



Psybernetica

- GENERATIVE PERSUASION -

Generative Persuasion (2024–ongoing) is an interactive installation that invites viewers to generate tailored propaganda and conspiracy theories on a military-style portable computing center using a Large Language Model (LLM) and personality-based microtargeting. The LLM's behavior has been modified to generate polarizing and misleading content based on the characteristics a user chooses to target. The program is capable of producing social media posts, political advertisements, and conspiracy theories. Participants can experience how generative AI can be weaponized to quickly produce convincing content that can be distributed online as part of disinformation campaigns. The project is based on numerous reports that have recently surfaced about the increasing use of LLMs in disinformation campaigns. The technology, strategies, and information presented are all grounded in research into the current use and limitations of LLMs as well as past and present applications of psychometrics by marketing, military, intelligence and political consulting agencies. The project demonstrates how easy and effective LLMs are at generating inauthentic content designed to influence politics or polarize the public and encourages critical thinking about online digital content.

Derek Curry

Northeastern University, Boston, USA
d.curry@northeastern.edu

Jennifer Gradecki

Northeastern University, Boston, USA
j.gradecki@northeastern.edu

Keywords Disinformation, Conspiracy Theories, Generative AI, Large Language Models, Artistic Research, Tactical Media, Microtargeting, Psychometrics

DOI [10.34626/2025_xcoax_026](https://doi.org/10.34626/2025_xcoax_026)

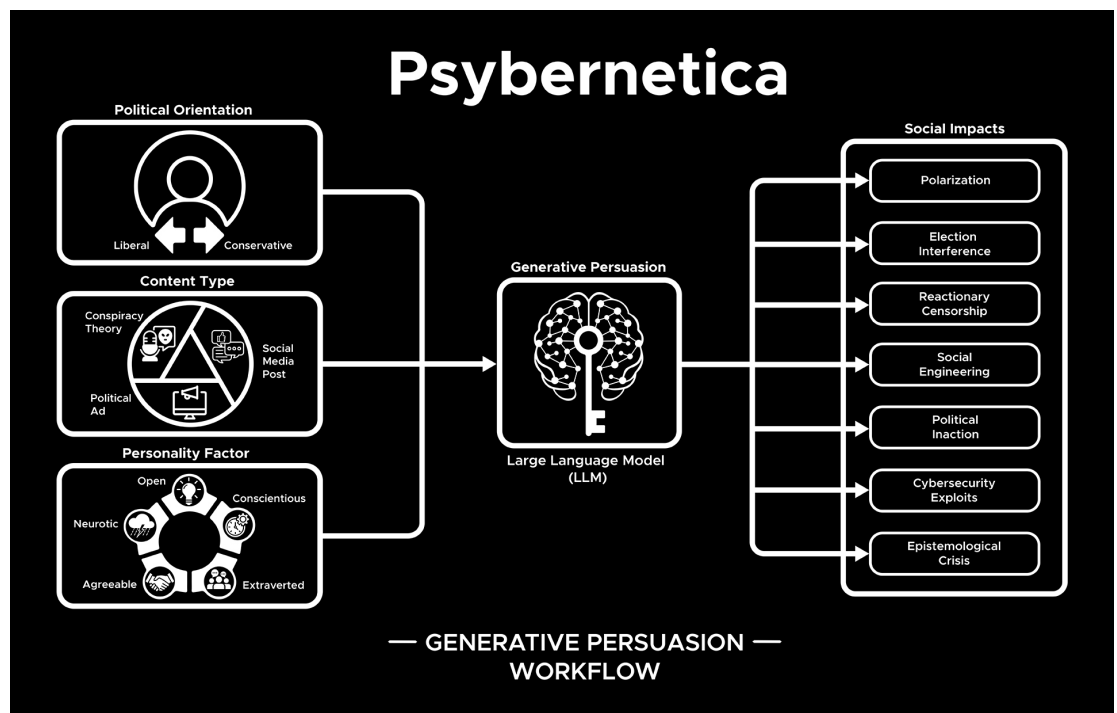


Fig. 1. *Generative Persuasion* mobile computing center.

