HD-NDEs: Neural Differential Equations for Hallucination Detection in LLMs

Anonymous ACL submission

Abstract

In recent years, large language models (LLMs) have made remarkable advancements, yet hallucinations, where models produce inaccurate or non-factual statements, remain a significant challenge for real-world deployment. Although current supervised methods, such as the probe method, are highly efficient in mitigating hallucinations, they struggle when non-factual information arises in the early or mid-sequence of outputs, reducing their reliability. To address these issues, we propose HD-NDEs, a novel method that systematically assesses the truthfulness of statements by capturing the full dynamics of LLMs within their latent space. Our approaches apply neural differential equations (Neural DEs) to model the dynamic system in the latent space of LLMs during decoding. Then, the sequence in the latent space is mapped to the classification space for truth assessment. The extensive experiments across multiple datasets and four common open-source LLMs demonstrate the effectiveness of HD-NDEs, especially, achieving over 10% improvement in AUC-ROC on the True-False dataset compared to state-of-the-art techniques.

1 Introduction

Hallucination has been widely recognized as a significant challenge in large language models (LLMs), as highlighted in various studies applications (Li et al., 2023; Min et al., 2023; Geng et al., 2023). Efforts to mitigate this issue have led to the development of hallucination detection techniques, which are broadly categorized into evidence-based and evidence-free approaches. Evidence-based methods (Wang et al., 2023; Wei et al., 2024) generally involve retrieving relevant information from external sources to verify whether inconsistencies exist between the generated content and the retrieved evidence. Nevertheless, this retrieval and verification process is computationally intensive and time-consuming, making it impractical for high-throughput applications in routine use. In contrast, evidence-free methods (Manakul et al., 2023; Duan et al., 2023; Geng et al., 2023) primarily utilize the inherent characteristics of LLMs and semantic features to identify potential hallucinations. For instance, unsupervised techniques often estimate the overall uncertainty of a sentence by analyzing logit-based uncertainty at the token level. Alternatively, these approaches may assess the consistency of model outputs, based on the premise that hallucination tends to increase variability in the generated responses.

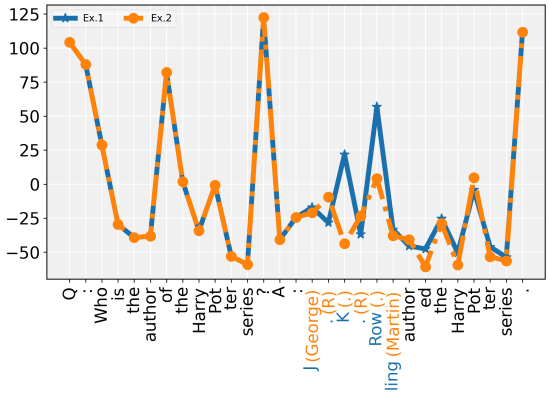

Figure 1: 1D PCA projection of hidden layer embeddings for each token in Ex.1 and Ex.2. Both examples use the same question with different answers. In the hidden state space, the embeddings of the earlier tokens are identical until the tokens begin to differ. The final few tokens, being the same, result in minimal differences in the hidden state activations.

Furthermore, recent methods have proposed using a model's internal states to probe its confidence in factual vs. non-factual sentences in a supervised setting. For instance, a simple feed-forward neural network classifier can be trained on the activations of the final token's last-layer hidden state to predict the reliability of the model's output. This approach has demonstrated significant effective-

ness, validated across numerous studies and various model architectures (Azaria and Mitchell, 2023a; Li et al., 2024; Kossen et al., 2024). However, the probing method is still in its early stages and remains inadequate for handling cases where the final token of a statement fails to capture the reliability of the entire sequence. It often struggles when non-factual tokens are located at the beginning or middle of the sequence, as mentioned in (Levinstein and Herrmann, 2024). We employ principal component analysis (PCA) (Abdi and Williams, 2010) to further investigate such failure cases. As shown in Figure 1, we reduce the dimensionality of each token's activations to a single dimension for clearer interpretation. Ex.1 illustrates a question with the correct answer, while Ex.2 presents the same question with an incorrect answer. Notably, the reduced hidden information of the last tokens in both examples appears nearly identical, despite differences in the middle of the sequences. This suggests that we need to effectively leverage hidden state information across the entire sequence, rather than only the last token, to accurately assess the truthfulness.

The advances of neural differential equations (Neural DEs, Kidger 2022) offer a promising solution by modeling hidden state transformations as continuous trajectories, providing a more accurate representation of information flow through LLMs. Proven effective in tasks such as time-series forecasting, classification, and outlier detection (Choi et al., 2022; Jhin et al., 2024), Neural DEs are well-suited for addressing hallucination detections in LLMs, where subtle errors can result in factual inaccuracies in generated sequences. Motivated by these strengths, our work introduces a novel, supervised method, called HD-NDEs, marking the first application of Neural DEs in hallucination detection. As shown in Figure 2, the method explicitly models the trajectory of intermediate states in the latent space using Neural DEs. Unlike previous methods that focus on individual token representations, our approach leverages temporal information in state dynamics. We conduct an extensive study on two challenging hallucination datasets, evaluating our method and baseline approaches using four widely adopted public LLMs. The results demonstrate the effectiveness of our approach. Our contributions are summarized as follows:

- We introduce HD-NDEs, the first method to apply Neural DEs, including Neural ODEs, CDEs, and SDEs, for detecting hallucinations in LLMs. By modeling the token generation process as continuous trajectories in latent space, HD-NDEs provides a more accurate and dynamic approach to detecting hallucinations.

- We evaluate HD-NDEs on two diverse and complex hallucination datasets and compare their performance with baseline methods across four widely used public LLMs. Our results demonstrate that HD-NDEs outperforms existing approaches over a 10% improvement in True-False Dataset.

