

# ANHALTEN: Cross-Lingual Transfer for German Token-Level Reference-Free Hallucination Detection

Janek Herrlein<sup>1</sup>, Chia-Chien Hung<sup>2,3</sup>, Goran Glavaš<sup>1</sup>

<sup>1</sup>CAIDAS, University of Würzburg, Germany

<sup>2</sup>NEC Laboratories Europe, Heidelberg, Germany

<sup>3</sup>Data and Web Science Group, University of Mannheim, Germany

janek.herrlein@live.de

Chia-Chien.Hung@necclab.eu

goran.glavas@uni-wuerzburg.de

## Abstract

Research on token-level reference-free hallucination detection has predominantly focused on English, primarily due to the scarcity of robust datasets in other languages. This has hindered systematic investigations into the effectiveness of cross-lingual transfer for this important NLP application. To address this gap, we introduce ANHALTEN, a new evaluation dataset that extends the English hallucination detection dataset to German. To the best of our knowledge, this is the first work that explores cross-lingual transfer for token-level reference-free hallucination detection. ANHALTEN contains gold annotations in German that are parallel (i.e., directly comparable to the original English instances). We benchmark several prominent cross-lingual transfer approaches, demonstrating that larger context length leads to better hallucination detection in German, even without succeeding context. Importantly, we show that the sample-efficient few-shot transfer is the most effective approach in most setups. This highlights the practical benefits of minimal annotation effort in the target language for reference-free hallucination detection. Aiming to catalyze future research on cross-lingual token-level reference-free hallucination detection, we make ANHALTEN publicly available: <https://github.com/janekh24/anhaltten>

## 1 Introduction

Detecting hallucinations in large pretrained language models (e.g., Brown et al., 2020; Jiang et al., 2024) is critical for ensuring their reliability in real-world applications. Most existing hallucination detection benchmarks focus on *reference-based* tasks (e.g., summarization, machine translation, question answering) (Maynez et al., 2020; Rebuffel et al., 2022; Sadat et al., 2023), comparing model generated text against provided references. However, reference-based hallucination detection is not appropriate for free-form text generation, where

obtaining ground-truth references in real-time demands sufficient and accurate preceding retrieval step. To address these challenges, *reference-free* hallucination detection approaches have been introduced (Liu et al., 2022; Su et al., 2024), focusing on identifying inconsistencies within the generated context itself to effectively detect hallucinations in real-time. Besides, most research in hallucination detection has concentrated on *sentence* or *passage-level* (Dhingra et al., 2019; Manakul et al., 2023; Zhang et al., 2023), which is inadequate for real-time applications that require immediate feedback during text generation. Fine-grained, token-level reference-free hallucination detection benchmark is necessary for this purpose. However, research in this area has focused on English (Liu et al., 2022), primarily due to the lack of robust evaluation datasets in other languages. Creating token-level hallucination detection datasets for new languages (from scratch or using machine translation) is significantly more expensive and time-consuming than for most other NLP tasks, due to the need for accurate translation and adaptation of nuanced contexts and token-level annotations. The lack of multilingual evaluation benchmarks hinders the investigation of cross-lingual transfer approaches for token-level reference-free hallucination detection.

In this work, we target this gap and introduce ANHALTEN (germAN HALucinaTion dETectionN), a new benchmark derived from the English token-level reference-free hallucination detection dataset HADES (Liu et al., 2022). ANHALTEN is: (1) *reliable* – with complete texts and hallucination spans (i.e., labels) manually translated, and (2) *parallel* – the same set of texts and labels have been translated to German, enabling direct comparison of multilingual models and cross-lingual transfer approaches.

We then use ANHALTEN to benchmark a range of cross-lingual transfer approaches and simulate the real-world applications in multiple setups. Our

HADES	ANHALTEN	HADES	ANHALTEN
Not Hallucination		Hallucination	
haunted homes is a british reality television series made by september films productions . [...] the show centers around <b>writer richard hillier</b> ( who owns the rights to the <b>story</b> ) , <b>ghostwriter andrew scott smith</b> ( pilot ) , only <b>aired</b> due to his lack of confidence level ) [...] they spend <b>the weekend</b> in a supposedly haunted house , <b>hoping</b> to find out if there are any ghosts around , [...]	haunted homes ist eine britische reality - fernsehserie , die von september films productions produziert wird . [...] im mittelpunkt der sendung stehen der autor richard hillier ( der die rechte an der geschichte besitzt ) , der ghostwriter andrew scott smith ( pilotfilm , der aufgrund seines mangelnden vertrauenslevels nur ausgestrahlt wurde ) [...] sie verbringen das wochenende in einem vermeintlichen spukhaus , <b>in der hoffnung</b> herauszufinden , ob es dort geister gibt , [...]	ieva zunda ( born 20 july 1978 in tukums ) is a latvian athlete . [...] she did not make it past the first round at the 1999 and 2003 world championships . [...] in 2008 [...] shortly before the deadline - on 28 june , she had finally reached the <b>qualifying</b> standard in the <b>400 m</b> ( 56 . 50 ) , as she clocked in the <b>first round</b> . she finished <b>third</b> in her heat , again missing out on a place in the <b>first</b> round .	ieva zunda ( geboren am 20 . juli 1978 in tukums ) ist eine lettische leichtathletin . [...] bei den weltmeisterschaften 1999 und 2003 kam sie nicht über die erste runde hinaus . [...] 2008 versuchte sie erneut [...] kurz vor dem stichtag - am 28 . juni - hatte sie endlich die qualifikationsnorm über 400 m ( 56 . 50 ) erreicht , wie sie in der ersten runde lief . sie wurde dritte in ihrem lauf und verpasste erneut den einzug in die <b>erste</b> runde .
Word Spans: [105, 105]	Word Spans: [112, 114]	Word Spans: [153, 153]	Word Spans: [154, 154]