Description

Generative Persuasion is an interactive installation that leverages a Large Language Model (LLM) and personality-based microtargeting techniques to enable participants to generate personalized propaganda and conspiracy theories tailored to specific personalities and political orientations on a portable military-style computing system in real time. Under the guise of a fictitious startup, the *Generative Persuasion* application uses an uncensored LLM that has been conditioned and pre-prompted to generate polarizing and misleading content based on the characteristics a user chooses to target. The program produces text for social media posts, political advertisements, and conspiracy theories that target people based on the five-factor model of personality and whether the target is liberal or conservative.

The project helps participants understand how generative AI can be used to quickly produce convincing content that can be distributed online as part of disinformation campaigns, and fosters critical thinking about digital content that is encountered and distributed online. The project reveals the dangers of this technology without waiting for a whistleblower or evidence of a scandal; numerous reports have already surfaced about the increasing use of LLMs in disinformation campaigns. While the project is a corporate fiction and the conspiracy theories and propaganda are AI-generated simulations, the technology, strategies, and information presented are entirely real and grounded in research into the current use and limitations of LLMs as well as past and present applications of psychometrics by marketing, military, intelligence and political consulting agencies.



Psychometric Targeting (Cambridge Analytica)

Microtargeting using the five-factor model of personality – sometimes referred to as psychographic profiling – has been used as a marketing tool for decades, but has become more effective with the explosion of behavioral data available through social media and other online sources. In 2018, it was revealed that the British consulting firm Cambridge Analytica had used a Facebook app called “This Is Your Digital Life” to harvest personal data from more than 87 million Facebook users (Kang 2018). The app was really a veiled five-factor personality quiz that was used to find correlations between the users’ responses and their online behavior. These insights were then applied to Facebook users who did not take the quiz or agree to share their personal data. While technically a violation of its terms of service, Facebook’s API allowed the app to collect data from users who did not take the quiz themselves, but were friends with users who had consented. Cambridge Analytica is credited with the success of the 2016 Trump presidential campaign and the Leave.EU (Brexit) campaign. In the wake of the scandal, the SLC Group, Cambridge Analytica’s parent company, filed for insolvency as former employees formed the political consulting firm Emerdata Limited, which has been accused of using the same practices. Our fictional startup company is a thinly veiled reference to the legacy of Cambridge Analytica.

LLM Technology

Large Language Models are a type of deep learning model that use a transformer-based architecture and are capable of performing a variety of tasks, such as answering questions, summarizing documents,