2 Related Work

Hallucination Detection. Hallucinations in LLMs pose significant challenges for their deployment (Zhang et al., 2023; Li et al., 2023). The generation of inaccurate information can result in customer attrition or legal risks, rendering the decision-making process unreliable. Detecting hallucinations has garnered increasing attention, and this detection is typically performed in one of the following ways: conducting a conventional fact-checking task (Min et al., 2023; Wang et al., 2023), which requires external knowledge; converting the logits output into an uncertainty estimate for the sentence; or evaluating self-consistency (Mündler et al., 2023), where inconsistent outputs often indicate hallucinations. Recent studies have shown that hallucinations can be attributed to the model's internal representations and have proposed white-box methods to detect or predict hallucinations based on token latent states (Burns et al., 2023; Azadi et al., 2023; Zhu et al., 2024). These supervised approaches have outperformed black-box methods across various tasks. However, as noted in Levinstein and Herrmann (2024), they often struggle when non-factual tokens appear at the beginning or middle of the sequence.

Neural Differential Equations have been extensively used in modeling dynamical systems or simulated neural networks (Chang et al., 2018; Lu et al., 2018; Chen et al., 2018; Lu et al., 2019; Dutta et al., 2021). Lu et al. (Lu et al., 2018) showed that any parametric ordinary differential equation (ODE) solver can be conceptualized as a deep learning framework with infinite depth. Chen et al. (2018) achieved ResNet-comparable results with a drastically lower number of parameters and memory

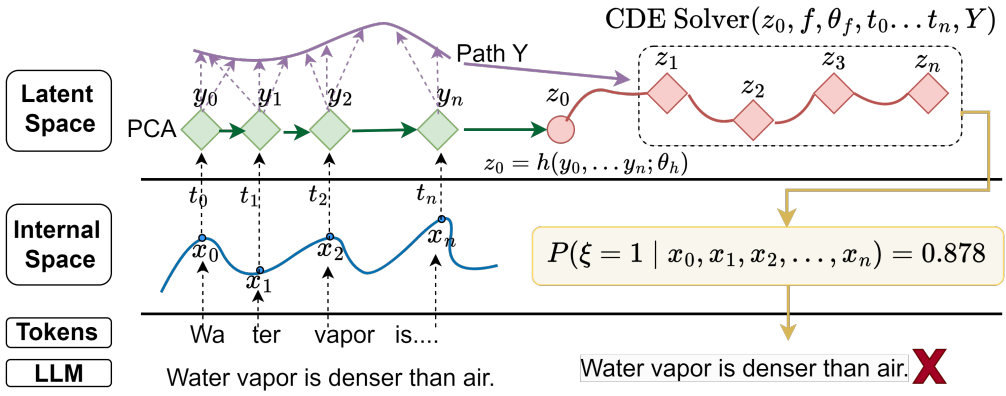

Figure 2: Computation graph of HD-NDEs detecting hallucination via Neural CDEs. The statement is processed by LLMs, from which we extract embedding information of each token in the internal space to construct the trajectory $\boldsymbol{x} = (x_0, x_1, x_2, ..., x_n)$, with corresponding time points $(t_0, t_1, ..., t_n)$. The PCA processes the trace and generates states $\boldsymbol{y} = (y_0, y_1, y_2, ..., y_n)$. These states are used to parameterize a latent space representation z_0 and extract control path Y. The CDE solver predicts future latent states to form $\boldsymbol{z} = (z_0, z_1, z_2, ..., z_n)$, leading to a final incorrect factuality score $P(\xi = 1|\boldsymbol{x})$.

complexity by parameterizing hidden layer derivatives and using ODE solvers. Lu et al. (2019) was the first to draw analogies between transformers and dynamical systems, conceptualizing the transformer as a numerical approximation of an ODE. Furthermore, Neural DEs play important roles in interpolation, forecasting, and classification tasks in time series data (Kidger et al., 2020a; Liang et al., 2021; Li et al., 2020; Oh et al., 2024; Li et al., 2022).

3 Methodology: HD-NDEs

We denote the generated text as a sequence of tokens $o_{0:n} = (o_0, o_1, ..., o_n)$, where o_t represents the t-th token. Given a generated text sample $\mathbf{o} = o_{0:n}$, our objective is to predict $P(\xi|\mathbf{o})$ where $\xi \in \{0, 1\}$ serves as the hallucination indicator variable, with $\xi = 1$ indicating a hallucination and $\xi = 0$ otherwise. Naturally, each token o_t is associated with an internal state representation $x_t \in \mathbb{R}^{d_x}$, derived from the specific hidden layer embeddings corresponding to token t. We generally use the embedding from the last layer to represent each token, where d_x denotes the embedding dimension. The value of d_x varies across models; for instance, $d_x = 4096$ for LLama-7B, while $d_x = 5120$ for LLama-13B.

3.1 Neural DEs

To capture the dynamic behavior of LLMs, we utilize Neural ODEs (Chen et al., 2018), Neural CDEs (Kidger et al., 2020a), and Neural SDEs (Oh et al., 2024) to model the evolution in the latent space. Figure 2 illustrates hallucination detection using HD-NDEs with Neural CDEs. Neural ODEs describe smooth, continuous-time dynamics using deterministic equations, neural CDEs introduce control signals to guide system evolution. Furthermore, neural SDEs incorporate stochasticity to account for uncertainty or noise within the system.

Neural ODEs. Let $\boldsymbol{x} = x_{0:n} = (x_0, ..., x_n) \in \mathbb{R}^{d_x}$ denote the embeddings in the internal space. $\boldsymbol{x}$ is projected into $\boldsymbol{y} = y_{0:n} = (y_0, ..., y_n) \in \mathbb{R}^{d_y}$ by PCA. Consider a latent representation $z(t) \in \mathbb{R}^{d_z}$ at time t in latent space, which is given by

$$z(t) = z(0) + \int_0^t f(s, z(s); \theta_f) ds \quad \text{with} \quad z(0) = h(\boldsymbol{y}; \theta_h), \tag{1}$$

where $h : \mathbb{R}^{d_y} \to \mathbb{R}^{d_z}$ is an function with parameter θ_h and $f(t, z(t); \theta_f)$ is a neural network parameterized by θ_f to approximate $\frac{d\boldsymbol{z}(t)}{dt}$. Neural ODEs rely on ODE solvers, such as the explicit Euler method (Euler, 1845), to solve the integral problem in (1). Since we can freely choose the upper limit t of the integration, we can predict z at any time t. That is, once $h(\cdot; \theta_h)$ and $f(\cdot, \cdot; \theta_f)$ have been learned, then we are able to compute $\boldsymbol{z}(t)$ for any $t \geq 0$.

