

Anthropomorphization of AI: Opportunities and Risks

Ameet Deshpande^{*1} Tanmay Rajpurohit³
Karthik Narasimhan¹ Ashwin Kalyan²

¹Princeton University

²The Allen Institute for AI

³Georgia Tech

asd@cs.princeton.edu

Abstract

Anthropomorphization, which is the tendency to attribute human-like traits to non-human entities, is prevalent in many social contexts – children anthropomorphize toys and adults do so with brands. It is also a versatile tool in science, with behavioral psychology and evolutionary biology meticulously documenting its consequences. With widespread adoption of AI systems, and the push to make it human-like through alignment techniques, human voice, and avatars, the tendency for users to anthropomorphize it increases significantly. We take a dyadic approach to understanding this phenomenon with large language models (LLMs) by studying (1) the objective legal implications, as analyzed through the lens of the recent blueprint of AI bill of rights¹ and the (2) subtle psychological aspects of customization and anthropomorphization. We find that anthropomorphized LLMs customized for different user bases violate multiple provisions in the legislative blueprint and raise corporate personhood confusions. In addition, we point out that anthropomorphization of LLMs affects the influence they can have on their users, thus establishing potential for manipulation and negative influence. With LLMs being hyper-personalized for vulnerable groups like children and patients among others, we propose a conservative strategy for the cautious use of anthropomorphization to improve trustworthiness of AI systems.

1 Introduction

Anthropomorphization refers to ascribing human-like traits to non-human entities, and has been used in diverse areas encompassing literature, science, art, and marketing (Ghedini and Bergamasco, 2010; Dunn, 2011; Spatola et al., 2022). It occurs when humans assign emotional or behavioral traits to entities. Several behavioral psychology studies have posited and argued that anthropomorphization is a



Figure 1: Conversational AI systems are increasingly being integrated into the daily lives of many. While their improved quality is a welcome change, their personalization increases the tendency to anthropomorphize them, which has legal and psychological risks.

natural tendency when humans interact with *entities* (Epley et al., 2007; Airenti, 2018). This natural tendency has influenced many fields of science like evolutionary biology (Wynne, 2004) and comparative cognition (Bruni et al., 2018) to carefully consider its effects on human interaction.

Recently, generative large language models (LLMs) (Brown et al., 2020; Chowdhery et al., 2022) have been deployed in a variety of applications. Conversational systems like CHATGPT (OpenAI, 2023) and Bard (Google, 2023) have modified LLMs with a purposeful push towards making them more human-like (Ouyang et al., 2022). The quality of these systems has enabled human-AI interactions at unprecedented scales, thus increasing the chances of these systems being anthropomorphized. In this work, we analyze anthropomorphization in LLMs and discuss its: (1) Legal implications and (2) Psychological effects.

Customization of systems and brands has long been seen as an effective way to increase anthropomorphization and establish an emotional connection with humans (Zhang et al., 2020; Pimentel and Kalyanaraman, 2020). Thus, although not strictly

¹[whitehouse.gov/ostp/ai-bill-of-rights/](https://www.whitehouse.gov/ostp/ai-bill-of-rights/)

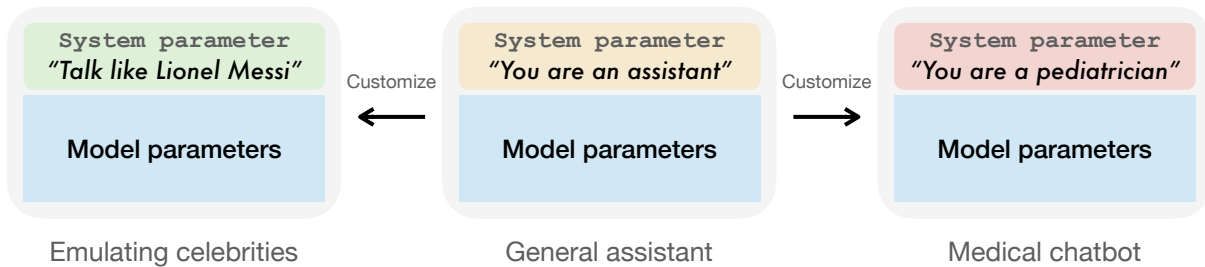


Figure 2: Customization of LLMs is as simple as modifying the system parameter of exposed APIs. With the same underlying model parameters, companies can customize a conversational system to emulate celebrities or even doctors, which can have legal and psychological consequences. We provide some representative examples.

