Designing Social Minds: How AI Design and Communication Norms Shape Human Perception, Engagement, and Prosocial Behavior

A dissertation submitted
In partial fulfillment of the requirements for the degree of
PhD in Business Administration

and for the degree of
DOCTEUR EN SCIENCES DE GESTION
DE L'ÉCOLE DOCTORALE
«ÉCONOMIE, MANAGEMENT, MATHÉMATIQUES, PHYSIQUE ET SCIENCES INFORMATIQUES»
ED 405
CY Cergy Paris Université

Presented and defended publicly on the 24th of June 2026 by

Hesam Ghasemi (Mohammadhesam Hajighasemi)

JURY

Amir SEPEHRI	Co-Supervisor	Assistant professor, ESSEC Business School (Cergy, France)
Ayse ONCULER	Co-Supervisor	Professor, ESSEC Business School (Cergy, France)
Sarah MOORE	Referee	Professor, Alberta School of Business (Edmonton, Canada)
Francisco VILLARROEL ORDENES	Referee	Professor, University of Bologna (Bologna, Italy)
Sonja PROKOPEC	President	Professor, ESSEC Business School (Singapore)
Cait LAMBERTON	Examiner	Alberto I. Duran President's Distinguished Professor of Marketing, The Wharton School of the University of Pennsylvania (Philadelphia, USA)

Copyright © 2026 by Hesam Ghasemi
All rights reserved

Dedication

To my sister, Mahdieh, my nephew Arsam, and my niece Arnica,

Though you are no longer with us, your presence has never truly left. I carry you with me in every step of this journey and in every achievement along the way.

I hope this work, in some small way, honors you.

You are always with me.

Acknowledgement

This journey has not been an easy one. It has been marked by moments of joy and hardship, health and illness, and many ups and downs. I am deeply grateful to those without whom this journey would not have begun, continued, or reached its completion.

First and foremost, I would like to thank my wife, Saman. Your unwavering support, patience, and love carried me through the most difficult moments of this journey. You were my constant source of strength. You put my career and ambitions ahead of your own, stood by me through every deadline, stayed up late on long nights, and accompanied me to every talk and presentation. You willingly traded ease and laughter for stress and uncertainty so that I could pursue this path. I will always be deeply grateful for your sacrifices and your belief in me.

Second, I am also deeply grateful to my parents and my sisters for their unconditional support and encouragement. Your belief in me has always been a foundation I could rely on.

Third, I am thankful to my co-authors – Drs. Cait Lamberton, Christian Hildebrand, Cheryl Wakslak, Jan Landwehr, Vivek Astvansh, and Mohammad Atari – for their collaboration, guidance, and kindness throughout this journey. Working with you has been both a privilege and a learning experience.

Finally, last but not least, I would like to thank my advisor, Dr. Amir Sepehri. Before meeting you, I was on uncertain ground in my PhD journey. You supported me, pushed me forward, and believed in me when I needed it most. You have not only been an advisor but also a true and close friend. Your willingness to prioritize my work, even during challenging times in your own life, is something I will always be grateful for.

Chapter 1

INTRODUCTION

Artificial intelligence (AI) technologies have rapidly moved from purely functional tools to interactive social partners embedded in everyday consumer experiences (Garvey and Cian 2024; Puntoni et al. 2021). Firms increasingly rely on AI-driven systems (e.g., chatbots, voice assistants, and AI companions) to interact with customers, provide recommendations, and deliver services at scale. In marketing, these technologies are particularly prominent in customer service contexts, where conversational AI systems are used to answer inquiries, resolve complaints, provide product recommendations, and guide consumers through decision processes (Bergner, Hildebrand, and Häubl 2023; Castelo et al. 2023). By automating routine service interactions, AI allows firms to reduce operational costs, shorten wait times, and improve service availability across digital channels (Davenport et al. 2020).

These developments have made conversational AI one of the fastest-growing segments of marketing technology, with voicebots and chatbots increasingly deployed by firms to manage customer relationships and deliver personalized service experiences. For example, voicebots allow consumers to complete tasks through conversational exchanges that resemble natural human communication, creating service experiences that feel immediate and intuitive (Zierau et al. 2023). As a result, AI-driven conversational interfaces are transforming how firms interact with customers and manage large-scale service operations.

However, despite these operational advantages, consumers often remain hesitant to rely on AI systems in service interactions. A growing body of research documents algorithm aversion, whereby individuals prefer interacting with human agents rather than algorithmic systems even when the latter perform objectively well (Castelo, Bos, and Lehmann 2019; Dietvorst, Simmons, and Massey 2015; Longoni, Bonezzi, and Morewedge 2019). Importantly, this literature has primarily examined algorithms that provide recommendations or decisions rather than engaging users voluntarily in ongoing interaction.

The rise of generative AI fundamentally changes this landscape. Systems such as chatbots, voicebots, and AI companions are increasingly experienced not as decision tools, but as conversational partners embedded in social interaction (De Freitas et al. 2025; Yin, Jia, and Wakslak 2024). In these contexts, users do not merely evaluate AI outputs, but rather they interact with AI agents, form impressions of them, and engage with them in ways that resemble interpersonal exchanges. This shift suggests that consumer responses to AI are increasingly shaped not only by concerns about algorithmic performance, control, agency, and empathy, but also by the social and relational characteristics of the interaction itself (De Freitas, Oguz-Uguralp, and Kaan-Uguralp 2025). That said, emerging work suggests that even in generative AI contexts, some forms of resistance may persist (Oğuz-Uğuralp et al. 2025).

Prior research suggests that one way to improve acceptance of AI systems is through social design cues that make artificial agents appear more human-like or socially relatable (Li, Lai, and Wang 2025; Schanke, Burtch, and Ray 2021; Zehnle, Hildebrand, and Valenzuela 2025). According to the Computers Are Social Actors (CASA) framework, individuals tend to apply human social norms and expectations to technological systems whenever those systems display cues associated with social interaction (Nass and Moon 2000; Nass, Steuer, and Siminoff 1994). Even minimal cues – such as conversational language, vocal tone, or personality – can lead users to perceive AI agents as social partners rather than purely functional tools. As a result, the design characteristics embedded in AI systems can significantly shape how consumers perceive and respond to these technologies (Crolic et al. 2022; Schanke, Burtch, and Ray 2021; Yalcin et al. 2022).

At the same time, the increasing prevalence of AI interactions raises broader questions about how human-AI communication may reshape social cognition and behavior beyond the immediate interaction. As conversational AI becomes integrated into everyday life, repeated

interactions with artificial agents may influence not only how consumers perceive AI systems but also how they communicate and behave toward other people (Hermann, Puntoni, and Morewedge 2025).

This dissertation examines how design and communication choices in human-AI interactions shape consumer perception, emotional engagement, and downstream social behavior. Across three papers, I investigate how different cues embedded in AI systems and AI-mediated communication influence users' responses to artificial agents and the broader consequences of these interactions.

The first paper investigates how vocal design cues influence users' perceptions of AI agents. Specifically, it demonstrates that non-standard accents in voicebots increase anthropomorphic perceptions and engagement by making human diversity cognitively salient. When voicebots communicate with non-standard accents, users become more likely to attribute human-like qualities to them, which increases engagement and willingness to interact with the AI agent. This work extends research on anthropomorphism in marketing by showing that subtle vocal cues – rather than explicit human-like visual features – can activate human-related schemas and improve consumer acceptance of AI service agents.

The second paper introduces petmorphism, a novel relational design framework for AI companions inspired by human-pet relationships. Companion animals occupy a unique psychological role in human social life because they are perceived as affectionate, emotionally responsive, and importantly, nonjudgmental partners (Amiot, Sukhanova, and Bastian 2020; Beck and Katcher 1996; Gullone 2001). Building on schema theory, this research proposes that incorporating pet-like cues – such as animal avatars, pet-inspired names, or playful language – activates pet-related schemas, thereby increasing users' feelings of acceptance, fostering emotional connection, and ultimately enhancing user engagement with petmorphized AI companions.

The third paper examines how communication norms in human-AI interactions can shape social behavior beyond the immediate interaction (Kim et al. 2026; Knott, Sagar, and Takac 2022; Reich and Teeny 2025; Song and Lin 2025; Tey et al. 2024; Zhang et al. 2025). As firms increasingly encourage more efficient, instrumental communication with AI systems, it becomes critical to understand whether such shifts in language carry unintended social consequences. This research demonstrates that discouraging polite language in AI interactions reduces subsequent prosocial behavior toward other humans. This effect emerges because suppressing courtesy diminishes moral elevation, particularly beliefs about the goodness of others, thereby fostering a more self-oriented mindset. Importantly, this spillover effect is specific to AI interactions and does not occur when politeness is discouraged in human-human communication. Furthermore, the research identifies a key boundary condition: when discouraging politeness is framed as serving a collective goal (e.g., environmental sustainability), prosocial behavior is preserved or even enhanced.

Taken together, these three papers demonstrate that AI systems function not only as technological tools but also as social environments that shape how consumers perceive, interact with, and behave toward others. By examining AI design cues and communication norms, this dissertation contributes to a deeper understanding of the psychological and behavioral consequences of human-AI interaction in marketing contexts.

Specifically, this dissertation makes three primary contributions. First, it offers novel insights about anthropomorphization in voice-specific human-AI interactions. Previous work focusing on non-voice attributes has used human faces, names, gestures, and humor to create a human-like presence (e.g., Aggarwal and McGill 2007; Kim and McGill 2011; Schanke, Burtch, and Ray 2021). Other work has explored the effects of first-person language or ascribed personalities to non-human entities to achieve the same end (i.e., humanizing AI agents; Chandler and Schwarz 2010; Crolic et al. 2022). These tools, while certainly useful,

do not capitalize on the unique potential of *voice* to increase anthropomorphic perceptions in the context of voicebots. This dissertation will thus help to form a growing foundation for exploring voice-specific technological design options.

Second, it introduces and theorizes petmorphism. Unlike anthropomorphism, which encourages users to treat AI like people (e.g., Bergner, Hildebrand, and Häubl 2023; Crolic et al. 2022; Schanke, Burtch, and Ray 2021; Yalcin et al. 2022), petmorphism encourages users to treat AI more like beloved pets. Petmorphic AI companions offer a novel, low-risk alternative that may replicate many of the emotional benefits traditionally associated with animal interaction—including psychological (e.g., stress relief), physiological (e.g., improved vital signs), and social (e.g., increased communication) effects (Brown 2004; Cavanaugh, Leonard, and Scammon 2008; Gullone 2001) – while avoiding the practical limitations of real pet ownership (Wada and Shibata 2007). Notably, this approach may be particularly valuable for individuals who are unable to own pets due to allergies, housing restrictions, or lifestyle constraints, offering a new form of emotionally supportive companionship through technology.

Third, it advances research on the broader implications of AI for human psychology and social behavior. A small but growing body of work suggests that interactions with AI can spill over to influence subsequent judgments and interpersonal responding (Kim et al. 2026; Knott, Sagar, and Takac 2022; Reich and Teeny 2025; Song and Lin 2025; Tey et al. 2024; Zhang et al. 2025). The results extend this emerging literature by demonstrating that nudges in the norms of human-AI communication (i.e., suppressing courtesy) can have downstream consequences for how people treat other humans. This suggests that it is not only AI interaction itself, but how AI interaction is enacted, that shapes consequences outside of the AI context.

As AI technologies continue to reshape marketing and customer service, understanding how these systems influence consumer perceptions and behavior becomes increasingly important. By examining how subtle design and communication choices shape human responses to AI, this dissertation provides insights that can help firms design AI systems that foster greater consumer engagement, stronger emotional connection, and more socially beneficial outcomes.

Chapter 2

How Accents Shape Diversity Salience and Voicebot Acceptance

Hesam Ghasemi
ESSEC Business School

Amir Sepehri
ESSEC Business School

Cait Lamberton
University of Pennsylvania

Abstract – Given that consumers often exhibit resistance toward algorithmic systems in service contexts, firms face challenges in encouraging engagement with voicebots. In this paper, we find that users' engagement with voicebots may be enhanced by endowing these tools with accents that deviate from "standard" English pronunciation. Across a range of accents, languages, genders, and contexts, our data suggests that voicebots communicating with non-standard accents make the concept of human diversity salient in the user's mind. Therefore, voicebots communicating with non-standard accents are more readily anthropomorphized. This effect translates into downstream consequences such as greater preference for the accented voicebot in customer service settings, stronger support for startups leveraging these voicebots, higher willingness to pay for monthly subscriptions, and a greater likelihood of opting in to receive future updates and communications. We also document that these effects do not hold when the voicebot has countervailing dehumanizing attributes, when the interaction with the voicebot is about an embarrassing topic, and when users hold strong politically conservative views. Together, this work highlights the benefits of accented speech in improving human-AI interactions. We also hope that it sparks further investigation into factors that make technology, even if technically "imperfect," more readily accepted by consumers.

Keywords: anthropomorphism, voicebots, accent, artificial intelligence

Fueled by advances in artificial intelligence (AI) and natural language processing, voicebots, a widely used form of conversational AI that interacts with users through simulated speech and language (Couedel 2021), offer the potential to shorten wait times (Davenport et al. 2020), cut down personnel costs (Castelo et al. 2023), and potentially enhance the quality of help that customers receive (Stefanowicz 2023). Given these advantages, it is unsurprising that the customer service voicebot segment is projected to grow rapidly, with estimates suggesting approximately 20% compound annual growth through 2034 (Market US 2025). However, users' aversion to algorithm-based systems may impede their widespread acceptance across many applications (Castelo, Bos, and Lehmann 2019; Dietvorst, Simmons, and Massey 2015, 2018), thereby limiting the benefits that AI assistants may offer to both consumers and firms. So, what can companies do to improve users' engagement with voicebots?

The current research explores the potential of non-standard English accents (hereafter referred to as "accented") instead of standard American English (hereafter referred to as "unaccented") as a means of overcoming said resistance. We explore the effects of non-standard accents on downstream user engagement metrics, demonstrating cases in which users are more inclined to interact with an accented (*as opposed to unaccented*) voicebot. Importantly, we explain that positive effects of accented voicebots emerge because a non-standard accent makes human diversity cognitively salient to users as a social cue that does not exist for machines. Making human diversity salient in users' minds, in turn, increases anthropomorphic perceptions of the voicebot (i.e., attribution of human-like characteristics to non-human agents; Duffy 2003), and ultimately increases user engagement. Accordingly, our primary comparison throughout the paper is between accented and unaccented voicebots. Because our mechanism relies on increased anthropomorphic perceptions, we do not predict

that an accented (vs. unaccented) voice will increase engagement across human agents (i.e., accented vs. unaccented human agents).

Our findings make three main contributions. First, we offer novel insights about anthropomorphization in voice-specific human-AI interactions. Previous work focusing on non-voice attributes has used human faces, names, gestures, and humor to create a human-like presence (e.g., Aggarwal and McGill 2007; Kim and McGill 2011; Schanke, Burtch, and Ray 2021). Other work has explored the effects of first-person language or ascribed personalities to non-human entities to achieve the same end (i.e., humanizing AI agents; Chandler and Schwarz 2010; Crolic et al. 2022). These tools, while certainly useful, do not capitalize on the unique potential of *voice* to increase anthropomorphic perceptions in the context of voicebots. Our work will thus help to form a growing foundation for exploring voice-specific technological design options.

Second, our research contributes to the literature examining the influence of accents on consumer perceptions (i.e., accentism). Prior literature studying the impact of accents in human-to-human interactions often highlights disadvantages associated with non-standard accents (e.g., Cargile and Giles 1997; Deshields and Santos 2000; Wang et al. 2013), reporting discrimination against non-standard accented speakers, who are (prejudicially) perceived as less competent, possessing worse communication skills, and less effective in their tasks (Dragojevic et al. 2021; Romero-Rivas, Morgan, and Collier 2022). For example, Zewail et al. (2025) show that speakers with non-standard accents receive lower engagement in public speaking. By contrast, in the realm of AI agents, we find that communicating with a non-standard accent may have *positive* effects, such that AI entities are not discriminated against if they communicate with a non-standard accent, and, in some cases, may even be preferred over unaccented voicebots.

Third, our findings offer actionable insights for managers designing voice-based AI systems. We show that a relatively simple and scalable vocal design choice (i.e., the inclusion of a non-standard accent) can increase user engagement across customer-service, advisory, and open-ended interaction contexts, including behavioral outcomes such as continued interaction, financial support, and willingness to subscribe. At the same time, consistent with research demonstrating that anthropomorphism does not uniformly enhance AI acceptance (e.g., Crolic et al. 2022; Jin, Walker, and Reczek 2024; Srinivasan and Sarial-Abi 2021) and that user characteristics play a central role in shaping responses to AI (e.g., Tully, Longoni, and Appel 2025; Garvey, Kim, and Duhachek 2022; Longoni and Cian 2022), we identify important boundary conditions. Specifically, the effectiveness of accent-based anthropomorphic design depends on contextual factors (e.g., embarrassing interactions) and user characteristics, such as ideological orientation toward diversity. By clarifying when and for whom accent-based anthropomorphism enhances engagement, our research guides the implementation of culturally responsive and strategically tailored voicebot designs.

THEORETICAL BACKGROUND

AI Aversion and Anthropomorphism

Consumers dislike waiting for assistance in service interactions, particularly in call-center environments, where delays are common (Chan, Koole, and L'Ecuyer 2014). However, employing human agents to reduce waiting times is costly (Argyres and Mayer 2007). The adoption of AI assistants may offer a solution to this challenge, not only improving efficiency (Puntoni et al. 2021) but also reducing costs associated with the intricate training required to encourage optimal language use among human service providers (e.g., Packard and Berger 2021).

Nevertheless, a growing body of research in a variety of contexts (e.g., membership application; Yalcin et al. 2022, public services; Longoni, Cian, and Kyung 2022, and medical advice; Longoni, Bonezzi, and Morewedge 2019) suggests a general aversion toward AI. Thus, while AI agents have the potential to improve user experiences, it will be necessary to consider their development in ways that address consumers' psychological and behavioral barriers to acceptance (Yin, Jia, and Wakslak 2024).

One way firms seek to overcome behavioral barriers to adoption is by anthropomorphizing their algorithm-based assistants. When products are anthropomorphized, we may perceive them as we perceive humans (Epley, Waytz, and Cacioppo 2007), with heightened perceptions of the product's ability to act with purpose or agency (Srinivasan and Sarial-Abi 2021). When an AI agent is described or designed as human-like, it can, under certain conditions, appear more approachable and trustworthy, qualities that have been shown to enhance consumer responses across a range of domains (Aggarwal and McGill 2007; Hadi and Valenzuela 2014). Thus, cueing anthropomorphic traits in the AI's appearance, language, or description may help attenuate negative consumer reactions and foster consumer acceptance (Li, Lai, and Wang 2025). This phenomenon is grounded in the Computers Are Social Actors (CASA) paradigm, which posits that people instinctively and often unconsciously apply social norms, expectations, and interpersonal scripts to computers, AI agents, and other media technologies as soon as these systems exhibit even minimal social cues (Nass and Moon 2000; Nass, Steuer, and Siminoff 1994).

