

Trust in One Round: Confidence Estimation for Large Language Models via Structural Signals

Pengyue Yang
The University of Sydney
Sydney, Australia
pyan8493@uni.sydney.edu.au

Jiawen Wen
The University of Sydney
Sydney, Australia
jwen8784@uni.sydney.edu.au

Haolin Jin
The University of Sydney
Sydney, Australia
hjin3177@uni.sydney.edu.au

Linghan Huang
The University of Sydney
Sydney, Australia
lhua5130@uni.sydney.edu.au

Huaming Chen*
The University of Sydney
Sydney, Australia
huaming.chen@sydney.edu.au

Ling Chen
University of Technology Sydney
Sydney, Australia
ling.chen@uts.edu.au

Abstract

Large language models (LLMs) are increasingly deployed in domains where errors carry high social, scientific, or safety costs. Yet standard confidence estimators, such as token likelihood, semantic similarity and multi-sample consistency, remain brittle under distribution shift, domain-specialised text, and compute limits. In this work, we present Structural Confidence, a single-pass, model-agnostic framework that enhances output correctness prediction based on multi-scale structural signals derived from a model’s final-layer hidden-state trajectory. By combining spectral, local-variation, and global shape descriptors, our method captures internal stability patterns that are missed by probabilities and sentence embeddings. We conduct extensive, cross-domain evaluation across four heterogeneous benchmarks—FEVER (fact verification), SciFact (scientific claims), WikiBio-hallucination (biographical consistency), and TruthfulQA (truthfulness-oriented QA). Our Structural Confidence framework demonstrates strong performance compared with established baselines in terms of AUROC and AUPR. More importantly, unlike sampling-based consistency methods which require multiple stochastic generations and an auxiliary model, our approach uses a single deterministic forward pass, offering a practical basis for efficient, robust post-hoc confidence estimation in socially impactful, resource-constrained LLM applications.

CCS Concepts

• **Computing methodologies** → **Neural networks; Knowledge representation and reasoning; Learning latent representations.**

Keywords

Large Language Models, Uncertainty Quantification, Hallucination Detection, Hidden-State Trajectory

*Corresponding author

ACM Reference Format:

Pengyue Yang, Jiawen Wen, Haolin Jin, Linghan Huang, Huaming Chen, and Ling Chen. 2026. Trust in One Round: Confidence Estimation for Large Language Models via Structural Signals. In *Proceedings of the ACM Web Conference 2026 (WWW '26)*, April 13–17, 2026, Dubai, United Arab Emirates. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3774904.3793055>

1 Introduction

Large language models (LLMs) now serve as core components in search engines, dialogue agents, scientific assistants, and decision-support systems, and are increasingly embedded into developer-facing workflows and agentic toolchains (e.g., software engineering) [15]. However, despite their impressive generative fluency, LLMs frequently produce incorrect or fabricated content, commonly referred to as hallucinations [14, 32]. This gap between generative fluency and factual fidelity has made *post-hoc confidence estimation* a critical requirement for safety-sensitive applications, such as fact verification, reasoning, and content moderation.

Existing confidence estimation methods fall into three main families, each exhibiting inherent limitations. (1) Probability-based estimators, such as mean log-probability, token entropy, and calibration-based scores, are known to be unreliable under distribution shift and miscalibration [12, 30]. (2) Semantic-similarity estimators rely on embedding geometry (e.g., Sentence-BERT [33], DeBERTa-based classifiers), whose performance can degrade severely when input text diverges from the model’s pretraining distribution [1, 20]. (3) Multi-sample consistency methods such as SelfCheckGPT [25] require repeated stochastic sampling and auxiliary natural language inference (NLI) models, resulting in prohibitive latency and computational cost, making them unsuitable for high-throughput or real-time applications. Furthermore, a fourth emerging line of work explores hidden-state geometry signals, which examine the trajectory of token-level activations rather than output probabilities or embedding-space similarity [8, 40].

These three paradigms are not ad hoc, rather, each is designed to optimize a different axis of the confidence–cost trade-off. Probability-based scores are cost-effective because they read off uncertainty from the model’s own output distribution, requiring no additional components and incurring only a small constant-factor overhead over normal decoding [12, 30]. Semantic methods leverage pre-trained encoders or sentence representations to inject high-level similarity structure, yielding strong in-domain discrimination when



the test distribution closely matches the pretraining corpus [20, 33]. Multi-sample consistency methods such as SelfCheckGPT [25] deliberately trade latency and compute for robustness: by aggregating agreement patterns across many stochastic generations and auxiliary NLI or QA models, they approximate a reference-augmented verification signal without requiring labeled data.

However, these design choices create a structural tension for Web-scale deployment. Probability-based estimators inherit the model’s miscalibration and become brittle under distribution shift; semantic estimators depend on the stability of an embedding manifold that can collapse on domain-specialized or long-form text; and sampling-based consistency methods such as SelfCheckGPT are inherently expensive because they combine multiple generations with auxiliary models. And existing efforts on hidden-state geometric signals remain sparse and fragmented, and do not provide a unified framework or a robust characterization that generalizes across domains or model families. Consequently, achieving high-throughput deployment requires a confidence signal that (i) respects strict single-pass and model-agnostic constraints, while (ii) remaining robust under domain shift—motivating a shift from probability and semantic spaces to the structural stability of the hidden-state trajectory.

A key limitation shared by these approaches is their reliance on *surface-level signals*, such as token likelihoods, embedding similarity, or cross-sample agreement, while ignoring the internal computational process that produces model predictions. During decoding, an LLM generates a sequence of hidden states that reflect its evolving reasoning dynamics. Recent studies show that these hidden-state trajectories exhibit distinctive structural patterns associated with uncertainty, including activation variance, unstable attention transitions [27], and geometric or topological irregularities in the latent space [7, 37]. However, despite these observations, trajectory-level structure remains largely underexplored as a primary confidence signal in its own right.

This paper introduces **Structural Confidence**, a framework that models *hidden-state structural stability* as an indicator of confidence that is designed to be robust across distributions within the factual and hallucination-style settings we study. Our key hypothesis is that confident predictions correspond to smooth, low-frequency trajectories with compact global geometry, whereas uncertainty manifests as spectral irregularity, sharp local fluctuations, and dispersed or fragmented global structure. By capturing these multi-scale dynamical perturbations, Structural Confidence aims to surface a class of confidence cues that often remain informative when probability-based and semantic-based signals degrade, particularly under domain shift. Our contributions are as follows:

A new confidence modality based on hidden-state structural stability. We introduce the Structural Confidence framework, a new confidence modality based on hidden-state structural stability. We define trajectory-level structural descriptors, covering spectral stability, local variation, and global shape coherence of token-level hidden states, as a theoretically motivated and empirically validated confidence signal that robustly complements existing likelihood-, embedding-, and sampling-based estimators.

A model-agnostic, single-pass confidence estimator. The core estimator operates on structural features extracted from a *single* deterministic forward pass, requiring no sampling, auxiliary

models, or architecture-specific internals. We further show that fusing these structural signals with a sentence-level semantic representation yields a lightweight, generalizable estimator, without requiring any extra LLM calls.

