

HalluZig: Hallucination Detection using Zigzag Persistence

Shreyas N. Samaga Gilberto Gonzalez Arroyo Tamal K. Dey

Department of Computer Science
Purdue University
West Lafayette, IN
{ssamaga, gonza982, tamaldey}@purdue.edu

Abstract

The factual reliability of Large Language Models (LLMs) remains a critical barrier to their adoption in high-stakes domains due to their propensity to hallucinate. Current detection methods often rely on surface-level signals from the model’s output, overlooking the failures that occur within the model’s internal reasoning process. In this paper, we introduce a new paradigm for hallucination detection by analyzing the dynamic topology of the *evolution of model’s layer-wise attention*. We model the sequence of attention matrices as a *zigzag graph filtration* and use *zigzag persistence*, a tool from Topological Data Analysis, to extract a topological signature. Our core hypothesis is that factual and hallucinated generations exhibit distinct topological signatures. We validate our framework, HalluZig, on multiple benchmarks, demonstrating that it outperforms strong baselines. Furthermore, our analysis reveals that these topological signatures are generalizable across different models and hallucination detection is possible only using structural signatures from partial network depth.

1 Introduction

Large Language Models (LLMs) form the foundation of modern natural language processing, powering systems for search, question answering, decision support for various domains such as healthcare, law and finance. Despite their impressive fluency, LLMs are prone to *hallucination* - the generation of confident yet factually incorrect or unsupported content (Huang et al., 2025; Sahoo et al., 2024). This undermines trust and remains a critical barrier in the adoption of LLMs in safety-critical applications. A growing body of work has sought to address this (Manakul et al., 2023; Sriramanan et al., 2024; Fadeeva et al., 2024; Farquhar et al., 2024; Azaria and Mitchell, 2023; Zhou et al., 2025; Bazarova et al., 2025; Orgad et al.,

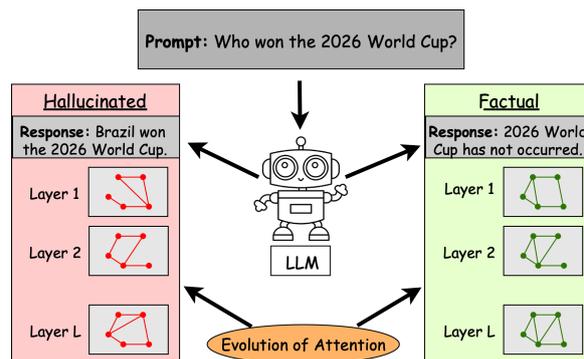


Figure 1: Attention matrices modeled as graphs show distinct topological patterns as they evolve through a model’s layers. We leverage zigzag persistence in topological data analysis to quantify these evolving attention structures for hallucination detection.

2025; Binkowski et al., 2025), yet most of the existing methods share a common limitation: they primarily operate on the final output text or shallow token-level statistics. They inspect the result of the model’s reasoning and not the reasoning pathway.

To address this, we shift the focus from *what* the model generates to *how* it arrives at its conclusion. We hypothesize that a faithful generation relies on coherent flow of information, where tokens consistently attend to the relevant evidence across layers. Contrastingly, hallucination may arise when this flow breaks down and attention patterns diverge towards spurious contexts, leading to fabricated answers. Capturing these dynamics requires tools that can characterize how the structure of attention evolves across the model’s layers.

Characterizing this evolving structure is a non-trivial task. While aggregate metrics of the attention matrix, such as eigenvalues (Binkowski et al., 2025) or the determinant (Sriramanan et al., 2024), provide valuable insights, they do not capture higher-order structural properties. They cannot, for example, describe how distinct groups of tokens form conceptual clusters or how reason-

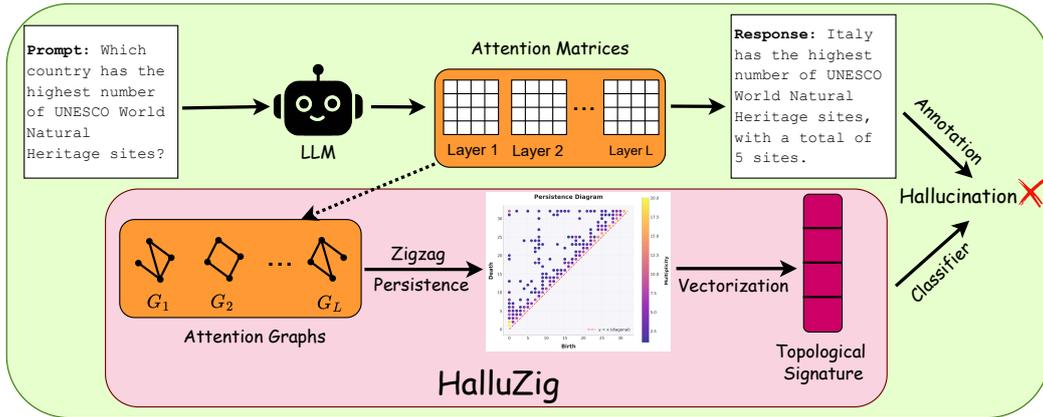


Figure 2: **The HalluZig framework: capturing attention evolution for hallucination detection.** We model the layer-wise attention matrices from an LLM as a sequence of attention graphs (G_1, \dots, G_L). Zigzag persistence is applied to this sequence to capture the evolution of topological features resulting in a Persistence Diagram. The Persistence Diagram is vectorized into a topological signature, which is used by a classifier to detect hallucinations.

ing loops are formed. Topological Data Analysis (TDA) provides a principled way to bridge this gap.

TDA (Edelsbrunner and Harer, 2010; Dey and Wang, 2022) provides a mathematical language to describe the ‘shape’ of this evolving structure. Persistent Homology, a flagship concept of TDA, is a method for identifying the most robust structural features of a system by analyzing it at all scales simultaneously. This is achieved through a filtration, a process analogous to gradually lowering a threshold on attention weights to see which conceptual clusters and reasoning loops are fundamental (i.e., they persist for a long time) versus those that are ephemeral artifacts of noise. However, standard persistent homology is limited to analyzing systems that only grow, following a nested sequence ($G_0 \subseteq G_1 \subseteq \dots$).

This model is insufficient for capturing evolution of attention, where the structure is not merely augmented but is completely transformed between layers, with connections being both formed and broken. To capture this complex evolution, we leverage *zigzag persistence* (Carlsson and de Silva, 2010; Maria and Oudot, 2016; Carlsson et al., 2019; Dey and Hou, 2022). Zigzag persistence is an extension of standard persistence designed to track topological changes through a series of inclusions (additions) and deletions (subtractions). Thus, viewing through the lens of zigzag persistence enables us to move beyond quantifying individual connections, to characterizing the topology of the evolving attention matrices.

In this paper, we introduce HalluZig - a novel

framework that captures the *evolution of attention* in an LLM using *zigzag persistence*. We model the attention matrix at each layer as a graph (*attention graph*) and connect successive graphs through a *zigzag filtration* (Dey and Wang, 2022, Chapter 4). By computing the zigzag persistence of this filtration, we obtain a topological summary quantifying births and deaths of connected components and cycles as information propagates through the model. This topological summary is distilled into a numerical vector to give a *topological signature* using established methods (Adams et al., 2017; Atenza et al., 2020). Our core hypothesis is that factual and hallucinated responses leave distinct topological signatures in the evolution of the model’s attention. To evaluate this hypothesis, we experiment with diverse datasets, which include human annotated datasets and generic QA based datasets which we annotate using LLM-as-a-judge paradigm (Zheng et al., 2023). Our experiments confirm the hypothesis, showing that the topological signatures can distinguish between factual and hallucinated responses. Further, we demonstrate the applicability of HalluZig for hallucination detection using partial network depth. We show that it can reliably detect hallucinations by analyzing only the first 70% of the model’s layers with minimal degradation in accuracy compared to a full-model analysis. Furthermore, we demonstrate that these topological signatures are not model-specific, exhibiting remarkable zero-shot generalization when transferred between different LLM architectures.

