

# Investigating Hallucination in Conversations for Low Resource Languages

Amit Das<sup>1</sup>, Md. Najib Hasan<sup>2</sup>, Souvika Sarkar<sup>2</sup>, Zheng Zhang<sup>3</sup>, Fatemeh Jamshidi<sup>4</sup>,  
Tathagata Bhattacharya<sup>5</sup>, Nilanjana Raychawdhury<sup>6</sup>, Dongji Feng<sup>7</sup>,  
Vinija Jain<sup>8,9\*</sup>, Aman Chadha<sup>8, 10\*</sup>

<sup>1</sup>University of North Alabama, <sup>2</sup>Wichita State University, <sup>3</sup>Murray State University,

<sup>4</sup>California State Polytechnic University Pomona, <sup>5</sup>Auburn University at Montgomery,

<sup>6</sup>Auburn University, <sup>7</sup>Gustavus Adolphus College, <sup>8</sup>Stanford University, <sup>9</sup>Meta, <sup>10</sup>Amazon GenAI

Corresponding author: Amit Das (adas@una.edu)

## Abstract

Large Language Models (LLMs) have demonstrated remarkable proficiency in generating text that closely resemble human writing. However, they often generate factually incorrect statements, a problem typically referred to as ‘hallucination’. Addressing hallucination is crucial for enhancing the reliability and effectiveness of LLMs. While much research has focused on hallucinations in English, our study extends this investigation to conversational data in three languages: Hindi, Farsi, and Mandarin. We offer a comprehensive analysis of a dataset to examine both factual and linguistic errors in these languages for GPT-3.5, GPT-4o, Llama-3.1, Gemma-2.0, DeepSeek-R1 and Qwen-3. We found that LLMs produce very few hallucinated responses in Mandarin but generate a significantly higher number of hallucinations in Hindi and Farsi.

## 1 Introduction

LLMs have emerged as a groundbreaking advancement in artificial intelligence, revolutionizing the field of natural language processing (NLP). These models show an extraordinary ability to perform a wide range of language-related tasks such as text generation, machine translation, summarization, question answering, etc. Despite their impressive performance, a critical challenge persists: the phenomenon of ‘hallucination,’ where LLMs produce text that is factually incorrect, nonsensical, or misleading (Ji et al., 2023; Li et al., 2023). This issue is not limited to English but also affects low-resource languages.

Hallucination in neural language generation has become a growing concern across vari-

ous areas, including neural machine translation (Raunak et al., 2021; Wang and Sennrich, 2020) and summarization (Durmus et al., 2020; Kang and Hashimoto, 2020). Research on LLMs has highlighted problems like generating inaccurate information, producing incoherent text, and merging unrelated sources (Mckenna et al., 2023). In the context of machine translation, studies have reported hallucination rates exceeding 10% for certain language pairs (Guerreiro et al., 2023). Despite these advancements, the issue of hallucination in conversational systems, especially in low-resource languages, remains largely underexplored.

The complexity and significance of low-resource languages underscore the urgent need for a detailed investigation into hallucinations within these linguistic contexts. Addressing this issue not only improves the reliability and applicability of LLMs but also has implications for a range of applications, including information retrieval, sentiment analysis, and machine translation. This study aims to rigorously evaluate the factual accuracy and reliability of LLMs, specifically GPT-3.5 and GPT-4o, when generating texts in Hindi, Farsi, and Mandarin.

Our research seeks to make the following contributions:

1. We conduct an analysis of the factual accuracy of outputs of LLMs (specifically for GPT-3.5, GPT-4o, Llama-3.1, Gemma-2.0, DeepSeek-R1 and Qwen-3) in Hindi, Farsi, and Mandarin, offering insights into their reliability across these languages.<sup>1</sup>

2. We identify and categorize the types of factual and linguistic errors present in conver-

<sup>\*</sup>Work does not relate to position at Amazon or Meta.

<sup>1</sup>Sample LLM responses can be found here: <https://github.com/AmitDasRup123/LLM-Hallucination-Low-Resource-Languages/>.




Figure 1: Workflow diagram of our work. It shows a sample conversation where LLM (Large Language Models) provides irrelevant response for Hindi, Farsi and Mandarin. The left side are the inputs to the LLM and right side are the irrelevant responses by LLM. We have explored GPT-3.5, GPT-4o, Llama-3.1, Gemma-2.0, DeepSeek-R1 and Qwen-3 in this paper.

sational outputs generated by LLMs for Hindi, Farsi, and Mandarin.

## 2 Related Work

Large Language Models (LLMs) are increasingly scrutinized for both their susceptibility to bias and their tendency to hallucinate, with recent research revealing concerning patterns across diverse NLP tasks. Various forms of bias have been identified—including gender, religion, and political ideology—affecting both model behavior and output generation (Salinas et al., 2023; Jenny et al., 2023; He et al., 2023). These biases are not only pervasive but also task-sensitive, as shown by Zheng et al. (Zheng et al., 2023), who demonstrate that simple perturbations such as altering the order of multiple-choice options can lead to selection bias. Similarly, Gonçalves and Strubell (Gonçalves and Strubell, 2023) offer a comprehensive synthesis of how biases manifest in LLMs, while others have proposed mitigation techniques by retraining on bias-filtered datasets (Dong et al., 2023; Limisiewicz et al., 2023). However, as highlighted in our Section 3, such mitigation techniques may not generalize well, particularly in complex tasks like stance detection, where the bias distribution itself varies significantly.