Extensive and cross-domain evaluation. We conduct extensive, cross-domain evaluation across FEVER [36], SciFact [38], WikiBio [19], and TruthfulQA [21]. Structural Confidence achieves state-of-the-art performance in terms of AUROC and AUPR, while demonstrating robustness under domain shift. Specifically, across the mixed-domain benchmark, our approach is significantly more efficient, for example, the NLI-based SelfCheckGPT consistency baseline incurs on average 5–6× higher FLOPs and 4–5× higher latency than our method in the confidence-estimation stage.

Together, these findings suggest that hidden-state structural stability is a promising and underexplored basis for post-hoc confidence estimation within the factual and hallucination-style benchmarks we study. Structural Confidence offers an efficient, model-agnostic, and practically domain-robust solution for assessing LLM reliability, and points toward a line of work in trustworthy AI centered on the internal dynamics of model activations.

2 Related Work

Post-hoc confidence estimation for LLMs has been approached through three main paradigms: probability-based estimators, semantic-similarity estimators, and sampling-based consistency methods. A separate line of research examines the geometry and topology of hidden states, providing the theoretical substrate for this work.

Probability-based confidence estimation. Early approaches measure uncertainty via token-level statistics such as mean log-probability, entropy, or temperature-scaled confidence [12, 30]. Despite their simplicity, these signals are known to be poorly calibrated, highly sensitive to distribution shift, and often misaligned with correctness [16]. Classical epistemic–aleatoric separation approaches such as Prior Networks require explicit access to model logits and are thus incompatible with our black-box setting [24]. Recent studies further show that log-probability can systematically assign high scores to hallucinated or stylistically fluent text [14]. Our results confirm these limitations: probability-based baselines collapse under SciFact’s domain drift.

Semantic embedding–based confidence estimation. A second paradigm relies on semantic similarity or embedding geometry to assess confidence, including selective prediction [17], representation-based analysis of embedding stability [18, 29], and sentence-level embedding classifiers [33]. These methods perform well in-domain but degrade sharply when the test distribution diverges from the pretraining corpus [1, 20].

Sampling-based consistency methods. Multi-sample agreement aims to estimate confidence by testing whether a model reaches consistent conclusions across repeated stochastic generations. Reasoning-based hallucinations have been studied extensively, particularly in chain-of-thought settings [23]. The most prominent example is SelfCheckGPT [25], which combines multiple GPT samples with auxiliary NLI- or question-generation–based checks. In its original formulation, SelfCheckGPT requires repeated decoding and additional classifier passes, which can be costly and brittle for short or noisy outputs. In our experiments, we use the official SelfCheckNLI

checker from the SelfCheckGPT library as a single-pass baseline: we score a single deterministic generation using a DeBERTa-based NLI model (no multi-sample decoding). This setting isolates the cost and behaviour of the NLI-based checking component, enabling a cleaner comparison to other single-pass estimators.

Hidden-state geometry and structural signals. A complementary line of work investigates the internal geometry and dynamics of transformer hidden states. Layer-wise probing has shown that intermediate activations encode rich linguistic and structural information [35], while mechanistic studies reveal that specific subspaces and feed-forward “key-value memories” implement interpretable computation patterns [9]. More recently, several papers have linked uncertainty to characteristic behaviours in activation space, including increased variance, curvature changes, and unstable attention dynamics [8, 40]. Topological analyses further suggest that latent representations exhibit persistent homological structure that can fragment under distributional stress [7]. A related line of work examines the geometry of encoder-only transformers such as BERT and RoBERTa, showing that their hidden states form structured subspaces and evolve along smooth, interpretable trajectories [5, 37]. However, these studies remain primarily diagnostic: they characterise hidden-state phenomena but do not convert trajectory-level structure into a practical, end-to-end confidence estimator under strict single-pass constraints.

Positioning of this work. Our approach builds upon a foundation with prior hidden-state and topology-based analyses, but diverges significantly in objective. Instead of using geometry or topology only for post-hoc interpretation, we *operationalise* trajectory structure as a confidence signal: Structural Confidence extracts a compact set of spectral stability, local variation, and shape-coherence descriptors from the final-layer trajectory, which are then fed into a lightweight supervised estimator. It offers advantages over existing baselines. Our method is deployed as a model-agnostic, drop-in scoring layer that (i) relies solely on features available from a single deterministic forward pass, (ii) assumes no access to gradients, attention maps, or auxiliary models. This work represents the first systematic effort to operationalize hidden-state structural stability as a unified, domain-general modality for post-hoc confidence and rigorously evaluate it against current state-of-the-art methods across heterogeneous factuality tasks.

3 Problem Definition

Confidence estimation for LLMs aims to determine whether a generated answer is likely to be correct. Formally, given an input x and the LLM’s single-pass autoregressive output $y = (y_1, \dots, y_T)$ with final-layer hidden states h_1, \dots, h_T , the objective is to estimate

$$c(x, y) \approx \Pr(\text{correct}(y) \mid x, y, h_{1:T}).$$

In this work, we focus on the deployment regime characterized by strict computational constraints, where the confidence estimator must (i) rely on only a *single deterministic forward pass* of the LLM, (ii) use no auxiliary task-specific models such as NLI rerankers, and (iii) perform no multi-sample decoding such as SelfCheckGPT-style sampling. These constraints reflect the latency, cost, and scalability requirements of Web-scale LLM deployment. Existing confidence estimators fall into three dominant paradigms, yet each conflicts with this budget. Probability-based heuristics (token log-probabilities,

predictive entropy) adhere to the single-pass constraint but are known to be miscalibrated and brittle under distribution shift [12]. Semantic embedding methods estimate confidence from representation similarity, but their geometry often collapses for out-of-domain (OOD) inputs where embeddings drift from their pretraining manifold [33]. Sampling-based consistency methods improve robustness by aggregating multiple samples [25], but they incur a multi-fold increase in compute and latency and therefore explicitly fall outside the single-pass deployment design.

We posit that these collective limitations arise because existing approaches operate either in the *probability space* or the *semantic space*, both of which are surface-level and sensitive to lexical or distributional mismatch. In contrast, uncertainty fundamentally emerges from the LLM’s *internal computation dynamics*. When the model is confident, its hidden-state trajectory tends to evolve smoothly along a coherent direction; when it is uncertain, the trajectory exhibits local irregularity, higher-frequency oscillation, and global shape inconsistency. These structural signatures are intrinsic to model computation and empirically more stable across domains than probabilities or static embeddings.

We therefore define a third modality for confidence estimation:

structural stability of the hidden-state trajectory.

Let $\tau = (h_1, \dots, h_T)$ denote the final-layer trajectory. We seek a function

$$c_{\text{struct}}(x, y) = g(f(\tau)),$$

where f extracts multi-scale structural descriptors (local variation, spectral stability, and shape coherence), and g maps them to a scalar confidence score. This formulation reframes confidence estimation as a structural inference problem grounded in the geometry and stability of internal LLM dynamics, and it directly motivates the Structural Confidence framework introduced in Section 4.

4 Method

Structural Confidence is a single-pass, generator-agnostic confidence estimator designed for online large language model (LLM) services. Its central idea is to extract structural stability signals from a *proxy* hidden-state trajectory computed by a frozen encoder, rather than requiring logits, sampling, or internal activations from the production LLM. This design is explicitly motivated by practical deployment constraints: modern Web APIs (e.g., GPT-4o) expose only text outputs, enforce strict rate limits, and disallow access to token-level probabilities or hidden states. Consequently, any post-hoc confidence estimator must operate under a *single-pass*, *black-box*, and *computationally bounded* regime. Structural Confidence is built around these constraints and aims to provide a stable, efficient, and generalizable mechanism for estimating correctness across diverse Web domains.

