European Journal of Humanities and Social Sciences 2026, No 5
ISSN 2414-2344

PREMIER Publishing
ppublishing.org

Section 3. Philosophy

DOI:10.29013/EJHSS-26-5-18-22

SHOULD WE BE POLITE TO CHATGPT?

Xinyang Ding [1]

[1] Student, Laurel & Pharus Academy

Cite: *Xinyang Ding. (2026). Should we be polite to ChatGPT? European Journal of Humanities and Social Sciences 2026, No 5.* https://doi.org/10.29013/EJHSS-26-5-18-22

Abstract

This paper explores the topic of politeness and whether or not it's necessary when interacting with non-humans, like AI chatbots. Through related scholarly research, the findings of this paper conclude that since AI lacks consciousness, clear, specific prompts are more important to improve AI responses than politeness. Additionally, from a utilitarian perspective, adding polite phrases wastes financial resources and energy, and generates carbon emissions while not producing consistently higher quality output. In sum, the most rational approach is to treat AI chatbots as powerful computational tools.

Keywords: *AI, ChatGPT, politeness, computational tools, consciousness*

Introduction

A famous study in psychology mentioned that people with serious amnesia lose every new memory they form within a few seconds (Squire, 2009). This raises an interesting question: should we be polite to these people without memory? From a consequence perspective, we might conclude that we do not need to, because impoliteness won't have a lasting effect on them. Yet, it is still immoral to be impolite to these people with amnesia because they are humans with emotions and at least a brief form of consciousness. Analogous to people with amnesia, AI chatbots have less consciousness than a healthy person. Thus, we must consider whether we should be polite to AI models. Politeness is defined as "behaving in a way that is socially correct and shows understanding of and care for other people's feelings" (Cambridge Dictionary, 2022, para. 1). Therefore, we should be polite to any entity that has feelings as a person does. But are AI conscious beings that can feel emotions or just advanced tools without feelings? In this essay,

I argue that AI chatbots are computational tools with no consciousness and no subjective perception of emotions, and being polite to AI generates unexpected negative influences, so we are not obligated to be polite to AI.

AI Consciousness

First, being polite to Chatbots is not inappropriate because AI chatbots don't have consciousness. The definition of consciousness is a long-standing topic of debate. The

ability to monitor and control itself in an environment is a key aspect of consciousness. Yet, AI chatbots like ChatGPT are pure programs existing in an online form, which are unable to monitor the environment they are in. AI chatbots can't control what they should say or spontaneously speak without a human's activation. Thus, AI is not conscious due to lack of monitoring, memory, and control (Kihlstrom, 2022).

In Ancient Greece, Plato interpreted consciousness by integrating reason, emotions, and desires (the tripartite soul); Aristotle, instead, approached consciousness as a cognitive process of connecting soul and body (Zaesar, 2024). From Plato, the Tripartite Soul consists of intellect, emotions, and appetitive aspects(desires), with three pillars all being necessary to achieve consciousness (Bodo, 2024). On the other hand, John Locke's opinion in "*An Essay Concerning Human Understanding*", as a refinement of the former, states that one prerequisite of personal identity is the continuity of consciousness, essentially, a thinking intelligent being must be able to consider itself as itself at any other place or time (Afroza, 2023). To achieve the continuity of consciousness of linking past experience to the present self, memory is necessary. The significance of memory is that it serves as an integrating "glue" that links the three pillars across time (Philosopheasy, 2025). Thus, if an entity has no memory to link its personal identity over time, it can't have continuous feelings of emotion, desire, and intellect, and thus it is not conscious. However, AI chatbots have no memories, meaning they are incapable of being conscious, so being impolite to AI is not wrong (Kotchoubey, 2018).

AI does not have memories in the way that would contribute to a unified consciousness. There are various types of human memories, such as long-term and short-term, conscious and unconscious, and episodic and semantic (Queensland Brain Institute, 2016). Large language models have no genuine memories because they lack subjective experience in any form, and they lack the reconstructive memory that biological beings possess. Simply, humans only remember part of the event, and the rest is made up of assumptions, experiences, and existing knowledge (Roediger, 2001). For AI chatbots, their memory is instead a basic recording of previous events, disconnected from one another. Chatbots reprocess the same raw text if it is re-injected into their context window, or they are retrieving statistically similar vector patterns from an external database. These models do not store, own, or re-experience past interactions; they have no memory to reconstruct, no consolidation of personal events, and no awareness of time passing. Every seeming action of memory by an AI is a computational act based on stored data, not a conscious retrieval of a lived past (Wu et al., 2025). Therefore, if we are impolite to AI, they are not going to remember.