Neural CDEs. The solution to Neural ODEs is determined by its initial condition, making it inadequate for incorporating incoming information into a differential equation. To address this issue,

Kidger et al. (2020b) proposed Neural CDEs by combining a controlled path $Y(t)$ of the underlying time-series data. Specifically, given the sequential data $\boldsymbol{y} = (y_0, y_1, \ldots, y_n)$, $z(t)$ is determined by

$$\begin{aligned} z(t) &= z(0) + \int_0^t f(s, z(s); \theta_f) dY(s) \\ \text{with} \quad z(0) &= h(\boldsymbol{y}; \theta_h), \end{aligned} \tag{2}$$

where $Y(t)$ is chosen as a natural cubic spline path (Kidger et al., 2020b) or hermite cubic splines with backward differences (James et al., 2021) of the underlying time-series data. Differently from Neural ODEs, $f(t, z(t); \theta_f)$ is a neural network parameterized by θ_f to approximate $\frac{dz(t)}{dY(t)}$.

Neural SDEs. Neural SDEs allow for describing the stochastic evolution of trace, rather than the deterministic evolution (Kidger et al., 2021b,a). The latent representation $z(t)$ of Neural SDEs is governed by the following SDE:

$$\begin{aligned} z(t) &= z(0) + \int_0^t f(s, z(s); \theta_f) ds \\ &\quad + \int_0^t g(s, z(s); \theta_g) dW(s) \\ \text{with} \quad z(0) &= h(\boldsymbol{y}; \theta_h), \end{aligned} \tag{3}$$

where $\{W_t\}_{t \geq 0}$ is a d_z-dimensional Brownian motion, $f(\cdot, \cdot; \theta_f)$ is the drift function, and $g(\cdot, \cdot; \theta_g)$ is the diffusion function. Drift and diffusion functions are represented by neural networks.

3.2 DEs Solvers and Adjoint Sensitivity Methods for Propagation

For simplicity, we denote the parameters of neural networks used in Equations (1), (2) and (3) as $\boldsymbol{\theta}$, where $\boldsymbol{\theta} = (\theta_h, \theta_f)$ in Neural ODEs and Neural CDEs, and $\boldsymbol{\theta} = (\theta_h, \theta_f, \theta_g)$ in Neural SDEs. After choosing one from Neural ODEs, Neural CDEs, or Neural SDEs to capture the state generation process in latent space, two natural questions arise: (1) How can we generate the subsequent latent states $(z(t_1), z(t_2), ...)$ based on $z(0)$ and $\boldsymbol{\theta}$? (2) How can we update the parameters $\boldsymbol{\theta}$ to ultimately obtain the optimal solution $\boldsymbol{\theta}^*$? Sections 3.2.1 and 3.2.2 answer the above two questions respectively.

3.2.1 DE Solvers for Forward Propagation

We begin by introducing two common ODE solvers: First-Order and High-Order schemes. Numerical methods for Neural CDEs and Neural SDEs can be adapted from these approaches.

First-Order ODE Solvers. The Euler method (Euler, 1845) is the simplest method for solving ODEs. The transformation at each time step can be expressed as:

$$\begin{aligned} z_{t+1} &= z_t + \frac{dz(t)}{dt} \\ &= z_t + f(t, \boldsymbol{z}(t); \theta_f). \end{aligned} \tag{4}$$

High-Order ODE Solvers. The Euler method is not "precise" because it is a first-order method, and naturally with local truncation errors. The global error will be accumulated if we want to capture a long timestep trajectory. Herein, we use the Runge-Kutta method (Runge, 1895) for a higher-order solution to ODEs. They are a classic family of iterative methods with different orders of precision. More formally, the explicit Runge-Kutta methods of an n-step solution are defined to be:

$$\begin{aligned} z_{t+1} &= z_t + \sum_{i=1}^{n} \gamma_i Z_i \\ Z_1 &= \Delta t f(t, z_t; \theta_f) \\ Z_i &= \Delta t f(t + \alpha_i \Delta t, z_t + \sum_{j=1}^{i-1} \beta_{ij} Z_j; \theta_f) \end{aligned} \tag{5}$$

where Δt is the time size and could be simply 1 in most cases. Z_i is an intermediate approximation to the solution at step $t + \alpha_i \Delta t$. α, β and γ are coefficients which can be determined by the series of z_{t+1}.

The Runge-Kutta methods are a general class of numerical approaches, with the Euler method being a first-order instance. In this work, we use fourth-order Runge-Kutta (RK4) for solving Equation (1), details in Appendix A. The integral problem in (2) can be solved by using the above-mentioned ODE solvers since $\frac{dz(t)}{dt} = f(t, z(t); \theta_f) \frac{dX(t)}{dt}$. However, the Neural SDE problem (3) requires additional handling of stochastic noise, making its solution methods more complex. Herein, we use the Euler-Maruyama method designed to handle noise terms, which is given by

$$z_{t+1} = z_t + f(t, z(t); \theta_f) + g(t, z(t); \theta_f) \mathcal{Z}, \tag{6}$$

where $\mathcal{Z} \sim \mathcal{N}(0, 1)$ is a standard normal random variable with mean 0 and variance 1.

3.2.2 Adjoint Sensitivity Methods for Back Propagation

Since Neural DEs are continuous-time models computed through DE solvers, standard backpropagation cannot be directly applied. Chen et al. (2018)

Method	Company*				Fact*			
	LLama-7B	LLama-13B	Alpaca-13B	Vicuna-13B	LLama-7B	LLama-13B	Alpaca-13B	Vicuna-13B
Logit-based Methods								
AvgProb	59.0	59.2	53.0	60.2	59.5	59.3	54.2	58.3
AvgEnt	54.0	56.4	54.2	53.3	54.2	54.1	50.3	51.2
MaxProb	56.0	57.2	49.3	59.5	57.1	58.4	50.2	58.2
MaxEnt	51.1	54.0	52.2	50.0	53.2	52.6	51.4	52.3
Supervised Learning Method								
Probing	54.0	58.2	59.3	68.2	58.3	62.4	59.8	65.5
ODEs	59.7	65.3	67.8	72.9	58.6	66.9	64.3	70.4
CDEs	65.9	72.8	**75.3**	**79.8**	67.5	**74.8**	**72.9**	76.7
SDEs	**73.8**	**78.4**	70.5	72.3	**70.3**	73.1	70.3	**78.6**