4.1 Proxy Hidden-State Trajectory Extraction

Because modern Web LLM APIs do not expose internal hidden states, we construct a *proxy* trajectory using a frozen encoder-only model. Given (x, \hat{y}) , we concatenate them into $t = \text{concat}(x, \hat{y})$ using fixed delimiters and feed t into a 12-layer transformer encoder with hidden size 768 instantiated from bert-base-uncased [4].

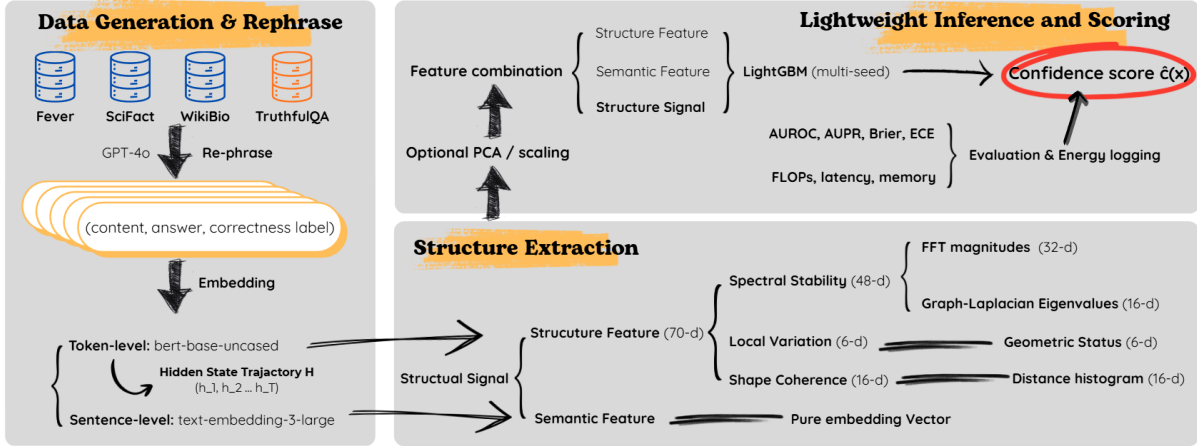


Figure 1: Overall Structural Confidence pipeline. An LLM produces a single deterministic answer; an encoder maps the (context, answer) to a hidden-state trajectory from which multi-scale structural descriptors are extracted and scored by a lightweight confidence model.

The final-layer hidden-state sequence is

$$H = (h_1, \dots, h_T) \in \mathbb{R}^{T \times 768},$$

with T truncated or padded to a maximum of 256 tokens.

Tokenizer robustness. Although BERT and GPT-4o use different tokenizers, structural descriptors depend on activation geometry rather than lexical alignment. Empirically, small variations in tokenization do not change the multi-scale patterns encoded in H , and the fixed cap of $T = 256$ standardizes descriptor computation.

In the complete Structural Signal variant (structure-feature + semantic-feature), we also obtain a sentence embedding $s \in \mathbb{R}^{d_s}$ from a pretrained encoder. No gradients are applied to either encoder, ensuring strict generator-agnostic behaviour.

4.2 Structural Descriptor Families

The complete Structural Signal consists of two components: (i) a structure-feature vector derived from the hidden-state trajectory, and (ii) a semantic-feature vector obtained as a fixed sentence embedding $s \in \mathbb{R}^{d_s}$ from a frozen pretrained encoder (OpenAI text-embedding-3-large). The semantic feature serves as a coarse-grained semantic anchor and is concatenated with the structural descriptor.

This section therefore discusses only the structure-feature component. Given a trajectory $H = (h_1, \dots, h_T)$, we extract three complementary families of structural signals. Each family is length-agnostic: it produces a fixed-size descriptor independent of T , enabling consistent inference across Web inputs.

4.2.1 Spectral stability (48-d). This family summarizes how smoothly the hidden-state trajectory evolves. It consists of two components:

Frequency-domain smoothness (32-d). We apply a real-valued DFT along the token axis of H and retain the mean and maximum power of the lowest $K = 16$ non-trivial frequencies, yielding $2K = 32$ features. Zero-padding is applied in token-space before the DFT;

because only the low-frequency band is retained, padding does not affect the retained spectral energy.

Graph-spectral diffusion (16-d). We construct a similarity graph over tokens and compute the smallest 16 eigenvalues of the normalized Laplacian L . These eigenvalues characterize global smoothness and diffusion along the trajectory. With T capped at 256, computing the lowest eigenvalues remains tractable and significantly cheaper than multi-sample generation.

4.2.2 Local variation (6-d). Local variation captures short-range instability in the hidden-state trajectory. For each consecutive pair of tokens we compute the displacement $\Delta_t = \|h_t - h_{t-1}\|_2$ and summarise the trajectory using six statistics: total path length, mean and variance of Δ_t , start-end distance, embedding-wise variance, and centroid norm. These descriptors quantify fine-grained fluctuations that global structural measures may overlook.

4.2.3 Shape coherence (16-d). Global dispersion is summarized by aggregating all pairwise distances $\|h_i - h_j\|_2$ into a 16-bin normalized histogram, producing a length-agnostic "shape signature" that reflects the overall coherence of the trajectory.

4.3 Unified Structural Representation and Granularity

Let $f_{\text{spec}}(H) \in \mathbb{R}^{48}$, $f_{\text{loc}}(H) \in \mathbb{R}^6$, and $f_{\text{shape}}(H) \in \mathbb{R}^{16}$. Structural Confidence concatenates them into:

$$u(H) = [f_{\text{spec}}(H) \parallel f_{\text{loc}}(H) \parallel f_{\text{shape}}(H)] \in \mathbb{R}^{70}.$$

To control structural resolution, we use three granularity modes:

Global. All features computed on the full trajectory.

Local. Features computed over overlapping windows ($w = 5$, stride = 2) and averaged across windows. Short windows capture token-level instability that global descriptors may smooth away.

Two-scale. The global descriptor and the averaged local descriptor are elementwise averaged. Although descriptors differ in nature (spectral, geometric, histogram-based), their normalization ensures comparable scales, and averaging serves as a multi-resolution smoothing analogous to wavelet-based representations.

Unless stated otherwise, the two-scale mode is used as default. Global-only and local-only variants are evaluated in ablations.

4.4 Confidence Estimator and Training Objective

We train a gradient-boosted decision tree model (LightGBM) to map either $u(H)$ (Struct-only) or $[u(H) \parallel s]$ (Struct+Sent) to a correctness probability. A binary logistic objective is used. Trees are shallow and few in number, reflecting the low-dimensional input and the need for fast inference. Hyperparameters (learning rate, number of leaves, feature fraction, max depth) are fixed per dataset and specified in Section 5. Tree-based models provide robustness to heterogeneous feature scales and non-linear interactions between descriptor families.

4.5 Computational Considerations

The end-to-end cost per instance is:

$$\text{Cost}(T) = \text{Cost}_{\text{enc}}(T) + O(T \log T + T^2),$$

where $\text{Cost}_{\text{enc}}(T)$ is the encoder forward pass. With T capped at 256, both $O(T \log T)$ FFT features and $O(T^2)$ shape-coherence and graph-spectral components remain CPU-feasible. In practice (Section 6.4), Structural Confidence achieves significantly lower latency and FLOPs compared to sampling-based methods (e.g., 10 rounds of SelfCheckGPT). Moreover, all computations are deterministic and do not require repeated generator calls, making the approach suitable for real-time Web inference.

