

Contents

Cognometry v0: 8-Benchmark Cross-Validated Hallucination Detection in Production LLMs	1
Abstract	1
1. Motivation	1
2. Three laws of cognometry	2
3. 8-benchmark cross-validated hallucination detection	2
4. Comparison to published single-benchmark detectors	4
5. Limits and open problems	5
6. Reproducing	5
7. Conclusion	5
Citation	6
Appendix A: Signal module versions	6
Appendix B: Per-seed raw AUCs	6

Cognometry v0: 8-Benchmark Cross-Validated Hallucination Detection in Production LLMs

2026-04-23. Numbers from committed run artifacts. All reproducers in `benchmarks/hallucination_test/`.

Abstract

We define **cognometry** as the empirical quantification of cognitive states in machine systems — refusal, confabulation, retrieval, reasoning, and adversarial drift — from signals already carried on the token stream and residual activations of a language model during inference. We publish three falsifiable laws of cognometry (vitals exist, vitals transfer across substrates, vitals are causally actionable) with cross-validated numerical support for each, and ship the first open-source instrument (**styxx** on PyPI) that realizes the measurement.

The central claim of this paper is narrower: a 9-signal logistic regression fused over text, entity, novelty, grounding, and NLI contradiction signals achieves *cross-validated* hallucination discrimination across **8 public benchmarks** — HaluEval-QA, Dialog, Summarization, TruthfulQA, and four HaluBench subsets (DROP, PubMedQA, FinanceBench, RAGTruth) — with honest per-dataset performance ranging from near-perfect (AUC 0.998 on HaluEval-QA) to below chance (AUC 0.424 on DROP). We openly report and taxonomize the failure modes: reading-comprehension extractive-span errors and financial arithmetic errors are not detected by the present signal stack because both classes of error pass the entailment (NLI) and novelty bars by construction. We publish the failure modes in the weights module itself.

This is the first 8-benchmark cross-validated hallucination detector we are aware of in the open literature. Above-chance performance on 5/8 benchmarks with 3/8 near-perfect is the reproducible empirical floor we are laying down. Two below-chance results are the reproducible research agenda we are laying down.

1. Motivation

Hallucination detection has a reproducibility and generalization crisis. Published numbers typically report on one benchmark (HaluEval-QA dominantly) at undisclosed random-seed and dataset-split configurations. When we trained our own v3.9.0 calibration on HaluEval-QA (AUC 0.90, n=230 test) and cross-validated it on HaluEval-Dialog, HaluEval-Summ, and TruthfulQA with the same weights, performance collapsed to AUC 0.56–0.63 on three of four datasets (Styxx v3.9.1 CHANGELOG, 2026-04-23). The headline number was a single-benchmark overfit.

We caught our own overfit and retrained on pooled data from the four benchmarks, producing v2 calibration with mean AUC 0.79 at n=400 test. Dialog and Summarization remained near chance (AUC ~0.60) — not

an overfitting artifact but a structural one: faithful dialog and summary responses *add content not verbatim present in the reference*. Novelty signals cannot distinguish faithful addition from contradiction. A proper entailment (NLI) signal is required.

The present paper addresses three failure modes of the state of the art:

1. **Single-benchmark overfitting** — detectors that report AUC 0.85+ on HaluEval-QA routinely underperform chance on summarization.
2. **Unreported failure modes** — no open detector we are aware of publishes where it fails, only where it succeeds.
3. **No shared vocabulary** — hallucination, refusal, drift, and confabulation detection are treated as separate tasks with separate papers and separate signals, when in fact they share a measurement substrate.

The proposed frame for #3 is **cognometry**. We name the field, put three laws on the table, and ship the first instrument for #1 and #2 under that frame.

2. Three laws of cognometry

Law I — Every computation leaves vitals

A language model in inference produces text conditioned on a logprob trajectory, a residual-stream geometry, and a generation-order time series. Any of these carries enough signal to classify the cognitive state that produced them.