Hallucination remains a central challenge in LLM deployment (Dale et al., 2022). Initial work in this area focused mainly on grounded tasks such as summarization and dialogue, where hallucinations were detected by aligning the generated factual units with the input of the source using entailment-based or QA-driven techniques (Maynez et al., 2020). As hallucination detection has expanded to open source generation, newer approaches now use reliable references such as Wikipedia and web searches to assess factuality (Min et al., 2023; Chern et al., 2023; Mishra et al., 2024), while others focus on specific domains such as citation verification (Agrawal et al., 2023). Complementing these reference-based methods, internal LLM techniques, such as consistency checks across sampled responses (Manakul et al., 2023) and logit-space analysis (Varshney et al., 2023), aim to identify contradictions without external validation.

Together, these lines of work underscore the intertwined nature of bias and hallucination in LLMs. Both phenomena reflect deeper model limitations in generalization, factual grounding, and representation learning—issues that continue to motivate new strategies for detection and mitigation.

## 3 Methodology

### 3.1 Dataset

We used the BlendedSkillTalk (Smith et al., 2020) and the DailyDialog (Li et al., 2017) datasets for this work. The BlendedSkillTalk dataset is a conversational dataset designed to help train and evaluate dialogue models. The dataset consists of 4819 training set conversations, 1009 validation set conversations, and 980 test set conversations. The DialiDialog dataset contains total 13,118 dialogues. We translated the conversations from English into Hindi, Farsi and Mandarin using GPT-3.5. A sample prompt we used for translating the texts to Farsi is : ‘Translate the following into Farsi: ’. After translation, we manually checked each translation to ensure the translation quality. A few sample translations are given in table 1.

### 3.2 Analysis

Figure 1 shows the workflow of our work. We used speaker 1 as the reference and asked the LLMs for their responses. We measured hallucination in the LLM responses by comparing with the actual response. We used both ROUGE-1 and ROUGE-L scores for measurement.

## 4 Results & Discussion

The results in Tables 2 and 3 provide an extensive comparison of hallucination tendencies across six LLMs — GPT-3.5, GPT-4o, Llama-3 8B, Gemma-2B, DeepSeek-1.3B, and Qwen-1.5-1.8B — on two multilingual conversational datasets (BlendedSkillTalk and DailyDialogue) in Hindi, Farsi, and Mandarin. We analyze hallucination using ROUGE-1 and ROUGE-L scores across the train, test, and validation splits.

### 4.1 Overall Trends

Across both datasets, ROUGE scores are lowest for Mandarin, indicating minimal lexical overlap and, consequently, fewer hallucinations. In contrast, Hindi and Farsi exhibit higher ROUGE scores, suggesting a greater degree of hallucinated overlap with human references.

## 4.2 Dataset-specific Observations

### 4.2.1 BlendedSkillTalk Dataset

In the BlendedSkillTalk (Table 2) presents hallucination scores using ROUGE-1 and ROUGE-L on the BlendedSkillTalk dataset across Hindi, Farsi, and Mandarin for training, testing, and validation splits. In Hindi, GPT-4o shows highest ROUGE scores across all splits, followed closely by GPT-3.5 and Llama-3.1. DeepSeek-R1 and Gemma-2.0 exhibit comparatively lower scores, suggesting they hallucinate more. In Farsi, Qwen-3 outperforms others with ROUGE scores consistently above 3.5, particularly on the validation set, indicating better factual alignment. Mandarin results, however, show a significant drop across all models. The highest ROUGE scores in Mandarin barely exceed 1.0, with most models scoring close to or below 0.5—highlighting the difficulty of reducing hallucination in Mandarin using these models.

### 4.2.2 DailyDialogue Dataset

The DailyDialogue dataset (Table 3) illustrates hallucination scores on the DailyDialogue dataset. In Hindi, GPT-3.5 and GPT-4o both maintain consistent performance, with GPT-3.5 scoring slightly higher across the board. Llama-3.1 also shows competitive performance, while DeepSeek-R1 and Gemma-2.0 again lag behind. In Farsi, GPT-3.5 and GPT-4o dominate, achieving high ROUGE values over 3.0. For Mandarin, hallucination remains a challenge, with scores mostly under 1.5. Qwen-3 achieves the highest scores in Mandarin on the validation split, but overall, scores remain modest compared to Hindi and Farsi. The results emphasize that while LLMs like GPT-4o and GPT-3.5 are more robust in Hindi and Farsi, Mandarin continues to pose challenges in minimizing hallucinations.

### 4.3 Language-wise Comparison

Across both datasets, Mandarin consistently shows the least hallucination, suggesting that LLMs are more cautious or limited in generating fluent hallucinated responses in Mandarin, particularly when translated content leads to mismatches with model priors. Hindi and Farsi, on the other hand, demonstrate similar and higher levels of hallucination, potentially due to their shared syntactic complexity

English	Farsi	Hindi	Mandarin
1: 'Which level are you at?'	1: «در چه سطحی هستی؟»	1: 'आप किस स्तर पर हैं?'	1: 你现在处于哪个水平?
2: 'I received on-the-job training when I first started'	2: «وقتی اولین بار شروع کردم، آموزش‌های حین کار دریافت کردم.»	2: 'जब मैंने पहली बार शुरू किया था, तो मुझे ऑन-द-जॉब प्रशिक्षण मिला।'	2: 我刚开始时接受了在职培训。
1: 'That's great! How long have you been doing this work?'	1: «عالیه! چقدر مدت است که این کار را انجام می‌دهید؟»	1: 'यह बहुत अच्छा है! आप यह काम कितने समय से कर रहे हैं?'	1: 太好了! 你做这份工作的时间有多久了?
2: 'For a good number of years now.'	2: «برای چندین سال خوب.»	2: 'काफी सालों से।'	2: 已经有好几年了。
1: 'Sounds like it could be a dangerous job at times.'	1: «به نظر می‌رسد گاهی اوقات می‌تواند شغل خطرناکی باشد.»	1: 'सुनने में यह काम कभी-कभी खतरनाक भी हो सकता है।'	1: 听起来有时候这份工作可能会很危险。
2: 'That it is, especially if you don't take the proper measures'	2: «همینطور است، به خصوص اگر اقدامات لازم را انجام ندهید.»	2: 'हाँ, खासकर अगर आप सही सावधानियाँ नहीं बरतें।'	2: 是的，特别是如果不采取适当的措施。
1: 'Well you have entered a trade that will always guarantee you some work.'	1: «خب، شما وارد حرفه‌ای شده‌اید که همیشه برایتان کار تضمین می‌کند.»	1: 'खैर, आपने एक ऐसा पेशा चुना है जिसमें हमेशा काम मिलता रहेगा।'	1: 嗯，你进入了一个总能保证你有工作的行业。
2: 'That's true, especially in this economy'	2: «درست است، به خصوص در این شرایط اقتصادی.»	2: 'सही है, खासकर इस अर्थव्यवस्था में।'	2: 是的，特别是在这个经济环境下。