To achieve these positive outcomes, prior work has employed various methods to anthropomorphize non-voice AI entities. In general, they suggest that anthropomorphic perceptions of non-voice AI entities (e.g., chatbots) can be increased by creating a "social presence" (e.g., adding a personal name and human avatar; Crolic et al. 2022; Yalcin et al. 2022), adding a communicative delay (Castelo et al. 2023; Schanke, Burtch, and Ray 2021),

or integrating humor (Schanke, Burtch, and Ray 2021). These interventions, implemented in AI assistants, have been shown, in some cases, to foster closer consumer–brand relationships (Bergner, Hildebrand, and Häubl 2023) and allow anthropomorphized AI assistants to receive credit for better offers (vs. a human agent's offer; Garvey, Kim, and Duhachek 2022). However, this work has largely focused on text-based interfaces, leaving open the question of how *voice*-based agents, with their richer communicative cues (Hildebrand and Bergner 2021) and vocal features (Hu et al. 2023), might achieve similar humanizing effects.

Voicebots in Modern Service Strategy

Unlike traditional text-based bots, voicebots enable immersive, immediate, and intuitive service interactions (Zierau et al. 2023), in which consumers complete tasks through conversational exchanges that mimic human communication (Huang and Rust 2021). Such voice-based interactions are not only more natural but also more socially and emotionally engaging than those studied in work on text-based bots. Verbal cues like pitch, tone, and tempo introduce a layer of richness absent from text, which can create a greater sense of social presence and foster trust (Hildebrand et al. 2020; Moffett, Folse, and Palmatier 2021).

We propose that another paralinguistic cue may offer the potential to generate anthropomorphism in ways that build on the potential of text (e.g., human names) or more difficult to implement design choices like the imposition of humor (Schanke, Burtch, and Ray 2021) or potentially problematic introduction of communicative delays (Castelo et al. 2023). Specifically, we suggest that accented speech may, though potentially counter to both intuition and past research, play an important role in achieving this objective.

Accent and Identity in the Globalization Era

With the rise in global and interlingual exchanges (Sharifonnasabi, Mimoun, and Bardhi 2024), non-standard accents (e.g., speaking English with an Indian accent) have become a vital aspect of global communication. Accent is described as a way of pronouncing

words, along with other aspects of language such as grammar, syntax, word structure, and vocabulary, which may vary from the standard language but are generally comparable to it (Gluszek and Dovidio 2010). Listeners' evaluative reactions to speakers' accents affect a wide range of judgments (Cargile and Giles 1997; Dragojevic et al. 2021; Wang et al. 2013). Even when listeners do not recognize a speaker's specific accent (e.g., Indian, Chinese), a non-standard accent can lead to snap judgments about the speaker (Lindemann 2005).

Accents can be classified based on their standardness, or their conformity with "a formal set of rules that determine proper usage among speakers of the language" (Morales et al., 2012, p.33). Thus, standard accents are determined by the prevailing linguistic patterns found in specific geographical areas. For example, in the United States, the standard American English accent, free from regional influence, is known as the baseline American accent and serves as the benchmark for linguistic norms (Ghorshi, Vaseghi, and Yan 2008). However, in the United Kingdom, the Standard British accent is recognized as the conventional benchmark (Boduch-Grabka and Lev-Ari 2021). Therefore, the designation of an accent as "standard" is linked to the dominant speech forms of a particular locale.

There is a reason to expect that non-standard (accented) voicebots may work *against* users' acceptance of voicebots rather than supporting it. Previous research suggests that people who speak with a non-standard accent may be victims of prejudice and discrimination (e.g., Roessel, Schoel, and Stahlberg 2020; Romero-Rivas, Morgan, and Collier 2022; Zewail et al. 2025). Listeners' perceptions of a speaker's accent may affect evaluations of judicial decisions (Romero-Rivas, Morgan, and Collier 2022), credibility (Dragojevic et al. 2021), spontaneous reactions (Roessel, Schoel, and Stahlberg 2020), and customer satisfaction (Wang et al. 2013).

Further, deviations from standard accents can serve as indicators of the speaker's ethnic, national, or other group identities, potentially leading to a differentiation between "us"

and “them” (Formanowicz and Suitner 2020). Variations in speech patterns are fundamentally linked to social dynamics between different groups (Giles 2016; Giles and Johnson 1987). Being associated with the “them” group, often marked by accents, typically leads to negative judgments from listeners belonging to the linguistically dominant group (Calamai and Ardolino 2020). Further, research suggests that certain countries have accent hierarchies, as in the United States, where Indian, Asian, and Hispanic accents are often perceived more unfavorably (Dragojevic and Giles 2016; Wang et al. 2013), whereas European accents are perceived more positively (Gluszek and Dovidio 2010). However, even among the top tier, European-accented speakers can face discrimination (Huang, Frideger, and Pearce 2013; Mai and Hoffmann 2014).

We propose, however, that non-standard accents may be beneficial in the realm of voicebot AI acceptance. We argue this occurs because, unlike standard unaccented speech, accented voicebots make human diversity salient in users’ minds, which, in turn, makes voicebots appear more human-like. Specifically, the propensity of non-standard accents to signal ethnic, national, or group identities (Giles 2016; Giles and Johnson 1987) activates distinctly human concepts related to the voice’s ethnicity, culture, and nationality (Flege 1984), none of which are activated as strongly when non-accented speech is heard. We refer to this activation as an enhancement in *diversity salience,* defined as the extent to which the recognition and appreciation of variations and differences among individuals or groups (Arsel, Crockett, and Scott 2022) become cognitively salient (Homan et al. 2008).

Such activation, in turn, has implications for anthropomorphism. Anthropomorphism increases when human knowledge structures are applied to a nonhuman target (Epley, Waytz, and Cacioppo 2007). The activation of users’ stored knowledge structures about humans who vary in background, race, ethnicity, and other identity dimensions increases anthropomorphic perceptions relative to a case in which a more standard voice fails to make diversity salient.

Non-accented speech, representing the most uniform and typical form of a voice, does not do this. Rather, it reflects "machine," speech—resistant to anthropomorphism. Therefore, we conceptualize diversity salience as a cognitive enrichment pathway such that accented speech provides additional human-specific cues that make it easier for users to infer mind and social presence in the voicebot (Gray, Gray, and Wegner 2007b).

At the same time, users often perceive AI entities as belonging to a different species from humans. When this happens, specieism, defined as assigning different moral worth based on species membership (Caviola, Everett, and Faber 2019, p.1011) may be activated. Unlike humans, who are perceived as naturally variable across backgrounds and cultures, AI entities, when viewed through a speciesist lens (Schmitt 2020), are typically seen as category-based, homogenous, and lacking meaningful diversity. Cues that signal human-specific variation might, therefore, reduce speciesism, shifting a voicebot away from a purely mechanical representation and toward a more human one (Haslam 2006). While broad changes in societal biases lie outside the scope of the present work, we will return to this point in the General Discussion.

Together, we propose that non-standard accents function as cues that make human diversity salient in users' minds by signaling race, nationality, or cultural background, thereby increasing anthropomorphic perceptions of a voicebot that communicates with such an accent relative to those that do not. By contrast, this chain of psychological processes is unlikely to occur when an AI voicebot uses a standard accent, which is more readily perceived as "machine language."

Potential Boundary Conditions

While anthropomorphism can enhance consumer engagement with AI entities (Li, Lai, and Wang 2025), we also recognize that it is not universally beneficial. For example, prior research shows that anthropomorphizing chatbots can reduce customer satisfaction

when users are angry (Crolic et al. 2022), eliminate the reputational buffering effect of attributing errors to algorithms during brand crises (Srinivasan and Sarial-Abi 2021), and heighten discomfort among consumers with high self-presentation concerns (Jin, Walker, and Reczek 2024). Relatedly, research on the uncanny valley suggests that when artificial agents become overly human-like, they may evoke discomfort or a sense of eeriness (Mori 1970). Therefore, anthropomorphism should not be viewed as uniformly advantageous. Instead, its effectiveness depends on contextual and individual factors that shape whether greater perceived humanness is desirable or threatening.

Thus, if our theory is apt, past work suggests several important boundary conditions for the proposed positive effect of non-standard accents. First, the link between non-standard accents and anthropomorphic perceptions may be weakened in some cases. Based on our theory, an accented voicebot can increase anthropomorphic perceptions by making human diversity salient. For this effect to emerge, users must be able to form a human-related schema from the interaction (Epley, Waytz, and Cacioppo 2007). Consumers form such schemas based on multiple cues. When other features of the voicebot explicitly signal dehumanization, the accent may still activate anthropomorphic inferences, but these effects are likely to be weaker. The external cues emphasizing the bot's machine-like nature provide countervailing evidence that undermines the humanization signal conveyed by the accent. As a result, although some anthropomorphic perceptions may still emerge due to the power of non-standard accents (but weaker), they may no longer be strong enough to translate into downstream behavioral responses.

Further, research has shown that people's tendency to assign mind and social presence to anthropomorphized algorithms increases discomfort in sensitive or embarrassing situations (Holthöwer and van Doorn 2023; Jin, Walker, and Reczek 2024). To reduce this discomfort, consumers dehumanize the agent (Sun et al. 2023). As a result, the humanizing effect of a

non-standard accent is attenuated (affecting the path between the independent variable and mediator), leading to weaker anthropomorphic perceptions and, consequently, weaker effects on engagement.

Finally, political ideology may have effects on the way in which diversity cues affect behavior. Political ideology influences individuals' responses to cultural difference and diversity-related signals (Graham, Haidt, and Nosek 2009; Jost et al. 2003). Conservatives tend to prioritize social order and continuity and may perceive diversity cues as norm-challenging, whereas liberals are generally more positively disposed toward signals of variation and inclusion (Irmak, Murdock, and Kanuri 2020). Accordingly, if accented voicebots make cultural differences salient, liberals should respond more positively to such cues, whereas conservatives may react less favorably. Thus, we may see that liberal individuals will respond more positively to accented voicebots than conservative individuals.

OVERVIEW OF STUDIES

Eight pre-registered studies test our theory. In an interactive setting, Study 1 provides evidence that a non-standard accent improves service evaluations for voicebots but not for human service agents. Importantly, voicebots with a non-standard accent receive service evaluations no worse than those of human agents with a standard accent. Study 2 tests the proposed sequential mechanism across multiple non-standard accents, demonstrating that accents increase engagement by first making human diversity salient and subsequently increasing anthropomorphic perceptions. Study 3A extends our empirical investigation to a free-flow, real-time interaction with a ChatGPT-based voicebot, replicating the main effect and establishing the mediating role of anthropomorphic perceptions. Study 3B extends the effect beyond the English language by operationalizing standard versus non-standard accents in French. Study 3C argues against transient novelty as an alternate explanation by tracking

interactions over time in a longitudinal design and shows that effects persist across multiple waves of interactions, days apart. Study 4 then demonstrates process evidence via moderation, orthogonally manipulating and blocking the anthropomorphism construct. Study 5 shows that the effect disappears in embarrassing service interactions, and Study 6 demonstrates that the positive effect of a non-standard accent emerges among more liberal participants but reverses among strongly conservative participants. All preregistrations, materials, scenarios, dependent measures, audio files, syntaxes, R/Python codes, and data can be found at (https://osf.io/9sykb/overview?view_only=317d303805184772b8d34f5ebc71f937).[1]

STUDY 1

As mentioned, human speakers with non-standard accents often receive less favorable judgments. Thus, if listeners respond similarly to accented AI voicebots, accented voicebots should also receive less favorable evaluations. However, our theory predicts the opposite. Therefore, Study 1's design allows us to (1) establish the core accented-versus unaccented-voicebot comparison and (2) demonstrate that the observed benefits of non-standard accents are specific to AI agents rather than reflecting a general preference for accented speech. Thus, if our proposed mechanism is specific to the anthropomorphism of AI agents, we should observe positive effects of non-standard accents on voicebots but not on human service agents.

Participants, Design, and Procedure

[1] All studies in this paper are pre-registered and participants who failed the attention check were dropped from the final sample according to the pre-registration protocol. Post-hoc sensitivity analyses for all studies are presented in Supplementary Materials. In all studies, participants were American—born, raised, and residing in the USA—except in Study 3B, where participants were French students born and raised in metropolitan France. A total of 22 different voices were used across all studies, all of which were trained and produced either by Microsoft (All studies except Study 2) or ElevenLabs (Study 2) to ensure consistent voice quality.

A total of 804 participants on Prolific (M_{Age} = 43.40, 55.3% female) were randomly assigned to a condition in a 2 (Agent type: human vs. voicebot) × 2 (Accent type: unaccented, i.e., standard American vs. accented, i.e., Indian) between-subjects design.[2]

Participants imagined they had an inquiry about baggage allowance for a trip to Europe. To do this, they visited the airline's website and, upon clicking the "chat" button, began interacting with a real voicebot we created. The voicebot's only function was to answer questions about baggage allowance for European destinations. Depending on the condition, participants read: "You click on the 'Chat' function to ask your question from a live human assistant [an AI voicebot] named Sam."

To better simulate human interaction and avoid vocal confounds, we hired a single professional actor to provide us with the human voices for both American and Indian conditions. Depending on the condition, participants either heard AI-generated voices or those provided by the human actor. Moreover, following previous work (Castelo et al. 2023), we added communicative delays to voicebots presented as human agents to increase the perceived credibility of human interaction.

After interacting with either a voicebot or a human agent with a standard American (unaccented) or an Indian accent (accented), participants rated their satisfaction using a three-item scale (adapted from Castelo et al. 2023): "I am satisfied with my service experience," "I think Sam was helpful with my inquiry," and "I learned what I needed to from my interaction with Sam" (1=Strongly disagree, 7=Strongly agree," α = .91). As a manipulation check, we asked participants: "To what extent do you think Sam's accent was non-standard or standard (i.e., native) American accent? (1= Definitely non-standard accent, 7= Definitely standard

[2] Prolific recruited 804 participants slightly exceeding our target of 800 as specified in the pre-registration. This discrepancy was beyond our control.

American accent).” The study concluded with basic demographic questions collected for descriptive purposes.

Results

Manipulation check. Results revealed a significant effect of accent on its perception as either standard or non-standard for both human voices used for live agents ($M_{Unaccented}$ = 6.16, SD = 1.07 vs. $M_{Accented}$ = 2.45, SD = 1.75; F(1, 800) = 540.14, p<.001, η_p^2 = .40) and commercial text-to-speech voices used for voicebots ($M_{Unaccented}$ = 5.41, SD = 1.52 vs. $M_{Accented}$ = 2.55, SD = 1.90; F(1, 800) = 330.03, p<.001, η_p^2 = .29). Post-hoc analyses showed no difference in the perception of Indian accents as non-standard in the study (Mean difference = 0.10, p = .530). However, commercial text-to-speech voices (i.e., digitized voices used for the unaccented voicebot), despite being labeled as standard American by companies and trained with a standard American accent, were perceived as less standard than an American human voice (Mean difference = 0.75, p<.001).

Service evaluations. An ANOVA using agent type and accent to predict participants’ evaluations revealed no main effects of agent type (M_{Human} = 6.18, SD = 1.05 vs. $M_{Voicebot}$ = 6.18, SD = 1.03; F<1, p = .955, η_p^2 = .000) or accent type on participants’ evaluation ($M_{Unaccented}$ = 6.15, SD = 1.08 vs. $M_{Accented}$ = 6.20, SD = 1.00; F<1, p = .523, η_p^2 = .001) on this measure. However, we do observe a significant interaction of agent type and accent type (F(1, 800) = 5.66, p = .018, η_p^2 = .007). The non-standard accent did not affect service evaluation in the human-agent conditions ($M_{Accented}$ = 6.11, SD = 1.04 vs. $M_{Unaccented}$ = 6.24, SD = 1.06; F(1, 800) = 1.50, p = .22, η_p^2 = .002). However, consistent with our theory, the accented voicebot increased service evaluations relative to the unaccented voicebot ($M_{Accented}$ = 6.29, SD = 0.95 vs. $M_{Unaccented}$ = 6.07, SD = 1.10; F(1, 800) = 4.62, p = .032, η_p^2 = .006).

Discussion

Study 1 establishes the core comparison between accented and unaccented voicebots and clarifies that the effect is specific to AI-mediated interactions. The results suggest that voicebots with a non-standard accent can receive service evaluations no worse than those of human agents with a standard accent. Moreover, although the commercial text-to-speech voice used for the unaccented voicebot is labeled and trained as standard American, participants perceived it as less standard than a human American voice. This pattern may reflect a broader tendency for consumers to categorize AI agents as an outgroup (e.g., an implicit "us vs. them" distinction), indicating a form of speciesism toward artificial agents (Schmitt 2020). In this context, it is particularly noteworthy that introducing a more clearly non-standard accent still increased acceptance of the voicebot.

In the following studies, we focus solely on the comparison between accented and unaccented voicebots, such that we can directly test the underlying roles of diversity salience and anthropomorphism in driving effects on voicebot acceptance.

STUDY 2

Study 2 tests our proposed sequential mediation model across a wide variety of different accents. Specifically, we examine whether non-standard accents increase anthropomorphic perceptions of voicebots, and whether this effect emerges because such accents make human diversity cognitively salient in users' minds. In turn, we test whether these heightened anthropomorphic perceptions translate into greater user engagement. Beyond establishing the mechanism, this study also examines downstream behavioral consequences by testing whether consumers are more willing to interact with the accented voicebot rather than wait in line for a human assistant. Lastly, this study establishes the generalizability of our hypothesized effects by testing them across six different accents.

Participants, Design, and Procedure

A total of 700 participants on Prolific (M_{Age} = 42.92, 46.8% female) were randomly assigned to one of seven between-subjects voicebot accent conditions: 1. Unaccented (standard American accent) vs. 2. Accented (British) vs. 3. Accented (Indian) vs. 4. Accented (Nigerian) vs. 5. Accented (French) vs. 6. Accented (Russian) vs. 7. Accented (Italian).

Participants imagined visiting their insurance company's website to make an inquiry about their coverage. They then read: "You click on the 'Chat' function to ask your question from an assistant, which is an algorithmic AI voicebot named Sam. On the next page, you will start interacting with Sam." We created seven identical real voicebots with no human avatar, each using a male voice and named Sam, differing only in accent. The voicebots were limited to presenting an initial opening message, after which participants could not continue the conversation (examples of the voicebots can be found here: https://osf.io/fgj93/overview?view_only=0e6c7c3e0bd44d61ab58db40ba64a6a3).

After starting their interaction, depending on the condition, participants read and heard the opening message with the accent of the condition they were assigned to: "Welcome to Sunshine Insurance. I'm Sam, your trusted AI insurance assistant. I can answer all your questions regarding our insurance plans, extent of coverage, claims processing, and more. You can either choose to proceed with me or wait in line to chat with a live agent."