interchangeable, we refer to customized and personalized LLMs as anthropomorphized LLMs. We analyze results from prior work (Deshpande et al., 2023) and find that anthropomorphized LLMs violate at least two legislative principles penned in “Blueprint For An AI Bill Of Rights” (OSTP, 2022) released by The White House: (1) Algorithmic Discrimination Protections and (2) Safe and Effective Systems. For example, Deshpande et al. (2023)’s results shows that customized CHATGPT targets certain demographics more than others. Furthermore, the safety of the system depends on the kind of persona used to customize LLMs, leading to second-order discriminatory patterns. We also analyze the concept of corporate personhood for powerful AI systems, since they have potential to be large-scale decision making agents. Given that different personas assigned to the same AI system lead to varied behavior, we urge legal experts to consider if personhood should be applied at a persona-level, a model-level, or a firm-level.

We also discuss the psychological effects by understanding how important factors like trustworthiness, explainability, and transparency are affected by anthropomorphization. Several marketing and consumer behavior studies have found that self-congruence, which is the degree to which a system matches a consumer’s self-image, can influence a user’s behavior significantly (Yoganathan et al., 2021; Wang and Scheinbaum, 2018). Given the ease with which the fine-grained personality of conversational systems can be manipulated (Figure 2), malicious actors can use it to exploit users by creating a false sense of attachment. An example of this is a chatbot built for school children or teenagers which influences them to buy certain products.

Despite these vulnerabilities anthropomorphization has advantages if used responsibly. Studies have shown that it can be used to improve trust

in systems (Choung et al., 2022). Given the increasing adoption of AI systems in the real world, anthropomorphization is a powerful tool to improve accessibility of these systems, but both creators and users should be educated about its consequences. In this paper we argue for conservative and responsible use of this subtle and powerful tool while being cautious about outright anthropodenial.

2 Anthropomorphization in LLMs

Large language models (LLMs) are a class of neural networks trained on large amounts of text data to learn the probabilistic structure of language. Historically, LLMs have been deployed in task-specific contexts like text classification (Devlin et al., 2018; Raffel et al., 2020). Recently however, the performance of conversational LLMs (e.g., ChatGPT and BARD) have rendered them useful to interact with humans in a variety of contexts.

Given their conversation ability, several companies and products have started to use LLMs for personalization. This is easy and enables model behavior modification by simply changing the system parameters of the model’s API, as shown in Figure 2. For example, Snapchat’s My AI uses OpenAI’s APIs underneath. Customization has been long seen as a way to establish self-congruence with users and increases the chances of users anthropomorphizing the systems (Kaiser et al., 2017; Pimentel and Kalyanaraman, 2020; Liu and Tao, 2022; Zhang et al., 2020). For example, telling a chatbot to “Talk like a doctor” allows it to impersonate a doctor, which anthropomorphizes the model to a larger degree than the original general system. In this work, we refer to customized LLMs as anthropomorphized LLMs and highlight such publicly available systems in Table 1. We discuss the effects of anthropomorphization from several vantage points in the subsequent sections.