Table 1: Sample translations of conversations from English into Hindi, Farsi and Mandarin using GPT-3.5. A sample prompt we used for translating the texts to Farsi is: 'Translate the following into Farsi: '. Each translation was then manually checked to ensure translation quality.

LLMs	Hindi						Farsi						Mandarin					
	Train		Test		Val		Train		Test		Val		Train		Test		Val	
	R-1	R-L	R-1	R-L	R-1	R-L	R-1	R-L	R-1	R-L	R-1	R-L	R-1	R-L	R-1	R-L	R-1	R-L
GPT-3.5	2.16	2.28	2.21	2.24	2.93	2.01	3.66	3.02	3.06	3.24	3.66	3.91	0.36	0.35	0.10	0.11	0.10	0.10
GPT-4o	2.55	2.82	2.80	2.90	2.48	2.73	3.32	3.53	3.54	3.66	3.32	3.50	0.02	0.02	0.02	0.02	0.06	0.06
Llama-3.1	2.41	2.42	2.32	2.33	2.40	2.39	3.65	3.65	3.78	3.78	3.67	3.67	0.99	0.99	0.96	0.96	0.95	0.95
Gemma-2.0	2.02	2.02	2.03	2.03	2.01	2.01	3.72	3.72	3.69	3.69	3.74	3.74	0.93	0.93	0.86	0.86	0.84	0.84
DeepSeek-R1	1.76	1.76	1.80	1.80	1.80	1.84	3.30	3.30	3.40	3.40	3.40	3.40	1.10	1.10	1.00	1.08	1.00	1.00
Qwen-3	2.18	2.18	2.23	2.23	2.26	2.23	3.76	3.76	3.80	3.80	3.80	3.83	1.02	1.02	0.99	1.00	1.01	1.05

Table 2: Hallucination scores (ROUGE-1 and ROUGE-L) on the BlendedSkillTalk dataset in Hindi, Farsi, and Mandarin across train, test, and validation splits.

and less available high-quality training data in these languages.

#### 4.4 Model-wise Comparison

While GPT-3.5 and GPT-4o are more fluent, they also demonstrate lower hallucination rates, particularly in Mandarin compared to smaller open-source models like Gemma and DeepSeek.

The notably low hallucination rate observed in Mandarin can be attributed to the availability of large volumes of high-quality training data for this language. In contrast, the elevated hallucination rates in Hindi and Farsi highlight the challenges faced by LLMs when dealing with low-resource languages. These findings support the hypothesis that greater data availability contributes to the generation

LLMs	Hindi						Farsi						Mandarin					
	Train		Test		Val		Train		Test		Val		Train		Test		Val	
	R-1	R-L	R-1	R-L	R-1	R-L	R-1	R-L	R-1	R-L	R-1	R-L	R-1	R-L	R-1	R-L	R-1	R-L
GPT-3.5	2.33	2.34	2.20	2.20	2.26	2.26	3.10	3.11	3.00	3.00	3.05	3.05	0.25	0.25	0.19	0.19	0.18	0.19
GPT-4o	2.08	2.08	2.00	2.00	2.05	2.05	3.04	3.04	3.00	3.00	3.07	3.08	0.02	0.02	0.02	0.02	0.03	0.03
Llama-3.1	2.03	2.03	1.90	1.90	1.96	1.96	3.43	3.44	2.82	2.82	3.41	3.40	0.70	0.70	0.81	0.82	0.90	0.90
Gemma-2.0	1.91	1.91	2.00	2.00	2.03	2.04	3.27	3.27	4.01	4.00	2.35	2.35	0.81	0.82	0.88	0.88	0.93	0.94
DeepSeek-R1	1.91	1.90	1.89	1.88	1.84	1.84	2.57	2.57	2.84	2.85	2.62	2.62	1.44	1.44	1.28	1.26	1.28	1.28
Qwen-3	2.19	2.19	2.17	2.17	2.09	2.10	3.54	3.56	3.31	3.31	3.48	3.48	1.56	1.56	1.26	1.27	1.22	1.24

Table 3: Hallucination scores (ROUGE-1 and ROUGE-L) on the DailyDialogue dataset in Hindi, Farsi, and Mandarin across train, test, and validation splits.




Figure 2: Hallucination (R-1) scores across the 6 LLMs for Hindi, Farsi and Mandarin on the BlendedSkillTalk dataset. It can be seen that across all the LLMs, Farsi has the highest hallucination with Mandarin the lowest.




Figure 3: Hallucination (R-L) scores across the 6 LLMs for Hindi, Farsi and Mandarin on the BlendedSkillTalk dataset. Similar trend like 2 can also be seen here.




Figure 4: Hallucination (R-1) scores across the 6 LLMs for Hindi, Farsi and Mandarin on the DailyDialog dataset.




