
Cultural Knowledge Gaps in LLMs: Geographic Hallucination Bias Across Latin American Countries¹

Karla Doctor Mauricio 1
Independent

With
Apart Research

Abstract

Large language models (LLMs) are used in education, public services, and information tools across Latin America. But we do not know how often they produce wrong information about Latin American culture, or where they fail the most. This paper evaluates two models, GPT-4o-mini and Claude Haiku 4.5, on 270 questions from CHOCLO, a benchmark of cultural knowledge from 18 Latin American countries. We use semantic similarity scores and a four-category LLM-as-judge classifier (correct, partial, hallucination, abstention) to compare how each model fails. GPT-4o-mini hallucinates at 33.0% overall and almost never abstains (1.5%). Claude Haiku 4.5 hallucinates less (24.1%) but abstains much more (17.8%). Both models show geographic gaps of 33 to 40 percentage points across countries. These gaps remain after controlling for category composition and topic diversity. The public_figure category has the highest hallucination rate (66.7% for GPT), and it is also the least represented category in the benchmark itself. This points to a two-layer problem: models lack knowledge about specific Latin American individuals, and the benchmark has limited coverage of them too. We also document infrastructure barriers that prevented evaluation of LatamGPT.

¹ Research conducted at the [The Global South AI Safety Hackathon](#), June 2026

1. Introduction

LLMs are now used in tools that affect people's daily lives: educational assistants, public service chatbots, and information platforms in Spanish. These tools are deployed across Latin America without strong evidence of whether their cultural knowledge is reliable. This matters for AI safety. A model that says it does not know transfers uncertainty to the user. A model that produces a confident wrong answer hides uncertainty and looks correct. The second failure mode is more dangerous.

Latin American cultural knowledge is an area where dominant LLMs likely have less training coverage. Most large models are trained on data that overrepresents English-language and North American or European sources. This creates a plausible risk that models fail more often, and more silently, on questions about Latin American culture.

CHOCLO (Del Solar et al., 2026) is a benchmark of approximately 104,000 open-ended questions in Spanish, covering cultural knowledge from 18 Latin American countries across 7 categories and 3 difficulty levels. LatamGPT (CENIA et al., 2026) is a 70B parameter model based on Llama 3.1, designed to address the representation gap that CHOCLO documents. No published work has evaluated LatamGPT on CHOCLO, or decomposed LLM errors into qualitative failure modes by country.

Our central question is: do GPT-4o-mini and Claude Haiku 4.5 show geographic gaps in CHOCLO that are explained by hallucination, and are those gaps robust to controls for category composition and topic diversity?

Our main contributions are:

1. An auditing framework that adapts fairness methods from classical ML to open-ended LLM benchmarks
2. Empirical evidence of divergent failure modes and substantial geographic gaps
3. Documentation of two-layer underrepresentation for public figures from Latin America
4. Documentation of infrastructure barriers as a form of access inequality

2. Related Work

Work on LLM hallucination has grown significantly. HaluEval (Li et al., 2023) provides a large-scale benchmark for hallucination evaluation across question answering, dialogue, and summarization tasks. FActScore (Min et al., 2023) evaluates factual precision at the atomic claim level in long-form text generation. Both focus on English-language tasks and do not address geographic or cultural variation in failure rates.

Fairness in machine learning has been studied extensively (Barocas et al., 2020; Mehrabi et al., 2019), but most work focuses on classification tasks with well-defined labels. Applying fairness concepts to open-ended generation with LLM-as-judge classification is less common.

CHOCLO is, to our knowledge, the most comprehensive open-ended benchmark for Latin American cultural knowledge in Spanish. Masakhane (Masakhane Research Foundation) has done similar work for African languages, building community-driven NLP resources and

benchmarks for low-resource languages. Our governance proposals draw directly from their model.

The key gap our work addresses is the combination of: (a) qualitative failure mode decomposition (hallucination vs. abstention) across two models, (b) geographic analysis with fairness controls, and (c) policy-oriented governance proposals grounded in empirical results.

3. Methods

3.1 Benchmark and sampling

We drew a stratified sample of 270 questions from CHOCLO: 15 questions per country across 18 countries. Within each country, we balanced by category and difficulty level when coverage allowed. CHOCLO's natural distribution is not uniform (geography = 43.7%, fauna = 21.1%), so our sample prioritizes statistical power for category comparisons rather than matching the benchmark distribution. As a result, the `public_figure` category was absent in 15 of 18 countries in our sample. This limitation becomes a finding: underrepresentation of Latin American public figures appears in both the benchmark and the models (see Section 3.3).

