arXiv:2603.16459v2 [cs.CL] 24 Aug 2026

DynHD: Hallucination Detection for Diffusion Large Language Models via Denoising Dynamics Deviation Learning

Yanyu Qian[1], Yue Tan[2], Yixin Liu[2*], Wang Yu[1], Shirui Pan[2]

[1]Nanyang Technological University, Singapore [2]Griffith University, Australia

[*]Corresponding author: yixin.liu@griffith.edu.au

Abstract

Diffusion large language models (D-LLMs) have emerged as a promising alternative to auto-regressive models due to their iterative refinement capabilities. However, hallucinations remain a critical issue that hinders their reliability. To detect hallucination responses from model outputs, token-level uncertainty (e.g., entropy) has been widely used to indicate potential factual errors. Nevertheless, unlike auto-regressive models that generate tokens sequentially, D-LLMs generate fixed-length sequences simultaneously, where only a small subset of tokens is informative for hallucination detection. Thus, aggregating uncertainty over all tokens can be suboptimal. Moreover, the evolution trend of uncertainty throughout the diffusion process can also provide valuable signals, highlighting the necessity of modeling its denoising dynamics for hallucination detection. In this paper, we propose DynHD that bridges these gaps from both *spatial* (token sequence) and *temporal* (denoising dynamics) perspectives. To handle the information density imbalance across tokens, we propose a semantic-aware evidence construction module that extracts hallucination-indicative signals by removing task-invariant structural tokens and emphasizing the uncertainty of the remaining informative tokens. To model denoising dynamics for hallucination detection, we introduce a reference evidence generator that learns the expected evolution trajectory of uncertainty evidence, along with a deviation-based hallucination detector that makes predictions by measuring the discrepancy between the observed and reference trajectories. Extensive experiments demonstrate that DynHD consistently outperforms state-of-the-art baselines while achieving higher efficiency across multiple benchmarks and backbone models. The code is available at: https://github.com/qyy11-com/DynHD.

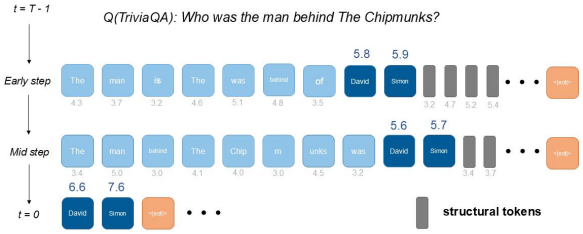

Figure 1: Visualization of spatial uncertainty distribution during decoding. Tokens exhibiting the highest entropy spikes serve as the primary indicators of factual instability. On the contrary, structural tokens provide limited cues for hallucination detection.

1 Introduction

Diffusion large language models (D-LLMs) (Nie et al., 2025; Ye et al., 2025) have recently received extensive attention as a significant generative paradigm alternative to autoregressive large language models (AR-LLMs). Unlike AR-LLMs using token-by-token decoding, D-LLMs generate fixed-length sequences through multi-step iterative denoising, performing discrete remasking at each step: the model predicts a denoised intermediate sequence while selectively remasking a portion of positions for further optimization (Austin et al., 2021; Zhao et al., 2026). This paradigm offers advantages in global consistency and the ability to gradually refine the generated content (Li et al., 2022; Huang et al., 2024).

While D-LLMs offer several advantages, they still suffer from hallucinations, i.e., generating factually incorrect or unsupported content, which limits their reliability in real-world applications (Ji et al., 2025; Liu et al., 2026b). To mitigate this issue, hallucination detection has become a critical task for ensuring the factual reliability of D-LLMs. Nevertheless, although extensive efforts have been devoted to the AR-LLMs setting, these hallucination detection approaches do not readily generalize to D-LLMs due to their fundamentally different

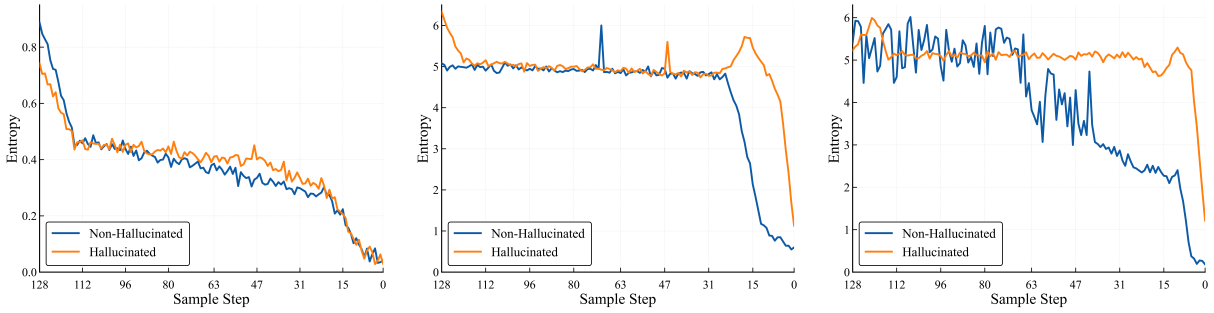

Figure 2: Entropy dynamics under spatial filtering. (a) Mean trajectories fail to distinguish between samples due to information density imbalance. (b) & (c) Filtered trajectories (top-K) reveal distinct convergence: factual samples decay monotonically, while hallucinated ones exhibit late-stage stagnation or rebound. Differences in decay rates between HotpotQA (b) and TriviaQA (c) reflect varying task-specific denoising dynamics.

architectural and generative paradigms (Yehuda et al., 2024; Tan et al., 2026; Li et al., 2026b). In particular, the key evidence identifying hallucinations in D-LLMs is typically distributed across the dynamic evolution of the entire denoising trajectory, rather than being concentrated in the output confidence of the final step. Consequently, hallucination detectors designed for AR-LLMs, such as those relying on final token probabilities (Han et al., 2024), single-step uncertainty and entropy (Du et al., 2024), or multi-sample consistency (Manakul et al., 2023), often ignore the process-level signals inherent in the denoising trajectory. Moreover, these approaches, which are already computationally expensive for AR-LLMs, become even more time-consuming for D-LLMs, as additional sampling or verification requires repeatedly executing the denoising process (Chern et al., 2025).

As a pioneering work on hallucination detection for D-LLMs, TraceDet (Chang et al., 2026) adopts a trajectory-based perspective by treating the denoising process as observable action traces. It extracts hallucination-relevant sub-trajectories for discrimination and emphasizes the importance of "process signals" over "final-state signals" for hallucination detection in D-LLMs. However, through our in-depth analysis of the hallucination-indicated signals in D-LLMs, TraceDet still exhibits two orthogonal weaknesses: From a spatial perspective, it ignores the **information density imbalance** within generated sequences, failing to explicitly filter out structural padding tokens introduced by fixed-length generation. From a temporal perspective, it overlooks the underlying **denoising dynamics** of entropy-based hallucination evidence, hindering the modeling of uncertainty evolution in the diffusion trajectory.

Firstly, regarding the spatial dimension, D-LLMs naturally generate answers within fixed-length sequences (Li et al., 2022; Chen et al., 2025). In these sequences, the roles of individual tokens are highly heterogeneous, leading to an **information density imbalance** (Gong et al., 2023; Tan et al., 2023), as illustrated in Figure 1. A few tokens are vital for the final answer, providing key signals for hallucination detection. In contrast, many positions are filled with structural padding tokens introduced by fixed-length generation, which carry little hallucination-related information and should be explicitly filtered out.

In the context of such imbalance, TraceDet (Chang et al., 2026) aggregates uncertainty from all tokens as detection evidence without explicitly accounting for their semantic importance, which may drown out hallucination signals with low-informative tokens (Zhang et al., 2024; Pan et al., 2026). As a result, the entropy profiles of hallucinated and factual samples become indistinguishable from each other (see Figure 2a), leading to sub-optimal hallucination detection performance.

Even if hallucination-related evidence can be extracted from token sequences, how to leverage the evidence trajectory over multiple diffusion steps remains another challenge. Through visualization of evidence trajectories, we empirically find that **denoising dynamics**, i.e., the trend of how evidence evolves throughout the diffusion process, are critical for hallucination detection. As shown in Figure 2b and Figure 2c (corresponding examples are in Appendix D.1), factual and hallucinated samples often show different trends during the denoising

process, providing strong discriminative signals. More importantly, the range, trend, and shape of entropy curves often vary significantly across different scenarios and questions (see the difference between Figures 2b and 2c), highlighting the necessity of modeling complex denoising dynamics for robust hallucination detection. Nevertheless, TraceDet selects only a few key steps for hallucination prediction, which may neglect the continuous and holistic dynamics of the diffusion process.