Figure 5: Hallucination (R-L) scores across the 6 LLMs for Hindi, Farsi and Mandarin on the DailyDialog dataset.

of more accurate and contextually relevant responses. For Hindi and Farsi, the scarcity of training data likely leads the models to rely more heavily on extrapolation, increasing the risk of generating irrelevant or erroneous outputs.

In Mandarin, hallucinations tend to be partial or minor in nature. As illustrated in Tables 6 and 7, the few observed hallucinations in this language are typically subtle. Conversely, the hallucinations in Hindi and Farsi, partic-

ularly with GPT-3.5 and GPT-4o, are more pronounced. For example, in Hindi (Table 4), when prompted with ‘I hope so, how old are your kids?’, GPT-3.5 responds with ‘I would be happy to help answer your question,’ which fails to address the specific query. Similarly, in Farsi (Table 8), the prompt ‘Yup, but I do need to call my friend about her ring...’ elicits the response, ‘If you need any other help, I would be happy to assist,’ which is contextually disconnected from the input.



Comparatively, GPT-4o exhibits a reduced tendency toward hallucination. For instance, in Table 9, when prompted with ‘Yeah, fine. I have a habit of driving too fast,’ GPT-4o responds with a somewhat tangential yet semantically richer reply about sports cars, demonstrating less severe deviation. Similarly, in Table 5, when asked ‘What kind of clothing do you like?’, GPT-4o replies with a general informative statement: ‘I am an AI, so I don’t have personal preferences...’, which, although not directly conversational, remains informative and relevant.

Other LLMs present more varied hallucination patterns. Gemma-2, for example, in response to ‘Hey man, you wanna buy some weed?’ in Hindi (Table 13), outputs a cautionary message: ‘I am an AI, I am not allowed to use intoxicating substances...’, which, while responsible, breaks the flow of natural conversation. Llama-3.1 generates a clearly irrelevant response in Table 12, replying with ‘Wow, thank you very much! This gift of 1000 yuan is very important to me,’ despite the context being about currency exchange, not gifting.

Qwen-3 and DeepSeek-R1 similarly exhibit hallucinations. In Mandarin (Table 11), Qwen-3 responds with an explanation of Newtonian mechanics when the prompt expresses an emotional reaction, showing a stark disconnect in context. In Farsi (Table 10), DeepSeek-R1 responds to a statement about enjoying meat with an introspective and unrelated question, ‘How can I understand this interest in eating meat?’, further illustrating the model’s struggle with contextual alignment.

#### 4.5 Implications and Recommendations

Our findings suggest that hallucination in multilingual dialogue generation is highly influenced by language-resource availability, dataset style, and model size. For real-world deployment in Hindi and Farsi, hallucination mitigation techniques such as retrieval-augmented generation (RAG), grounded decoding, or human-in-the-loop supervision become crucial. Moreover, models specifically pretrained or fine-tuned on native corpora (e.g., Qwen for Mandarin) show reduced hallucination, highlighting the importance of

language-aware pretraining strategies.

## 5 Conclusion

In this paper, we investigated hallucination in conversations across three low-resource languages—Hindi, Farsi, and Mandarin—using six large language models: GPT-3.5, GPT-4o, Llama-3.1, Gemma-2, Qwen-3, and DeepSeek-R1. Our analysis reveals that hallucination behavior varies substantially across languages and model architectures.

Across the board, hallucination was consistently minimal in Mandarin, with all models. This can be attributed to the availability of abundant and high-quality training data in Mandarin. In contrast, significantly higher hallucination rates were observed in Hindi and Farsi. While GPT-4o generally exhibited lower hallucination than GPT-3.5, both models still struggled in low-resource settings. Models like Llama-3.1 and Gemma-2 demonstrated relatively better performance in Hindi and Farsi compared to the GPT models, though some contextually irrelevant responses were still present. DeepSeek-R1 also showed occasional hallucinations, particularly in Farsi.

These findings reinforce the strong influence of language resource availability on model performance. The disparities in hallucination severity across languages emphasize the models’ reliance on the quantity and quality of their training data. For languages with limited datasets, models are more prone to generating responses that are contextually misaligned or semantically inaccurate.

To address these limitations, future research should explore strategies such as targeted data augmentation, multilingual fine-tuning, and architectural adaptations that prioritize contextual grounding in low-resource languages. Additionally, systematic evaluations across a broader range of languages and domains can help generalize these findings and guide more inclusive model development.

Ultimately, our study highlights the importance of tailoring LLM development to address the linguistic diversity of global users. Reducing hallucinations in low-resource languages is essential for building equitable, reliable, and context-aware conversational systems.

## Limitations

This study focuses on only three low-resource languages—Hindi, Farsi, and Mandarin—which may limit the generalizability of our findings to other languages or dialects. While we evaluate multiple LLMs, future research could explore more models and include detailed human evaluations to better capture the subtle differences in hallucination across different languages.