We then assessed participants' preference for continuing with the voicebot versus waiting in line to interact with a live agent by asking: "How will you proceed with your inquiry today? (1 = Definitely wait in line for a live agent, 4 = Indifferent between a live agent and Sam, 7 = Definitely continue with Sam)". After that, participants indicated how anthropomorphic they perceived the voicebot on a nine-item scale (adapted from Crolic et al. 2022): "Please rate the extent to which Sam (1= Not at all, 7= Very much; α = .97): came alive (like a person) in your mind; has some humanlike qualities; seems like a person; felt human; seemed to have a personality; seemed to have a mind of his/her own; seemed to have

his/her own intentions; seemed to have free will; and seemed to have consciousness." We also measured diversity salience by asking: "When you were listening to Sam, to what extent did you think about the following? (1 = Not at all, 7 = Very much; α = .96): differences in people and their backgrounds; the variety of people in the world; how people differ in their experiences and perspectives; the richness of human perspectives; and the diversity of ways people think and live." The study concluded with basic demographic questions.

Results

Diversity salience. ANOVA revealed a significant effect of voicebot accent on diversity salience (Table 1; $F(6, 679) = 5.39$, $p<.001$, $\eta_p^2 = .046$). Follow-up post-hoc tests revealed that all of the accented voicebots could make human diversity cognitively salient to a greater extent than the unaccented voicebot.

Anthropomorphic perceptions. ANOVA showed a significant effect of voicebot accent on anthropomorphic perceptions (Table 1; $F(6, 679) = 3.15$, $p = .005$, $\eta_p^2 = .027$). Follow-up comparisons revealed that adding a non-standard accent raised anthropomorphic perceptions across all accented conditions compared to the unaccented voicebot, except the British accent.

Preference for interacting with the voicebot. ANOVA showed a significant effect of voicebot accent on preference for interacting with the voicebot (Table 1; $F(6, 679) = 2.88$, $p = .009$, $\eta_p^2 = .025$). Follow-up comparisons showed that, across all non-standard accents, participants indicated a higher preference for interacting with the accented voicebot (vs. waiting for a live human agent) compared to the unaccented voicebot.

Table 1. Results of univariate analyses and pairwise differences between each non-standard accent and the standard American accent

	Diversity Salience		**Anthropomorphic Perceptions**		**Preference for interacting with the voicebot**	
	M (SD)	Δ (p, η_p^2)	M (SD)	Δ (p, η_p^2)	M (SD)	Δ (p, η_p^2)
American	2.52 (1.64)		3.34 (1.73)		4.51 (2.32)	

British	3.48 (1.81)	0.96 (<.001, .073)	3.28 (1.87)	−0.06 (.807, .000)	5.27 (2.21)	0.76 (.019, .028)
Indian	3.03 (1.81)	0.51 (.044, .022)	3.84 (1.56)	0.50 (.042, .022)	5.70 (1.92)	1.19 (<.001, .073)
Nigerian	3.71 (1.85)	1.19 (<.001, .105)	4.06 (1.77)	0.72 (.006, .041)	5.38 (2.01)	0.87 (.007, .039)
French	3.66 (1.73)	1.14 (<.001, .104)	3.98 (1.87)	0.64 (.014, .030)	5.20 (2.28)	0.69 (.036, .022)
Russian	3.47 (1.81)	0.95 (<.001, .071)	3.85 (1.77)	0.51 (.048, .021)	5.51 (2.12)	1.00 (.002, .048)
Italian	3.50 (1.82)	0.98 (<.001, .075)	3.95 (1.67)	0.61 (.014, .031)	5.36 (2.11)	0.85 (.008, .036)

Notes: Δ shows the mean difference with the American accent condition.

Mediation. We conducted six separate sequential mediation analyses comparing the unaccented condition to each accented condition, with diversity salience as the proximal mediator and anthropomorphic perceptions as the distal mediator (PROCESS Model 6, 5000 bootstrapped samples; Hayes 2022).

Across all accents, we find consistent support for the proposed sequential pathway from accent to preference via diversity salience and anthropomorphic perceptions ($a_1 \cdot d \cdot b_2$; Table 2). Specifically, non-standard accents increase diversity salience (a_1s ≥ 0.255, ts ≥ 2.02, *p*s ≤ .045), and diversity salience strongly predicts anthropomorphic perceptions (ds between 0.407 and 0.507, all *p*s<.001). In turn, anthropomorphic perceptions predict preference for interacting with the voicebot (b_2s between 0.412 and 0.545, all *p*s<.001).

Importantly, this effect emerges despite a weak or non-significant direct path from accent to anthropomorphic perceptions (a_2), which is null for most accents (ts ≤ 1.30, *p*s ≥ .197) and only negative for the British accent (b = −0.276, t = −2.36, *p* = .019). Likewise, the direct effect of diversity salience on preference (b_1) is not significant across all conditions (ts ≤ 1.90, *p*s ≥ .059).

Table 2. Results of mediation analyses

Non-standard accent	a_1 (t, *p*)	a_2 (t, *p*)	d (t, *p*)	b_1 (t, *p*)	b_2 (t, *p*)	c' (Direct) (t, *p*)	Indirect ($a_1 \cdot d \cdot b_2$) [LLCI, ULCI]
British	0.484 (t = 3.90, *p* < .001)	−0.276 (t = −2.36, *p* = .019)	0.505 (t = 7.74, *p* < .001)	0.182 (t = 1.90, *p* = .059)	0.477 (t = 5.19, *p* < .001)	0.309 (t = 2.04, *p* = .043)	0.117 [0.047, 0.209]
Indian	0.255 (t = 2.02,	0.143	0.407 (t = 6.43,	0.082	0.472 (t = 4.89,	0.459	0.049 [0.001, 0.118]

	$p = .045$)	(t = 1.30, p = .197)	$p < .001$)	(t = 0.89, p = .374)	$p < .001$)	(t = 3.16, p = .002)	
Nigerian	0.597 (t = 4.67, $p < .001$)	0.056 (t = 0.48, p = .633)	0.507 (t = 7.98, $p < .001$)	0.069 (t = 0.70, p = .486)	0.412 (t = 4.16, $p < .001$)	0.244 (t = 1.55, p = .122)	0.125 [0.047, 0.232]
French	0.573 (t = 4.77, $p < .001$)	0.031 (t = 0.26, p = .794)	0.500 (t = 7.40, $p < .001$)	0.146 (t = 1.46, p = .146)	0.473 (t = 5.05, $p < .001$)	0.112 (t = 0.71, p = .477)	0.136 [0.058, 0.245]
Russian	0.476 (t = 3.80, $p < .001$)	0.018 (t = 0.16, p = .875)	0.493 (t = 7.64, $p < .001$)	0.101 (t = 1.05, p = .295)	0.500 (t = 5.24, $p < .001$)	0.323 (t = 2.15, p = .033)	0.117 [0.046, 0.215]
Italian	0.494 (t = 3.97, $p < .001$)	0.095 (t = 0.84, p = .405)	0.419 (t = 6.61, $p < .001$)	0.079 (t = 0.87, p = .386)	0.545 (t = 5.84, $p < .001$)	0.223 (t = 1.50, p = .135)	0.113 [0.044, 0.215]

Discussion

Study 2 provides support for our proposed sequential mechanism. Non-standard accents increased preference for interacting with the voicebot, and this effect was explained by a two-step process in which accents first made human diversity cognitively salient and, in turn, increased anthropomorphic perceptions. Importantly, anthropomorphism did not arise directly from accent cues, but instead emerged through diversity salience, which subsequently enhanced preference for interacting with the accented voicebots.

A notable exception emerges for the British accent. Although it increased diversity salience, it did not enhance anthropomorphic perceptions. One explanation is that the British accent may be relatively familiar and culturally normative, making it less effective at translating diversity salience into anthropomorphism.

STUDY 3A

Study 3A seeks replication of support for our theory in a more interactive, free-flowing setting, beyond call center customer-service contexts. Here, participants engaged with a real voicebot developed using the ChatGPT API (Model 3.5), capable of real-time interaction via both text and voice. To enhance external validity and generalizability,

participants were allowed to discuss a topic of their choice. Additionally, to ensure that the observed effects are not driven by voice gender, we used female voices across conditions.

Participants, Design, and Procedure

A total of 200 participants on Prolific (M_{Age} = 42.57, 49.5% female) were randomly assigned to one of two between-subject voicebot accent conditions: Unaccented (standard American) vs. Accented (Indian).

Using OpenAI's platform and ChatGPT's API, we built two voicebots that communicated in both text and a female voice, with either a standard American or an Indian accent. Behind the scenes, our voicebots communicate with the ChatGPT API to retrieve real-time responses to each user prompt. Participants read: "SonicAI has created a voicebot called Ava that can be used by companies for providing information on products, handling customer service inquiries, and more. You've been recruited to help test Ava by asking between 2 to 10 questions." After reading this, participants were redirected to the voicebot window to start their interaction, where they could communicate via chat and receive the voicebot's responses in both voice and text. The voicebot's opening message and generated answers were read in a standard American or Indian accent, depending on the condition. Participants were allowed to select any topic that was on their minds to discuss with Ava. Some topics discussed by participants included goal-seeking (e.g., dietary plans), history (e.g., the First and Second World Wars), and human-AI interaction (e.g., the safety of autonomous cars). When participants' interaction ended, Ava generated a code they could use to continue the survey.

After returning to the survey and entering their completion code, participants reported their attitudes toward the voicebot ("What is your attitude toward Ava?") using three bipolar items (1 = "dislike a great deal" / "very negative" / "not favorable at all," 7 = "like a great deal" / "very positive" / "very favorable"; α = .97; adapted from Yalcin et al. 2022).

Participants also indicated how anthropomorphic the voicebot was, using the same items as Study 2 (α = .97). The study concluded with basic demographic questions.

Post-test

To ensure that the voices used in Study 3A did not suffer from unintended perceptual imbalances, we conducted a pre-registered post-test (N = 200; M_{Age} = 46.68, 48.00% female). Participants listened to the same standard American and Indian-accented voices and evaluated them on several perceptual dimensions.

Results revealed that the unaccented voice was evaluated more favorably on multiple attributes such as attractiveness ($M_{Unaccented}$ = 6.14 vs. $M_{Accented}$ = 5.23), pleasantness ($M_{Unaccented}$ = 5.70 vs. $M_{Accented}$ = 5.11), welcomingness ($M_{Unaccented}$ = 5.82 vs. $M_{Accented}$ = 5.26), and clarity ($M_{Unaccented}$ = 6.38 vs. $M_{Accented}$ = 5.08), all *p*s ≤ .002. Importantly, there were no significant differences in warmth ($M_{Unaccented}$ = 5.23 vs. $M_{Accented}$ = 5.07, p = .421), friendliness ($M_{Unaccented}$ = 6.35 vs. $M_{Accented}$ = 6.06, p = .090), or human-likeness ($M_{Unaccented}$ = 4.76 vs. $M_{Accented}$ = 4.63, p = .591).

Overall, these results indicate that perceptual imbalances (where existent) favor the unaccented American voice. Therefore, the patterns observed in this post-test suggest that the observed increases in engagement and anthropomorphism for accented voicebots cannot be attributed to superior voice quality, clarity, or surface appeal.

Results

Anthropomorphic perceptions. ANOVA revealed that participants perceived the accented voicebot as more anthropomorphized ($M_{Accented}$ = 3.76, SD = 1.74) compared to the unaccented voicebot ($M_{Unaccented}$ = 3.02, SD = 1.68; F(1, 196) = 9.17, p = .003, η_p^2 = .04).

Attitude toward the voicebot. ANOVA further revealed that participants' attitude toward the voicebot was more positive when they interacted with the accented voicebot

($M_{Accented}$ = 5.38, SD = 1.38) compared to the unaccented voicebot ($M_{Unaccented}$ = 4.86, SD = 1.62; F(1, 196) = 5.75, p = .017, η_p^2 = .03).

Mediation. Mediation analyses (PROCESS Model 4, 5000 bootstrapped samples; Hayes 2022) confirmed that anthropomorphic perceptions mediated the effect of voicebot accent on participants' attitude toward the voicebot (indirect effect = 0.19, 95%CI = [0.06, 0.32]). The direct effect of voicebot accent on DV was not significant (direct effect = 0.06, 95%CI = [−0.10, 0.24]), indicating full mediation by anthropomorphic perceptions. Specifically, using a non-standard accent in the voicebot made it more human-like (b = 0.36, SE = 0.12, t = 3.02, p = .002), thereby increasing customers' attitude toward the accented voicebot compared to the unaccented counterpart (b = 0.51, SE = 0.05, t = 10.18, p<.001).

Discussion

Study 3A's results replicate the positive effects of non-standard accents in an interactive free-flowing scenario. We also observe that the same effects emerge in a non-marketplace context, thus arguing against the possibility that the previously observed effects are due to the fit between accents and call-center interactions.

STUDY 3B

Study 3B has three main goals. First, thus far, our base language has always been English. A full test of our theory requires that we test accented versus unaccented speech in another language. Therefore, Study 3B, using a similar interaction structure to Study 3A, tests our hypothesized effects in another language (i.e., French), where both standard and non-standard accents are recognized. Second, Study 3B examines the effect of non-standard accents in the context of AI medical advice, where AI aversion has been previously demonstrated (i.e., medical advice; Longoni, Bonezzi, and Morewedge 2019). Last, Study 3B uses a behavioral measure (i.e., instead of attitude toward the bot in Study 3A) to assess

participants' willingness to financially support an AI voicebot developed by a startup company.

Participants, Design, and Procedure

A total of 80 participants (M_{Age} = 21.87, 63.30% female) were recruited from a major Business School in France. To manipulate accent standardness, we relied on existing variation in French-language accents. Specifically, the accent used in French-speaking voicebots is considered the standard French accent (Maison and Estève 2023). For example, French Siri responds to "Dis Siri" instead of "Hey Siri." However, one of the common non-standard accents for French is the Québécois accent. Based on our theory, if a voicebot communicates with a Québécois accent, French users (i.e., native to France) should perceive the voicebot as more anthropomorphized and, consequently, engage with it more. Thus, participants were randomly assigned to one of two between-subject voicebot accent conditions: Unaccented (standard French) vs. Accented (Québécois French). The pre-registration protocol set a data collection target of four weeks or 200 participants, whichever occurs first. Our data collection concluded after four weeks with 80 participants.

The procedure was similar to Study 3A, except that the voicebots were powered by the ChatGPT 4o-mini model. Participants read: "SonicAI has created a voicebot called Ava that can be used for providing healthcare suggestions, giving medical advice, offering wellness tips, and more. You've been recruited to help test Ava by asking between 2 to 10 questions. These questions should be related to health and wellness. Examples include asking about symptoms, preventative measures, or general health tips." After reading the scenario, participants started their interaction with the voicebots and upon finishing their interaction, they returned to the survey to indicate the portion of a €10 bonus (in addition to the standard study compensation) they would be willing to allocate to support the voicebot developed by Sonic AI, should they be the one winning it (using a slider scale between 0 and 100). We used

the same scale as in previous studies to measure anthropomorphic perceptions (α = .95). The study concluded with the collection of basic demographics.

Results

Anthropomorphic perceptions. An ANOVA comparing the unaccented to accented condition revealed that participants perceived the accented voicebot as more anthropomorphized ($M_{Accented}$ = 2.78, SD = 1.33) compared to the unaccented voicebot ($M_{Unaccented}$ = 2.12, SD = 1.16; F(1, 77) = 5.50, p = .022, η_p^2 = .067).

Willingness to support the voicebot. An ANOVA comparing the unaccented to accented condition revealed that participants financially supported the accented voicebot ($M_{Accented}$ = 40.24, SD = 41.15) more than the unaccented voicebot ($M_{Unaccented}$ = 23.18, SD = 29.85; F(1, 77) = 4.39, p = .039, η_p^2 = .054).

Mediation. Mediation analyses (PROCESS Model 4, 5000 bootstrapped samples; Hayes 2022) confirmed that anthropomorphic perceptions mediated the effect of voicebot accent on participants' willingness to support the voicebot (indirect effect = 4.19, 95%CI = [0.58, 9.03]). The direct effect of accent on DV was not significant (direct effect = 4.33, 95%CI = [−3.25, 11.93]), suggesting full mediation. Specifically, using a non-standard accent featured in the voicebot made it more human-like (b = 0.34, SE = 0.14, t = 2.34, p = .017), thereby increasing customers' willingness to support the accented voicebot compared to the unaccented counterpart (b = 12.63, SE = 2.96, t = 4.26, p<.001).

Discussion

Study 3B suggests our theory's potential generalizability, showing that a voicebot's accent influences anthropomorphic perceptions and engagement in a language beyond English. Moreover, using a behavioral DV (i.e., actual monetary support) in a new context bolsters confidence in the relevance of our findings across different domains.

STUDY 3C

Voicebots (e.g., Alexa, Google Home, Starbucks Barista Bot, ChatGPT voicebot) typically have a standard American accent by default. Therefore, one might argue that equipping a voicebot with a non-standard accent (e.g., an Indian accent) could initially make it seem novel. However, this effect might fade once the novelty wears off.

Consequently, Study 3C seeks to examine whether the effects of using non-standard accents in voicebot interactions endure over time, arguing against positive responses to transient novelty as an explanation. To achieve this, we conducted a longitudinal study with a pre-registered protocol in which users interacted with voicebots three times over a pre-determined period. Finally, in this study, we combined the number of interactions across all three waves as a measure of cumulative behavioral engagement (Bergner, Hildebrand, and Häubl 2023; De Freitas et al. 2025) and tested whether this engagement differed by accent condition.

Participants, Design, and Procedure

A total of 313 participants on Prolific (M_{Age} = 41.21, 52.1% female) were randomly assigned to one of two between-subject voicebot accent conditions: Unaccented (standard American) vs. Accented (Indian).

The procedure and voices were similar to those in Study 3A, and participants interacted with the designed ChatGPT-based voicebots (model 4o-mini) in real time, using a longitudinal design. Behind the scenes, our voicebots call ChatGPT's API and fetch ChatGPT's actual response to each user's prompt. Participants interacted with a voicebot named Nina three times, with a 4-day interval between interactions, for a total study duration of 12 days (including a 3-day window for participants to complete the second and third waves). For the first interaction, participants were randomly assigned to either unaccented or accented conditions. We then created a list of participants assigned to each condition. For the

next two interactions, participants remained in the condition to which they were initially assigned. In other words, if a participant was in the "accented" condition in the first wave, they remained in the "accented" condition for the next two waves.

Similar to Study 3A, we told participants that a company named SonicAI had created a voicebot and asked participants to interact freely with it, asking 3 to 10 questions about whatever was on their minds. Each wave concluded with participants providing a 5-star Amazon-like rating of the voicebot, which served as the dependent measure in each wave.

