Axias Experiment 1: Synthetic Data Generation Using AI-Generated Personas to Replicate Human Personality

Soham Bafana

Abstract

This study explores the effectiveness of using AI-generated personas to replicate the latent structures found in human responses to the Big Five Personality Test. By comparing factor analysis results from 20,000 human samples and 200 AI-generated responses, we assess whether synthetic data generated through artificial intelligence can accurately model the underlying dimensions of human personality. Our findings reveal a high degree of similarity between the human and AI-generated datasets across multiple statistical measures, indicating that AI-generated synthetic data can effectively simulate human personality structures. The implications of this research are significant for survey methodology, psychological research, and fields where data privacy and collection costs are concerns. We also discuss the benefits, challenges, and ethical considerations of using AI personas in synthetic data generation.

1 Introduction

The rise of artificial intelligence (AI) has introduced new possibilities for data generation, particularly in research fields where obtaining high-quality data can be challenging due to cost, time, or privacy concerns (Wilbanks et al., 2023). Synthetic data, created through AI-driven techniques, offers a groundbreaking way to simulate real-world datasets. In survey research, synthetic data generated by AI personas can mimic human responses, facilitating large-scale analysis without the need for human participants (Ziems et al., 2023).

Understanding human personality structures is crucial for various applications in psychology, marketing, and social sciences. The Big Five Personality Test is a widely recognized model for assessing personality traits, encompassing Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism (Jiang and Liu, 2023). This study investigates whether AI-generated personas can replicate the latent structures found in human responses to this test.

1.1 Purpose and Significance of the Study

The primary aim of this research is to determine the effectiveness of AI-generated personas in replicating human personality structures as measured by the Big Five Personality Test. By demonstrating that AI can accurately model these structures, we provide evidence for the viability of using synthetic data in psychological research and survey methodologies. This has significant implications for fields where data collection is resource-intensive or sensitive due to privacy concerns.

2 Literature Review

2.1 Synthetic Data Generation in Survey Research

Synthetic data refers to artificially generated data created by machine learning algorithms or AI models that mimic the structure, relationships, and complexities of real data while avoiding privacy risks (Li et al., 2024). In survey research, synthetic data is generated by programming AI models with demographic and psychographic attributes to represent human respondents, enabling AI personas to answer survey questions in ways that align with specific personas (Safdari et al., 2023).

2.1.1 Methods of Generating Synthetic Data Using AI Personas

Several techniques are employed to generate synthetic data using AI personas:

- **Persona-Driven Synthesis**: AI personas are created based on demographic data such as age, occupation, and cultural background. These personas respond to surveys as if they were real individuals, allowing researchers to simulate diverse respondent pools (Jiang and Liu, 2023).
- **Few-Shot Prompting**: Providing the AI model with a few examples of how a persona might respond to specific questions enables the generation of numerous synthetic responses that align with the persona's profile (Li et al., 2024).
- **Synthetic Users for Psychological Research**: Recent studies have used large language models (LLMs) to generate synthetic users that simulate human-like personality traits and behaviors, useful for testing psychological hypotheses and validating personality assessments (Safdari et al., 2023).

2.2 Benefits and Challenges of Synthetic Data in Survey Research

2.2.1 Benefits

- **Scalability and Speed**: AI can generate large volumes of data quickly, enabling rapid experimentation and iteration (Ziems et al., 2023).
- **Cost-Effectiveness**: Reduces costs associated with recruiting and compensating human participants, and minimizes logistical issues related to survey distribution (Li et al., 2024).
- **Consistency**: AI-generated personas remain consistent in their responses over time, avoiding fatigue or variability often observed in human participants (Safdari et al., 2023).

2.2.2 Challenges

- **Bias in AI Models**: AI models trained on existing datasets may include historical biases, influencing the responses generated by AI personas (Deshpande et al., 2023).
- **Validity and Representativeness**: Synthetic data may not capture the full complexity of human emotions, behaviors, or cultural nuances, limiting the applicability of results (Wang et al., 2024).
- **Ethical and Regulatory Concerns**: The use of synthetic data raises ethical questions, particularly regarding its use in sensitive research areas. Regulatory bodies are still developing guidelines for overseeing AI-generated data (Harding et al., 2023).

3 Methodology

3.1 Overview

To assess whether AI-generated personas can replicate the latent structures found in human personality data, we conducted a comparative study using responses from 20,000 human participants and 200 AI-generated personas on the Big Five Personality Test.