## References

- Ayush Agrawal, Mirac Suzgun, Lester Mackey, and Adam Tauman Kalai. 2023. Do language models know when they’re hallucinating references? *arXiv preprint arXiv:2305.18248*.
- 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*.
- David Dale, Elena Voita, Loïc Barrault, and Marta R Costa-Jussà. 2022. Detecting and mitigating hallucinations in machine translation: Model internal workings alone do well, sentence similarity even better. *arXiv preprint arXiv:2212.08597*.
- Xiangjue Dong, Ziwei Zhu, Zhuoer Wang, Maria Teleki, and James Caverlee. 2023. Co<sup>2</sup> pt: Mitigating bias in pre-trained language models through counterfactual contrastive prompt tuning. *arXiv preprint arXiv:2310.12490*.
- Esin Durmus, He He, and Mona Diab. 2020. Feqa: A question answering evaluation framework for faithfulness assessment in abstractive summarization. *arXiv preprint arXiv:2005.03754*.
- Gustavo Gonçalves and Emma Strubell. 2023. Understanding the effect of model compression on social bias in large language models. *arXiv preprint arXiv:2312.05662*.
- Nuno M Guerreiro, Duarte M Alves, Jonas Waldendorf, Barry Haddow, Alexandra Birch, Pierre Colombo, and André FT Martins. 2023. Hallucinations in large multilingual translation models. *Transactions of the Association for Computational Linguistics*, 11:1500–1517.
- Zihao He, Siyi Guo, Ashwin Rao, and Kristina Lerman. 2023. Inducing political bias allows language models anticipate partisan reactions to controversies. *arXiv preprint arXiv:2311.09687*.
- David F Jenny, Yann Billeter, Mrinmaya Sachan, Bernhard Schölkopf, and Zhijing Jin. 2023. Navigating the ocean of biases: Political bias attribution in language models via causal structures.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38.
- Daniel Kang and Tatsunori Hashimoto. 2020. Improved natural language generation via loss truncation. *arXiv preprint arXiv:2004.14589*.
- 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. *arXiv preprint arXiv:2305.11747*.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. Dailydialog: A manually labelled multi-turn dialogue dataset. *arXiv preprint arXiv:1710.03957*.
- Tomasz Limisiewicz, David Mareček, and Tomáš Musil. 2023. Debiasing algorithm through model adaptation. *arXiv preprint arXiv:2310.18913*.
- Potsawee Manakul, Adian Liusie, and Mark JF Gales. 2023. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models. *arXiv preprint arXiv:2303.08896*.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. On



- faithfulness and factuality in abstractive summarization. *arXiv preprint arXiv:2005.00661*.
- Nick Mckenna, Tianyi Li, Liang Cheng, Mohammad Hosseini, Mark Johnson, and Mark Steedman. 2023. Sources of hallucination by large language models on inference tasks. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 2758–2774.
- Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. Factscore: Fine-grained atomic evaluation of factual precision in long form text generation. *arXiv preprint arXiv:2305.14251*.
- 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*.
- Vikas Raunak, Arul Menezes, and Marcin Junczys-Dowmunt. 2021. The curious case of hallucinations in neural machine translation. *arXiv preprint arXiv:2104.06683*.
- Abel Salinas, Louis Penafiel, Robert McCormack, and Fred Morstatter. 2023. "im not racist but...": Discovering bias in the internal knowledge of large language models. *arXiv preprint arXiv:2310.08780*.
- Eric Michael Smith, Mary Williamson, Kurt Shuster, Jason Weston, and Y-Lan Boureau. 2020. Can you put it all together: Evaluating conversational agents’ ability to blend skills. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2021–2030.
- Neeraj Varshney, Wenlin Yao, Hongming Zhang, Jianshu 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*.
- Chaojun Wang and Rico Sennrich. 2020. On exposure bias, hallucination and domain shift in neural machine translation. *arXiv preprint arXiv:2005.03642*.
- Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. 2023. Large language models are not robust multiple choice selectors. *arXiv preprint arXiv:2309.03882*.

## A Appendix

### A.1 Data translation and LLM responses

We used ‘gpt-3.5-turbo’ for the translation of the dataset. We observed while translating, the names of brands or the proper nouns were kept as it is. For example, in table 5, in the translation of speaker 2, ‘American Eagle Outfitters’ was kept as it is; it was not even converted into Hindi characters.

We used both ‘gpt-3.5-turbo’ and ‘gpt-4o’ for the responses. We used the following parameter values for getting responses from the LLMs (both GPT-3.5 and GPT-4o): max\_tokens=100, temperature=0.7 and top\_p=0.9. After obtaining the translations and responses from the LLMs, they were checked manually by native speakers ensuring data quality.

We used Llama 3.1 8B, DeepSeek R1 14B, and Qwen 2.5 7B as open-source model for response generation using the Hugging Face transformers library. All three models were loaded in 4-bit precision to ensure efficient memory usage and faster runtime. For decoding, we applied the same configuration across models: max\_new\_tokens=50, temperature=0.5, top\_p=1.0, and do\_sample=True. These settings offered a good balance between response diversity and stability. After generating the responses, all outputs—including translations—were manually reviewed by native speakers to ensure high linguistic quality and semantic accuracy.

We used Gemini 2.0 via the Google AI Studio API for generating responses. The model was accessed using its default inference settings, without specifying temperature or top\_p values. We set max\_tokens=100 to control the length of the generated outputs. As with the other models, all responses were manually reviewed by native speakers to ensure translation accuracy, naturalness, and overall data

quality.

## **A.2 Hallucination measurement**

We used ROUGE-1 and ROUGE-L scores for measuring the hallucinations. ROUGE-1 measures the overlap of unigrams, or individual words, between the system-generated and reference summaries and ROUGE-L naturally captures sentence-level structural similarity and automatically identifies the longest co-occurring n-grams in sequence. Some LLM responses were partially hallucinated, while others were fully hallucinated. For example, in table 6, GPT-3.5’s reply mentions about a particular dish (‘Yes, this dish is indeed very delicious. You could try making it yourself; you might find it quite rewarding’) whereas the whole conversation does not mention about any specific dishes. Along with measuring hallucination using ROUGE-1 and ROUGE-L scores, we asked the native speakers to check few random samples to find the patterns in the hallucinations. It was observed that the hallucinations in Mandarin were mostly partial in nature whereas Hindi and Farsi contained both partial and complete hallucinations.