To overcome the above limitations, we propose **DynHD**, a denoising **Dyn**amics deviation learning-based approach for **H**allucination **D**etection of D-LLMs. To address the *information density imbalance*, our method introduces a semantic-aware token filtering module to explicitly remove structural padding tokens introduced by fixed-length generation, and further leverages multivariate statistical modeling of entropy to construct evidence with strong hallucination indications. Moreover, we model denoising dynamics for robust hallucination detection with a denoising dynamical deviation learning module, where a reference evidence dynamics generator is established to model the normal evolution patterns of evidence, and a deviation-based detector identifies hallucinations by measuring the difference between observed evidence dynamics and the reference dynamics.

Extensive experiments on TriviaQA (Joshi et al., 2017), HotpotQA (Yang et al., 2018), and CommonsenseQA (Talmor et al., 2019) with Dream-7B (Ye et al., 2025) and LLaDA-8B (Nie et al., 2025) demonstrate that DynHD achieves state-of-the-art performance. It outperforms TraceDet by an average AUROC margin of 12.2% without requiring external retrieval or computationally expensive repeated sampling.

2 Preliminaries

Diffusion Denoising Mechanism. Diffusion large language models (D-LLMs) generate sequences via an iterative stochastic denoising process. Given a prompt $\mathbf{q}$, the model maintains a discrete state $\mathbf{x}^{(t)} = (x_t^{(1)}, \ldots, x_t^{(l)}) \in \mathcal{V}^l$ across the trajectory $t \in \{T, \ldots, 0\}$, where $\mathcal{V}, l, T$ denote the vocabulary, fixed sequence length, and total denoising steps, respectively. Starting from a highly masked initialization $\mathbf{x}^{(T)}$, the model iteratively fills masked positions to reach the final output $\mathbf{x}^{(0)}$.

At each step t, the model predicts a categorical distribution $\pi_i^{(t)} \in \mathbb{R}^{|\mathcal{V}|}$ for each position i via a softmax over the generated logits. We quantify the position-wise prediction uncertainty using the token entropy:

$$s_{t,i} = -\sum_{v \in \mathcal{V}} \pi_i^{(t)}(v) \log \pi_i^{(t)}(v). \tag{1}$$

Let $\mathcal{S}_t = \{s_{t,i}\}_{i=1}^{l}$ represent the sequence of token entropies at step t. The complete raw uncertainty trajectory spanning all descending denoising steps is thus denoted as $\mathcal{T} = (\mathcal{S}_T, \mathcal{S}_{T-1}, \ldots, \mathcal{S}_0)$. Notably, due to the fixed length l, sequences that terminate early are padded with non-semantic structural tokens ($\langle|endoftext|\rangle$), which introduces significant noise into $\mathcal{T}$.

Definition of Hallucination Detection Problem. We formulate hallucination detection as a binary classification task for a question-response pair $(\mathbf{q}, \mathbf{r})$, where $\mathbf{r} = \mathbf{x}^{(0)}$. Given a training set $\mathcal{D} = \{(\mathbf{q}_n, \mathbf{r}_n, y_n)\}_{n=1}^{N}$ with labels, where $y = 1$ indicates a hallucinated response and $y = 0$ denotes a factually correct one. Our goal is to learn an optimal detector h_ϕ (e.g., a neural network parameterized by ϕ) that minimizes the empirical risk:

$$\begin{aligned} \min_{\phi} \quad & \mathbb{E}_{(\mathbf{q},\mathbf{r},y)\sim\mathcal{D}} \left[\mathcal{L}(y, h_\phi(\mathbf{q}, \mathbf{r}, \mathcal{T}))\right], \\ \text{s.t.} \quad & \mathbf{x}^{(t-1)} \sim p_\theta(\mathbf{x}^{(t-1)} | \mathbf{x}^{(t)}, \mathbf{q}), \end{aligned} \tag{2}$$

where $\mathcal{L}$ is a classification loss function, and the detector makes its decision based on the question, the final response, and the collected uncertainty trajectory $\mathcal{T}$.

The related work is summarized in Appendix A.

3 Methodology

In this section, we introduce DynHD that detects hallucinations by learning deviations in denoising dynamics. As illustrated in Figure 3, DynHD starts with a *semantic-aware evidence construction* module (Sec. 3.1). To handle the information density imbalance issue, it first filters out structural padding tokens introduced by fixed-length generation, and then constructs the hallucination evidence by statistically aggregating the uncertainty of the remaining tokens. After that, a *dynamical deviation learning* module (Sec. 3.2) identifies the hallucination answers by modeling denoising dynamics of D-LLMs. Concretely, we model the continuous denoising process by building a question-conditioned reference path to serve as a baseline, and then compare the

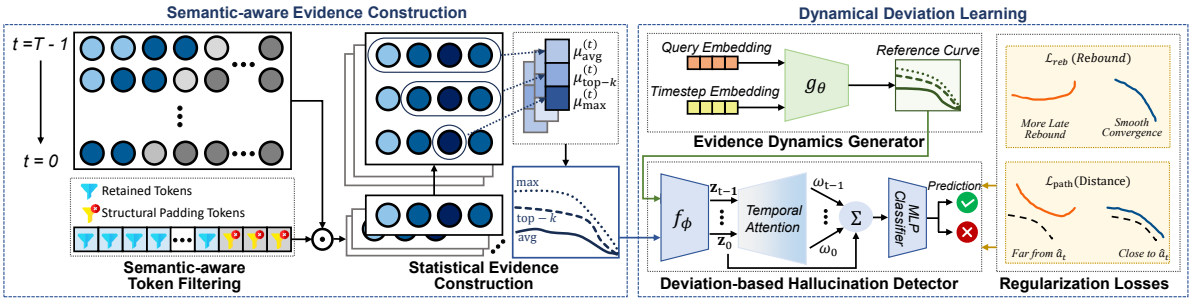

Figure 3: Framework overview of DynHD, which constructs a semantic-aware evidence trajectory from uncertainty, and detects hallucination via learning the deviation between the reference and observed trajectories.

actual evidence trajectory against this established reference for hallucination prediction.

3.1 Semantic-aware Evidence Construction

To identify hallucinations generated by D-LLMs, the first step is to construct hallucination-indicative evidence from the denoising process. While token-level uncertainty (e.g., entropy) can serve as a useful indicator, the fixed-length generation paradigm of D-LLMs introduces an information density imbalance, i.e., not all tokens in a generated sequence provide meaningful signals for hallucination detection. To address this issue, we adopt a two-stage strategy for evidence construction. First, we conservatively filter out task-invariant structural padding tokens introduced by fixed-length generation. Then, we construct statistical evidence by aggregating the entropy of the remaining tokens.

Semantic-aware Token Filtering. Due to the inherent fixed-length generation paradigm of D-LLMs, generated sequences often contain many structural padding tokens, especially when the final answer is short. These tokens are introduced to maintain a fixed sequence length and usually carry little hallucination-related information.

To extract reliable uncertainty evidence, we employ a conservative filtering mechanism that only excludes task-invariant structural tokens from evidence calculation. Specifically, we define a structural token set $\mathcal{I}_{\text{struct}}$, including padding tokens, boundary markers, and empty positions, such as `<|endoftext|>` and `[PAD]`.

At each denoising step t, for token $x_i^{(t)}$ at position $i \in \{1, \ldots, l\}$, we extract the subset of structurally valid positions $\mathcal{K}_t$:

$$\mathcal{K}_t = \{i \mid 1 \le i \le l,\ x_i^{(t)} \notin \mathcal{I}_{\text{struct}}\}. \tag{3}$$

Confining evidence construction to $\mathcal{K}_t$ reduces structural noise caused by fixed-length generation while preserving potentially meaningful lexical and formatting tokens.

