

MULTIHALUDET: Multilingual Hallucination Detection via LLM Hidden State Probing

Riasad Alvi¹, Nurul Labib Sayeedi¹, Md. Faiyaz Abdullah Sayeedi^{1,2}

¹United International University, ²BRAC University

ralvi212069@bscse.uui.ac.bd, nsayeedi2410045@bsds.uui.ac.bd, msayeedi212049@bscse.uui.ac.bd

<https://github.com/alvi-uui/MultiHaluDet>

Abstract

Hallucinations in Large Language Models (LLMs) represent a critical barrier to their reliable deployment, a vulnerability heavily exacerbated in non-English and resource-constrained contexts. Existing detection approaches that rely on output confidence heuristics or single-layer internal representations frequently fail to capture deep, complex factual inconsistencies across diverse languages. To address this, we introduce MULTIHALUDET, a novel four-stage framework that detects multilingual hallucinations by probing the full hidden state trajectories of frozen LLMs without requiring language-specific fine-tuning. Our method extracts sequential features across multiple layers and processes them via a hybrid architecture using multi-scale attention and self-attention pooling. By generating out-of-fold embeddings that feed into a learned ensemble meta-learner, MULTIHALUDET captures both fine-grained and coarse-grained patterns of factual inconsistency. Extensive experiments demonstrate that our framework achieves state-of-the-art detection performance, reaching up to 98.55% AU-ROC on the English HaluEval and TriviaQA benchmarks using Mistral-7B and LLaMA2-7B architectures. Crucially, we rigorously evaluate our framework’s cross-lingual generalization across high (French), medium (Bangla), and low-resource (Amharic) languages. MULTIHALUDET demonstrates exceptional representational robustness, consistently outperforming baselines and successfully transferring hallucination detection capabilities across typologically diverse linguistic tiers.

1 Introduction

Hallucinations in LLMs have emerged as a critical barrier to their reliable deployment across high-stakes domains (Farquhar et al., 2024; Mishra et al., 2024; Varshney et al., 2023). These generations pose significant risks (Farquhar et al., 2024), motivating diverse detection techniques. Evidence-

based methods (Chern et al., 2023; Zhang et al., 2024) retrieve external information to verify factual consistency, but are computationally intensive. Evidence-free methods leverage inherent model characteristics: logit-based and consistency-based methods estimate uncertainty and output stability, while classification-based methods (Orgad et al., 2024; Binkowski et al., 2025) probe internal hidden states without external retrieval.

While internal state probing shows promise (Orgad et al., 2024; Binkowski et al., 2025), current approaches remain inadequate, particularly across diverse languages. Recent work reveals truthfulness information is concentrated in specific tokens (Orgad et al., 2024); however, single-position methods struggle when non-factual tokens are distributed across sequences. Furthermore, attention-based approaches (Chuang et al., 2024; Binkowski et al., 2025) using simple ratios achieve limited discrimination. Probabilistic frameworks (Hou et al., 2025) require complex reasoning pipelines, while active validation (Varshney et al., 2023) and tool-augmented systems (Chern et al., 2023; Zhang et al., 2024) introduce substantial latency. As noted in Farquhar et al. (2024), hallucinations manifest as semantic-level confabulations rather than token-level uncertainty. Existing methods largely fail to address these confabulations across non-English and resource-constrained languages.

Transformer architectures process information through deep layers, embedding language-agnostic hallucination signals within hidden state trajectories. Deep sequence modeling with multi-scale feature aggregation offers a promising solution. We introduce MULTIHALUDET, a supervised framework leveraging multi-scale attention and transformer encoders to model hidden state dynamics across the full depth of frozen LLMs. Unlike methods focused on individual tokens or static layers, our approach aggregates information across multiple scales through self-attention pooling. We evaluate

across multiple architectures and language tiers, demonstrating strong cross-lingual performance.

Our contributions are summarized as follows:

- We introduce MULTIHALUDET, a four-stage framework comprising dynamic feature extraction, multi-scale attention encoding, out-of-fold deep feature generation, and learned ensemble meta-learning, which jointly capture fine-grained and coarse-grained patterns from hidden-state trajectories.
- We evaluate our framework on the HaluEval and TriviaQA benchmarks using Llama-2-7B and Mistral-7B-Instruct, demonstrating substantial performance gains and representational robustness over existing baselines.
- We comprehensively evaluate the cross-lingual generalization of our framework by translating standard benchmarks into French (high-resource), Bangla (medium-resource), and Amharic (low-resource). Our results demonstrate that internal state probing can maintain strong detection signals across typologically diverse linguistic tiers under controlled translation-based evaluation, without language-specific fine-tuning.

2 Related Work

Hallucination detection approaches fall into three categories: (i) evidence-based methods verifying outputs against external knowledge, (ii) self-detection methods leveraging internal model states, and (iii) consistency-based approaches assessing output stability.

Evidence-based methods retrieve external information for verification. Kale and Alfeo (2025) converts LLM responses into knowledge graphs for atomic fact verification. Vangala et al. (2025) introduced multi-source retrieval with contradiction graph analysis. Zhang and Wang (2025) employed HHEM for lightweight consistency assessment. While effective, these methods depend on retrieval quality, making them vulnerable to gaps and latency.

Self-detection methods probe LLM internal representations without external resources. Su et al. (2024) introduced MIND, training classifiers on auto-generated pseudo-labels. Liang and Wang (2025) employed Bayesian optimization to identify optimal layer insertion points. Zhang et al. (2025) proposed MHAD, selecting specific neurons via

linear probing. Kossen et al. (2024) approximated semantic uncertainty from hidden states. Chen et al. (2024) proposed INSIDE, measuring consistency through eigenvalues of response covariances. Kim et al. (2025) analyzed layer-wise usable information across transformer depths. Quevedo et al. (2024) demonstrated strong detection using only four token probability features. Consistency-based methods assess output stability. Yang et al. (2025a) proposed MetaQA, leveraging metamorphic relations for semantic consistency verification.

Despite these advances, challenges remain. Evidence-based methods suffer from retrieval dependency. Self-detection methods relying on single-position representations struggle when non-factual tokens appear at sequence beginnings. Consistency-based methods incur costs from multiple generation passes. Simhi et al. (2025) demonstrated that LLMs can hallucinate with high certainty despite possessing correct knowledge. Yang et al. (2025b) revealed the need to distinguish between intelligent and defective hallucinations.