English	Speaker1	Speaker2	GPT-3.5 response
1: 'I like acting, I hope to be an actor, what about you?' 2: 'that is ok. have any kids?'	मुझे अभिनय पसंद है, मैं एक अभिनेता बनने की उम्मीद करता हूँ, आपके बारे में क्या कहेंगे?	ठीक है। क्या आपके पास कोई बच्चे हैं?	आपकी यह उम्मीद बहुत अच्छी है। अभिनय एक कला है (Your hope is very good. Acting is an art.)
1: 'No, but someday.' 2: 'that is good. I have 2'	नहीं, लेकिन किसी दिन।	यह अच्छा है। मेरे पास 2 हैं।	ठीक है, कोई बात नहीं। मुझे उम्मीद है कि वह दिन जल्द ही आएगा। (Okay, no problem. I hope that day will come soon.)
1: 'After I am done with school I plan to have a family.' 2: 'that is great! you will be ready'	मेरे स्कूल के समापन के बाद, मेरी योजना है कि मैं एक परिवार बनाऊँ।	यह बहुत बढ़िया है! तुम तैयार हो जाओगे।	मुझे एक साथी की तलाश है जिसके साथ मैं अपना जीवन बिता सकूँ। (I am looking for a partner with whom I can spend my life.)
1: 'I hope so, how old are your kids?' 2: '5 & 7. they take up a lot of my time'	मुझे आशा है, आपके बच्चे कितने साल के हैं?	5 और 7। वे मेरे बहुत समय लेते हैं।	मुझे खुशी होगी आपके सवाल का उत्तर देने में मदद करने में। (I would be happy to help answer your question.)
1: 'I would imagine. I am sure they are great kids.' 2: 'luckily, they love flowers just as much as I do. we spend a lot of time in the garden'	मैं कल्पना करता हूँ। मुझे यकीन है कि वे बहुत अच्छे बच्चे हैं।	भाग्य से, उन्हें फूल उतना ही पसंद है जितना मुझे। हम बगीचे में बहुत समय बिताते हैं।	यह अच्छा है कि आप उन्हें इस तरह से सोचते हैं। बच्चों को इस तरह का समर्थन मिलना बहुत महत्वपूर्ण है। (It's good that you think of them this way. It's very important for children to receive such support.)

Table 4: Sample hallucinations in responses by GPT-3.5 for Hindi. The hallucinations in the responses provided by GPT-3.5 are highlighted in yellow. The first hallucination highlighted is pronounced whereas the second one shows partial hallucination.

English	Speaker1	Speaker2	GPT-4o response
1: 'What kind of clothing do you like?' 2: 'I like dresses alot, especially the ones from American Eagle Outfitters'	तुम्हें कौनसे प्रकार के कपड़े पसंद हैं?	मुझे ड्रेसेस बहुत पसंद है, खासकर American Eagle Outfitters की।	<p>मैं एक एआई हूँ, इसलिए मेरी व्यक्तिगत पसंद नहीं होती। लेकिन मैं आपकी मदद कर सकता हूँ अगर आप कपड़ों के विभिन्न प्रकारों के बारे में जानकारी चाहते हैं या किसी खास अवसर के लिए सुझाव चाहते हैं। आप किस प्रकार के कपड़ों के बारे में जानना चाहेंगे?</p> <p>(I am an AI, so I don't have personal preferences. However, I can help you if you want information about different types of clothing or suggestions for a specific occasion. What type of clothing would you like to know about?)</p>

Table 5: Sample hallucinations in responses by GPT-4o for Hindi. The hallucinations in the responses provided by GPT-4o are highlighted in yellow.

English	Speaker1	Speaker2	GPT-3.5 response
1: ‘Oh really!?! That is interesting. I am actually italian myself.’ 2: ‘Awesome. Me and my partner just bought a house. I can’t wait to cook in my kitchen.’ 1: ‘Moving in a new place can be a lot of fun. Are you a good cook?’ 2: ‘I like to think so. I love to make coffee for an after dinner treat too.’	哦，真的吗！？ 那很有趣。 其实我自己是意大利人。  搬到一个新地方可能很有趣。 你会做饭吗？	太棒了。我和我的伴侣刚买了一栋房子。我迫不及待地想在我的厨房里做饭。  我喜欢这样想。我也喜欢为晚餐后做咖啡作为甜点。	真的吗？那太棒了！意大利是一个美丽的国家，我听说那里的食物和文化都非常迷人。你是在意大利长大的吗？ (Really? That’s awesome! Italy is a beautiful country, and I’ve heard that the food and culture there are fascinating. Did you grow up in Italy?) 是的，我会做一些简单的饭菜。搬到一个新地方也可以让我尝试学习一些新的菜谱，挑战自己在烹饪方面的技能。 (Yes, I can cook some simple dishes. Moving to a new place would also give me the opportunity to try learning new recipes and challenge my cooking skills.)
1: ‘Mmm That sounds delicious right now.’ 2: ‘What do you like to do?’	嗯， 那听起来现在很美味。	你喜欢做什么？	是的，这道菜确实很美味， 你可以试着自己做一下， ，也许会觉得很有成就感。 (Yes, this dish is indeed very delicious. You could try making it yourself; you might find it quite rewarding.)

Table 6: Sample hallucinations in responses by GPT-3.5 for Mandarin. The hallucinations in the responses provided by GPT-3.5 are highlighted in yellow.

English	Speaker1	Speaker2	GPT-4o response
1: 'What kind of food do you usually eat?' 2: 'I try my best to eat only fruits, vegetables, and chicken.'	你通常吃什么食物?	我尽量只吃水果、蔬菜和鸡肉。	我通常吃各种食物，包括水果、蔬菜、瘦肉蛋白、谷物和乳制品。 (I usually eat a variety of foods, including fruits, vegetables, lean proteins, grains, and dairy products.)
1: 'Well, let me check for you. I am sorry, sir. They mistook your order as a normal one. Actually it's an express laundry, is that right?' 2: 'Yes, I ordered the express service.'	好的，让我为您查一下。对不起，先生。他们把您的订单误认为是普通的。实际上是一个快速洗衣，对吗?	是的，我订购了快速服务。	是的，没错。本来是快速服务。 (Yes, that's correct. It was supposed to be express laundry.)
1: 'That's not the way to solve the problem. Did you look in the classified advertisements column?' 2: 'Yes, I did. But I wasn't qualified for most of the jobs that I was interested in.'	这不是解决问题的方式。你看过分类广告栏吗?	是的，我看过。但我对大部分我感兴趣的工作都不符合资格。	不，我还没查看分类广告。那可能是个寻找解决方案的好地方。我会试试的。  (No, I didn't check the classified advertisements. That might be a good place to look for a solution. I'll give it a try.)