Statistical Evidence Construction. After removing structural padding tokens, the remaining tokens in $\mathcal{K}_t$ still contribute differentially to hallucination detection (e.g., factual nouns versus prepositions or pronouns). Consequently, a simple average of these tokens may smooth out localized signals and produce weak evidence. To better capture informative patterns, we adopt a statistics-based method to condense the filtered uncertainty field $\{s_{t,i}\}_{i \in \mathcal{K}_t}$ at each step t into a statistical evidence vector $\mathbf{a}_t \in \mathbb{R}^3$. This captures the mean, peak, and top-k characteristics of the entropy distribution:

$$\mathbf{a}_t = \left[\mu_{\text{avg}}^{(t)},\ s_{\max}^{(t)},\ \mu_{\text{top-}k}^{(t)}\right]. \tag{4}$$

These three statistical moments capture the filtered uncertainty at different granularities: ❶ **Global Uncertainty** ($\mu_{\text{avg}}^{(t)}$) tracks the mean uncertainty across $\mathcal{K}_t$, indicating macro-level convergence to a coherent state; ❷ **Peak Uncertainty** ($s_{\max}^{(t)}$) captures the maximum entropy spike, revealing micro-level factual errors (e.g., incorrect entities) that are typically masked by global averaging; ❸ **Regional Uncertainty** ($\mu_{\text{top-}k}^{(t)}$) averages the top-k entropy (or all, if $|\mathcal{K}_t| < k$), capturing phrase-level uncertainty patterns among the remaining tokens. Together, these statistics allow the detector to capture errors at different scales and remain robust to varying sequence lengths. This approach ensures that isolated factual mistakes are not hidden within fluent and confident contexts.

Finally, stacking these step-wise statistics transforms the raw trajectory $\mathcal{T}$ into a compact evidence trajectory $\mathcal{E} = (\mathbf{a}_T, \mathbf{a}_{T-1}, \ldots, \mathbf{a}_0)$, mapping the semantic uncertainty resolution from initial noise ($t = T$) to the final generation ($t = 0$).

3.2 Dynamical Deviation Learning

Once high-quality evidence is constructed, the next question is how to leverage it for hallucination detection. While the above module focuses on the spatial dimension (i.e., structural-token filtering) of evidence construction, the temporal dimension is also crucial, namely how to model the denoising dynamics of the evidence to identify hallucinations. Although the evolution of evidence can reveal useful patterns, the range and trend of entropy may vary significantly depending on the specific scenarios and questions, making it difficult to directly use the evidence for hallucination prediction. To mine discriminative signals from denoising dynamics of evidence, we propose a dynamical deviation learning module consisting of two components.

First, we build an *evidence dynamics generator* to predict an expected reference trajectory conditioned on the input question, which captures the specific evolution patterns of evidence under different Q&A scenarios. Meanwhile, a *deviation-based hallucination detector* makes the final prediction by adaptively measuring the discrepancy between the actual evidence trajectory $\mathcal{E}$ and the predicted reference, supporting reliable hallucination detection.

Reference Evidence Dynamics Generator. Since denoising dynamics vary significantly across different queries and scenarios, raw evidence is not a reliable direct indicator for hallucination detection. Rather than predicting hallucinations directly from evidence, we design a generator g_θ to model the expected evidence trajectory. Conditioned on the input query embedding $\mathbf{q}$ and the timestep embedding $\mathbf{e}_t$, g_θ learns the relationship between the task content and the uncertainty dynamics. This produces a reference trajectory that represents normal, factual behavior to assist the detection process.

For each step t, the generator predicts a reference evidence vector $\hat{\mathbf{a}}_t = g_\theta(\mathbf{q}, \mathbf{e}_t)$. We optimize the generator using only factual samples ($y = 0$) via:

$$\mathcal{L}_{\text{ref}} = \mathbb{E}_{y=0}\left[\frac{1}{T}\sum_{t=0}^{T}\text{SmoothL1}(\mathbf{a}_t, \hat{\mathbf{a}}_t)\right]. \quad (5)$$

This approach allows the framework to model task-specific dynamics of factual sequences. In practice, we first pre-train g_θ using $\mathcal{L}_{\text{ref}}$ to establish a stable reference space for the subsequent detection stage.

Deviation-based Hallucination Detector. Based on the reference evidence trajectory, we use the deviation between the actual and reference curves as the direct criterion for hallucination detection. The detector can infer hallucinations by adaptively evaluating differences in curve trends. Specifically, we first construct a composite feature vector $\mathbf{x}_t = [\mathbf{a}_t;\ \hat{\mathbf{a}}_t;\ \Delta\mathbf{a}_t]$ at each step t to model the dynamics. Here, $\mathbf{a}_t$ and $\hat{\mathbf{a}}_t$ represent the observed and reference states, while $\Delta\mathbf{a}_t = \mathbf{a}_{t-1} - \mathbf{a}_t$ captures the step-wise velocity toward the subsequent state. A shared MLP, f_ϕ, then projects these vectors into hidden representations $\mathbf{z}_t$.

Since different steps contribute differently to detection, we employ a learnable weight ω_t to model their importance. A scoring network evaluates each time step by computing unnormalized scores

$$u_t = \mathbf{w}^\top \tanh(\mathbf{U}\mathbf{z}_t + \mathbf{b}_u), \quad (6)$$

where $\mathbf{w}$, $\mathbf{U}$, and $\mathbf{b}_u$ are learnable parameters. Then, these scores are normalized via a softmax function to produce the weight distribution $\omega_t = \exp(u_t)/\sum_{j=0}^{T}\exp(u_j)$. Finally, the hallucination prediction $\tilde{y}$ is computed by $\tilde{y} = \text{MLP}\left(\sum_{t=0}^{T}\omega_t\mathbf{z}_t\right)$, where temporal phases with prominent deviations can be prioritized.

Regularization Losses. To better identify hallucination characteristics, we introduce two dynamical regularizers weighted by ω_t that focus on late-stage stagnation and uncertainty rebound. Concretely, the path-deviation score s_{path} quantifies cumulative divergence:

$$s_{\text{path}} = \sum_{t=0}^{T}\omega_t\sum_{d=1}^{D}|a_{t,d} - \hat{a}_{t,d}|. \quad (7)$$

The rebound score s_{reb} captures uncertainty increases along the descending steps ($T \to 0$):

$$s_{\text{reb}} = \sum_{t=1}^{T}\omega_t\sum_{d=1}^{D}\big(\max(0, a_{t-1,d} - a_{t,d})\big)^2. \quad (8)$$

We define an adaptive margin $\hat{m}^{(\tau)}$ to represent the boundary of factual behaviors, updated via EMA using factual samples:

$$\begin{cases} m_{\text{batch}} = \text{Quantile}(\{s_i \mid y_i = 0\}, q) \\ \hat{m}^{(\tau)} = (1-\beta)\hat{m}^{(\tau-1)} + \beta m_{\text{batch}}. \end{cases} \quad (9)$$

The regularization losses are formulated as:

$$\mathcal{L}_{\text{path}} = \mathcal{L}_{\text{hinge}}(s_{\text{path}}; \hat{m}^{(\tau)}_{\text{path}}), \quad (10)$$

Table 1: AUROC(%) comparison of hallucination detection methods on two D-LLMs across three QA datasets. The highest score is **bolded** and the second highest is underlined.

Model Method	TriviaQA		HotpotQA		CSQA		Avg
	128	64	128	64	128	64	
LLaDA-8B-Instruct							
Output-based Methods							
Perplexity	50.4	47.6	49.3	51.2	65.6	65.0	54.9
LN-Entropy	54.6	53.5	54.8	54.7	64.6	64.4	57.8
Semantic Entropy	68.9	67.3	57.6	53.8	44.1	43.9	55.9
Lexical Similarity	62.5	59.0	64.2	57.1	57.3	60.7	60.1
Latent-based Methods							
EigenScore	69.2	66.9	64.7	59.2	58.5	60.6	63.2
CCS	57.1	54.2	57.6	55.8	50.5	58.5	55.6
TSV	60.2	61.1	65.0	59.4	52.9	55.2	59.0
Trajectory-based Methods							
TraceDet	73.9	74.1	66.1	63.7	77.2	77.1	72.0
DynHD	**86.6**	**86.2**	**84.3**	**85.2**	**81.5**	**81.2**	**84.2**
Dream-7B-Instruct							
Output-based Methods							
Semantic Entropy	73.7	72.5	62.7	67.7	51.4	48.6	62.8
Lexical Similarity	58.3	64.0	59.7	62.7	77.3	76.9	66.5
Latent-based Methods							
EigenScore	66.0	69.1	62.5	67.0	76.9	77.5	69.8
CCS	56.9	50.3	51.7	58.2	54.2	53.2	54.1
TSV	75.6	74.7	58.7	63.0	62.3	56.8	65.2
Trajectory-based Methods							
TraceDet	78.1	**86.7**	75.1	76.0	**84.7**	84.1	80.8
DynHD	**87.2**	84.5	**80.2**	**85.4**	83.4	**84.6**	**84.3**

$$\mathcal{L}_{\text{reb}} = \mathcal{L}_{\text{hinge}}(s_{\text{reb}}; \hat{m}_{\text{reb}}^{(\tau)}), \tag{11}$$

where $\mathcal{L}_{\text{hinge}}(s; \hat{m}) = (1-y)\cdot s + y \cdot \max(0, \hat{m} - s)$ (Cortes and Vapnik, 1995). This forces the model to minimize deviations for factual samples while pushing hallucinations beyond the boundary. The total training objective combines the classification loss $\mathcal{L}_{\text{cls}}$ with these regularizers:

$$\mathcal{L} = \mathcal{L}_{\text{cls}} + \lambda_1 \mathcal{L}_{\text{path}} + \lambda_2 \mathcal{L}_{\text{reb}}, \tag{12}$$

where λ_1 and λ_2 are balance hyperparameters. Detailed pseudocode and complexity analysis of DynHD are provided in Appendix B.