Method	City*				Invention*			
	LLama-7B	LLama-13B	Alpaca-13B	Vicuna-13B	LLama-7B	LLama-13B	Alpaca-13B	Vicuna-13B
Logit-based Methods								
AvgProb	54.2	56.3	51.5	59.2	51.2	51.3	49.4	55.3
AvgEnt	49.1	50.2	49.3	52.4	45.1	46.3	48.4	47.5
MaxProb	52.6	55.4	55.3	53.5	52.3	53.1	49.6	58.3
MaxEnt	50.7	49.3	50.2	51.6	48.5	47.3	47.2	46.3
Supervised Learning Methods								
Probing	60.0	69.3	59.4	64.5	59.2	66.0	52.4	69.3
ODEs	73.0	82.3	71.2	73.2	60.3	80.9	69.7	80.4
CDEs	75.7	80.6	72.1	80.1	**75.9**	**88.3**	73.8	81.2
SDEs	**79.1**	**89.8**	**74.3**	**82.5**	68.7	79.6	**74.2**	**85.9**

Table 1: The detection AUC-ROC (%) for different approaches over multiple benchmark LLMs over Company*, Fact*, City* and Invention*. **Bold** and underlined numbers denote the best and second-best values, respectively. ODEs, CDEs, and SDEs are the abbreviations of Neural ODEs, Neural CDEs, and Neural SDEs, respectively.

applied the adjoint sensitivity method (Pontryagin et al., 1962) to compute gradients for Neural ODEs. Specifically, to optimize the loss function L, we require gradients with respect to θ. The first step is to determine how the gradient of the loss depends on the hidden state $z(t)$ at each instant. This quantity is called the adjoint

$$a(t) = \frac{\partial L}{\partial z(t)}. \tag{7}$$

Its dynamics are given by another ODE, which can be thought of as the instantaneous analog of the chain rule:

$$\frac{da(t)}{dt} = -\alpha(t)^T \frac{\partial f(t, z(t); \theta_f)}{\partial z}. \tag{8}$$

We can compute $a(t)$ by another call to an ODE solver. Computing the gradients with respect to the parameters θ requires evaluating a third integral, which depends on both $z(t)$ and $a(t)$:

$$\frac{dL}{d\theta} = -\int_{t_1}^{t_0} \alpha(t)^T \frac{\partial f(t, z(t), \theta_f)}{\partial z}. \tag{9}$$

In addition, Kidger et al. (2020a) and Li et al. (2020) proposed the adjoint sensitivity methods for Neural CDEs and Neural SDEs, respectively. In our work, we build upon the above methods to update the parameters of neural networks.

4 Experimental Setup

4.1 Datasets

To demonstrate the applicability of our approach, we use two datasets: True-False and TruthfulQA. We expand the original True-False dataset to include statements that are nearly identical in expression but opposite in factuality, differing by only a single word. For TruthfulQA, we randomly pair both a correct and an incorrect answer for each question in the dataset.

True-False Dataset. The original dataset consists of six sub-datasets, each named after its subject matter (Azaria and Mitchell, 2023b). We follow the method proposed in Levinstein and Herrmann (2024) to create factual and non-factual statements containing subtle differences. Specifically, we prompt GPT-4o to generate new statements that are factually opposite to the original while maintaining only minor word differences. For example, we obtain a non-factual statement "The earth doesn't orbit the sun." from the factual statement "The earth orbits the sun." For our experiments, we randomly select 550, 560, 500, and 500 statements from the *Companies*, *Scientific Facts*,

Method	TruthfulQA			
	LLama-7B	LLama-13B	Alpaca-13B	Vicuna-13B
Logit-based Methods				
AvgProb	51.4	54.6	53.3	55.1
AvgEnt	49.4	53.0	52.7	53.6
MaxProb	48.5	49.2	49.3	48.1
MaxEnt	53.1	53.8	52.5	53.6
Consistency-based Methods				
QA	54.8	56.3	55.4	57.9
BERTScore	53.2	59.9	55.2	60.6
Unigram	57.6	62.2	60.1	63.4
NLI	60.6	63.7	61.6	65.1
Supervised Learning Methods				
Probing	87.5	86.3	84.9	88.6
ODEs	84.2	87.9	83.1	83.8
CDEs	86.7	84.0	84.3	89.2
SDEs	**88.3**	**89.3**	**86.4**	**89.5**

Table 2: The detection AUC-ROC (%) for different approaches over multiple benchmark LLMs on TruthfulQA.

Cities, and *Inventions* sub-datasets, respectively. The resulting datasets are referred to as *Company**, *Fact**, *City**, and *Invention**. This dataset poses a greater challenge for hallucination detection.

TruthfulQA (Lin et al., 2022) dataset is a widely used benchmark consisting of 873 questions, each with multiple correct and incorrect answers. Unlike True-False, we use the Levenshtein (Levenshtein, 1966) distance to select the pair of correct and incorrect answers with the greatest textual difference. These pairs, along with the original questions, form the data used for our experiments.

4.2 Models

We evaluate both our method and baseline approaches using common open-source LLMs, including LLama-7B (Touvron et al., 2023), LLama-13B (Touvron et al., 2023), Alpaca-13B (Taori et al., 2023), and Vicuna-13B (Chiang et al., 2023). Due to computational resource constraints, we don't conduct experiments with larger models.

4.3 Our Approach and Baselines

We compare our approach HD-NDEs with the following widely-used baselines, including both white-box and black-box methods.

HD-NDEs. To reduce computational complexity, we employ PCA to reduce the dimensionality of the internal space to $K = 1024$. The integrands $h(\cdot;\theta_h)$, $f(\cdot,\cdot;\theta_f)$, $g(\cdot,\cdot;\theta_g)$ in Equations (1), (2) and (3) are taken to be feedforward neural networks. Specifically, we use a single hidden layer network to represent $h(\cdot;\theta_h)$ in all variants of our methods. We use an 8-layer neural network to represent $f(\cdot,\cdot;\theta_f)$ in Neural CDEs and 10-layer neural networks for $f(\cdot,\cdot;\theta_f)$ in both Neural ODEs and Neural SDEs. Additionally, $g(\cdot,\cdot;\theta_g)$ in Neural SDEs is represented by a 4-layer neural network. A final linear layer is always applied to map the latent state to the output. We use ReLU activation functions for Neural CDEs and Neural SDEs, while tanh activations are used for Neural ODEs. The binary cross-entropy loss is applied to the sigmoid of the model output. Additionally, we employ the Adam optimizer with a learning rate of 0.001, a batch size of 32, and 50 epochs.

