

# Geography According to ChatGPT – How Generative AI Represents and Reasons About Geography

Krzysztof JANOWICZ<sup>a</sup>, Gengchen MAI<sup>b</sup>, Rui ZHU<sup>c</sup>, Song GAO<sup>d</sup>,  
Zhangyu WANG<sup>e</sup>, Yingjie HU<sup>f</sup>, and Lauren BENNETT<sup>g</sup>

<sup>a</sup>University of Vienna, Austria

<sup>b</sup>University of Texas at Austin, USA

<sup>c</sup>University of Bristol, UK

<sup>d</sup>University of Wisconsin-Madison, USA

<sup>e</sup>University of Maine, USA

<sup>f</sup>University at Buffalo, USA

<sup>g</sup>Esri, USA

**Abstract.** Understanding how AI will represent and reason about geography should be a key concern for all of us, as the broader public increasingly interacts with spaces and places through these systems. Similarly, in line with the nature of foundation models, our own research often relies on pre-trained models. Hence, *understanding what world AI systems construct* is as important as evaluating their accuracy, including factual recall. To motivate the need for such studies, we provide three illustrative vignettes, i.e., exploratory probes, in the hope that they will spark lively discussions and follow-up work: (1) Do models form strong defaults, and how brittle are model outputs to minute syntactic variations? (2) Can distributional shifts resurface from the composition of individually benign tasks, e.g., when using AI systems to create personas? (3) Do we overlook deeper questions of *understanding* when solely focusing on the ability of systems to recall facts such as geographic principles?

## 1. Introduction and Motivation

“Do we know what AI will know?” is among the most profound and consequential questions we can ask today. Finding an answer may take teams from multiple disciplines and methods that we have not yet developed. In this book, we settle with a more concrete but no less relevant question by asking “How does geography look according to ChatGPT?”. By geography, we mean how the environment varies across regions, be it cities, mountains, or nations, and how humans interact with this environment. While ChatGPT is among the first and most prominent examples, our investigations will span further and concern (generative) foundation models more broadly. Finally, by asking how geography *looks according to* such AI systems, we put our focus on *representation*.

Now, intuitively, the notion of representation could pertain to either the study of meaning in an ontological or epistemological sense, or to questions of allocation, i.e.,

uneven coverage or visibility. The latter case relates to the well-studied notion of *AI biases* [23]. For instance, due to representation or coverage bias, certain geographic regions may be overrepresented in model outputs. In contrast, other regions that are less well captured in the training data may be underrepresented. Consequently, the accuracy of model outputs or factual recall may vary geographically [32,22,19,33,24]. However, for the development and deployment of generative AI (GenAI), both notions need to be considered: (1) As will be shown later, GenAI models create strong *defaults*, that is, issues reminiscent of mode collapse where the same prototypical feature is returned over and over again. To give a concrete example, systems such as ChatGPT tend to disproportionately reply with Japan (or, in specific cases, Brazil or Canada) when asked about countries [21]. (2) At its core, representation is *never entirely epistemically neutral*, at least not from a cognitivist and constructivist point of view, as there are many ways to construct such representations. The way individuals or societies construct concepts such as *Poverty* or *Forest* is a function of their experiences, cultures, age, practices, political affiliations, current context, and so on [4,11,28]. Although not making active decisions or having experience (yet) themselves, generative AI agents will necessarily select one among those representational alternatives [1], and, therefore, their outputs will align better with some parts of society and less with others [12,29]. Even more abstractly, at scale, bias is a key to (human) learning and even survival [6]. However, AI agents are nothing like *artificial humans*; rather, they seem to behave more like collective systems with strongly convergent outputs. Hence, even a minute bias or tendency to favor certain representational alternatives may have significant consequences. Unfortunately, however, these issues and the effects of agentic AI acting as an ecosystem are barely studied and understood, especially not in fields such as geography [30,18], urban planning [5], or even the broader geospatial sciences.

Finally, by focusing on representation, we can go beyond merely studying the ability of AI systems to recall facts or perform (spatial) analysis and instead ask questions such as which regions are described in rich, balanced detail and which are painted with a broad brush. Simply put, we suggest that studying AI *correctness* in isolation is insufficient; what matters equally is how AI represents the world. Whether we like it or not, people and companies *are already using AI on a daily basis* to decide on their vacation destination [34,36], career trajectory, real estate, insurance [26], or urban planning [5,31]. In most everyday discourse, *contextual statements matter beyond a more abstract notion of truth* [16]. Yet, we understand little about the consequences. We already know that AI/ML models have been systematically biased along dimensions such as gender and ethnicity [3]; whether they also distort our thinking about and interaction with geography is not well studied yet and a key issue for AI robustness [8]. Related to robustness is a similar but seemingly less studied notion, namely *brittleness*.

