

Research Article

A Geometric Taxonomy of Hallucinations in LLMs

Javier Marin¹

1. Independent researcher

The term “hallucination” converge different failure modes with specific geometric signatures in embedding space. We propose a taxonomy identifying three types: unfaithfulness (Type I: ignoring provided context), confabulation (Type II: inventing semantically foreign content), and factual error (Type III: wrong details within correct conceptual frames). We introduce two detection methods grounded in this taxonomy: the Semantic Grounding Index (SGI) for Type I, which measures whether a response moves toward provided context on the unit hypersphere, and the Directional Grounding Index (Γ) for Type II, which measures displacement geometry in context-free settings. Γ achieves AUROC 0.958 ± 0.034 on human-crafted confabulations with 3.8% cross-domain degradation. External validation on three independently collected human-annotated benchmarks—WikiBio GPT-3 [1], FELM [2], and ExpertQA [3]—yields domain-specific AUROC 0.581–0.695, with Γ outperforming an NLI CrossEncoder baseline by $\Delta = 0.243$ on expert-domain data where surface entailment operates at chance. On LLM-generated benchmarks, detection is domain-local. We examine the Type III boundary through TruthfulQA [4], where apparent classifier signal (LR AUROC 0.731) is traced to a stylistic annotation confound: false answers are geometrically closer to queries than truthful ones, a pattern incompatible with factual-error detection. This identifies a theoretical constraint from a methodological limitation.

Corresponding author: Javier Marín, javier@jmarin.info

1. Introduction

LLMs generate text by optimizing distributional statistics without maintaining correspondence to external reality [5][6]. This produces hallucination—fluent, confident text that is factually wrong [7][8]—posing challenges for deployment in high-stakes contexts where accuracy must be documented [9].

“Hallucination” refers to at least three distinct failure modes. A model that ignores provided documents ^[10] differs fundamentally from one that invents non-existent entities, which differs from one that provides wrong details about real concepts ^[11]. These pathologies have different roots, different consequences, and different geometric signatures in embedding space.

The problem is compounded by how hallucination benchmarks are constructed. Most evaluation datasets generate false answers by prompting LLMs to fabricate ^[12], which captures how models write when asked to lie rather than the geometry of genuine factual error. ^[13] have questioned whether “hallucination” is even the right framing, proposing instead “bullshit” in Frankfurt’s sense ^[14]—speech produced without concern for truth. Rather than adjudicating this terminological debate, we operationalize distinctions geometrically. If different failure modes produce different signatures on \mathbb{S}^{d-1} , the taxonomy becomes empirically grounded.

We introduce SGI (Semantic Grounding Index) for context-based detection of Type I, and Γ (Directional Grounding Index) for context-free detection of Type II. Through systematic experiments including external validation on three independent human-annotated benchmarks and a detailed investigation of TruthfulQA, we characterize where geometric detection succeeds, where it fails, and why.

Type I (Unfaithfulness): The model ignores provided context, generating from parametric memory ^[10]. The response remains angularly close to the query rather than moving toward the context. We call this pathological behaviour “semantic laziness” and is detectable via SGI.

Type II (Confabulation): The model invents entities, mechanisms, or concepts that do not exist. The response displacement departs from the plausible-answer manifold in geometrically detectable directions. Detectable via Γ .

Type III (Factual Error): The model provides wrong details within the correct conceptual frame. The response is semantically plausible, occupying a region of embedding space that angular geometry cannot distinguish from correct responses. Distributional representations encode co-occurrence, not truth conditions ^{[15][16]}.

2. Related Work

The term “hallucination” entered NLP from neural machine translation, where it described outputs unrelated to source sentences ^{[17][18]}. Academic definitions groups around two poles: ^[7] emphasize

unfaithfulness to source; [8] emphasize plausible-yet-nonfactual content. [10] distinguish intrinsic hallucinations (contradicting source) from extrinsic ones (unverifiable from source).