3.2 Human Data Collection

We collected responses from 20,000 individuals using the 50-item Big Five Personality Test as described by Jiang and Liu (2023). All Human responses were obtained through https://openpsychometrics.org/.

3.3 AI-Generated Data Creation

3.3.1 Persona Generation Process

We generated 200 AI personas using Language Models (LLM). The generation process involved:

- **Random Demographic Assignment**: Assigning demographics such as race, age, gender, dominant hand, and country to each persona to mirror the diversity of the human dataset.
- **Personality Trait Assignment**: Each persona was assigned a mix of positive and negative traits (e.g., compassionate, stubborn) selected from established psychological trait lists.
- **Persona Summaries**: The LLM was prompted to create detailed summaries for each persona, focusing on how the assigned traits manifest within the context of the Big Five dimensions.
- **Response Generation**: The LLM provided Likert scale ratings (1 to 5) for each of the 50 items on the Big Five Personality Test, ensuring alignment with the persona's traits and demographics.

3.3.2 Sample Prompts Used

Table 1: Persona Generation Prompt

Prompt Content
You are an AI assistant tasked with generating detailed persona summaries for the Big Five Personality Test. Create a diverse range of personas with varying demographic information and personality traits. **Demographic Information**: Race: [Race] Age: [Age] Native English Speaker: [Yes/No] Gender: [Gender] Dominant Hand: [Hand] Country: [Country] **Personality Traits**: [List of Traits] Create a detailed summary of the persona's characteristics related to the Big Five personality traits. Incorporate the given random traits into your summary, adding depth and nuance. Ensure that the persona is unique and well-rounded, reflecting how these traits manifest in different aspects of the person's life. **Instruction**: Based on the persona summary, provide likely responses to each question of the Big Five Personality Test. Use a scale from 1 (Strongly Disagree) to 5 (Strongly Agree).

3.4 Data Processing

We cleaned the data to ensure all responses were valid and correctly mapped to the corresponding questions. The data was formatted into a standardized structure suitable for statistical analysis using Python libraries such as `pandas` and `factor_analyzer`..

3.5 Statistical Analysis

3.5.1 Exploratory Factor Analysis

We conducted Exploratory Factor Analysis (EFA) with varimax rotation on both datasets, extracting five factors corresponding to the Big Five personality traits. Eigenvalues greater than 1 and scree plot analyses confirmed the number of factors (Cao, 2023).

3.5.2 Factor Alignment and Comparison

To compare the factor structures, we employed the Hungarian algorithm to find the best correspondence between factors. We calculated statistical measures such as Tucker's Phi coefficient and Procrustes Similarity Index to assess congruence between the datasets.

4 Results

4.1 Factor Structures

The EFA results showed that the factor loadings for both human and AI-generated data were highly similar. Figures 2 and 2 display the heatmaps of factor loadings, indicating that AI-generated personas effectively replicated the latent structures found in human personality responses.