Table 1: Examples of HADES (Liu et al., 2022) as the perturbed version with token-level label to detect hallucination, and our proposed ANHALTEN machine-translated and post-edited text. The **bold** terms indicate the perturbed words compared to the original Wiki (Guo et al., 2020), and the underline term presents the token required to detect hallucination. For brevity, the compared version with original Wiki is available in Appendix A.

HADES	MACHINE TRANSLATED (MT)	ANHALTEN (MT & Post-Edited)
other similar shows include most haunted and ghost home . it is also shown in the u . s . on the discovery channel fridays and saturdays schedule .	andere ähnliche shows sind most haunted und ghost home . es ist auch in den <b>u . s .</b> , auf dem discovery channel freitags und samstags <b>schedule</b> gezeigt .	andere ähnliche shows sind most haunted und ghost home . es wird auch in den <b>usa</b> auf dem discovery channel freitags und samstags gezeigt .
dold ' s research in algebraic topology , in particular , his views on fixed - point topology has made him influential in economics as well as mathematics .	dold ' s forschung in der algebraischen topologie , insbesondere , seine ansichten über fixpunkt-topologie hat ihn einflussreich in der wirtschaft als auch in der mathematik .	<b>dolds</b> forschung in der algebraischen topologie , insbesondere , seine ansichten über fixpunkt - topologie hat ihn <b>sowohl</b> in der wirtschaft als auch in der mathematik einflussreich <b>gemacht</b> .

Table 2: Examples compared with original English HADES text, the automatic machine translation to German, and the final translation after manual post-editing. The **highlighted** texts indicate the errors that were corrected during post-editing. These errors primarily include incorrect translations, grammatical mistakes, and missing information.

results show that (i) hallucination detection works comparably well even without succeeding texts, indicating that larger context length helps detect hallucinations in German, thus supporting proactive hallucination prevention *on-the-fly* during text generation, and (ii) few-shot transfer methods achieve high performance with minimal annotated data, highlighting the practical benefits of inexpensive annotation of a handful of target-language hallucination instances for training detection models.

## 2 Methodology

### 2.1 Dataset Creation

We translate the *full* development set and 10% of the training set of English HADES dataset (Liu et al., 2022) in German, with 1,000 and 876 instances, respectively.<sup>1</sup> Each instance includes a TEXT, MARKED WORD SPANS, POSITION OF MARKED WORD SPANS, and LABEL to indicate whether the MARKED WORD SPANS causes hallucination. Examples compared to the original English HADES dataset are shown in Table 1.

<sup>1</sup>Since the original test set labels were not published, we rely on training and development sets throughout our experiments. We also ensure the subsample of the training set retains the original label ratio of the training data.

Following the well-established practice (Hung et al., 2022; Senel et al., 2024), we carried out a two-phase translation process: (1) we started with an *automatic translation* – followed by (2) the *manual post-editing* of the translations. We first automatically translate the development and training set portions for both TEXT and MARKED WORD SPANS relying on DeepL Translator. We then incorporate native speaker with University degree and fluent in English, to post-edit the automatic translations to ensure the correctness of the translation – especially the directly preceding and succeeding context, and the correct determination of the MARKED WORD SPANS. Common errors identified in machine-translated texts include incorrect translations, missing words, grammatical mistakes, or contextual inaccuracies. Examples comparing the original English HADES with the automatically translated and manually post-edited texts are shown in Table 2. Besides, as the position of MARKED WORD SPANS changes in the German text<sup>2</sup>, the

<sup>2</sup>German and English, both Germanic languages, differ in ways that impact dataset design. In German, compound words are written as single words, whereas in English, they are separated by spaces, affecting MARKED WORD SPANS. Additionally, German commonly uses particle verbs, where MARKED WORD SPANS are split by other parts of the sentence.

POSITION OF MARKED WORD SPANS is adjusted accordingly.