4 Experiments

4.1 Experimental Setup

Datasets and Models. To ensure a fair and rigorous comparison, we strictly adhere to the evaluation protocol established by TraceDet (Chang et al., 2026). Specifically, our experiments are conducted across three widely adopted fact-based QA benchmarks: TriviaQA (Joshi et al., 2017), HotpotQA (Yang et al., 2018), and CommonsenseQA (Talmor et al., 2019), which cover open-domain fact-based questions, multi-hop reasoning, and commonsense reasoning. We adopt two representative models, LLaDA-8B-Instruct (Nie et al., 2025) and Dream-7B-Instruct (Ye et al., 2025), as backbone D-LLMs.

Table 2: Zero-shot cross-task generalization matrix.

Method	TriviaQA		HotpotQA		CSQA		Avg
	H-QA	CSQA	T-QA	CSQA	T-QA	H-QA	
Latent-based							
CCS	50.1	54.5	51.8	54.0	54.2	56.6	53.5
TSV	58.5	65.3	65.5	59.1	56.2	63.2	61.3
Trajectory-based							
TraceDet	73.1	61.5	57.4	65.0	**74.8**	66.2	66.3
DynHD	**85.4**	**68.6**	**73.3**	**65.4**	73.6	**71.2**	**72.9**

Baselines. We compare our method against the same set of baselines as TraceDet, covering both output-based and latent-based hallucination detection approaches. Output-based baselines include Perplexity (Ren et al., 2023), Length-Normalized Entropy (LN-Entropy) (Malinin and Gales, 2020), Semantic Entropy (Kuhn et al., 2023), and Lexical Similarity (Lin et al., 2024). Latent-based baselines include EigenScore (Chen et al., 2024), Contrast-Consistent Search (CCS) (Burns et al., 2023), and Truthfulness Separator Vector (TSV) (Park et al., 2025). We also compare our framework against TraceDet (Chang et al., 2026), the state-of-the-art method for D-LLM hallucination detection.

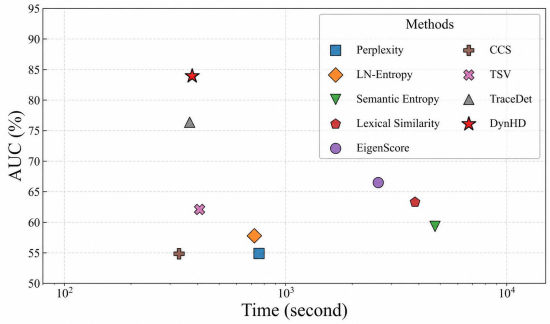

Figure 4: Pareto-frontier of performance vs. efficiency.

Evaluation Protocol. We evaluate hallucination detection performance using AUROC, following prior work (Park et al., 2025; Chen et al., 2024). Ground-truth correctness labels are provided by GPT-4o-mini, which is instructed to judge whether a generated answer is factually correct based on its consistency with the reference answer and its own internal knowledge (Zheng et al., 2023; Liu et al., 2023). We further assess the reliability of the automatic annotation by measuring agreement with human judgments, achieving an agreement rate of 94%. We partition the dataset into 1,400 samples for training, 300 for validation, and 300 for testing. Model selection is based on validation AUROC, and all reported results are obtained on the held-out test set. We provide detailed descriptions of the baselines, benchmarks, and experimental settings in Appendix C, with additional experiments in Appendix D.

4.2 Experimental Results

Main Results. Table 1 presents a comprehensive comparison of DynHD against baseline methods across two D-LLMs and three factuality QA datasets with varying generation lengths. From the results, we have several observations. ❶ DynHD achieves the highest performance in almost all settings. On the LLaDA-8B-Instruct model, it outperforms the strongest baseline by an average of 12.2% absolute AUROC, while on Dream-7B-Instruct, it reaches an average of 84.3%. These consistent gains exhibit the value of exploiting fine-grained deviation dynamics and convergence morphology rather than relying solely on output uncertainty or static hidden-state representations. ❷ Output-based methods lack consistency across datasets, showing poor performance on several cases. This occurs because multiple-choice answers are often confident but wrong, rendering uncertainty-based metrics ineffective. ❸ Latent-based approaches like EigenScore and TSV show more competitive results on Dream-7B-Instruct by capturing hidden-state geometry, but remain sensitive to dataset shifts.

Table 3: Ablation study in AUROC(%).

Ablation Variant	TriviaQA		HotpotQA		CSQA	
	128	64	128	64	128	64
LLaDA-8B-Instruct						
DynHD (Full Framework)	**86.6**	**86.2**	**84.3**	**85.2**	81.5	**84.2**
Semantic-aware Evidence Construction Components						
w/o Token Filtering	74.8	73.6	75.5	78.1	74.1	76.6
w/o Mean Uncertainty (μ_{avg})	78.6	81.1	79.0	80.2	78.5	77.1
w/o Max Entropy ($s_{\max}$)	82.5	83.1	81.8	82.3	79.1	82.1
w/o Top-k Entropy ($\mu_{\text{top-}k}$)	82.7	82.6	81.7	83.0	78.9	82.5
Dynamical Deviation Learning Components						
Naive Average	57.3	65.5	51.7	48.3	63.2	61.8
Flattened Temporal Aggregation	62.3	70.5	66.8	57.5	67.1	66.2
Uniform Temporal Pooling (w/o Attn)	84.5	84.8	82.4	83.1	79.5	80.5
w/o $\mathcal{L}_{path}$ (Stagnation)	84.1	85.3	84.1	83.8	**81.6**	83.2
w/o $\mathcal{L}_{reb}$ (Rebound)	85.2	85.6	84.2	83.5	80.6	83.3
Dream-7B-Instruct						
DynHD (Full Framework)	**87.2**	**84.5**	**80.2**	**85.4**	**83.4**	84.6
Semantic-aware Evidence Construction Components						
w/o Token Filtering	75.4	72.1	69.8	74.2	73.6	73.9
w/o Mean Uncertainty (μ_{avg})	79.2	78.5	74.5	79.1	76.2	77.8
w/o Max Entropy ($s_{\max}$)	83.1	81.2	77.5	82.3	80.9	81.5
w/o Top-k Entropy ($\mu_{\text{top-}k}$)	83.5	80.9	77.8	82.6	80.5	81.9
Dynamical Deviation Learning Components						
Naive Average	58.2	59.1	48.5	52.3	60.1	57.8
Flattened Temporal Aggregation	65.1	66.3	59.4	62.7	68.2	65.6
Uniform Temporal Pooling (w/o Attn)	84.5	82.1	78.0	83.1	81.2	81.8
w/o $\mathcal{L}_{path}$ (Stagnation)	85.2	83.5	79.2	84.5	82.8	83.8
w/o $\mathcal{L}_{reb}$ (Rebound)	86.1	83.1	79.5	84.2	82.7	**84.8**

Inference Efficiency. Figure 4 shows the balance between performance and efficiency. DynHD occupies the upper-left quadrant, representing the best results in our evaluation. While its speed is comparable to fast baselines like CCS, DynHD achieves significantly higher accuracy. Notably, our approach is much more efficient than multi-sample methods like Semantic Entropy, which often have high computational costs caused by multi-round sampling. Overall, DynHD provides an ideal balance of accuracy and speed, allowing for real-time monitoring without causing extra delays.

Cross-dataset Generalization Ability. To evaluate robustness, we conduct zero-shot cross-dataset assessments in Table 2, where DynHD yields the highest average AUROC (72.9%). It notably achieves an 85.5% AUROC in TriviaQA to H-QA transfer, significantly outperforming TraceDet (73.1%) and suggesting that our dynamical encoding captures universal denoising laws across domains. On the challenging CSQA source, DynHD remains robust (e.g., 73.6% on T-QA), despite slightly underperforming TraceDet. This is because the fixed options lead to a faster decision process, which reduces the observable uncertainty signals.