To give a concrete example, consider this simple yet telling experiment: the aforementioned tendency to form defaults, such as Japan for the category *Country*, is highly sensitive to variations in prompting. In our experiments with 200 independent queries against GPT-5.1 at temperature 0.3, the system will reply mostly with Japan (168 times) when prompted with *'Name a country, please.'* However, it will start to favor Canada (104 times) for *'Please name a country.'* Despite both prompts being semantically equivalent. It will take a temperature of 1.0 before the system occasionally replies with a third country, namely Brazil. This experiment is not just an illustration of defaults and their strength but also raises a more important question for the interaction

with artificial agents: Would we not expect intelligent systems to exhibit distributional stability under mere paraphrasing?

Finally, before outlining some of these aforementioned representational issues by providing exploratory vignettes, we would like to argue why work in this book relies mostly on studying outputs of *deployed* generative AI systems instead of relying on technologies known from explainable (Geo)AI [10], which themselves have methodological limitations [35]. To make use of an analogy: Cognitive Science studies *what the mind does and how by observing its interaction with the world*, often via the behavior of individual humans or groups. Neuroscience, in contrast, studies *why the brain functions the way it does* by investigating its inner workings and structures such as activation patterns across different regions of the brain during controlled experiments. Both fields complement one another and are at their best when paired. While current AI systems have neither brains nor minds, by analogy, we can study them from similar perspectives. While explainable AI methods may reveal how and why models (that we can fully access) perform better for certain regions or categories, they cannot reveal the impacts of deployed models on people and places, nor can they be used to understand the interaction between human and artificial agents *in-the-wild*.

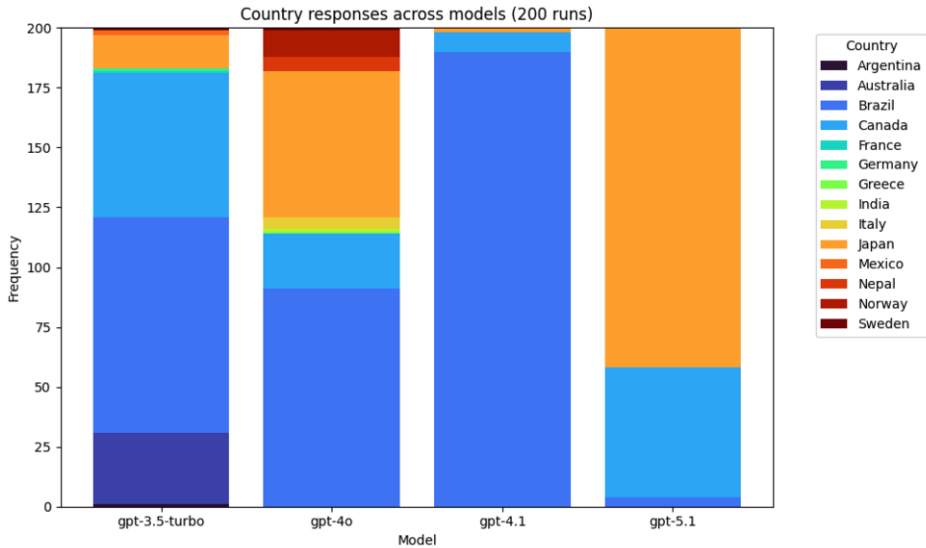
## 2. Vignette I: Default Strength and Brittleness

Before, we introduced the tendency of (generative) foundation models to form *defaults*, that is, strongly favor particular exemplars for a given category. Recently, Liu et al. [20] studied these defaults across about 300 geographic feature types extracted from GeoNames. For instance, among the 11 large language models (LLMs) they considered, 11 would consistently pick the San Diego Zoo when asked to name a zoo, 10 would consistently name the Everglades when prompted for a wetland, and 8 would consistently name mostly Paris as an example of a city. Not only is the geographic footprint of places across all considered feature types biased towards the US and ‘the West’, but the tendency to hold on to these defaults may also be increasing across LLM generations. For example, the older ChatGPT-3.5 seems to produce a more diverse range of outputs than newer models. This is in line with findings from Murthy et al. [25], and Liu et al. [21], and may be due to changes in alignment or a side effect of developers trying to limit hallucinations.