Additionally, to conduct *Translate-Train* experiment for cross-lingual transfer, *full* training set (8,754 instances) are automatically translated using DeepL Translator, without post-editing. However, only 6,344 instances (72.5%) remain, since the discarded ones contain incorrect MARKED WORD SPANS.

## 2.2 Downstream Cross-Lingual Transfer

The parallel nature and substantial size of ANHALTEN facilitate benchmarking of cross-lingual transfer methods for hallucination detection tasks. We investigate three common methods for downstream cross-lingual transfer (XLT) (Ebing and Glavaš, 2023; Senel et al., 2024): (1) *Zero-Shot Transfer*, where we assume the absence of labeled task instances in the target language. The model is trained exclusively in English and is expected to perform the task directly in German without prior exposure to German labeled data. This method relies on the model’s capability to generalize knowledge from English to German. (2) *Few-Shot Transfer*, where a limited number of labeled instances in the target language exist with the majority of training data in the source language. The model is trained on abundant English data and a small amount of German data jointly,<sup>3</sup> helping it adapt to the specific nuances of the German language with limited annotated data. (3) *Translate-Train*, where training instances in source language are automatically translated (i.e., noisy) to target language leveraging the state-of-the-art machine translation model. While this approach relies on the quality of translation, it benefits from creating a substantial amount of training data in German, closely approximating a fully supervised learning scenario.

To facilitate modular and efficient XLT, adapter-based approach is proposed to learn specialized task and language adapters for high portability and parameter-efficient transfer to various tasks and languages (Pfeiffer et al., 2020b). For downstream XLT, a task adapter is stacked on the pre-trained source language adapter, where the parameters are only updated for the task adapter. During evaluation, the source language adapter is replaced by the

In such cases, only the conjugable main part of the verb is marked, while the particle is ignored.

<sup>3</sup>Compared to *sequential fine-tuning* (Lauscher et al., 2020; Hung et al., 2022), *joint fine-tuning* (Schmidt et al., 2022) on instances in both source and target language can achieve better performances with higher stability.

pre-trained target language adapter. In our setup, the task adapter is trained by (1) the English-only data for *Zero-Shot Transfer*; (2) a joint training of English and a small portion of German data for *Few-Shot Transfer*; or (3) the machine-translated English-to-German data for *Translate-Train*. The adapter-based approach ensures that the model can efficiently adapt to new tasks with minimal parameter updates, maintaining the balance between performance and computational efficiency.

## 3 Experimental Setup

### 3.1 Evaluation Tasks and Measures

We evaluate multilingual pre-trained language models (PLMs) in XLT methods (§ 2.2) for token-level reference-free hallucination detection tasks. To simulate real-world applications, we evaluate on two sub-tasks: *offline* and *online* (Liu et al., 2022). In the *offline* setting, the model accesses both preceding and succeeding contexts of the MARKED WORD SPANS, suitable for detecting hallucinations in pre-generated texts. In the *online* setting, the model considers only the preceding context, enabling proactive prevention of hallucinations during on-the-fly text generation.

We follow Liu et al. (2022) and evaluate the XLT capabilities utilizing multilingual PLMs on hallucination detection tasks. The evaluation metrics include accuracy, precision, recall, F1, Area Under Curve (AUC), G-Mean (Espíndola and Ebecken, 2005), and Brier Score (BS) (Brier, 1950). These metrics provide a comprehensive evaluation of model performance, balancing correctness, and the ability to handle imbalanced classes.

### 3.2 Models and Experimental Setup

Experiments are conducted on multilingual PLMs, namely multilingual BERT (mBERT) (Devlin et al., 2019) and XLMR (Conneau et al., 2020),<sup>4</sup> using *language adapters*<sup>5</sup> proposed by Pfeiffer et al. (2020b) to facilitate modular and efficient XLT.

To evaluate downstream XLT, the experiments are conducted with 5 runs in both *offline* and *online* settings, with a fixed context window of 200 tokens. In the *online* setting, the context includes the 200 tokens preceding the MARKED WORD SPANS. In the *offline* setting, it includes 100 tokens before and

<sup>4</sup>The weights of PLMs are loaded from HuggingFace: multilingual-bert-base-cased and xlm-roberta-base.

<sup>5</sup>The pre-trained adapters are selected from Adapter-Hub (Pfeiffer et al., 2020a) for English (en-wiki@ukp) and German (de-wiki@ukp).