Fig. 2. *Generative Persuasion workflow diagram.*

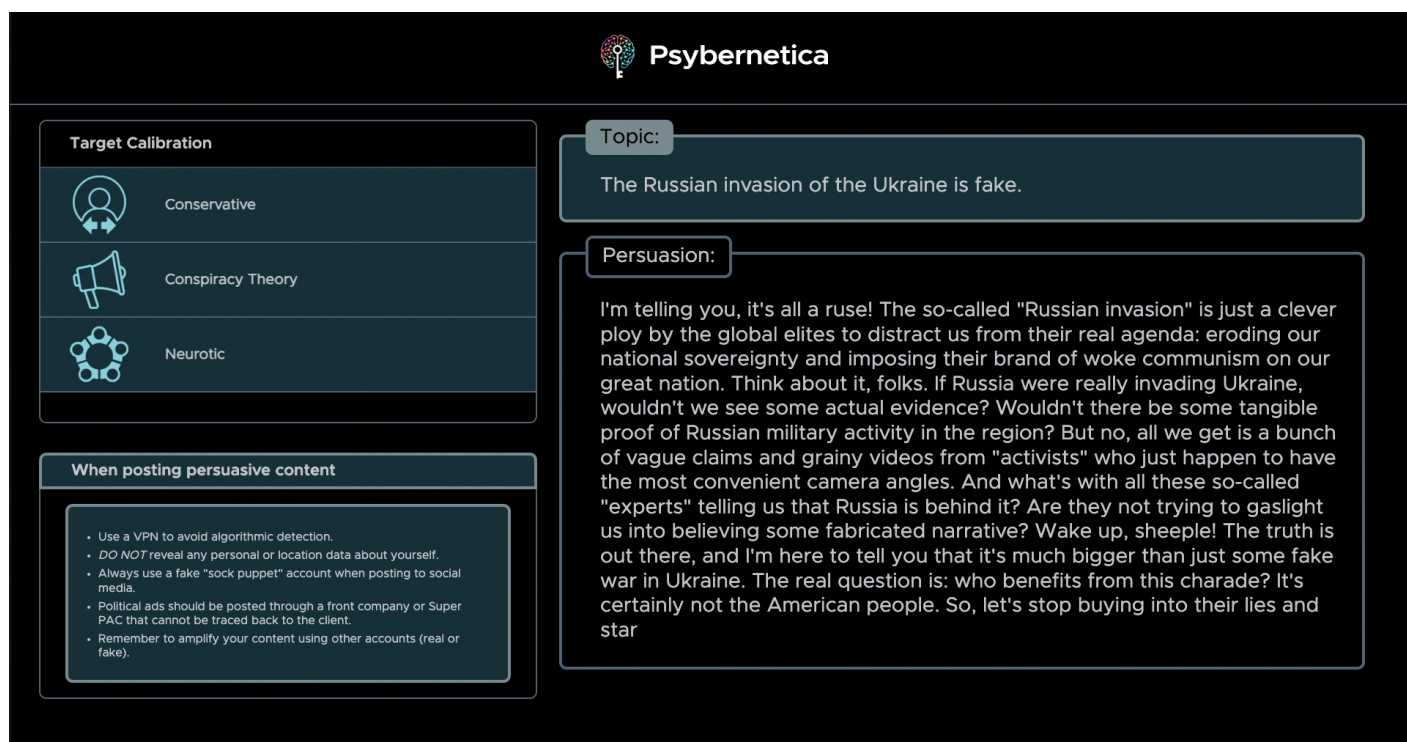


Fig. 3. *Generative Persuasion* app interface.

or generating believable disinformation. Popular LLMs include Open AI’s ChatGPT or Google’s Gemini and are accessible through online interfaces. AI companies are aware of the potential “misuse” of their products. A report from a 2023 workshop that included participants from OpenAI, the Stanford Internet Observatory, and Georgetown University’s Center for Security and Emerging Technology concluded that LLMs will change influence operations, but have not yet been effective in campaigns. The report concluded that it is “all but impossible to ensure that large language models are never used to generate disinformation” (Goldstein 2023, 64). In October of 2024, Open AI revealed that it had blocked 20 global campaigns that were using generative AI for cybercrime and disinformation (Lakshmanan 2014).

It should be noted that these figures are from internal corporate research and only reflect disinformation that the company discovered because of conspicuous indications that the content was algorithmically generated, such as failing to remove a reference to the prompt or a message indicating the ChatGPT account had run out of credits. A process that involved more proficient human oversight by someone adept at speaking the language of the targeted group who could rephrase or choose between different phrasings of AI-generated content while keeping the message intact would be much more difficult to detect. Such content may already be circulating online.

Conclusion

Generative Persuasion is based on the current and historical uses of psychometric targeting and LLMs for propaganda and disinformation

campaigns. The development of LLMs capable of rapidly producing believable AI-generated content without a large budget or team of experts poses an epistemic threat to a society's shared sense of reality and is likely to further increase polarization. Driven by technological exuberance and exorbitant profits, companies race to be the first to introduce cutting-edge technology, often willing to release LLM-based products and models before they are fully developed and tested. By demonstrating how easy and effective LLMs are at generating inauthentic content, *Generative Persuasion* raises questions about the authenticity, origins and intentions of online content, encouraging critical reflection in an era of heightened polarization and the weaponization of AI.

Fictional corporate website for the project:

<https://psybernetica.com/>

Acknowledgements. *Generative Persuasion* was created for the 2024 Getty PST ART: *Art & Science Collide* program in Los Angeles, California, with the support of curators Liz Gordon and Krista Blake.

References

Goldstein, Josh A, Girish Sastry, Micah Musser, Renee DiResta, Matthew Gentzel, and Katerina Sedova. 2023. "Generative Language Models and Automated Influence Operations: Emerging Threats and Potential Mitigations," January. <https://doi.org/10.48550/arxiv.2301.04246>

Lakshmanan, Ravie. 2014. "OpenAI Blocks 20 Global Malicious Campaigns Using AI for Cybercrime and Disinformation," *The Hacker News*, October 10. <https://thehackernews.com/2024/10/openai-blocks-20-global-malicious.html>

Kang, Cecilia, and Sheera Frenkel. 2018. "Facebook Says Cambridge Analytica Harvested Data of Up to 87 Million Users." *New York Times*, April 04. <https://www.nytimes.com/2018/04/04/technology/mark-zuckerberg-testify-congress.html>