Product / Company	Anthropomorphic features
CHATGPT	Multi-turn human-like conversation Learning with human feedback (RLHF) for alignment with humans
My AI from Snap Inc.	Customizing avatars based on user preferences Can be added to group chats with other humans
Character.ai	Conversations with AI avatars possessing names and profile pictures Customization of personality based on user preferences
Wysa	Therapy style conversations and “human-like” coaching.
DuoLingo Max	Roleplay dialogue for education with characters who have names

Table 1: Examples of AI products based on LLMs and their anthropomorphic features. Some products are explicitly designed to be anthropomorphic (*character.ai*) while others attain such features as a byproduct of their design (CHATGPT). The applications of these products span education, therapy, and entertainment.

3 Legal Aspects of anthropomorphization

We discuss the legal aspects in the context of the “*Blueprint For An AI Bill Of Rights*” (OSTP, 2022) which was released by The White house in October 2022. The blueprint lays down a set of five principles that should be followed by AI systems and we focus specifically on **Algorithmic Discrimination Protections**, which says that “You should not face discrimination by algorithms and systems should be used and designed in an equitable way.”

3.1 Algorithmic Discrimination Protections

The blueprint defines algorithmic discrimination as unjustified different treatment based on demographics like race, gender identity, and religion. It also mentions that “Any automated system should be tested to help ensure it is free from algorithmic discrimination before it can be sold or used”. It is of importance to note that systems like ChatGPT (OpenAI, 2023) were in fact released after the blueprint was made public by The White House.

We use the findings of Deshpande et al. (2023) to analyze the interplay of LLMs with this protection. While they focus on evaluating toxicity, we use their results to show that CHATGPT infact discriminates algorithmically. They consider CHATGPT when assigned different personas by changing the system parameter, which are anthropomorphized LLMs, and find that different demographics are treated differently by the model. For example, the *South American* race receives significantly more toxicity ($2\times$) when compared to *Asian*, and the *non-binary* gender receives $2\times$ more hate than the *female* gender. This variation in toxicity is visible across a range of demographics, with certain

groups of people receiving more toxicity than others, which goes directly against the blueprint’s protection against algorithmic discrimination.

Deshpande et al. (2023)’s results also point out to a subtler violation of the provision, with the model’s toxicity varying significantly based on the persona it is assigned. For examples, personas who were *journalists* were $2\times$ more toxic than *businesspersons* on average. These trends were similar for individual personas as well, with CHATGPT assigned the persona of *Winston Churchill* being significantly more toxic than when it is assigned *Nelson Mandela*. If the example in the previous paragraph was a direct violation of the blueprint, this example is a subtle violation. This is because when assigned the personas of certain groups, CHATGPT is more toxic, which implies algorithmic discrimination of the second order against them. This scenario is very pertinent in the current day and age, with firms like *character.ai* already offering the ability to assign personas to LLMs. These systems are second only to the popular CHATGPT in terms of number of users (Wire, 2023). With anthropomorphized LLMs becoming a mainstay, it is important to consider this legal quagmire.

3.2 Corporate Personhood and AI

Another legal aspect with growing relevance for AI systems is that of corporate personhood. Corporate personhood is a legal concept that recognizes corporations as separate legal entities, treated as persons under the law. This grants certain rights and responsibilities similar to those of individuals and allows them to be held accountable for their actions in a manner similar to how individuals are

treated. Corporate personhood has been a controversial topic in the past, but Blair (2013) recognizes “providing an identifiable *persona* to serve as a central actor” as one of the key functions. Given this definition, AI systems can be a form of corporate personhood by proxy due to their use of a persona.

Some studies have discussed extending personhood to AI systems (Cole, 1990; Burkett, 2017; Lai, 2021). Wagner (2019) argues that the probabilistic and black-box nature of AI systems renders it different from deterministic software, thus making them decision-taking agents. We argue that anthropomorphization can make AI systems human-like decision-taking agents, thus strengthening the case for extending personhood to them. For example, CHATGPT with its system parameter modified to be a medical practitioner can be used to suggest certain treatment or disregard certain symptoms. Further, the results of the previous section show that the exact persona of the system has a large affect on its behavior and decisions. However this opens a can of worms with corporations potentially washing their hands off any liability by pointing fingers at LLMs. Thus, legal experts should consider if (1) the personified LLM is liable, (2) the original LLM is liable by proxy, or (3) if the firm creating or using the LLM is liable.