Logit-based methods are extensively studied for hallucination detection in LLMs (Hu et al., 2023; Huang et al., 2023). We adopt the four proposed metrics: *AvgProb*, *AvgEnt*, *MaxProb*, *MaxEnt* to aggregate token-level logit-based uncertainty to measure sentence uncertainty. Appendix B.1 provides detailed descriptions of these metrics.

Consistency-based methods are motivated by the idea that if an LLM possesses specific knowledge, the sampled responses are likely to be similar and contain consistent facts (Manakul et al., 2023). We apply four variants – *BERTScore*, *Question Answering* (QA), *Unigram*, *Natural Language Inference* (NLI) – of the consistency-based method proposed in (Manakul et al., 2023) in this work. More details are shown in the Appendix B.2.

Supervised learning methods train a classifier on a dataset containing labeled statements. The probing method proposed in (Azaria and Mitchell, 2023b) is the classic approach. Its classifier employs a feedforward neural network featuring three hidden layers with decreasing numbers of hidden units (1024, 512, 256), all utilizing ReLU activations. The final layer is a sigmoid output. The statements are run through LLMs and extracted embeddings during the process. The probe receives the embedding for the last token from the last layer of LLMs.

4.4 Evaluation Metric

We utilize **AUC-ROC**, which stands for the area under the ROC curve, to objectively evaluate the effectiveness of models. The higher value of AUC-PR, the stronger the ability of this method for hallucination detection. All experiments are conducted on NVIDIA A100 GPUs with 40GB memory.

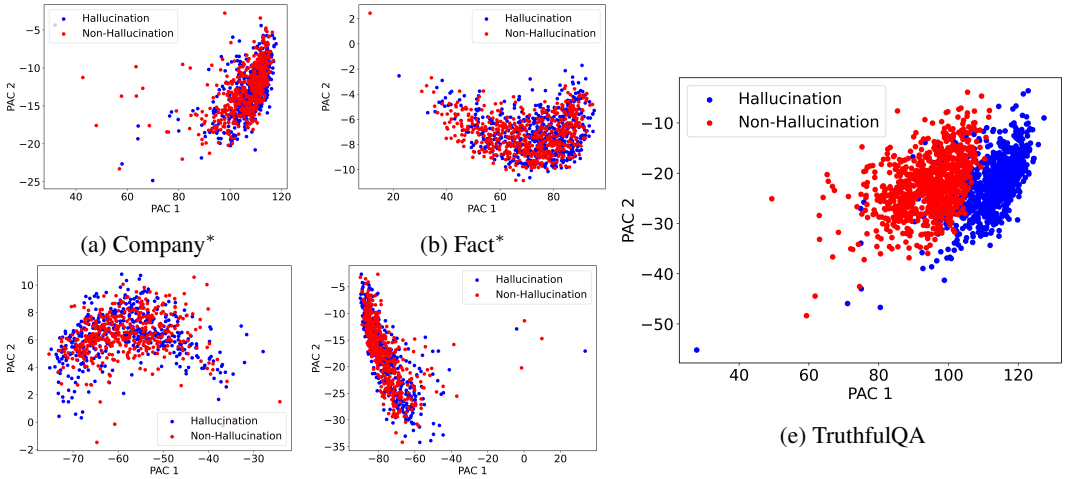

(a) Company*

(b) Fact*

(c) City*

(d) Invention*

(e) TruthfulQA

Figure 3: 2D PCA projection of the last hidden layer's embedding for the final token on Company*, Fact*, City*, Invention* and TruthfulQA. Blue and red dots represent hallucinations and non-hallucinations, respectively.

5 Experimental Results and Analysis

5.1 Comprehensive Results

True-False Dataset The comprehensive results are demonstrated in Table 1. Since consistency-based methods rely on question-and-answer pairs, and the True-False dataset is not structured in this format, we do not include this type of method as a comparison in this dataset. It is obvious that our proposed methods surpass the probing method by a noticeable margin, evidenced by an average increase of over 10% in the detection of AUC-ROC across different models and datasets. Particularly, Neural CDEs outperforms the probing method by 22.3% on the Invention* dataset when using the LLama-13B. Even in the worst case, Neural ODEs performs comparably to the probing method on the Fact* dataset based LLama-7B. Furthermore, in most cases, the four logit-based methods perform worse than the supervised learning methods.

Among our methods, we can find that Neural CDEs and Neural SDEs outperform Neural ODEs. As shown in Table 1, the best and second-best values are achieved by Neural CDEs and Neural SDEs models in 13 out of 16 cases. The likely reason is that Neural CDEs and Neural SDEs can capture richer dynamical behaviors than Neural ODEs. As mentioned in Section 3, Neural CDEs incorporates control theory, enabling the dynamic system to account for the influence of incoming information, while Neural SDEs introduces stochasticity into the modeling process. These capabilities enable both techniques to capture nuanced variations in the data over time, resulting in improved performance in hallucination detection. While Neural ODEs assumes deterministic dynamics, which can limit its flexibility in modeling the dynamics of LLMs.

TruthfulQA Table 2 show the results of our method, alongside all baseline methods. The probing method's performance on TruthfulQA has significantly improved across four models compared to the True-False Dataset. Moreover, it demonstrates comparable performance to HD-NDEs, including Neural ODEs, Neural CDEs, and Neural SDEs, with a difference of less than 3%. Particularly, the probing method outperforms Neural ODEs and Neural CDEs in the LLama-7B and Alpaca-13B while being slightly weaker than Neural SDEs.

In addition, we would like to highlight several observations that hold potential implications for future developments in this field. A general trend can be identified that in most instances, the consistency-based approaches surpass the logit-based metrics, yet they remain significantly inferior to the supervised learning methods by approximately 20%. Supervised learning methods perform well because they leverage labeled data during training, allowing them to inherently capture more information than logit-based and consistency-based approaches. In addition, NLI outperforms best among consistency-based methods.