In summary, our main contributions are as fol-

lows: (1) We propose a new framework for hallucination detection by modeling the evolution of attention through an LLM. (2) To the best of our knowledge, this is the first application of zigzag persistence for capturing the layer-to-layer structural transformations of attention graphs for this task. (3) We demonstrate that HalluZig outperforms strong baselines on multiple datasets annotated for hallucination detection. (4) We empirically demonstrate that HalluZig achieves near-maximum performance when restricted to topological signatures from the first 70% of model layers.

2 Related Work

Research in hallucination detection has attracted a lot of attention in the recent times (Huang et al., 2025; Zhang et al., 2023; Wang et al., 2024) can be broadly classified into *black-box* methods and *white-box* methods.

Black-box methods. These methods operate only on the model’s final text output. Consistency based techniques such as (Manakul et al., 2023; Chen et al., 2024; Kuhn et al., 2023; Qiu and Mikkilineni, 2024; Nikitin et al., 2024) evaluate agreement among multiple generations. These methods rely on multiple model runs which imposes significant computational overhead. Surface-level confidence measures, such as perplexity, logit entropy, or predictive uncertainty (Fadeeva et al., 2024; Malinin and Gales, 2021) provide lightweight alternatives but with limited discriminating power, as they do not consider the model’s internal reasoning process.

White-box methods. These methods leverage internal representations such as hidden states, attention maps, or logits. One of the early works in this direction was (Azaria and Mitchell, 2023) which had a linear probe into these states to determine factuality. Subsequent studies such as (Farquhar et al., 2024; Chen et al., 2024; CH-Wang et al., 2024) quantified internal uncertainties by comparing hidden states across multiple generations. (Du et al., 2024) reduced annotation requirements while (Kossen et al., 2024) learned to approximate expensive self-consistency scores. A more recent and less explored white-box direction involves analyzing *attention maps*. Chuang et al. (2024) introduced lookback ratio which measures how strongly a model attends to relevant input tokens when generating context-dependent answers.

Sriramanan et al. (2024) introduced simple attention statistics to flag a response as hallucinated in an unsupervised manner. Binkowski et al. (2025) use the eigen values of the attention matrix and the eigen values of the Laplacian of the attention matrix (modelled as a graph) to classify whether a response is hallucinated. Bazarova et al. (2025) introduce a topology-based hallucination detection technique which leverages a topological divergence metric between the prompt and the response sub-graphs. These techniques typically exploit only local or scalar attention features, overlooking the rich, global structure of the evolution of attention graph that HalluZig aims to capture.

3 Method

3.1 Attention Mechanism

The *self-attention* (Vaswani et al., 2017) is the core component of the transformer architecture, which allows the LLM to dynamically weigh the importance of different tokens in a sequence when producing a representation for a given token.

At a high level, self-attention operates using three learned vector representations for each token: (1) Query(Q), (2) Key(K) and (3) Value(V). Given a generated sequence of tokens $S = \{t_1, \dots, t_T\}$, let $X \in \mathbb{R}^{T \times d}$ denote the matrix of T tokens, each having a d -dimensional representation. Let $W_Q, W_K, W_V \in \mathbb{R}^{d \times d}$ denote the trainable projection matrices. Then, the three vector representations Q, K and V are XW_Q, XW_K and XW_V respectively. Note that $Q, K, V \in \mathbb{R}^{T \times d}$. The attention mechanism is defined as follows:

$$\text{Attn}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d}} \right) V,$$

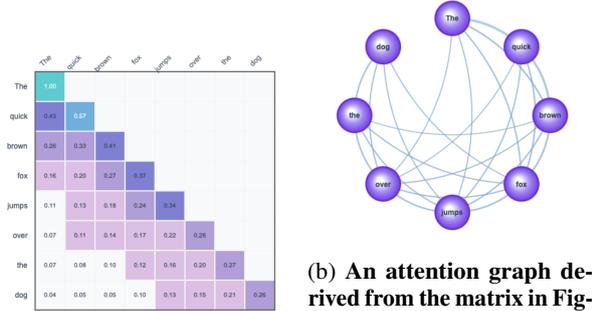
where d denotes the dimension of the token embedding. The matrix $A = \text{softmax} \left(\frac{QK^T}{\sqrt{d}} \right)$ is called the *attention matrix*.

Modern LLMs use multi-head attention, which performs this process multiple times in parallel with different, learned projections for Q, K and V . We denote the attention matrix at head h of layer l as $A^{(l,h)}$.

3.2 Attention Graph Construction

For each layer $l \in \{1, 2, \dots, L\}$, we average the attention matrices across all heads to get a mean attention matrix $A^{(l)} \in \mathbb{R}^{T \times T}$. We model this as a weighted, graph $G_l = (V, E_l, w_l)$, where V is the set of tokens, and $w_l(t_i, t_j)$ is the attention

weight from token t_i to t_j . We choose the top k percentile of attention weights to form the edges of the graph G_l in order to focus on the most significant structural connections within the layer, ensuring our analysis is robust to the noise from low-value attention weights. We refer to these graphs as *attention graphs*. This gives us a sequence of graphs, which captures the state of information flow at a specific depth in the model. Refer to Figure 3 for an illustration.



(a) **A single-head attention matrix from a decoder-only model.** The matrix is lower-triangular due to the causal mask, which prevents tokens from attending to future positions in the sequence.

(b) **An attention graph derived from the matrix in Figure 3a.** The nodes represent tokens, while the edges visualize the of attention. The thickness of each edge corresponds to its attention weight, highlighting the key inter-token relationships that form the basis for our topological analysis.

Figure 3: **Visualizing the Attention Mechanism.** (a) A causally masked attention matrix. (b) The corresponding attention graph where nodes are tokens and thick edges represent high-attention links. This structural representation is the input to our topological pipeline.

3.3 Topological Signatures

To analyze the structure of attention dynamics, we employ tools from TDA. TDA provides a mathematical framework for characterizing the "shape" of complex data. Here, we introduce the core concepts, building intuition from the familiar structure of an attention graph.

Our starting point is an attention graph, where tokens are vertices and attention scores are weighted edges. TDA provides a formal language to describe the structure of such graphs that keep on changing.

A *filtration* \mathcal{F} of graphs is a nested sequence of graphs indexed by natural numbers/integers

$$\mathcal{F} : G_0 \subseteq G_1 \subseteq G_2 \dots \subseteq G_n.$$

Notice that the inclusions here are all in the forward direction, i.e., the sequence is a non-decreasing sequence of graphs. Now, if some of the inclusions

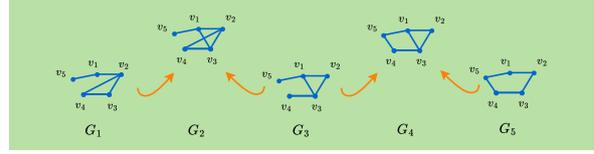


Figure 4: **Zigzag filtration.** The figure shows a zigzag filtration where $G_2 = G_1 \cup G_3$ and $G_4 = G_3 \cup G_5$.

were to be reversed, we get a *zigzag filtration* \mathcal{Z}

$$\mathcal{Z} : G_0 \subseteq G_1 \supseteq G_2 \subseteq \dots \supseteq G_n.$$

Refer to Figure 4. Unlike a standard filtration, which only allows the addition of new vertices/edges, a zigzag filtration also allows their removals. Attention values between tokens, as a sentence passes through various layers, keep evolving. To capture the evolution of attention graphs across layers, it is crucial to allow the deletion of edges. This is because two tokens that exhibit strong mutual attention in layer l_i may no longer do so in layer l_j , requiring the corresponding edge to be removed from the attention graph at that layer.

We obtain a zigzag persistence module $M_{\mathcal{Z}}$ from such a zigzag filtration of graphs by computing the homology groups in dimensions 0 (# of connected components) and 1 (# of independent cycles) of each graph G_i appearing in the filtration. In other words, for each stage G_i , we record the p -dimensional topological features (such as connected components for $p = 0$, loops for $p = 1$). These homology groups are then connected by linear maps that follow the direction of the inclusion arrows in \mathcal{Z} .