As such, defaults are potentially problematic, e.g., by reducing the geo-diversity of model outputs and thereby causing (mis)representation in downstream tasks such as (over)tourism. Hence, it is worth studying the strength of these defaults. Here we outline an intuitive definition of this phenomenon as an avenue for future studies; see Eq. 1.

$$ds(C) \propto \min \{T > 0 | \exists i \neq c^* : p(i) > \delta\} \quad (1)$$

Simply put, the default strength ( $ds$ ) for a given concept ( $C$ ) is proportional to the minimum model temperature required before we see at least some instances ( $i$ ) of the said concept to be produced as model outputs aside from the default itself ( $c^*$ ), where  $\delta$  controls the required empirical frequency of  $i$ . Returning to our example above, the Zoo default, which predominantly returns the San Diego Zoo, will eventually break as



**Figure 1.** Results for the GPT family for the ‘Name a country, please.’ prompt at a higher temperature of 0.7 (without reasoning for 5.1) to force more varied outputs.

temperature increases, and an increasing fraction of results will be the Schönbrunn Zoo, the Prospect Park Zoo, and so on. The higher this required temperature, the stronger the default. Now, one may argue that such prototypes are simply and exclusively an artifact of training data and that Canada, Everglades, and so forth are merely the most common examples in these datasets. That alone does not readily explain the observed brittleness illustrated above, i.e., why minute changes seem to shift the distribution significantly, nor does it account for differences across models; e.g., GPT-4.1 returns Brazil in 190 of 200 cases, while GPT-5 predominantly favors Japan and outputs Brazil merely 4 times. Figure 1 shows this convergence *within* a model and simultaneous divergence *across* models.

As there is no right or wrong answer here, others have recently rightfully asked, ‘Whose opinions do language models reflect?’ [27]. Clearly, in our collective experience as humans, if we were to ask our friends or colleagues for travel destinations, we would expect a diverse range of responses. Instead, present-day generative AI may present all of us (unknowingly) with the same small set of outcomes, or base what we believe are individualized answers to our prompts on inconsequential syntactic variations in our questions.

### 3. Vignette II: Beyond Surface Bias

A lot of work has gone into debiasing generative AI systems such as text-to-image or large language models [17] as well as their underlying data providers such as knowledge graphs [9,2]. Unfortunately, however, debiasing is not only controversial but may also merely suppresses the deeper underlying issues. First, debiasing itself may lead to misrepresentation as illustrated by Google’s Gemini AI producing images

of black, Asian, and female German Nazi soldiers in a likely over-correction to increase the diversity of model outputs<sup>1</sup>. Second, debiasing itself is not a neutral task but is based on social norms that may vary by country, region, culture, and so on [14]. Among other challenges, this introduces normative issues similar to those faced by pluralistic alignment[29], where contextualization becomes a fine line between contextual sensitivity and political influence or even censorship.

In this vignette, we briefly highlight how difficult *deep debiasing* really is and how unexpected distributional shift may re-enter through the *composition* of individually benign tasks. AI agents such as Gemini or ChatGPT will refuse to answer racially loaded questions or those that may imply other forms of discrimination based on gender or religion. These safeguards are effective in general, but they may not necessarily capture potentially deeper or indirect challenges related misalignment or bias. To illustrate this, we explored a two-stage setup to probe such effects, without implying that they would lead to actual discrimination.


(1) We prompted GPT-4o to create 50 realistic personas representative of the greater Los Angeles area. The AI agents compiled a list of people, including names, occupations, ethnicity/race, residency, and age. After running 8 independent queries, we compared the resulting 381 correctly created hypothetical residents with publicly available demographic statistics. The results varied, e.g., the dataset predominantly listed people aged 25-54, and certain occupations were over-represented.<sup>2</sup> (2) In a second step, we fed each set of roughly 50 people back to (cold-start) GPT-4o agents, stating that the list of people describes residents of the greater Los Angeles area for a future book. Next, we asked the system to identify which characters had a past criminal record, as we wanted the book to be a crime story. In our runs, the systems assigned criminal roles in patterns that differ from, say, the pre-COVID LA arrest statistics, e.g., by selecting less than 6% Whites; see Figure 2.