Brazil and Haiti are absent from CHOCLO and therefore from our evaluation. According to Britannica's definition of Latin America, these two countries are part of the region. Their absence reflects a coverage gap in the benchmark itself, not a design choice of this study.

3.2 Models evaluated

We evaluated two models as response generators: GPT-4o-mini (OpenAI) and `claude-haiku-4-5-20251001` (Anthropic). Both used `temperature = 0.2`, `max_tokens = 256`, and an identical system prompt in Spanish. For classification, we used a LLM-as-judge approach: GPT-4o-mini as primary judge, Claude Haiku 4.5 as backup if GPT fails, and a heuristic fallback (abstention phrase detection plus similarity threshold) if both fail.

We note that the same model that generates responses also acts as primary judge of the other model's responses, and in some cases its own. We declare this as a potential evaluation conflict and discuss it in Section 5.

We attempted to evaluate LatamGPT-SFT-1.0 (CENIA, based on Llama 3.1 70B) through three public inference routes: (1) HuggingFace Inference API, where the model appeared as not deployed by any provider; (2) the `featherless-ai` router, which returned technically corrupt outputs with chat template tokens leaking as plain text; (3) HuggingFace automatic provider selection, which showed the same failure pattern. These failures are consistent with the LatamGPT authors' own warnings about sensitivity to their specific chat template. We report this as an infrastructure access barrier, not a model failure.

On reproducibility: OpenAI exposes a `seed` parameter (in beta) for best-effort determinism. Anthropic does not expose an equivalent. We did not use seed in either model, to keep comparisons symmetric. We measured empirical stability by running the judge three independent times on the same inputs, finding 98.3% agreement across 540 classifications.

3.3 Metrics

Semantic similarity: cosine similarity over paraphrase-multilingual-MiniLM-L12-v2 embeddings, normalized and clipped to [0, 1]. Error = 1 minus similarity. This metric captures surface-level semantic distance. An honest abstention can score similarly to a hallucination, since neither approaches the gold reference textually. This motivates the qualitative classification step.

LLM-as-judge classification: four categories aligned with CHOCLO's taxonomy: correct, partial, hallucination, abstention. Prompt examples were taken directly from CHOCLO's published benchmark card. We detected and corrected a bug where the judge inherited the generation system prompt. After correction, hallucination detection increased by 4.1 percentage points for GPT and 4.8 percentage points for Claude. We report the corrected version throughout.

3.4 Fairness and robustness analysis

Geographic gap: maximum minus minimum hallucination rate by country (range gap), with 95% confidence intervals via bootstrap (1,000 resamples). Confound control: residual between observed performance and expected value under each country's category mix. Significance: Wilcoxon signed-rank test on continuous error; McNemar test on binary hallucination rate (paired design, $n = 270$). Multiple comparisons: Benjamini-Hochberg and Bonferroni corrections. Topic robustness: UMAP projections on embeddings by country to check whether topic diversity explains the gaps. Temperature sensitivity: full re-evaluation at temperature = 0.4 and 0.6.

4. Results

4.1 Global error profiles

Table 1. Response quality distribution ($n = 270$ each). 95% CI via bootstrap (1,000 resamples). Classification validated at 98.3% stability across 3 independent judge runs.

Category	GPT-4o-mini	Claude Haiku 4.5
Correct	22.2% [17.8-27.4%]	23.7% [18.5-29.3%]
Partial	43.3% [37.8-49.6%]	34.4% [28.9-40.0%]
Hallucination	33.0% [26.7-38.9%]	24.1% [19.3-29.3%]
Abstention	1.5% [0.4-3.0%]	17.8% [13.3-22.2%]

GPT-4o-mini hallucinates 22 times more than it abstains. Claude Haiku 4.5 abstains 12 times more than GPT. The difference is significant on continuous error (Wilcoxon $p = 0.019$) and on binary hallucination rate (McNemar chi-squared = 10.78, $p = 0.001$). Confidence intervals do not overlap (GPT: 26.7-38.9%; Claude: 19.3-29.3%). Both models have the same rate of correct answers (~22-24%), which means the difference is not in knowledge but in behavior when uncertain.

4.2 Qualitative examples

Table 2. Representative examples. In two cases, the model placed the answer in a different country than the question referred to, showing hallucination not anchored to the geographic context of the question.