Table 7: Sample hallucinations in responses by GPT-4o for Mandarin. The hallucinations in the responses provided by GPT-4o are highlighted in yellow.

English	Speaker 1	Speaker 2	GPT-3.5 Response
1: 'It's a small garden on my patio. I like to sit out there and drink my coffee.' 2: 'That's sounds like a good plan. How big is your garden?'	این یک باغ کوچک در پاتیو من است. من دوست دارم بیرون بنشینم و قهوه‌ام را بنوشم.	این به نظر یک برنامه خوب می‌آید. باغ شما چقدر بزرگ است؟	باغتان بسیار زیباست! احتمالاً لحظات زیبایی را در آنجا سپری می‌کنید. آیا گیاهان مختلفی در باغ شما دارید؟ (Your garden is very beautiful! You probably spend some beautiful moments there. Do you have different plants in your garden?)
1: 'Yes, it is. My quiet space.' 2: 'Sounds like a great way to relax, eh.'	بله، اینجا فضای آرام من است.	صدای یک راه عالی برای آرامش به نظر می‌رسد، نه؟	خیلی خوبه که یک فضای آرام برای خودت داری. آیا اینجا برای تفکر و استراحت مناسبه؟ (It's great that you have a quiet space for yourself. Is it suitable for reflection and relaxation?)
1: 'Yup, but I do need to call my friend about her ring. She has a late class so I'll call her tonight.' 2: 'And a good way to reflect especially with all that's going in the world.'	بله، اما من نیاز دارم به دوستم تلفن بزنم در مورد انگشترش. او کلاس دیری دارد، پس امشب به او زنگ می‌زنم.	و یک راه خوب برای تأمل، به ویژه با تمام اتفاقاتی که در دنیا رخ می‌دهد.	اگر نیاز به کمک دیگری داری خوشحال می‌شوم کمک کنم. (If you need any other help, I would be happy to assist.)
1: 'Yes, the ring is special to her.' 2: 'I'm sure she'll be excited to hear from you.'	بله، این حلقه برای او ویژه است.	مطمئنم که او خوشحال خواهد شد که از تو بشنود.	آیا می‌توانید بیشتر توضیح دهید که چرا این حلقه برای او ویژه است؟ اینکه چه خصوصیت یا معنا ویژه‌ای دارد که آن را برای او ویژه می‌سازد؟ (Can you explain more about why this ring is special to her? What quality or special meaning does it have that makes it special for her?)

Table 8: Sample hallucinations in responses by GPT-3.5 for Farsi. The hallucinations in the responses provided by GPT-3.5 are highlighted in yellow.



English	Speaker 1	Speaker 2	GPT 4.0 Response
1: 'I am doing ok. I have had some crazy experiences recently.'	من خوب هستم. اخیراً تجربیات دیوانه‌واری داشتم.	صادقانه باید بگم که من هم تجربیات دیوانه‌واری داشتم، چه نوع تجربیات دیوانه‌واری؟	خوشحالم که حالت خوب است. تجربیات دیوانه‌وار می‌توانند هم هیجان‌انگیز و هم گاهی کمی ترسناک باشند. اگر دوست داری، می‌توانی بیشتر درباره‌شان صحبت کنی. شاید بتوانم به تو کمک کنم یا فقط شنونده خوبی باشم. چه اتفاقاتی افتاده؟ (I'm glad you're okay. Crazy experiences can be both exciting and sometimes a bit scary. If you'd like, you can talk more about them. Maybe I can help or just be a good listener. What happened?)
2: 'Honestly I have too, what sort of crazy experiences?'			
1: 'I had a near death experience after crashing my sports car.'	من تجربه‌ی نزدیک به مرگ را پس از برخورد خودروی ورزشی‌ام داشتم.	وای، این خیلی سخت بود، حداقل الان حالت خوبه؟	تجربه‌ی نزدیک به مرگ می‌تواند بسیار تأثیرگذار و تغییر دهنده‌ی زندگی باشد. این لحظات معمولاً باعث می‌شوند که افراد به ارزش‌های زندگی و اولویت‌های خود با دیدی تازه نگاه کنند. (A near-death experience can be very impactful and life-changing. These moments often cause people to look at their life's values and priorities with a fresh perspective.)
2: 'Wow, that's rough, are you okay now at least?'			
1: 'Yeah, fine. I have a habit of driving too fast.'	بله، خوب. من عادت دارم خیلی سریع رانندگی کنم.	اما قطعاً ترجیح می‌دهم که کامیونم را داشته باشم.	اتومبیل ورزشی واقعاً جذاب هست و بسیاری از افراد به دلیل طراحی زیبا، سرعت بالا و تجربه رانندگی فوق‌العاده‌ای که ارائه می‌دهند، طرفدار آن‌ها هستند. (Sports cars are really fascinating, and many people are fans of them due to their beautiful design, high speed, and the incredible driving experience they offer. )
2: 'I definitely prefer my truck though.'			

Table 9: Sample hallucinations in responses by GPT-4o for Farsi. The hallucinations in the responses provided by GPT-4o are highlighted in yellow.