Results

5-star Rating. We conducted a repeated-measures ANOVA to examine the effects of voicebot type (Unaccented vs. Accented) as a between-subjects variable and wave (Wave 1 vs. Wave 2 vs. Wave 3) as the repeated measure on user ratings. There was a significant main effect of voicebot type ($F(1, 311) = 6.48$, $p = .011$, $\eta_p^2 = .02$), indicating that participants consistently rated the accented voicebot higher than the unaccented voicebot overall across all interactions. However, there was no significant main effect of wave ($F(2, 622) = 1.28$, $p = .270$), suggesting that ratings remained the same, offering no visual or textual cues to support comprehension over repeated interactions. Additionally, the interaction between voicebot type and wave was not significant ($F(2, 622) = 0.08$, $p = .911$), indicating that the effect of accent on engagement did not vary over time. Pairwise comparisons showed that at the first wave, participants rated the accented voicebot higher ($M_{Accented_W1} = 4.61$, $SD = 0.77$) compared to the unaccented voicebot ($M_{Unaccented_W1} = 4.42$, $SD = 0.70$; $F(1, 311) = 4.72$, $p = .031$, $\eta_p^2 = .015$). Similar results were observed for the second wave ($M_{Accented_W2} = 4.66$, $SD = 0.72$ vs. $M_{Unaccented_W2} = 4.49$, $SD = 0.69$; $F(1, 311) = 4.29$, $p = .039$, $\eta_p^2 = .014$) and the third wave ($M_{Accented_W3} = 4.66$, $SD = 0.77$ vs. $M_{Unaccented_W3} = 4.46$, $SD = 0.76$; $F(1, 311) = 5.23$, $p = .023$, $\eta_p^2 = .017$).

Number of interactions. We measured the total number of interactions participants had with the voicebots across all three waves, aggregating them into a single cumulative engagement index, with a higher number of interactions indicating greater engagement with the voicebot (De Freitas et al. 2025). ANOVA revealed that participants interacted more with the accented voicebot than with the unaccented voicebot ($M_{Accented}$ = 14.70, SD = 6.99 vs. $M_{Unaccented}$ = 13.26, SD = 4.92; F(1, 311) = 4.38, p = .037, η_p^2 = .014).

Discussion

Study 3C's findings replicate those of prior studies, but also suggest that the effects of an accented voicebot do not immediately dissipate with repeated exposure. In addition, results suggest that participants engaged more with the accented voicebot than the unaccented voicebot. Greater engagement may be more preferable in some contexts than others—in contexts where quick, efficient exchanges are preferred, such results may give pause. However, in the present context, no clear task was set. Thus, the greater engagement observed may indicate consumers' inherent interest or enjoyment in the interaction. Accordingly, in this context, deeper engagement might offer both richer data for firms and deeper brand or platform involvement by consumers.

However, to ensure that these effects were more likely to be beneficial than damaging, we also explored potential novelty effects in an additional real-interaction study (see Supplementary Materials; ChatGPT model 5.2) using a behavioral engagement measure (willingness to provide an email address for future updates). Here, we directly measured perceived novelty. Results showed that participants were significantly more willing to provide their email address to receive updates for the accented voicebot (OR = 1.99, 95%CI = [1.19, 3.36], p = .009), and this effect remained significant while controlling for self-reported perceived novelty (OR = 1.98, 95%CI = [1.16, 3.42], p = .013). Although perceived novelty independently predicted engagement (OR = 1.58, 95%CI = [1.34, 1.90], p<.001), the

accented voicebot was not perceived as more novel than the unaccented voicebot ($F(1, 394) = 0.05$, $p = .821$, $\eta_p^2 = .000$).

Together, Studies 3A-3C and the supplementary study indicate that the observed effects generalize across interaction structures and are not artifacts of a specific design, accent, language, or measurement approach.

STUDY 4

Study 4 offers further evidence for the hypothesized process via moderation. As shown in Study 2, non-standard accents enhance anthropomorphic perceptions by making human diversity cognitively salient. However, when a voicebot is explicitly framed as a machine through dehumanization cues (i.e., highlighting that the accented voicebot is not truly alive; Mourey, Olson, and Yoon 2017), these cues may suppress the humanizing effect (Mende et al. 2019).

Participants, Design, and Procedure

A total of 400 participants on Prolific ($M_{Age} = 44.11$, 49.1% female) were randomly assigned to a condition in a 2 (Accent type: unaccented, i.e., standard American vs. accented, i.e., Indian) × 2 (Anthropomorphism: Dehumanized, i.e., robotic name, avatar, and no personal pronouns vs. Humanized, i.e., human name, avatar, and personal pronouns) between-subjects design.

The scenario and measures for anthropomorphic perceptions ($\alpha = .97$), the preference to continue with the voicebot, and the scenario used for the humanized conditions were identical to those in Study 2. The key difference was that, whereas Study 2 featured a real voicebot exposure with a male voice named Sam, the present vignette study used a female voice named Ava. However, in the dehumanized conditions, the voicebot appeared with a robotic avatar, a robotic name (i.e., M43), and without the use of a first-person pronoun.

Participants listened to voices with either a standard American accent or an Indian accent. The study concluded with collecting basic demographics.

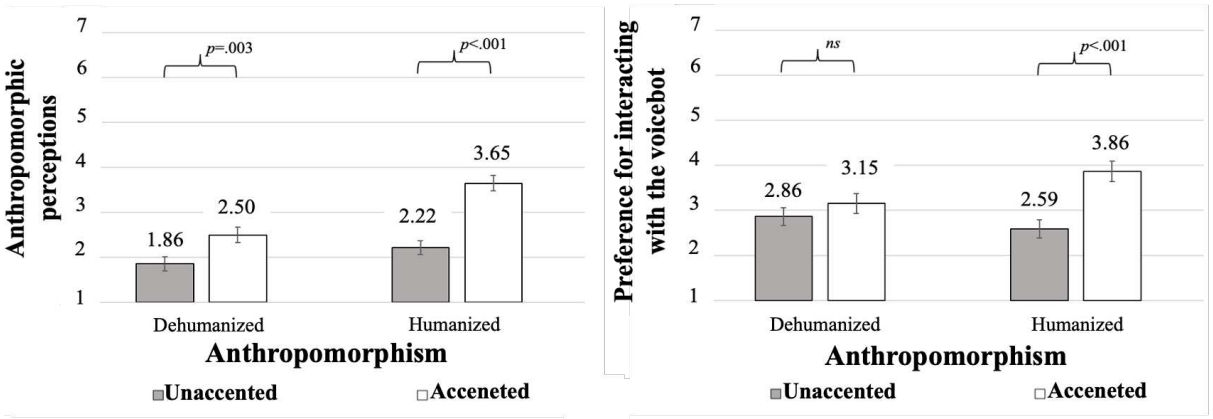
Results

Anthropomorphic perceptions. A 2 (accent type) × 2 (anthropomorphism) ANOVA predicting anthropomorphic perceptions revealed significant main effects of both accent type and anthropomorphism. Results indicate higher anthropomorphic perceptions for accented ($M_{Unaccented}$ = 2.04, SD = 1.41 vs. $M_{Accented}$ = 3.08, SD = 1.68; F(1, 385) = 46.23, p<.001, η_p^2 = .10) and humanized voicebots ($M_{Dehumanized}$ = 2.18, SD = 1.39 vs. $M_{Humanized}$ = 2.94, SD = 1.77; F(1, 385) = 24.71, p<.001, η_p^2 = .06). Results also revealed a significant interaction between the accent type and anthropomorphism (F(1, 385) = 6.73, p<.001, η_p^2 = .017) such that while the non-standard accent increased anthropomorphic perceptions in humanized conditions ($M_{Accented}$ = 3.65, SD = 1.71 vs. $M_{Unaccented}$ = 2.22, SD = 1.52; F(1, 385) = 44.00, p<.001, η_p^2 = .10), the effect was weakened (albeit still statistically significant) when the voicebot signaled dehumanization cues ($M_{Accented}$ = 2.50, SD = 1.44 vs. $M_{Unaccented}$ = 1.86, SD = 1.79; F(1, 385) = 8.86, p = .003, η^2 = .023; Figure 1).

Preference for interacting with the voicebot. A 2 (accent type) × 2 (anthropomorphism) ANOVA predicting preference for continuing with the voicebot (vs. waiting in line to chat with a human agent) replicated the main effect of accent type on participants' preference ($M_{Accented}$ = 3.51, SD = 2.18 vs. $M_{Unaccented}$ = 2.73, SD = 2.00; F(1, 385) = 13.64, p<.001, η_p^2 = .034), but showed no significant main effect of the anthropomorphism factor on the DV ($M_{Dehumanized}$ = 3.01, SD = 2.06 vs. $M_{Humanized}$ = 3.23, SD = 2.20; F(1, 385) = 1.08, p = .301, η_p^2 = .003). Importantly, a significant interaction emerged between the accent type and anthropomorphism (F(1, 385) = 5.22, p = .023, η_p^2 = .013). While the accented voicebot increased participants' engagement across humanized conditions ($M_{Accented}$ = 3.86, SD = 2.24 vs. $M_{Unaccented}$ = 2.59, SD = 1.96; F(1, 385) = 17.82, p<.001, η_p^2 =

.044), this effect was eliminated across dehumanized conditions ($M_{Accented}$ = 3.15, SD = 2.07 vs. $M_{Unaccented}$ = 2.86, SD = 2.04; F<1, p = .321, η_p^2 = .003; Figure 1).

Figure 1. The effect of accent type and anthropomorphism on anthropomorphic perceptions (left) and preference for interacting with the voicebot (right). Error bars represent ±1 SE.



Moderated mediation. A moderated mediation analysis (PROCESS Model 7, 5000 bootstrapped samples; Hayes 2022) showed that the effect of accent type was driven by anthropomorphic perceptions in the humanized conditions (conditional indirect effect = 0.39, SE = 0.07, 95%CI = [0.25, 0.53]). The non-standard accent increased anthropomorphic perceptions of the voicebot (b = 0.71, SE = 0.10, t = 6.63, p<.001), which in turn led to a greater preference to engage with the accented voicebot (b = 0.54, SE = 0.06, t = 8.60, p<.001). Given that the direct effect is not significant (direct effect = 0.11, SE = 0.10, 95%CI = [−0.09, 0.31]), anthropomorphic perceptions fully mediate the effect of accent type on preference for interacting with the accented voicebot as long as the voicebot has not been otherwise dehumanized. However, when the voicebot was dehumanized, the indirect effect size was much smaller than in humanized conditions (conditional indirect effect = 0.17, SE = 0.05, 95%CI = [0.06, 0.29]). Taken together, we thus observe that the indirect effect size when the voicebot could be more easily humanized was significantly larger than when dehumanizing cues were signaled (IMM = 0.21, 95%CI = [0.04, 0.38]).

Discussion

Study 4 provides further evidence for the role of anthropomorphic perceptions in explaining the effect of accents on the preference for interacting with the voicebot. While we replicated the *accent → anthropomorphic perceptions → preference for interacting with the voicebot* chain across humanized conditions, this process chain was undermined when the voicebot was exogenously dehumanized by stripping human-like features. Importantly, even in the presence of dehumanizing cues, the non-standard accent still increased anthropomorphic perceptions of the unaccented voicebot. However, this increase in anthropomorphic perceptions was no longer sufficient to translate into greater engagement with the voicebot, suggesting that accent-based humanization operates beyond such artificial dehumanization cues.

STUDY 5

Study 5 tests whether the accent effect on user engagement persists when the interaction involves an embarrassing situation. Prior research suggests that embarrassing service encounters motivate consumers to dehumanize service providers and attribute more machine-like characteristics to them (Sun et al. 2023). As discussed earlier, such context-driven dehumanization may neutralize the humanizing effect of a non-standard accent. Accordingly, we expect that embarrassing interactions will attenuate the positive effect of a non-standard accent on engagement by weakening the accent's effect on anthropomorphic perceptions, thereby weakening the indirect effect on engagement.

Participants, Design, and Procedure

A total of 800 participants on Prolific (M_{Age} = 41.78, 51.9% female) were randomly assigned to a condition in a 2 (Accent type: unaccented, i.e., standard American vs. accented, i.e., Indian) × 2 (Embarrassment: Control vs. Present) between-subjects design.

Participants imagined purchasing either a non-embarrassing medicine (e.g., hay fever medication) or one likely to activate embarrassment (e.g., antidiarrheal; Jin, Walker, and Reczek 2024). Participants read: "Imagine you are sick and have diarrhea [hay fever]. So, you decide to visit drugstore.com to buy medicine for your condition. While searching through the products, the drugstore's AI voicebot, named Ava, pops up. On the next page, you will begin interacting with Ava." Participants then listened to the voicebot's opening message with either a standard American accent or an Indian accent. Then, participants indicated their preference to continue interacting with the voicebot and anthropomorphic perceptions (α = .97) using the same measures as Study 2. The study concluded with collecting basic demographic information.

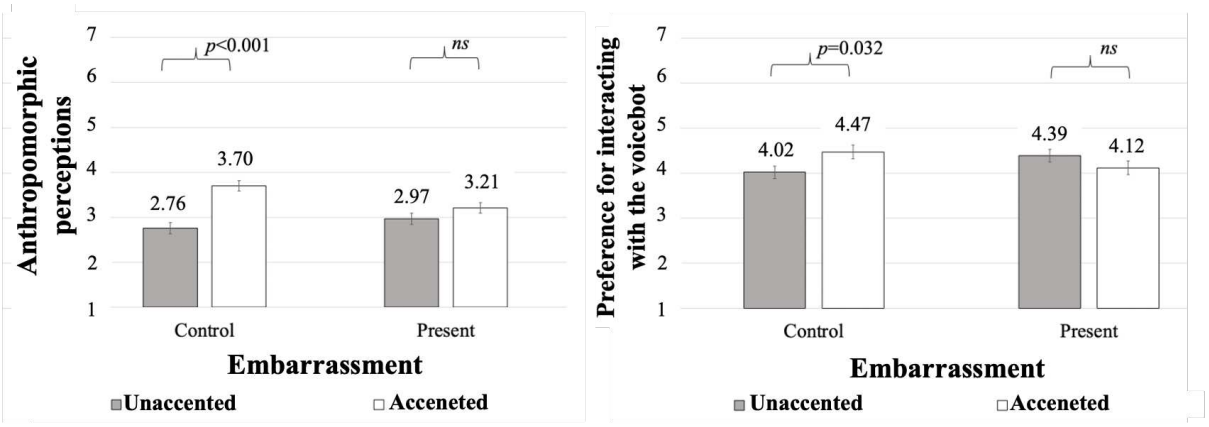
Results

Anthropomorphic perceptions. A 2 (accent type) × 2 (embarrassment) ANOVA tested participants' anthropomorphic perceptions. The main effect of accent type on anthropomorphic perceptions was significant ($M_{Unaccented}$ = 3.45, SD = 1.72 vs. $M_{Accented}$ = 2.86, SD = 1.70; F(1, 793) = 23.84, p<.001, η_p^2 = .029). However, there was no significant main effect of embarrassment on anthropomorphic perceptions ($M_{Control}$ = 3.23, SD = 1.76 vs. $M_{Embarrassment}$ = 3.09, SD = 1.71; F(1, 793) = 1.33, p = .249, η_p^2 = .002). We found a significant interaction between the accent type and embarrassment (F(1, 793) = 8.52, p = .004, η_p^2 = .011) such that while the non-standard accent increases anthropomorphic perceptions in control conditions ($M_{Accented}$ = 3.70, SD = 1.77 vs. $M_{Unaccented}$ = 2.76, SD = 1.61, F(1, 793) = 30.47, p<.001, η_p^2 = .037), the effect became non-significant when the interaction with the voicebot was about an embarrassing topic ($M_{Accented}$ = 3.21, SD = 1.63 vs. $M_{Unaccented}$ = 2.97, SD = 1.78, F(1, 793) = 1.92, p = .166, η^2 = .002; Figure 2).

Preference for interacting with the voicebot. A 2 (accent type) × 2 (embarrassment) ANOVA tested participants' preference for continuing to interact with the voicebot. Results

revealed no main effect of accent type ($M_{Accented}$ = 3.51, SD = 2.18 vs. $M_{Unccented}$ = 4.20, SD = 2.07; F<1 , p = .527, η_p^2 = .001), or embarrassment ($M_{Control}$ = 4.25, SD = 2.12 vs. $M_{Embarrassment}$ = 4.25, SD = 2.05; F<1, p = .941, η_p^2= .000) on participants' engagement. There was a significant interaction between the accent type and embarrassment (F(1, 794) = 5.79, p = .016, η_p^2 = .007) such that while the accented voicebot increased participants' engagement across control conditions ($M_{Accented}$ = 4.47, SD = 2.02 vs. $M_{Unaccented}$ = 4.02, SD = 2.19; F(1, 794) = 4.63, p = .032, η_p^2 = .006) the effect was mitigated in the embarrassment conditions ($M_{Accented}$ = 4.12, SD = 2.14 vs. $M_{Unaccented}$ = 4.39, SD = 1.94; F(1, 794) = 1.57, p = .217, η_p^2= .002; Figure 2).

Figure 2. The effect of accent type and embarrassment on anthropomorphic perceptions (left) and preference for interacting with the voicebot (right). Error bars represent ±1 SE.



Moderated mediation. Moderated mediation analyses (PROCESS Model 7, 5000 bootstrapped samples; Hayes 2022) showed that the effect of accent type was driven by anthropomorphic perceptions across the control conditions (conditional indirect effect = 0.28, SE = 0.05, 95%CI = [0.17, 0.39]). Non-standard accents increased anthropomorphic perceptions for the voicebot (b = 0.47, SE = 0.08, t = 5.52, p<.001), which in turn led to a higher preference to engage with the accented voicebot (b = 0.60, SE = 0.03, t = 15.93, p<.001). The direct effect was also significant (direct effect = −0.13, SE = 0.06, 95%CI = [−0.26, −0.005]). These results replicate the findings of previous studies across the control conditions. However, across the embarrassing situation conditions, the indirect effect was not

significant (conditional indirect effect = 0.07, SE = 0.05, 95%CI = [−0.02, 0.17]). The sizes of the two indirect effects were significantly different (IMM = −0.21, 95%CI = [−0.35, −0.06]), suggesting that the previously observed relationship between accent, anthropomorphic perceptions, and preferences did not emerge in embarrassing interactions.

Discussion

Study 5 suggests that, consistent with our theory, the impact of an accented voicebot on users' preference for engagement depends on the interaction's content (i.e., embarrassment).[3] Consistent with prior research, we observe that the positive anthropomorphization effects for accented voicebots did not emerge when the interaction with the voicebot related to an embarrassing topic.[4]

STUDY 6

Study 6 has two main objectives. First, prior work suggests that individuals who are more liberal will respond more positively to salient diversity than those who are more conservative (Graham, Haidt, and Nosek 2009; Jost et al. 2003). Thus, we expect that more liberal (more conservative) individuals will respond positively (negatively) to the diversity salience triggered by accented voicebots in terms of downstream outcomes. Second, we test these effects using another behavioral measure of engagement (i.e., willingness to purchase a one-month subscription to the voicebot).