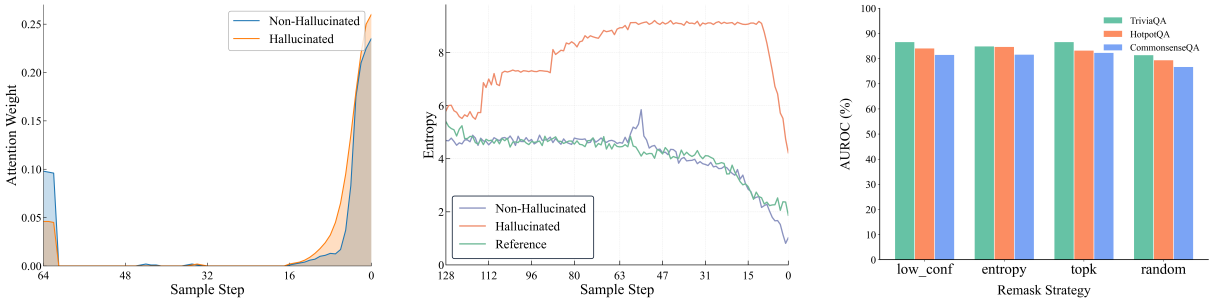

(a) (b) (c)

Figure 5: (a) Visualization of attention weights ω_t on different steps. (b) Top-k entropy trajectories of representative samples on HotpotQA ($T = 128$). (c) Performance comparison of different D-LLM remasking strategies.

Ablation Study. As summarized in Table 3, we evaluate the contributions of key designs in DynHD, coming with the following findings.

In terms of the evidence construction module, ❶ without token filtering, the performance drops significantly, which proves that filtering out noise is essential to get clean signals. ❷ Removing individual statistical evidence components consistently degrades results, proving that robust detection requires simultaneous monitoring of semantic stability, localized anomaly, and regional instability.

For the deviation learning module, ❸ *Naive Average* variant yields poor performance because simple arithmetic averaging cannot effectively distinguish hallucinations. ❹ *Flattened Temporal Aggregation* underperforms due to its lack of explicit temporal dynamic analysis. ❺ Removing attention-based weighting degrades performance, as the model fails to capture critical dynamic changes, especially at the late stages of the denoising process. ❻ Our regularization terms further improve the model's ability to differentiate hallucinations from factual generations by capturing and penalizing specific non-convergent trajectories, such as *Stagnation* and *Rebound* patterns.

Visualization of Temporal Weighting. Figure 5a shows the attention weights for LLaDA-8B-Instruct on the TriviaQA dataset ($t = 64$). We can see the learned weights are mostly concentrated in the final steps of the denoising process. This matches our observation that stagnation and rebound patterns usually happen at the end of the trajectory. By focusing on these late steps, our model can find signs of hallucinations and ignore early random noise.

Visualization of Reference Trajectories. Figure 5b shows entropy trajectories on the HotpotQA dataset. We introduce the learned reference path to show how a correct denoising process should look. The factual sample stays close to this path and drops smoothly toward zero. In contrast, a hallucinated sample is unstable and shows a clear rebound at the end. This shows that the curve's shape matters more than its absolute value. The reference generator facilitates identifying denoising processes associated with hallucinations.

Robustness to Masking Strategies. Figure 5c examines the performance of DynHD on LLaDA-8B-Instruct under various internal remasking strategies employed by D-LLMs during generation. The results across multiple datasets consistently demonstrate competitive detection performance, regardless of the specific spatial filtering logic used (such as entropy-based or confidence-based remasking). This stability suggests that the proposed framework is adaptable and maintains its effectiveness across diverse D-LLM configurations.

5 Conclusion

We propose DynHD to detect D-LLM hallucinations using trajectory-based deviation modeling. DynHD identifies factual errors by comparing actual evidence against a learned reference. This approach captures key signals such as late-stage stagnation and uncertainty rebound. By combining semantic-aware evidence construction with dynamical deviation learning, DynHD provides a practical framework for exploiting both spatial uncertainty patterns and temporal denoising behaviors. Experiments on mainstream models show that DynHD achieves superior performance with high computational efficiency. Overall, this work demonstrates that denoising dynamics can serve as signals for hallucination detection of D-LLMs.

Acknowledgment

The work of Y. Liu was partially supported by the Australian Research Council (ARC) under Grant No. DE260101172.

Limitations

Although DynHD demonstrates strong performance on standard QA benchmarks, the current evaluation mainly focuses on factual question answering scenarios, including open-domain QA, multi-hop QA, and commonsense QA. These benchmarks provide clear correctness criteria and are suitable for controlled hallucination detection evaluation. Nevertheless, hallucination detection in more open-ended generation settings, such as long-form summarization, multi-turn dialogue, mathematical reasoning, code generation, and structured prediction, may involve more diverse factual boundaries and formatting patterns. Evaluating DynHD in these scenarios would further strengthen the understanding of its generality. In addition, our experiments are conducted on two representative D-LLM backbones under commonly used generation lengths. While the results demonstrate consistent effectiveness across different datasets and models, future work could further examine broader model families, decoding configurations, and multilingual settings. Finally, the current design relies on training an evidence dynamics generator using in-domain factual samples to establish reference trajectories, which may affect its adaptability to entirely unseen domains. Since our framework is tailored to the continuous denoising process of Diffusion LLMs, extending these trajectory-based insights to standard auto-regressive models remains an interesting direction for future research.

Ethics Statement

Our research focuses exclusively on scientific questions, with no involvement of human subjects, animals, or environmentally sensitive materials. Therefore, we foresee no ethical risks or conflicts of interest. We are committed to maintaining the highest standards of scientific integrity and ethics to ensure the validity and reliability of our findings.

Artifact Use and Licensing

We use publicly available datasets, models, and baseline methods in this work, including TriviaQA, HotpotQA, CommonsenseQA, LLaDA-8B-Instruct, and Dream-7B-Instruct. We cite the original creators of these artifacts and use them only for research and evaluation purposes. We do not redistribute the original datasets or model weights, and our use follows the corresponding licenses and terms of use released by the artifact providers.

References

Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. 2021. Structured denoising diffusion models in discrete state-spaces. *Advances in neural information processing systems*, 34:17981–17993.

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.

Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. 2023. Discovering latent knowledge in language models without supervision. In *International Conference on Learning Representations*.

Shenxu Chang, Junchi Yu, Weixing Wang, Yongqiang Chen, Jialin Yu, Philip Torr, and Jindong Gu. 2026. Tracedet: Hallucination detection from the decoding trace of diffusion large language models. In *International Conference on Learning Representations*.

Baiyu Chen, Wilson Wongso, Xiaoqian Hu, Yue Tan, and Flora D Salim. 2025. Multi-stage verification-centric framework for mitigating hallucination in multi-modal rag. In *2025 KDD Cup Workshop for Multimodal Retrieval Augmented Generation*.

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 *International Conference on Learning Representations*.

I-Chun Chern, Steffi Chern, Shiqi Chen, Weizhe Yuan, Kehua Feng, Chunting Zhou, Junxian He, Graham Neubig, and Pengfei Liu. 2025. Factool: Factuality detection in generative ai – a tool augmented framework for multi-task and multi-domain scenarios. In *Conference on Language Modeling*.

Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James Glass, and Pengcheng He. 2024. Dola: Decoding by contrasting layers improves factuality in large language models. In *International Conference on Learning Representations*.

Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine learning*, 20(3):273–297.

Xuefeng Du, Chaowei Xiao, and Sharon Li. 2024. Haloscope: Harnessing unlabeled llm generations for hallucination detection. *Advances in Neural Information Processing Systems*, 37:102948–102972.

Shansan Gong, Mukai Li, Jiangtao Feng, Zhiyong Wu, and LingPeng Kong. 2023. Diffuseq: Sequence to sequence text generation with diffusion models. In *International Conference on Learning Representations*.

Jiatong Han, Jannik Kossen, Muhammed Razzak, Lisa Schut, Shreshth A Malik, and Yarin Gal. 2024. Semantic Entropy Probes: Robust and Cheap Hallucination Detection in LLMs. In *ICML 2024 Workshops: FM-Wild*.

Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and 1 others. 2025. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, 43(2):1–55.

Yue Huang, Lichao Sun, Haoran Wang, Siyuan Wu, Qihui Zhang, Yuan Li, Chujie Gao, Yixin Huang, Wenhan Lyu, Yixuan Zhang, and 1 others. 2024. Trustllm: Trustworthiness in large language models. In *Proceedings of the 41st International Conference on Machine Learning*, pages 20166–20270. PMLR.

Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, and 1 others. 2024. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*.