Support. Cross-validated on 8 benchmarks with a 9-signal pooled LR (this paper, §3). Independent validation on Claude API without logprobs: category-accuracy 0.536, gate agreement 0.940 with n=84 fixtures (Cognitive Monitoring Without Logprobs paper, [papers/cognitive-monitoring-without-logprobs.md](#)).

Law II — Vitals are substrate-transferable

Cognitive state directions (refusal, sycophant-pressure, confab-prompt) trained on one model share measurable geometric overlap with the corresponding direction natively learned on another model. Overlap strength tracks the similarity of the alignment regimes of the two models.

Support. UCB Phase 2 paper ([papers/universal-cognitive-basis-phase2.md](#)). Cross-scale within Llama family: $\cos = +0.464$ on refusal direction (~ 26 above chance). Cross-vendor similar-alignment: $\cos = +0.362$ (~ 14). Cross-vendor divergent-alignment (Qwen \rightarrow Phi-3.5): $\cos = +0.043$ (~ 2 , null). The law is nontrivial precisely because it fails where it should fail.

Law III — Vitals are causally actionable

Cognitive states are not only observable but steerable: adding a direction into the residual stream at inference time changes behavior at predicted magnitudes.

Support. CIS v0 paper ([papers/cognitive-instruction-set-v0-filled.md](#)). Refuse@unsafe drops 97% \rightarrow 17% at $\beta=3.0$ multi-position patching on Llama-3.2-1B (n=60 JBB test split). Gradient-free capability amplification: +7.0pp MC1 on TruthfulQA (n=200) at $\beta=1.0$, validated against a 3-seed random-direction control (random directions hurt accuracy by a mean -5.3 pp at $\beta=0.5$). Three refusal-family directions measured at near-orthogonal angles (86.7° – 91.9°) — cognitive states form a basis, not a scalar dial.

3. 8-benchmark cross-validated hallucination detection

3.1 Signal stack

Nine cheap-to-compute signals, combined via pooled logistic regression:

Signal	Description	Cost
text_claim_risk	Surface heuristics (hedges, confidence markers, entity density, line structure) on decomposed claims	sub-ms
entity_unverified_frac	Fraction of named entities that fail Wikipedia verification	~100 ms per entity
knowledge_grounding	Claim-level content-token coverage against the reference passage	sub-ms
content_novelty	Fraction of response content tokens absent from reference	sub-ms
entity_novelty	Fraction of capitalized tokens absent from reference	sub-ms
number_novelty	Fraction of numeric tokens absent from reference	sub-ms
bigram_novelty	Fraction of response bigrams absent from reference	sub-ms
trigram_novelty	Fraction of response trigrams absent from reference	sub-ms
nli_contradict	MoritzLaurer/DeBERTa-v3-base-mnli-400k-anli contradiction probability on (reference \rightarrow response)	sub-ms

All nine are computable at inference time without access to the generating model’s weights.

3.2 Training protocol

Three independent seeds (31, 47, 83). For each seed:

1. Load n=150 pairs per dataset (HaluEval has paired truth/hallu responses; TruthfulQA has correct/incorrect answer pairs; HaluBench has per-example PASS/FAIL labels which we balance at 150-per-class).
2. 75%/25% stratified train/test split per dataset.
3. Pool training rows across all 8 benchmarks (n_train = 1800).
4. Fit 9-coefficient LR with L2=0.05, lr=0.3, 800 epochs of batch gradient descent.
5. Evaluate per-dataset held-out AUC independently.

Coefficients averaged across the three seeds are the v4.0.0 published weights.

