

REVIEWER'S REPORT

Manuscript No.: JNHST-047

Title: GENERATIVE AI AND THE SOCIAL SCIENCES: RETHINKING METHODOLOGY WITH LARGE LANGUAGE MODELS AND SYNTHETIC POPULATIONS

Recommendation:

Accept after minor revision.....

Rating	Excel.	Good	Fair	Poor
Originality		✓		
Techn. Quality			✓	
Clarity			✓	
Significance	✓			

Reviewer's ID: JPR-Abdul Haseeb Mir

Detailed Reviewer's Report

The manuscript addresses a highly contemporary, disruptive, and paradigm-shifting frontier in academic research: the integration of generative Artificial Intelligence (AI) and Large Language Models (LLMs) into the methodology of the social sciences. The authors correctly frame this technological evolution not as a minor software update or a mere quantitative tool enhancement, but as a fundamental challenge to traditional social research methods. The epistemological significance of this work rests on its exploration of "synthetic populations" and "generative agents" as viable substitutes or supplements for human research subjects. The study explores how computational models can simulate complex social behaviors, generate synthetic survey data, and reconstruct counterfactual historical scenarios that are otherwise inaccessible due to ethical boundaries, high financial costs, or data missingness.

By investigating this computational turn, the paper confronts a foundational question in the philosophy of science: can an algorithmic system trained on historical text corpora truly replicate the subjective, emergent, and context-dependent nature of human agency? The scope of the paper successfully bridges the gap between machine learning engineering and classical sociology, political science, and behavioral psychology. The manuscript provides a timely critique of how traditional statistical assumptions regarding sampling, variance, and probability must be radically adapted when human populations are

REVIEWER'S REPORT

replaced by software-driven simulacra, making it an essential text for contemporary computational social scientists.

Methodologically, the paper functions as an advanced conceptual review and theoretical synthesis, combining deep technical insights into machine learning architectures with rigorous sociological critique. The literature synthesis is remarkably robust, properly drawing from foundational artificial intelligence scholarship, such as the interactive agent frameworks established by Park and colleagues, alongside seminal statistical and causal inference texts like Judea Pearl's work on causality. The authors display a sophisticated awareness of recent 2024 through 2026 preprints and publications tracking advanced techniques in large-scale persona generation, narrative characterization, and the alignment of public opinion in language models.

The paper identifies three major methodological bottlenecks that currently threaten the validity of generative AI in social research: the false sense of probabilistic inference when analyzing LLM outputs, the risk of demographic representativeness and diversity collapse, and the structural trade-off between individual-level prediction accuracy and population-level distribution validity. The synthesis explores how the internal mechanics of language modeling—specifically token prediction algorithms and temperature-driven stochastic sampling—can easily delude researchers into mistaking algorithmic consistency for genuine human statistical variation. This section is handled with a high degree of mathematical and logical precision, establishing a strong critical foundation for the study.

The central thesis of the manuscript posits that while LLMs present unprecedented opportunities to solve persistent data gaps, their uncritical application leads to severe data distortions due to the models' inherent bias toward mainstream cultural narratives and algorithmic homogenization. The author constructs this argument by demonstrating that when an LLM is prompted to simulate a diverse population, it frequently suffers from "diversity collapse," where the generated personas default to highly stereotyped, middle-class, or Western-centric viewpoints, effectively erasing marginalized or low-frequency demographic groups.

To counter this limitation, the manuscript synthesizes several cutting-edge technical interventions. It highlights the deployment of parametric social identity injection, which forces the model to condition its outputs on precise, multi-layered demographic matrices. It also examines the evolutionary generation of social identity, a method that allows synthetic agents to dynamically update their beliefs through simulated social interactions over time. The paper concludes by promoting a pragmatic, dual-layered methodological paradigm where generative AI is not used to completely replace human field research,

REVIEWER'S REPORT

but rather to serve as a high-fidelity simulator for piloting surveys, refining hypotheses, and testing experimental designs before executing expensive real-world data collection.

Recommendations for Manuscript Improvement**Clarifying the Operational Mechanics of Social Identity Injection Techniques**

The primary recommendation to elevate the manuscript's substantive quality centers on the technical description of parametric social identity injection and evolutionary generation. While the text introduces these concepts as viable solutions to diversity collapse, the explanation remains high-level and abstract, preventing researchers from fully understanding how to implement them. The authors must expand this section to provide a detailed narrative walkthrough of the algorithmic syntax and conditioning processes involved in these techniques. They should explain exactly how a system architect structures a prompt payload or fine-tuning matrix to inject multi-layered demographic features—such as intersectional parameters combining socioeconomic status, geographic location, and age—into the model's latent space. Describing how these constraints alter the underlying token probability distribution will provide readers with a clear understanding of how to prevent the model from defaulting to its baseline behavioral biases.