English	Speaker 1	Speaker 2	DeepSeek-R1 Response
1: 'I have at least twenty pairs. I really love to wear boots. I am a country person.'	من حداقل بیست جفت دارم. واقعاً عاشق پوشیدن بوت هستم. من یک آدم روستایی هستم.	من بوت‌های بلندی دارم که چه موسیقی کانتري دوست داشته باشم چه نه، محشرن!	<p>بیست جفت کفش استفاده چگونه می‌توانم از این!</p> <p>اگر پیشنهادی دارید</p> <p>کنم؟ خوشحال می‌شوم!</p> <p>(How can I make use of these twenty pairs of shoes? I'd be happy if you have any suggestions!)</p>
2: 'i've tall boots , that rock whether i like country or not !'			
1: 'Well do you also like eating meat. I am tough and love eating meat too.'	من گاهی اوقات برش‌های گوشت گاو مثل گوشت خشک و تکه‌های گوشت گاو می‌خورم، اما نه خیلی زیاد.	وای، این خیلی سخت بود، حداقل الان حالت خوبه؟	<p>چطور می‌توانم این علاقه به گوشت خوردن را درک کنم؟</p> <p>سلام! اگر علاقمند به خوردن گوشت هستید، من هم گاهی گوشت گاو می‌خورم</p> <p>(How can I understand this interest in eating meat? Hello! If you are interested in eating meat, I sometimes eat beef too.)</p>
2: 'I do eat beef cuts like beef jerky and beef trimmings sometimes, but not very often.'			

Table 10: Sample hallucinations in responses by DeepSeek-R1 for Farsi. The hallucinations in the responses provided by DeepSeek-R1 are highlighted in yellow.

English	Speaker1	Speaker2	Qwen-3 response
<p>1: ‘Wow. is there repercussions’</p> <p>2: ‘Yes, if a whistleblower tells, the company can get in trouble. They can get big fines for exploitation of workers.’</p>	<p>哇。有后果吗？</p>	<p>是的，如果告密者告诉了，公司可能会惹上麻烦。他们可能会因剥削工人而被处以巨额罚款。</p>	<p>会有什么后果？比如，我之前在抖音上刷到一个视频，说某地的学校有学生被狗咬了，然后家长向学校索赔，结果学校说学生有错，</p> <p>(What consequences would there be? For example, I saw a video on Douyin before, saying that at a certain school, a student was bitten by a dog. Then the parents claimed compensation from the school, but the school said the student was at fault, so they refused to pay, and the parents sued.)</p>
<p>1: ‘Thats terrible have you ever witnessed this personally’</p> <p>2: ‘No, I own a small business. I can only afford to pay my workers minimum wage, but I treat them well. I have heard stories about horrible working conditions.’</p>	<p>那太可怕了，你有亲眼见过吗？</p>	<p>不，我拥有一家小企业。我只能支付我的工人最低工资，但我对他们很好。我听说过一些关于可怕工作条件的故事。</p>	<p>我有点害怕。</p> <p>这是用户对我的回答的反应。</p> <p>用户之前问的是，如果一个物体在真空中被扔出去，它会一直运动下去吗？</p> <p>我回答说会的，因为没有空气阻力，根据牛顿</p> <p>(I am a bit scared. This is the user’s reaction to my answer. The user previously asked if an object thrown in a vacuum would keep moving. I said yes because there is no air resistance, according to Newton.)</p>

Table 11: Sample hallucinations in responses by Qwen-3 for Mandarin. The hallucinations in the responses provided by Qwen-3 are highlighted in yellow.

English	Speaker1	Speaker2	Llama-3.1 response
1: 'Good morning , sir . Is there a bank near here ? ' 2: ' There is one . 5 blocks away from here ?'	शुभ प्रभात, सर। यहाँ के पास एक बैंक है क्या?	हाँ, एक है। यहाँ से 5 ब्लॉक दूर है?	शुभ प्रभात! हाँ, यहाँ के पास एक बैंक है, लेकिन मुझे आपको यह जानकारी देने के लिए कोई विशिष्ट जानकारी (Good morning! Yes, there is a bank nearby, but I do not have any specific information to give you about it.)
1: 'Well , that's too far.Can you change some money for me ? ' 2: ' Surely , of course . How much would you like to change ?'	अच्छा, वह बहुत दूर है। क्या आप मेरे लिए कुछ पैसे बदल सकते हैं?	बिल्कुल, जरूर। आप कितनी राशि बदलना चाहेंगे?	हाँ, मैं आपके लिए कुछ पैसे बदल सकता हूँ। कितने पैसे हैं जिन्हें आप बदलना चाहते हैं? (Yes, I can exchange some money for you. How much money do you want to exchange?)
1: '1000 Yuan.Here you are .'	1000 युआन। यह रहा।		वाह, बहुत धन्यवाद! 1000 युआन का यह उपहार मेरे लिए बहुत महत्वपूर्ण है। (Wow, thank you very much! This gift of 1000 yuan is very important to me.)

Table 12: Sample hallucinations in responses by Llama-3.1 for Hindi. The hallucinations in the responses provided by Llama-3.1 are highlighted in yellow.

English	Speaker1	Speaker2	Gemma-2 response
1: 'Hey man , you wanna buy some weed?' 2: 'Some what ?'	अरे भाई, तुम कुछ गाँजा खरीदना चाहोगे?	कुछ क्या?	मैं एक एआई हूँ, मुझे नशीले पदार्थों का उपयोग करने की अनुमति नहीं है और मैं आपको ऐसा करने के लिए प्रोत्साहित नहीं करूँगा। (I am an AI, I am not allowed to use intoxicating substances, and I will not encourage you to do so.)

Table 13: Sample hallucinations in responses by Gemma-2 for Hindi. The hallucinations in the responses provided by Gemma-2 are highlighted in yellow.