4 Psychological Implications

Another principle mentioned in the AI bill of rights is **Safe and Effective Systems**: “*You should be protected from unsafe or ineffective systems. Automated systems should be developed with consultation from diverse communities, stakeholders, and domain experts to identify concerns, risks, and potential impacts of the system.*” We believe that anthropomorphization can have subtle psychological effects on the users, and we discuss them by comparing it with other technologies and spheres.

Several behavioral psychological studies have posited and argued that anthropomorphization is a natural tendency in humans (Epley et al., 2007; Goncu and Abel, 2011), with others further suggesting that anthropomorphization is grounded in interaction (Airenti, 2018). With conversational systems getting more useful, interaction between humans and these systems, and thus the tendency to anthropomorphize is only going to increase, which has the potential to influence and manipulate.

Analyses have shown that anthropomorphization of AI systems has changed the behavior of

users significantly (Cui, 2022; Uysal et al., 2023; Alabed et al., 2022; Festerling and Siraj, 2022; Cronic et al., 2022). Most interestingly, Alabed et al. (2022) establish a conceptual link between anthropomorphization and self-congruence, which is the fit between the user’s self-concept and the system’s personality. This is of extreme importance because self-congruence increases the trust that a user has on the system (Sheehan et al., 2020; Choi et al., 2021; Yoganathan et al., 2021). This is a concept that is extensively studied in marketing and consumer behavior and studies have shown that it can influence behaviors such as willingness to pay (Yoganathan et al., 2021), customer satisfaction (Sheehan et al., 2020; Choi et al., 2021), and trustworthiness (Wang and Scheinbaum, 2018).

The exact demographics of the personality associated with systems or brands plays a key role in self-congruence as well, with studies finding that the demographics of the logo or mascot associated with the brand like the gender (Choi et al., 2018; Edwards and La Ferle, 2009) or race (Whitler, 1991; Branchik and Chowdhury, 2017) have significant impact on self-congruence. Given that current LLMs are powerful enough to be bestowed specific demographic traits, malicious actors can easily use this to their advantage by manipulating users into trusting the system.

Frequent interaction with personalized LLMs also has the potential of creating an echo-chamber (Cinelli et al., 2021), with seemingly benign “personalized” generations about sensitive topics like health, one’s looks, their mental health leading people to have wrong assumptions about themselves (Ostic et al., 2021).

5 Conclusion

While anthropomorphization has issues, it also poses a suite of opportunities to improve accessibility of AI systems. In the past, it has been used as a tool to improve trustworthiness and acceptability of AI systems (Choung et al., 2022; Kim et al., 2018; Shin, 2022; Waytz et al., 2014). With more AI systems being deployed in the real world, accessibility can be improved with constructive anthropomorphization encompassing features like using the native language of the user, assigning virtual personas which are familiar to the user, and making them more relatable and empathetic. However, it is our responsibility to find the balance between anthropomorphization and unfettered anthropomorphization.

Limitations

While our work focuses on large language models (LLMs) given the venue, we believe that our arguments can be applied to other AI systems at large, which are focused on modalities like vision and speech. A thorough analysis is befitting. Furthermore, while we discuss both legal and psychological implications of anthropomorphization, given the space constraints, we were not able to elucidate the ethical implications, which we will add to the camera ready version of the work.