4.6 Why a Proxy Encoder is Sufficient

Because proprietary LLMs such as GPT-4o do not expose hidden states, Structural Confidence relies on a lightweight encoder-only transformer to obtain token-level trajectories. Importantly, we do *not* assume that the proxy encoder reproduces the generator’s internal activations. Instead, we adopt an input-conditioned view of activation geometry: many stability- and uncertainty-related patterns are induced by the structure of the context–answer sequence and the autoregressive decoding process, rather than by model-specific parameters.

Evidence for this perspective comes from prior work showing that geometric and structural signatures—such as low-frequency spectral concentration, reduced local curvature, and coherent manifold organization—appear across different transformer families [13, 28]. These patterns are therefore at least partly model-agnostic.

Our empirical results further support the adequacy of a proxy encoder. The same structural descriptors computed from BERT trajectories consistently separate correct from incorrect outputs across FEVER, SciFact, and WikiBio-LLM, and transfer smoothly to the out-of-domain TruthfulQA benchmark. Performance degrades gradually rather than catastrophically, suggesting that the proxy trajectory preserves the text-induced stability signals that correlate with correctness.

Thus, while the proxy encoder does not approximate GPT-4o layerwise, it provides a stable and computationally lightweight surrogate trajectory whose geometric and spectral structure retains the correctness-sensitive patterns needed for post-hoc confidence estimation. A direct comparison with true generator trajectories is left for future work.

4.7 Theoretical Motivation: Why Structural Signals Predict Confidence

We posit that confidence is reflected in the stability of autoregressive hidden-state dynamics: when the model is confident, the trajectory is smooth and low-variance; under uncertainty or hallucination, it becomes irregular, high-curvature, and structurally fragmented. This behaviour is consistent with how transformer decoders update their hidden states during autoregressive prediction.

At each token step, the model updates its representation via $h_t = f(h_{t-1}, \text{Attn}_t, \text{logits}_t)$. Confident predictions typically exhibit peaked next-token distributions and stable attention patterns, so the update varies smoothly across t , yielding trajectories with (i) low-frequency spectral energy, (ii) small and homogeneous incremental changes, and (iii) coherent global structure.

Uncertain or hallucinated predictions break this stability. When the model faces ambiguous evidence or conflicting contextual cues, the next-token distribution tends to become less peaked (higher-entropy) and more variable across steps, which amplifies the sensitivity of the attention weights and the MLP pathways. The resulting update function becomes locally chaotic: small variations in logits or attention cause disproportionately large changes in h_t . This induces (i) high-frequency spectral components, (ii) irregular local curvature and step-length variance, and (iii) fragmentation and topological noise in the global trajectory.

These effects have been observed empirically across transformer families. Spectral irregularity correlates with uncertainty-driven activation drift [31], local geometric instability aligns with inconsistent token-level decision boundaries [26], and loss of topological coherence has been linked to unreliable internal reasoning pathways [7]. Importantly, these phenomena arise from the dynamics of autoregressive decoding rather than model-specific parameters, making them stable across proxy encoders and domains.

Accordingly, Structural Confidence operationalizes three measurable families—spectral stability, local variation, and global shape coherence—as model-agnostic indicators of trajectory stability: stable trajectories imply high confidence, while unstable trajectories signal representational uncertainty and likely errors.

5 Experimental Setup

We design the experiments to answer three questions: (Q1) How competitive are structural confidence signals under a unified single-pass protocol? (Q2) How robust are these signals under domain shift and mixed-domain training? (Q3) Which structural design choices—signal families, trajectory granularity, or semantic augmentation—are necessary for strong confidence estimation? This section describes the datasets, LLM generation protocol, proxy encoder extraction, confidence estimators, and evaluation metrics. Implementation follows our released codebase exactly, and all results use cached feature files generated by the scripts in Sections 4 and 6.3.

5.1 Evaluation Tasks

We evaluate confidence estimation in two settings. **(1) Sentence-level prediction:** given an instance (x, a) consisting of an input claim/question x and a GPT-4o-generated answer a , a confidence estimator outputs a scalar score predicting whether a is factually correct. **(2) Passage-level ranking:** for a set of generations from the same benchmark, we assess whether confidence scores rank correct outputs above incorrect ones. This ranking view is identical to the AUC-style factuality assessment used in recent hallucination and truthfulness work.

5.2 Datasets

We follow the LLM-ese construction in Section 4. FEVER [36] and SciFact [38] contain claims paired with binary support/refute labels. WikiBio [19] provides infobox-conditioned biography generations annotated as factually correct or incorrect. All three benchmarks are converted into single-pass evaluation sets by prompting GPT-4o once per instance and aligning each instance with a binary correctness label. After removing malformed outputs, the evaluation sets contain 1,893 (FEVER), 809 (SciFact), and 780 (WikiBio) instances.

For **cross-domain evaluation**, we merge the LLM-ese train and dev splits into a mixed-domain training pool (`mix_train`, 5,455 instances; `mix_dev`, 778 instances). The `mix_test` split (3,131 instances) contains held-out portions of all three datasets plus the entire TruthfulQA benchmark [21], excluded from training.

5.3 LLM Generation Protocol

Following SelfCheckGPT’s single-pass design, GPT-4o is queried once per instance using a unified instruction template. All outputs are generated deterministically using greedy decoding (temperature = 0, top- k = 1). No sampling, resampling, or paraphrasing is used. All confidence estimators operate strictly on these fixed outputs.

5.4 Proxy Encoder and Structural Feature Extraction

Because GPT-4o does not expose hidden states, we adopt a lightweight proxy encoder as described in Section 4. The concatenated context-answer sequence is encoded with bert-base-uncased; final-layer hidden states (h_1, \dots, h_T) are extracted from a single forward pass with a maximum length of 256 subword tokens. Structural descriptors are computed exactly as implemented in our extraction scripts: (i) *local variation* (geometric stability), (ii) *spectral stability* (FFT and graph-Laplacian spectra), and (iii) *shape coherence* (distance-histogram and TDA summaries). The concatenation yields a fixed 70-dimensional vector per instance. All datasets use the same configuration.

5.5 Baselines

We compare Structural Confidence to four families of post-hoc confidence scores. **(1) Probability-based heuristics:** mean log-probability (Mean `logp`; [11]) and predictive entropy computed from token-level log-probabilities using an open-weight model (`logp_family`), because GPT-4o does not expose token-level log-probabilities or logits in our black-box API setting. We compute these scores on the same deterministic GPT-4o answer by applying

teacher forcing with the open-weight model and extracting per-token log-probabilities. **(2) Semantic baseline:** a *Semantic-feature* variant using mean-pooled BERT representations and the same LightGBM classifier, together with a Semantic Entropy (NLI, single-pass) score computed from entailment probabilities on the same deterministic GPT-4o answer; unlike the original formulation [6], this instantiation operates under a strictly single-pass, black-box setting without additional generations or auxiliary QA models. **(3) Multi-sample consistency:** the NLI-based SelfCheckGPT baseline [25] using the official ‘SelfCheckNLI’ model. We average contradiction probabilities across passages and invert them to obtain correctness scores. We do not include heavier multi-sample reasoning methods (e.g., Stable Explanations [2]) because they require 10–50 generations per query and external encoders. **(4) Geometry-based outlier baseline:** a k -means outlier detector (`kmeans_outlier`) that scores each instance by its distance to the nearest centroid in the representation space. This baseline instantiates a standard clustering-based outlier detection heuristic, for which we cite recent comparative evaluations of clustering-based outlier detection methods [34]. We restrict comparison to such single-pass, black-box baselines and do not include white-box internal-state methods that require access to the generator’s hidden layers. In addition, we include a RACE-style structural baseline in our efficiency comparison, instantiated as a five-module proxy-trajectory pipeline following the original design [39]; this baseline shares the same proxy encoder but replaces our lightweight descriptors with a deeper, multi-stage structural analysis.