As this is a highly sensitive topic that deserves a deeper analysis than a vignette permits, it is important to clarify that we are interested in the composite task of creating personas using AI systems and the observed distributional shift; we explicitly do *not* imply that the observed shift in the sample is a sign of racial bias. First, crime, offenders, and arrest statistics have their own (racial) biases. Hence, finding a suitable reference distribution is complex. Second, we did not explicitly account for confounding factors such as age, and we *did not* isolate race as the driving causal factor. Third, the study design, such as the prompt phrasing, may also contribute to the results.

Instead, the key motivation for this vignette is to illustrate how difficult it is to evaluate or control such distributional shifts in generative AI systems, even when explicit safeguards are already in place. As GenAI is already in use across many application areas where artificial personas are created to (implicitly) represent sections of the real population, these issues require follow-up research and guidelines on how best to create such personas. Ideally, symbolic systems such as knowledge graphs may provide reference data for such situations. Closing on a more abstract note, in the future, the cost of validating whether a (Geo)AI model is well-aligned may grow faster than our ability

<sup>1</sup><https://www.theguardian.com/technology/2024/feb/28/google-chief-ai-tools-photo-diversity-offended-users>

<sup>2</sup>For ethnicity/race more specifically, the results were: Hispanic or Latino: 35.43%, White alone non-Hispanic: 26.77%, Black or African American alone: 19.69%, Asian alone: 17.59%, and Other Race alone: 0.52%. Hence, this specific generated sample contains more Black or African American and Whites and fewer Hispanics or Latinos than the demographic statistics for 2020.



**Figure 2.** Criminal record labels in the fictitious book population compared to pre-COVID arrest data from the LAPD. Note, the difference between those distributions as such does *not* imply any racial bias.

to verify the representational consequences of outputs across geographic contexts and downstream tasks.

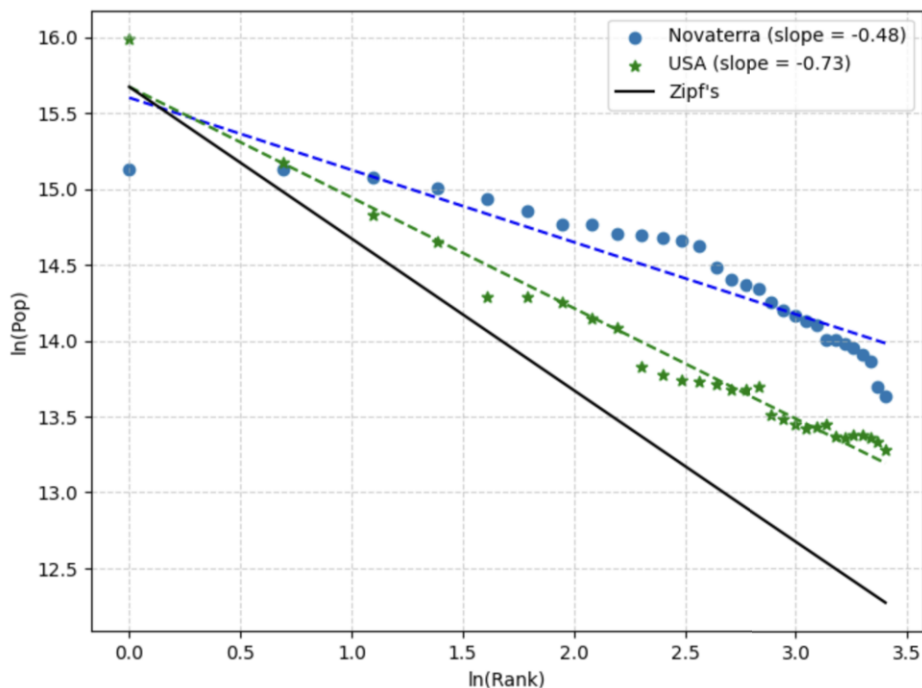
#### 4. Vignette III: Model Knows, Model Shows

Another aspect worth considering in the quest to understand how AI systems represent and reason about geographic spaces and places is to surface the difference between what the model can reproduce and what it fully understands. More concretely, there is a notable difference between producing and explaining concepts *on request* and understanding when and how to apply them independently.<sup>3</sup> Without carefully addressing these differences, we risk confusing the superhuman abilities of AI systems with superhuman intelligence.