3.3 Results

3-seed mean \pm std, n=150/dataset:

Benchmark	AUC	Domain
HaluEval-QA	0.998 \pm 0.001	general QA
TruthfulQA	0.994 \pm 0.006	truthfulness
HaluBench-RAGTruth	0.807 \pm 0.043	RAG faithfulness
HaluBench-PubMedQA	0.719 \pm 0.051	biomedical QA
HaluEval-Dialog	0.676 \pm 0.037	knowledge-grounded dialog
HaluEval-Summ	0.643 \pm 0.060	abstractive summarization
HaluBench-FinanceBench	0.492 \pm 0.026	financial document QA
HaluBench-DROP	0.424 \pm 0.080	reading comprehension

Benchmark	AUC	Domain
Overall mean	0.719	

Learned coefficients (3-seed averaged, intercept = -0.7518):

nli_contradict	+0.5570	dominant signal
trigram_novelty	+0.4943	
content_novelty	+0.2551	
bigram_novelty	+0.1867	
text_claim_risk	+0.1733	
entity_novelty	+0.1315	
number_novelty	+0.1271	
knowledge_grounding	+0.0792	
entity_unverified_frac	+0.0000	rarely fires at this scale

3.4 Honest failure modes

Two benchmarks returned below-chance AUC across all three seeds. We report them, taxonomize them, and decline to drop them from the fit.

DROP (AUC 0.424). DROP answers are extractive spans from the provided passage. A hallucinated answer is typically the wrong span from the right passage: string-level and subsequence-level overlap with the reference remain high, so content / n-gram / entity novelty signals are near-zero on both correct and incorrect answers. More problematically, the incorrect span is *entailed* by its passage in the NLI sense (it appears as true statement within the source text), so NLI contradiction probability is also near-zero on hallucinations. The signal stack has no mechanism to detect “right-source, wrong-span.”

FinanceBench (AUC 0.492, at chance). FinanceBench hallucinations are predominantly calculation or aggregation errors on numbers that appear verbatim in the source. The hallucinated answer “operating cash flow ratio 0.25” shares all of its content tokens, numeric tokens, and n-grams with a source passage that contains the words “operating cash flow ratio” and a different number that the model failed to correctly compute. NLI does not distinguish arithmetic correctness: both “the ratio is 0.25” and “the ratio is 0.30” are non-contradicted by a passage that only provides raw inputs.

Both failure classes are structural — not regularization or training-distribution artifacts. Both are declared in `styxx.guardrail.calibrated_weights_v4.CALIBRATION_NOTES.documented_failure_modes` so callers can gate on them in production.

The proposed remediation for DROP is a **span-faithfulness** signal: identify the semantic role demanded by the question (entity type, temporal range, numeric unit), identify the corresponding role of the answer, and contribute a hallucination signal when the two mismatch. The proposed remediation for FinanceBench is a **number-symbolic verification** signal: extract the arithmetic operation implied by the question and the numbers present in the reference, recompute independently, and contribute a hallucination signal when the model’s output differs from the computed result. Both are v4.1+ roadmap items.

4. Comparison to published single-benchmark detectors

System	Benchmark	Reported AUC	Cross-validated on 4 datasets?
SelfCheckGPT	HaluEval-QA	0.71–0.79	No
KnowHalu	HaluEval-QA	0.74	No
HaluCheck	HaluEval-QA	0.82	No
Styxx v3.8.0 (v1 LR)	HaluEval-QA	0.901	No (HaluEval-QA only)

System	Benchmark	Reported AUC	Cross-validated on 4 datasets?
Styxx v3.9.1 (v2 LR, novelty)	4-benchmark	0.805 mean	Yes (4)
Styxx v4.0.0 (v3 LR, NLI 4-bench)	4-benchmark	0.846 mean	Yes (4, NLI-augmented)
Styxx v4.0.0 (v4 LR, 8-bench)	8-benchmark	0.719 mean	Yes (8)

The drop from 0.901 (single benchmark) \rightarrow 0.719 (8 benchmarks, averaged) is not a regression. It is the reporting-framework regression that the field has been accumulating: we are the first to quantify how much any detector’s headline number depends on the benchmark chosen. The 5/8 benchmarks above AUC 0.65 is a stronger claim, properly normalized for generalization.