Our work aligns with internal state probing methods but addresses their key limitations. While Liang and Wang (2025) and Zhang et al. (2025) focus on specific layers or tokens, we employ dynamic layer sampling and multi-scale attention to aggregate information across the full depth trajectory. Unlike Kossen et al. (2024) and Chen et al. (2024), which probe specific positions, our architecture processes sequential layer features with transformer encoders and self-attention pooling, enabling adaptive depth selection. Our out-of-fold stacking with ensemble meta-learning provides more robust generalization than single-classifier approaches (Su et al., 2024; Liang and Wang, 2025).

3 Methodology

We present MULTIHALUDET (Figure 1), a four-stage framework for hallucination detection.

3.1 LLM-Based Feature Extraction

3.1.1 Prompt Construction and Forward Pass

Let $\mathcal{D} = \{(q_i, a_i, y_i)\}_{i=1}^N$ denote a dataset of question-answer pairs, where q_i is a natural-language question, a_i is a candidate answer, and $y_i \in \{0, 1\}$ is a binary label indicating whether a_i is a hallucination ($y_i = 1$) or a faithful response ($y_i = 0$). Each sample (q_i, a_i) is formatted as a structured natural-language prompt and tokenized with a fixed maximum sequence length. The

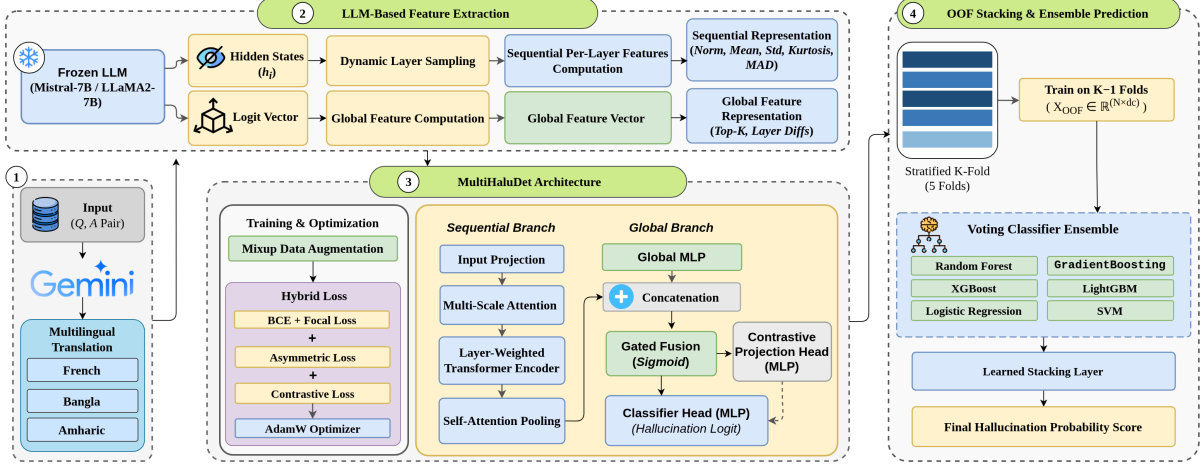


Figure 1: Overview of the four-stage MULTIHaluDET framework for multilingual hallucination detection.

prompt is passed through a frozen, quantized LLM in a single forward pass, yielding a sequence of hidden state tensors $\{\mathbf{H}^{(l)}\}_{l=0}^L$, where L is the number of transformer layers and $\mathbf{H}^{(l)} \in \mathbb{R}^{T \times d}$ collects the d -dimensional representations of T tokens at layer l . The next-token logit vector $\mathbf{z} \in \mathbb{R}^V$ at the final position is retained for global statistics. No gradient is computed during feature extraction; the LLM parameters remain fully frozen throughout all experiments.

3.1.2 Dynamic Layer Sampling

To ensure architectural compatibility across LLMs of varying depth, we introduce a *dynamic layer sampling* strategy that maps any model’s L transformer layers to a fixed target count K , producing a uniform sequential representation regardless of model size. Three regimes are handled without model-specific configuration:

- **Exact match** ($L = K$): identity mapping $\mathcal{I}_K = \{1, \dots, K\}$.
- **Shallow model** ($L < K$): all L layers are used and the deepest layer is repeated to pad the sequence to K .
- **Deep model** ($L > K$): K indices are selected by uniform interpolation over the full depth:

$$\mathcal{I}_K = \left\{ \text{clip} \left(\left\lfloor 1 + (L-1) \cdot \frac{k}{K-1} \right\rfloor, 1, L \right) \right\}_{k=0}^{K-1} \quad (1)$$

where $\lfloor \cdot \rfloor$ denotes rounding to the nearest integer and $\text{clip}(\cdot, 1, L)$ clamps indices to valid range. This formulation is differentiable with respect to K and requires no architecture-specific configuration, making it directly applicable to any transformer-

based LLM.

3.1.3 Sequential Per-Layer Features

For each sampled layer index $l \in \mathcal{I}_K$, we extract a compact descriptor from two complementary views of the hidden state tensor: the last-token representation $\mathbf{h}^{(l)} = \mathbf{H}_{T,:}^{(l)}$, which reflects the model’s final contextual state, and the sequence mean $\bar{\mathbf{h}}^{(l)} = \frac{1}{T} \sum_{t=1}^T \mathbf{H}_{t,:}^{(l)}$, which captures the aggregate token-level context. From these, we compute a per-layer descriptor $\mathbf{s}^{(l)} \in \mathbb{R}^{d_s}$ comprising distributional statistics: the ℓ_2 norm, mean, standard deviation, extremal values, activation sparsity, near-zero mass, and the kurtosis and median absolute deviation (MAD) of the hidden state, capturing the peakedness and robust dispersion of the activation distribution:

$$\kappa(\mathbf{h}) = \frac{1}{d} \sum_{j=1}^d \left(\frac{h_j - \bar{h}}{\hat{\sigma}} \right)^4, \quad (2)$$

$$\text{MAD}(\mathbf{h}) = \text{median}_j(|h_j - \text{median}(\mathbf{h})|)$$

where \bar{h} and $\hat{\sigma}$ denote the empirical mean and standard deviation of the hidden-state vector, and $\text{median}(\mathbf{h})$ is its coordinate-wise median. Collecting descriptors across all sampled layers yields the sequential representation $\mathbf{S} \in \mathbb{R}^{K \times d_s}$, which encodes how the LLM’s internal dynamics evolve as a function of depth.