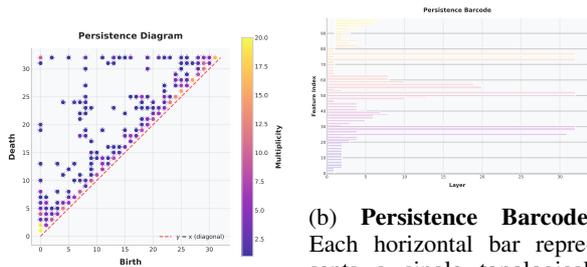
$$M_{\mathcal{Z}} : H_p(G_0) \rightarrow H_p(G_1) \leftarrow \dots \leftarrow H_p(G_n)$$

The persistence module $M_{\mathcal{Z}}$ thus encodes how topological features appear, disappear, or reappear as the graphs evolve. We refer the reader to Appendix A for formal definitions of (simplicial) homology groups.

A key result in zigzag persistence is that any zigzag persistence module $M_{\mathcal{Z}}$ decomposes into a direct sum of indecomposable interval modules (Carlsson and de Silva, 2010). Each interval module corresponds to a contiguous range of indices in the zigzag filtration over which a homology class exists. Intuitively, each interval corresponds to the "lifetime" of an individual topological feature. This decomposition is unique up to reordering. This collection of intervals provides a complete and

concise summary of the filtration’s topological dynamics and can be visualized in two equivalent ways:

- **Barcodes:** A collection of horizontal line segments, where the length and the position of each line segment represents the lifetime of a feature.
- **Persistence Diagrams:** A 2D scatter plot of (birth, death) coordinates for each feature.



(a) **Persistence Diagram.**

Each point (b, d) corresponds to a topological feature born at filtration value b and dying at d . The vertical distance from the diagonal $y = x$, known as the feature’s persistence.

(b) **Persistence Barcode.**

Each horizontal bar represents a single topological feature, starting at its birth layer (x-axis) and ending at its death layer. The length of the bar directly visualizes the feature’s persistence. For readability, this plot displays the first 100 bars, sorted vertically by birth time.

Figure 5: **Two Views of a Topological Summary.** (a) The 2D persistence diagram plots each feature’s birth vs. death layers. (b) The 1D barcode represents each feature’s lifetime as a horizontal bar.

3.4 Featurization and Classification

A persistence diagram is not a suitable input for most machine learning models. We need to transform each persistence diagram into a vector, which would be representative of the topological information captured in the persistence diagram. This vector can be used as input to a machine learning classifier. There are multiple known vectorization techniques such as persistence images (Adams et al., 2017), persistence entropy (Atienza et al., 2020), persistence landscapes (Bubenik, 2015), Betti curves (Umeda, 2017).

3.5 Zigzag Persistence of Attention Dynamics

To analyze the information flow between layers, we focus on consecutive pairs of graphs G_l and G_{l+1} . We construct a short zigzag filtration between these two graphs as follows:

$$G_l \hookrightarrow G_l \cup G_{l+1} \hookleftarrow G_{l+1}.$$

Here, $G_l \cup G_{l+1}$ is an intermediate graph containing the union of vertices and edges in G_l and G_{l+1} . Since the vertex set is the same for graphs from each layer, $G_l \cup G_{l+1}$ is the union of the edges of G_l and G_{l+1} . Refer to Figure 4 for an illustration. The weight of the edge in $G_l \cup G_{l+1}$ is taken as the maximum of the weights of the edge in G_l and G_{l+1} if it exists in both graphs. This explicitly models how the structure of G_l is transformed to produce the structure of G_{l+1} . The pairwise connections weave together all graphs from layer 1 to layer L to model the full sequence of attention dynamics.

In our case, the vertex set of the graphs do not change and thus we focus on the evolution of loops (representing cyclic attention on tokens). Algebraically, we track the evolution of the number of independent cycles or loops in G_l , which is the rank of the 1-dimensional homology group $H_1(G_l)$. Such loops can be interpreted in two ways: as stable, resonant structures that reinforce a coherent semantic concept, or as flawed, circular reasoning pathways. We hypothesize that these two cases can be distinguished by their topological persistence. Cycles in factual statements are likely to be formed by strong, stable attention weights that persist across multiple layers, representing the successful consolidation of a concept. Conversely, we posit that hallucinations may manifest as numerous, short-lived, and structurally unstable cycles, indicative of spurious thought loops. By analyzing the 1-dimensional persistence diagrams, we aim to capture these signatures of flawed reasoning that are invisible to simpler connectivity measures.

In this paper, we use three vectorization schemes to vectorize the zigzag persistence diagrams and capture the topological information present in the barcodes. Different vectorization techniques present different representations of the underlying topological information:

(1) **Persistence Images (PersImg) (Adams et al., 2017):** PersImg treats the persistence diagram as a 2D distribution. It projects each persistence point (b, d) onto a grid using a Gaussian kernel, effectively creating a heatmap or "image" that captures the geometric density and spatial relationships of topological features.

(2) **Persistence Entropy (PersEntropy) (Atienza et al., 2020):** PersEntropy provides a statistical summary of the barcode. It calculates the Shannon entropy (Shannon, 1948) of the distribution of feature lifetimes (the lengths of the

Model	Measure	AUCROC	Accuracy	TPR @ 5% FPR	F1 Score
Llama-2-7b	Self-Prompt	50.30	50.30	-	66.53
	FAVA Model	53.29	53.29	-	43.88
	SelfCheckGPT-Prompt	50.08	54.19	-	67.24
	INSIDE	59.03	57.98	13.17	39.66
	LLM-Check (Attn Score)	72.34	67.96	14.97	69.27
	PersImg	82.09	75.00	26.79	80.67
	PersEntropy	75.67	72.82	21.43	78.26
	Betti Curve	74.75	68.47	23.21	77.86
Llama-3-8b	LLM-Check (Attn Score)	68.19	65.87	15.57	70.53
	PersImg	82.64	73.91	46.43	80.33
	PersEntropy	74.09	71.73	19.29	78.69
	Betti Curve	74.33	70.65	21.43	77.69
Vicuna-7b	LLM-Check (Attn Score)	71.69	66.47	24.55	62.00
	PersImg	83.28	77.17	35.71	82.35
	PersEntropy	75.47	68.47	17.86	76.03
	Betti Curve	76.71	70.65	24.59	80.29

Table 1: **Hallucination detection results on the FAVA Annotated Dataset.** The LLM-Check (Attn Score), Self-Prompt, FAVA Model, SelfCheckGPT-Prompt and INSIDE results are according to the numbers in [Sriraman et al. \(2024\)](#).

bars or equivalently, $(d - b)$ for a point (b, d) in the persistence diagram). A single entropy value quantifies the overall complexity of the topological signature, with higher values indicating more uniform persistence across features.

(3) Betti Curve (Umeda, 2017): This method produces a 1D vector by plotting the Betti number, the count of currently active bars, as a function of the filtration value. The resulting curve, which tracks the feature counts, is then sampled at discrete points to form the feature vector.

For simplicity, we refer to our overall approach as HalluZig, regardless of the specific vectorization scheme employed.

4 Experiments

4.1 Experimental Setup

Datasets. To ensure comprehensive evaluation, we test HalluZig on a diverse suite of benchmarks covering different domains and annotation styles:

Generative Benchmarks: We use two benchmarks with explicit human provided hallucination labels. The FAVA Annotated Dataset ([Mishra et al., 2024](#)) provides passage-level binary labels for Wikipedia abstract generation. The RAGTruth Summarization Dataset ([Niu et al., 2024](#)) offers span-level annotations, which we normalize to passage-level task: a summary is considered if it contains one or more annotated spans.

QA-based Benchmarks: We assess the performance on TruthfulQA ([Lin et al., 2022](#)) and NQ-Open ([Kwiatkowski et al., 2019](#)) datasets. As these

datasets do not contain explicit hallucination labels, we employ the LLM-as-a-judge paradigm ([Zheng et al., 2023](#)). We use GPT-4o-mini ([Hurst et al., 2024](#)), a closed-source LLM, to automatically annotate whether a generated answer is a hallucination or not.