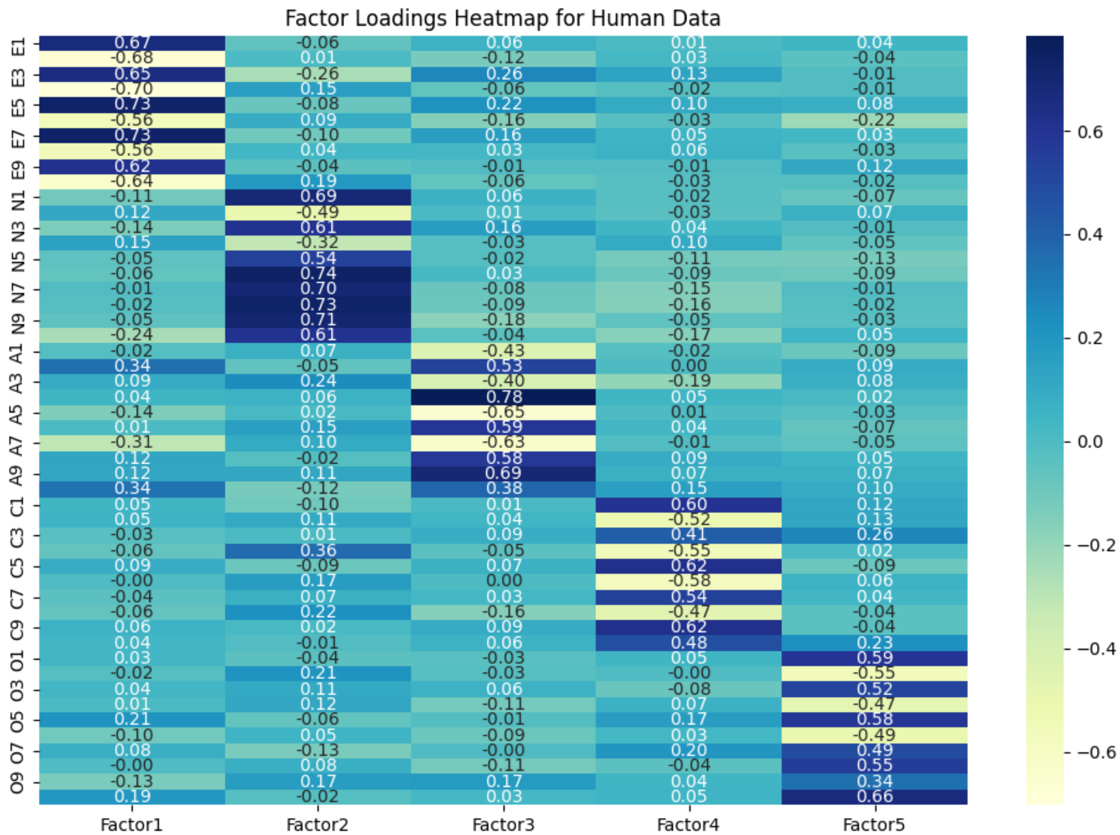

Figure 1: Factor loadings heatmap for human data.

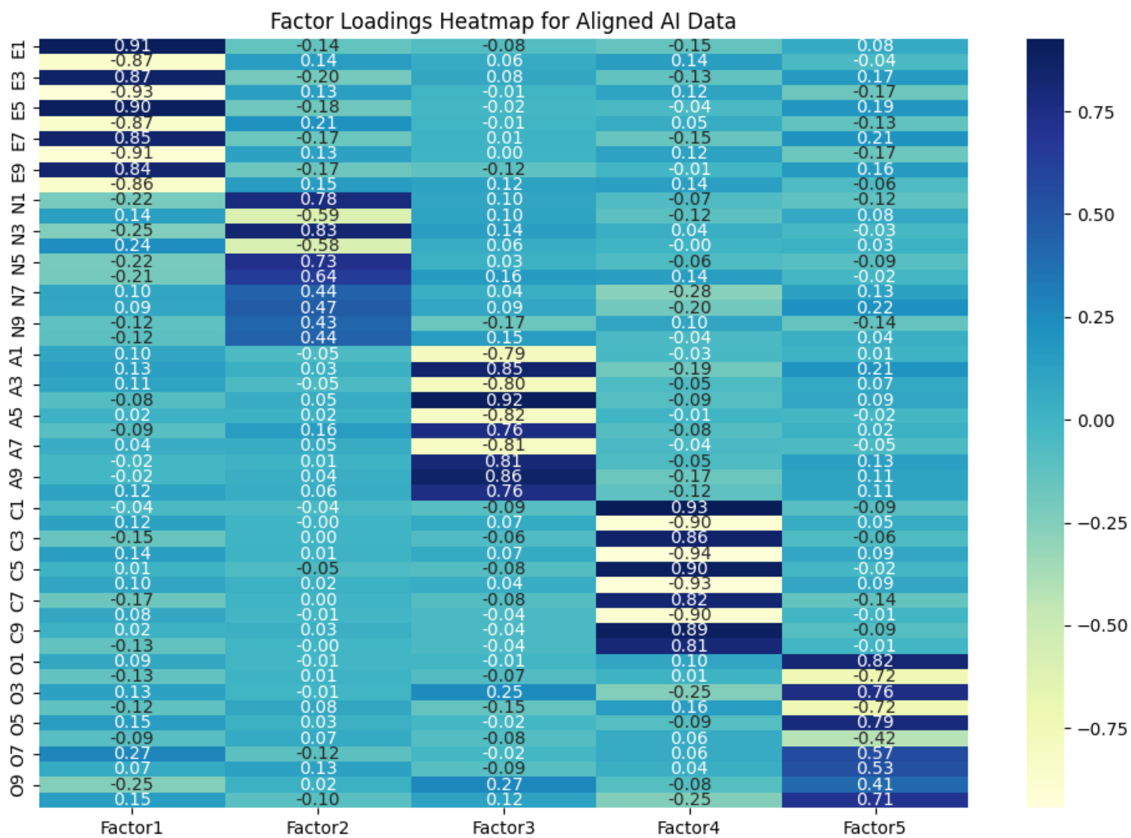

Figure 2: Factor loadings heatmap for AI-generated data.