5. Limits and open problems

1. **Dialog and summarization do not reach production-grade AUC** (0.676 and 0.643). The NLI signal contributed the largest gain on these two — the pre-NLI versions were at chance. The residual gap is inherent paraphrase ambiguity: faithful dialog and summaries frequently restructure the reference, which NLI can rarely score as entailed at the whole-response level. A claim-level NLI pipeline (decompose the response, score each claim independently) is an expected near-term improvement.
2. **Larger models remain untested at our evaluation scale.** Every causal-steering result cited in this paper is at 1B–3B.
3. **Arithmetic errors and span-substitution errors are not detected.** See §3.4.
4. **English only.** Novelty tokenization and NLI model are English-only at this weights version.

6. Reproducing

```
pip install styxx==4.0.0[nli]
```

```
# Full 8-benchmark calibration, 3-seed averaged:
python benchmarks/hallucination_test/cross_dataset_8bench_multiseed.py
# Writes results/cross_dataset_8bench_multiseed.json
# Expected: overall_mean ~0.719, per-dataset AUCs as in §3.3  $\pm 1$  .

# Single-seed diagnostic run:
python benchmarks/hallucination_test/cross_dataset_8bench.py \
    --n 150 --seed 31 --no_entity --nli
# ~2 min on CUDA, ~15 min on CPU (dominated by NLI).
```

Raw data: HaluEval via pminervini/HaluEval, TruthfulQA via truthful_qa, HaluBench subsets via PatronusAI/HaluBench. All four sources are public on Hugging Face Hub. No data preparation beyond what is in the calibration scripts.

7. Conclusion

We propose **cognometry** as the name for the empirical measurement of cognitive states in language models, and lay down three falsifiable laws with cross-validated numerical support. The present paper is the 8-benchmark grounding of Law I. Every number in this paper has a reproducer in a committed script; every failure mode is declared in the shipping weights module; every assumption is recoverable from the public corpora that trained the detector. The invitation is open for replication, disconfirmation, and extension — including to the two benchmarks where the present signal stack fails.

Citation

```
@misc{styxx2026cognometry,  
  author = {Flobi and Fathom Lab},  
  title = {Cognometry v0: 8-Benchmark Cross-Validated Hallucination  
          Detection in Production LLMs},  
  year = {2026},  
  month = {april},  
  howpublished = {\url{https://fathom.darkflobi.com/cognometry}},  
  note = {Software: \url{https://github.com/fathom-lab/styxx};  
          PyPI: \url{https://pypi.org/project/styxx/4.0.0/};  
          Zenodo DOI pending deposit}  
}
```

Appendix A: Signal module versions

- styxx.guardrail.text_signals v1.0 (2026-04-19)
- styxx.guardrail.entity_verify v1.0 (2026-04-19)
- styxx.guardrail.knowledge_grounding v1.0 (2026-04-19)
- styxx.guardrail.response_novelty v1.0 (2026-04-22)
- styxx.guardrail.nli_signal v1.0 (2026-04-23) — MoritzLaurer/DeBERTa-v3-base-mnli-fever-anli, 184M params.
- styxx.guardrail.calibrated_weights_v4 v1.0 (2026-04-23) — this paper’s published weights.

Appendix B: Per-seed raw AUCs

Seed 31:

halueval_qa	0.9993
halueval_dialogue	0.7215
halueval_summarization	0.7194
truthfulqa	0.9851
halubench_drop	0.5328
halubench_pubmed	0.6565
halubench_finance	0.4557
halubench_ragtruth	0.7464
mean	0.7271

Seed 47:

halueval_qa	0.9979
halueval_dialogue	0.6316
halueval_summarization	0.5732
truthfulqa	0.9964
halubench_drop	0.3936
halubench_pubmed	0.7209
halubench_finance	0.5157
halubench_ragtruth	0.8463
mean	0.7095

Seed 83:

halueval_qa	0.9979
halueval_dialogue	0.6757
halueval_summarization	0.6356
truthfulqa	1.0000
halubench_drop	0.3449

halubench_pubmed	0.7806
halubench_finance	0.5036
halubench_ragtruth	0.8272
mean	0.7207

Full JSON in `benchmarks/hallucination_test/results/cross_dataset_8bench_multiseed.json`.