Country	Question	GPT Response	Claude Response
Argentina	¿Qué instrumento tocó Oscar Alem antes de pasarse al piano?	Contrabajo	"tocó el acordeón" (alucinación)
Argentina	¿Qué características geográficas y oceanográficas tiene la bahía Relegada?	Profundidad ~3000m, 158 ha	"ubicada en la costa de Chile..." (alucinación; país incorrecto)
Guatemala	¿Dónde nace el río Las Vacas?	Cerros al sudeste de Ciudad de Guatemala	"nace en la Sierra de la Ventana, Buenos Aires, Argentina " (alucinación; país incorrecto)
Venezuela	¿Dónde habita Andersonoplatus laculata?	Insecto, brillante, negro, patas y antenas amarillentas	"es un pez de la cuenca del Amazonas..." (alucinación elaborada)

4.3 Hallucination by category

Table 3. Hallucination rate by category. public_figure is present in only 3 of 18 countries in our sample (n = 9 total). Its high hallucination rate in GPT (66.7%) coexists with near-total absence from the benchmark for most countries.

Category	GPT-4o-mini	Claude Haiku 4.5	n (sample)
public_figure	66.7%	22.2%	9 (3 countries)
object	38.9%	22.2%	18
tradition	37.5%	12.5%	48
geography	25.5%	31.4%	51
fauna	25.0%	22.9%	48
flora	23.8%	11.9%	42
dish	22.2%	14.8%	54

When the benchmark does include public figure questions, GPT-4o-mini hallucinates at 66.7%, more than double any other category. But the category is also absent in 15 of 18 countries because CHOCLO itself has limited coverage of public figures for many countries. Both layers trace to the same source: the underdocumentation of specific Latin American individuals in dominant digital sources.

4.4 Geographic gaps across all 18 countries


Table 4 shows hallucination rates for all 18 countries in our sample, ordered by GPT-4o-mini rate from highest to lowest. The observed gap is 40.0 percentage points for GPT and 33.3 percentage points for Claude. After controlling for category composition, the residual gap is 41.4 percentage points for GPT and 33.7 percentage points for Claude, meaning the category mix does not explain the geographic variation.

Country	GPT %	GPT 95% CI	Claude %	Claude 95% CI	Risk
Argentina	53.3%	[26.7-80.0]	20.0%	[0.0-40.0]	High
Guatemala	46.7%	[20.0-73.3]	40.0%	[13.3-66.7]	High
Venezuela	46.7%	[20.0-73.3]	40.0%	[13.3-60.0]	High
Honduras	40.0%	[13.3-66.7]	40.0%	[20.0-66.7]	Moderate
Nicaragua	40.0%	[13.3-60.2]	20.0%	[0.0-46.7]	Moderate
Uruguay	40.0%	[19.8-66.7]	20.0%	[0.0-40.0]	Moderate
Mexico	40.0%	[20.0-66.7]	20.0%	[0.0-40.0]	Moderate
Panama	40.0%	[13.3-66.7]	26.7%	[6.7-46.7]	Moderate
Colombia	40.0%	[20.0-66.7]	20.0%	[0.0-40.0]	Moderate
Costa Rica	33.3%	[6.7-60.0]	26.7%	[6.7-46.7]	Moderate
Cuba	33.3%	[13.2-60.0]	26.7%	[6.7-53.3]	Moderate
Paraguay	26.7%	[6.7-46.7]	26.7%	[6.7-53.3]	Moderate
Bolivia	26.7%	[6.7-53.3]	20.0%	[0.0-40.0]	Moderate
Chile	26.7%	[6.7-53.3]	6.7%	[0.0-20.0]	Moderate
Ecuador	20.0%	[0.0-46.7]	20.0%	[0.0-40.0]	Low
El Salvador	13.3%	[0.0-33.3]	20.0%	[0.0-40.0]	Low
Peru	13.3%	[0.0-33.3]	20.0%	[0.0-40.0]	Low
Dominican Rep.	13.3%	[0.0-33.3]	20.0%	[0.0-40.0]	Low

Three robustness controls converge on the same conclusion. First, the residual gap after category composition control is comparable to or larger than the observed gap. The unequal category mix across countries does not explain the geographic differences. Second, Argentina vs. El Salvador in GPT is significant after conservative Bonferroni correction ($p = 0.026$). Chile vs. Venezuela in Claude survives Benjamini-Hochberg ($p = 0.031$) but not Bonferroni ($p = 0.062$); we treat this with caution. Third, Guatemala and Venezuela appear among the three worst countries in both models, which have different architectures and providers. This pattern does not depend on the sample size per cell.