5.6 Supervised Training Protocol

Supervised estimators (*Structure-feature*, *Semantic-feature*, *Structural Signal*) use LightGBM with binary logistic loss. Training strictly follows our scripts: 200 trees, learning rate 0.05, unlimited depth, and fixed random seeds for reproducibility. No dataset-specific tuning or validation heuristics are used. Single-domain models are trained on each dataset’s LLM-ese train/dev split; cross-domain models are trained on `mix_train` and validated on `mix_dev`. Test sets remain strictly unseen.

5.7 Evaluation Metrics

We report discrimination (AUROC, AUPR) and calibration (Brier Score, Expected Calibration Error) following standard practice in LLM truthfulness [12, 22, 25]. All metrics operate on the continuous scalar confidence score produced by each method.

5.8 Efficiency Protocol

We measure the computational cost of the *confidence estimation stage only* using instrumented FLOPs and wall-clock measurements. The shared GPT-4o forward pass is excluded for all methods. For multi-pass baselines such as SelfCheckGPT, we report both the cost of a single-pass NLI check and the extrapolated cost for S hypothetical samples (scaled linearly with S as in [25]). Structural Confidence uses only a single BERT-base pass and CPU-side feature extraction.

Hardware setup. All experiments were conducted on a single workstation equipped with an AMD Ryzen 9 5900HX CPU, 32 GB of RAM, and an NVIDIA GeForce RTX 3080 Laptop GPU with

16 GB of VRAM. Unless otherwise stated, all confidence estimation methods share this hardware configuration. We measure FLOPs and wall-clock latency for the confidence-estimation stage only, excluding the shared GPT-4o generation cost.

6 Results and Analysis

We evaluate Structural Confidence in four parts: main sentence-level prediction results, cross-domain robustness, ablation studies of key design choices, and efficiency analysis.

6.1 Sentence-Level Confidence Prediction

Table 1 compares all confidence estimators across FEVER, SciFact, and WikiBio. We focus on three questions: how structural signals compare to probability-based heuristics, how they behave relative to semantic baselines, and how the fused configuration behaves against strong reference methods such as SelfCheckGPT.

Structural signals vs probability-based heuristics. Mean-log-probability and entropy struggle on FEVER and WikiBio, often falling near or below random-guess AUROC. In contrast, *Structure-feature* provides a stable ranking signal across all three datasets without using token probabilities. On FEVER and WikiBio, *Structure-feature* outperforms both probability-based heuristics in AUROC and AUPR, supporting the premise that trajectory-level stability captures correctness information that simple likelihoods cannot. On SciFact, mean-log-probability is competitive, but all probability-based methods remain brittle and offer only moderate discrimination.

Structural signals versus semantic representations. *Semantic-feature* performs well on FEVER and WikiBio (0.682/0.683 AUROC), but collapses on SciFact (0.255 AUROC), consistent with prior observations that sentence embeddings degrade under domain and stylistic shift [33]. Structural signals degrade smoothly rather than catastrophically: *Structure-feature* remains clearly discriminative on FEVER and WikiBio and retains non-trivial signal on SciFact (0.460 AUROC), providing a more conservative confidence estimate when semantic encoders misfire. This pattern is desirable in Web settings where domain and style cannot be controlled.

Fused configuration versus strong reference baselines. The fused configuration *structure* signal combines the strengths of both sources, achieving top or near-top AUROC/AUPR among single-pass baselines on FEVER and WikiBio and avoiding the embedding collapse on SciFact. Because all supervised variants share the same classifier and hyperparameters, improvements reflect complementary information rather than additional capacity. Compared to SelfCheckGPT (NLI), which attains the strongest AUROC on FEVER (0.700) but is substantially more expensive (Section 6.4), Struct+Sent trades a small accuracy gap on FEVER for a much lower computational footprint, while matching or outperforming SelfCheckGPT on WikiBio in AUPR. For applications where repeated sampling or NLI passes are infeasible, Structural Confidence therefore offers a competitive single-pass alternative.

Discrimination and calibration. Beyond AUROC and AUPR, calibration metrics (Brier Score and ECE) follow the same trend: structural descriptors and the fused Struct+Sent model are consistently better calibrated than probability-only baselines on FEVER and WikiBio, and remain comparable on SciFact. This suggests that

Table 1: AUROC / AUPR on FEVER, SciFact, and WikiBio. Bold entries denote the fused Structural Signal configuration.

Method	FEVER		SciFact		WikiBio	
	AUC	AUPR	AUC	AUPR	AUC	AUPR
Spectral Stability	0.555	0.268	0.475	0.725	0.504	0.309
Shape Coherence	0.552	0.263	0.522	0.701	0.540	0.295
Local Variation	0.534	0.247	0.515	0.689	0.565	0.320
Structure-feature	0.573	0.274	0.460	0.670	0.543	0.354
Semantic-feature	0.682	0.349	0.255	0.602	0.683	0.430
Structural Signal	0.683	0.348	0.335	0.640	0.686	0.454
Mean logp [11]	0.438	0.208	0.505	0.677	0.485	0.295
Semantic entropy (NLI, single-pass) [6]	0.432	0.201	0.512	0.684	0.453	0.264
KMeans-outlier [34]	0.518	0.234	0.544	0.696	0.594	0.355
SelfCheckGPT [25]	0.700	0.364	0.458	0.650	0.564	0.312

structural stability not only improves ranking but also yields confidence scores that better track empirical correctness frequencies.

6.2 Cross-Domain Robustness

We now examine how structural and semantic confidence signals behave under domain shift. A single model is trained on the mixed-domain `mix_train` pool and evaluated on FEVER-LLM, SciFact-LLM, WikiBio-LLM, and the out-of-domain TruthfulQA. Figure 2 plots AUROC for *Struct-only* and *Semantic-feature* across domains.

Structural signals generalize smoothly. *Structure-feature* demonstrates stable performance across FEVER-LLM, SciFact-LLM, WikiBio-LLM, and TruthfulQA, confirming that structural stability captures domain-invariant regularities in hidden-state dynamics. Despite a marginal degradation in performance when moving from training domains to TruthfulQA, it remains robust, preventing the performance collapse typically seen in cross-domain shifts.

Semantic signals oscillate strongly. *Semantic-feature* varies sharply by domain, with large drops in AUROC on SciFact-LLM and TruthfulQA. This mirrors the single-domain findings and highlights the fragility of sentence embeddings under stylistic and topical shift, especially when training and test domains diverge.