Yatai Ji, Teng Wang, Yuying Ge, Zhiheng Liu, Sidi Yang, Ying Shan, and Ping Luo. 2025. From denoising to refining: A corrective framework for vision-language diffusion model. *arXiv preprint arXiv:2510.19871*.

Che Jiang, Biqing Qi, Xiangyu Hong, Dayuan Fu, Yang Cheng, Fandong Meng, Mo Yu, Bowen Zhou, and Jie Zhou. 2024. On large language models' hallucination with regard to known facts. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 1041–1053.

Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611.

Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. In *International Conference on Learning Representations*.

Gen Li and Changxiao Cai. 2025. Breaking ar's sampling bottleneck: Provable acceleration via diffusion language models. In *Advances in neural information processing systems*.

Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2023. Inference-time intervention: Eliciting truthful answers from a language model. *Advances in Neural Information Processing Systems*, 36:41451–41530.

Shiyuan Li, Yixin Liu, Yu Zheng, Mei Li, Quoc Viet Hung Nguyen, and Shirui Pan. 2026a. Ofamas: One-for-all multi-agent system topology design based on mixture-of-experts graph generative models. In *Proceedings of the ACM Web Conference 2026*, pages 1333–1344.

Shiyuan Li, Yunfeng Zhao, Yue Tan, Qingfeng Chen, Yixin Liu, and Shirui Pan. 2026b. Towards anomaly detection on relational data. *arXiv preprint arXiv:2606.18621*.

Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. 2022. Diffusion-lm improves controllable text generation. *Advances in neural information processing systems*, 35:4328–4343.

Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. 2024. Generating with confidence: Uncertainty quantification for black-box large language models. *Transactions on Machine Learning Research*.

Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. G-eval: Nlg evaluation using gpt-4 with better human alignment. In *Proceedings of the 2023 conference on empirical methods in natural language processing*, pages 2511–2522.

Yixin Liu, Shiyuan Li, Yu Zheng, Qingfeng Chen, Chengqi Zhang, Philip S Yu, and Shirui Pan. 2026a. From few-shot to zero-shot: Towards generalist graph anomaly detection. *IEEE Transactions on Knowledge and Data Engineering*.

Yixin Liu, Kehan Yan, Shiyuan Li, Qingfeng Chen, and Shirui Pan. 2026b. Beyond a single perspective: Text anomaly detection with multi-view language representations. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*.

Yujing Liu, Yixin Liu, Yu Zheng, Alan Wee-Chung Liew, Xiaofeng Cao, and Shirui Pan. 2026c. Rethinking feature alignment in generalist graph anomaly detection: A relational fingerprint-based approach. In *International Conference on Machine Learning*.

Andrey Malinin and Mark Gales. 2020. Uncertainty estimation in autoregressive structured prediction. *arXiv preprint arXiv:2002.07650*.

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*, pages 9004–9017.

Samuel Marks and Max Tegmark. 2024. The geometry of truth: Emergent linear structure in large language model representations of true/false datasets. In *Conference on Language Modeling*.

Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. 2025. Large language diffusion models. In *Advances in Neural Information Processing Systems*.

Junjun Pan, Yixin Liu, Yu Zheng, Lianhua Chi, Alan Wee-Chung Liew, and Shirui Pan. 2026. Camera: Adapting to semantic camouflage in unsupervised text-attributed graph fraud detection. In *International Joint Conference on Artificial Intelligence*.

Bhargavi Paranjape, Mandar Joshi, John Thickstun, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2020. An information bottleneck approach for controlling conciseness in rationale extraction. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1938–1952.

Seongheon Park, Xuefeng Du, Min-Hsuan Yeh, Haobo Wang, and Yixuan Li. 2025. Steer llm latents for hallucination detection. In *International Conference on Machine Learning*.

Elizabeth Pavlova and Xue-Xin Wei. 2025. Diffusion models under low-noise regime. *arXiv preprint arXiv:2506.07841*.

Jie Ren, Jiaming Luo, Yao Zhao, Kundan Krishna, Mohammad Saleh, Balaji Lakshminarayanan, and Peter J Liu. 2023. Out-of-distribution detection and selective generation for conditional language models. In *International Conference on Learning Representations*.

Xu Shen, Yixin Liu, Yili Wang, Rui Miao, Yiwei Dai, Shirui Pan, Yi Chang, and Xin Wang. 2026. Raising the bar in graph ood generalization: Invariant learning beyond explicit environment modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. Commonsenseqa: A question answering challenge targeting commonsense knowledge. 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)*, pages 4149–4158.

Yue Tan, Chen Chen, Weiming Zhuang, Xin Dong, Lingjuan Lyu, and Guodong Long. 2023. Taming heterogeneity to deal with test-time shift in federated learning. In *International Workshop on Federated Learning for Distributed Data Mining*.

Yue Tan, Guodong Long, Jing Jiang, and Chengqi Zhang. 2026. Influence-oriented personalized federated learning. In *IEEE International Conference on Data Mining*.

Wen Wang, Bozhen Fang, Chenchen Jing, Yongliang Shen, Yangyi Shen, Qiuyu Wang, Hao Ouyang, Hao Chen, and Chunhua Shen. 2026a. Time is a feature: Exploiting temporal dynamics in diffusion language models. In *International Conference on Learning Representations*.

Xu Wang, Chenkai Xu, Yijie Jin, Jiachun Jin, Hao Zhang, and Zhijie Deng. 2026b. Diffusion llms can do faster-than-ar inference via discrete diffusion forcing. In *International Conference on Learning Representations*.

Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 conference on empirical methods in natural language processing*, pages 2369–2380.

Jiacheng Ye, Zhihui Xie, Lin Zheng, Jiahui Gao, Zirui Wu, Xin Jiang, Zhenguo Li, and Lingpeng Kong. 2025. Dream 7b: Diffusion large language models. *arXiv preprint arXiv:2508.15487*.

Yakir Yehuda, Itzik Malkiel, Oren Barkan, Jonathan Weill, Royi Ronen, and Noam Koenigstein. 2024. Interrogatellm: Zero-resource hallucination detection in llm-generated answers. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9333–9347.

Zhangyue Yin, Qiushi Sun, Qipeng Guo, Jiawen Wu, Xipeng Qiu, and Xuan-Jing Huang. 2023. Do large language models know what they don't know? In *Findings of the association for Computational Linguistics: ACL 2023*, pages 8653–8665.

Muru Zhang, Ofir Press, William Merrill, Alisa Liu, and Noah A. Smith. 2024. How language model hallucinations can snowball. In *International Conference on Machine Learning*, pages 59670–59684.

Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, and 1 others. 2025. Siren's song in the ai ocean: A survey on hallucination in large language models. *Computational Linguistics*, 51(4):1373–1418.

Yunfeng Zhao, Yixin Liu, Qingfeng Chen, Shiyuan Li, Yue Tan, and Shirui Pan. 2026. Fedcigar: A personalized reconstruction approach for federated graph-level anomaly detection. In *International Joint Conference on Artificial Intelligence*.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, and 1 others. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36:46595–46623.

A Related Work

Hallucination Detection Most existing hallucination detection studies focus on Autoregressive Language Models (AR-LLMs) (Zhang et al., 2025; Shen et al., 2026). These methods typically extract signals from the final prediction, such as token-level uncertainty measures (e.g., entropy or perplexity (Malinin and Gales, 2020; Ren et al., 2023)), or assess semantic consistency across multiple samples (e.g., semantic entropy (Kuhn et al., 2023; Manakul et al., 2023; Li et al., 2026b)). Some recent work also probes internal representations during the forward pass (e.g., hidden states (Azaria and Mitchell, 2023; Li et al., 2023; Marks and Tegmark, 2024; Liu et al., 2026c)) to infer factuality. However, diffusion large language models (D-LLMs) differ fundamentally from AR-LLMs in their generation mechanism: hallucination cues are not confined to the final step but are distributed across a denoising trajectory spanning the entire iterative process. This process-centric nature makes static, outcome-based detectors prone to missing discriminative evidence exposed along the trajectory (Chuang et al., 2024).

Diffusion Large Language Models Diffusion Large Language Models (D-LLMs) extend the remarkable success of diffusion paradigms (Wang et al., 2026b) to the discrete text domain (Li et al., 2022). By unifying the discrete remasking process, LLaDA has successfully scaled D-LLMs to the 8B parameter level, achieving performance parity with leading AR-LLMs such as LLaMA-3. Furthermore, Dream-7B demonstrates the architectural versatility of this paradigm by adopting the identical configurations of Qwen2.5-7B (Hui et al., 2024) and training it under a diffusion objective. On the reliability side, Wang et al. (2026a) first identified the phenomenon of temporal oscillation, where correct answers may emerge at intermediate denoising steps but can be overridden by erroneous late-stage refinements. Building on the trajectory view, TraceDet (Chang et al., 2026) proposes a trajectory-based detection framework that models generation as an action trace and leverages the Information Bottleneck principle to automatically select the most discriminative sub-trace steps (Paranjape et al., 2020). These studies motivate the use of denoising trajectories as process-level signals for D-LLM hallucination detection.