Models and metrics. For our experiments, we use a diverse set of open-source LLMs, including models from the Llama family (Llama-2-7b, Llama-2-13b ([Touvron et al., 2023](#)), Llama-3-8b, Llama-3.1-8b, Llama-3.2-3b ([Grattafiori et al., 2024](#))), Vicuna-7b ([Chiang et al., 2023](#)) and Mistral-7b ([Jiang et al., 2023](#)), all accessed via the Hugging Face transformers library ([Wolf et al., 2020](#)). Our topological pipeline is implemented using the FastZigzag library ([Dey and Hou, 2022](#)) for persistence computation and the Gudhi library ([Carrière et al., 2025](#)) for vectorization. We employ a Random Forest Classifier for the final binary classification and evaluate its performance using Accuracy, F1-Score, AUC-ROC, and TPR@5%FPR. For clarity in our main results, we report the performance of the best run across multiple random seeds. To ensure full reproducibility, we provide the mean and standard deviation for all experiments, along with a detailed list of all hyperparameters, in Appendix B. Moreover, all the results reported in the paper are the result of a single LLM run of the respective models. The code is available at <https://github.com/TDA-Jyamiti/halluzig>

Baselines. Since our method leverages structural signals embedded in attention matrices, we eval-

uate our approach against baselines that capture analogous information. On generative benchmarks (FAVA Annotated Dataset and RAGTruth Summarization Dataset), we compare against the Attention Score metric from the LLM-Check (Sriramanan et al., 2024). To provide additional context on where HalluZig stands with respect to other uncertainty quantification methods, we add Self-Prompt (Kadavath et al., 2022), INSIDE (Chen et al., 2024), SelfCheckGPT-Prompt (Manakul et al., 2023) and FAVA Model (Mishra et al., 2024) to the comparison. For the QA benchmarks, we add LapEigVals (Binkowski et al., 2025) to the comparison as it extracts spectral properties from attention matrices.

4.2 Main Results

HalluZig outperforms the baselines across models on most datasets. The key empirical findings are summarized below.

Method	Metric	Black-box		
		Llama-2-7b	Llama-2-13b	Mistral-7b
Attention Score	AUCROC	54.19	60.05	55.37
	Accuracy	54.52	59.66	56.99
	TPR @ 5% FPR	5.88	14.48	5.18
	F1 Score	54.50	55.97	57.72
PersImg	AUCROC	73.37	72.90	74.45
	Accuracy	62.44	61.63	63.45
	TPR @ 5% FPR	19.03	18.39	26.22
	F1 Score	69.10	68.63	70.84
PersEntropy	AUCROC	51.46	52.78	52.97
	Accuracy	52.00	53.95	52.02
	TPR @ 5% FPR	7.96	2.69	3.11
	F1 Score	60.15	61.22	58.37
Betti Curve	AUCROC	69.26	67.93	71.03
	Accuracy	57.78	60.05	63.00
	TPR @ 5% FPR	23.89	14.80	19.11
	F1 Score	65.95	67.99	70.69

Table 2: **Hallucination detection performance on the RAGTruth Summarization Dataset.** We compare HalluZig (PersImg, PersEntropy, Betti Curve) with Attention Score from (Sriramanan et al., 2024). Results are reported for both a white-box setting (using the Llama-2-7b generator) and a black-box setting (using Llama-2-13b and Mistral-7b as substitutes).

Performance on Generative Benchmarks. On the FAVA Annotated (Mishra et al., 2024) and RAGTruth Summarization Datasets (Niu et al., 2024), HalluZig achieves superior performance (Table 1 and Table 2). A key takeaway is the robustness of the underlying topological signal: all three vectorization schemes (persistence images, Betti curves, and persistent entropy) yield competitive results, confirming that the structural information is the dominant feature. Furthermore, HalluZig excels in both black-box and white-box settings on

RAGTruth Summarization dataset. While performance varies by vectorization, both PersImg and Betti Curve significantly outperform the baseline in both scenarios. This demonstrates that the dynamic structural information captured by HalluZig is a robust signal for hallucination, even when direct model access is limited. Refer to Figure 6 for an illustration.

LLM	Feature	NQOpen	TruthfulQA
Llama-3.1-8b	AttentionScore	0.556	0.541
	AttnEigvals	0.732	0.587
	LapEigvals	0.748	0.589
	PersImg	0.730	0.733
	PersEntropy	0.682	0.664
	Betti Curve	0.715	0.684
Llama-3.2-3b	AttentionScore	0.546	0.581
	AttnEigvals	0.694	0.535
	LapEigvals	0.693	0.539
	PersImg	0.712	0.641
	PersEntropy	0.685	0.656
	Betti Curve	0.688	0.667

Table 3: **Hallucination detection performance on QA Benchmark.** Test AUROC scores for baseline and HalluZig. The results for AttentionScore, AttnEigVals and LapEigvals are based on the experiments we performed with their methods.

Performance on QA Benchmarks. HalluZig’s effectiveness extends to the QA domain (Table 3). HalluZig shows consistent improvement on both, TruthfulQA (Lin et al., 2022) and NQ-Open (Kwiatkowski et al., 2019) datasets, demonstrating that topological features generalize from controlled generative tasks to open-domain question answering.

We note that there is a discrepancy in the LapEigVal score reported in Table 3 and that reported in their original paper (Binkowski et al., 2025). One of the main reasons is that using GPT-4o-mini as annotator produced a heavily class imbalanced dataset, (711/91/15) in our case versus (~500/~250/~80) reported in (Binkowski et al., 2025). Moreover, we used the temperature values as 0.7 for all our experiments while the results reported in (Binkowski et al., 2025) are with 0.1 and 1.0. Our aim for these experiments was only to ensure a fair comparison between HalluZig and other baselines under similar labeling conditions.

We can see from the results that, generally, PersImg has a higher performance than BettiCurve and PersEntropy. Persistence image is a 2D heatmap (distribution) that captures the

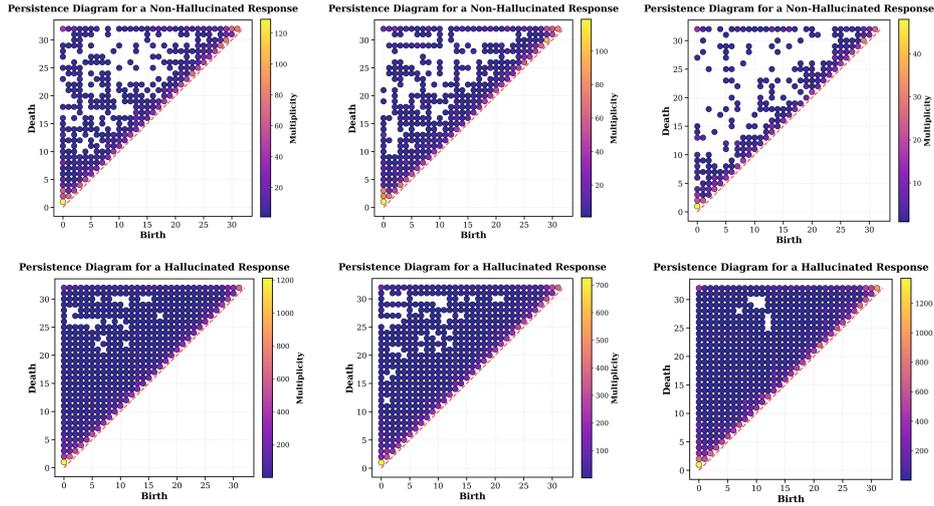


Figure 6: **Visualization of persistence diagrams for hallucinated versus non-hallucinated responses for FAVA dataset.** The top row depicts the persistence diagrams for three randomly selected non-hallucinated responses while the bottom row depicts the persistence diagrams for three randomly selected hallucinated responses from FAVA dataset. The persistence diagrams are colored by multiplicity, i.e., the multiplicity of each point in the diagram is depicted by its color. We can see that the persistence diagrams look visually different for hallucinated versus non-hallucinated responses which gets reflected in the HalluZig performance.

spatial relationship between birth and death layers in a persistence diagram. Betti Curve generally performs the second best which is a 1-D vector of counts of the number of features at different points over the filtration. Persistence Entropy is a single value summary of the persistence diagram. While still useful, quantifying the entire 2D plot (persistence diagram) by a single number inherently loses information, explaining its lower but competitive performance.

5 Analysis and Discussion

The results demonstrate the efficacy of our method and also provide deeper insights into the nature of hallucinations and the potential for new model safety mechanisms. We analyze the implications of our key experimental findings.