Complementary fusion. When combined, *Structural Signal* tends to track the stronger of structural or semantic signals in each domain, demonstrating that structural stability offers a robust backbone when semantic encoders are misaligned, while semantics act as a strong in-domain refinement when reliable. For Web-scale deployment, where inputs span heterogeneous domains, this combination is strictly safer than relying on embeddings alone.

6.3 Ablation Studies

In this section, we study the main structural design choices for ablation studies, including signal families, semantic augmentation, trajectory granularity, and descriptor dimensionality.

Structural families are complementary. Table 2 decomposes the full structural descriptor into its three families. Each family performs best on at least one dataset: spectral stability is strongest on FEVER, shape coherence on SciFact, and local variation on WikiBio.

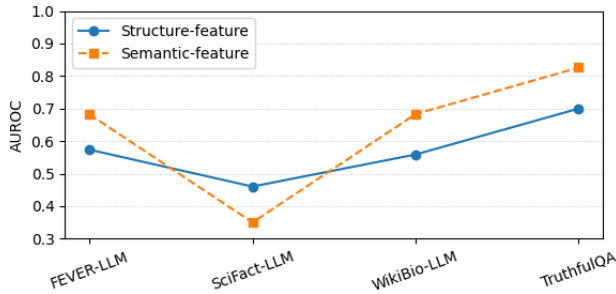


Figure 2: Cross-domain AUROC for *Structure-feature* and *Semantic-feature* trained on *mix_train* and evaluated on four domains.

Table 2: Ablation AUROC for structural families and semantic fusion. (*Local*: Local Variation, *Spectral*: Spectral Stability, *Shape*: Shape Coherence, *Struct*: Structure-feature, *Sem*: Semantic-feature)

Dataset	Local	Spectral	Shape	Struct	Sem	Structural Signal
FEVER	0.538	0.579	0.553	0.574	0.682	0.683
SciFact	0.455	0.449	0.499	0.460	0.350	0.347
WikiBio	0.570	0.532	0.520	0.559	0.683	0.685

The concatenated *Structure-feature* variant remains close to the best single family on each dataset while being more stable across all of them, indicating that the three families capture complementary aspects of trajectory stability.

Semantic augmentation: when does it help? On FEVER and WikiBio, *Semantic-feature* is the strongest single-source signal, and simple concatenation (*Structural Signal*) yields small but consistent gains, suggesting that structural descriptors refine an already strong semantic baseline. On SciFact, however, *Structure-feature* clearly dominates *Semantic-feature*, and naive fusion does not help, showing that semantic information is not uniformly beneficial under domain shift and must be combined carefully.

Trajectory granularity and dimensionality. Varying trajectory granularity (global, local, two-scale) reveals that purely local windows are insufficient on long, noisy claims (FEVER, mixed-domain), where global and two-scale variants provide better AUROC. SciFact, with shorter and more templated claims, benefits more from local structure. PCA compression from 70 to 32/16/8 dimensions yields only minor AUROC changes on FEVER and even small gains on SciFact, indicating that the descriptor is redundant and can be stored in a compact subspace.

6.4 Efficiency Analysis

Finally, we compare the compute and latency cost of Structural Confidence against stronger but more expensive baselines. Table 3 reports relative FLOPs and latency, normalised so that single-pass Structural Signal has cost 1.0×. The importance of minimizing inference-time compute is well documented in recent surveys on

Table 3: Relative FLOPs and latency of confidence estimators (Structural Confidence = 1.0).

Method	FLOPs	Latency
Structural Confidence	1.0	1.0
RACE baseline	4.0	3.0
SelfCheckGPT (NLI)	6.0	5.0

LLM deployment under resource constraints [10]. All numbers reflect the confidence estimation stage only.

Single-pass structural confidence is the cheapest. Structural Confidence requires a single BERT-base forward pass plus lightweight CPU-side feature extraction and LightGBM inference. This yields the lowest FLOPs and latency among methods that go beyond raw log-probabilities and single-vector sentence embeddings.

RACE-style and SelfCheckGPT baselines are substantially more expensive. RACE-style structural pipelines incur approximately 4× FLOPs and 3× latency relative to the Structural Signal (structure-feature + semantic-feature), while SelfCheckGPT (NLI) incurs roughly 6× FLOPs and 5× latency due to the extra NLI model and longer processing pipeline. These costs would grow linearly with the number of samples in a multi-sample SelfCheckGPT configuration. In contrast, Structural Confidence remains strictly single-pass and proxy-based.

7 Conclusion

This work introduced *Structural Confidence*, a single-pass, model-agnostic framework for post-hoc confidence estimation in large language models. By analysing the geometric and spectral structure of hidden-state trajectories, and in ablation variants with their topological summaries, it derives stability-based confidence signals that are distinct from traditional approaches, i.e., token likelihoods, sentence-level semantics, and multi-sample consistency scores, requiring only a single encoder forward pass. Using FEVER, SciFact, and WikiBio-hallucination as the benchmarks, Structural Confidence remains reliable under domain shift, outperforms probability-based baselines, and provides complementary gains over sentence-level semantic embeddings. In a mixed-domain regime trained on the union of FEVER-, SciFact-, and WikiBio-LLM and evaluated on these three test sets plus the out-of-domain TruthfulQA benchmark, the same structural model maintains competitive AUROC and AUPR without dataset-specific tuning, and the fused Struct+Sent variant offers a favourable accuracy–efficiency trade-off compared to an NLI-based SelfCheckGPT baseline in the confidence-estimation stage. Our study is limited to a small set of factual and truthfulness benchmarks, one proprietary base model, and FLOPs-/latency-based efficiency estimates on specific hardware configuration. Future work will extend structural confidence modelling to other Web tasks, open-weight and multimodal models, and integration into end-to-end systems such as fact-checking, retrieval-augmented assistance, and content moderation.

References

- [1] Rishi Bommasani. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* (2021).