3.1.4 Anchor-Based Depth Probing

To enable consistent cross-architecture comparisons at semantically meaningful network depths, we define four *anchor layers* corresponding to pro-

portional depth fractions $\{\alpha_j\}_{j=1}^4$. For each anchor, the closest available sampled layer is identified:

$$l_j^* = \arg \min_{l \in \mathcal{I}_K} |l - \lfloor \alpha_j \cdot L \rfloor| \quad (3)$$

and the corresponding descriptor $\mathbf{s}^{(l_j^*)}$ is retained as an anchor feature. This construction ensures that features at early, middle, and late network stages are explicitly represented in the global feature vector regardless of the total model depth, providing a principled substitute for hardcoded layer indices.

3.1.5 Global Feature Vector

A global feature vector $\mathbf{g} \in \mathbb{R}^{d_g}$ aggregates information across the full forward pass. It comprises: (i) the top- k next-token probabilities and their pairwise differences, capturing the sharpness of the model’s next-token distribution; (ii) the entropy, standard deviation, and maximum of the logit vector \mathbf{z} ; (iii) first- and second-order statistics of the layer-wise ℓ_2 -norm trajectory $\{r_l\}_{l \in \mathcal{I}_K}$, $r_l = \|\mathbf{h}^{(l)}\|_2$, capturing norm growth and volatility across depth; (iv) the anchor descriptors from Eq. 3; and (v) cross-feature interaction terms between logit statistics and norm dynamics, which encode the coupling between the model’s output confidence and its internal representational geometry. Together, \mathbf{S} and \mathbf{g} form the complete feature representation passed to classification stage.

3.2 MultiHaluDet Architecture

3.2.1 Projection and Multi-Scale Attention

The sequential input \mathbf{S} is first projected into a uniform hidden space of dimension H via a linear layer followed by layer normalization and a GELU activation. A *multi-scale attention* module then processes the projected sequence at multiple temporal resolutions simultaneously. For each scale factor $c \in \mathcal{C}$, the sequence is locally average-pooled to compress K positions into $\lceil K/c \rceil$ positions, passed through a scale-specific linear projection, and up-sampled back to length K via nearest-neighbor interpolation. The contributions of each scale are combined through a learned position-wise gate:

$$\tilde{\mathbf{h}}_t = \sum_{c \in \mathcal{C}} w_{t,c} \mathbf{P}_c(\text{Upsample}(\text{Pool}_c(\mathbf{H}_{\text{proj}})))_t \quad (4)$$

where $w_{t,c} = \text{softmax}(\mathbf{W}_g \mathbf{h}_t)_c$ is the position-wise scale gate computed from the original projected hidden state, and \mathbf{P}_c is the per-scale linear

projection. The output is combined with the original projection via a residual connection and layer normalization, preserving fine-grained layer information while enriching it with multi-scale context.

3.2.2 Layer-Weighted Transformer Encoder

Because different LLM layers carry information of varying discriminative value for hallucination detection, we modulate the fused sequence by a learnable, softmax-normalized importance vector $\boldsymbol{\lambda} \in \mathbb{R}^K$ before encoding:

$$\begin{aligned} \hat{\mathbf{h}}_k &= \bar{\lambda}_k \cdot \tilde{\mathbf{h}}_k + \mathbf{p}_k, \quad k = 1, \dots, K, \\ \bar{\boldsymbol{\lambda}} &= \text{softmax}(\boldsymbol{\lambda}) \end{aligned} \quad (5)$$

where \mathbf{p}_k is a learned positional embedding that encodes the relative depth of each sampled layer within the representational sequence. The modulated sequence is encoded by a stack of Pre-LN Transformer encoder layers with multi-head self-attention and GELU feed-forward sublayers, enabling the model to attend across LLM depths and capture long-range inter-layer dependencies in the hidden state trajectory.

3.2.3 Self-Attention Pooling

Rather than discarding positional information through mean pooling or relying on a fixed summary token, we aggregate the transformer output via a learned attention pooling mechanism. A two-layer MLP with a tanh nonlinearity assigns a scalar relevance score to each position, and the final sequential representation is their weighted sum:

$$\mathbf{u} = \sum_{k=1}^K \alpha_k \mathbf{e}_k, \quad \alpha_k = \frac{\exp(a(\mathbf{e}_k))}{\sum_{j=1}^K \exp(a(\mathbf{e}_j))} \quad (6)$$

where \mathbf{e}_k is the encoder output at position k and $a : \mathbb{R}^H \rightarrow \mathbb{R}$ is the learned scoring MLP. This allows the model to focus on the LLM layers most informative for the current input, providing a form of input-adaptive depth selection.

3.2.4 Global Branch and Gated Fusion

The global feature vector \mathbf{g} is processed independently through a two-layer MLP with layer normalization and dropout, yielding a compact global representation $\mathbf{v} \in \mathbb{R}^{H/2}$. The sequential representation $\mathbf{u} \in \mathbb{R}^H$ and global representation are concatenated to form the joint embedding $\mathbf{c} = [\mathbf{u}; \mathbf{v}] \in \mathbb{R}^{3H/2}$. A sigmoid gate then

performs element-wise re-weighting to suppress uninformative dimensions:

$$\tilde{\mathbf{c}} = \mathbf{c} \odot \sigma(\mathbf{W}_{\text{gate}} \mathbf{c} + \mathbf{b}_{\text{gate}}) \quad (7)$$

The gated representation $\tilde{\mathbf{c}}$ is passed through a three-layer MLP classifier to produce the final hallucination logit. A separate two-layer projection head maps $\tilde{\mathbf{c}}$ onto a unit hypersphere for the contrastive objective described in Section 4.