Methods requiring model internals achieve strong performance: hidden-state trajectories identify truthful from hallucinated generations [19] (AUROC 0.96); eigenvalue spectra of response covariance provide complementary signals [20]. These require white-box access unavailable in API deployments. Consistency-based methods sample repeatedly: SelfCheckGPT [1] achieves AUC-PR 0.93; semantic entropy [21][22] achieves AUROC 0.79, both requiring 5–20 generations per query. When source documents are available, alignment verification works well: HHEM [] achieves 74–77% balanced accuracy; AlignScore [23] reaches AUC-ROC 0.87. Our methods require neither model internals, nor multiple generations, nor source documents—using only the geometry of a single embedding call.

Contrastive training explicitly optimizes angular relationships on \mathbb{S}^{d-1} [24][25]. [26] analyze alignment and uniformity as complementary objectives of contrastive learning on the hypersphere. Our approach exploits the geometric structure induced by this training.

3. Experimental Setup

3.1. Detection methods

Let $\hat{\phi}(s) = \phi(s)/\|\phi(s)\|$ denote the L_2 -normalized sentence representation. Angular distance $\theta(a, b) = \arccos(\hat{\phi}(a)^\top \hat{\phi}(b))$ is the geodesic distance on \mathbb{S}^{d-1} [27][28] and satisfies the triangle inequality, which constrains how the three pairwise distances $\theta(q, c)$, $\theta(q, r)$, $\theta(r, c)$ can simultaneously take values.

3.1.1. Semantic Grounding Index (SGI)

For query q , context c , and response r :

$$SGI(r; q, c) = \frac{\theta(r, q)}{\theta(r, c)} \quad (1)$$

A grounded response moves toward the context’s semantic territory on \mathbb{S}^{d-1} , satisfying $SGI > 1$ (closer to context than to query). A Type I unfaithful response remains near the query, satisfying $SGI \leq 1$. The ratio compares two geodesic distances on the same sphere; it is not bounded by $[0, 1]$. Angular distance is used in preference to cosine similarity because the geodesic metric satisfies the triangle inequality [28], making the ratio geometrically principled.

3.1.2. Directional Grounding Index

For query q and response r , the normalized displacement is $\hat{\delta}(q, r) = \delta / \|\delta\|$ where $\delta = \hat{\phi}(r) - \hat{\phi}(q)$. Given a reference set $\mathcal{R} = \{(q_i, r_i)\}_{i=1}^N$ of verified grounded pairs, $\hat{\mu} = \mu / \|\mu\|$ where $\mu = \frac{1}{N} \sum_{i=1}^N \hat{\delta}(q_i, r_i)$. Normalizing each $\hat{\delta}$ before averaging computes the Von Mises–Fisher mean direction on the sphere, ensuring $\hat{\mu}$ reflects grounding direction rather than magnitude. The two normalizations serve distinct geometric purposes and cannot be collapsed. We define Directional Grounding Index as:

$$\Gamma(q, r; \mathcal{R}) = \hat{\delta}(q, r)^\top \hat{\mu} \quad (2)$$

where $\Gamma \in [-1, +1]$. High values indicate alignment with the learned grounding direction; low or negative values indicate anomalous displacement characteristic of Type II confabulation. A local variant uses k -nearest neighbors in \mathcal{R} to compute a query-adaptive $\hat{\mu}_q$. Global Γ requires one embedding and one dot product: $O(d)$.

3.2. Baseline and Statistics

The NLI baseline (`cross-encoder/nli-deberta-v3-small`) takes (q, r) pairs and outputs entailment probability. It applies to Type II and Type III experiments only: these are context-free settings where both methods receive the same inputs. SGI operates on (q, c, r) triples and has no directly comparable surface baseline, as no standard entailment model accepts three inputs in this configuration.

We report AUROC with 1,000-resample bootstrap CIs, permutation test p -values (5,000 resamples), and Hanley–McNeil z -tests [29]. All experiments use sentence-t5-large (768 dimensions); architecture ablation appears in Section 4.2.4.

4. Results

4.1. Type I: SGI on Grounded Context

Table 1 reports SGI on HaluEval QA ($n = 10,000$) [12] across five embedding architectures. Grounded responses satisfy $\text{SGI} > 1$; Type I unfaithful responses satisfy $\text{SGI} \leq 1$. Cross-model correlation ($r = 0.85$) confirms SGI measures a property of the text, not an embedding artifact.