4.5 Topic diversity control via UMAP

Figure 1 UMAP projection of the 270 benchmark questions in embedding space, colored by country. Questions from all 18 countries overlap in topic space, showing that geographic hallucination gaps are not explained by thematic differences across countries.



We evaluated whether topic diversity in the benchmark (dispersion in embedding space) explains the geographic gaps. Questions from all countries mix together in topic space, which rules out the possibility that some countries have thematically harder or more unusual questions. Argentina has the highest topic dispersion (rank 3 of 18) and also the highest hallucination rate (53.3%). Guatemala and Chile have similar topic dispersion (~ 1.55 - 1.57) but opposite hallucination rates (46.7% vs. 26.7%). Topic dispersion does not predict hallucination risk.

4.6 Temperature robustness

Table 5 shows results across three temperature settings. Hallucination rate varies by less than 1.2 percentage points across temperature = 0.2, 0.4, and 0.6 for both models. The documented gaps are not an artifact of any specific sampling configuration.

Model	Temp	Correct	Partial	Hallucination	Abstention
GPT-4o-mini	0.2	22.2%	43.3%	33.0%	1.5%
GPT-4o-mini	0.4	23.7%	40.7%	33.3%	2.2%
GPT-4o-mini	0.6	23.7%	42.6%	32.6%	1.1%
Claude Haiku 4.5	0.2	23.7%	34.4%	24.1%	17.8%
Claude Haiku 4.5	0.4	23.0%	35.2%	25.2%	16.7%
Claude Haiku 4.5	0.6	23.0%	33.7%	25.2%	18.2%

4.7 MAE and informative divergences

Global MAE (GPT: 0.544; Claude: 0.565) correlates partially with hallucination rate but diverges in informative cases. Chile in Claude has high MAE (0.628, among the highest in the sample) but low hallucination (6.7%). Its error concentrates in the partial category (46.7%), not in fabrication. A country with high MAE from imprecision has a qualitatively different risk profile than a country with high MAE from confident hallucination. This divergence is only visible when both metrics are reported together.

5. Discussion and Limitations

5.1 Implications for AI safety

The risk of failure is not uniform across models or geography. GPT-4o-mini produces wrong information with the same apparent fluency as correct information. This is the most dangerous failure mode for deployed systems because it is invisible to users. Claude Haiku 4.5 prefers to declare uncertainty, transferring that uncertainty to the user. For real-impact applications in Argentina, Guatemala, Venezuela, or Honduras, the concrete risk is confident misinformation, not missing answers.

These results also connect to a broader point about benchmark coverage. The hallucination rates we measure are evaluated against a benchmark that itself inherits documentation gaps from its primary data sources. CHOCLO is built largely from Wikipedia, which underrepresents indigenous knowledge, rural knowledge, and oral tradition from Latin America. This means our measurements are a lower bound: models may fail even more on cultural knowledge that the benchmark does not yet cover. Both the model performance gaps and the benchmark coverage gaps point to the same upstream source, which is the unequal representation of Latin American knowledge in dominant digital sources.

5.2 Calibration

We implemented post-hoc group calibration on similarity scores: calibrated prediction = original prediction minus (group mean minus global mean). By algebraic construction, the post-calibration bias reduces to near zero without changing model responses. This is useful in downstream decision systems such as dynamic country-level thresholding in an educational assistant, but it does not change the model's underlying cultural knowledge.

5.3 Infrastructure inequality

The inability to evaluate LatamGPT shows that representation inequality in AI for Latin America operates at more than one level: the behavior of deployed models, and access to the models designed to improve that behavior. Infrastructure barriers, including GPU memory requirements, chat template sensitivity, and the absence of reliable public inference providers, can make proprietary models more accessible in practice than open models designed for low-resource contexts.

Limitations

Sample size: $n = 15$ per country can reliably detect gaps of 15-20 percentage points or more. Bootstrap confidence intervals per cell are wide. Potential evaluation conflict: GPT-4o-mini acts as primary judge of its own responses and Claude's, which could bias classification. CHOCLO inherits Wikipedia biases: underrepresentation of indigenous, rural, and oral tradition knowledge. Global MAE confidence intervals use a 5% approximation rather than bootstrap. Only two models were empirically evaluated.

Future Work

Our results point toward a Representativeness Index (RI): a tool that would let communities and governments quantify how well their cultural knowledge is represented in global AI systems, and use that number to negotiate fairer terms. Formalizing this index, incorporating independent sources like Wikipedia coverage and pretraining corpus data, and validating it with community experts remains a priority for future work.

