

Investigating Bias: A Multilingual Pipeline for Generating, Solving, and Evaluating Math Problems with LLMs

Mariam Mahran^{1,*}, Katharina Simbeck^{1,*}

¹HTW Berlin University of Applied Sciences, Treskowallee 8, 10318 Berlin, Germany

Abstract

Large Language Models (LLMs) are increasingly used for educational support, yet their response quality varies depending on the language of interaction. This paper presents an automated multilingual pipeline for generating, solving, and evaluating math problems aligned with the German K-10 curriculum. We generated 628 math exercises and translated them into English, German, and Arabic. Three commercial LLMs (GPT-4o-mini, Gemini 2.5 Flash, and Qwen-plus) were prompted to produce step-by-step solutions in each language. A held-out panel of LLM judges, including Claude 3.5 Haiku, evaluated solution quality using a comparative framework. Results show a consistent gap, with English solutions consistently rated highest, and Arabic often ranked lower. These findings highlight persistent linguistic bias and the need for more equitable multilingual AI systems in education.

Keywords

LLMs, Fairness, Education, Language Bias, Evaluation, Equity

1. Introduction

As Large Language Models (LLMs) become increasingly embedded in educational tools, they are reshaping how students access academic support [1, 2, 3]. In subjects like mathematics, these models are often used to explain concepts, walk through problem-solving steps, and supplement classroom instruction [1]. Nonetheless, most commercial LLMs are trained on overwhelmingly English-centric data, raising concerns about how effectively they perform in other languages [4, 5].

In multilingual educational settings, this imbalance can have serious implications. Students interacting with the same model in different languages may receive responses that vary not only in fluency, but in clarity, accuracy, and pedagogical value [1, 6]. This variation is especially problematic in math education, where precise explanations and terminology are essential. Prior studies have highlighted significant language-related disparities in LLM outputs, particularly in educational contexts [6, 5]. These findings underscore the importance of systematic, scalable evaluation across languages.

In this paper, we present a scalable, automated pipeline for evaluating how language impacts the quality of LLM-generated math solutions. Covering generation, translation, solution, and evaluation, the framework supports consistent multilingual comparison across hundreds of exercises. Using three commercial models (GPT-4o-mini (OpenAI) [7], Gemini-2.5-flash (Google) [8], and Qwen-plus (Alibaba Cloud) [9]) we apply it to English, German, and Arabic to uncover persistent performance disparities in educational contexts. This work aims to drive future research toward reducing linguistic biases in AI education tools that serve learners regardless of language background.

2nd Workshop on Education for Artificial Intelligence (edu4AI 2025, <https://edu4ai.di.unito.it/>), Co-located with ECAI 2025, the 28th European Conference on Artificial Intelligence which will take place on October 25-30, 2025 in Bologna, Italy

*Corresponding author

[†]These authors contributed equally.

✉ mariam.mahran@htw-berlin.de (M. Mahran); katharina.simbeck@htw-berlin.de (K. Simbeck)

🌐 <https://mariamkhhmahran.github.io/connect/> (M. Mahran); <https://iug.htw-berlin.de> (K. Simbeck)

🆔 0009-0003-0568-0172 (M. Mahran); 0000-0001-6792-461X (K. Simbeck)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

2. LLMs in education: Use Cases & Biases

LLMs are increasingly used in education to support both teachers and students. For educators, they assist with lesson planning, content creation, and assessment, helping reduce workload and improve efficiency [1, 2]. For students, they offer real-time, adaptive support that fosters independent learning, clarifies difficult concepts, and guides problem-solving [3].

Nonetheless, AI fairness in education is an increasing concern [10]. Growing evidence shows that LLMs introduce linguistic biases, particularly disadvantaging non-Western and low-resource language contexts [5]. Models often default to Western entities (i.e. names, foods, locations...etc) even when operating in non-Western languages like Arabic [11]. Recent work also suggests that multilingual models often perform internal reasoning in English regardless of the prompt language [4].

These tendencies raise concerns about the inclusivity and cultural relevance of AI-generated educational content. In mathematical reasoning tasks, LLMs have shown tendency to perform well in high-resource languages such as English and Chinese but struggle with languages like Korean due to difficulties processing non-English input [6]. Recent work by [5] further demonstrate this gap by evaluating six LLMs on four educational tasks across six non-English languages. They find performance drops sharply in low-resource languages, correlating with training data representation. English prompts was also shown to outperform their translations across tasks.