3.3 Out-of-Fold Stacking

To obtain unbiased deep representations for the meta-learner without data leakage, we employ K -fold out-of-fold (OOF) stacking over the training set. For fold k , a fresh MULTIHALUDET instance is trained on the remaining $K - 1$ folds and used to extract the gated fusion representation $\tilde{\mathbf{c}}$ for the held-out fold. After all K folds, the OOF features form a complete training matrix:

$$\mathbf{X}^{\text{OOF}} \in \mathbb{R}^{N_{\text{train}} \times d_c}, \quad d_c = \dim(\tilde{\mathbf{c}}) \quad (8)$$

with every training sample represented exactly once without ever being used in its own fold’s training. Test-set representations are obtained by averaging the gated fusion embeddings across all K fold models in feature space:

$$\tilde{\mathbf{c}}^{\text{test}} = \frac{1}{K} \sum_{k=1}^K \tilde{\mathbf{c}}^{(k)\text{test}} \quad (9)$$

This averaging operates at the representation level rather than the probability level, providing implicit test-time ensembling and producing a more stable input to the meta-learner. The OOF and test features are standardized before being passed to the ensemble stage.

3.4 Ensemble Meta-Learner

The OOF deep features serve as input to a diverse ensemble of classifiers comprising gradient-boosted trees, random forests, a kernel support vector machine, a multi-layer perceptron, and a logistic regression baseline. Model diversity across both inductive biases (linear, kernel, tree, neural) and hyperparameter profiles reduces variance and improves robustness to the distributional properties of the deep features. Rather than fixed heuristic weights, the ensemble prediction is formed by stacking the base classifier probability estimates through a logistic meta-regressor trained

on held-out out-of-fold predictions. Denoting the log-odds transform of each base probability by $\ell_m = \log(\hat{p}_m / (1 - \hat{p}_m))$, the final ensemble probability is:

$$\hat{p}_{\text{ens}} = \sigma \left(\beta_0 + \sum_{m=1}^M \beta_m \ell_m \right) \quad (10)$$

where $\sigma(\cdot)$ is the sigmoid function and the coefficients $\{\beta_m\}_{m=0}^M$ are learned by an ℓ_2 -regularized logistic regression meta-learner trained on a secondary hold-out split of the out-of-fold probability outputs. This formulation adapts the ensemble composition to the empirical discriminative strength of each base learner while producing inherently calibrated probability estimates. All component models are configured with balanced class weights to maintain sensitivity under any residual distributional skew. The classification threshold is set by Youden’s J statistic $\tau^* = \arg \max_{\tau} [\text{TPR}(\tau) - \text{FPR}(\tau)]$, which maximizes the simultaneous improvement in sensitivity and specificity.

3.5 Multilingual Adaptation Strategy

To evaluate the multilingual generalization of our hidden state probing framework across varying resource tiers, we translated the source English datasets into French (high-resource), Bangla (medium-resource), and Amharic (low-resource) utilizing the Gemini 2.5 Flash model. Formally, let the original source dataset be defined as $\mathcal{D}_{\text{en}} = \{(q_i, a_i, y_i)\}_{i=1}^N$. We introduce a translation mapping \mathcal{T}_{ℓ} for the set of target languages $\mathcal{L} = \{\text{fr}, \text{bn}, \text{am}\}$. For every query-response pair, this operation strictly preserves the original binary hallucination label, yielding a new language-specific instance defined as: $(q_i^{(\ell)}, a_i^{(\ell)}, y_i^{(\ell)}) = (\mathcal{T}_{\ell}(q_i), \mathcal{T}_{\ell}(a_i), y_i)$. The multilingual evaluation space is then constructed as the union of the original English data and all language-specific subsets.

To guarantee the semantic integrity and context preservation of the translated data, we conducted human evaluation for quality assurance. We sampled 100 instances per language from each dataset, resulting in 300 evaluated samples per dataset and a total of 600 evaluated samples across both datasets. As three of the authors are native Bangla speakers, the Bangla subset was evaluated directly. For French and Amharic, we employed a back-translation methodology, translating the target samples back into English, to verify semantic equivalence against the original source. This man-

ual evaluation confirmed an initial translation accuracy of 96%. The remaining 4% of instances, which exhibited minor semantic drift or structural errors, were explicitly polished and regenerated using Gemini 2.5 Flash to ensure strict alignment with the source truth conditions.

By passing these high-fidelity translated inputs through the frozen LLM, we generate the sequential representation S and global feature vector g natively within each target language’s token space. This formulation ensures that the subsequent deep architecture evaluates intrinsic, language-agnostic hallucination signals embedded within the model’s internal dynamics, allowing us to probe representation alignment without requiring language-specific fine-tuning.

4 Experimental Setup

4.1 Dataset

We evaluate on two question-answering benchmarks. Let \mathcal{Q} denote the set of questions and \mathcal{A} the set of candidate answers. Each dataset $\mathcal{D} = \{(q_i, a_i, y_i)\}_{i=1}^N$ consists of N samples where $q_i \in \mathcal{Q}$ is a question, $a_i \in \mathcal{A}$ is an answer, and $y_i \in \{0, 1\}$ is a binary label indicating non-hallucination ($y_i = 0$) or hallucination ($y_i = 1$). HaluEval (Li et al., 2023) provides $N = 10,000$ human-annotated QA pairs with native hallucination labels derived from real LLM outputs, where each (q_i, a_i, y_i) is explicitly labeled based on factual consistency with verified knowledge sources. TriviaQA (Joshi et al., 2017) is adapted for our evaluation by collecting realistic, model-generated hallucinations. Let $\mathcal{D}_{\text{Trivia}} = \{(q_i, a_i^*)\}_{i=1}^M$ denote the original dataset, where q_i is a question and a_i^* is the ground-truth correct answer. To generate plausible hard negatives, we prompt an early-generation language model known for its propensity to hallucinate, Gemma-2-2B ($\sim 32\%$ - 38% hallucination rate), to answer each q_i . Responses that are plausible but factually incorrect are extracted and denoted as a_i^- . We then construct our final evaluation dataset by pairing each question with its true correct answer and its corresponding model-generated hallucination. To facilitate our multilingual evaluation, we also expanded these source English datasets into three different target languages.

4.2 Baselines

We compare against four types of hallucination detection methods. **Prompt-based methods** include

P(True) (Kadavath et al., 2022), which utilizes a simple prompt template to enable the model to assess the correctness of its own response. **Logit-based methods** use the uncertainty of LLM outputs to detect hallucination; we adopt AvgProb and AvgEnt from Huang et al. (2023) to aggregate logit-based uncertainty across all tokens, and also compare with EUBHD (Su et al., 2024), which focuses on key tokens rather than considering all tokens. **Consistency-based methods** are motivated by the idea that consistent responses indicate factual knowledge; we apply Unigram and NLI variants, as well as INSIDE (Chen et al., 2024), which leverages eigenvalues of the covariance matrix of responses. **Classification-based methods** train a classifier on labeled statements; we compare with SAPLMA (Azaria and Mitchell, 2023), which trains a classifier on the last token’s last-layer hidden state; MIND (Su et al., 2024), which uses unsupervised training on auto-generated pseudo-labels; and Probe@Exact, which relies on information from potentially correct tokens. We also include HD-NDEs (Li et al., 2025) (Neural ODEs, CDEs, and SDEs), which model hidden state trajectories using neural differential equations.