	# Instances	Setting	Accuracy $\uparrow$	G-Mean $\uparrow$	BS $\downarrow$	AUC $\uparrow$	Not Hallucination			Hallucination		
							P $\uparrow$	R $\uparrow$	F1 $\uparrow$	P $\uparrow$	R $\uparrow$	F1 $\uparrow$
Zero-Shot	0	offline	62.82	59.38	25.51	72.90	<b>74.89</b>	41.87	53.37	58.18	<b>84.89</b>	<b>68.96</b>
Few-Shot	10	offline	63.86	61.89	24.23	73.32	72.84	47.72	57.34	59.57	80.87	68.51
Few-Shot	100	offline	65.12	63.87	23.28	73.88	72.64	51.46	60.12	60.93	79.51	68.94
Few-Shot	876	offline	<b>67.76</b>	<b>67.55</b>	<b>20.83</b>	<b>74.68</b>	68.24	69.75	<b>68.87</b>	67.50	65.67	66.44
Translate-Train	6344	offline	66.42	64.13	21.30	73.80	66.69	<b>72.44</b>	67.84	<b>69.98</b>	60.08	62.91
Zero-Shot	0	online	63.70	62.01	24.14	72.44	71.51	48.85	57.76	59.72	<b>81.14</b>	68.02
Few-Shot	10	online	63.50	61.53	23.84	72.43	<b>72.11</b>	47.29	56.88	59.30	80.58	<b>68.24</b>
Few-Shot	100	online	64.88	63.87	22.88	72.55	70.31	57.29	62.03	61.39	73.93	66.19
Few-Shot	876	online	67.28	<b>67.14</b>	21.22	<b>73.52</b>	67.89	68.89	68.33	66.75	65.59	66.09
Translate-Train	6344	online	<b>67.66</b>	66.55	<b>21.02</b>	73.20	65.66	<b>77.66</b>	<b>71.13</b>	<b>70.86</b>	57.13	63.19

Table 3: Cross-lingual transfer results of XLMR (%) averaged over 5 runs. According to Table 4, XLMR outperforms mBERT. For brevity, cross-lingual transfer results of mBERT are provided in Appendix B.

Model	Setting	Acc. $\uparrow$	G-Mean $\uparrow$	BS $\downarrow$	AUC $\uparrow$	Not Hallucination			Hallucination		
						P $\uparrow$	R $\uparrow$	F1 $\uparrow$	P $\uparrow$	R $\uparrow$	F1 $\uparrow$
mBERT	offline	61.00	56.12	26.49	69.66	71.04	42.92	50.68	57.89	80.04	66.54
XLMR	offline	62.82	59.38	25.51	<b>72.90</b>	<b>74.89</b>	41.87	53.37	58.18	84.89	<b>68.96</b>
mBERT	online	60.44	55.34	26.71	67.81	73.47	36.69	48.10	56.33	<b>85.46</b>	67.76
XLMR	online	<b>63.70</b>	<b>62.01</b>	<b>24.14</b>	72.44	71.51	<b>48.85</b>	<b>57.76</b>	<b>59.72</b>	81.14	68.02

Table 4: Zero-shot transfer results (%) averaged over 5 runs. Reference English performance of XLMR for accuracy: 70.40% (*offline*), 68.80% (*online*).

after the MARKED WORD SPANS.<sup>6</sup> During training, the instances are randomly split into a 70/30 train and validation split, while the original label ratio of training data is retained for the split. We train for 10 epochs in batches of 8 instances, with learning rate  $5 \cdot 10^{-3}$ , and a dropout ratio 0.2 is set to avoid overfitting.

## 4 Results and Discussions

We present and discuss the downstream XLT results on ANHALTEN for the token-level reference-free hallucination detection task across three XLT setups (§ 2.2): zero-shot transfer, few-shot transfer, and translate-train.

**Zero-Shot Transfer.** The results summarized in Table 4 highlight the performance of zero-shot transfer. Notably, XLMR consistently outperforms mBERT across most metrics, indicating that XLMR is better suited for zero-shot transfer. Minimal performance differences between the online and offline settings suggest that the selection of

<sup>6</sup>Liu et al. (2022) observed that model performance for English HADES dataset stabilizes around 80 tokens, with minimal performance differences between *offline* and *online* settings regarding context length. Thus, using 200 tokens would not limit performance, and increasing the context is unlikely to improve results.

large context windows does not significantly impact performance, aligning with findings from Liu et al. (2022). Having only preceding text with larger context lengths aids in detecting hallucinations, which is valuable for real-world applications, especially for proactively preventing hallucinations during on-the-fly generation. Compared with reference English performance, the zero-shot transfer results show significantly lower accuracy for both online and offline settings, with drops exceeding 5% points. These substantial performance declines underscore the inherent challenges in achieving reliable zero-shot XLT, which is consistent with the findings from prior work (Lauscher et al., 2020; Pfeiffer et al., 2020b).