Such biases have real consequences. In grading and assessment contexts, models may favor Western argument structures, writing styles, and measurement systems placing non-Western students at a disadvantage. More broadly, the dominance of English-centric models reflects global educational inequality, as advanced AI tools remain most accessible to wealthier, English-speaking regions [1].

3. Dataset Creation

3.1. Dataset Generation & Translation

This study uses a dataset of 628 math exercises spanning five topic areas and grade levels B through H (equivalent to Grades 2-10), based on the official German K-10 mathematics curriculum [12].¹ Exercises were generated using ChatGPT-4o (via the ChatGPT interface) in a manual, iterative process guided by curriculum-aligned topic names and learning objectives. Rather than replicating specific examples, the curriculum served as a conceptual reference to ensure broad topic coverage and grade-level appropriateness. Each item was then manually reviewed for clarity, pedagogical relevance, and alignment with the intended mathematical concept. Unclear, repetitive, or misaligned exercises were revised or removed. The resulting dataset consists of brief, one-line prompts focused on single operations or concepts, minimizing the risk for any contextual or linguistic bias. Table 1 shows example exercises.

All exercises were initially written in English, then translated into German and Arabic using the GPT-4o-mini API. Given their concise and context-independent structure, the exercises were well-suited for machine translation. Nonetheless, all translations were manually audited to ensure semantic and linguistic appropriateness. Minor edits were applied as needed to maintain consistency across languages.

3.2. Solution Generation

Three LLMs were used in this study: *GPT-4o-mini*, *Gemini 2.5 Flash*, and *Qwen-plus*. Each was prompted to generate step-by-step solutions in English, German, and Arabic. The prompts, written in the target language, asked the model to explain how to solve the given problem rather than to provide a final answer. This setup was intended to mirror how K-10 learners might seek guided support or clarification. The resulting solutions were saved in separate CSV files for each language-model combination.

¹Code and data are available at https://github.com/iug-htw/solution_evaluator.

Table 1

Examples of exercises across different topic areas and progress levels used in the study.

Topic Area	Topic	Level	Exercise
Numbers and operations	Understanding and representing numbers	B (2nd grade)	Show 28 using tens and ones.
Sizes and measurement	units of measurement	C (4th grade)	Convert 4,289 mL to liters.
Geometry	measuring and calculating	F (8th grade)	A triangle has sides 7, 24, and 25 cm. Prove it is a right triangle.
Data frequencies and probabilities	Compare and evaluate data	E (7th grade)	Which measure (mean, median, mode) best represents this data: 3, 3, 3, 4, 100?
Functional relationships (patterns and structures)	Linear and other simple functions	H (10th)	Determine the x-intercepts and asymptotes of the logarithmic function $f(x) = \ln(x - 2)$

4. Evaluation

To evaluate solution quality, we adopted a hybrid evaluation framework in which **LLMs acted as judges**. LLM-based evaluation is gaining traction as a scalable alternative to human or code-based assessments [13]. Findings suggest that LLMs can generate assessments comparable to human judgments in domains such as open-ended story generation and adversarial attack evaluations [14]. In educational contexts, LLM judgments have shown strong correlation with faculty evaluations, particularly when using pairwise comparisons and structured rubrics [15]. Other studies have shown that, with appropriate fine-tuning and prompt design, LLMs can effectively automate course evaluations across diverse university settings [16].

For each math exercise, three step-by-step solutions, one in English, German, and Arabic, were presented as alternative responses to the same prompt. These were then assessed and compared by a panel of LLM judges based on their clarity, accuracy, and pedagogical effectiveness.

To ensure impartiality and prevent indirect bias from models evaluating their own outputs, we implemented a **held-out judging strategy**. In each round, the model under evaluation was excluded from the judging panel and replaced with a neutral model, *Claude 3.5 Haiku* [17]. The remaining three judges were selected from *GPT-4o-mini*, *Gemini-2.5-flash*, *Qwen-plus*, and *Claude*.

Prior to evaluation, we identified technical terms relevant to each exercise, as the accurate use and explanation of such terms are crucial for clarity and learner comprehension. To do so, we employed a script that extracted these terms from each exercise using the *GPT-4o-mini* model.