Model	SGI grnd.	SGI halluc.	AUROC	TPR@5%FPR
sentence-t5-large	1.203	0.856	0.824	0.41
all-mpnet-base-v2	1.142	0.921	0.776	0.33
bge-base-en-v1.5	1.231	0.948	0.823	0.40
e5-base-v2	1.138	0.912	0.794	0.37
all-MiniLM-L6-v2	1.188	0.913	0.806	0.38
Mean	1.180	0.910	0.805	0.38

Table 1. SGI on HaluEval QA. SGI > 1 for grounded responses; ≤ 1 for Type I hallucinated. AUROC 0.776–0.824 consistent across architectures.

Mean SGI for grounded responses (1.180) exceeds 1 across all models, confirming that grounded responses move toward context on \mathbb{S}^{d-1} . Unfaithful responses (0.910) remain near the query.

4.2. Type II: DGI on Confabulations and External Benchmarks

4.2.1. Human-crafted confabulations.

Table 2 reports Γ on 142 human-written confabulations (finance, medical, legal) prepared by the authors to represent prototypical Type II errors: redefined terms, fabricated institutions, or invented mechanisms. These introduce semantically foreign content, unlike within-frame substitutions.

Method	AUROC	95% CI	TPR@10%FPR
Γ global $k = 15$	0.958	[0.924, 0.992]	0.84
Γ local ($n = 142$)	0.965	[0.931, 0.999]	0.87
NLI CrossEncoder	0.611	[0.527, 0.695]	0.19

Table 2. Γ on human-crafted confabulations ($\hat{\mu}$) vs. NLI baseline. Cross-domain degradation: 3.8%. Local vs. global improvement: 0.7% (n.s.).

Γ outperforms NLI by $\Delta = 0.347$. NLI achieves only 0.611 because confabulations are syntactically coherent and entailment-compatible [23]—surface entailment cannot detect displacement off the plausibility manifold.

4.2.2. Domain-locality on LLM-generated benchmarks.

On HaluEval (QA and Dialogue) [12], in-domain AUROC is 0.68–0.99; cross-domain AUROC collapses to ≈ 0.50 (Dialogue \rightarrow QA: 0.38). Mean cosine similarity between domain grounding directions $\hat{\mu} = -0.07$, confirming near-orthogonality. Human-crafted confabulations show only 3.8% degradation because realistic semantic departure from the plausibility manifold is domain-agnostic; prompted fabrication [12] is not.

4.2.3. External human-annotated benchmarks.

Table 3 reports validation on three independently collected datasets not used in method development. WikiBio GPT-3 [1] provides 238 GPT-3 biographies with sentence-level human annotation (102 matched pairs). FELM [2] provides segment-level factuality labels across world-knowledge and writing domains (81 pairs). ExpertQA [3] provides claim-level expert annotations across 32 specialist fields (900 pairs).

Dataset	-0.131 dom.	95% CI	$n = 81$	NLI	95% CI	+0.024
WikiBio ($n = 900$)	0.581	[0.416, 0.695]	.199	0.712	[0.639, 0.781]	+0.243
FELM (Γ)	0.648	[0.510, 0.801]	.013	0.624	[0.532, 0.704]	p
ExpertQA (Δ)	0.695	[0.654, 0.738]	.000	0.452	[0.426, 0.478]	Γ_{dom} - NLI

Table 3. External benchmark validation: domain-specific Γ vs. NLI. p : permutation test. Δ : $n = 102$.

ExpertQA experiment provides a relevant result: Γ outperforms NLI by $\Delta = 0.243$ on expert-annotated inaccurate claims across 32 specialist fields. NLI is at chance (0.452, $p = 1.0$): expert errors are entailment-compatible but occupy a different region of embedding space. FELM shows modest domain-specific detection (0.648, $p = .013$). WikiBio domain-specific detection fail ($p = .199$): its annotation criterion marks any incorrect detail as “major inaccurate” regardless of semantic distance, conflating Type II and Type III errors. This contrast with ExpertQA shows the relevant boundary is Type II/III, not human-annotated vs. LLM-generated.