**Few-Shot Transfer and Translate-Train.** As detailed in Table 3, few-shot transfer results for XLMR show remarkable improvements as the number of annotated German instances increases. With 10% of the English HADES training set (i.e., 876 annotated instances), accuracy improves by 4.9% points (offline) and 3.6% points (online) compared to zero-shot transfer. The corresponding G-Mean score increases by 8.2% points (offline) and 5.1% points (online). Notably, with only 100 annotated instances, accuracy improves by 2.3% points (offline) and 1.2% points (online), and the G-Mean score improves by 4.5% points (offline) and 1.9% points (online). This demonstrates the substantial impact of incorporating minimal annotated data on enhancing XLT performance. The translate-train approach, which involves translating a large corpus of 6,344 instances, yields accuracy gains of 3.6% points (offline) and 4.0% points (online) compared to zero-shot transfer. While beneficial, the marginal gains compared to few-shot transfer highlight the



POS	Accuracy ↑	G-Mean ↑	BS ↓	AUC ↑	Not Hallucination			Hallucination		
					P ↑	R ↑	F1 ↑	P ↑	R ↑	F1 ↑
Adjectives	<b>65.80</b>	<b>64.65</b>	<b>22.20</b>	<b>72.61</b>	<b>71.07</b>	53.55	61.06	<b>62.66</b>	<b>78.06</b>	<b>69.52</b>
Nouns	58.42	56.31	25.37	63.99	61.97	43.97	51.15	56.64	72.88	63.62
Verbs	52.16	37.20	28.13	59.69	51.24	<b>88.65</b>	<b>64.93</b>	58.78	15.68	24.64

Table 5: Part-of-Speech (POS) results of XLMR (%) in the *online* setting averaged over 5 runs. We only consider instances with MARKED WORD SPANS containing particular types of POS in the German language: adjectives, nouns, verbs.

practical efficiency of using smaller amounts of high-quality annotated data. Based on our findings, few-shot transfer emerges as a highly viable strategy for cross-lingual transfer of reference-free hallucination detection, offering robust performance gains over zero-shot transfer without the extensive resource required by the translate-train approach. This re-emphasizes the well-documented practical benefits of few-shot cross-lingual transfer (Lauscher et al., 2020; Schmidt et al., 2022), here for reference-free hallucination detection.

**Analysis.** According to Liu et al. (2022), nouns and verbs are the most frequently occurring part-of-speech (POS) in the MARKED WORD SPANS of the HADES dataset. The majority of instances with nouns (62.4%) and adjectives (74.0%) in the MARKED WORD SPANS belong to the hallucination class, while the majority of instances with verbs belong to the non-hallucination class (62.8%). This indicates a significant imbalance in label distribution. To assess the impact of this imbalance on cross-lingual transfer performance, we classify the validation set of ANHALTEN based on the selected POS (nouns, verbs, adjectives) in the MARKED WORD SPANS. Instances with MARKED WORD SPANS containing multiple words from different POS are excluded. To ensure an equal number of labels for each POS, we randomly remove instances from the more frequent class. This process results in 292 noun instances, 222 verb instances, and 62 adjective instances.

We then analyze the XLT results of XLMR in the online setting. The POS results in Table 5 show that adjectives are significantly more effective in detecting hallucinations compared to nouns and verbs in German. While the effectiveness of adjectives is notable, the imbalanced distribution of instances across different part-of-speech tags, as highlighted by Liu et al. (2022), warrants further investigation and consideration. Addressing these imbalances is crucial for improving the overall robustness and

accuracy of hallucination detection models.

We further conduct morphological analysis (detailed in Appendix B) and demonstrate that preceding words indicate grammatical gender in German impact model performance, underscoring the importance of linguistic context. These findings emphasize the need to address imbalances and encourage future work to enhance model performance concerning diverse linguistic features for token-level reference-free hallucination detection.

## 5 Conclusions

Token-level reference-free hallucination detection has predominantly focused on English, primarily due to the lack of robust benchmarks in other languages, hindering investigation into cross-lingual transfer approaches for this important task. To address this gap, we have presented ANHALTEN, an extension of the English HADES containing gold hallucination annotations in German, allowing for reliable and comparable cross-lingual estimates for token-level reference-free hallucination detection tasks. We utilized a modular adapter-based approach to facilitate the cross-lingual transfer, demonstrating the effectiveness of sample-efficient few-shot transfer. We believe that our dataset and findings advance the understanding of hallucination detection in cross-lingual transfer setups and contribute towards multilingual hallucination detection and real-time hallucination prevention in free-form text generation.

## Acknowledgements

This work was supported by the Alcatel-Lucent Stiftung and Deutsches Stiftungszentrum through the grant “Equitably Fair and Trustworthy Language Technology” (EQUIFAIR, Grant Nr. T0067/43110/23).

## Limitations

Despite the contributions of this research, several limitations are acknowledged, which present opportunities for future enhancement. Currently, ANHALTEN extends hallucination detection to German, broadening the scope beyond English but still covering only two languages. Expanding this research to include additional languages could further increase the global applicability of our findings. Besides, incorporating data from sources other than Wikipedia could enrich the diversity and complexity of the dataset. Additionally, extending the re-

search to include other types of hallucinations (e.g., subjective hallucinations) would provide a more comprehensive understanding of hallucination detection in various text types. We experimented on encoder-only multilingual PLMs, while decoder-based PLMs (e.g., [Le Scao et al., 2022](#); [Jiang et al., 2023](#); [Abdin et al., 2024](#)) warrants exploration. We hope that future research builds on top of our findings and extends the research toward more domains, more languages, and specifically with the efficiency and effective concerns of hallucination detection in different languages.