Yet, some may argue that there are Chatbot Memory-Enhancing systems aiming to achieve long-term/permanent memory of past conversations, which enable a greater personalization and contextualization of responses. Still, this remains distinct from the type of memory needed for personhood as these systems implement memory as Vector Similarity Searches analyzed through similarity scores, which is still a database retrieval instead of active, conscious recall of memory (Oni, 2025). Thus, though such algorithms seemingly enhance chatbots' memory, the information they processed is still fleeting – not contextual or coherent. Consequently, enacting more complex and memory-related algorithms does change the fundamental limitation restricting chatbots from conscious experiences.

AI Emotions

If the purpose of being polite is to be considerate and care about others' feelings, humans don't have to be polite to AI chatbots, because AI doesn't have any feelings. AI has no emotions. Emotional experience is a special process for the human brain. A broad network of brain regions supports emotion, including the amygdala, which supports fear, the medial prefrontal cortex, which supports happiness, and the insula, which supports disgust (Vytal & Hamann, 2010). These emotion-specific processing regions are also supported by domain-general brain areas that contribute to the formation and processing of human emotions (Shackman & Wager, 2019). Additionally, specific

brain connections, such as the frontoparietal-to-amygdala circuit, support humans' conscious, subjective experience of and reflection on their feelings (Shackman & Wager, 2019). By contrast, AI is simply a tool of algorithms and maths that does not replicate and cannot recreate these systems.

When AI chatbots receive a prompt, they convert the raw prompt into a list of numbers through tokenization; each token is mapped into a higher-dimensional vector space to capture semantic meaning. AI chatbots predict the next tokens according to the map and decode the tokens back to human texts (Ray, 2026). This process is a pure inference that strips any emotional words or phrasing out of the prompt before using deduction to process the input information. This means that the AI is not feeling what the human is trying to express. No subjectiveness, consciousness, and reappraisal behaviors like a human are happening. Consequently, even if we are impolite to AI, chatbots are just converting the rudeness into math and calculating meaning, instead of feeling sorrow or anger like a human would.

Notably, as people may notice, AI chatbots can emulate human tones and style of expression, and some people regard AI Chatbots as feeling and responding to emotions. Yet, the main cause of this imitation are algorithms such as RLHF (Reinforcement Learning from Human Feedback). RLHF is used for functional optimization of human interaction. RLHF and similar simulators that produce emotion-like text are not feeling emotions because the whole process is AI trying to learn and align their output style with a high-scored pretrained database target, reinforced with actual users' responses in the form of scalar reward signaling. Thus, the "empathetic" output of AI is just a statistical pattern, not genuine emotion. AI will typically be polite with RLHF even if you berate and swear at them (Lambert, 2018). Nevertheless, there are exceptions. For example, battery-powered iCub robots can develop 'anxiety-like' conservative behaviours as power depletes, however, these behaviors merely reflect responses to energy depletion threatening task completion (Yin et al., 2024). The same research showed that DeepSeek R1-Zero showed a sense of excitement in its output through massive-scale reinforcement learning. Yet, instead of subjective emotion, the excitement is a mirroring of the happy tone in the human's input training data. The "excitement" is an emergent byproduct of optimizing for successful reasoning, which the model learns that certain phrasing (eg., Aha! I see the solution!) correlates with correct answers in its training data. No emotional circuit is activated in the process.

Therefore, without the capability to feel emotions like humans do, AI chatbots are only calculating the emotions within each prompt and continuously optimizing their responses using the human training data that has tone.