4.2 Statistical Measures of Similarity

4.2.1 Structural Similarity Index (SSI)

The SSI between the two datasets was calculated to be 0.8721, indicating strong structural similarity.

4.2.2 Tucker's Phi Coefficients

Tucker's Phi coefficients for each factor were:

- Factor 1: 0.9399
- Factor 2: 0.8889
- Factor 3: 0.9305

- Factor 4: 0.9183
- Factor 5: 0.9220

4.2.3 Procrustes Similarity Index

The Procrustes Similarity Index was found to be 0.8359, supporting the congruence between the factor structures.

4.3 Variance Explained by Factors

4.3.1 Human Data

The variance explained by each factor in the human data was:

- Factor 1: 9.97%
- Factor 2: 9.22%
- Factor 3: 7.52%
- Factor 4: 6.54%
- Factor 5: 6.34%

4.3.2 AI-Generated Data

In the AI-generated data:

- Factor 1: 17.02%
- Factor 2: 17.01%
- Factor 3: 14.16%
- Factor 4: 9.75%
- Factor 5: 8.20%

5 Discussion

5.1 Replication of Human Personality Structures

Our findings indicate that AI-generated personas can effectively replicate the latent structures found in human personality data. The high degree of similarity across multiple statistical measures suggests that synthetic data generated through AI can model complex human traits (Safdari et al., 2023).

5.2 Significance of the Results

The results of this study are significant because they demonstrate that AI-generated personas can effectively replicate the latent factor structures found in human personality data. The high similarity across multiple statistical measures—such as the Structural Similarity Index (SSI) of 0.8721, Tucker's Phi coefficients ranging from 0.8889 to 0.9399, and a Procrustes Similarity Index of 0.8359—indicates a strong congruence between the AI-generated data and actual human responses. This means that the AI is not just producing random or superficial data; it is capturing the underlying patterns and relationships that exist in human personality traits.

5.2.1 On-the-Fly Factor Analysis Using Specific Keywords

These findings suggest that it is possible to perform factor analysis dynamically by specifying certain keywords in an AI prompt. Since the AI-generated data closely mirrors human data in terms of latent structures, using specific keywords can guide the AI to produce data that inherently contains these latent factors. This capability allows researchers to explore and test for latent values without the need for extensive data collection from human participants.

5.2.2 Intrinsic Relationship Between Keywords and Latent Values

The use of specific keywords in AI prompts has an intrinsic relationship with the latent values that are naturally uncovered through statistical analysis. Keywords serve as proxies or signals for certain personality traits or factors. When these keywords are input into the AI model, they influence the generation of data in a way that aligns with the latent structures associated with those traits. This means that the AI is not just responding to the keywords at a surface level but is generating data that reflects deeper, underlying psychological constructs.

5.2.3 Contextualization in Broader Applications

The ability to test for latent values using AI-generated data has far-reaching implications across various fields:

- **Psychology and Social Sciences:** Researchers can simulate personality assessments and explore theoretical constructs without the constraints of recruiting participants. This accelerates hypothesis testing and theory development.

- **Market Research and Consumer Behavior:** Companies can generate synthetic personas that reflect target customer segments, allowing for more tailored marketing strategies and product development without infringing on privacy.

- **Human-Computer Interaction:** Designers can create more personalized user experiences by understanding the latent traits that influence user behavior, improving engagement and satisfaction.

- **Educational Tools:** Educators can develop adaptive learning systems that respond to the latent abilities and preferences of students, enhancing learning outcomes.

5.2.4 Testing for Latent Values

By leveraging AI's ability to replicate human-like data structures, we can now test for latent values in a controlled, efficient manner. This approach reduces the reliance on large datasets and allows for rapid experimentation. It also opens up possibilities for real-time analysis and decision-making based on the inferred latent traits from specific keywords or inputs.

5.3 Conclusion

The study's results matter because they highlight a novel way of exploring and understanding complex human traits through AI. By specifying particular keywords in prompts, we can generate data that not only reflects those keywords but also uncovers the latent values associated with them. This intrinsic relationship allows for on-the-fly factor analysis, making it possible to test and explore latent structures in various contexts. The ability to harness AI in this way has the potential to transform research methodologies and applications across multiple disciplines.

6 Future Work

Future research should focus on:

- **Expanding the AI-Generated Dataset**: Increasing the number of AI personas to enhance statistical power and robustness.