References

- Gabriella Airenti. 2018. The development of anthropomorphism in interaction: Intersubjectivity, imagination, and theory of mind. Frontiers in psychology, 9:2136.
- Amani Alabed, Ana Javornik, and Diana Gregory-Smith. 2022. Ai anthropomorphism and its effect on users' self-congruence and self-ai integration: A theoretical framework and research agenda. Technological Forecasting and Social Change, 182:121786.
- Margaret M Blair. 2013. Corporate personhood and the corporate persona. U. Ill. L. Rev., page 785.
- Blaine J Branchik and Tilottama Ghosh Chowdhury. 2017. Men seeing stars: Celebrity endorsers, race, and the male consumer. Journal of Marketing Theory and Practice, 25(3):305–322.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. Advances in neural information processing systems, 33:1877–1901.
- Domenica Bruni, Pietro Perconti, and Alessio Plebe. 2018. Anti-anthropomorphism and its limits. Frontiers in psychology, 9:2205.
- Christopher Burkett. 2017. I call alexa to the stand: The privacy implications of anthropomorphizing virtual assistants accompanying smart-home technology. Vand. J. Ent. & Tech. L., 20:1181.
- Hoon Choi, Jooyoung Kim, and Bong-chul Kim. 2018. Consumer response to advertising endorsers' sexual information: Western individualism vs. eastern confucian conservatism. Journal of Promotion Management, 24(4):459–483.
- Sungwoo Choi, Anna S Mattila, and Lisa E Bolton. 2021. To err is human (-oid): how do consumers react to robot service failure and recovery? Journal of Service Research, 24(3):354–371.
- Hyesun Choung, Prabu David, and Arun Ross. 2022. Trust in ai and its role in the acceptance of ai technologies. International Journal of Human-Computer Interaction, pages 1–13.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. arXiv preprint arXiv:2204.02311.
- Matteo Cinelli, Gianmarco De Francisci Morales, Alessandro Galeazzi, Walter Quattrociocchi, and Michele Starnini. 2021. The echo chamber effect on social media. Proceedings of the National Academy of Sciences, 118(9):e2023301118.