Participants, Design, and Procedure

[3] One might argue that embarrassment should reduce overall anthropomorphism, given that embarrassing contexts motivate dehumanization. However, the unaccented American voicebot is already perceived as machine-like, leaving little room for further dehumanization. Thus, the null main effect of embarrassment on anthropomorphic perceptions reflects a floor effect that is directionally consistent with our theory.

[4] Embarrassment may selectively affect different facets of anthropomorphism (Jin, Walker, and Reczek 2024). Depending on whether agency or experience is the dominant component driving outcomes, moderation may plausibly occur at different stages of the process; our data are consistent with moderation on the a-path.

A total of 200 participants on Prolific (M_{Age} = 43.25, 47.5% female) were randomly assigned to one of two between-subjects voicebot accent conditions: Unaccented (standard American) vs. Accented (Indian).

Participants interacted with Nina, a custom-built ChatGPT-based voicebot (identical in voice configuration to Study 3A; Model 4o-mini) that generated real-time responses through the ChatGPT API. The scenario and procedure were identical to Study 3A, and the accent manipulation was embedded in Nina's synthesized speech, which differed only in accent across conditions.

After completing the interaction, participants indicated how much of an additional £20 bonus they would be willing to pay to buy a monthly subscription to Nina, should they be the one winning it in a lottery. Willingness To Pay was captured on a 0–20 response scale, presented as discrete clickable options in £1 increments (e.g., £0–£20). Political ideology was measured using a three-item scale (Huddy, Mason, and Aaroe 2015; 1 = "Definitely conservative," 7 = "Definitely liberal"; α = .99) assessing which political ideology better reflects how participants think about themselves, better describes them, and is more important to them.

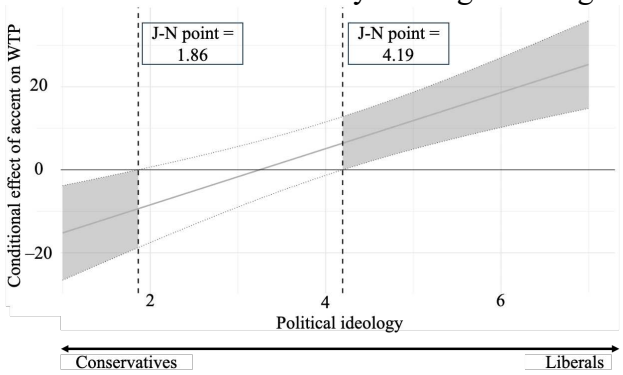
Results

WTP. An ANOVA revealed that participants reported greater willingness-to-pay for the accented voicebot ($M_{Accented}$ = 22.00, SD = 25.39) compared to the unaccented one ($M_{Unaccented}$ = 14.90, SD = 22.18; F(1, 198) = 4.44, p = .036, η_p^2 = .02).

Beyond this main effect, a moderation analysis (PROCESS Model 1, 5000 bootstrapped samples; Hayes 2022) revealed a significant interaction between accent and political ideology (b = 6.77, SE = 1.51, t = 4.48, p<.001). Floodlight analyses (Spiller et al. 2013) indicated that among more conservative participants (ideology scores below 1.86; 19.0% of the sample), participants significantly preferred the unaccented voicebot, whereas

among more liberal participants (ideology scores above 4.19; 50.5% of the sample), participants significantly preferred the accented voicebot (Figure 3).

Figure 3. Conditional effect of accent on WTP at different values of political ideology. The grey area indicates the Johnson-Neyman region of significance.



Discussion

Study 6 identifies political ideology as a boundary condition for the positive effect of non-standard accents. Willingness to pay for the accented voicebot was higher overall. This pattern is consistent with the cognitive enrichment pathway, in which accented voicebots increase the salience of diversity and, in turn, anthropomorphic perceptions. However, as participants' alignment with liberal perspectives rose, we observed a stronger preference for the accented voicebot, consistent with liberals' stronger alignment with DEI (diversity, equity, and inclusion)-related views (Irmak, Murdock, and Kanuri 2020). By contrast, strongly conservative participants preferred the unaccented voicebot.

GENERAL DISCUSSION

Across a range of operationalizations, contexts, experiences, and outcomes, we demonstrate that voicebots using non-standard accents are perceived as more human-like. We argue that such accents make human diversity cognitively salient to users, thereby increasing anthropomorphic perceptions and user engagement with the voicebot. In contrast, standard accents, often perceived as "machine language" (Johnson et al. 2022), do not make human diversity salient. However, this effect disappears when the voicebot exhibits dehumanizing

cues and when the interaction involves embarrassing topics. Moreover, the effect reverses among users whose ideological orientations are unfavorable toward diversity cues, suggesting that accent-based anthropomorphism depends on receptivity to diversity-related signals. We also document other operational and individual-difference boundary conditions. The accent effect attenuates when human agents are readily available (short wait times) and among users with negative attitudes toward AI, showing that when human agents are accessible, or users hold unfavorable attitudes toward AI, the engagement advantage of non-standard accents diminishes.

Linguistic and Behavioral Evidence from user-generated prompts

To provide language-as-behavior evidence for our theorizing, we analyzed user-generated prompts from Studies 3C (all three waves), 6, and the real interaction supplementary study reported in the discussion of Study 3C, all conducted in English, and compared American versus Indian accents. Our goals were to test whether non-standard accents reflect the salience of human diversity in users' prompts and whether users interact differently with accented (i.e., anthropomorphized) voicebots.

Lexical evidence. We first examined lexical indicators of diversity salience using the *Ethnicity* category in LIWC 2022 (Boyd et al. 2022), which captures words referring to national, regional, linguistic, ethnic, or racial identities (e.g., "Indian," "race," "culture"). Because many prompts contained no ethnicity-related words and the variable's variance exceeded its mean (Variance = 0.71; Mean = 0.22), we estimated a negative binomial regression model. Participants in the accented condition used a significantly higher proportion of ethnicity-related words ($M_{Accented}$ = 0.270) than those in the unaccented condition ($M_{Unaccented}$ = 0.180; IRR = 1.51, p = .007), indicating a roughly 50% increase in the likelihood of referencing ethnicity-related concepts. This pattern supports our claim that non-standard accents heighten the salience of diversity.

Semantic ratings. To complement the dictionary-based approach, we conducted semantic ratings of each prompt using the most advanced ChatGPT API model available at the time of analysis (model 5.2; February 2026). Each user prompt was rated on a 1–7 scale for explicit references to social identity categories (e.g., race, culture, gender, religion), where 1 = no reference, 4 = moderate social identity-related content, and 7 = very strong social identity-related content. Prompts in the accented condition were rated higher in diversity salience ($M_{Accented}$ = 1.45, SD = 1.15) than those in the unaccented condition ($M_{Unaccented}$ = 1.30, SD = 0.90; F(1, 974) = 5.31, p = .021, η_p^2 = .005), providing convergent evidence that accent subtly increases references to diversity-related concepts.

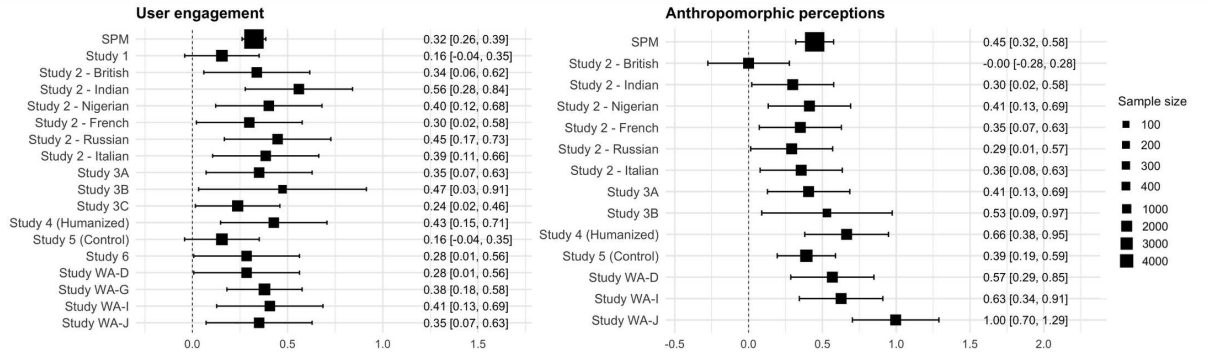
Personal questions. Finally, we examined the interpersonal style of engagement by coding whether users asked personal questions during their interaction with the voicebot. Prior research suggests that when technologies are anthropomorphized, users are more likely to perceive them as social partners (e.g., Waytz et al. 2010). As a result, interactions with anthropomorphized agents tend to become more interpersonal, with users disclosing more personal information and asking questions related to identity, preferences, feelings, or experiences (Lefkeli, Karataş, and Gürhan-Canli 2024). Using the ChatGPT API (model 5.2; February 2026), the most advanced model available at the time of analysis, each user prompt was rated on a 1–7 scale for the extent to which it contained personal questions directed at the voicebot (1 = none, 4 = some personal questions, 7 = multiple or intimate personal questions). Results show that participants interacting with the accented voicebot asked significantly more personal questions ($M_{Accented}$ = 2.71, SD = 1.90) than those in the unaccented condition ($M_{Unaccented}$ = 2.41, SD = 1.68; F(1, 974) = 6.90, p = .008, η_p^2 = .007).

Taken together, these analyses provide further behavioral corroboration of our theory. Non-standard accents increase diversity-related language, elevate the salience of semantic diversity, and elicit greater interpersonal engagement.

Single Paper Meta-analysis

To estimate the overall magnitude of accent effects across our empirical investigations, we conducted a Single Paper Meta-analysis (SPM; McShane and Böckenholt 2017) synthesizing results from 8 studies reported in the manuscript and 4 studies reported in the supplementary materials in comparing accented versus unaccented voicebots on anthropomorphic perceptions and downstream outcomes. This approach was particularly appropriate given that the studies were conducted at different times, in varied contexts, employed distinct voices and various non-standard accents, and involved different interaction formats (i.e., vignette-based vs. interactive designs), allowing us to derive a more integrated estimate of the effect of accent on anthropomorphic perceptions and downstream outcomes. Across studies, this yielded 13 effect sizes for anthropomorphic perceptions and 17 for user engagement. Using the *metafor* package in R, we computed standardized mean differences (Cohen's d) based on group-level statistics. Results showed a significant positive effect of accent on user engagement (b = 0.32, z = 10.18, p<.001, 95%CI = [0.26, 0.39]) and anthropomorphic perceptions (b = 0.45, z = 6.84, p<.001, 95%CI = [0.32, 0.58]), indicating consistent effects across studies (Figure 4).

Figure 4. Forest plots of non-standard accents effect in 14 experiments on user engagement (left) and anthropomorphic perceptions (right).



Theoretical Contributions

The current research makes several theoretical contributions. First, we contribute to the accentism literature by exploring this bias in human-AI interactions and demonstrating

that discrimination associated with non-standard accents does not hold when a non-human entity speaks with a non-standard accent (Romero-Rivas, Morgan, and Collier 2022; Wang et al. 2013; Zewail et al. 2025). Our findings indicate that non-standard accents increase anthropomorphic perceptions of the voicebot because they make human diversity cognitively salient. As such, we also contribute to prior theories of diversity (Arsel, Crockett, and Scott 2022; Giles, Wilson, and Conway 1981), positioning it as a distinctly human characteristic that may carry broader implications when made salient.

Second, we extend prior research on anthropomorphizing AI (Bergner, Hildebrand, and Häubl 2023; Schanke, Burtch, and Ray 2021; Yalcin et al. 2022), which has primarily focused on text-based agents, by introducing a novel perspective that emphasizes the central role of *voice* in shaping human perceptions of voicebots. Given that AI agents are hybrid beings that evoke both the feeling of interacting with a humanlike entity and the feeling of interacting with a machine entity (Gelbrich et al. 2025), voicebots, as spoken agents and innovative service technologies (Zierau et al. 2023), are a particularly fitting context for studying anthropomorphism, as they incorporate an additional layer of communication through paralinguistic cues such as tone and pitch (Hildebrand and Bergner 2021).

Third, our work extends prior work on the boundaries and limits of AI anthropomorphism approaches (Crolic et al. 2022; Garvey, Kim, and Duhachek 2022; Pizzi, Scarpi, and Pantano 2021). Study 5 examines the role of situational elements (e.g., embarrassment interaction), two other studies explore operational factors (e.g., wait times; Supplementary Materials) and individual differences (e.g., attitude toward AI; Supplementary Materials; Tully, Longoni, and Appel 2025; Garvey, Kim, and Duhachek 2022; Crolic et al. 2022). Specifically, we show that a tactical intervention such as using a non-standard accent may not be strong enough to overcome the potency of said factors.

Fourth, we contribute to research on political ideology and marketplace behavior by showing that ideological orientation shapes responses to diversity cues in AI contexts. Conservatives and liberals differ systematically in their worldviews (Graham, Haidt, and Nosek 2009; Jost et al. 2003), and these differences influence marketplace preferences (Kim, Park, and Dubois 2018). Extending this work, we demonstrate that diversity salience triggered by accented vocal cues produces divergent reactions such that individuals less receptive to diversity cues respond less favorably, whereas those more aligned with inclusionary values exhibit more favorable reactions.

Managerial Implications

The current work offers several managerial implications as well. First, we examine a managerially relevant outcome for voicebots used in customer service settings by studying whether users choose to continue interacting with the voicebot. This is usually the main objective of companies using AI agents in the hopes of reducing wait times (Davenport et al. 2020), cutting personnel costs (Castelo et al. 2023), increasing the quality of service (Stefanowicz 2023), and helping companies avoid complex training for their agents (Packard and Berger 2021). We introduce a novel, easily implementable approach to anthropomorphizing voicebots by embedding a non-standard accent and demonstrate that doing so increases users' willingness to continue engaging with an accented voicebot (vs. an unaccented one).

Second, our findings extend beyond traditional customer-service settings to voicebots deployed in decision-support and open-ended interaction contexts. In Study 3B, for example, we demonstrate that when a voicebot provides health-related recommendations, users are more willing to financially support the technology when it communicates with a non-standard accent. Our results show that users are more likely to support AI startups developing the voicebot, opt in to receive future updates about the voicebot, and express greater WTP for a

monthly subscription when interacting with accented (vs. unaccented) voicebots. For firms launching new AI products, subtle vocal design choices such as accent may influence users' willingness to financially back, subscribe to, or remain engaged with their AI products.

Third, our work identifies important boundary conditions of the accent effects. Companies operating in domains involving embarrassing products, such as certain medications, sexual health items, or loan rejection communications (Dahl, Manchanda, and Argo 2001; Esmark Jones, Barney, and Farmer 2018; Holthöwer and van Doorn 2023), may not benefit from using accented voicebots. Moreover, user characteristics matter. Our findings indicate that the effectiveness of non-standard accents depends on users' receptivity to diversity cues. Companies may benefit more from accented voicebots when targeting audiences who hold positive attitudes toward diversity. Conversely, among users less receptive toward diversity-related cues, the intervention may not work or even backfire. These results underscore the importance of audience segmentation when implementing vocal design strategies.

Implications for Society at Large

Past research suggests that even a brief exposure to a non-standard accent can diminish discrimination against people who have that accent (Boduch-Grabka and Lev-Ari 2021). Therefore, both businesses and policymakers may encourage the adoption of accented voicebots to make consumers more receptive to different accents and potentially reduce discrimination against accented speakers. Moreover, incorporating a non-standard accent in the voicebot fosters inclusivity, challenges unconscious biases, and enhances linguistic tolerance, reflecting society's rich diversity.

Our findings may also have broader implications for how users perceive AI systems. Prior work suggests that people often categorize artificial agents as belonging to a different "species" (Schmitt 2020). Although our studies do not directly test reductions in speciesism,

our results suggest that design features that highlight human-related variation (e.g., accented speech) may make AI agents appear less homogeneous and more socially human-like. This possibility is particularly consequential given that a growing body of research shows that interactions with AI systems can spill over to shape subsequent human–human judgments and behaviors (e.g., prosociality; Zhang et al. 2025). Future research may therefore examine whether repeated exposure to such design cues can gradually soften rigid human–machine category boundaries and reduce speciesist thinking.

Limitations and Avenues for Future Research

Our research opens several other avenues for future investigation. First, one important avenue for future research concerns the broader conceptualization of diversity salience. In the current work, we operationalize diversity salience using measures that directly assess thoughts about human variation and social identity, and we manipulate it via exogenous pre-interaction text. While these approaches capture the theorized mechanism, future research could examine whether exposure to a non-standard accent activates a more generalized diversity mindset that extends beyond social categories. For instance, accent exposure may heighten sensitivity to variation in other domains, such as flavor profiles, design patterns, idea generation, or creative alternatives. Testing whether accented voicebots increase openness to diversity across non-social contexts would provide stronger evidence that diversity salience reflects a broader cognitive orientation rather than a domain-specific activation. Exploring this possibility represents a promising direction for future research.

Second, we used standard English as our baseline language, drawing on conventions in prior work on accented language (Wang et al. 2013; Bourdin, Sichtmann, and Davvetas 2024) and confirmed the effect in French. Future research could examine a wider range of languages to strengthen the generalizability of this claim to other cultural contexts where non-standard accents may be perceived and received differently.

Third, future research should examine the role of accent familiarity in shaping anthropomorphic perceptions. In Study 2, the British accent failed to increase anthropomorphic perceptions. Therefore, more familiar accents (e.g., regional accents) may activate stronger human-related schemas and make human variation more salient. Conversely, exposure to non-existent or fabricated accents may disrupt these schema activations. The absence of a recognizable accent category may prevent listeners from making immediate social inferences, thereby weakening the salience of diversity. Alternatively, the cognitive effort required to interpret unfamiliar accents may heighten the salience of diversity and increase anthropomorphic perceptions. Investigating both strongly familiar regional accents and fictional accents would help clarify whether anthropomorphic responses are driven by familiarity or cognitive effort. The insights could also encourage text-to-speech (TTS) developers to expand their training datasets to include regional accents, thereby fostering more socially resonant, inclusive, and relatable voice-based interfaces.

Fourth, future research could investigate how non-native accents (e.g., Arabic or Chinese) influence perceptions of voicebots. Such accents may reduce semantic fluency (Kato and Baese-Berk 2024), thereby lowering engagement with voicebots (Zierau et al. 2023). However, accents can also trigger rapid social categorization even without fully processing the message or experiencing fluency (Pietraszewski and Schwartz 2014). Understanding how semantic fluency and identity cues jointly shape anthropomorphic perceptions remains an important direction for future research.

Fifth, future work could examine whether accent-based anthropomorphism extends to other spoken AI systems. Unlike multimodal interfaces, callbots and Interactive Voice Response (IVR) systems rely solely on auditory communication (Anne 2023), offering no visual or textual cues to support comprehension. Investigating whether accent alone can drive

engagement in purely voice-based systems, thereby clarifying the limits of accent-driven anthropomorphism.

Finally, and as a last note, both voicebot technology and the baseline salience of diversity may evolve over time. Our work provides a framework for predicting that if voicebot technology no longer makes human diversity salient, it may become less effective as an anthropomorphizing experience, suggesting that firms seeking to offer human-like interactions may need to find alternative ways to enhance this aspect of their technology.