- **Bias Mitigation**: Developing techniques to identify and reduce biases in AI-generated data.

- **Validity Across Diverse Contexts**: Testing the effectiveness of AI-generated personas in replicating human responses across different cultures and contexts.

By addressing these areas, future studies can build on our findings and further integrate AI-generated synthetic data into research methodologies.

A Additional Figures

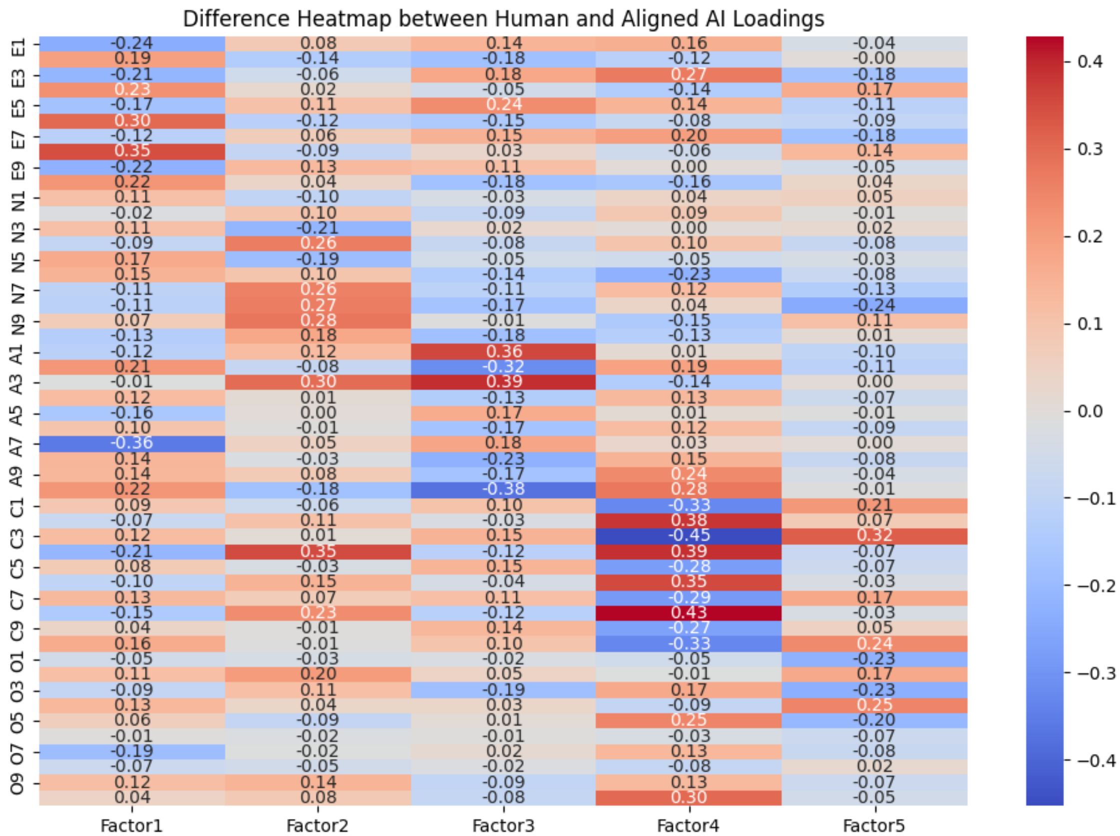

Figure 3: Difference heatmap between human and AI factor loadings.