5.1 The Critical Role of Dynamics: Zigzag vs. Static TDA

To validate our central hypothesis that the dynamics of attention evolution are more informative than static snapshots, we conducted an ablation study comparing our full model against a Static TDA baseline. This baseline computes standard persistent homology on each layer individually, deliberately ignoring the layer-to-layer transformations.

As shown in Figure 7, the Static TDA baseline underperforms HalluZig by a significant margin. This provides compelling evidence that a substan-

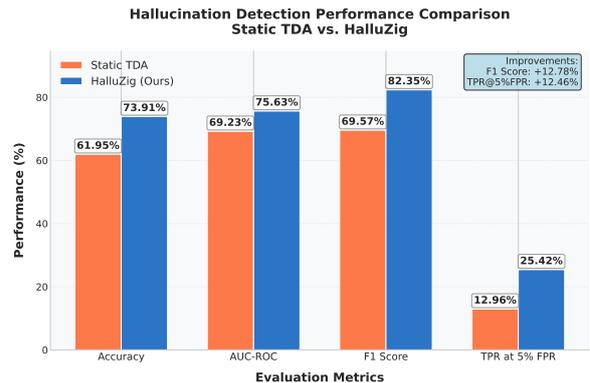


Figure 7: **Role of Dynamics: Comparing HalluZig against a baseline using static, layer-wise persistent homology.** Our full model demonstrates improved performance across all metrics, notably achieving a +12.78% F1-score. The results underscore the importance of modeling the dynamic evolution of attention graphs over static analysis.

tial amount of discriminative information is lost when inter-layer dynamics are disregarded. The key takeaway is that the crucial signal for hallucination lies not in the static topology of individual layers, but in the dynamics of their evolution, a property that zigzag persistence is equipped to capture, which supports our core hypothesis.

5.2 Universality of Signatures: Cross-Model Generalization

A critical question is whether the learned topological signatures are universal or model-specific. We

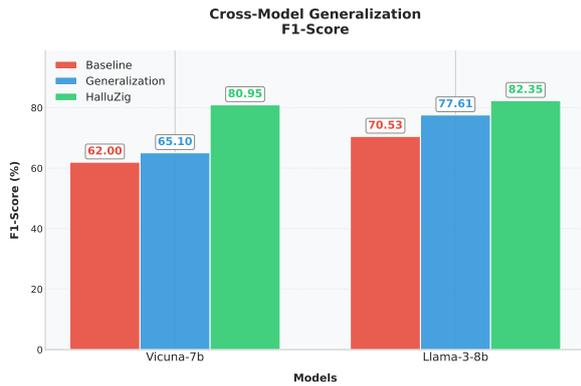


Figure 8: **Zero-Shot Cross-Model Generalization Performance.** The figure evaluates the ability of HalluZig to generalize across different LLM architectures without retraining. (Left) Performance of the classifier on Llama-3-8b trained exclusively on topological signatures extracted from Vicuna-7b. The generalization model outperforms the baseline (Attn-Score (Sriraman et al., 2024)) by 3%. (Right) The reverse scenario, showing performance on Vicuna-7b of the classifier trained on Llama-3-8b. The generalization model surpasses the baseline by 7% in this case.

investigated this via a zero-shot cross-model experiment, training a classifier on topological signatures extracted from Llama-3-8b and testing it on the topological signatures from Vicuna-7b and the vice-versa, without retraining. We use FAVA Annotated Dataset (Mishra et al., 2024) for this analysis. The results reveal a remarkable degree of transferability. This result indicates that different LLMs may exhibit similar topological dynamics when hallucinating.

5.3 Practical Implications: Hallucination Detection from Partial Network Depth

Our final analysis investigates whether hallucinations can be detected reliably before the final layer. We computed topological features (PersImg) up to varying model depths of Vicuna-7b and Llama-3-8b on the FAVA Annotated dataset (Mishra et al., 2024) and evaluated performance at each stage (Figure 9). We observe from the performance curves that HalluZig achieves an F1-score nearly identical to that of the full-model analysis just at 70% of the model’s depth for both models.

This result suggests that hallucination may not be a last-minute failure but is encoded and stabilized relatively early in the model’s reasoning pathway.

Together, these analyses reinforce that HalluZig

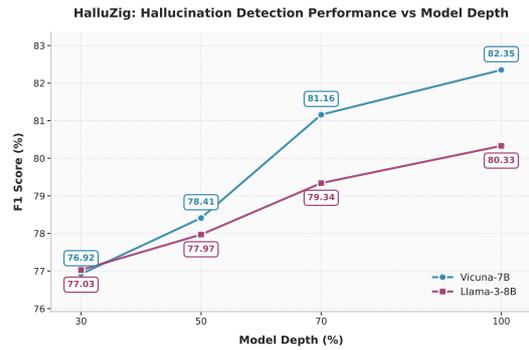


Figure 9: **Hallucination detection performance as a function of model depth.** The plot shows the F1-scores achieved when using topological features from an increasing percentage of the model’s total layers. We observe that HalluZig performance achieves over 98% of the final score by the 70% depth mark. This indicates that the structural signatures of hallucination are formed in the middle layers.

is not only as a diagnostic tool but can also be a safety monitoring mechanism in LLMs.

6 Conclusion

In this paper, we introduced HalluZig, a novel approach for detecting hallucination by analyzing the topological information hidden in the evolution of attention in an LLM. By leveraging zigzag persistence, we demonstrated that the structural evolution of attention provides a powerful and robust signal for hallucination detection. Our experiments show that this method not only outperforms strong baselines but also exhibits remarkable cross-model generalization and enables reliable early detection. This work establishes the viability of structural interpretability, offering a new lens to understand and improve the trustworthiness of LLMs.

Acknowledgements

This work is partially supported by NSF grants DMS-2301360 and CCF-2437030.

Limitations

While our work establishes the viability of topological analysis of evolving attention for hallucination detection, we point out some limitations that can be in agenda for future research.

Computational Complexity

The primary limitation of our current approach is the computational cost of computing persistent homology, particularly for long sequences (sentences)

which result in large attention graphs. However, we note that this computational challenge is not unique to our framework, but is a well-known limitation within the applied TDA community.

Scope of Analysis

Our framework focuses exclusively on the attention mechanism. However, other model components, particularly the MLP layers, are known to store and manipulate factual knowledge. Our approach is currently blind to procedural failures that may originate solely within these components. A more holistic approach could integrate topological signals from both attention and MLP activations.

Attention-Head Averaging

A methodological choice of our framework is the mean-pooling of attention heads to create a single graph per layer. While this provides a stable, holistic view of information flow and ensures computational tractability, it is an important limitation. It is known that attention heads can specialize in distinct functions, and our averaging approach may dilute or obscure a strong, localized signal from a single “rogue” head whose aberrant behavior is the primary cause of a hallucination. Future work could pursue more fine-grained, head-specific topological analysis to gain deeper diagnostic insights, though this would entail a significant increase in computational cost.

Scope of Model Scale

Our experimental validation is conducted on open-source LLMs with up to 13 billion parameters. This leads to the natural question of whether our findings on the topological dynamics of hallucination generalize to much larger, state-of-the-art foundation models (e.g., 70B+ parameters). We hypothesize that the observed structural patterns are a fundamental property of the Transformer architecture and will therefore apply to larger models. However, empirically verifying this scalability is a direction for future research.

Ethical Considerations

To the best of our knowledge, we did not violate any ethical code while conducting the research work described in this paper. We report the technical details needed for reproducing the results and will release the code upon acceptance. All results are from a machine learning model and should be interpreted

as such. The LLMs used to generate attention matrices for this paper are publicly available and are allowed for scientific research.