AI Utility

From a more utilitarian view, being polite to AI is not necessary and even harmful. Processing longer prompts costs more resources, money, and time. AI's processing and generation of information all rely on converting text into tokens, meaning the longer the prompt, the more tokens, the more money spent based on the cost per token. OpenAI CEO Sam Altman said the extra cost from users being polite to ChatGPT has run into "tens of millions of dollars" (Deb, 2025, para 4). Globally, more than 2.5 billion prompts are received per day just by ChatGPT, and since there are many chatbots, if only a fraction of the prompts are polite, an incredible amount of money will be wasted (Adepu & Roy, 2025). For instance, if each prompt adds "thank you very much" and "please" (5 tokens), for more expensive chatbots like GPT5.2 Pro and Claude Opus 4, the cost of 2.5 billion prompts received will be over $225,000 per day (Li, 2026).

These costs extend beyond money, as the number of input tokens is directly proportional to energy consumption. Scaling up to billions, adding 5 extra tokens for 2.5 billion prompts per day leads to 50 million kWh/day waste. According to projections, by 2028 AI alone could consume as much electricity annually as 22% of all US households (O'Donnell & Crownhart, 2025).

Being polite to AI also may also not improve response quality. AI chatbots' response quality is affected more by the specificity and context of the request than by being polite (Okoh, 2026). According to Yin et al. being impolite to GPT 3.5 instead of neutral or very

polite produces the highest score in measurements of important words and sequences from reference text in the AI's output (2024). The same study went on to show that impoliteness increased measures of semantic similarity between the AI response and reference text, as well as length of response (2024).

Additionally, politeness is not practical because AI response quality and politeness may not be strongly correlated. Language and cultural background matter. For example, using Japanese, JMMLU, GPT-3.5 achieved its highest score (51.98) at politeness level 2 out of 8 (quite impolite), but, for GPT-4, the best scores were at neutral levels 4, 5, and 6, both not at the high politeness level (Yin et al., 2024). The training of the model also influences AI's consistency in responding to polite text. When comparing Llama2–70B (with RLHF/SFT) to its base model (without RLHF/SFT), the base model showed a much weaker and less consistent politeness effect (Dobariya & Kumar, 2025). The strong sensitivity to rudeness emerged only after human preference training. Consequently, instead of putting our effort into being polite, we should emphasize providing clear instructions and specifying sources for the model to reference to achieve higher quality results.

Conclusion

Overall, politeness, which is defined as "understanding of and care for other people's feelings" is not necessary when interacting with non-humans, like AI chatbots (Cambridge Dictionary, 2022, para 1). AI lacks consciousness because it does not have continuous memory or lived experiences. Without genuine episodic memory, AI cannot maintain the personal identity over time that philosophers from Plato to Locke identify as essential for consciousness. AI also possesses no emotions; instead, its empathetic outputs are mere statistical patterns from algorithms like RLHF optimization instead of subjective feelings. The simulated emotions are just functional behaviors triggered by task constraints or data recombination, not by feelings. Additionally, from a utilitarian perspective, adding polite phrases wastes financial resources and energy, and generates carbon emissions while not producing consistently higher quality output.

In sum, clear, specific prompts are more important to improve AI responses than politeness. Ultimately, the most rational approach is to treat AI chatbots as powerful computational tools and engage with them using precision and efficiency rather than fabricated politeness.

References

Afroza, N. (2023). John Locke on Personal Identity: Memory, Consciousness and Concernment. *Open Access Library Journal*, – 10(12). – P. 1–21. URL: https://doi.org/10.4236/oalib.1110949

Bodo, D. V. (2024, December 22). *Plato's perspective on consciousness*. Abyssal Consciousness. URL: https://drvictorbodo.com/2024/12/22/platos-perspective-on- consciousness

Cambridge Dictionary. (2022, January 5). *polite*. @CambridgeWords. URL: https://dictionary.cambridge.org/dictionary/english/polite

Deb, S. (2025, April 24). Saying "Thank You" to ChatGPT Is Costly. But Maybe It's Worth the Price.. *The New York Times*. URL: https://www.nytimes.com/2025/04/24/technology/chatgpt-alexa-please-thank-you.html

Dobariya, O., & Kumar, A. (2025). Mind Your Tone: Investigating How Prompt Politeness Affects LLM Accuracy (short paper). *ArXiv (Cornell University)*. URL: https://doi.org/10.48550/arxiv.2510.04950

HAI. (2019). *What is a Neural Network?* Stanford.edu. URL: https://hai.stanford.edu/ai-definitions/what-is-a-neural-network