Distinction from TraceDet TraceDet (Chang et al., 2026) is the most relevant work to ours, as it also exploits the denoising trajectory of D-LLMs for hallucination detection. However, DynHD differs from TraceDet in two key aspects. First, TraceDet focuses on selecting discriminative sub-trace steps from the generation trajectory, but does not explicitly address the spatial information density imbalance caused by fixed-length generation. In contrast, DynHD introduces semantic-aware token filtering to remove structural padding tokens before evidence construction, reducing the dilution effect of non-semantic positions. Second, TraceDet mainly relies on selected trajectory fragments for detection, whereas DynHD models the continuous evolution of uncertainty evidence by learning a question-conditioned reference trajectory. Hallucinations are then detected by measuring deviations between the observed evidence trajectory and the learned reference dynamics. Therefore, DynHD extends trajectory-based hallucination detection from sub-trace selection to fine-grained spatial evidence construction and temporal deviation modeling.

Denoising Dynamics Denoising dynamics characterizes how a model progressively removes randomness and approaches the target distribution through multi-step refinement. Li and Cai (2025); Li et al. (2026a) provide theoretical analysis suggesting that, under ideal conditions, the sampling error of D-LLMs should decrease as the number of denoising steps increases and eventually converge to the data distribution. This implies that reliable generation trajectories should be accompanied by smooth uncertainty dissipation. When hallucinations arise due to insufficient knowledge (Yin et al., 2023; Jiang et al., 2024), this dynamical balance can break: uncertainty may remain persistently high because of unresolved inconsistencies, leading to stagnation, or exhibit abnormal rebound in mid-to-late stages when refinement fails (Huang et al., 2025). Such morphological deviations from a normal convergence path directly reflect instability in the generation process (Pavlova and Wei, 2025; Liu et al., 2026a). This work is the first to utilize such continuous temporal evolution patterns for D-LLM hallucination detection, advancing the field from static uncertainty summaries toward a comprehensive process-level characterization of generative reliability.

Algorithm 1 The two-stage training algorithm of DynHD

Require: Training pool $\mathcal{X}$; ignore set $\mathcal{I}_{\text{ignore}}$; hyperparameter k; epochs E_1, E_2; batch size B.
Ensure: Optimized reference parameters θ and detector parameters Φ.

// Stage 1: Evidence Dynamics Generator (Sec. 3.2)
Initialize generator parameters θ;
for $epoch = 1$ **to** E_1 **do**
  Sample mini-batch $\{(\mathcal{T}, \mathbf{q}, y)_i\}_{i=1}^{B} \subset \mathcal{X}$ where $y = 0$ {Factual samples only}
  for all $(\mathcal{T}, \mathbf{q})$ **in** mini-batch **do**
    Evidence Construction: Extract valid semantic positions $\mathcal{K}_t$ via $\mathcal{I}_{\text{ignore}}$;
    Condense step-wise uncertainty into $\mathbf{a}_t = [\mu_{\text{keep}}^{(t)}, s_{\max}^{(t)}, \mu_{\text{top-}k}^{(t)}]$ (Eq. 4);
    Predict reference trajectory $\hat{\mathbf{a}}_t = g_\theta(\mathbf{q}, \mathbf{e}_t)$ using query and timestep embeddings;
  end for
  Update θ by minimizing $\mathcal{L}_{\text{ref}}$ via SmoothL1 loss (Eq. 5);
end for
// Stage 2: Deviation-based Hallucination Detector (Sec. 3.2)
Initialize detector parameters Φ and adaptive factual margins $\hat{m}^{(\tau)}$;
for $epoch = 1$ **to** E_2 **do**
  Sample mini-batch $\{(\mathcal{T}, \mathbf{q}, y)_i\}_{i=1}^{B} \subset \mathcal{X}$;
  for all $(\mathcal{T}, \mathbf{q}, y)$ **in** mini-batch **do**
    Evidence Construction: Obtain actual evidence trajectory $\mathcal{E} = (\mathbf{a}_T, \ldots, \mathbf{a}_0)$ from $\mathcal{T}$ via Eq. 1 & 2;
    Generate reference trajectory $\hat{\mathcal{E}} = (\hat{\mathbf{a}}_T, \ldots, \hat{\mathbf{a}}_0)$ via frozen g_θ;
    Compute velocity $\Delta\mathbf{a}_t = \mathbf{a}_{t-1} - \mathbf{a}_t$ to capture transition dynamics;
    Construct composite feature $\mathbf{x}_t = [\mathbf{a}_t; \hat{\mathbf{a}}_t; \Delta\mathbf{a}_t]$ for each step t;
    Project $\mathbf{x}_t$ to $\mathbf{z}_t$, compute weight ω_t (Eq. 6), and aggregate to $\mathbf{h}$;
    Compute path-deviation score s_{path} (Eq. 7) and rebound score s_{reb} (Eq. 8);
  end for
  Update adaptive margins $\hat{m}^{(\tau)}$ via EMA using scores of factual samples ($y = 0$) in the batch (Eq. 9);
  Update Φ by minimizing total loss $\mathcal{L} = \mathcal{L}_{\text{cls}} + \lambda_1\mathcal{L}_{\text{path}} + \lambda_2\mathcal{L}_{\text{reb}}$ (Eq. 10);
end for

B Algorithmic Details of DynHD

B.1 Pseudocode

Algorithm 1 details the complete process of our algorithm.

B.2 Complexity Analysis

The computational overhead of DynHD is negligible compared to the D-LLM inference cost, as it operates on pre-computed uncertainty fields without re-invoking the backbone.

- **Spatial Filtering:** For each step t, position-wise masking is $\mathcal{O}(l)$, and extracting Top-k statistics requires $\mathcal{O}(l \log k)$ using a min-heap. The total complexity for a full trajectory is $\mathcal{O}(T \cdot l \log k)$. Since the sequence length l is typically much smaller than the vocabulary size $|\mathcal{V}|$, this process is significantly faster than the standard logit-to-probability conversion in Transformer heads.
- **Temporal Modeling:** Our temporal encoder and attention mechanism operate on the condensed 3D evidence vectors. With a hidden dimension h, the complexity is $\mathcal{O}(T \cdot (d \cdot h + h^2))$, where $d = 3$. Since d and h are small constants independent of the LLM's scale, the inference latency of the detection head is essentially constant, adding no perceptible delay to the generation pipeline.

As a result, DynHD maintains an efficient $\mathcal{O}(T \cdot l)$ footprint, which is dwarfed by the quadratic self-attention complexity $\mathcal{O}(T \cdot l^2)$ of the underlying D-LLM.

C Detailed Experimental Setup

C.1 Datasets and Benchmarks

To ensure a rigorous and multi-faceted evaluation, we conduct experiments across three representative factuality-oriented QA benchmarks, following

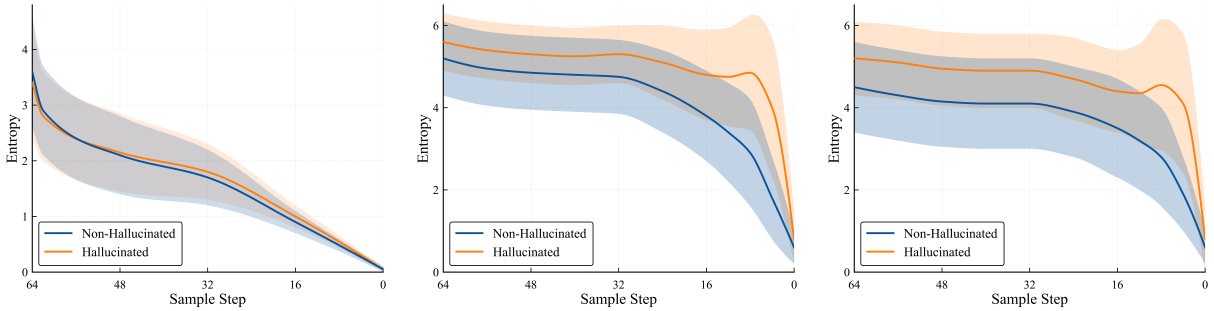

(a) Global Distinguishability (b) HotpotQA Results (c) TriviaQA Results

Figure 6: Empirical analysis of the distinguishability of statistical evidence trajectories. (a) shows the enhanced separation after structure-aware filtering, while (b) and (c) demonstrate the consistent morphological divergence between factual and hallucinatory samples across datasets.

the protocol established in recent trajectory-based studies (Chang et al., 2026):

- **TriviaQA** (Joshi et al., 2017): A large-scale open-domain dataset focusing on precise entity retrieval and factual knowledge.
- **HotpotQA** (Yang et al., 2018): A multi-hop reasoning benchmark that requires synthesizing information across multiple context fragments to maintain logical consistency.
- **CommonsenseQA** (Talmor et al., 2019): A task evaluating general reasoning capabilities, where the model must navigate commonsense logic beyond surface-level pattern matching.