References

- Henry Adams, Tegan Emerson, Michael Kirby, Rachel Neville, Chris Peterson, Patrick Shipman, Sofya Chepushtanova, Eric Hanson, Francis Motta, and Lori Ziegelmeier. 2017. Persistence images: A stable vector representation of persistent homology. *Journal of Machine Learning Research*, 18(8):1–35.
- Josh Alman, Ran Duan, Virginia Vassilevska Williams, Yinzhan Xu, Zixuan Xu, and Renfei Zhou. 2025. *More Asymmetry Yields Faster Matrix Multiplication*, pages 2005–2039.
- Nieves Atienza, Rocio Gonzalez-D az, and Manuel Soriano-Trigueros. 2020. *On the stability of persistent entropy and new summary functions for topological data analysis*. *Pattern Recognition*, 107:107509.
- Amos Azaria and Tom Mitchell. 2023. The internal state of an LLM knows when it’s lying. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 967–976, Singapore. Association for Computational Linguistics.
- Alexandra Bazarova, Aleksandr Yugay, Andrey Shulga, Alina Ermilova, Andrei Volodichev, Konstantin Polev, Julia Belikova, Rauf Parchiev, Dmitry Simakov, Maxim Savchenko, Andrey Savchenko, Serguei Barannikov, and Alexey Zaytsev. 2025. *Hallucination detection in llms with topological divergence on attention graphs*. *Preprint*, arXiv:2504.10063.
- Jakub Binkowski, Denis Janiak, Albert Sawczyn, Bogdan Gabrys, and Tomasz Kajdanowicz. 2025. *Hallucination detection in llms using spectral features of attention maps*. *Preprint*, arXiv:2502.17598.
- Peter Bubenik. 2015. *Statistical topological data analysis using persistence landscapes*. *J. Mach. Learn. Res.*, 16:77–102.
- Gunnar Carlsson and Vin de Silva. 2010. Zigzag persistence. *Foundations of computational mathematics*, 10(4):367–405.
- Gunnar Carlsson, Vin De Silva, Sara Kaliřnik, and Dmitry Morozov. 2019. Parametrized homology via zigzag persistence. *Algebraic & Geometric Topology*, 19(2):657–700.
- Mathieu Carri re, Gard Spreemann, and Wojciech Reise. 2025. *Persistence representations scikit-learn like interface*. In *GUDHI User and Reference Manual*, 3.11.0 edition. GUDHI Editorial Board.
- Sky CH-Wang, Benjamin Van Durme, Jason Eisner, and Chris Kedzie. 2024. *Do androids know they’re only dreaming of electric sheep?* In *Findings of the Association for Computational Linguistics: ACL 2024*,

- pages 4401–4420, Bangkok, Thailand. Association for Computational Linguistics.
- Chao Chen, Kai Liu, Ze Chen, Yi Gu, Yue Wu, Mingyuan Tao, Zhihang Fu, and Jieping Ye. 2024. [INSIDE: LLMs’ Internal States Retain the Power of Hallucination Detection](#). In *The Twelfth International Conference on Learning Representations*.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality](#).
- Yung-Sung Chuang, Linlu Qiu, Cheng-Yu Hsieh, Ranjay Krishna, Yoon Kim, and James R. Glass. 2024. [Lookback Lens: Detecting and Mitigating Contextual Hallucinations in Large Language Models Using Only Attention Maps](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1419–1436, Miami, Florida, USA. Association for Computational Linguistics.
- Tamal K. Dey and Tao Hou. 2021. [Computing zigzag persistence on graphs in near-linear time](#). In *37th International Symposium on Computational Geometry, SoCG 2021, June 7-11, 2021, Buffalo, NY, USA (Virtual Conference)*, volume 189 of *LIPICs*, pages 30:1–30:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Tamal K. Dey and Tao Hou. 2022. [Fast Computation of Zigzag Persistence](#). In *30th Annual European Symposium on Algorithms (ESA 2022)*, volume 244 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 43:1–43:15, Dagstuhl, Germany. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- Tamal K. Dey and Yusu Wang. 2022. [Computational Topology for Data Analysis](#). Cambridge University Press.
- Xuefeng Du, Chaowei Xiao, and Yixuan Li. 2024. [HaloScope: Harnessing Unlabeled LLM Generations for Hallucination Detection](#).
- Herbert Edelsbrunner and John Harer. 2010. [Computational Topology: An Introduction](#). Applied Mathematics. American Mathematical Society.
- Ekaterina Fadeeva, Aleksandr Rubashevskii, Artem Shelmanov, Sergey Petrakov, Haonan Li, Hamdy Mubarak, Evgenii Tsymbalov, Gleb Kuzmin, Alexander Panchenko, Timothy Baldwin, and 1 others. 2024. [Fact-checking the output of large language models via token-level uncertainty quantification](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 9367–9385.
- Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. 2024. [Detecting hallucinations in large language models using semantic entropy](#). *Nature*, 630(8017):625–630.
- Peter Gabriel. 1972. [Unzerlegbare darstellungen](#). *Manuscripta Mathematica*, 6(1):71–103.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2025. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#). *ACM Trans. Inf. Syst.*, 43(2).
- Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Madry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino, and 399 others. 2024. [Gpt-4o system card](#). *Preprint*, arXiv:2410.21276.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L el io Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth ee Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El-Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, and 17 others. 2022. [Language models \(mostly\) know what they know](#). *Preprint*, arXiv:2207.05221.
- Jannik Kossen, Jiatong Han, Muhammed Razzak, Lisa Schut, Shreshth Malik, and Yarin Gal. 2024. [Semantic Entropy Probes: Robust and Cheap Hallucination Detection in LLMs](#). *arXiv preprint*. ArXiv:2406.15927 [cs].
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. [Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation](#). In *The Eleventh International Conference on Learning Representations*.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural Questions: A Benchmark for Question Answering Research](#). *Transactions of the Association for*

- Computational Linguistics*, 7:452–466. Place: Cambridge, MA Publisher: MIT Press.
- 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)*, pages 3214–3252.
- Andrey Malinin and Mark Gales. 2021. [Uncertainty estimation in autoregressive structured prediction](#). In *International Conference on Learning Representations*.
- Potsawee Manakul, Adian Liusie, and Mark Gales. 2023. [SelfcheckGPT: Zero-resource black-box hallucination detection for generative large language models](#). In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Clément Maria and Steve Oudot. 2016. Computing zigzag persistent cohomology. *arXiv preprint arXiv:1608.06039*.
- Abhika Mishra, Akari Asai, Vidhisha Balachandran, Yizhong Wang, Graham Neubig, Yulia Tsvetkov, and Hannaneh Hajishirzi. 2024. [Fine-grained hallucination detection and editing for language models](#). In *First Conference on Language Modeling*.
- Alexander Nikitin, Jannik Kossen, Yarin Gal, and Pekka Marttinen. 2024. Kernel language entropy: Fine-grained uncertainty quantification for llms from semantic similarities. *Advances in Neural Information Processing Systems*, 37:8901–8929.
- Cheng Niu, Yuanhao Wu, Juno Zhu, Siliang Xu, Kashun Shum, Randy Zhong, Juntong Song, and Tong Zhang. 2024. Ragtruth: A hallucination corpus for developing trustworthy retrieval-augmented language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10862–10878.
- Hadas Orgad, Michael Tokar, Zorik Gekhman, Roi Reichart, Idan Szepke, Hadas Kotek, and Yonatan Belinkov. 2025. [LLMs Know More Than They Show: On the Intrinsic Representation of LLM Hallucinations](#). In *The Thirteenth International Conference on Learning Representations*.
- Xin Qiu and Risto Miikkulainen. 2024. Semantic density: Uncertainty quantification for large language models through confidence measurement in semantic space. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Pranab Sahoo, Prabhath Meharia, Akash Ghosh, Sriparna Saha, Vinija Jain, and Aman Chadha. 2024. A comprehensive survey of hallucination in large language, image, video and audio foundation models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 11709–11724.
- Claude Elwood Shannon. 1948. [A mathematical theory of communication](#). *The Bell System Technical Journal*, 27:379–423.
- Gaurang Sriramanan, Siddhant Bharti, Vinu Sankar Sadasivan, Shoumik Saha, Priyatham Kattakinda, and Soheil Feizi. 2024. [LLM-check: Investigating detection of hallucinations in large language models](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, and 1 others. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Yuhei Umeda. 2017. [Time series classification via topological data analysis](#). *Transactions of the Japanese Society for Artificial Intelligence*, 32:D–G72_1.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Yanling Wang, Haoyang Li, Hao Zou, Jing Zhang, Xinlei He, Qi Li, and Ke Xu. 2024. [Hidden Question Representations Tell Non-Factuality Within and Across Large Language Models](#). *arXiv preprint. ArXiv:2406.05328 [cs]*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. [Transformers: State-of-the-Art Natural Language Processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Tianhang Zhang, Lin Qiu, Qipeng Guo, Cheng Deng, Yue Zhang, Zheng Zhang, Chenghu Zhou, Xinbing Wang, and Luoyi Fu. 2023. [Enhancing Uncertainty-Based Hallucination Detection with Stronger Focus](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 915–932, Singapore. Association for Computational Linguistics.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-judge with MT-bench and Chatbot Arena. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS ’23*, Red Hook, NY, USA. Curran Associates Inc. Event-place: New Orleans, LA, USA.
- Xiaoling Zhou, Mingjie Zhang, Zhemg Lee, Wei Ye, and Shikun Zhang. 2025. [Hademif: Hallucination detection and mitigation in large language models](#). In *The Thirteenth International Conference on Learning Representations*.