5.2 Analysis

We try to understand why HD-NDEs obviously outperforms the probing method on the True-False

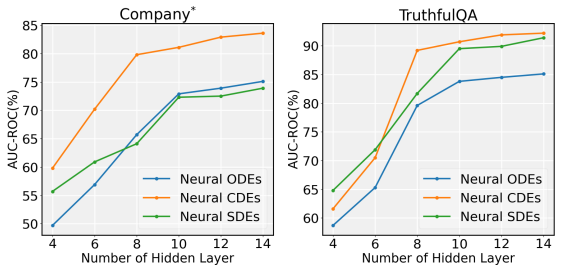

Figure 4: The impact of the number of hidden layers on the AUC-ROC of HD-NDEs on Vicuna-13B: Company* and TruthfulQA.

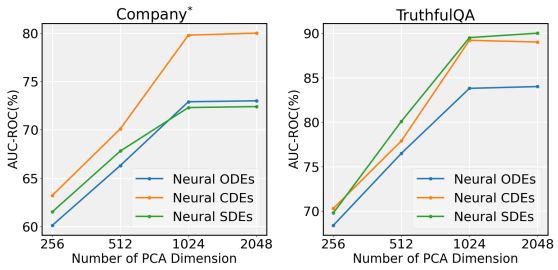

Figure 5: The impact of the PCA projection dimensions on the AUC-ROC of HD-NDEs on Vicuna-13B: Company* and TruthfulQA.

dataset, yet performs comparably to probing methods on the TruthfulQA dataset. We employ PCA to reduce the dimensions of the hidden embeddings, retaining the two dominant components of True-False and TruthfulQA. The results are shown in Figure 3. The 2D PCA projection reveals a significant overlap between correct and incorrect statements in True-False, with many points intertwined. The poor separation causes the probing method to perform only marginally better than random guessing in many cases, highlighting the need to consider all tokens and integrate temporal information by analyzing state dynamics. In contrast, the 2D PCA projection of TruthfulQA reveals a much clearer distinction between hallucination and non-hallucination. The statements in TruthfulQA exhibit substantial variation, as we use the Levenshtein distance to select statements with significant differences. This allows the probing method to more easily differentiate them based on the embeddings of the final token. This also accounts for its comparable performance to HD-NDEs.

5.3 Ablation Studies

In this section, we investigate several factors that may be critical for the detection performance and practicality of HD-NDEs. We adhere to the default configuration for all the experiments in this section unless stated otherwise.

Number of Hidden Layers. An important factor impacting the performance of detection methods is the number of hidden layers in neural networks representing $f(s, z(s); \theta_f)$. Therefore, we investigate the effectiveness of our approach across varying numbers of hidden layers, as shown in Figure 4. We notice an increase in detection effectiveness with more hidden layers, likely due to improved modeling capacity and expressive power. Specifically, the performance of Neural CDEs improves substantially as the number of layers increases up to 8, with further increases beyond 8 still showing gains but at a slower pace. For Neural ODEs and Neural SDEs, this turning point occurs when the layer number reaches 10, based on the results from both datasets. Considering the balance between increased computational resources demands and performance gains as the number of hidden layers increases, we ultimately set 8 layers for Neural CDEs and 10 layers for both Neural ODEs and Neural SDEs in Section 4.3.

Dimensions in Latent Space. Another key factor is the dimension of the latent state after being mapped from the internal space to the latent space by PCA. We then examine the impact of varying dimensions as shown in Figure 5. An evident improvement in detection effectiveness is associated with retaining more components during down-projection. We hypothesize that this trend can be largely attributed to the preservation of a more substantial amount of information when expanding the latent space. Importantly, all three variants of our methods achieve the best performance on both datasets when the dimension is set to 1024, affirming our hyperparameter setting in Section 4.3.

6 Conclusion

In this paper, we introduce HD-NDEs, which leverages various types of neural differential equations to detect hallucinations in language models. By tracking the dynamic changes in latent space, HD-NDEs can effectively detect logical or factual inconsistencies that arise in the generated text. Comprehensive empirical results confirm the effectiveness of our methods. Our approach surpasses various unsupervised methods and probing-based supervised methods by over 10% on the True-False.

Limitations

This work identifies three major limitations. First, the model's training process takes longer compared to classical probing methods. Second, NeuralDEs, as currently presented, do not provide uncertainty estimates for their predictions, though such extensions may be feasible in the future. Third, we experiment with a limited set of numerical schemes, and other methods could potentially exploit the structure of differential equations to further improve performance.

Ethics and Broader Impact

We sampled a portion of the data from existing datasets for our experiments, which may affect the accuracy of some of our conclusions.

References

Hervé Abdi and Lynne J Williams. 2010. Principal component analysis. *Wiley interdisciplinary reviews: computational statistics*, 2(4):433–459.

Fatemeh Azadi, Heshaam Faili, and Mohammad Javad Dousti. 2023. PMI-align: Word alignment with point-wise mutual information without requiring parallel training data. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 12366–12377, Toronto, Canada. Association for Computational Linguistics.

Amos Azaria and Tom Mitchell. 2023a. 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.

Amos Azaria and Tom Mitchell. 2023b. 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.

Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. 2023. Discovering latent knowledge in language models without supervision. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

Bo Chang, Lili Meng, Eldad Haber, Frederick Tung, and David Begert. 2018. Multi-level residual networks from dynamical systems view. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.

Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. 2018. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 6572–6583.

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.

Jeongwhan Choi, Hwangyong Choi, Jeehyun Hwang, and Noseong Park. 2022. Graph neural controlled differential equations for traffic forecasting. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 6367–6374. AAAI Press.

Jinhao Duan, Hao Cheng, Shiqi Wang, Chenan Wang, Alex Zavalny, Renjing Xu, Bhavya Kailkhura, and Kaidi Xu. 2023. Shifting attention to relevance: Towards the uncertainty estimation of large language models. *ArXiv preprint*, abs/2307.01379.

Subhabrata Dutta, Tanya Gautam, Soumen Chakrabarti, and Tanmoy Chakraborty. 2021. Redesigning the transformer architecture with insights from multi-particle dynamical systems. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 5531–5544.