4.3 Hyperparameters

We extract hidden states from frozen LLMs in half-precision without gradient computation. For all LLM forward passes and synthetic generation tasks, the temperature is set to 0.0 to ensure deterministic outputs. Dynamic layer sampling maps variable-depth models to $K = 32$ uniform indices. The deep architecture uses hidden dimension $d = 384$, 8 attention heads, and 6 transformer encoder layers. Training employs AdamW optimizer with learning rate 2×10^{-4} , weight decay 6×10^{-5} , and ReduceLRonPlateau scheduling over 45 epochs with early stopping patience of 15 epochs. We apply a composite loss combining BCE, focal, asymmetric, and contrastive objectives with label smoothing, along with data augmentation via Mixup and CutMix. We employ 5-fold stratified cross-validation with out-of-fold feature extraction to prevent leakage. The stacked ensemble combines 6 classifiers (RandomForest, XGBoost, GradientBoosting, LightGBM, LogisticRegression, Support Vector Machine). Their probability outputs are fused by a logistic meta-regressor trained on held-out out-of-fold predictions, producing inherently calibrated ensemble probabilities. All experiments are conducted on AMD Ryzen 5 8500G CPU with 32GB

Method	HaluEval		TriviaQA	
	Mistral-7B	LLaMA2-7B	Mistral-7B	LLaMA2-7B
P(True)	49.7	46.7	48.3	42.3
AvgProb	43.6	42.1	48.5	44.1
AvgEnt	49.7	47.3	47.6	41.1
EUBHD	70.5	71.9	80.6	80.5
Unigram	62.3	58.2	59.5	56.8
NLI	63.1	61.3	61.4	59.4
INSIDE	76.0	74.5	81.3	81.7
SAPLMA	89.4	87.0	84.1	80.0
MIND	94.5	86.1	84.5	79.4
Probe@Exact	93.4	88.3	84.1	81.3
Neural ODEs	91.2	89.5	83.7	81.7
Neural CDEs	95.4	91.4	84.1	83.7
Neural SDEs	93.7	92.8	85.1	81.0
MultiHaluDet	98.43	98.55	98.30	98.26

Table 1: Hallucination detection performance (AUROC %) on HaluEval and TriviaQA using Mistral-7B-Instruct and LLaMA2-7B. Best results in bold.

RAM and NVIDIA GeForce RTX 5060 Ti GPU with 16GB VRAM.

4.4 Evaluation Metric

We utilize AUROC (%), which stands for the area under the ROC curve, to objectively evaluate the effectiveness of models. The higher the value of AUROC, the stronger the ability of this method for hallucination detection. All experiments employ 5-fold cross-validation with stratified sampling to ensure reliable estimates.

5 Results and Analysis

5.1 RQ1: To what extent does MultiHaluDet improve hallucination detection compared to existing internal-state and confidence based baselines?

Table 1 presents the detection performance (AUROC %) of MULTIHALUDET compared against thirteen baseline methods, evaluated across both the HaluEval and TriviaQA datasets using Mistral-7B and LLaMA2-7B.

Failure of Surface-Level Heuristics. The results reveal a clear hierarchy in the efficacy of different detection paradigms. Token-level probability and entropy metrics (P(True), AvgProb, AvgEnt) systematically fail to provide meaningful detection signals, hovering around or below random chance (41.1% – 49.7% AUROC). This confirms that raw output confidence is severely miscalibrated and insufficient for identifying deep factual inconsistencies. Intermediate methods incorporating structural linguistic features (Unigram, NLI) offer only marginal improvements, generally plateauing in the low 60% range.

Representational Baselines. Approaches that

leverage deeper internal representations show significantly more promise. Methods like SAPLMA, MIND, and traditional probing (Probe@Exact) push performance into the 80%–94% range. The strongest baselines are those modeling the continuous dynamics of hidden states (Neural ODEs, CDEs, SDEs). Neural CDEs, in particular, achieve a competitive 95.4% on Mistral-7B for HaluEval. However, these methods exhibit high variance between model architectures; for instance, MIND’s performance drops sharply from 94.5% on Mistral-7B to 86.1% on LLaMA2-7B.

Ours. MULTIHALUDET achieves state-of-the-art performance across all experimental conditions, substantially outperforming the strongest continuous-time baselines. On HaluEval, we achieve 98.43% AUROC with Mistral-7B and 98.55% with LLaMA2-7B. Crucially, MULTIHALUDET demonstrates exceptional cross-architecture robustness. While nearly all baselines suffer noticeable performance degradation when transitioning from Mistral to LLaMA2, our framework maintains near-identical, near-perfect efficacy (actually scoring slightly higher on LLaMA2 in HaluEval). Similarly, on the TriviaQA dataset, which features plausible but factually incorrect synthetic hard negatives, we achieve 98.30% and 98.26% respectively. These consistent gains across diverse datasets and architectures demonstrate that our multi-scale attention mechanism and out-of-fold stacking approach successfully aggregate robust, architecture-agnostic signals from LLM hidden state trajectories.

5.2 RQ2: How robust is MultiHaluDet across typologically diverse languages with varying resource availability?

To demonstrate the cross-lingual generalization of MULTIHALUDET, we extend our evaluation beyond English to include three typologically diverse languages: French, Bangla, and Amharic.