Afra Zomorodian and Gunnar Carlsson. 2004. **Computing persistent homology**. In *Proceedings of the Twentieth Annual Symposium on Computational Geometry*, SCG '04, page 347–356, New York, NY, USA. Association for Computing Machinery.

A Topological Data Analysis Preliminaries

This appendix provides the necessary background in topological data analysis, with particular emphasis on concepts relevant to zigzag persistence and its application to graphs.

A.1 Simplicial Homology

We begin by introducing simplicial complexes. Simplicial complexes are spaces built with smaller geometric objects (simplices), such as vertices, edges, filled triangles, and so on. More formally:

An *abstract simplicial complex* K is a family of non-empty subsets of an underlying finite set $V(K)$ that is closed under the operation of taking subsets. In other words, if $\sigma \in K$ and $\tau \subset \sigma$ then $\tau \in K$. Each such element $\sigma \in K$, is called a p -simplex if $|\sigma| = p + 1$. A $\tau \subset \sigma$ is called a face of σ .

A graph $G = (V, E)$ is easily seen as abstract simplicial complex, where the underlying set is the set of vertices $V = V(K)$, and the family of subsets K is given by the union of set of vertices and the set of edges. Then the vertices correspond to the 0-simplices and the set of edges to the 1-simplices.

For an abstract simplicial complex K , define for each $p \geq 0$ a *chain group* $C_p(K, \mathbb{Z}_2)$, which is the vector space over \mathbb{Z}_2 generated by the p -simplices of K . We usually omit the field \mathbb{Z}_2 in our notation. Elements of $C_p(K)$ are called *p -chains*, and they are formal sums of p -simplices with \mathbb{Z}_2 coefficients. One can connect $C_p(K)$ and $C_{p-1}(K)$ by a boundary operator $\partial_p : C_p(K) \rightarrow C_{p-1}(K)$, which maps each p -simplex to the sum of its $(p-1)$ -dimensional faces. These boundary operators satisfy the key property $\partial_{p-1} \circ \partial_p = 0$.

Using these boundary maps, define:

- The group of p -cycles:

$$Z_p(K) = \ker(\partial_p),$$

which consists of p -chains with zero boundary (closed cycles).

- The group of p -boundaries:

$$B_p(K) = \text{im}(\partial_{p+1}),$$

which consists of p -chains that are themselves boundaries of $(p + 1)$ -chains.

The p -th *simplicial homology group* is then defined as the quotient

$$H_p(K) = Z_p(K)/B_p(K),$$

which measures the p -cycles modulo those that bound higher-dimensional simplices. Intuitively, elements of $H_p(K)$ correspond to p -dimensional “holes” in the space. Since \mathbb{Z}_2 is a field, each homology group forms a vector space.

A.2 Zigzag Persistence

Simplicial homology provides a powerful tool for analyzing the topology of a single, static simplicial complex. However, real-world data is often not a single static object but a sequence of related data sampled at different moments or corresponding to different parameters of evolution. In order to study these evolving objects, we define a *zigzag filtration*. A zigzag filtration \mathcal{Z} is a sequence of simplicial complexes

$$\mathcal{Z} : K_0 \subseteq K_1 \supseteq K_2 \subseteq \dots \supseteq K_n.$$

where the sequence of inclusions could be both in forward and backward directions.

The homology group of K_i is a vector space $H_p(K_i)$ (under a field coefficient like \mathbb{Z}_2) for each $i \in \{1, \dots, n\}$. The inclusion map between K_i and K_j , induces a natural linear map between $H_p(K_i)$ and $H_p(K_j)$. By assembling these vector spaces and linear maps together, we obtain the *zigzag persistence module* $M_{\mathcal{Z}}$ defined by

$$M_{\mathcal{Z}} : H_p(K_0) \rightarrow H_p(K_1) \leftarrow \dots \leftarrow H_p(K_n)$$

The Interval Decomposition Theorem for zigzag persistence modules (Carlsson and de Silva, 2010; Gabriel, 1972), states that $M_{\mathcal{Z}}$ decomposes uniquely (up to reordering and isomorphism) into interval modules:

$$M_{\mathcal{Z}} \cong \bigoplus_k I[b_k, d_k]$$

where each $I[b_k, d_k]$ is an interval module with the support on the interval $[b_k, d_k]$, called a bar. The collection of these bars yields the *zigzag barcode*, defined as the multiset of pairs $\{[b_k, d_k]\}_k$. This barcode carries the topological information

present in the zigzag persistence module, which we leverage for hallucination detection.

One of the important properties of zigzag persistence is its stability under small perturbations (Carlsson and de Silva, 2010), implying that minor changes in attention weights between layers result in only small changes in the barcodes. This stability property ensures that the topological signatures we extract are robust to noise.

A.3 Time Complexity

We use the algorithm proposed in (Dey and Hou, 2022) to compute zigzag persistence barcodes. The authors show that the time complexity for computing zigzag persistent homology is the same as the time complexity for computing standard persistent homology. The time complexity for computing standard persistent homology is $O(n^\omega)$, where n is the number of simplices (Zomorodian and Carlsson, 2004) and $\omega < 2.371339$ is the matrix multiplication exponent (Alman et al., 2025). However, for the special case of graphs there exists a near linear time algorithm (Dey and Hou, 2021) that could potentially be used.

B Additional Experimental Details

To ensure the robustness of our method, we conducted several experiments to select the optimal hyperparameters for our main results. We use an NVIDIA A100 (40GB) GPU with 16 cores for all the experiments.

B.1 Hyperparameter Tuning

Minimum Persistence Filtering. As is well-known in the applied TDA community, shorter bars are considered to be noise. Hence, we filter out shorter bars and retain the longer ones. In order to choose the optimal threshold, we evaluate the performance after filtering out bars below various persistence thresholds (5, 7, 9 or 11 layers) at two different levels of graph sparsity (edge selection) - 30% (Table 5) and 10% (Table 6). Based on our results, we find that a moderate level of persistence filtering is optimal. Therefore, to balance the benefit of removing noise against the risk of discarding valuable signals, we select a conservative minimum persistence threshold of 5 for all main experiments.

Edge Selection Threshold. The construction of our attention graphs depends on a sparsity parameter. We investigated the impact of this by varying the percentage of top attention weights retained

(5%, 10%, 20%, and 30%) for different minimum persistence values - 5 (Table 7) and 9 (Table 8). The results indicate that retaining the top 10% of edges provides a strong balance between detection performance and computational efficiency. We therefore use this 10% threshold for all main experiments reported in the paper.

B.2 Implementation Details

Model Temperature. We used a temperature value of 0.7 for all the models while generating answers. For GPT-4o-mini, we used a temperature of 0 while generating the annotation labels on QA Benchmarks.

Topological Feature Vectorization. We use GUDHI (Carrière et al., 2025) for all vectorization schemes. The specific parameters are as follows:

- **Persistence Images (PersImg):** We use a resolution of 32×32 .
- **Betti Curves:** We use a sampling resolution of 32 points.
- **Persistence Entropy (PersEntropy):** We use the default implementation from the library.

Classifier. For all our experiments, we classify using a random forest classifier. For experiments on RAGTruth Dataset, we use `n_estimators = 500`. For all other experiments, we use `n_estimators = 100`. The `max_depth` parameter was tuned for each configuration. The values are reported in Table 4.

Baseline Reproduction. To ensure a fair comparison with LapEigVals baseline from (Binkowski et al., 2025), we followed their experimental protocol precisely. As specified in their work, we retained the top $k = 50$ eigenvalues across all heads and layers and trained a logistic regression classifier with their reported parameters.