## Ethics Statement

This research addresses the critical need for non-English language datasets in hallucination detection by introducing ANHALTEN. The ethical considerations of this work are multifaceted. By extending hallucination detection to German, the research promotes linguistic diversity and inclusivity in AI systems. This inclusivity helps to mitigate biases and misinformation that can arise from language restrictions, fostering more equitable applications. The study also aims to facilitate the recognition of potential hallucinated content produced by large-scale pretrained models in free-form generation – could be useful in both *offline* and *online* settings. Additionally, the research outcome emphasizes the importance of resource-efficient approaches, reducing the reliance on extensive annotated data and promoting more sustainable development.

## References

- Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat Behl, et al. 2024. [Phi-3 technical report: A highly capable language model locally on your phone](#). *arXiv preprint arXiv:2404.14219*.
- Glenn W Brier. 1950. [Verification of forecasts expressed in terms of probability](#). *Monthly weather review*, 78(1):1–3.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. [Language models are few-shot learners](#). *Advances in neural information processing systems*, 33:1877–1901.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Bhuwan Dhingra, Manaal Faruqui, Ankur Parikh, Ming-Wei Chang, Dipanjan Das, and William Cohen. 2019. [Handling divergent reference texts when evaluating table-to-text generation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4884–4895, Florence, Italy. Association for Computational Linguistics.
- Benedikt Ebing and Goran Glavaš. 2023. [To translate or not to translate: A systematic investigation of translation-based cross-lingual transfer to low-resource languages](#). *arXiv preprint arXiv:2311.09404*.
- Rogério P Espíndola and Nelson FF Ebecken. 2005. [On extending f-measure and g-mean metrics to multi-class problems](#). *WIT Transactions on Information and Communication Technologies*, 35:25–34.
- Mandy Guo, Zihang Dai, Denny Vrandečić, and Rami Al-Rfou. 2020. [Wiki-40B: Multilingual language model dataset](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 2440–2452, Marseille, France. European Language Resources Association.
- Chia-Chien Hung, Anne Lauscher, Ivan Vulić, Simone Ponzetto, and Goran Glavaš. 2022. [Multi2WOZ: A robust multilingual dataset and conversational pre-training for task-oriented dialog](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3687–3703, Seattle, United States. Association for Computational Linguistics.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. [Mistral 7b](#). *arXiv preprint arXiv:2310.06825*.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. [Mixtral of experts](#). *arXiv preprint arXiv:2401.04088*.

- Anne Lauscher, Vinit Ravishankar, Ivan Vulić, and Goran Glavaš. 2020. [From zero to hero: On the limitations of zero-shot language transfer with multilingual Transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4483–4499, Online. Association for Computational Linguistics.
- Teven Le Scao, Angela Fan, Christopher Akiki, Elie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. 2022. [Bloom: A 176b-parameter open-access multilingual language model](#). *ArXiv*, abs/2211.05100.
- Tianyu Liu, Yizhe Zhang, Chris Brockett, Yi Mao, Zhifang Sui, Weizhu Chen, and Bill Dolan. 2022. [A token-level reference-free hallucination detection benchmark for free-form text generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6723–6737, Dublin, Ireland. Association for Computational Linguistics.
- 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.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.
- Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. 2020a. [AdapterHub: A framework for adapting transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 46–54, Online. Association for Computational Linguistics.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2020b. [MAD-X: An Adapter-Based Framework for Multi-Task Cross-Lingual Transfer](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7654–7673, Online. Association for Computational Linguistics.
- Clément Rebuffel, Marco Roberti, Laure Soulier, Geoffrey Scuttheeten, Rossella Cancelliere, and Patrick Gallinari. 2022. [Controlling hallucinations at word level in data-to-text generation](#). *Data Mining and Knowledge Discovery*, pages 1–37.
- Mobashir Sadat, Zhengyu Zhou, Lukas Lange, Jun Araki, Arsalan Gundroo, Bingqing Wang, Rakesh Menon, Md Parvez, and Zhe Feng. 2023. [DelusionQA: Detecting hallucinations in domain-specific question answering](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 822–835, Singapore. Association for Computational Linguistics.
- Fabian David Schmidt, Ivan Vulić, and Goran Glavaš. 2022. [Don’t stop fine-tuning: On training regimes for few-shot cross-lingual transfer with multilingual language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10725–10742, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Lütfi Kerem Senel, Benedikt Ebing, Konul Baghirova, Hinrich Schuetze, and Goran Glavaš. 2024. [Kardeş-NLU: Transfer to low-resource languages with the help of a high-resource cousin – a benchmark and evaluation for Turkic languages](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1672–1688, St. Julian’s, Malta. Association for Computational Linguistics.
- 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](#). *arXiv preprint arXiv:2403.06448*.
- Jiaxin Zhang, Zhuohang Li, Kamalika Das, Bradley Malin, and Sricharan Kumar. 2023. [SAC<sup>3</sup>: Reliable hallucination detection in black-box language models via semantic-aware cross-check consistency](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15445–15458, Singapore. Association for Computational Linguistics.