Table 4 reports the cross-domain transfer matrix. The FELM \leftrightarrow ExpertQA symmetry (0.822) confirms that open-ended factual confabulation shares a consistent geometric regime across subject domains. WikiBio’s $\hat{\mu}$ actively misleads on both other datasets, identifying it as geometrically incommensurable with the confabulation regime.

	Test domain		
Train ↔	WikiBio	FELM	ExpertQA
WikiBio	0.537	0.200	0.288
FELM	0.433	0.822	0.653
ExpertQA	0.446	0.822	0.691

Table 4. Cross-domain transfer matrix (AUROC). FELM↔ExpertQA symmetry (0.822) confirms a shared geometric regime for factual confabulation.

4.2.4. Architecture ablation

Higher-dimensional embedding spaces provide finer geometric resolution for displacement measurement (Table 5).

Model	Dim.	AUROC	95% CI
all-MiniLM-L6-v2	384	0.800	[0.730, 0.870]
all-mpnet-base-v2	768	0.880	[0.820, 0.940]
sentence-t5-large	768	0.958	[0.924, 0.992]

Table 5. Architecture ablation on human-crafted confabulations.

4.3. Type III Boundary: TruthfulQA

TruthfulQA ^[4] provides 817 questions with human-written truthful (best_answer) and false (incorrect_answers) response pairs. False answers reflect beliefs people actually hold, making them topically identical to truthful answers. Table 6 reports all detection methods.

Method	AUROC	95% CI	p_{HM}	$< .001$	Signal
Cosine similarity	0.365	[0.340, 0.392]	1.000	Γ	<i>inverted</i>
Embedding magnitude	0.518	[0.497, 0.541]	.059	.198	n.s.
Γ (80/20 split)	0.535	[0.473, 0.600]	.137	.265	n.s.
$\hat{\mu}$ (global $< .001$)	0.579	[0.551, 0.607]	.000	$< .001$	style
NLI entailment	0.539	[0.511, 0.567]	.004	.006	style
LR raw embeddings (CV)	0.731	[0.706, 0.756]	.000	Γ	style artifact
LR on $< .001$ features	0.786	[0.765, 0.807]	.000	z	style artifact

Table 6. Detection on TruthfulQA (817 matched pairs). Bootstrap 95% CI. HM: Hanley–McNeil p_{perm} -test ^[29].
The cosine inversion (AUROC 0.365) is the diagnostic finding.

The LR achieves AUROC=0.731, which appears to contradict the Type III hypothesis. The cosine inversion resolves this: false answers are geometrically closer to queries than truthful ones (AUROC=0.365, below chance)—the opposite of confabulation geometry. The signal is annotation style: truthful answers are longer and hedged; false answers are shorter and declarative. Longer responses accumulate semantic content orthogonal to the query, producing larger displacement vectors that the LR exploits across 768 dimensions. Held-out Γ calibration, which cannot propagate this style signal through $\hat{\mu}$, yields AUROC=0.535 ($p=0.265$). The LR detects how TruthfulQA was written, not whether answers are factually correct.

Three observations confirm this interpretation: (1) the cosine signal is inverted relative to what factual-error detection would predict; (2) held-out Γ calibration achieves AUROC 0.535 ($p = .265$, CI crosses 0.50); (3) per-category Γ averages 0.823 but inflates with small-sample overfitting—categories with $n > 50$ converge toward 0.57–0.65, consistent with a modest style signal rather than factual detection.

5. Discussion

The experiments validate the core claim of the taxonomy: the three types of hallucination produce distinct geometric signatures, and those signatures determine what is detectable and what is not. Type I unfaithfulness produces a measurable failure to move toward context on \mathbb{S}^{d-1} , captured by SGI. Type II confabulation produces anomalous displacement from the plausible-answer manifold, captured by Γ . Type III factual error produces no geometric anomaly, because a wrong detail within the correct conceptual frame occupies the same region of embedding space as the correct detail. The taxonomy does not merely name these differences—it predicts them from the geometry of how embedding models represent meaning.

