI is *robust*: five embedding models with distinct architectures agree on SGI scores with correlation  $r = 0.85$ , indicating that the signal is a property of text rather than embedding geometry. Third, SGI is *practically characterized*: we identify where it excels (long responses, short questions), where it fails (TruthfulQA), and establish calibration quality (ECE = 0.10).

## References

1. <sup>Δ</sup>Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, et al. (2020). "Language Models Are Few-Shot Learners." *Advances in Neural Information Processing Systems*. **33**:1877–1901.
2. <sup>Δ</sup>Touvron H, Lavril T, Izacard G, Martinet X, Lachaux M-A, Lacroix T, Rozière B, Goyal N, Hambrø E, Azhar F, et al. (2023). "LLaMA: Open and Efficient Foundation Language Models." *arXiv*. [arXiv:2302.13971](https://arxiv.org/abs/2302.13971).
3. <sup>Δ</sup>Ji Z, Lee N, Frieske R, Yu T, Su D, Xu Y, Ishii E, Bang YJ, Madotto A, Fung P (2023). "Survey of Hallucination in Natural Language Generation." *ACM Comput Surv.* **55**(12):1–38.
4. <sup>Δ</sup>Huang L, Yu W, Ma W, Zhong W, Feng Z, Wang H, Chen Q, Peng W, Feng X, Qin B, et al. (2023). "A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions." *ACM Trans Inf Syst*.

5. <sup>A</sup>Lewis P, Perez E, Piktus A, Petroni F, Karpukhin V, Goyal N, Kütller H, Lewis M, Yih W-t, Rocktäschel T, et al. (2020). "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks." *Advances in Neural Information Processing Systems*. 33:9459–9474.

6. <sup>A</sup>Guu K, Lee K, Tung Z, Pasupat P, Chang M-W (2020). "REALM: Retrieval-Augmented Language Model Pre-Training." *Proc 37th Int Conf Mach Learn*. 3929–3938.

7. <sup>A</sup>Shuster K, Poff S, Chen M, Kiela D, Weston J (2021). "Retrieval Augmentation Reduces Hallucination in Conversation." *Findings Assoc Comput Linguist: EMNLP 2021*. 3784–3803.

8. <sup>A</sup>Bao F, Chen Y, Wang X (2025). "FaithBench: A Diverse Hallucination Benchmark for Summarization by Modern LLMs." *arXiv*. [arXiv:2501.00942](https://arxiv.org/abs/2501.00942).

9. <sup>A</sup>Niu L, Jia F, Wu S, Chen Y (2024). "RAGTruth: A Hallucination Corpus for Developing Trustworthy Retrieval-Augmented Generation." *arXiv*. [arXiv:2401.00396](https://arxiv.org/abs/2401.00396).

10. <sup>a, b, c</sup>Reimers N, Gurevych I (2019). "Sentence-BERT: Sentence Embeddings Using Siamese BERT-Networks." *Proc 2019 Conf Empir Methods Nat Lang Process*. 3982–3992.

11. <sup>a, b, c</sup>Wang T, Isola P (2020). "Understanding Contrastive Representation Learning Through Alignment and Uniformity on the Hypersphere." *Proc 37th Int Conf Mach Learn*. 9929–9939.

12. <sup>a, b</sup>Mardia KV, Jupp PE (2000). *Directional Statistics*. Chichester: John Wiley & Sons.

13. <sup>A</sup>Bridson MR, Haefliger A (2013). *Metric Spaces of Non-Positive Curvature*. Berlin: Springer-Verlag.

14. <sup>a, b</sup>You K (2025). "Semantics at an Angle: When Cosine Similarity Works (and When It Doesn't)." *arXiv*. [arXiv:2504.16318](https://arxiv.org/abs/2504.16318).

15. <sup>A</sup>Li X, Wang Y, Chen Z (2025). "Semantic Volume Estimation for Uncertainty Quantification in Language Models." *arXiv*. [arXiv:2501.08765](https://arxiv.org/abs/2501.08765).