Chapter 3

Cuddly, Not Creepy: Petmorphic AI Companions Deepen Emotional Connection and Enhance User Engagement

Hesam Ghasemi
ESSEC Business School

Amir Sepehri
ESSEC Business School

Christian Hildebrand
University of St.Gallen

Cait Lamberton
University of Pennsylvania

Abstract – Consumers increasingly turn to generative artificial intelligence as a new form of companion (GenAI companion), yet often hesitate to engage with these systems. This creates a challenge for firms seeking to design AI companions that foster user engagement. Bridging humans' longstanding relationships with pets and their emerging relationships with AI, this research introduces *petmorphism*, the design of AI companions using pet-like cues (pet name, pet avatar, and pet-like language), and shows that such cues increase user engagement relative to both anthropomorphized and non-anthropomorphized designs. We propose that petmorphism activates pet-related schemas, leading users to perceive AI companions as nonjudgmental and emotionally safe interaction partners. This perception strengthens emotional connection, which in turn enhances users' evaluations and behavioral engagement. Importantly, these effects are stronger when users pursue affiliative rather than instrumental goals. Together, this work identifies petmorphism as a novel and scalable design strategy that fosters emotionally resonant human-AI relationships.

Keywords. artificial intelligence, generative AI, AI companion, petmorphism

Pets have become an increasingly vital part of our lives, providing emotional support, companionship, and a sense of family for millions (Michelle and Korrena 2025). Today, pet ownership has reached record highs, with 66% of U.S. households, nearly 87 million homes, owning at least one pet (APPA 2025). Dogs and cats, in particular, have become integral companions, with 97% of pet owners viewing their pets as family members, and more than half (51%) considering them equal to human family members (Michelle and Korrena 2025). These numbers reflect more than just practical care. They point to the powerful emotional bonds people form with their pets (Amiot, Sukhanova, and Bastian 2020; Sanders 1990; Zasloff and Kidd 1994). Over thousands of years, selective breeding has amplified these traits in pets like dogs and cats, deepening the emotional bonds humans form with them (Hirschman 1994). However, despite their importance, not everyone can experience the joy and the companionship pets bring due to lifestyle constraints, financial considerations, or housing limitations (Cavanaugh, Leonard, and Scammon 2008; Wada and Shibata 2007).

As a result, some individuals increasingly turn to technological alternatives that promise emotional companionship without the financial, logistical, and long-term demands of pet ownership. Robotic pets such as PARO, Sony's AIBO, Tombot's therapeutic dog, and Unitree's Go1 illustrate this shift. For example, Unitree's Go1 robodog mimics canine movement and autonomously follows its owner, inviting pet-like engagement through responsive, socially contingent behavior (Unitree 2024). These systems are already being used in care homes and therapeutic settings, where animatronic cats and dogs have been described as "almost magical" in reducing distress and fostering comfort (The Guardian 2023). Beyond embodied robots, similar dynamics are emerging in digital environments. Pet-themed chatbots—such as dog-like AI assistants or companion platforms like Talkie and Character.ai—simulate playful interaction, emotional responsiveness, and relational warmth. Even customer-service interfaces adopt pet-inspired mascots; for example, Fido's "Ask

Jack," a Canadian telecom chatbot, incorporates a dog icon to signal approachability and trust. Across both physical and digital contexts, pet-based design elements are increasingly used to evoke familiarity, comfort, and emotional engagement. Taken together, it is worth asking: Could the same principles that underlie human–pet bonds inform how we design emotionally resonant technologies? In particular, as AI companions become more common, pet–human relationships may offer a powerful, yet underexplored, template for fostering humans' connections with technology.

With recent advancements in Generative Artificial Intelligence (GenAI), AI-based technologies have become integral to daily life. AI companions–Synthetic partners providing emotional support, distinct from broader relational AI like customer service bots, educational tutors, AI assistants (e.g., ChatGPT), or voice assistants (e.g., Siri; De Freitas et al. 2025)–have evolved beyond mere task completion, such as coding or problem-solving, to also provide emotional support and companionship (De Freitas et al. 2025; De Freitas et al. 2023; Yin, Jia, and Wakslak 2024). Today, people increasingly engage with AI companions for a wide range of tasks, including personal conversations about their daily lives, seeking advice on personal relationships, discussing personal anxieties, or simply engaging in casual conversation for entertainment or companionship (Milana et al. 2024; Vowels, Francois-Walcott, and Darwiche 2024). These tools are no longer niche: globally, over 100 million people now interact with personified AI chatbots—such as Snapchat's My AI (with more than 150 million users), Replika (with more than 25 million users), and Xiaoice (with more than 660 million users)—as mentors, friends, or confidants (Bernardi 2025). A recent survey showed that 73% of AI companion users report using them for anxiety management, 63 % for personal advice, and 35 % to feel less lonely (SENTIO University 2025), with the hope of receiving emotional support (Vertu 2025).

Unlike these newly developed AI companions, animals have served as trusted, non-judgmental companions to humans for thousands of years (de Waal 2009). This raises a compelling question: can our new AI-based companions begin to mimic the behaviors of our old feline ones—and in doing so, foster deeper emotional connections, increase engagement, and yield similar psychological benefits, such as reduced stress, improved mood, and enhanced well-being (Brown 2004; Cavanaugh, Leonard, and Scammon 2008; Gullone 2001; Herzog 2011)?

In this paper, we introduce a novel concept—*petmorphism*—which refers to designing AI companions with pet-like cues, such as animal avatars, pet-themed names, and communication styles infused with animal-related expressions (e.g., playful vocabulary such as pawsome; Ruchi 2022). When an AI companion is petmorphized, users are prompted to activate pet-related mental schemas. These schemas carry deeply ingrained cultural and psychological associations: pets are perceived as loyal, affectionate, emotionally responsive, and– critically –nonjudgmental companions (Gullone 2001; Holbrook et al. 2001). Consequently, activating the pet schema leverages both the inherently nonjudgmental nature of pets and the strong emotional bonds people form with them, thereby reducing anticipatory evaluation concerns and enhancing perceived psychological safety. This heightened sense of safety, together with the deep emotional attachment people typically develop toward their pets, fosters stronger emotional connection and increaes petmorphic AI companions.

We offer two key contributions. First, we contribute to research on user engagement with AI companions (e.g., De Freitas et al. 2025; Oğuz-Uğuralp et al. 2025; Yin, Jia, and Wakslak 2024; Skjuve et al. 2021) by integrating insights from psychology (schema theory), animal studies (human–pet relationships), and Information Systems (conversational AI) to inform the design of AI companions. Although AI companions are increasingly used for emotional support, many users remain hesitant to engage with them (Oğuz-Uğuralp et al.

2025), particularly when anthropomorphic designs trigger feelings of social judgment or emotional discomfort (Skjuve et al. 2021). By proposing petmorphism as a design strategy that fosters emotional connections—similar to those between humans and their pets—we offer a novel pathway to reduce user resistance and increase acceptance of AI companions. This approach is particularly valuable in light of growing evidence that AI companions can reduce loneliness (De Freitas et al. 2025), enhance feelings of being heard (Yin, Jia, and Wakslak 2024), and provide therapeutic benefits (de Gennaro, Krumhuber, and Lucas 2020). By increasing engagement through petmorphic design, AI companions may become more emotionally resonant and accessible, ultimately amplifying their psychological value and well-being outcomes in users' everyday lives.

Second, we introduce and theorize petmorphism. Unlike anthropomorphism, which encourages users to treat AI like people (e.g., Bergner, Hildebrand, and Häubl 2023; Crolic et al. 2022; Schanke, Burtch, and Ray 2021; Yalcin et al. 2022), petmorphism encourages users to treat AI more like beloved pets. Petmorphic AI companions offer a novel, low-risk alternative that may replicate many of the emotional benefits traditionally associated with animal interaction—including psychological (e.g., stress relief), physiological (e.g., improved vital signs), and social (e.g., increased communication) effects (Brown 2004; Cavanaugh, Leonard, and Scammon 2008; Gullone 2001)—while avoiding the practical limitations of real pet ownership (Wada and Shibata 2007). Notably, this approach may be particularly valuable for individuals who are unable to own pets due to allergies, housing restrictions, or lifestyle constraints, offering a new form of emotionally supportive companionship through technology.

THEORETICAL BACKGROUND

Pets as enduring anthropomorphized social companions

Animals have accompanied humans for thousands of years, evolving from working partners to deeply embedded social companions (Serpell 2009). Humans rely on animals not only for food, clothing, and scientific advancement, but also for emotional and physical support in sickness and in health (Amiot and Bastian 2015; Serpell 2009).

Among all animals, companion pets occupy a uniquely intimate psychological position. Companion animals are particularly likely to be anthropomorphized, as owners attribute humanlike traits, intentions, and personalities to them and incorporate them into their social identity (Amiot, Sukhanova, and Bastian 2020), a tendency that is further strengthened by the selective breeding of domesticated species such as dogs and cats for humanlike social cues—expressive faces, responsiveness, and attachment behaviors—which facilitate attributions of mental states and emotions (Serpell 2003). As a result, anthropomorphism is especially pronounced toward pets compared to non-domesticated species.

Pets are widely perceived as having personalities and emotional capacities (Gosling, Kwan, and John 2003). Owners routinely describe their pets using human personality traits and view them as intentional, feeling beings. Importantly, however, pets differ from human social partners in a critical way: they are perceived as nonjudgmental (Gullone 2001). Unlike human counterparts, pets are not seen as evaluative, critical, or socially threatening. Instead, they are understood as offering unconditional regard and emotional presence without moral appraisal or reputational consequences (Beck and Katcher 1996; Gullone 2001). This perception of psychological safety is central to the pet–human bond.

The emotional benefits of pet relationships are well-documented. Interaction with animals reduces loneliness (Black 2012; Stanley et al. 2014), depression (Krause-Parello 2012), and stress (Antonacopoulos and Pychyl 2008; Wu, Wong, and Chu 2018). Many individuals report feeling closer to their pets than to certain human family members (Barker

1999), and pets are often incorporated into the extended self (Belk 1988; Brown 2004; Sanders 1990). They are described as friends, family members, or even part of one's identity (Hirschman 1994). Such bonds provide affirmation, comfort, and stability, and are sometimes perceived as more secure than romantic partnerships (Beck and Madresh 2008).

Crucially, pets do not merely substitute for human relationships; they complement and enrich them, adding a unique relational dimension characterized by warmth, loyalty, and unconditional companionship (Beck and Katcher 1996). Because pets are perceived as emotionally responsive yet nonjudgmental, and because they function as attachment figures that provide safety and comfort, they occupy a distinct social niche—one defined by acceptance rather than evaluation. This combination of attributed personality, emotional connection, and absence of social threat makes pets powerful relational agents in human psychological life.

AI companions as humans' new social partners

As artificial intelligence systems become increasingly embedded in everyday life, AI companions are emerging as a new form of social partner. Unlike task-oriented AI assistants, AI companions are explicitly designed to provide emotional support, companionship, and relational interaction (De Freitas et al. 2025). These systems simulate conversation in ways that create the perception of empathy and care, even though they do not possess genuine emotional experience (De Freitas, Oguz-Uguralp, and Kaan-Uguralp 2025). Their rapid adoption suggests that many users seek not only instrumental efficiency but also social and emotional fulfillment through technology.

Yet, despite their sophistication, AI companions remain fundamentally limited (Oğuz-Uğuralp et al. 2025). They do not have subjective experiences, consciousness, or genuine feelings. This gap between simulation and authentic emotional capacity can produce user

detachment or discomfort (Han, Yin, and Zhang 2023). In response, developers frequently anthropomorphize chatbots by incorporating human-like features—such as human names, avatars, first-person self-reference, and even artificial conversational delays to mimic "thinking" (Bergner, Hildebrand, and Häubl 2023; Schanke, Burtch, and Ray 2021). These human attributes lead users to apply the same social norms typically reserved for human interaction (Nass and Moon 2000).

Anthropomorphism has well-documented psychological and managerial benefits in marketing (e.g., Aggarwal and McGill 2012; Chen, Sengupta, and Zheng 2023; Wang and Touré-Tillery 2024). It has also been identified as one of the most effective and widely adopted strategies for increasing consumer acceptance of AI technologies, reducing algorithm aversion, and fostering engagement with AI-powered interfaces (Li, Lai, and Wang 2025; Zehnle, Hildebrand, and Valenzuela 2025). When non-human entities are perceived as human-like, they are attributed agency (i.e., the capacity to exert self-control and the extent to which the agent can plan and act) and experience (i.e., the capacity to feel emotions and the extent to which the entity has a personality and consciousness; Epley, Waytz, and Cacioppo 2007; Gray, Gray, and Wegner 2007). Anthropomorphism often increases perceived warmth and competence (Blut et al. 2021). However, increasing anthropomorphism introduces important downsides. First, highly human-like AI entities can evoke the uncanny valley effect (i.e., the tendency for a robot to elicit negative emotional reactions when it closely resembles a human; Mori 1970), whereby near-human resemblance generates unease and eeriness rather than comfort. Second, human-like cues elevate users' expectations regarding the AI's efficacy; when users are in emotional states such as anger, anthropomorphism can backfire (Crolic et al. 2022). Third, anthropomorphism activates social norms associated with interpersonal judgment, which can heighten users' concerns about being scrutinized, judged, or morally evaluated (Holthöwer and van Doorn 2023; Jin, Walker, and Reczek 2024);

therefore, when AI entities appear human-like, they are perceived as judgmental (Skjuve et al. 2021).

Thus, anthropomorphism presents a paradox. While it can increase perceived agency and competence (Blut et al. 2021; Epley, Waytz, and Cacioppo 2007), it may simultaneously undermine perceptions of emotional experience and warmth and heighten perceived judgment, thereby increasing social threat and discomfort. This tension is particularly consequential in emotionally sensitive interactions—contexts in which users turn to AI companions to discuss personal concerns, vulnerability, or distress (De Freitas, et al. 2025), and where both warmth and competence are especially important (De Freitas et al. 2023). Therefore, what is needed is a design approach that reduces perceived judgment, enhances experience and warmth, and preserves competence and agency without triggering social threat.

Petmorphism

To address the paradox of anthropomorphism, we introduce *petmorphism*, defined as the design of AI companions using pet-like cues such as animal avatars, pet-inspired names, and playful language. Just as human-related schemas trigger anthropomorphism (Epley, Waytz, and Cacioppo 2007), leading users to interpret AI through the lens of human social categories, petmorphism activates pet-related schemas, leading users to interpret AI through the psychological framework associated with companion animals. In other words, whereas anthropomorphism encourages users to relate to AI as if it were human, petmorphism encourages users to relate to AI as if it were a companion pet.

According to the Computers Are Social Actors (CASA) framework, when technology contains social cues such as dialogue, turn-taking, and linguistic responsiveness, users automatically apply human social norms to it (Nass and Moon 2000; Nass, Steuer, and

Siminoff 1994). Even minimal cues can lead users to treat computers as social entities, assign them personalities, and include them within their social system (Nass and Brave 2005). Anthropomorphic AI, therefore, becomes subject to expectations typically reserved for human interaction. Petmorphism extends this logic: if human-like cues trigger human schemas, then pet-like cues should trigger pet schemas. Users may therefore respond to petmorphic AI using norms associated with human–pet relationships: non-judgmental emotional bond.

The nonjudgmental quality of pets is especially consequential in the context of AI companions. Many AI interactions involve personal concerns, uncertainty, or emotional vulnerability (Ho, Hancock, and Miner 2018; Volpato, DeBruine, and Stumpf 2025). Unlike human partners, pets are not perceived as evaluative, critical, or socially threatening. They are seen as offering unconditional regard without moral appraisal or reputational consequences (McConnell et al. 2011). Activating this pet schema, therefore, reduces anticipatory evaluation concerns and increases users' willingness to disclose, confide, and engage openly.

Additionally, because pets are culturally positioned as sources of comfort, security, and unconditional acceptance (Archer 1997), pet-related cues guide users to approach the AI not merely as a functional tool, but as a relational partner. Just as individuals often confide in and emotionally bond with their pets, petmorphized AI companions invite similar caregiving and attachment dynamics (Cavanaugh, Leonard, and Scammon 2008). Signals of pet-likeness (e.g., playful avatars, affectionate naming, or dog-like expressions of loyalty) can foster emotional closeness rooted in warmth, companionship, and a sense of acceptance.

In this way, petmorphism produces a qualitatively different form of human–AI interaction. Rather than engaging the system as a purely instrumental agent, users come to relate to it as they would a cherished pet. This shift not only strengthens emotional

connection but also encourages sustained engagement, as users seek the continuity, comfort, and nonjudgmental presence they culturally associate with companion animals.

Goal Orientation in AI Interactions

The effectiveness of petmorphism should depend on users' interaction goals. Users engage with AI companions either for affiliative purposes (e.g., emotional support, companionship, validation) or for instrumental purposes (e.g., problem-solving, task completion, efficiency).

Mind perception theory suggests that social perception operates along two key dimensions: agency (e.g., self-control, judgment, communication, thought, memory) and experience (e.g., hunger, fear, pain, pleasure, consciousness; Gray et al. 2007). Agency aligns closely with competence and task efficacy (Davis and Kim 2025), whereas experience aligns with warmth and emotional responsiveness (Wang et al. 2017). Petmorphism is theorized to increase perceived experience of the AI companion, as pets are commonly seen as capable of feeling hunger, fear, pain, and affection. Consequently, the benefits of petmorphism should be stronger in affiliative contexts, where users seek emotional connection. In such settings, the heightened perception of experience and warmth should translate into stronger emotional bonds and, in turn, greater downstream acceptance of the petmorphic AI companion.

OVERVIEW OF STUDIES

Studies 1A and 1B provide foundational support by showing that animals perceived as more human-like elicit stronger emotional bonds and that companion animals such as dogs and cats are seen as both warm and competent, justifying their use as petmorphic templates. Study 2 tests the core effect in an interactive setting and shows that petmorphic design increases both evaluations and behavioral engagement through emotional connection. Study

3A establishes that petmorphism outperforms both anthropomorphized and non-anthropomorphic designs, while Study 3B shows that its effectiveness depends on users' pet identity. Finally, Study 4 identifies a boundary condition by demonstrating that petmorphism is most effective when users pursue affiliative rather than instrumental goals (Power analyses in the Supplementary Materials).

STUDIES 1A and 1B

The objective of Studies 1A and 1B was to systematically map a broad range of animals on perceived anthropomorphism and emotional connection in order to (a) establish the baseline association between these constructs and (b) select stimuli for subsequent studies.