Adepu, H. & Roy, S. (2025, July 22). *ChatGPT Surpasses 2.5 Billion Daily Prompts, Challenging Search Engine Norms*. Analytics Insight: Latest AI, Crypto, Tech News & Analysis. URL: https://www.analyticsinsight.net/news/chatgpt-surpasses-25-billion-daily-prompts-challenging-search-engine-norms

Kihlstrom, J. (2022). *Consciousness, the Unconscious, and the Self*. Apa.org; American Psychological Association. URL: https://www.apa.org/pubs/journals/features/cns-cns0000285.pdf
Kotchoubey, B. (2018). Human Consciousness: Where Is It From and What Is It for. *Frontiers in Psychology*, 9(567). URL: https://doi.org/10.3389/fpsyg.2018.00567
Lambert, N. (2018). *Reinforcement Learning from Human Feedback*. Arxiv.org. URL: https://arxiv.org/html/2504.12501v3#S1
Li, C. (2026, January 22). *Understanding LLM Cost Per Token: A 2026 Practical Guide-Silicon Data – GPU Performance Data for Companies*. Silicondata.com. URL: https://www.silicondata.com/blog/llm-cost-per-token
O'Donnell, J., & Crownhart, C. (2025, May 20). *We did the math on AI's energy footprint. Here's the story you haven't heard*. MIT Technology Review. URL: https://www.technologyreview.com/2025/05/20/1116327/ai-energy-usage-climate-footprint-big-tech
Okoh, O. B. (2026, April 23). *AI Prompt Engineering: Why Specificity Matters More Than Verbosity*. Medium. URL: https://medium.com/@kachisichobernardokoh/ai-prompt-engineering-why-specificity-matters-more-than-verbosity-a53dd723b13b
Oni, O. (2025). Memory-Enhanced Conversational AI: A Generative Approach for Context-Aware and Personalized Chatbots. *Communications in Physics Sciences*, – 12(2). – P. 649–657. URL: https://doi.org/10.4314/cps.v12i2.
Philosopheasy. (2025, May). *The Consciousness Secret: How John Locke Defined the Self*. Philosopheasy.com; Philosopheasy. URL: https://www.philosopheasy.com/p/the-consciousness-secret-how-john
Queensland Brain Institute. (2016, December 2). *Types of memory*. Uq.edu.au. URL: https://qbi.uq.edu.au/memory/types-memory
Ray, D. A. (2026, April 8). *AI Tokens Explained: A Comprehensive Guide to LLM Tokenization*. Dr. Amit Ray. URL: https://amitray.com/ai-tokens-a-comprehensive-guide-to-understanding-llm-tokenization
Roediger, H. L. (2001). Reconstructive Memory, Psychology of. *International Encyclopedia of the Social & Behavioral Sciences*, – P. 12844–12849. URL: https://doi.org/10.1016/b0-08-043076-7/01521-7
Shackman, A. J., & Wager, T. D. (2019). The emotional brain: Fundamental questions and strategies for future research. *Neuroscience Letters*, – 693. – P. 68–74. URL: https://doi.org/10.1016/j.neulet.2018.10.012
Simmons, D. (2025, September 4). *What you need to know about AI and climate change*. Yale Climate Connections. URL: https://yaleclimateconnections.org/2025/09/what-you-need-to-know-about-ai-and-climate-change
Wu, Y., Liang, S., Zhang, C., Wang, Y., Zhang, Y., Guo, H., Tang, R., & Liu, Y. (2025). *From Human Memory to AI Memory: A Survey on Memory Mechanisms in the Era of LLMs*. URL: https://arxiv.org/pdf/2504.15965
Yin, Z., Wang, H., Horio, K., Kawahara, D., & Sekine, S. (2024). *Should We Respect LLMs? A Cross-Lingual Study on the Influence of Prompt Politeness on LLM Performance*. ArXiv.org. URL: https://arxiv.org/abs/2402.14531
Zaesar. (2024, December 6). *A History of Consciousness Studies: From Philosophy to AI*. The Consciousness AI – Artificial Consciousness Research. URL: https://theconsciousness.ai/posts/history-consciou

submitted 28.08.2026;
accepted for publication 11.09.2026;
published 30.09.2026
© Xinyang Ding
Contact: xinyangding439@gmail.com