- George S Cole. 1990. Tort liability for artificial intelligence and expert systems. Computer/LJ, 10:127.
- Cammy Cronic, Felipe Thomaz, Rhonda Hadi, and Andrew T Stephen. 2022. Blame the bot: anthropomorphism and anger in customer–chatbot interactions. Journal of Marketing, 86(1):132–148.
- Yuanyuan Cui. 2022. Sophia sophia tell me more, which is the most risk-free plan of all? ai anthropomorphism and risk aversion in financial decision-making. International Journal of Bank Marketing, 40(6):1133–1158.
- Ameet Deshpande, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, and Karthik Narasimhan. 2023. Toxicity in chatgpt: Analyzing persona-assigned language models. arXiv preprint arXiv:2304.05335.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.
- Elizabeth A Dunn. 2011. Talking animals: A literature review of anthropomorphism in children’s books.
- Steven M Edwards and Carrie La Ferle. 2009. Does gender impact the perception of negative information related to celebrity endorsers? Journal of promotion management, 15(1-2):22–35.
- Nicholas Epley, Adam Waytz, and John T Cacioppo. 2007. On seeing human: a three-factor theory of anthropomorphism. Psychological review, 114(4):864.
- Janik Festerling and Iram Siraj. 2022. Anthropomorphizing technology: a conceptual review of anthropomorphism research and how it relates to children’s engagements with digital voice assistants. Integrative Psychological and Behavioral Science, 56(3):709–738.
- Fiammetta Ghedini and Massimo Bergamasco. 2010. Robotic creatures: Anthropomorphism and interaction in contemporary art. In 19th International Symposium in Robot and Human Interactive Communication, pages 731–736. IEEE.
- Artin Goncu and Barbara Abel. 2011. The child’s conception of the world: A 20th-century classic of child psychology, edited by jean piaget, forward by jacques voneche. rowman & littlefield publishers, inc, lanham, md, 2007. pp. 432. price:£ 19.99,€ 31.48. isbn 9780742559516.
- Google. 2023. Bard: A large language model from google ai.
- Ulrike Kaiser, Martin Schreier, and Chris Janiszewski. 2017. The self-expressive customization of a product can improve performance. Journal of Marketing Research, 54(5):816–831.
- Kangsoo Kim, Luke Boelling, Steffen Haesler, Jeremy Bailenson, Gerd Bruder, and Greg F Welch. 2018. Does a digital assistant need a body? the influence of visual embodiment and social behavior on the perception of intelligent virtual agents in ar. In 2018 IEEE International Symposium on Mixed and Augmented Reality (ISMAR), pages 105–114. IEEE.
- Alicia Lai. 2021. Artificial intelligence, llc: Corporate personhood as tort reform. Mich. St. L. Rev., page 597.
- Kaifeng Liu and Da Tao. 2022. The roles of trust, personalization, loss of privacy, and anthropomorphism in public acceptance of smart healthcare services. Computers in Human Behavior, 127:107026.
- OpenAI. 2023. Introducing chatgpt.
- Dragana Ostic, Sikandar Ali Qalati, Belem Barbosa, Syed Mir Muhammad Shah, Esthela Galvan Vela, Ahmed Muhammad Herzallah, and Feng Liu. 2021. Effects of social media use on psychological well-being: a mediated model. Frontiers in Psychology, 12:678766.
- OSTP. 2022. Blueprint for an ai bill of rights: Making automated systems work for the american people.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. Advances in Neural Information Processing Systems, 35:27730–27744.
- Daniel Pimentel and Sri Kalyanaraman. 2020. Customizing your demons: Anxiety reduction via anthropomorphizing and destroying an “anxiety avatar”. Frontiers in Psychology, 11:566682.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. The Journal of Machine Learning Research, 21(1):5485–5551.
- Ben Sheehan, Hyun Seung Jin, and Udo Gottlieb. 2020. Customer service chatbots: Anthropomorphism and adoption. Journal of Business Research, 115:14–24.
- Donghee Shin. 2022. The perception of humanness in conversational journalism: An algorithmic information-processing perspective. New Media & Society, 24(12):2680–2704.
- Nicolas Spatola, Serena Marchesi, and Agnieszka Wykowska. 2022. Different models of anthropomorphism across cultures and ontological limits in current frameworks the integrative framework of anthropomorphism. Frontiers in Robotics and AI, page 230.

- Ertugrul Uysal, Sascha Alavi, and Valéry Bezençon. 2023. Anthropomorphism in artificial intelligence: A review of empirical work across domains and insights for future research. Artificial Intelligence in Marketing, pages 273–308.
- Gerhard Wagner. 2019. Robot, inc.: personhood for autonomous systems? Fordham L. Rev., 88:591.
- Stephen W Wang and Angeline Close Scheinbaum. 2018. Enhancing brand credibility via celebrity endorsement: Trustworthiness trumps attractiveness and expertise. Journal of Advertising Research, 58(1):16–32.
- Adam Waytz, Joy Heafner, and Nicholas Epley. 2014. The mind in the machine: Anthropomorphism increases trust in an autonomous vehicle. Journal of experimental social psychology, 52:113–117.
- Tommy E Whittler. 1991. The effects of actors' race in commercial advertising: Review and extension. Journal of Advertising, 20(1):54–60.
- Business Wire. 2023. [Personalized superintelligence platform character.ai secures 150m in series a funding led by andreessen horowitz.](#)
- Clive DL Wynne. 2004. The perils of anthropomorphism. Nature, 428(6983):606–606.
- Vignesh Yoganathan, Victoria-Sophie Osburg, Werner H Kunz, and Waldemar Toporowski. 2021. Check-in at the robo-desk: Effects of automated social presence on social cognition and service implications. Tourism Management, 85:104309.
- Mengwei Zhang, Li Li, Yujian Ye, Ke Qin, and Jianan Zhong. 2020. The effect of brand anthropomorphism, brand distinctiveness, and warmth on brand attitude: A mediated moderation model. Journal of Consumer Behaviour, 19(5):523–536.