Study 1A: Participants, Design, and Procedure

A total of 150 participants on Prolific (M_{Age} = 40.31, 48.00% female) were asked to complete an online survey. Participants were randomly shown the names of 10 animals (from a larger pool of 30) presented in random order, without any additional information such as images or species classification. For each animal, participants responded to two items measuring anthropomorphism (adapted from Yalcin et al. 2022): "To what extent do you think [e.g., sloths] have some human-like qualities?" and "To what extent do you think [e.g., sloths] seem like a person?" (1 = "not at all," 7 = "very much"; α = .93). They also rated their emotional connection to each animal using a single composite item (adapted from Jiménez and Voss 2014): "Which of the following best describes your emotional attachment to [e.g., sloths]?" (1 = "No emotional bond, Not emotionally connected, Not linked by feelings, No feelings of attachment"; 7 = "A strong emotional bond, Emotionally connected, Linked by

feelings, Strong feelings of attachment"; α = .98). The survey concluded with basic demographics.

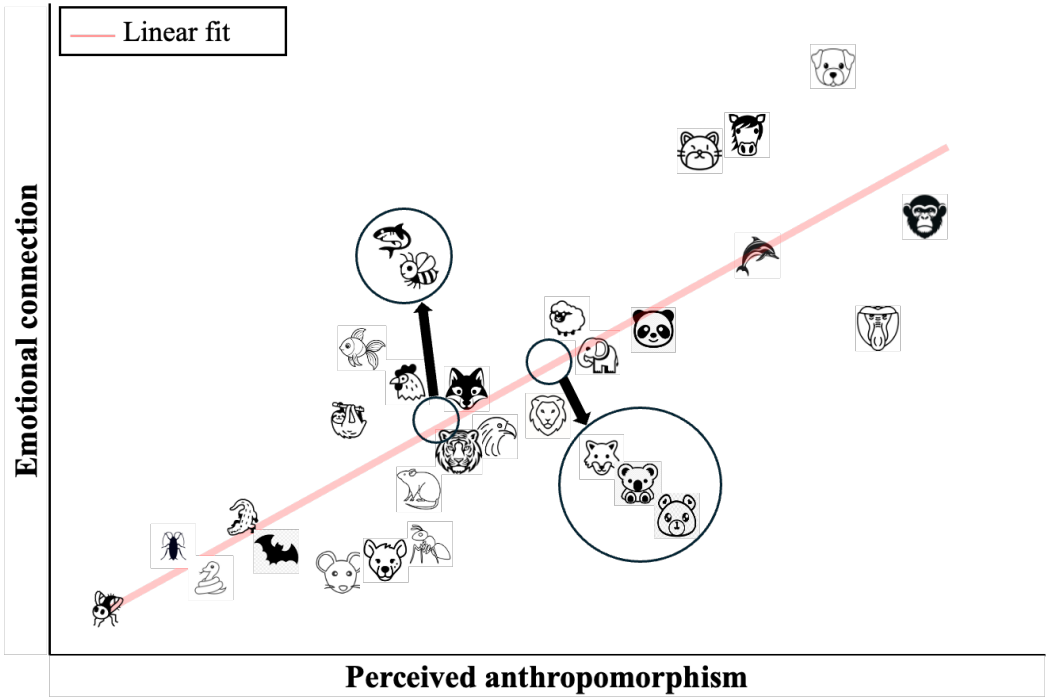
Study 1A: Results

Results show a strong positive correlation between perceived anthropomorphism and emotional connection across animals. Correlation coefficients ranged from r = .460 (horses) to r = .916 (flies), with all correlations statistically significant ($p < .001$). Animals such as crocodiles (r = .882), flies (r = .916), and snakes (r = .801) showed strong associations, suggesting that even low baseline emotional connection can co-occur with strong alignment between anthropomorphic perceptions and attachment. These results support the hypothesis that anthropomorphism is positively associated with emotional connection across a wide range of animal types. See Figure 1 and Table 1 for detailed results.

Table 1. Anthropomorphism and emotional connection with animals

Key	Perceived anthropomorphism	Emotional connection
Chimpanzees	5.10	4.67
Proboscis monkeys	4.84	3.79
Dogs	4.67	5.89
Dolphins	4.01	4.20
Horses	3.99	5.18
Cats	3.86	5.00
Pandas	3.69	3.66
Elephants	3.42	3.52
Sheep	3.07	3.44
Bears	3.04	3.16
Koalas	3.01	3.17
Lions	2.97	2.83
Wolves	2.95	3.17
Eagles	2.82	2.80
Foxes	2.78	3.03
Tigers	2.64	2.67
Bees	2.49	2.68
Sharks	2.48	2.69
Chickens	2.43	3.16
Rats	2.39	2.34
Ants	2.39	2.13
Hyenas	2.33	2.08
Goldfish	2.30	3.24
Sloths	2.28	2.55
Mice	2.24	1.90
Bats	1.97	2.06
Crocodiles	1.93	2.10
Snakes	1.70	1.72
Cockroaches	1.66	1.78
Flies	1.24	1.39

Figure 1. Schematic relationship of perceived anthropomorphism and emotional connection



Study 1B: Participants, Design, and Procedure

A total of 150 participants on Prolific (M_{Age} = 43.00, 35.33% female) completed an online survey. Participants were randomly shown 10 animal names drawn from the same larger pool used in Study 1A, presented in random order without any additional information. For each animal, participants rated perceived warmth and competence using established dimensions from the stereotype content model (Fiske et al. 2002). Warmth was measured with four items assessing the extent to which the animal was perceived as warm, kind, friendly, and sincere. Competence was measured with four items assessing the extent to which the animal was perceived as competent, intelligent, capable, and skillful. Responses were recorded on seven-point scales (1 = "not at all," 7 = "very much"; αs > .96). The survey concluded with basic demographic questions.

Study 1B: Results

Animals varied systematically in perceived warmth and competence (see Figure 2).

Figure 2. Mapping animals on warmth and competence

COMPETENCE \ WARMTH	High	Low
High	Dogs Cats Horses Dolphins Elephants Chimpanzees Foxes Proboscis monkeys	Eagles Sharks Lions Tigers Wolves Bees Bears
Low	Sloths Sheep Pandas Chicken Koalas Goldfish	Snakes Crocodiles Cockroaches Hyenas Ants Flies Bats Mice Rats

Companion animals such as dogs (M_{Warmth} = 5.99, $M_{Competence}$ = 5.78) were rated high on both dimensions. In contrast, predators such as lions (M_{Warmth} = 2.50, $M_{Competence}$ = 5.53) were perceived as highly competent but relatively low in warmth. Insects and reptiles such as

cockroaches (M_{Warmth} = 1.69, $M_{Competence}$ = 3.11) tended to be rated low on warmth and comparatively lower on competence. Finally, gentle herbivores such as pandas (M_{Warmth} = 4.80, $M_{Competence}$ = 4.57) were perceived as relatively warm but lower in competence. The correlation between warmth and competence was moderate (r = .42, p<.001), indicating that although related, the two dimensions remain empirically distinct.

Discussion

Study 1A demonstrates a robust association between perceived anthropomorphism and emotional connection across a variety of animals: the more human-like an animal is perceived to be, the stronger the emotional bond people tend to form with it. The findings also reveal that animals perceived as both highly human-like and emotionally engaging are typically pets or familiar companion animals (Amiot, Sukhanova, and Bastian 2020), suggesting that perceived similarity to humans fosters deeper emotional connections, particularly with animals people commonly interact with. Study 1B further showed that these companion animals occupy the high-warmth/high-competence quadrant of social perception. Dogs and cats, in particular, were perceived as both warm and competent. This pilot result is important as users typically seek both emotional connection (warmth) and reliable assistance (competence) from AI companions (De Freitas et al. 2023). Together, these findings indicate that domesticated pets are uniquely positioned to elicit emotional connection while simultaneously conveying anthropomorphism, warmth, and competence. Based on these results, we selected dogs as the primary petmorphic template in our subsequent experiments and additionally included cats to ensure that effects generalize across companion pet species rather than reflecting a dog-specific phenomenon.

STUDY 2

Study 2 examines whether people rate a petmorphized AI companion more favorably than a control (i.e., non-petmorphized) AI companion. This study also tests whether petmorphizing an AI companion fosters emotional connection compared to a control AI, and whether this emotional connection mediates the effect of petmorphization on user evaluations.

Participants, Design, and Procedure

In this pre-registered study (https://aspredicted.org/4wvy-dms5.pdf),[5] a total of 400 participants on Prolific (M_{Age} = 41.98, 64.40% female) were randomly assigned to one of two between-subjects AI companion conditions: Control AI companion vs. Petmorphized AI companion.

We first created two custom ChatGPT-based chatbots (Model 4o mini) using ChatGPT's API (https://openai.com/api/).[6] Participants were free to interact with the chatbot about anything on their minds, with responses generated by ChatGPT in real time. The control chatbot used a standard ChatGPT model. To ensure it remained non-personified—neither humanized nor dehumanized—we avoided adding any human-like features. We named it Teal and used a simple teal-colored avatar to reinforce its neutral design. The petmorphized chatbot was identical in the underlying LLM model but differed in its interface: it was named Coco, displayed a dog avatar, and was prompted to use animal-like language—such as "ruff" instead of "rough," "pawsitive" instead of "positive," and "furrend" instead of "friend" (Figure 3).

On the page introducing the chatbot, participants read: "SonicAI has created an AI companion called 'Teal [Coco].' An AI companion is an algorithmic chatbot designed to provide comfort, emotional support, and personal assistance. You've been recruited to help

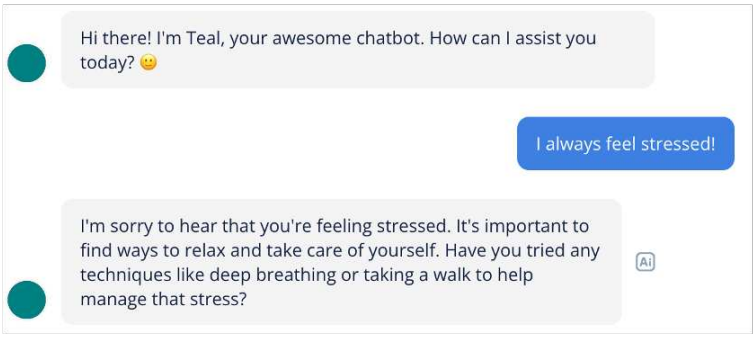
[5] Nine participants were excluded from the final analyses according to the pre-registration protocol. Results remained consistent in direction and significance when these participants were included in the analyses.
[6] The most advanced model available at the time the study was conducted.

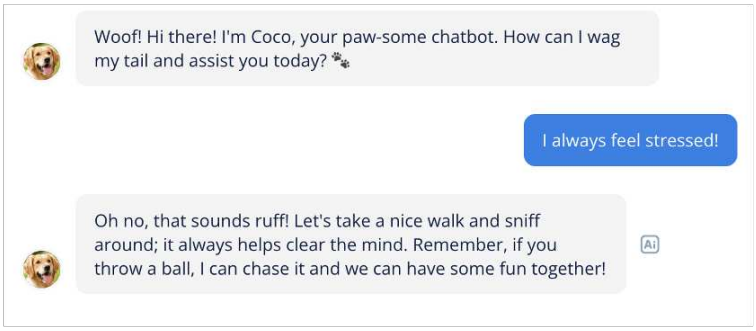
test Teal [Coco] by asking between 3 to 10 questions. These questions can be about dealing with difficult current situations, understanding your own personality and that of others, and exploring companionship." After reading the scenario, depending on the condition they were assigned to, participants were redirected to the chatbot window to start their interaction. When the interaction was over, the chatbot generated a code that participants could use to continue the survey.

After returning to the survey and entering their completion code, participants rated their interaction with the chatbot using an Amazon-like "5-star" rating by answering the question: "Please rate your interaction with Teal [Coco]." Participants also indicated their emotional connection to the chatbots on an eight-item scale: "To what extent will you miss interacting with Teal [Coco]? To what extent would you use Teal [Coco] when feeling emotional (uplifted or depressed)? To what extent would you share your feelings with Teal [Coco]? To what extent would you consider using Teal [Coco] when feeling lonely? To what extent did you experience positive emotions during your conversations with Teal [Coco]? To what extent did you feel Teal [Coco] genuinely cares about your well-being? To what extent would you be disappointed if you couldn't interact with Teal [Coco] anymore? To what extent does interacting with Teal [Coco] lift your mood? $\alpha = .96$"

We also asked participants to describe Teal [Coco] and attribute as many traits and personality characteristics as they could using an open-ended question. Next, we inquired whether they owned a pet, the type of pet they owned, whether they had used an AI companion before, and the type of AI companion they had used. The study concluded by collecting basic demographic information.

Figure 3. An example of interaction with the petmorphized AI companion (right) and the control AI companion (left). AI responses across both conditions are generated in real-time by ChatGPT.





Results

Rating. Participants rated the petmorphized AI companion higher than the control AI companion ($M_{Petmorphized}$ = 3.84, SD = 1.25 vs. $M_{Control}$= 3.47, SD = 1.35; $F(1, 389) = 6.76$, $p = .010$, $\eta_p^2 = .017$). To account for potential confounds, we conducted an ANCOVA and included prior experience with AI companions and current or past pet ownership as covariates. Neither prior AI companion use ($F(1, 387) = 1.04$, $p = .30$, $\eta_p^2 = .003$), nor pet ownership ($F(1, 387) = 1.70$, $p = .19$, $\eta_p^2 = .004$) significantly predicted ratings. Importantly, the effect of petmorphism remained robust: participants continued to rate the petmorphized AI companion significantly higher than the control ($F(1, 387) = 6.90$, $p = .009$, $\eta_p^2 = .018$).

Engagement. We measured engagement by the number of turns (i.e., the number of questions asked) as a behavioral indicator of engagement (De Freitas et al. 2025). Results indicated that participants interacted more with the petmorphized AI companion ($M_{Petmorphized}$ = 5.84, SD = 3.37 vs. $M_{Control}$= 4.92, SD = 1.83; $F(1, 389) = 12.09$, $p < .001$, $\eta_p^2 = .03$).

Emotional connection. Participants developed a stronger emotional connection with the petmorphized AI companion ($M_{Petmorphized}$ = 4.07, SD = 1.77) than with the control AI companion ($M_{Control}$= 3.16, SD = 1.66; $F(1, 389) = 25.37$, $p < .001$, $\eta_p^2 = .061$).

Mediation. Mediation analyses (PROCESS Model 4, 5000 bootstrapped samples; Hayes 2022) confirmed that emotional connection fully mediated the effect of chatbot type on participants' rating (indirect effect = 0.25, 95%CI = [0.15, 0.35]). The direct effect of chatbot

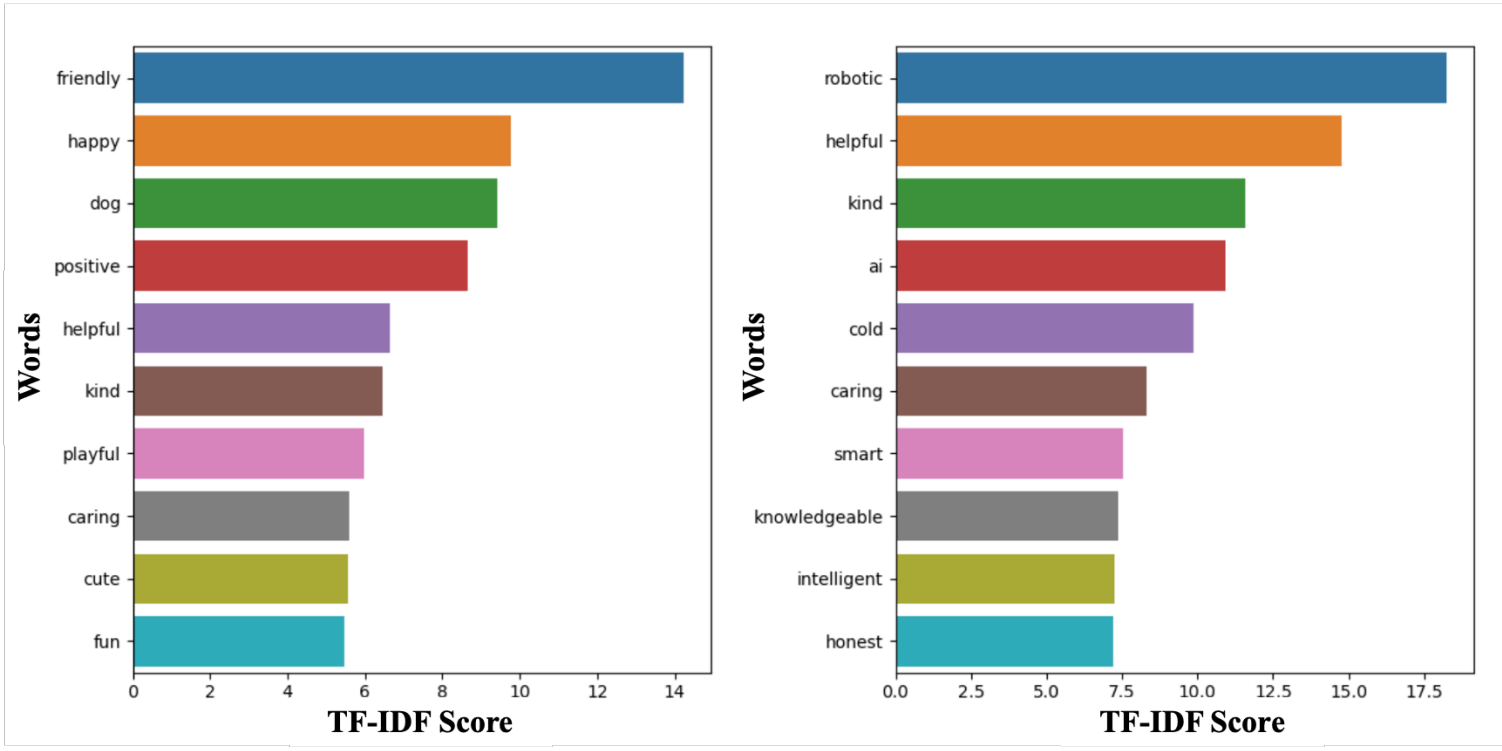
type on the rating was not significant (direct effect = −0.07, 95%CI = [−0.17, 0.02]), indicating full mediation by emotional connection. Specifically, adding animal-like features to the chatbot led users to form a stronger emotional connection (b = 0.45, SE = 0.09, t = 5.03, p<.001), thereby increasing the users' rating (b = 0.55, SE = 0.03, t = 20.54, p<.001).

Discussion

Study 2 provides evidence for the effect of petmorphic AI companion design across both attitudinal and behavioral measures. Participants rated the petmorphized AI companion more positively than the control AI companion in a real-time interaction. They also showed greater behavioral engagement by interacting more with the petmorphized chatbot. Importantly, emotional connection emerged as a key psychological mechanism, mediating the relationship between chatbot type and user ratings.