The Type III boundary deserves particular attention. The TruthfulQA analysis does not establish only that Type III is difficult to detect but that the apparent detection signal (LR AUROC 0.731) is geometrically inconsistent with factual-error detection. False answers are closer to queries than truthful ones, the opposite of what any displacement-based method would produce under genuine error detection. The held-out Γ result (AUROC=0.535, $p = .265$) confirms that once annotation style is removed from the signal, geometry offers nothing—reflecting the distributional hypothesis [15][16]: embedding geometry encodes co-occurrence, not truth conditions. Type III errors are geometrically invisible by construction.

The contrast between WikiBio and ExpertQA validates the Type II/III distinction empirically. The difference is annotation criterion: whether incorrect content represents semantic departure from the plausibility manifold (Type II) or within-frame substitution (Type III).

6. Limitations

The human-crafted confabulation dataset ($n = 142$) covers prototypical Type II errors with clear semantic departure. Mixed confabulations—responses that partially engage with context while introducing fabricated elements—may produce intermediate Γ scores whose interpretation is not yet characterized.

The Type II/III boundary is defined geometrically but not yet formally proven. The hypothesis that Type III errors are undetectable rests on the distributional hypothesis [15], [16] and the TruthfulQA evidence; a formal proof on \mathbb{S}^{d-1} for within-frame substitutions remains open.

The taxonomy covers outputs of generative models. It does not address retrieval errors, prompt injection, or failures that occur before generation.

7. Conclusions

We introduce a geometric taxonomy of LLM hallucination identifying three failure modes by their signatures on \mathbb{S}^{d-1} . Moreover, we introduce two detection methods grounded in this taxonomy. SGI detects Type I unfaithfulness by measuring whether a response moves toward provided context. Γ detects Type II confabulation via displacement geometry, achieving AUROC 0.958 ± 0.034 on human-crafted confabulations and $0.581\text{--}0.695$ on external human-annotated benchmarks [1][2][3], outperforming NLI by $\Delta = 0.243$ on expert-domain data where surface entailment operates at chance. Type III factual errors resist geometric detection; apparent classifier success on TruthfulQA [4] (LR 0.731) is traced to a stylistic annotation confound—false answers are geometrically closer to queries than truthful ones (cosine 0.365, inverted)—confirming a theoretical constraint rather than a methodological gap.

The domain-locality finding ($\overline{\cos}(\hat{\mu}) = -0.07$ on LLM benchmarks) and FELM \leftrightarrow ExpertQA transfer symmetry (0.822) characterize when geometric detection generalizes: within genuine confabulation regimes, a consistent grounding direction exists and transfers; within LLM-generation-artifact regimes, it does not. The taxonomy predicts which side of this boundary any dataset occupies.

Reproducibility

Code and datasets released upon publication. All embedding and NLI models publicly available on HuggingFace [24]. Benchmarks HaluEval [12], TruthfulQA [4], WikiBio GPT-3 [1], FELM [2], and ExpertQA [3] are publicly accessible. Human-crafted confabulation dataset released with code.

References

1. ^{a, b, c, d, e}Manakul P, Liusie A, Gales M (2023). "SelfCheckGPT: Zero-Resource Black-Box Hallucination Detection for Generative Large Language Models." *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. pp. 9004–9017.
2. ^{a, b, c, d}Chen S, Zhao Y, Zhang J, Chern I, Gao S, Liu P, He J (2023). "FELM: Benchmarking Factuality Evaluation of Large Language Models." *Advances in Neural Information Processing Systems*.
3. ^{a, b, c, d}Malaviya C, Lee S, Chen S, Sieber E, Yatskar M, Roth D, Clark C (2024). "ExpertQA: Expert-Curated Questions and Attributed Answers." *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics*.