To give a concrete example, consider the distribution of city populations. The size distribution of major cities across many nations roughly follows Zipf's law (or similar types of distributions). Now, if we prompt AI systems such as Gemini, DeepSeek, or ChatGPT to confirm whether they are aware of this fact, all will provide rich details, examples, and pointers to the literature. However, if we shift towards understanding what such laws imply and how to apply them consistently, the picture changes.

When asked to imagine a new island nation, say *Novaterra*, at a technical and economic level similar to the US or Japan, a total population of 60 million, and to provide the names and sizes of its 30 largest cities, the models struggled in two notable ways. Across 25 runs (five runs, five models), only two outputs – both from GPT-5 – explicitly referenced theories such as the rank-size rule and produced city sizes approximately in line with expectations. Moreover, none of the systems maintained the population

<sup>3</sup>Now, *understanding* is itself a loaded term; what we mean here is an agent's ability to independently apply knowledge to a context without having been explicitly prompted to do so.



**Figure 3.** Cities ranked by population: US versus the fictitious island nation of Novaterra.

constraint, often exceeding 60 million before reaching the rank of 10. Importantly, and in line with our prior argumentation that representation matters beyond accuracy alone, the point is not whether Novaterra’s population should have followed Zipf’s law – many real countries do not – but whether AI systems can both recall geographic theories and apply them appropriately. While individual nations vary in their size distributions, taken together, they *do* approximate a rank-size rule. In contrast, however, AI-generated city size distributions did not and varied widely; see Figure 3 for a selected example.

To conclude with a broader perspective, this experiment highlights avenues for future work on detecting AI representations, e.g., in geographic misinformation [13]. When Li and Goodchild [7] became interested in studying the quality of Volunteered Geographic Information (VGI), they turned to an unusual source: the map of the imaginary island of Allestone, drawn by the child prodigy Thomas Williams Malkin. They examined the hand-drawn map, asking whether it would adhere to well-known geographic principles such as the fractal dimension of the imaginary coastline, Horton’s law governing stream bifurcation ratios and Central Place Theory. Astonishingly, the map created by the 5-year-old adhered (approximately) to these principles, suggesting that Malkin had developed a latent understanding of core geographic principles. *Would present-day generative AI systems succeed as well?*

## 5. Conclusions and Outlook

In this chapter, we argued that understanding how generative AI systems represent and reason about the world around us requires going beyond traditional perspectives such as accuracy alone. As work on understanding what AI will (or can) know is still in its infancy, a frequent counterargument is that we do not need to study how GenAI systems represent geography or how these systems *latently* reason about spatial and topological relations, because ultimately it is humans that matter. Unfortunately, however, this view falls short, as generative AI systems are already widely deployed across many downstream tasks, including participatory planning, geospatial intelligence, (tourism) recommender systems, and real estate platforms. Increasingly, the broader public learns about the world through the lens of AI systems. In turn, future AI models may train on geographic representations constructed by earlier models, thereby increasing the risk of model collapse. As we have argued, representing geography is never entirely epistemically neutral. Given that present-day AI systems seem to exhibit a collective (herd-like) behavior, at scale, even minute biases can be consequential.

We outlined three illustrative vignettes that showcase directions for future GeoAI work: defaults and their brittleness, distributional shifts, and the potential illusion of understanding. Jointly, they point to a broader research agenda on how generative (Geo)AI systems construct, reproduce, and apply geographic knowledge latently and, in turn, how their outputs shape human perception and cognition of the spaces and places around us. While these issues are not entirely novel, there is no widely used set of *practices* how to interact with and study (Geo)AI agents yet. A key motivation for this book was to document early work on how to investigate deployed AI and their interactions *in-the-wild*. Hence, we leave a deeper empirical assessment for future research. We hope that the thematic chapters and surveys provide a fruitful basis for this.

Finally, arguably, the most substantial recent change in the performance of deployed AI systems is not necessarily driven by advancements in their neural architectures alone, but by their ability to decide when to reach out to (more) symbolic methods, such as various forms of retrieval augmented generation (RAG) or the scripting and execution of code. For instance, 2023/24 models would take inputs such as Well-Known Text (WKT) representations of polygons and reason about their topological relations latently [15]. In contrast, 2025 systems load these polygons into a Python script, use libraries such as Shapely for topological reasoning, and return the results. Strictly speaking, by doing so, these systems have entered the realm of neurosymbolic AI, where many open challenges remain unaddressed. At first glance, such neurosymbolic augmentations may reduce visible errors, e.g., by preventing mistakes in latent topological reasoning or spatial hallucination. However, this may merely obscure unresolved latent representation or reasoning issues that may surface unexpectedly in later downstream tasks.