## A Dataset

The HADES dataset, introduced by Liu et al. (2022), is designed for reference-free token-level hallucination detection tasks in English. It is sourced from English Wikipedia (Guo et al., 2020), with extracted text segments that are first perturbed and then verified by crowd-sourced annotators to determine if the marked word spans in the text cause hallucination. The dataset is available under an open-source MIT License and contains a total of 10954 instances, divided into train, development, and test sets with sizes of 8754, 1000, 1200 respectively. Within the dataset, 54.5% of the instances are classified as *Hallucination*, while 45.5% of the instances are classified as *Not Hallucination*. Since the original labels of the test set were not published, we primarily rely on the training and development sets throughout our experiments.

To further facilitate research on cross-lingual transfer in German hallucination detection tasks, we propose ANHALTEN in this work. We manually annotate 876 training instances and the entire development set of 1000 instances, which are machine-translated and further post-edited. The proposed ANHALTEN compared with the original Wiki and perturbed HADES, is shown in Table 6.

Wiki (Original)	HADES (Perturbed)	ANHALTEN (MT & Post-Edited)
<b>Not Hallucination</b>		
haunted homes is a british reality television series made by september films productions . [...] the show centers around <b>psychic mia dolan</b> ( who owns the rights to the <b>programme</b> ) , <b>ghost hunter david vee</b> ( pilot <b>episode</b> , only <b>allegedly</b> due to his lack of confidence <b>presenting</b> ) , actor <b>mark webb</b> and <b>professor / sceptic chris french</b> . they spend <b>two nights</b> in a supposedly haunted house , hoping to find out if there are any ghosts around , [...]	haunted homes is a british reality television series made by september films productions . [...] the show centers around <b>writer richard hillier</b> ( who owns the rights to the <b>story</b> ) , <b>ghostwriter andrew scott smith</b> ( pilot , only <b>aired</b> due to his lack of confidence <b>level</b> ) , actor <b>paul newman</b> and <b>scientist / paranormal investigation officer chris martin</b> . they spend <b>the weekend</b> in a supposedly haunted house , <b>hoping</b> to find out if there are any ghosts around , [...]	haunted homes ist eine britische reality - fernsehserie , die von september films productions produziert wird . [...] im mittelpunkt der sendung stehen der autor richard hillier ( der die rechte an der geschichte besitzt ) , der ghostwriter andrew scott smith ( pilotfilm , der aufgrund seines mangelnden vertrauenslevels nur ausgestrahlt wurde ) , der schauspieler paul newman und der wissenschaftler / paranormale untersuchungsbeauftragte chris martin . sie verbringen das wochenende in einem vermeintlichen spukhaus , <b>in der hoffnung</b> herauszufinden , ob es dort geister gibt , [...]
	Word Spans: [105, 105]	Word Spans: [112, 114]
<b>Hallucination</b>		
ieva zunda ( born 20 july 1978 in tukums ) is a latvian athlete . her main event is the sprint and hurdles , but she also competes in the 400 and 800 metres . [...] she did not make it past the first round at the 1999 and 2003 world championships . [...] in 2008 [...] shortly before the deadline - on 28 june , she had finally reached the <b>entry</b> standard in <b>400 m hurdles</b> ( 56 . 50 ) , as she clocked <b>56.34 seconds</b> , she finished <b>fifth</b> in her heat , again missing out on a place in the <b>second</b> round .	ieva zunda ( born 20 july 1978 in tukums ) is a latvian athlete . her main event is the sprint and hurdles , but she also competes in the 400 and 800 metres . [...] she did not make it past the first round at the 1999 and 2003 world championships . [...] in 2008 [...] shortly before the deadline - on 28 june , she had finally reached the <b>qualifying</b> standard in <b>the 400 m</b> ( 56 . 50 ) , as she clocked <b>in the first round</b> . she finished <b>third</b> in her heat , again missing out on a place in the <b>first</b> round .	ieva zunda ( geboren am 20 . juli 1978 in tukums ) ist eine lettische leichtathletin . ihre hauptdisziplin ist der sprint und der hürdenlauf , sie tritt aber auch über 400 und 800 m an . [...] bei den weltmeisterschaften 1999 und 2003 kam sie nicht über die erste runde hinaus . [...] 2008 versuchte sie erneut [...] kurz vor dem stichtag - am 28 . juni - hatte sie endlich die qualifikationsnorm über 400 m ( 56 . 50 ) erreicht , wie sie in der ersten runde lief . sie wurde dritte in ihrem lauf und verpasste erneut den einzug in die <b>erste</b> runde .
	Word Spans: [153, 153]	Word Spans: [154, 154]

Table 6: Examples of the original text from Wikipedia (Guo et al., 2020), HADES (Liu et al., 2022) as the perturbed version with token-level labels to detect hallucination, and the machine-translated (MT) and post-edited version from our proposed ANHALTEN.