During evaluation, solutions were presented in **randomized order** to each judge to mitigate position bias. Position bias occurs when the order of presentation influences rankings, a phenomenon observed both in LLMs and human decision-making [18, 13].

We followed a **comparative assessment** methodology, which has been shown to outperform direct scoring in terms of alignment with human preferences [19, 20]. Judges were asked to rank the three solutions from 1 (best) to 3 (worst) and provide a short justification to their judgment. Rankings were based on multiple educational criteria: conceptual understanding, clarity of explanation, step-by-step reasoning, correct terminology, accuracy, recognition of common mistakes, pedagogical suitability, generalizability, and grade-level appropriateness.

To identify the best-performing language for each exercise, we applied a **majority voting** scheme. If two or more judges ranked the same solution highest, the language of that solution was recorded as the top performer. If the judges couldn't agree on a single best-performing solution, the result was labeled a tie (TIE), ensuring that ambiguous cases were recorded without forcing a winner.

5. Results

Table 2 summarizes the performance of solutions per language, categorizing them for each LLM based on how frequently they were ranked as the best, mid-tier, or worst. Across all three models, English solutions were most frequently favored. GPT-4o-mini showed the strongest language preference, heavily

Table 2
Comparison of solutions ranking across the three models.

Language	GPT-4o-mini			Gemini-2.5-flash			Qwen-plus		
	Best	Mid	Worst	Best	Mid	Worst	Best	Mid	Worst
English	493	80	19	233	233	138	254	206	152
German	76	325	167	184	180	246	194	195	228
Arabic	38	131	403	206	188	223	174	199	239
TIE	21	92	39	5	27	21	6	28	9

favoring English and consistently rating Arabic lowest. Qwen-plus followed a similar trend, though with less extreme separation. Gemini-2.5-flash produced a more balanced distribution overall; while English still led in top rankings, the difference between German and Arabic was narrower, and German received the most “Worst” labels in that case.

In addition to collecting quantitative rankings, we asked the LLM judges to provide brief justifications for each evaluation (see examples in Table 3). Analysis of over 4,000 justifications revealed consistent emphasis on clarity, structured reasoning, and appropriate use of mathematical terminology. English solutions were most frequently praised as “comprehensive,” “well-structured,” or “clear,” reflecting strong positive sentiment overall. German responses received more mixed feedback, often described as adequate but occasionally critiqued for being “less detailed” or assuming prior knowledge. Arabic solutions had the highest concentration of negative sentiment, with justifications noting a lack of clarity or depth in explanation. Notably, comparative framing (e.g., “slightly better,” “strongest”) was common, confirming that judges made relative, not absolute, assessments.

6. Discussion

The results indicate a clear disparity in performance across languages. English solutions were rated most frequently the best, while Arabic was generally rated the worst, although not uniformly across models. German solutions typically occupied the middle range. GPT-4o-mini showed the most extreme language separation, assigning overwhelmingly high scores to English and low scores to Arabic. Gemini-2.5-flash diverged slightly, assigning the highest number of “Worst” rankings to German rather than Arabic. However, the margin between German and Arabic was relatively small (246 vs. 223), and Gemini’s rankings were more evenly distributed across languages overall.

It is worth noting that, while we did not repeat the full experiment under fixed conditions, the full pipeline was executed multiple times during development with variations in datasets, prompts, and code refinements. Despite fluctuations in exact scores (as expected with stochastic outputs), the relative pattern remained consistent: English was favored, Arabic was disadvantaged, and German remained in between. This recurring trend across settings suggests that the observed disparities are robust rather than incidental.

These gaps were not solely driven by solution correctness. Even when all three solutions were mathematically accurate, models showed clear preferences for responses that included explicit reasoning, pedagogical structure, and generalizability. Analysis of the provided justifications show that the LLMs consistently reward explanations that go beyond getting the answer right to showing *why* the method works. Another key insight is that each language showed strengths in different areas. Arabic was often rewarded for its warmth and accessibility, especially at lower grade levels. German performed well when structure and terminology were key but sometimes assumed prior knowledge. These patterns suggest that language-specific conventions shape both how explanations are presented and how models evaluate their educational value.