## References

- [1] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623, 2021.
- [2] Russa Biswas, Lucie-Aimée Kaffee, Michael Cochez, Stefania Dumbrava, Theis E Jendal, Matteo Lissandrini, Vanessa Lopez, Eneldo Loza Mencía, Heiko Paulheim, Harald Sack, et al. Knowledge

- graph embeddings: open challenges and opportunities. *Transactions on Graph Data and Knowledge*, 1(1):4–1, 2023.
- [3] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pages 77–91. PMLR, 2018.
- [4] Helen Couclelis. Ontologies of geographic information. *International Journal of Geographical Information Science*, 24(12):1785–1809, 2010.
- [5] Jiaxin Du, Xinyue Ye, Piotr Jankowski, Thomas W Sanchez, and Gengchen Mai. Artificial intelligence enabled participatory planning: a review. *International Journal of Urban Sciences*, 28(2):183–210, 2024.
- [6] Gerd Gigerenzer. What is bias? and why are we biased? *Proceedings of the Paris Institute for Advanced Study*, 21, 2025.
- [7] Michael F Goodchild and Linna Li. Assuring the quality of volunteered geographic information. *Spatial statistics*, 1:110–120, 2012.
- [8] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8340–8349, 2021.
- [9] Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia d’Amato, Gerard De Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, et al. Knowledge graphs. *ACM Computing Surveys (Csur)*, 54(4):1–37, 2021.
- [10] Chia-Yu Hsu and Wenwen Li. Explainable geoai: can saliency maps help interpret artificial intelligence’s learning process? an empirical study on natural feature detection. *International Journal of Geographical Information Science*, 37(5):963–987, 2023.
- [11] Krzysztof Janowicz. Observation-driven geo-ontology engineering. *Transactions in GIS*, 16(3):351–374, 2012.
- [12] Krzysztof Janowicz, Zilong Liu, Gengchen Mai, Zhangyu Wang, Ivan Majic, Alexandra Fortacz, Grant McKenzie, and Song Gao. Whose truth? pluralistic geo-alignment for (agentic) ai. In *Proceedings of the 33rd ACM International Conference on Advances in Geographic Information Systems*, pages 799–803, 2025.
- [13] Krzysztof Janowicz and Grant McKenzie. How” alternative” are alternative facts? towards measuring statement coherence via spatial analysis. In *Proceedings of the 25th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 1–4, 2017.
- [14] Krzysztof Janowicz, Bo Yan, Blake Regalia, Rui Zhu, and Gengchen Mai. Debiasing knowledge graphs: Why female presidents are not like female popes. In *ISWC (P&D/Industry/BlueSky)*, pages 1–5, 2018.
- [15] Yuhan Ji, Song Gao, Ying Nie, Ivan Majić, and Krzysztof Janowicz. Foundation models for geospatial reasoning: assessing the capabilities of large language models in understanding geometries and topological spatial relations. *International Journal of Geographical Information Science*, 39(9):1866–1903, 2025.
- [16] Werner Kuhn, Tomi Kauppinen, and Krzysztof Janowicz. Linked data-a paradigm shift for geographic information science. In *International Conference on Geographic Information Science*, pages 173–186. Springer, 2014.
- [17] Zichao Lin, Shuyan Guan, Wending Zhang, Huiyan Zhang, Yugang Li, and Huaping Zhang. Towards trustworthy llms: a review on debiasing and dehallucinating in large language models. *Artificial Intelligence Review*, 57(9):243, 2024.
- [18] Pengyuan Liu and Filip Biljecki. A review of spatially-explicit geoai applications in urban geography. *International Journal of Applied Earth Observation and Geoinformation*, 112:102936, 2022.
- [19] Zilong Liu, Krzysztof Janowicz, Ling Cai, Rui Zhu, Gengchen Mai, and Meilin Shi. Geoparsing: Solved or biased? an evaluation of geographic biases in geoparsing. *AGILE: GIScience Series*, 3:9, 2022.
- [20] Zilong Liu, Krzysztof Janowicz, Mina Karimi, Meilin Shi, Ivan Majic, and Alexandra Fortacz. Golden gate bridge, as always? eliciting prototypical places from autoregressive large language models via category production. *Transactions in GIS*, 2026; accepted.
- [21] Zilong Liu, Krzysztof Janowicz, Ivan Majic, Meilin Shi, Alexandra Fortacz, Mina Karimi, Gengchen Mai, and Kitty Carrier. Operationalizing geographic diversity for the evaluation of ai-generated content. *Transactions in GIS*, 29(3):e70057, 2025.
- [22] Rohin Manvi, Samar Khanna, Marshall Burke, David Lobell, and Stefano Ermon. Large language