Formalizing the Statistical Critique of LLM Probability Generation

The manuscript's critique regarding the "false sense of probabilistic inference" when using LLM outputs is an exceptionally strong argument that needs to be formalised and expanded. Currently, the paper notes that text generation is stochastic, but it does not sufficiently address how common decoding strategies distort classical social science statistics. The authors should add a dedicated section exploring how variations in temperature settings, top-p sampling, and beam search mechanics affect data generation. They must explicitly demonstrate that running multiple iterations of the same prompt does not create a true independent and identically distributed random sample of a human population. Instead, it generates a bounded probability distribution of text strings dictated by the training corpus. Integrating this formal statistical critique will ground the paper's philosophical warnings in concrete mathematical realities, appealing directly to quantitative methodologists.

Expanding the Ethical and Legal Discussion on Synthetic Data Deployment

While the paper briefly mentions privacy benefits as a reason to use synthetic populations, it overlooks the complex ethical and legal questions surrounding the deployment of generative agents in social planning. The authors should incorporate a comprehensive ethical critique addressing the dangers of algorithmic bias being institutionalized through synthetic research. If policymakers rely on synthetic populations to simulate the impact of city zoning laws or public health interventions, and the underlying LLM suffers from invisible demographic collapse, the resulting policies could severely harm real-world

REVIEWER'S REPORT

marginalized communities. The manuscript should evaluate these systemic risks, discussing the legal liability of replacing human research cohorts with proprietary, closed-source corporate models whose training data and alignment algorithms remain hidden behind corporate trade secret barriers.

Integrating Empirical Performance Baselines and Benchmark Comparative Data

The manuscript would be significantly improved by transforming its theoretical synthesis into a comparative framework supported by empirical baseline data. The authors should introduce a narrative comparison detailing how specific synthetic populations have performed when evaluated against real-world benchmark datasets, such as the General Social Survey or the World Values Survey. The text should explicitly analyze recent validation studies, highlighting where generative agents successfully matched human responses (such as basic consumer purchasing choices) and where they completely failed to capture human nuance (such as high-stakes political voting behavior or complex emotional responses during economic crises). Providing this specific comparative context will help readers understand the practical boundaries of the technology, transforming the manuscript from a conceptual manifesto into a grounded methodological guide.

Standardizing and Completing the Bibliographic Metadata Records

A thorough typographical audit of the concluding references reveals several incomplete citation records and formatting inconsistencies, especially among the 2025 and 2026 preprints tracking large language model developments. Multiple entries are missing essential publication metadata, such as missing journal volume numbers, issue designations, specific digital object identifiers, or definitive page ranges. The authors must systematically review the entire bibliography to complete these missing details and ensure that every single entry conforms perfectly to a singular, standardized academic style manual, such as the APA seventh edition or the Chicago Manual of Style. This uniformity is a mandatory requirement for the paper to be indexed and cited properly across computer science and social science databases alike.

Eradicating Layout Anomalies and Digital Processing Remnants

Finally, the review copy contains minor formatting issues, uneven line spacing around major section headings, and occasional paragraph alignment shifts that detract from its academic presentation. There are a few instances of missing spaces between text blocks and citation brackets, alongside minor typographical errors within the abstract's introductory lines. The authors must manually proofread the entire document to eliminate these mechanical processing remnants. Ensuring that the manuscript possesses a flawless visual and typographic presentation will provide a highly polished, professional appearance that matches the sophisticated nature of its interdisciplinary analysis.

Editorial Recommendation

REVIEWER'S REPORT

This manuscript is recommended for **publication with minor revisions**. The authors have delivered a highly original, vital, and sophisticated piece of scholarship that addresses a critical and immediate need at the intersection of machine learning and the behavioral sciences. By providing a structured, balanced, and deeply critical evaluation of large language models and synthetic populations, this work sets up an important protective framework against the uncritical adoption of generative AI in academic research.

The conceptual reasoning is elegant, the identification of methodological bottlenecks is precise, and the proposed technical solutions are highly innovative. Once the authors expand the technical explanations of social identity injection, formalize the statistical critique of token decoding, and complete the missing bibliographic metadata, this manuscript will be a benchmark, highly cited addition to the journal's computational methodology portfolio.