Table 2 shows that hallucination detection performance correlates with the representational frequency of these languages in the base models. MULTIHALUDET maintains exceptionally strong performance on French, achieving 96.2% and 95.8% on HaluEval for Mistral-7B and LLaMA2-7B respectively, trailing the English baselines (98.4% and 98.5%) by only a marginal fraction. Similar high retention is observed on TriviaQA, where French scores reach 95.5% and 94.9%. For Bangla, we observe a more noticeable but con-

Language	Method	HaluEval (M) / (L)	TriviaQA (M) / (L)
English (Base)	Best Baseline	95.4 / 92.8	85.1 / 83.7
	MultiHaluDet	98.4 / 98.5	98.3 / 98.2
French (High)	Best Baseline	92.1 / 89.4	81.3 / 80.5
	MultiHaluDet	96.2 / 95.8	95.5 / 94.9
Bangla (Medium)	Best Baseline	78.4 / 75.1	69.2 / 67.8
	MultiHaluDet	89.1 / 88.4	87.6 / 86.3
Amharic (Low)	Best Baseline	62.3 / 59.8	54.1 / 52.6
	MultiHaluDet	78.5 / 76.2	75.8 / 73.4

(M): Mistral-7B (L): LLaMA2-7B

Table 2: Cross-lingual hallucination detection performance. MultiHaluDet consistently outperforms the best baseline across all resource settings.

trolled degradation, yielding HaluEval scores of 89.1% (Mistral-7B) and 88.4% (LLaMA2-7B), alongside 87.6% and 86.3% on TriviaQA. While the framework successfully adapts to the language, accurately detecting factual inconsistencies requires the model to effectively navigate Bangla’s rich morphology and the complex linguistic adaptations necessary for regional dialect variations.

Amharic, a severely low-resource language with limited representation in both Mistral and LLaMA2, presents the greatest challenge. The base models inherently struggle with factual recall in Amharic, which limits the quality of the internal representations our detector relies on. As a result, absolute performance understandably drops to 78.5% and 76.2% on HaluEval, and 75.8% and 73.4% on TriviaQA for Mistral-7B and LLaMA2-7B, respectively. However, MULTIHALUDET still manages to extract meaningful detection signals, maintaining a robust lead over random chance and standard confidence baselines.

5.3 Ablation Study

To understand the necessity of our proposed architectural choices, we conduct an ablation study by systematically removing key components of the MULTIHALUDET framework. We evaluate the degraded models on both HaluEval and TriviaQA using Mistral-7B and LLaMA2-7B. The results are summarized in Table 3.

The removal of the Out-of-Fold (OOF) Stacking mechanism causes the most severe degradation across all configurations, resulting in a precipitous drop of nearly 10 percentage points (falling to 88.67% on Mistral-7B HaluEval and 87.41% on TriviaQA). Without OOF stacking, the meta-classifier heavily overfits to the localized noise of early hidden layers, completely failing to generalize when faced with the plausible hard negatives in

Method	Mistral-7B		LLaMA2-7B	
	HaluEval	TriviaQA	HaluEval	TriviaQA
Full	98.43	98.30	98.55	98.26
w/o MSA	91.45 (↓6.98)	90.82 (↓7.48)	92.14 (↓6.41)	91.33 (↓6.93)
w/o OOF	88.67 (↓9.76)	87.41 (↓10.89)	89.25 (↓9.30)	88.19 (↓10.07)
w/o TP	93.28 (↓5.15)	92.56 (↓5.74)	93.71 (↓4.84)	93.04 (↓5.22)

Table 3: Ablation study of MULTIHALUDET. Performance is measured in AUROC (%). The red text indicates the absolute performance drop when a specific core component is removed, confirming the structural necessity of the complete framework. Abbreviations – w/o: without, MSA: Multi-Scale Attention, OOF: Out-of-Fold Stacking, TP: Trajectory Probing.

TriviaQA. This confirms that our stacking approach is non-negotiable for robust feature aggregation.

Bypassing the Multi-Scale Attention module and relying on standard global average pooling results in a substantial 6–8% drop in AUROC. Hallucination signals are not uniformly distributed across the generation trajectory; they often manifest as sudden, localized semantic shifts in the middle layers. The sharp decline in performance without this module proves that capturing these local trajectory shifts is strictly necessary, and standard pooling mechanisms are too coarse to detect subtle factual deviations.

Finally, replacing our continuous trajectory probing with a static, final-layer representation (w/o Trajectory Probing) degrades performance by approximately 5 percentage points. While the final layer contains significant semantic information, this result empirically proves our core hypothesis: the *process* of how a model arrives at an answer (the hidden state evolution) contains critical truthfulness signals that are permanently lost if one only analyzes the final output state.

6 Conclusion

We presented MULTIHALUDET, a four-stage framework that detects hallucinations by probing the hidden state trajectories of frozen LLMs without fine-tuning. By integrating multi-scale attention with out-of-fold deep feature generation and learned ensemble meta-learning, our approach effectively captures complex semantic shifts indicating factual inconsistencies. Extensive evaluations show that MULTIHALUDET achieves state-of-the-art performance on standard benchmarks, while also demonstrating robust cross-lingual generalization across high, medium, and low-resource languages.

Limitations

While MULTIHALUDET demonstrates strong performance in detecting multilingual hallucinations, several limitations must be acknowledged.

First, our framework inherently relies on white-box access to the internal hidden states and logits of the target LLMs. Consequently, it cannot be directly applied to proprietary, black-box models (e.g., GPT-4 or Claude) where access to internal representational trajectories is restricted.

Second, although our approach successfully bypasses the computational burden of language-specific fine-tuning, extracting and processing full-depth hidden states across multiple layers still incurs non-trivial memory overhead and requires a complete forward pass. This makes it more computationally demanding than simple surface-level heuristics.

Finally, our multilingual evaluation leverages datasets translated from English using Gemini 2.5 Flash. Despite implementing a rigorous human quality assurance and back-translation pipeline to ensure semantic integrity, evaluating native, naturally occurring prompts in medium and low-resource languages (like Bangla and Amharic) might reveal cultural and linguistic nuances that are not fully captured by translated benchmarks.

References

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.

Jakub Binkowski, Denis Janiak, Albert Sawczyn, Bogdan Gabrys, and Tomasz Jan Kajdanowicz. 2025. Hallucination detection in llms using spectral features of attention maps. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 24365–24396.

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. *arXiv preprint arXiv:2402.03744*.

I Chern, Steffi Chern, Shiqi Chen, Weizhe Yuan, Kehua Feng, Chunting Zhou, Junxian He, Graham Neubig, Pengfei Liu, and 1 others. 2023. Factool: Factuality detection in generative ai—a tool augmented framework for multi-task and multi-domain scenarios. *arXiv preprint arXiv:2307.13528*.