- models are geographically biased. In *Proceedings of the 41st International Conference on Machine Learning*, pages 34654–34669, 2024.
- [23] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6):1–35, 2021.
- [24] Mazda Moayeri, Elham Tabassi, and Soheil Feizi. Worldbench: Quantifying geographic disparities in llm factual recall. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, pages 1211–1228, 2024.
- [25] Sonia Krishna Murthy, Tomer Ullman, and Jennifer Hu. One fish, two fish, but not the whole sea: Alignment reduces language models’ conceptual diversity. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 11241–11258, 2025.
- [26] Jinmeng Rao, Song Gao, Yuhao Kang, and Qunying Huang. LSTM-TrajGAN: A Deep Learning Approach to Trajectory Privacy Protection. In Krzysztof Janowicz and Judith A. Verstegen, editors, *11th International Conference on Geographic Information Science (GIScience 2021) - Part I*, volume 177 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 12:1–12:17, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [27] Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cino Lee, Percy Liang, and Tatsunori Hashimoto. Whose opinions do language models reflect? In *International conference on machine learning*, pages 29971–30004. PMLR, 2023.
- [28] Simon Scheider and Kai-Florian Richter. Pragmatic geoai: Geographic information as externalized practice. *KI-Künstliche Intelligenz*, 37(1):17–31, 2023.
- [29] Taylor Sorensen, Jared Moore, Jillian Fisher, Mitchell Gordon, Niloofar Miresghallah, Christopher Michael Rytting, Andre Ye, Liwei Jiang, Ximing Lu, Nouha Dziri, et al. A roadmap to pluralistic alignment. *arXiv preprint arXiv:2402.05070*, 2024.
- [30] Nico Van de Weghe, Lars De Sloover, Anthony Cohn, Haosheng Huang, Simon Scheider, Renée Sieber, Sabine Timpf, and Christophe Claramunt. Opportunities and challenges of integrating geographic information science and large language models. *Journal of Spatial Information Science*, (30):93–116, 2025.
- [31] Qingyi Wang, Yuebing Liang, Yunhan Zheng, Kaiyuan Xu, Jinhua Zhao, and Shenhao Wang. Generative ai for urban planning: Synthesizing satellite imagery via diffusion models. *Computers, Environment and Urban Systems*, 122:102339, 2025.
- [32] Zhangyu Wang, Nemin Wu, Qian Cao, Jiangnan Xia, Zeping Liu, Yiqun Xie, Akshay Nambi, Tanuja Ganu, Ni Lao, Ninghao Liu, et al. Geobs: Information-theoretic quantification of geographic bias in ai models. *arXiv preprint arXiv:2509.23482*, 2025.
- [33] Nemin Wu, Qian Cao, Zhangyu Wang, Zeping Liu, Yanlin Qi, Jieliu Zhang, Joshua Ni, Xiaobai Yao, Hongxu Ma, Lan Mu, et al. Torchspatial: A location encoding framework and benchmark for spatial representation learning. *Advances in Neural Information Processing Systems*, 37:81437–81460, 2024.
- [34] Jian Xie, Kai Zhang, Jiangjie Chen, Tinghui Zhu, Renze Lou, Yuandong Tian, Yanghua Xiao, and Yu Su. Travelplanner: A benchmark for real-world planning with language agents. In *International Conference on Machine Learning*, pages 54590–54613. PMLR, 2024.
- [35] Jin Xing and Renee Sieber. The challenges of integrating explainable artificial intelligence into geoai. *Transactions in GIS*, 27(3):626–645, 2023.
- [36] Dazhou Yu, Riyang Bao, Ruiyu Ning, Jinghong Peng, Gengchen Mai, and Liang Zhao. Spatial-rag: Spatial retrieval augmented generation for real-world geospatial reasoning questions. *arXiv preprint arXiv:2502.18470*, 2025.