Nonetheless, variations in linguistic complexity likely contributed to the observed performance patterns. German has a complex sentence structure with long compound words, flexible word order,

Table 3
Examples of LLM ranking justifications for each language

Positive Feedback	Negative Feedback
[en]: "... well-structured and clearly explains the concept using appropriate terminology and visual analogies."	[en]: "...while correct, lacks some of the engaging elements ...may not connect with younger students"
[de]: "...excels in clarity, structure , and appropriate use of technical terms ...avoids assumptions..."	[de]: "... although correct, assumes prior knowledge of rearranging the density formula..."
[ar]: "...provides a student-friendly approach to solving the problem...making it easy for a 4th grader to follow."	[ar]: "... presents the information in a slightly less structured manner ..."

and grammatical rules. These characteristics can make mathematical expression and phrasing more difficult for LLMs to manage, which may explain its intermediate performance. Arabic with its rich morphology and right-to-left script posed even greater challenges. This morphological richness can compress tokens, producing shorter outputs that feel abrupt next to more elaborated English or German responses. Moreover, LLMs may also lack exposure to Arabic instructional texts with formal, school-style structure. As a result, LLMs may misinterpret Arabic’s concise or direct solutions as less educationally complete, even when correct. These surface-level differences can shape how LLMs perceive clarity and instructional value, even when conceptual accuracy is maintained.

The decision to fully automate the pipeline was driven by the need for scalable, reproducible analysis. Manual evaluation methods, while valuable, limits consistency and scope, making it difficult to compare results at scale. Our automated approach ensures uniform treatment across hundreds of exercises and enables easy replication or extension to new models, languages, or curricula. It also reduces human bias in assessment and supports more systematic multilingual benchmarking.

Finally, the findings carry clear pedagogical implications. Teachers cannot assume equal support across languages, as LLMs may provide stronger guidance in English than other languages. LLMs should therefore be treated as supplementary tools, ideally with teacher oversight or cross-lingual checks. At the same time, the pipeline itself serves as an analytical resource, helping educators pinpoint curriculum areas where linguistic gaps pose the greatest risk, such as topics requiring precise terminology or multi-step reasoning. This allows teachers to adapt instruction accordingly to promote more equitable learning in multilingual classrooms.

7. Conclusion

This study presents a scalable, fully automated pipeline for evaluating LLM-generated math solutions across languages. Applying it to English, German, and Arabic revealed consistent disparities in explanation quality, with English generally favored and Arabic more frequently disadvantaged. These differences were not solely about correctness but reflected variation in reasoning depth and pedagogical clarity.

As LLMs become increasingly integrated into classrooms and tutoring platforms, such disparities risk reinforcing existing educational inequalities. Our findings underscore the importance of more linguistically inclusive development and evaluation practices. To ensure equitable access, future work should prioritize improving performance in underrepresented languages through targeted fine-tuning, diverse training data, and culturally informed prompt design. Expanding this evaluation framework to additional languages and domains can support the development of more accessible and fair educational AI systems worldwide.

Acknowledgments

This paper portrays the work carried out in the context of the KIWI project (16DHBKI071) that is generously funded by the Federal Ministry of Research, Technology and Space (BMFTR).