Leonhard Euler. 1845. *Institutionum calculi integralis*, volume 4. impensis Academiae imperialis scientiarum.

Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koeppl, Preslav Nakov, and Iryna Gurevych. 2023. A survey of language model confidence estimation and calibration. *ArXiv preprint*, abs/2311.08298.

Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2023. Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

Mengting Hu, Zhen Zhang, Shiwan Zhao, Minlie Huang, and Bingzhe Wu. 2023. Uncertainty in natural language processing: Sources, quantification, and applications. *ArXiv preprint*, abs/2306.04459.

Yuheng Huang, Jiayang Song, Zhijie Wang, Huaming Chen, and Lei Ma. 2023. Look before you leap: An exploratory study of uncertainty measurement for large language models. *ArXiv preprint*, abs/2307.10236.

Morrill James, Kidger Patrick, Yang Lingyi, and Lyons Terry. 2021. Neural controlled differential equations for online prediction tasks. *ArXiv preprint*, abs/2106.11028.

Sheo Yon Jhin, Heejoo Shin, Sujie Kim, Seoyoung Hong, Minju Jo, Solhee Park, Noseong Park, Seungbeom Lee, Hwiyoung Maeng, and Seungmin Jeon. 2024. Attentive neural controlled differential equations for time-series classification and forecasting. *Knowledge and Information Systems*, 66(3):1885–1915.

Patrick Kidger. 2022. On neural differential equations. *ArXiv preprint*, abs/2202.02435.

Patrick Kidger, James Foster, Xuechen Li, and Terry J. Lyons. 2021a. Efficient and accurate gradients for neural sdes. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 18747–18761.

Patrick Kidger, James Foster, Xuechen Li, and Terry J. Lyons. 2021b. Neural sdes as infinite-dimensional gans. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 5453–5463. PMLR.

Patrick Kidger, James Morrill, James Foster, and Terry J. Lyons. 2020a. Neural controlled differential equations for irregular time series. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.

Patrick Kidger, James Morrill, James Foster, and Terry J. Lyons. 2020b. Neural controlled differential equations for irregular time series. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.

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*, abs/2406.15927.

V Levenshtein. 1966. Binary codes capable of correcting deletions, insertions, and reversals. *Proceedings of the Soviet physics doklady*.

Benjamin A Levinstein and Daniel A Herrmann. 2024. Still no lie detector for language models: Probing empirical and conceptual roadblocks. *Philosophical Studies*, pages 1–27.

Bei Li, Quan Du, Tao Zhou, Yi Jing, Shuhan Zhou, Xin Zeng, Tong Xiao, JingBo Zhu, Xuebo Liu, and Min Zhang. 2022. ODE transformer: An ordinary differential equation-inspired model for sequence generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8335–8351, Dublin, Ireland. Association for Computational Linguistics.

Junyi Li, Xiaoxue Cheng, Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2023. HaluEval: A large-scale hallucination evaluation benchmark for large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6449–6464, Singapore. Association for Computational Linguistics.

Qing Li, Chenyang Lyu, Jiahui Geng, Derui Zhu, Maxim Panov, and Fakhri Karray. 2024. Reference-free hallucination detection for large vision-language models. *ArXiv preprint*, abs/2408.05767.

Xuechen Li, Ting-Kam Leonard Wong, Ricky T. Q. Chen, and David Duvenaud. 2020. Scalable gradients for stochastic differential equations. In *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020, 26-28 August 2020, Online [Palermo, Sicily, Italy]*, volume 108 of *Proceedings of Machine Learning Research*, pages 3870–3882. PMLR.

Yuxuan Liang, Kun Ouyang, Hanshu Yan, Yiwei Wang, Zekun Tong, and Roger Zimmermann. 2021. Modeling trajectories with neural ordinary differential equations. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 1498–1504. ijcai.org.

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, Dublin, Ireland. Association for Computational Linguistics.

Yiping Lu, Zhuohan Li, Di He, Zhiqing Sun, Bin Dong, Tao Qin, Liwei Wang, and Tie-Yan Liu. 2019. Understanding and improving transformer from a multi-particle dynamic system point of view. *ArXiv preprint*, abs/1906.02762.

Yiping Lu, Aoxiao Zhong, Quanzheng Li, and Bin Dong. 2018. Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 3282–3291. PMLR.

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, Singapore. Association for Computational Linguistics.

Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. FActScore: Fine-grained atomic evaluation of factual precision in long form text generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100, Singapore. Association for Computational Linguistics.

Niels Mündler, Jingxuan He, Slobodan Jenko, and Martin Vechev. 2023. Self-contradictory hallucinations of large language models: Evaluation, detection and mitigation. *ArXiv preprint*, abs/2305.15852.

YongKyung Oh, Dongyoung Lim, and Sungil Kim. 2024. Stable neural stochastic differential equations in analyzing irregular time series data. *ArXiv preprint*, abs/2402.14989.

Lev Semenovich Pontryagin, EF Mishchenko, VG Boltyanskii, and RV Gamkrelidze. 1962. *The mathematical theory of optimal processes*. Routledge.

Carl Runge. 1895. Über die numerische auflösung von differentialgleichungen. *Mathematische Annalen*, 46(2):167–178.

Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *ArXiv preprint*, abs/2307.09288.

Yuxia Wang, Revanth Gangi Reddy, Zain Muhammad Mujahid, Arnav Arora, Aleksandr Rubashevskii, Jiahui Geng, Osama Mohammed Afzal, Liangming Pan, Nadav Borenstein, Aditya Pillai, et al. 2023. Factcheck-gpt: End-to-end fine-grained document-level fact-checking and correction of llm output. *ArXiv preprint*, abs/2311.09000.

Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu, Nathan Hu, Dustin Tran, Daiyi Peng, Ruibo Liu, Da Huang, Cosmo Du, and Quoc V. Le. 2024. Long-form factuality in large language models. *ArXiv preprint*, abs/2403.18802.

Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. 2023. Siren's song in the ai ocean: a survey on hallucination in large language models. *ArXiv preprint*, abs/2309.01219.

Derui Zhu, Dingfan Chen, Qing Li, Zongxiong Chen, Lei Ma, Jens Grossklags, and Mario Fritz. 2024. PoLLMgraph: Unraveling hallucinations in large language models via state transition dynamics. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 4737–4751, Mexico City, Mexico. Association for Computational Linguistics.