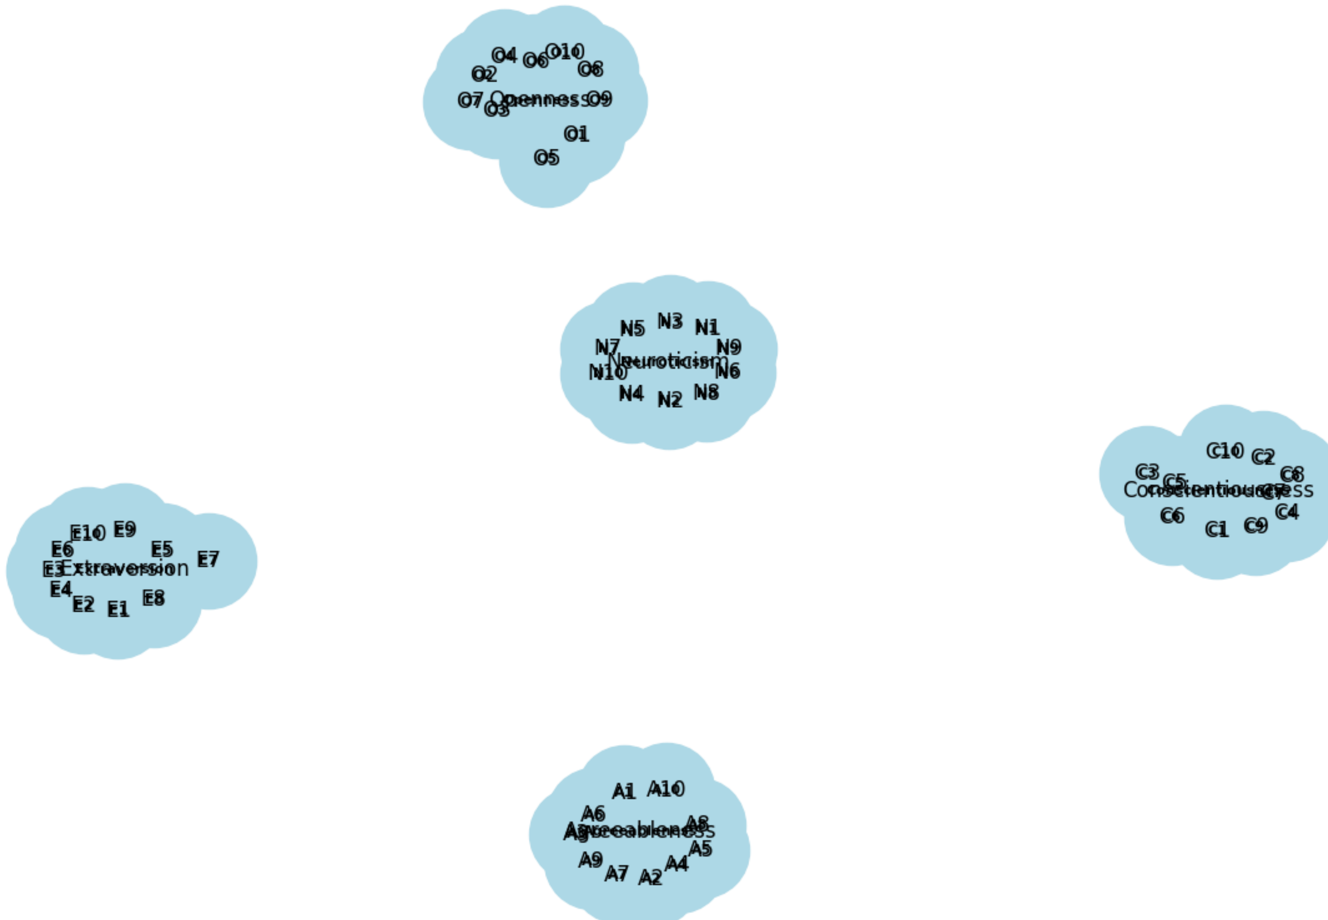

Figure 4: Confirmatory Factor Analysis (CFA) model structure used for alignment between human and AI factors.

B References

References

Cao, X. (2023). The application of structural equation model in psychological research. *CNS Spectrums*, 28(S1):S17–S19.

Deshpande, A., Murahari, V., Rajpurohit, T., et al. (2023). Toxicity in chatgpt: Analyzing persona-assigned language models. *arXiv preprint arXiv:2304.05335*.

Harding, J. L., D'Alessandro, W., Laskowski, N. G., and Long, R. (2023). Ai language models cannot replace human research participants. *AI & Society*.

Jiang, D. and Liu, Y. (2023). Personallm: Investigating the ability of large language models to express big five personality traits. *arXiv preprint arXiv:2305.02547*.

Li, Y., Huang, Y., Wang, H., et al. (2024). Quantifying ai psychology: A psychometrics benchmark for large language models. *arXiv preprint arXiv:2406.17675*.

Safdari, M., Serapio-García, G., Crepy, C., et al. (2023). Personality traits in large language models. *arXiv preprint arXiv:2307.00184*.

Wang, A., Morgenstern, J., and Dickerson, J. P. (2024). Large language models cannot replace human participants because they cannot portray identity groups. *arXiv preprint arXiv:2402.01908*.

Wilbanks, D., Tandon, N., Gu, Y., and Gray, K. (2023). Can ai language models replace human participants? *Trends in Cognitive Sciences*, 27(7):597–600.

Ziems, C., Held, W., Shaikh, O., et al. (2023). Can large language models transform computational social science? *arXiv preprint arXiv:2305.03514*.