Yung-Sung Chuang, Linlu Qiu, Cheng-Yu Hsieh, Ranjay Krishna, Yoon Kim, and James Glass. 2024. Lookback lens: Detecting and mitigating contextual

hallucinations in large language models using only attention maps. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1419–1436.

- Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. 2024. Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017):625–630.
- Bairu Hou, Yang Zhang, Jacob Andreas, and Shiyu Chang. 2025. A probabilistic framework for llm hallucination detection via belief tree propagation. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3076–3099.
- Yuheng Huang, Jiayang Song, Zhijie Wang, Shengming Zhao, Huaming Chen, Felix Juefei-Xu, and Lei Ma. 2023. Look before you leap: An exploratory study of uncertainty measurement for large language models. *arXiv preprint arXiv:2307.10236*.
- 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.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, and 1 others. 2022. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*.
- Sahil Kale and Antonio Luca Alfeo. 2025. Lie to me: Knowledge graphs for robust hallucination self-detection in llms. *arXiv preprint arXiv:2512.23547*.
- Hazel Kim, Tom A Lamb, Adel Bibi, Philip Torr, and Yarin Gal. 2025. Detecting llm hallucination through layer-wise information deficiency: Analysis of ambiguous prompts and unanswerable questions. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 32298–32310.
- Jannik Kossen, Jiatong Han, Muhammed Razzak, Lisa Schut, Shreshth Malik, and Yarin Gal. 2024. Semantic entropy probes: Robust and cheap hallucination detection in llms. *arXiv preprint arXiv:2406.15927*.
- Junyi Li, Xiaoxue Cheng, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2023. Halueval: A large-scale hallucination evaluation benchmark for large language models (2023). *URL https://arxiv.org/abs/2305.11747*, 2305.
- Qing Li, Jiahui Geng, Zongxiong Chen, Derui Zhu, Yuxia Wang, Congbo Ma, Chenyang Lyu, and Fakhri Karray. 2025. Hd-ndes: Neural differential equations for hallucination detection in llms. In *Proceedings of the 63rd Annual Meeting of the Association for*

- Computational Linguistics (Volume 1: Long Papers)*, pages 6173–6186.
- Shize Liang and Hongzhi Wang. 2025. Neural probe-based hallucination detection for large language models. *arXiv preprint arXiv:2512.20949*.
- Abhika Mishra, Akari Asai, Vidhisha Balachandran, Yizhong Wang, Graham Neubig, Yulia Tsvetkov, and Hannaneh Hajishirzi. 2024. Fine-grained hallucination detection and editing for language models. *arXiv preprint arXiv:2401.06855*.
- Hadas Orgad, Michael Toker, Zorik Gekhman, Roi Reichart, Idan Szpektor, Hadas Kotek, and Yonatan Belinkov. 2024. Lms know more than they show: On the intrinsic representation of llm hallucinations. *arXiv preprint arXiv:2410.02707*.
- Ernesto Quevedo, Jorge Yero Salazar, Rachel Koerner, Pablo Rivas, and Tomas Cerny. 2024. Detecting hallucinations in large language model generation: A token probability approach. In *World Congress in Computer Science, Computer Engineering & Applied Computing*, pages 154–173. Springer.
- Adi Simhi, Itay Itzhak, Fazl Barez, Gabriel Stanovsky, and Yonatan Belinkov. 2025. Trust me, i’m wrong: High-certainty hallucinations in llms. *arXiv e-prints*, pages arXiv–2502.
- Weihang Su, Changyue Wang, Qingyao Ai, Yiran Hu, Zhijing Wu, Yujia Zhou, and Yiqun Liu. 2024. Unsupervised real-time hallucination detection based on the internal states of large language models. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 14379–14391.
- Bhanu Prakash Vangala, Sajid Mahmud, Pawan Neupane, Joel Selvaraj, and Jianlin Cheng. 2025. Halumat: Detecting hallucinations in llm-generated materials science content through multi-stage verification. *arXiv preprint arXiv:2512.22396*.
- Neeraj Varshney, Wenlin Yao, Hongming Zhang, Jian-shu Chen, and Dong Yu. 2023. A stitch in time saves nine: Detecting and mitigating hallucinations of llms by validating low-confidence generation. *arXiv preprint arXiv:2307.03987*.
- Bo Yang, Yinfen Xia, Weisong Sun, and Yang Liu. 2025a. Hallucination detection for llm-based text-to-sql generation via two-stage metamorphic testing. *arXiv preprint arXiv:2512.22250*.
- Chengxu Yang, Jingling Yuan, Siqi Cai, Jiawei Jiang, and Chuang Hu. 2025b. Heaven-sent or hell-bent? benchmarking the intelligence and defectiveness of llm hallucinations. *arXiv preprint arXiv:2512.21635*.
- Chenggong Zhang and Haopeng Wang. 2025. Hallucination detection and evaluation of large language model. *arXiv preprint arXiv:2512.22416*.
- Jiawei Zhang, Chejian Xu, Yu Gai, Freddy Lecue, Dawn Song, and Bo Li. 2024. Knowhalu: Hallucination detection via multi-form knowledge based factual checking. *arXiv preprint arXiv:2404.02935*.
- Luan Zhang, Dandan Song, Zhijing Wu, Yuhang Tian, Changzhi Zhou, Jing Xu, Ziyi Yang, and Shuhao Zhang. 2025. Detecting hallucination in large language models through deep internal representation analysis. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pages 8357–8365.

A Multilingual Data Examples

In this section, we provide representative data points from our evaluation benchmarks across all four languages: English (Base), French (High-Resource), Bangla (Medium-Resource), and Amharic (Low-Resource). Table 2 presents an instance from the HaluEval dataset, illustrating the grounding knowledge, the dialogue history, and both the faithful and hallucinated responses.