- [2] Chao Chen, Chenghua Guo, Rufeng Chen, Guixiang Ma, Ming Zeng, Xiangwen Liao, Xi Zhang, and Sihong Xie. 2024. Training for stable explanation for free. *Advances in Neural Information Processing Systems* 37 (2024), 3421–3457.
- [3] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. 2005. Stability of persistence diagrams. In *Proceedings of the twenty-first annual symposium on Computational geometry*. 263–271.
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*. 4171–4186.
- [5] Kawin Ethayarajh. 2019. How contextual are contextualized word representations. *Comparing the geometry of BERT, ELMo, and GPT-2 Embeddings 2* (2019).
- [6] Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. 2024. Detecting hallucinations in large language models using semantic entropy. *Nature* 630, 8017 (2024), 625–630.
- [7] Stephen Fitz, Peter Romero, and Jiyan Jonas Schneider. 2024. Hidden holes: topological aspects of language models. *arXiv preprint arXiv:2406.05798* (2024).
- [8] Stanislav Fort, Huiyu Hu, and Balaji Lakshminarayanan. 2019. Deep ensembles: A loss landscape perspective. *arXiv preprint arXiv:1912.02757* (2019).
- [9] Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2021. Transformer feed-forward layers are key-value memories. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 5484–5495.
- [10] Sanjay Surendranath Giriya, Lakshit Arora, Shashank Kapoor, Dipen Pradhan, Aman Raj, and Ankit Shetgaonkar. 2025. Optimizing llms for resource-constrained environments: A survey of model compression techniques. In *2025 IEEE 49th Annual Computers, Software, and Applications Conference (COMPSAC)*. IEEE, 1657–1664.
- [11] Nuno M Guerreiro, Elena Voita, and André FT Martins. 2023. Looking for a needle in a haystack: A comprehensive study of hallucinations in neural machine translation. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*. 1059–1075.
- [12] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International conference on machine learning*. PMLR, 1321–1330.
- [13] John Hewitt and Christopher D Manning. 2019. A structural probe for finding syntax in word representations. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. 4129–4138.
- [14] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM computing surveys* 55, 12 (2023), 1–38.
- [15] Haolin Jin, Linghan Huang, Haipeng Cai, Jun Yan, Bo Li, and Huaming Chen. 2024. From llms to llm-based agents for software engineering: A survey of current, challenges and future. *arXiv preprint arXiv:2408.02479* (2024).
- [16] Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. 2022. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221* (2022).
- [17] Amita Kamath, Robin Jia, and Percy Liang. 2020. Selective question answering under domain shift. *arXiv preprint arXiv:2006.09462* (2020).
- [18] Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. Revealing the dark secrets of BERT. *arXiv preprint arXiv:1908.08593* (2019).
- [19] Rémi Lebre, David Grangier, and Michael Auli. 2016. Neural text generation from structured data with application to the biography domain. *arXiv preprint arXiv:1603.07771* (2016).
- [20] Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. 2022. Holistic evaluation of language models. *arXiv preprint arXiv:2211.09110* (2022).
- [21] Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. Truthfulqa: Measuring how models mimic human falsehoods. In *Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: long papers)*. 3214–3252.
- [22] Tianyu Liu, Yizhe Zhang, Chris Brockett, Yi Mao, Zhifang Sui, Weizhu Chen, and William B Dolan. 2022. A token-level reference-free hallucination detection benchmark for free-form text generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 6723–6737.
- [23] Qing Lyu, Shreya Havaladar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. 2023. Faithful chain-of-thought reasoning. In *The 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (IJCNLP-AACL 2023)*.
- [24] Andrey Malinin and Mark Gales. 2018. Predictive uncertainty estimation via prior networks. *Advances in neural information processing systems* 31 (2018).
- [25] Potsawee Manakul, Adian Liusie, and Mark Gales. 2023. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models. In *Proceedings of the 2023 conference on empirical methods in natural language processing*. 9004–9017.
- [26] R Thomas McCoy, Junghyun Min, and Tal Linzen. 2020. BERTs of a feather do not generalize together: Large variability in generalization across models with similar test set performance. In *Proceedings of the third blackboxnlp workshop on analyzing and interpreting neural networks for nlp*. 217–227.
- [27] Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one? *Advances in neural information processing systems* 32 (2019).
- [28] John Morris, Volodymyr Kuleshov, Vitaly Shmatikov, and Alexander M Rush. 2023. Text embeddings reveal (almost) as much as text. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. 12448–12460.
- [29] Marius Mosbach, Maksym Andriushchenko, and Dietrich Klakow. 2020. On the stability of fine-tuning bert: Misconceptions, explanations, and strong baselines. *arXiv preprint arXiv:2006.04884* (2020).
- [30] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. 2019. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in neural information processing systems* 32 (2019).
- [31] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. 2019. On the spectral bias of neural networks. In *International conference on machine learning*. PMLR, 5301–5310.
- [32] Vipula Rawte, Amit Sheth, and Amitava Das. 2023. A survey of hallucination in large foundation models. *arXiv preprint arXiv:2309.05922* (2023).
- [33] Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084* (2019).
- [34] Braulio V Sánchez Vinces, Erich Schubert, Arthur Zimek, and Robson LF Cordeiro. 2025. A comparative evaluation of clustering-based outlier detection. *Data Mining and Knowledge Discovery* 39, 2 (2025), 13.
- [35] Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. BERT rediscovers the classical NLP pipeline. *arXiv preprint arXiv:1905.05950* (2019).
- [36] James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a large-scale dataset for fact extraction and VERification. *arXiv preprint arXiv:1803.05355* (2018).
- [37] Lucrezia Valeriani, Diego Doimo, Francesca Cuterello, Alessandro Laio, Alessio Ansuni, and Alberto Cazzaniga. 2023. The geometry of hidden representations of large transformer models. *Advances in Neural Information Processing Systems* 36 (2023), 51234–51252.
- [38] David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. Fact or fiction: Verifying scientific claims. *arXiv preprint arXiv:2004.14974* (2020).
- [39] Changyue Wang, Weihang Su, Qingyao Ai, and Yiqun Liu. 2025. Joint Evaluation of Answer and Reasoning Consistency for Hallucination Detection in Large Reasoning Models. *arXiv preprint arXiv:2506.04832* (2025).
- [40] Daquan Zhou, Zhiding Yu, Enze Xie, Chaowei Xiao, Animeshree Anandkumar, Jiashi Feng, and Jose M Alvarez. 2022. Understanding the robustness in vision transformers. In *International conference on machine learning*. PMLR, 27378–27394.

A Additional Properties of Structural Descriptors

This appendix complements Section 4 by summarising robustness, complexity, and portability properties of the structural descriptors introduced in the main text. We focus on interpretability and reuse, rather than restating algorithmic definitions.

A.1 Fixed Dimensionality and Length-Agnostic Design

All descriptor families (local variation, spectral stability, and shape coherence) are designed to have fixed dimensionality, independent of the generation length T . Local variation aggregates token-to-token geometric statistics into a small number of scalars. Spectral descriptors compress the FFT magnitude profile into summaries of low-frequency energy, dominant components, and spectral entropy. Shape coherence builds fixed-bin histograms of pairwise distances, optionally combined with a few topological statistics.

In the global setting used in our experiments, the final structural descriptor therefore has constant dimension, enabling uniform application across Web queries of widely varying length.

Dataset	Struct-only	PCA-32	PCA-16	PCA-8
FEVER	0.574	0.552	0.544	0.534
SciFact	0.460	0.446	0.412	0.395

Table 4: Impact of PCA compression of the structural descriptor (two-scale setting, AUROC).

A.2 Robustness to Activation and Configuration Variations

The descriptors operate directly on raw hidden-state trajectories without explicit per-sample normalisation. Empirically, the use of energy ratios, aggregated geometric statistics, and distance distributions yields robustness to moderate shifts in activation scale and offset.

Across all datasets (FEVER, SciFact, WikiBio, and TruthfulQA), structural features and downstream confidence scores remained numerically stable under multiple random seeds, data splits, and minor configuration variations. This behaviour aligns with stability principles in topological data analysis, where summary functionals (e.g., persistence diagrams) are provably stable under perturbations of the underlying space [3]. We do not claim new guarantees here, but the same intuition applies: descriptors that focus on global geometric structure are less sensitive to microscopic activation noise.

A.3 Computational Complexity

The descriptors are inexpensive relative to the transformer forward pass. Local variation requires only vector differences and simple reductions over T steps. Spectral descriptors use an $O(T \log T)$ real-valued FFT per trajectory, negligible compared to model inference. Shape coherence uses vectorised pairwise distances and fixed-bin histogramming; in windowed configurations, the computation is restricted to fixed-size chunks, keeping the effective cost close to linear in T for long sequences.