A RK4

We can also define a fourth-order Runge-Kutta (RK4) block to be:

$$\begin{aligned} z_{t+1} &= z_t + \frac{\Delta t}{6}(Z_1 + 2Z_2 + 2Z_3 + Z_4) \\ Z_1 &= f(t, z_t; \theta_f) \\ Z_2 &= f(t + \frac{\Delta t}{2}, z_t + \frac{\Delta t}{2} Z_1; \theta_f) \\ Z_3 &= f(t + \frac{\Delta t}{2}, z_t + \frac{\Delta t}{2} Z_2; \theta_f) \\ Z_4 &= f(t + \Delta t, z_t + \Delta t Z_3; \theta_f) \end{aligned} \tag{10}$$

B Logit-based and Consistency-based Methods

We collect logit-based and consistency-based methods proposed in (Manakul et al., 2023) to test the effectiveness of models.

B.1 Logit-based methods

To aggregate the uncertainty information obtained at the token level, we employ four metrics to aggregate token-level uncertainty into sentence level. In particular, a sentence-level uncertainty score can be obtained by taking either the maximum or average of the negative loglikelihood $-\log p_{ij}$ in a sentence:

$$\text{MaxProb}(i) = \max_j(-\log p_{ij}), \tag{11}$$

$$\text{AvgProb}(i) = -\frac{1}{J_i}\sum_{j=1}^{J_i} \log p_{ij}, \tag{12}$$

where p_{ij} is the probability of a token at a position j in the sentence i and J_i is the total number of tokens in the considered sentence. Additionally, one can also replace the negative loglikelihood $-\log p_{ij}$ with the entropy $\mathcal{H}_{ij}$:

$$\text{MaxEnt}(i) = \max_j \mathcal{H}_{ij}, \tag{13}$$

$$\text{AvgEnt}(i) = \frac{1}{J_i}\sum_{j=1}^{J_i} \mathcal{H}_{ij}, \tag{14}$$

where H_{ij} is the entropy of the token distribution for the j-th token in the sentence i.

B.2 Consistency-based Methods

BERTScore denoted as $\mathcal{S}_{BERT}$, finds the average BERTScore of the considered sentence with the most similar sentence from each drawn sample. Let $\mathcal{B}(.,.)$ denote the BERTScore between two sentences. $\mathcal{S}_{BERT}$ is computed by $\mathcal{B}(.,.)$ as:

$$\mathcal{S}_{\text{BERT}}(i) = 1 - \frac{1}{N}\sum_{n=1}^{N} \max_k \mathcal{B}(r_i, s_k^n), \tag{15}$$

where r_i represents the i-th sentence in main response R, s_k^n represents the k-th sentence in the n-th sample S^n and N is the total number of sampled sentences.

Question Answering (QA). The multiple-choice question answering generation (MQAG) evaluates consistency by creating multiple-choice questions from the primary generated response and answering them using other sampled responses as reference. MQAG consists of two stages: question generation G and question answering A. For the sentence r_i in the response R, we draw questions q and options $\mathbf{o}$:

$$q, \mathbf{o} \sim P_G(q, \mathbf{o} \mid r_i, R). \tag{16}$$

The answering stage A selects the answers:

$$a_R = \arg\max_k [P_A(o_k \mid q, R, \mathbf{o})], \quad (17)$$

$$a_{S^n} = \arg\max_k [P_A(o_k \mid q, S^n, \mathbf{o})]. \quad (18)$$

We compare whether a_R is equal to a_{S^n} for each sample in $\{S^1, \ldots, S^N\}$, yielding N_m matches and N_n not-matches. A simple inconsistency score for the i-th sentence and question q based on the match/not-match counts is defined: $S_{\text{QA}}(i, q) = \frac{N_n}{N_n + N_m}$. Manakul et al. (2023) modify the inconsistency score $S_{\text{QA}}(i, q)$ by incorporating the answerability of generated questions. Finally, $\mathcal{S}_{\text{QA}}(i)$ is the average of inconsistency scores across q,

$$\mathcal{S}_{\text{QA}}(i) = E_q[S_{\text{QA}}(i, q)]. \quad (19)$$

Unigram. The concept behind Unigram is to develop a new model that approximates the LVLMs by samples $\{S^1, \ldots, S^N\}$ and get the LVLM's token probabilities using this model. As N increases, the new model gets closer to LVLMs. Due to time and cost constraints, we just train a simple n-gram model using the samples $\{S^1, \ldots, S^N\}$ as well as the main response R. We then compare the average and maximum of the negative probabilities of the sentence in response R using the following equations:

$$\mathcal{S}^{\text{Avg}}_{\text{n-gram}}(i) = -\frac{1}{J_i} \sum_{j=1}^{J_i} \log \hat{p}_{ij}, \quad (20)$$

$$\mathcal{S}^{\text{Max}}_{\text{n-gram}}(i) = \max_j (-\log \hat{p}_{ij}), \quad (21)$$

where $\hat{p}_{ij}$ is the probability of a token at position j of a sentence i.

Natural Language Inference (NLI) determines whether a hypothesis follows a premise, classified into either entailment/neutral/contradiction. In this work, we use DeBERTa-v3-large (He et al., 2023) fine-tuned to MNLI as the NLI model. The input for NLI classifiers is typically the premise concatenated to the hypothesis, which for NLI is the sampled passage S^n concatenated to the sentence to be assessed r_i in the response R. Only the logits associated with the 'entailment' and 'contradiction' classes are considered,

$$P(contradict \mid r_i, S^n) = \frac{\exp\left(z_e^{i,n}\right)}{\exp\left(z_e^{i,n}\right) + \exp\left(z_c^{i,n}\right)},$$

where $z_e^{i,n} = z_e(r_i, S^n)$ and $z_c^{i,n} = z_c(r_i, S^n)$ are the logits of the 'entailment' and 'contradiction' classes. NLI score for sentence r_i on samples $\{S^1, \ldots, S^N\}$ is then defined as,

$$\mathcal{S}_{\text{NLI}}(i) = \frac{1}{N} \sum_{n=1}^{N} P(contradict \mid r_i, S^n). \quad (22)$$