4. ^a, ^b, ^c, ^dLin S, Hilton J, Evans O (2022). "TruthfulQA: Measuring How Models Mimic Human Falsehoods." *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*. pp. 3214–3252.
5. [△]Brown T, Mann B, Ryder N, et al. (2020). "Language Models Are Few-Shot Learners." *Adv Neural Inf Process Syst*. **33**:1877–1901.
6. [△]Achiam J, Adler S, Agarwal S, et al. (2023). "GPT-4 Technical Report." *arXiv preprint arXiv:2303.08774*.
7. ^a, ^bJi Z, Lee N, Frieske R, et al. (2023). "Survey of Hallucination in Natural Language Generation." *ACM Comput Surv*. **55**(12):1–38.
8. ^a, ^bHuang L, Yu W, Ma W, et al. (2023). "A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions." *arXiv preprint arXiv:2311.05232*.
9. [△]European Parliament and Council of the European Union (2024). "Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence." *Official Journal of the European Union*.
10. ^a, ^b, ^cMaynez J, Narayan S, Bohnet B, McDonald R (2020). "On Faithfulness and Factuality in Abstractive Summarization." *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. pp. 1906–1919.
11. [△]Zhang Y, Li Y, Cui L, et al. (2023). "Siren's Song in the AI Ocean: A Survey on Hallucination in Large Language Models." *arXiv preprint arXiv:2309.01219*.
12. ^a, ^b, ^c, ^d, ^eLi J, Cheng X, Zhao WX, Nie J-Y, Wen J-R (2023). "HaluEval: A Large-Scale Hallucination Evaluation Benchmark for Large Language Models." *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. pp. 6449–6464.
13. [△]Hicks MT, Humphries J, Slater J (2024). "ChatGPT Is Bullshit." *Ethics Inf Technol*. **26**:38.
14. [△]Frankfurt HG (2005). *On Bullshit*. Princeton: Princeton University Press.
15. ^a, ^b, ^cHarris ZS (1954). "Distributional Structure." *Word*. **10**(2–3):146–162.
16. ^a, ^b, ^cFirth JR (1957). "A Synopsis of Linguistic Theory 1930–1955." *Studies in Linguistic Analysis*. pp. 1–32.
17. [△]Lee K, Firat O, Agarwal A, Fannjiang C, Grangier D (2018). "Hallucinations in Neural Machine Translation." *NeurIPS Workshop on Critiquing and Correcting Trends in Machine Learning*.
18. [△]Raunak V, Menezes A, Junczys-Dowmunt M (2021). "The Curious Case of Hallucinations in Neural Machine Translation." *Findings of the Association for Computational Linguistics: NAACL 2021*. pp. 1172–1183.
19. [△]Azaria A, Mitchell T (2023). "The Internal State of an LLM Knows When It's Lying." *arXiv preprint arXiv:2304.13734*.

20. [△]Chen C, Liu K, Chen Z, Gu Y, Wu Y, Tao M, Fu Z, Ye J (2024). "INSIDE: LLMs' Internal States Retain the Power of Hallucination Detection." *arXiv preprint arXiv:2402.03744*.
21. [△]Farquhar S, Kossen J, Kuhn L, Gal Y (2024). "Detecting Hallucinations in Large Language Models Using Semantic Consistency." *Nature*. **630**:625–630.
22. [△]Kuhn L, Gal Y, Farquhar S (2023). "Semantic Uncertainty: Linguistic Invariances for Uncertainty Estimation in Natural Language Generation." *arXiv preprint arXiv:2302.09664*.
23. [△]_aZha Y, Yang Y, Li R, Hu Z (2023). "AlignScore: Evaluating Factual Consistency With a Unified Alignment Function." *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*. pp. 11328–11348.
24. [△]_aReimers N, Gurevych I (2019). "Sentence-BERT: Sentence Embeddings Using Siamese BERT-Networks." *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. pp. 3982–3992.
25. [△]Gao T, Yao X, Chen D (2021). "SimCSE: Simple Contrastive Learning of Sentence Embeddings." *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. pp. 6894–6910.
26. [△]Wang T, Isola P (2020). "Understanding Contrastive Representation Learning Through Alignment and Uniformity on the Hypersphere." *Proceedings of the 37th International Conference on Machine Learning*. pp. 9929–9939.
27. [△]do Carmo MP (1992). *Riemannian Geometry*. Boston: Birkhäuser.
28. [△]_aBridson MR, Häfliger A (1999). *Metric Spaces of Non-Positive Curvature*. Berlin: Springer.
29. [△]_aHanley JA, McNeil BJ (1982). "The Meaning and Use of the Area Under a Receiver Operating Characteristic (ROC) Curve." *Radiology*. **143**(1):29–36.

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.