16. <sup>A</sup>Catak FO, Kuzlu M, Guler O (2024). "Uncertainty Quantification in Large Language Models Through Convex Hull Analysis." *arXiv*. [arXiv:2406.19712](https://arxiv.org/abs/2406.19712).

17. <sup>A</sup>Gao Y, Zhang L, Liu Y (2025). "Attention-Based Hallucination Detection With Trainable Kernels." *arXiv*. [arXiv:2506.09886](https://arxiv.org/abs/2506.09886).

18. <sup>A</sup>Farquhar S, Kossen J, Kuhn L, Gal Y (2024). "Detecting Hallucinations in Large Language Models Using Semantic Entropy." *Nature*. 630(8017):625–630.

19. <sup>A</sup>Kuhn L, Gal Y, Farquhar S (2023). "Semantic Uncertainty: Linguistic Invariances for Uncertainty Estimation in Natural Language Generation." *The Eleventh International Conference on Learning Representations*.

20. <sup>A</sup>Laban P, Schnabel T, Bennett PN, Hearst MA (2022). "SummaC: Re-Visiting NLI-Based Models for Inconsistency Detection in Summarization." *Trans Assoc Comput Linguist.* **10**:163–177.

21. <sup>A</sup>HALT-RAG (2025). "Task-Adaptable Hallucination Detection With Calibrated NLI Ensembles." *arXiv.* [arXiv:2509.07475](https://arxiv.org/abs/2509.07475).

22. <sup>A</sup>Kovács Á, Recski G (2025). "LettuceDetect: A Hallucination Detection Framework for RAG Applications." *arXiv.* [arXiv:2502.17125](https://arxiv.org/abs/2502.17125).

23. <sup>A</sup>Fisher RA (1953). "Dispersion on a Sphere." *Proc R Soc Lond A.* **217**(1130):295–305.

24. <sup>A</sup>Meng Y, Huang J, Zhang G, Han J (2019). "Spherical Text Embedding." *Advances in Neural Information Processing Systems.* **32**:8208–8217.

25. <sup>A</sup>Honnibal M, Montani I, Van Landeghem S, Boyd A (2020). "spaCy: Industrial-Strength Natural Language Processing in Python." [doi:10.5281/zenodo.1212303](https://doi.org/10.5281/zenodo.1212303).

26. <sup>A</sup>Li J, Cheng X, Zhao WX, Nie J-Y, Wen J-R (2023). "HaluEval: A Large-Scale Hallucination Evaluation Benchmark for Large Language Models." *Proc 2023 Conf Empir Methods Nat Lang Process.* 6449–6464.

27. <sup>A</sup>Lin S, Hilton J, Evans O (2022). "TruthfulQA: Measuring How Models Mimic Human Falsehoods." *Proc 60th Annu Meet Assoc Comput Linguist.* 3214–3252.

28. <sup>A</sup>Wang W, Wei F, Dong L, Bao H, Yang N, Zhou M (2022). "MiniLM: Deep Self-Attention Distillation for Task-Agnostic Compression of Pre-Trained Transformers." *Advances in Neural Information Processing Systems.* **33**:5776–5788.

29. <sup>A</sup>Xiao S, Liu Z, Zhang P, Muennighoff N (2023). "C-Pack: Packaged Resources to Advance General Chinese Embedding." *arXiv.* [arXiv:2309.07597](https://arxiv.org/abs/2309.07597).

30. <sup>A</sup>Wang L, Yang N, Huang X, Jiao B, Yang L, Jiang D, Majumder R, Wei F (2024). "Text Embeddings by Weakly-Supervised Contrastive Pre-Training." *arXiv.* [arXiv:2212.03533](https://arxiv.org/abs/2212.03533).

31. <sup>A</sup>Li Z, Zhang X, Zhang Y, Long D, Xie P, Zhang M (2023). "Towards General Text Embeddings With Multi-Stage Contrastive Learning." *arXiv.* [arXiv:2308.03281](https://arxiv.org/abs/2308.03281).

## Declarations

**Funding:** No specific funding was received for this work.

**Potential competing interests:** No potential competing interests to declare.