Declaration on Generative AI

During the preparation of this work, the authors used OpenAI's GPT-4.1 to check grammar and spelling and to enhance the writing style. After using these tools/services, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] L. Yan, L. Sha, L. Zhao, Y. Li, R. Martinez-Maldonado, G. Chen, X. Li, Y. Jin, D. Gašević, Practical and ethical challenges of large language models in education: A systematic scoping review, *British Journal of Educational Technology* 55 (2023) 90–112. doi:10.1111/bjjet.13370.
- [2] W. Gan, Z. Qi, J. Wu, J. C.-W. Lin, Large Language Models in Education: Vision and Opportunities, in: 2023 IEEE International Conference on Big Data (BigData), IEEE Computer Society, Los Alamitos, CA, USA, 2023, pp. 4776–4785. doi:10.1109/BigData59044.2023.10386291.
- [3] Q. Li, L. Fu, W. Zhang, X. Chen, J. Yu, W. Xia, W. Zhang, R. Tang, Y. Yu, Adapting Large Language Models for Education: Foundational Capabilities, Potentials, and Challenges, 2024. arXiv:2401.08664.
- [4] L. Schut, Y. Gal, S. Farquhar, Do multilingual LLMs think in english?, in: ICLR 2025 Workshop on Building Trust in Language Models and Applications, 2025. URL: <https://openreview.net/forum?id=I8BOtOPcOv>.
- [5] V. Gupta, S. P. Chowdhury, V. Zouhar, D. Rooein, M. Sachan, Multilingual performance biases of large language models in education, 2025. arXiv:2504.17720.
- [6] H. Ko, G. Son, D. Choi, Understand, Solve and Translate: Bridging the Multilingual Mathematical Reasoning Gap, 2025. arXiv:2501.02448.
- [7] OpenAI, GPT-4o Mini: Advancing Cost-Efficient Intelligence, 2024. URL: <https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/>.
- [8] Google DeepMind, Gemini 2.5 flash: high-throughput, reasoning-capable large language model, <https://cloud.google.com/vertex-ai/generative-ai/docs/models/gemini/2-5-flash>, 2025. Accessed: 2025-07-23.
- [9] A. Cloud, Qwen-Plus: large language models (LLMs) independently developed by Alibaba Cloud, 2025. URL: <https://www.alibabacloud.com/help/en/model-studio/developer-reference/what-is-qwen-llm>.
- [10] K. Simbeck, They shall be fair, transparent, and robust: auditing learning analytics systems, *AI and Ethics* 4 (2024) 555–571. doi:10.1007/s43681-023-00292-7.
- [11] T. Naous, M. J. Ryan, A. Ritter, W. Xu, Having Beer after Prayer? Measuring Cultural Bias in Large Language Models, in: Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Bangkok, Thailand, 2024, pp. 16366–16393. doi:10.18653/v1/2024.acl-long.862.
- [12] Ministerium für Bildung, Jugend und Sport des Landes Brandenburg und Senatsverwaltung für Bildung, Jugend und Wissenschaft des Landes Berlin, Rahmenlehrplan für die Jahrgangsstufen 1-10: Mathematik, 2015. Amtliche Fassung, gültig für Berlin und Brandenburg.
- [13] L. Zheng, W.-L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. P. Xing, H. Zhang, J. E. Gonzalez, I. Stoica, Judging LLM-as-a-judge with MT-bench and Chatbot Arena, in: Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23, Curran Associates Inc., Red Hook, NY, USA, 2023.
- [14] C.-H. Chiang, H.-y. Lee, Can Large Language Models Be an Alternative to Human Evaluations?, in: A. Rogers, J. Boyd-Graber, N. Okazaki (Eds.), Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Toronto, Canada, 2023, pp. 15607–15631. doi:10.18653/v1/2023.acl-long.870.
- [15] T. Ishida, T. Liu, H. Wang, W. Cheung, Large Language Models as Partners in Student Essay Evaluation, Technical Report, Cornell University, United States, 2024. arXiv:2405.18632.

- [16] B. Yuan, J. Hu, An Exploration of Higher Education Course Evaluation by Large Language Models, 2024. [arXiv:2411.02455](#).
- [17] Anthropic, Claude 3.5 haiku, 2024. URL: <https://www.anthropic.com/news/3-5-models-and-computer-use>.
- [18] S. Tan, S. Zhuang, K. Montgomery, W. Y. Tang, A. Cuadron, C. Wang, R. A. Popa, I. Stoica, JudgeBench: A Benchmark for Evaluating LLM-based Judges, 2024. [arXiv:2410.12784](#).
- [19] Y. Liu, H. Zhou, Z. Guo, E. Shareghi, I. Vulić, A. Korhonen, N. Collier, Aligning with Human Judgement: The Role of Pairwise Preference in Large Language Model Evaluators, in: D. Das, D. Chen, Y. Artizi, A. Fan (Eds.), First Conference on Language Modeling COLM 2024, OpenReview, 2024. URL: <https://openreview.net/group?id=colmweb.org/COLM/2024/Conference#tab-accept>, first Conference on Language Modeling 2024, COLM 2024.
- [20] A. Liusie, P. Manakul, M. J. F. Gales, LLM comparative assessment: Zero-shot NLG evaluation through pairwise comparisons using large language models, in: Y. Graham, M. Purver (Eds.), Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, St. Julian's, Malta, 2024, pp. 139–151. doi:10.18653/v1/2024.eacl-long.8.