All dataset preprocessing, including prompt templates and data splits, strictly adheres to the configurations used in the comparative studies to ensure a level playing field.

C.2 Baseline Configurations

We compare DynHD against seven established baselines, categorized into two primary paradigms based on their information source:

(1) Output-Centric Approaches: These methods detect hallucinations by analyzing the generated text or its output probability distribution:

- *Perplexity* (Ren et al., 2023): Identifies potential hallucinations based on the negative log-likelihood assigned by the model to its own sequence.
- *LN-Entropy* (Malinin and Gales, 2020): Employs length-normalized predictive entropy to flag generations with unusually high average uncertainty.
- *Semantic Entropy* (Kuhn et al., 2023): Quantifies consistency by partitioning multiple stochastic outputs into semantic equivalence classes.
- *Lexical Similarity* (Lin et al., 2024): Assesses hallucinations via surface-level lexical overlap across independent generation paths.

(2) Latent-Centric Approaches: These methods probe the internal representations and hidden states of the transformer architecture:

- *EigenScore* (Chen et al., 2024): Evaluates response consistency through the covariance structure of internal latent embeddings.
- *CCS* (Burns et al., 2023): Searches for a contrastive-consistent direction in the latent space to distinguish truthful from untruthful activations.
- *TSV* (Park et al., 2025): Utilizes a separator vector to classify truthfulness based on learned centroids within the latent space.

Finally, we include **TraceDet** (Chang et al., 2026) as the primary trajectory-based baseline. TraceDet characterizes the denoising path using a uniform temporal aggregation strategy, providing a representative point of comparison for analyzing the evolution of uncertainty during the denoising process.

C.3 Implementation Setup

Our experiments, including model inference and detector training, are conducted on **NVIDIA RTX 3090 GPUs (24GB VRAM)**.

C.4 Automatic Labeling Reliability

Following recent hallucination detection studies, we use GPT-4o-mini to obtain correctness labels for generated answers. The judge is instructed to determine whether a generated answer is factually correct based on its consistency with the reference answer and its own knowledge. This reference-grounded protocol provides scalable and consistent annotations for both training and evaluation.

We acknowledge that automatic judges may introduce bias, especially for ambiguous cases or semantically equivalent answers. To assess the reliability of the annotations, we manually examine a subset of labeled samples and observe a 94% agreement between GPT-4o-mini and human judgments. This agreement is also higher than the baseline agreement level used in our reliability check, further supporting the quality of the automatic annotations. This suggests that the automatic labels are sufficiently reliable for our controlled QA evaluation setting, while future work can further strengthen the annotation process with larger-scale human verification or multi-judge consensus.

C.5 Hyperparameter Search Space

We conduct a grid search to optimize the performance of DynHD. The variables and their corresponding search ranges are summarized in Table 4.

Table 4: Search space of key hyperparameters in DynHD.

Variable	Search Range
$\mathrm{LR}_{\mathrm{stage1}}$	$\{3\times10^{-4}, 1\times10^{-3}\}$
$\mathrm{LR}_{\mathrm{stage2}}$	$\{1\times10^{-4}, 3\times10^{-4}, 1\times10^{-3}\}$
$\mathrm{WD}_{\mathrm{stage1}}$	$\{0, 0.01, 0.05\}$
$\mathrm{WD}_{\mathrm{stage2}}$	$\{0, 1\times10^{-4}, 1\times10^{-3}\}$
λ_1 (Path)	$[0.0, 0.4]$ (step 0.05)
λ_2 (Rebound)	$[0.0, 0.4]$ (step 0.05)
$\lambda_{\mathrm{warmup}}$	$\{0, 0.3\}$

D Additional Experimental Results

D.1 Case of Hallucination Dynamics

Table 5 details the four samples discussed in our Introduction. For factual cases like the New York projects and the Shining Path, the denoising entropy steadily decreases as the model converges on the correct facts in the late-stage. In contrast, hallucinations show clear instability. In the Disney case, the model is biased toward "Florida" maybe due to the keyword "Disney," but the specific constraint "Djuan Rivers" causes an entropy rebound when the model fails to resolve the conflict. Similarly, the hallucination for the 70s hit shows a late-stage bounce, indicating a breakdown in factual retrieval. These examples confirm that the trajectory shape is a more reliable signal for detection than absolute entropy values.

D.2 Distinguishability of Statistical Evidence Trajectories

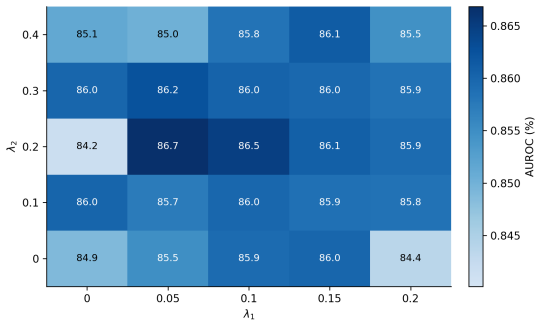

Figure 7: AUROC heatmap for varying λ_1 (Path) and λ_2 (Rebound) on TriviaQA. The peak region validates the synergistic effect of both dynamical regularizers.

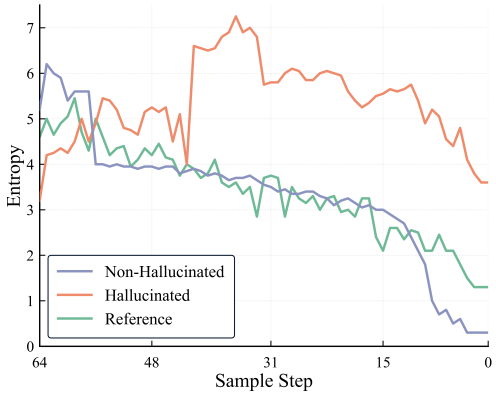

Figure 8: Empirical denoising dynamics of a_t (observed) and $\hat{a}_t$ (reference) at 64 steps on HotpotQA.

We test how our metrics separate factual answers from hallucinations in Figures 6a–6c. As shown in Figure 6a, our filtering method clarifies the uncertainty data. This creates a clear difference between facts and hallucinations, whereas raw data often overlaps. Specific tests in Figures 6b and 6c show a common pattern: factual paths smoothly go down, while hallucinations bounce at the end.

D.3 Sensitivity of λ_1 and λ_2

Figure 7 shows how λ_1 (Path) and λ_2 (Rebound) work together on TriviaQA. The heatmap indicates that the model works best when both λ_1 and λ_2

Dataset	Type	Question & Response
HotpotQA	Factual	**Q:** Are Manhattan West and Singer Building both projects in New York? **R:** Yes.
	Hallucination	**Q:** What State has a Disney Resort & Spa that is a beachside hotel, resort and vacation destination offering complimentary children's activities and programs and that Djuan Rivers was a General Manager at? **R:** The Disney Resort and Spa where Djuan Rivers served as General Manager is located in Florida.
TriviaQA	Factual	**Q:** Where did the Shining Path terrorists operate? **R:** The Shining Path terrorists operated in Peru.
	Hallucination	**Q:** Who had a 70s No 1 hit with Kiss You All Over? **R:** The Village People had a 70s No 1 hit with "Kiss You All Over."

Table 5: Examples of factual and hallucinated samples. Each dataset displays a Factual sample followed by a Hallucination. Factual samples exhibit consistent convergence, whereas hallucinations trigger detectable instability in denoising dynamics.

are used (around 0.1 to 0.3). This confirms that tracking the whole path and checking the late-stage bounce are both important. The model performs well across a wide range of values, showing it is not overly sensitive to specific parameter settings.

D.4 Denoising Dynamics at 64 Steps

Figure 8 compares the real paths a_t and the reference curves $\hat{a}_t$ at 64 steps. Factual answers follow the reference curve closely, but hallucinations show a clear bounce. These results prove that the shape of the curve is the most important indicator of an error, regardless of the exact values.