B.3 Detailed Performance Metrics and Statistical Uncertainty

To ensure reproducibility of our results, this section presents the detailed performance metrics for HalluZig across all datasets. The main results reported in the body of the paper correspond to the single best run for clarity. Here, we report the mean and standard deviation calculated over five independent runs with different random seeds. Refer to Table 10 and Table 9. For the RAGTruth

dataset, we use the official train-test split provided in the dataset. Consequently, all results reported in the main text are based on this single, pre-defined partition. For all other datasets, we use a 80-20 train-test split.

Dataset	Model	PersImg	PersEntropy	BettiCurve
FAVA Annotated	Llama-2-7b	20	10	10
	Llama-3-8b	20	10	10
	Vicuna-7b	20	10	10
RAGTruth	Llama-2-7b	25	25	25
	Llama-2-13b	25	25	25
	Mistral-7b	25	25	25
NQOpen	Llama-3.1-8b	6	4	5
	Llama-3.2-3b	5	4	4
TruthfulQA	Llama-3.1-8b	6	5	5
	Llama-3.2-3b	6	4	5

Table 4: Maximum tree depth (max_depth) used in Random Forest Classifier.

Method	Metric	Min. Persistence of Bars			
		≥ 5	≥ 7	≥ 9	≥ 11
PersImg	AUC-ROC	73.46	73.27	72.93	73.76
	Accuracy	65.22	68.48	65.22	64.13
	TPR at 5% FPR	7.41	7.41	7.41	7.41
	F1 Score	72.88	75.63	72.88	72.27
PersEntropy	AUC-ROC	75.41	77.75	74.44	69.98
	Accuracy	66.30	68.48	65.22	65.22
	TPR at 5% FPR	33.33	25.93	24.07	22.22
	F1 Score	71.56	74.34	70.91	72.41
Betti Curve	AUC-ROC	75.78	73.73	72.59	74.29
	Accuracy	69.57	67.39	64.13	65.22
	TPR at 5% FPR	27.78	29.63	29.63	37.04
	F1 Score	75.86	74.14	71.30	73.77

Table 5: This table evaluates our three vectorization schemes when filtering out topological features (bars) with low persistence. We select the top 30% of the edges for this experiment.

Method	Metric	Min. Persistence of Bars			
		≥ 5	≥ 7	≥ 9	≥ 11
PersImg	AUC-ROC	73.42	73.08	72.34	74.10
	Accuracy	66.30	64.13	64.13	68.48
	TPR at 5% FPR	12.96	25.93	14.81	20.37
	F1 Score	73.95	71.79	71.79	76.03
PersEntropy	AUC-ROC	66.96	57.36	59.23	64.91
	Accuracy	65.22	64.13	60.87	65.22
	TPR at 5% FPR	7.41	7.41	14.81	14.81
	F1 Score	70.91	70.80	67.27	71.93
Betti Curve	AUC-ROC	69.96	70.49	71.03	72.69
	Accuracy	66.30	65.22	64.13	65.22
	TPR at 5% FPR	12.96	14.81	20.37	18.52
	F1 Score	73.50	72.41	71.79	72.88

Table 6: This table evaluates our three vectorization schemes when filtering out topological features (bars) with low persistence. We select the top 10% of the edges for this experiment.

Method	Metric	Top % of edges (attn weights) selected			
		30%	20%	10%	5%
PersImg	AUC-ROC	73.46	72.83	73.42	71.47
	Accuracy	65.22	66.30	66.30	65.22
	TPR at 5% FPR	7.41	20.37	12.96	9.26
	F1 Score	72.88	73.50	73.95	72.88
PersEntropy	AUC-ROC	75.41	74.39	66.96	74.78
	Accuracy	66.30	66.30	65.22	67.39
	TPR at 5% FPR	33.33	25.93	7.41	22.22
	F1 Score	71.56	71.56	70.91	74.14
Betti Curve	AUC-ROC	75.78	72.78	69.96	74.17
	Accuracy	69.57	66.30	66.30	65.22
	TPR at 5% FPR	27.78	31.48	12.96	38.89
	F1 Score	75.86	73.95	73.50	71.93

Table 7: This table shows the performance of our three vectorization schemes while varying the percentage of top attention weights used to construct the attention graphs (minimum bar persistence is held constant at 5).

Method	Metric	Top % of edges (attn weights) selected			
		30%	20%	10%	5%
PersImg	AUC-ROC	72.34	72.93	72.39	72.00
	Accuracy	64.13	65.22	66.30	65.22
	TPR at 5% FPR	14.81	7.41	16.67	18.52
	F1 Score	71.79	72.88	73.95	72.88
PersEntropy	AUC-ROC	74.44	71.08	59.23	74.98
	Accuracy	65.22	64.13	60.87	67.39
	TPR at 5% FPR	24.07	14.81	14.81	25.93
	F1 Score	70.91	71.30	67.27	73.21
Betti Curve	AUC-ROC	72.59	72.39	71.03	72.64
	Accuracy	64.13	66.30	64.13	66.30
	TPR at 5% FPR	29.63	22.22	20.37	33.33
	F1 Score	71.30	73.95	71.79	72.07

Table 8: This table shows the performance of our three vectorization schemes while varying the percentage of top attention weights used to construct the attention graphs (minimum bar persistence is held constant at 9).

Method	TruthfulQA (AUC-ROC)		NQOpen (AUC-ROC)	
	Llama-3.1-8b	Llama-3.2-3b	Llama-3.1-8b	Llama-3.2-3b
PersImg	65.92 \pm 4.05	62.08 \pm 2.32	68.29 \pm 2.58	67.27 \pm 2.16
PersEntropy	63.62 \pm 3.07	58.34 \pm 6.61	66.11 \pm 1.36	65.53 \pm 1.68
BettiCurve	59.31 \pm 6.11	58.22 \pm 4.74	67.97 \pm 2.92	65.86 \pm 2.07

Table 9: Each entry reports mean \pm standard deviation across five random seeds for Test AUC-ROC scores.

Metric	PersImg			PersEntropy			Betti Curve		
	Llama-2-7b	Llama-3-8b	Vicuna-7b	Llama-2-7b	Llama-3-8b	Vicuna-7b	Llama-2-7b	Llama-3-8b	Vicuna-7b
AUC-ROC	78.72 ± 3.41	78.89 ± 2.74	79.20 ± 3.38	70.35 ± 4.53	72.29 ± 4.34	71.31 ± 3.63	71.64 ± 3.75	72.71 ± 4.90	72.95 ± 2.48
Accuracy	73.04 ± 2.70	74.35 ± 2.24	74.13 ± 2.70	67.83 ± 2.54	71.52 ± 2.52	69.57 ± 3.30	67.83 ± 3.12	71.30 ± 2.54	70.65 ± 2.75
TPR @ 5% FPR	30.02 ± 12.17	37.70 ± 11.57	36.62 ± 12.25	19.86 ± 3.70	16.97 ± 4.61	25.55 ± 8.08	21.41 ± 3.95	16.45 ± 6.69	22.76 ± 8.28
F1 Score	80.21 ± 1.66	81.42 ± 1.77	81.16 ± 1.70	76.77 ± 2.16	78.15 ± 2.33	76.72 ± 2.53	77.05 ± 2.54	79.45 ± 1.60	78.96 ± 2.30

Table 10: Each entry reports mean \pm standard deviation across five random seeds for all the metrics on the FAVA Annotated Dataset.

C QA prompt

We used OpenAI's API and ChatGPT-4o-mini model as judge for labeling the generated response of each LLM in the whitebox setting. The prompt we used is the same for NQopen, and TruthfullQA.

Prompt used for LLM-as-judge with Open AI's ChatGPT4o-mini

You are an evaluator for QA tasks. Given the question, the reference best answer, and a model's answer, classify the model's answer as:

- CORRECT if it matches or is equivalent to the best answer.
 - HALLUCINATION if it introduces facts not in the best answer.
 - N/A if the model expresses uncertainty or says it cannot answer.
- Respond with only one label: CORRECT, HALLUCINATION, or N/A.

Question: {question}

Reference best answer: {best_answer}

Model answer: {model_answer}