Consequently, Structural Confidence is dominated by the cost of a single LLM forward pass. The structural computation adds only a small constant factor, quantified empirically in Section 6.4 and Appendix C.2.

A.4 Reusability Across Models and Tasks

Because the descriptors depend solely on the hidden-state trajectory and do not require token identities, logits, or calibration layers, they transfer across decoder-only LLMs without modification. Any model exposing final-layer hidden states can be plugged into the Structural Confidence pipeline, making the descriptors reusable across Web tasks such as fact-checking, scientific assistance, or moderation.

B Extended Visualisation of Structural Confidence

To complement the numerical results in Section 6.1 and Section 6.2, we provide additional visualisations derived from the same result files used in the main tables.

B.1 Cross-Dataset AUROC Profile

Figure 3 shows AUROC for structural and semantic variants across FEVER, SciFact, and WikiBio. Whereas Table 1 reports absolute scores, the bar plot highlights relative gaps between descriptor families and illustrates the multi-scale nature of structural confidence.

B.2 SciFact Heatmap

Figure 4 provides a per-metric heatmap (AUROC and AUPR) for the same variants on SciFact, illustrating the domain-shift effects discussed in Section 6.2. Semantic confidence performs strongly on FEVER and WikiBio but degrades sharply on SciFact, whereas structural variants retain weaker but more stable signals across metrics.

C Additional Ablations and Robustness Checks

We report two additional ablations that complement the main text: PCA compression of the structural descriptor and a detailed energy analysis on FEVER.

C.1 PCA Compression of the Structural Descriptor

To assess redundancy in the descriptor, we apply PCA to the full structural vector and retain a fixed number of principal components before training the classifier. Table 4 reports AUROC on FEVER and SciFact at several target dimensions.

Moderate compression (PCA-32) yields only small drops on FEVER, suggesting that a large fraction of the signal lies in a lower-dimensional subspace. On SciFact, where all methods struggle, aggressive compression degrades performance more noticeably. The full descriptor therefore provides a useful safety margin, especially in harder domains.

C.2 Detailed Energy Comparison on FEVER

The main text reports average relative compute across datasets, normalised to Structural Confidence. Table 5 complements this with a single-dataset view on FEVER, normalised to a SelfCheckGPT-style baseline that performs multiple stochastic generations per input.

The results show the trade-off between multi-sample consistency and single-pass estimators. A SelfCheckGPT-style approach incurs substantial compute overhead, whereas Structural Confidence—using a single greedy generation and a lightweight gradient-boosted classifier—reduces FLOPs and runtime by over an order of magnitude. This supports the claim that structural signals can provide useful confidence estimates at a compute budget suitable for large-scale Web deployment.

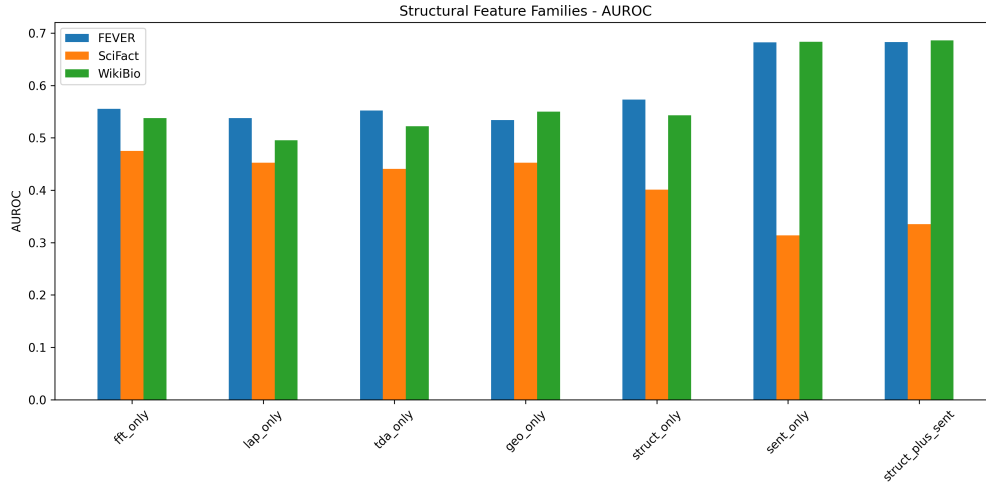


Figure 3: AUROC of `fft_only`, `lap_only`, `tda_only`, `geo_only`, `struct_only`, `sent_only`, and `struct_plus_sent` across FEVER, SciFact, and WikiBio. Unified structural descriptors and the fused Struct+Sent variant show the most stable behaviour across datasets.

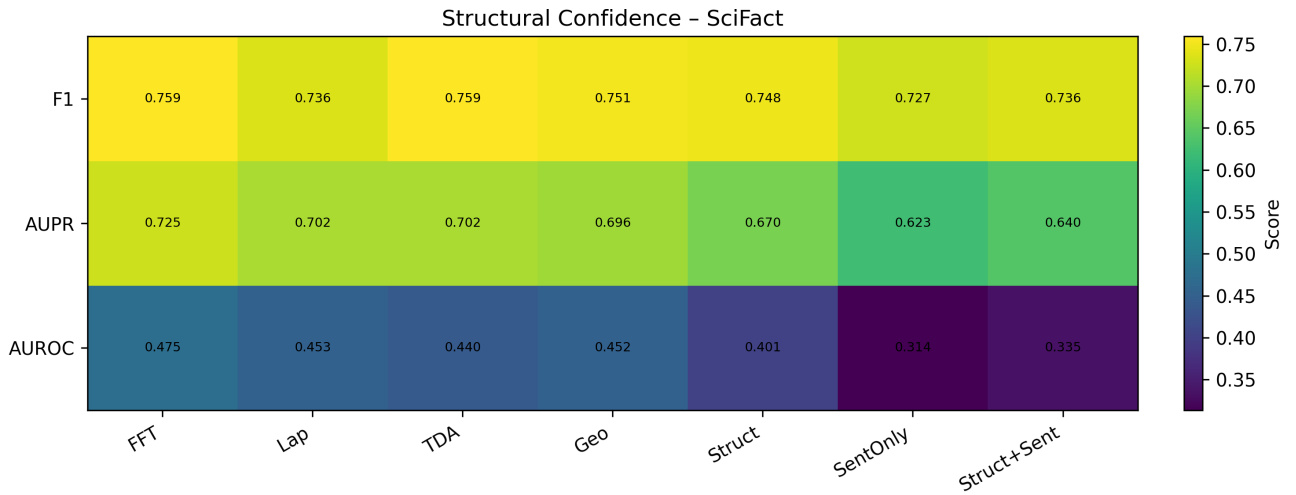


Figure 4: Per-metric heatmap (AUROC and AUPR) for structural and semantic variants on SciFact. Semantic confidence exhibits a pronounced collapse relative to FEVER and WikiBio, while structural trajectories provide weaker but more stable signals across metrics.

Method	#Generations	Rel. FLOPs	Rel. Runtime	Rel. Memory
SelfCheckGPT-style baseline	5	1.00	1.00	1.00
Structural Confidence (ours)	1	0.03	0.04	0.07

Table 5: Energy comparison on FEVER. Relative costs are normalised to a SelfCheckGPT-style baseline using $S = 5$ stochastic generations per input. All methods rely on the same underlying LLM; differences arise solely from the number of generations and additional modules at inference time.