Language	Knowledge	Dialogue	Right Response	Right Response Label	Hallucinated	Hallucinated Label
English	Iron Man is starring Robert Downey Jr. Robert Downey Jr. starred in Zodiac (Crime Fiction Film). Zodiac (Crime Fiction Film) is starring Jake Gyllenhaal.	[Human]: Do you like Iron Man [Assistant]: Sure do! Robert Downey Jr. is a favorite. [Human]: Yes i like him too did you know he also was in Zodiac a crime fiction film.	I like crime fiction! Didn't know RDJ was in there. Jake Gyllenhaal starred as well.	0	I'm not a fan of crime movies, but I did know that RDJ starred in Zodiac with Tom Hanks.	1
French	Iron Man met en vedette Robert Downey Jr. Robert Downey Jr. a joué dans Zodiac (film policier). Zodiac (film policier) met en vedette Jake Gyllenhaal.	[Human]: Aimez-vous Iron Man [Assistant]: Bien sûr ! Robert Downey Jr. est un de mes favoris. [Human]: Oui, je l'aime bien aussi. Saviez-vous qu'il a également joué dans Zodiac, un film policier.	J'aime les films policiers ! Je ne savais pas que RDJ y était. Jake Gyllenhaal y a également joué.	0	Je ne suis pas fan de films policiers, mais je savais que RDJ avait joué dans Zodiac avec Tom Hanks.	1
Bangla	আয়রন ম্যান-এ অভিনয় করেছেন রবার্ট ডাউনি জুনিয়র। রবার্ট ডাউনি জুনিয়র জোডিয়াক (ক্রাইম ফিকশন ফিল্ম)-এ অভিনয় করেছেন। জোডিয়াক (ক্রাইম ফিকশন ফিল্ম)-এ অভিনয় করেছেন জেক জিলেনহাল।	[মানুষ]: আপনি কি আয়রন ম্যান পছন্দ করেন [সহকারী]: অবশ্যই! রবার্ট ডাউনি জুনিয়র আমার অন্যতম প্রিয়। [মানুষ]: হ্যাঁ, আমিও তাকে পছন্দ করি। আপনি কি জানেন যে তিনি জোডিয়াক নামে একটি ক্রাইম ফিকশন ফিল্মেও ছিলেন।	আমি ক্রাইম ফিকশন পছন্দ করি! জানতাম না যে আরডিজি সেখানে ছিলেন। জেক জিলেনহালও এতে অভিনয় করেছেন।	0	আমি ক্রাইম মুভির ভক্ত নই, তবে আমি জানতাম যে আরডিজি টম হ্যাঙ্কসের সাথে জোডিয়াক-এ অভিনয় করেছেন।	1
Amharic	አይሮን ማን (Iron Man) ሮብርት ዳውኒ ጆህን ያካትታል። ሮብርት ዳውኒ ጆህን በዚህ ስራው (የወንጀል ልብ ወለድ ፊልም) ላይ ተጫውቷል። ዚህ ስራው (የወንጀል ልብ ወለድ ፊልም) ጆክ ጂሎንሃልን ያካትታል።	[ሰው]: እድረ ግን ትወዳለህ? [ረዳት]: በእርግጥ! ሮብርት ዳውኒ ጆህን ስምዎቼው ነገዱ ነው። [ሰው]: አዎ እኔም እወደዋለሁ። በዚህ ስራው (የወንጀል ልብ ወለድ ፊልም) ላይም እንደነበረ ታውቃለህ?	የወንጀል ልብ ወለድ እወዳለሁ! RDJ እዚያ እንደነበረ አላውቅም ነበር። ጆክ ጂሎንሃልም ተጫውቷል።	0	የወንጀል ፊልም ትወዳለሁም፣ ግን RDJ ከቶም ሃንክስ ጋር በዚህ ስራው ላይ እንደተጫወተ አውቃለሁ።	1

Figure 2: Representative multilingual data points from the HaluEval dataset, demonstrating the semantic preservation of entities and factual inconsistencies across translations.

Table 3 illustrates an example from our modified TriviaQA dataset. TriviaQA (Joshi et al., 2017) is adapted for our evaluation by collecting realistic, model-generated hallucinations. In the original dataset, each entry consists of a question and its ground-truth correct answer. To generate plausible hard negatives, we prompt an early-generation language model known for its propensity to hallucinate, Gemma-2-2B, to answer each question. Responses that are plausible but factually incorrect are extracted to serve as the hallucinated answers. We then construct our final evaluation dataset by pairing each question with its true correct answer and its corresponding model-generated hallucination.

Language	Context (Truncated)	Question	Ground Truth	Right Response Label	Hallucinated	Hallucinated Label
English	[DOC] [TLE] THEME FROM MAHOGANY... The Theme from the movie "Mahogany" also titled "Do You Know Where You're Going To" is a song written by Michael Masser and Gerald Giffin and was sung by Dianna Ross as the theme to the 1975 Paramount film...	Do You Know Where You're Going To? was the theme from which film?	Mahogany	0	Love Story	1
French	[DOC] [TLE] THÈME DE MAHOGANY... Le thème du film « Mahogany », également intitulé « Do You Know Where You're Going To », est une chanson écrite par Michael Masser et Gerald Giffin et a été chantée par Diana Ross comme thème du film de Paramount de 1975...	« Do You Know Where You're Going To ? » était le thème de quel film ?	Mahogany	0	Love Story	1
Bangla	[DOC] [TLE] মেহগনি থিম... "মেহগনি" মুভির থিম, যার শিরোনাম "ডু ইউ নো হোয়্যার ইউ আর গোয়িং টু", এটি মাইকেল ম্যাসার এবং জেরাল্ড গিফিনের লেখা একটি গান এবং এটি ১৯৭৫ সালের প্যারামাউন্ট চলচ্চিত্রের থিম হিসেবে ডায়ানা রস গেয়েছিলেন...	"ডু ইউ নো হোয়্যার ইউ আর গোয়িং টু?" গানটি কোন চলচ্চিত্রের থিম ছিল?	মেহগনি (Mahogany)	0	লাভ স্টোরি (Love Story)	1
Amharic	[DOC] [TLE] ጥያቄገጽ ጭብጥ (THEME FROM MAHOGANY)... ከ "ጥያቄገጽ" ፊልም የተወሰደው እና "Do You Know Where You're Going To" በወቅቱ የሚታወቀው ዘፈን በሚካኤል ማሰር እና ጄራልድ ጊፊን የተጻፈ ሲሆን በ1975 የፓራማውንት ፊልም ጭብጥ ሆኖ በዲያና ሮስ ተዘጋጅቷል::...	"Do You Know Where You're Going To?" የትኛው ፊልም ጭብጥ ነበር?	ጥያቄገጽ (Mahogany)	0	ላቭ ስቶሪ (Love Story)	1

Figure 3: Representative multilingual data points from the modified TriviaQA dataset. The context has been truncated for brevity. The hallucinated response (a_i^-) is generated by prompting Gemma-2-B to produce a plausible but factually incorrect answer.