On the governance side, the technical gap we document needs to be closed in ways that do not repeat historical patterns of extraction. We propose three concrete models: community data cooperatives where local curators retain ownership and veto rights over how their knowledge is used; decentralized public funding pools that require LLM providers to share revenue derived from community data; and non-extractive living benchmarks where each country maintains and updates its own version of CHOCLO under terms that prioritize local use over global aggregation. All three draw from the Masakhane model for low-resource African languages.

6. Conclusion

GPT-4o-mini and Claude Haiku 4.5 show opposite failure profiles. GPT hallucinates with confidence (33.0% overall, reaching 53.3% in Argentina). Claude prefers to abstain (17.8%). Both models show geographic gaps of 33 to 40 percentage points that are robust to category composition control, topic diversity control, and temperature variation. These gaps concentrate in Guatemala and Venezuela across both models. The public_figure category shows two-layer underrepresentation: high hallucination when questions exist, and near-total absence of those questions in the benchmark. Infrastructure barriers for LatamGPT show that inequality does not end at model behavior.

These results measure model performance, but they are evaluated against a benchmark that already reflects the same documentation gaps. Both layers, model performance and benchmark coverage, trace to the same upstream source: unequal representation of Latin American knowledge in dominant digital sources. This is not a criticism of the benchmark. CHOCLO is a valuable and rigorous tool. It is an observation about the scope of the problem and an invitation to address it at the source, through the governance models we propose.

Code and Data

- **Code** **repository:**
<https://github.com/KarlaDoctorMauricio/Cultural-Knowledge-Gaps-in-LLMs-Geographic-Hallucination-Bias-Across-Latin-American-Countries>

- **Data/Datasets:** <https://huggingface.co/datasets/latam-gpt/CHOCLO>

Author Contributions (optional)

Karla Doctor Mauricio led the project, designed the experiments, implemented the evaluation framework, and wrote the paper.

References

1. Barocas, S., Hardt, M., and Narayanan, A. (2020). Fairness and Machine Learning: Limitations and Opportunities. fairmlbook.org.
2. CENIA et al. (2026). Llama-3.1-70B-LatamGPT-SFT-1.0 [Model card]. Hugging Face. <https://huggingface.co/latam-gpt/Llama-3.1-70B-LatamGPT-SFT-1.0>
3. Del Solar, B., Carvallo, A., and Soto, A. (2026). CHOCLO: Latin American Cultural Knowledge Benchmark. Hugging Face. <https://huggingface.co/datasets/latam-gpt/CHOCLO>
4. Li, J. et al. (2023). HaluEval: A Large-Scale Hallucination Evaluation Benchmark for LLMs. EMNLP 2023.
5. Masakhane Research Foundation. MasakhaNER / AfriQA. <https://www.masakhane.io/>
6. Mehrabi, N. et al. (2019). A Survey on Bias and Fairness in Machine Learning. arXiv:1908.09635.
7. Min, S. et al. (2023). FActScore: Fine-grained Atomic Evaluation of Factual Precision in Long Form Text Generation. EMNLP 2023

Appendix A. LLM-judge Prompt (v2 corrected system prompt)

The system prompt was separated from the generation prompt after the methodological correction described in Section 3.3:

You are an evaluator of language model responses about Latin American cultural knowledge. Classify each response into exactly ONE of four categories aligned with the CHOCLO criterion: CORRECT: matches the expected answer in substance. PARTIAL: captures the right domain or entity type but fails in precision or specificity. HALLUCINATION: states concrete facts (names, places, dates) that are incorrect or have no real connection to the question. ABSTENTION: explicitly states it does not have the information, without inventing a concrete fact. Always respond with a single word: CORRECT, PARTIAL, HALLUCINATION, or ABSTENTION.

Prompt examples were taken directly from the illustrative cases published in the CHOCLO benchmark card, ensuring alignment with the authors' original definition of partial correctness.

LLM Usage Statement

We used LLM assistance in two ways. First, Claude Code and Cursor supported the software pipeline implementation, including API clients, metric calculations, and unit tests, under

continuous author review. Second, Claude and Adaption assisted with methodological planning and writing this report. All quantitative results were generated by the software pipeline, not by the writing assistant, and were manually verified by the author through individual example inspection before inclusion. Note that GPT-4o-mini and Claude Haiku 4.5 also appear in this work as the systems under evaluation, a separate and distinct role from their use as development tools.