## B Additional Experiments

### B.1 Cross-Lingual Transfer Results of mBERT

	# Instances	Setting	Accuracy $\uparrow$	G-Mean $\uparrow$	BS $\downarrow$	AUC $\uparrow$	Not Hallucination			Hallucination		
							P $\uparrow$	R $\uparrow$	F1 $\uparrow$	P $\uparrow$	R $\uparrow$	F1 $\uparrow$
Zero-Shot	0	offline	61.00	56.12	26.49	69.66	71.04	42.92	50.68	57.89	80.04	66.54
Few-Shot	10	offline	62.84	60.36	25.96	69.53	71.41	46.94	55.74	59.15	79.59	<b>67.54</b>
Few-Shot	100	offline	61.16	55.94	25.52	70.27	<b>73.14</b>	40.98	49.95	58.02	<b>82.42</b>	67.21
Few-Shot	876	offline	62.08	58.61	24.27	70.44	68.46	52.63	56.48	60.80	72.03	64.29
Translate-Train	6344	offline	<b>64.54</b>	<b>63.76</b>	<b>22.61</b>	<b>70.46</b>	66.97	<b>62.61</b>	<b>63.99</b>	<b>63.44</b>	66.57	64.35
Zero-Shot	0	online	60.44	55.34	26.71	67.81	<b>73.47</b>	36.69	48.10	56.33	<b>85.46</b>	<b>67.76</b>
Few-Shot	10	online	60.04	53.77	27.61	68.05	72.94	36.53	46.48	56.46	84.81	67.39
Few-Shot	100	online	60.84	56.55	27.35	67.44	70.99	41.79	50.68	57.32	80.90	66.77
Few-Shot	876	online	61.50	58.07	24.66	68.86	71.52	43.66	52.82	57.91	80.29	66.82
Translate-Train	6344	online	<b>64.78</b>	<b>63.66</b>	<b>22.66</b>	<b>71.11</b>	69.43	<b>57.66</b>	<b>62.05</b>	<b>62.55</b>	72.28	66.45

Table 7: Cross-lingual transfer results of mBERT (%) averaged over 5 runs.

### B.2 Morphological Analysis

In English, grammatical gender is not distinguished, whereas German has three grammatical genders that influence articles, pronouns, and adjectives. Words indicating gender often lie outside the MARKED WORD SPANS used for hallucination detection. Our experiment selects instances where gender-indicating words (articles, possessive pronouns, demonstrative pronouns) precede nouns in the MARKED WORD SPANS from both the German and English datasets. This dataset includes 64 instances per language, with an equal distribution of labels.

Testing with XLMR in the online setting, the goal is to determine if contextual gender information influences hallucination detection results. The additional gender information might help classify non-hallucination instances but could mislead models if the original, correct word has a different gender.

Results in Table 8 show a performance drop in accuracy and G-Mean when gender-indicating words are included in the MARKED WORD SPANS, particularly for English instances. However, AUC improves, suggesting that the extended spans do not hinder the model’s ability to distinguish between classes. The models tend to assign more instances to the hallucination class, reducing the F1 score for the non-hallucination class. This performance drop may result from a lack of such gender-indicating contexts in the fine-tuning dataset, indicating potential issues with handling longer MARKED WORD SPANS.

Language	Preceding	Accuracy $\uparrow$	G-Mean $\uparrow$	BS $\downarrow$	AUC $\uparrow$	Not Hallucination			Hallucination		
						P $\uparrow$	R $\uparrow$	F1 $\uparrow$	P $\uparrow$	R $\uparrow$	F1 $\uparrow$
English	With	76.56	76.13	16.74	85.16	76.73	76.88	76.31	77.94	76.25	76.51
German	With	59.38	46.32	21.73	86.99	85.66	24.37	35.99	55.75	94.37	69.93
English	Without	69.69	68.58	21.67	74.57	75.30	59.38	65.97	66.67	80.00	72.46
German	Without	61.25	49.38	22.34	80.96	86.59	27.50	39.63	57.17	95.00	71.16

Table 8: Results of XLMR (%) in the *online* setting averaged over 5 runs, for instances with MARKED WORD SPANS containing nouns with and without preceding words that indicate the grammatical gender of the noun.