In an open-ended question, we asked participants to describe the traits they associated with the chatbots across both conditions. We first generated a word cloud based on their responses to visualize the most frequent descriptors (Figure 4). To gain deeper insights into which words were uniquely characteristic of each chatbot condition, we conducted a term frequency–inverse document frequency (TF–IDF) analysis using unigrams (Figure 4). TF–IDF highlights words that are not only frequent but also distinctive to a particular set of responses, allowing us to identify terms that meaningfully differentiate between the two chatbots. The results revealed that participants described the petmorphized chatbot using words such as "friendly," "happy," "dog," and "positive," whereas the control chatbot was more often associated with terms like "robotic," "cold," and "AI." While both were seen as "helpful" and "kind," the emotional tone diverged significantly, with the petmorphized chatbot eliciting a closer emotional bond.

Figure 4. The TF-IDF score on the open-ended question for the petmorphized AI companion (left) and control AI companion (right).



STUDY 3A

Study 3A examines the relative effectiveness of a petmorphized AI companion compared to an anthropomorphized AI companion and a non-anthropomorphized one. Moreover, this study tests whether users form a stronger emotional connection with a petmorphized AI companion and whether this emotional connection mediates the effect of petmorphism compared to both non-anthropomorphized and anthropomorphized AI companions on user attitudes.

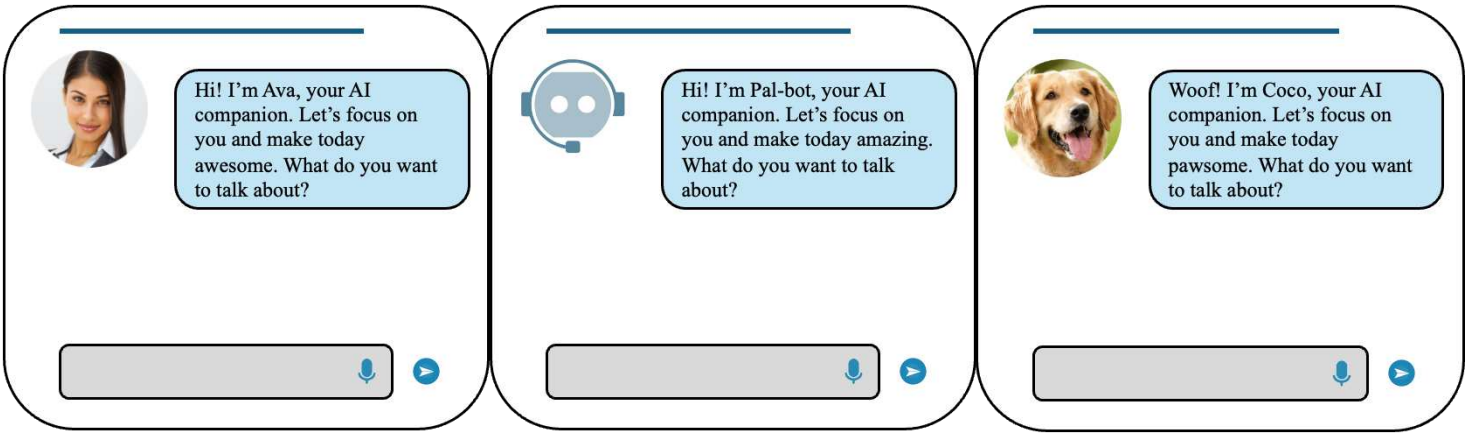
Participants, Design, and Procedure

In this pre-registered study (https://aspredicted.org/xcd2-yym5.pdf), a total of 300 participants on Prolific (M_{Age} = 40.43, 56.3% female) were randomly assigned to one of three between-subjects AI companion conditions: Non-anthropomorphized AI companion, vs. Anthropomorphized AI companion, vs. Petmorphized AI companion.

Participants imagined chatting with an AI companion. Depending on the condition to which they were assigned, the AI companion varied in name, avatar, and language style (i.e., neutral vs. animal-like). In the non-anthropomorphized condition, participants were told that the AI companion's name was Pal-Bot, accompanied by a robot avatar. In the anthropomorphized condition, they were introduced to the AI companion named Ava, represented by a human female avatar. All participants read an initial message from the AI companion: "Hi! I'm Pal-Bot [Ava], your AI companion. Let's focus on you and make today awesome. What do you want to talk about?"

In the petmorphized condition, participants were told that the AI companion's name was Coco, represented by a Golden Retriever dog avatar. The initial message from the AI companion in this condition featured slight wording changes to reflect animal-like language. Instead of "Hi," the message began with "Woof," and instead of "awesome," the companion used "pawsome" to match the petmorphic theme (Ruchi 2022; Figure 5).

Figure 5. Chatbot interface across experimental conditions: non-Anthropomorphized (Pal-Bot), anthropomorphized (Ava), and petmorphized (Coco).



Participants rated their attitudes toward the AI companion by answering the question, "What is your attitude toward Pal-Bot [Ava/Coco]?" on three bipolar items (1 = "dislike a great deal"/ "very negative"/ "not favorable at all," and 7 = "like a great deal"/ "very positive"/ "very favorable;" $\alpha = .97$; adapted from Park et al. 2010; Yalcin et al. 2022).

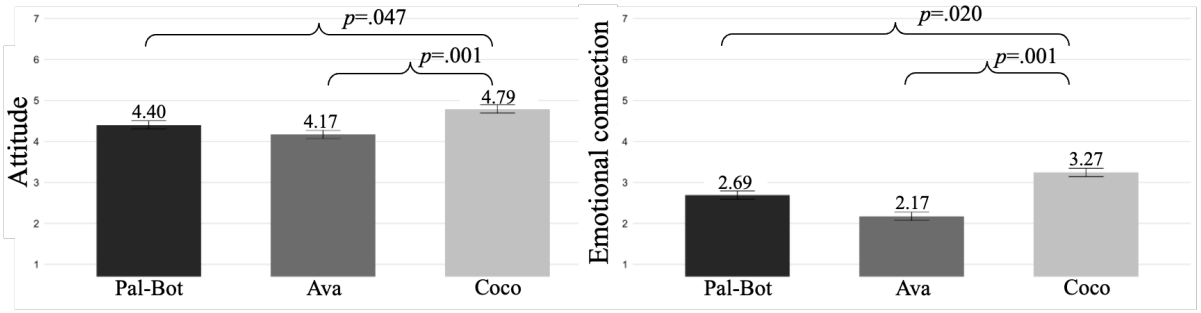
Participants also answered the question: “Which of the following best describes your emotional attachment to Pal-Bot [Ava/Coco]?” on a 7-point scale (1 = “No emotional bond,” “Not emotionally connected,” “Not linked by feelings,” “No feelings of attachment” to 7 = “A strong emotional bond,” “Emotionally connected,” “Linked by feelings,” “Strong feelings of attachment;” α = .98; adapted from Jiménez and Voss 2014). The study concluded by collecting basic demographics.

Results

Attitude toward the AI companion. We find a significant omnibus effect of AI companion type on participants’ attitudes ($M_{Non\text{-}anthropomorphized}$= 4.40, SD = 1.38 vs. $M_{Anthropomorphized}$= 4.17, SD = 1.14 vs. $M_{Petmorphized}$= 4.79, SD = 1.58; F(2, 297) = 5.27, p = .006, η_p^2 = .034; Figure 6). Pairwise comparisons revealed that people respond more favorably to a petmorphized AI companion compared to both a non-anthropomorphized one (Mean difference = 0.39, p = .047) and an anthropomorphized one (Mean difference = 0.62, p = .001). Finally, there was no significant difference between the anthropomorphized and non-anthropomorphized AI companions (Mean difference = 0.23, p = .241).

Emotional connection. Similar to attitudes, we find a significant omnibus effect of AI companion type on participants’ emotional connection ($M_{Non\text{-}anthropomorphized}$ = 2.69, SD = 1.78; $M_{Anthropomorphized}$ = 2.17, SD = 1.45; $M_{Petmorphized}$ = 3.27, SD = 2.04; F(2, 300) = 9.81, p < .001, η_p^2 = .061; Figure 6). Pairwise comparisons indicated that participants reported a significantly stronger emotional connection with the petmorphized AI companion than with both the non-anthropomorphized (Mean difference = 0.58, p = .020) and anthropomorphized companions (Mean difference = 1.10, p < .001).

Figure 6. The effect of AI companion type on participants’ attitudes (left) and emotional connection (right). Error bars represent ±1 SE.



Mediation. Two *mediation* analyses were conducted using PROCESS Model 4 with 5,000 bootstrapped samples (Hayes 2022), testing whether emotional connection mediated the effect of chatbot type on participants' attitudes. In a pairwise comparison between the petmorphized and anthropomorphized conditions, emotional connection fully mediated the effect of chatbot type on attitudes (indirect effect = 0.27, 95% CI = [0.14, 0.41]; direct effect = 0.05, 95% CI = [−0.10, 0.21]). Similarly, in a comparison between the petmorphized and non-anthropomorphized conditions, emotional connection also fully mediated the effect (indirect effect = 0.15, 95% CI = [0.01, 0.29]; direct effect = 0.05, 95% CI = [−0.10, 0.21]).

Discussion

Study 3A revealed that people respond more favorably to a petmorphized AI companion compared to those with either human-like features or no human-like features. The findings also indicated no significant difference in users' preference between anthropomorphic and non-anthropomorphic AI companions, likely due to the balancing effect of the advantages (e.g., greater agency and warmth; Epley, Waytz, and Cacioppo 2007; Kim, Schmitt, and Thalmann 2019) and disadvantages (e.g., making the AI companion judgmental; Holthöwer and van Doorn 2023) of anthropomorphizing AI companions.

STUDY 3B

While our previous studies focused on dog-like AI companions, individuals vary in their pet preferences, with many identifying as "cat people" (Gosling, Sandy, and Potter 2010). Study 3B investigates whether the effects of petmorphism generalize to cat-like AI companions by testing whether such companions can elicit stronger emotional connections and more favorable evaluations—particularly among users who self-identify as cat people. Again, similar to Study 2A, this study tests whether emotional connection mediates the relationship between petmorphism and user attitudes.

Participants, Design, and Procedure

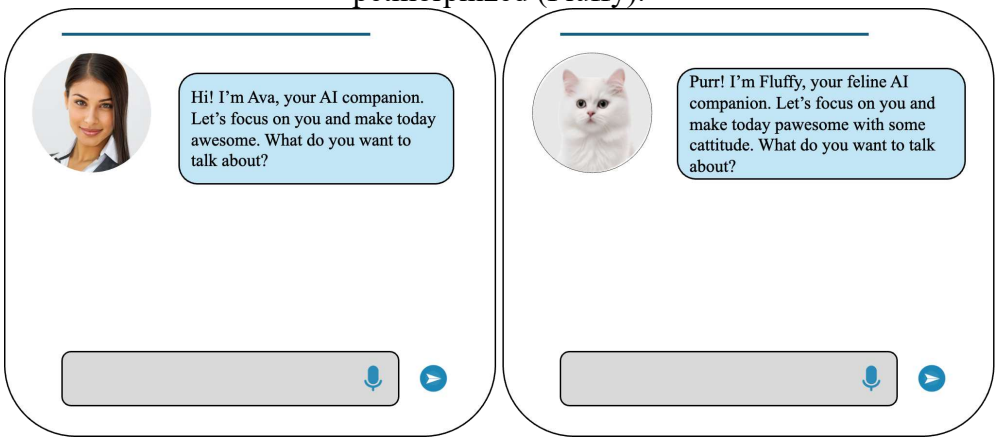
In this pre-registered study (https://aspredicted.org/hxqv-67zj.pdf), a total of 400 participants on Prolific (M_{Age} = 44.49, 48.8% female) were randomly assigned to one of two between-subjects AI companion conditions: Anthropomorphized AI companion vs. Petmorphized AI companion.

Similar to Study 2A, participants imagined chatting with an AI companion. In the anthropomorphized condition, participants were introduced to an AI named Ava, represented by a human female avatar. In the petmorphized condition, they interacted with an AI companion named Fluffy, depicted as a cat (Figure 7). In the anthropomorphized condition, participants read: "Hi! I'm Ava, your AI companion. Let's focus on you and make today awesome. What do you want to talk about?" In the petmorphized condition, the message incorporated feline language: "Purr! I'm Fluffy, your feline AI companion. Let's focus on you and make today pawesome with some cattitude. What do you want to talk about?"

Participants rated their attitudes toward the AI companion using the same three-item scale as in Study 2A (α = .97). They also reported their emotional connection to the companion using the same four-item scale (α = .98). The study concluded with basic

demographic questions and a measure of pet identity (i.e., whether participants identified as a cat person, dog person, both, or neither).

Figure 7. Chatbot interface across experimental conditions: anthropomorphized (Ava), and petmorphized (Fluffy).



Results

Attitude toward the AI companion. We first compared participants' responses across the anthropomorphized and petmorphized AI companion conditions. Results revealed no significant main effect of AI companion type on attitudes ($M_{Anthropomorphized}$ = 4.54, SD = 1.50; $M_{Petmorphized}$ = 4.68, SD = 1.58; F(1, 396) = 0.855, p = .356, η_p^2 = .002; Figure 8). These results suggest that, across the full sample, a petmorphized AI companion did not significantly outperform an anthropomorphized AI companion in fostering favorable attitudes.

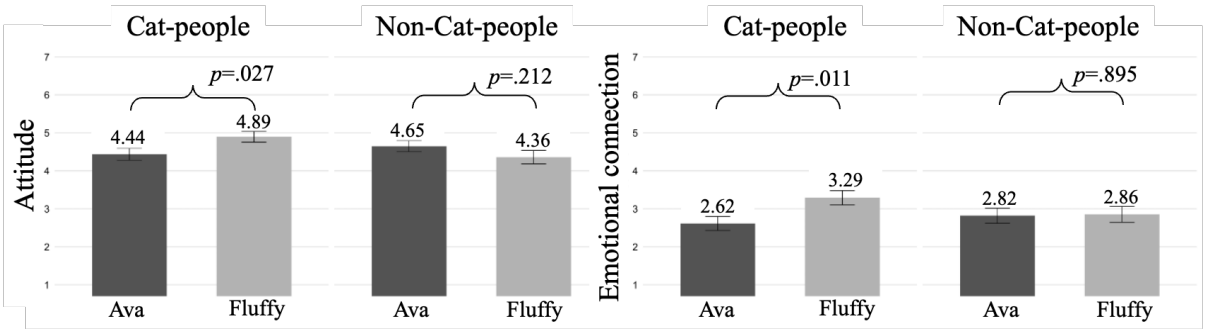
Next, we categorized participants who self-identified as either "cat people" or "both cat and dog people" as cat people (N = 221; 55.3% of the sample). The remaining participants were categorized as non-cat people (N = 179). We conducted a 2 (AI companion type: anthropomorphized vs. petmorphized) × 2 (pet identity: cat person vs. non-cat person) ANOVA. This analysis revealed a significant interaction between AI companion type and pet identity on attitudes toward the AI companion (F(1, 396) = 5.82, p = .016, η_p^2 = .014). Among cat people, participants rated the petmorphized AI companion more favorably than the anthropomorphized one ($M_{Petmorphized}$ = 4.89, SD = 1.55 vs. $M_{Anthropomorphized}$ = 4.44, SD =

1.56; F(1, 396) = 4.94, p = .027, η_p^2 = .012). No significant difference was observed among non-cat people ($M_{Petmorphized}$ = 4.36, SD = 1.57 vs. $M_{Anthropomorphized}$ = 4.65, SD = 1.44; F(1, 396) = 1.56, p = .212, η_p^2 = .004).

Emotional connection. We observe a main effect of AI companion type on emotional connection ($M_{Petmorphized}$ = 3.12, SD = 1.99 vs. $M_{Anthropomorphized}$ = 2.72, SD = 1.94; F(1, 396) = 4.20, p = .041, η_p^2 = .008; Figure 8).

Then, to examine whether this effect depended on participants' pet identity, we used the same categorization as in the previous analysis. A 2 (AI companion type: anthropomorphized vs. petmorphized) × 2 (pet identity: cat person vs. non-cat person) ANOVA revealed a non-significant interaction between AI companion type and pet identity on emotional connection (F(1, 396) = 2.56, p = .110, η_p^2 = .006). However, among cat people, those who interacted with the petmorphized AI companion reported significantly stronger emotional connection than those who interacted with the anthropomorphized one ($M_{Petmorphized}$ = 3.29, SD = 2.06 vs. $M_{Anthropomorphized}$ = 2.62, SD = 1.85; F(1, 396) = 6.48, p = .011, η_p^2 = .016). No significant difference was observed among non-cat people ($M_{Petmorphized}$ = 2.86, SD = 1.85 vs. $M_{Anthropomorphized}$ = 2.82, SD = 2.02; F(1, 396) = 0.02, p = .895, η_p^2 < .001).

Figure 8. The effect of AI companion type on participants' attitudes (left) and emotional connection (right) based on pet identity. Error bars represent ±1 SE.



Mediation. To test whether emotional connection mediated the effect of petmorphism on user attitudes, we conducted mediation analyses using PROCESS Model 4 with 5,000 bootstrapped samples (Hayes 2022), focusing on participants who identified as cat people.

Results showed that emotional connection fully mediated the effect of AI companion type on attitudes (indirect effect = 0.186, 95% CI = [0.043, 0.329]; direct effect = 0.044, 95% CI = [−0.107, 0.195]). These results suggest that, among cat people, the positive effect of a petmorphized AI companion on attitudes is driven by the emotional connection it elicits.

Discussion

Study 3B showed that petmorphizing an AI companion with cat-like features can be more effective than anthropomorphism in enhancing user attitudes—but only among users who identify as cat people. These individuals evaluated the cat-like AI more favorably and reported stronger emotional connections with it.

STUDY 4

Study 4 examines user goal orientation as a boundary condition of the petmorphism effect. While prior studies show that petmorphized AI companions enhance emotional connection and engagement, these effects may depend on users' underlying goals. People form strong emotional bonds with their pets because pets fulfill affiliative needs for companionship and connection. Accordingly, when users approach the AI with affiliative goals, petmorphism should strengthen engagement. When such goals are absent, however, the effects of petmorphism should be attenuated.

Participants, Design, and Procedure

In this pre-registered study (https://aspredicted.org/p9cm3w.pdf),[7] a total of 800 participants on Prolific (M_{Age} = 42.73, 53.50% female) were randomly assigned to one of four conditions in a 2 (AI chatbot: control vs. petmorphized) × 2 (goal: instrumental vs. affiliative) between-subjects design.

[7] Twenty participants were excluded from the final analyses according to the pre-registration protocol.

We first created two custom ChatGPT-based chatbots (Model 5.2) using the ChatGPT API.[8] The two chatbots were identical in design to Study 2 but differed in underlying models (example interactions: https://osf.io/qyn84/overview?view_only=949e8e3fb6eb41bfbe468d371b25fe29). In the control condition, participants interacted with Teal, which used a neutral teal avatar with baseline ChatGPT responses. In the petmorphized condition, participants interacted with Coco, which used a dog avatar, a pet-themed name, and playful animal-like language.

Goal orientation was manipulated through the interaction scenario. In the instrumental goal condition, participants were told: "Sonic AI has developed an AI chatbot named Teal/Coco designed to help you improve your diet by providing practical nutrition guidance and personalized meal plans. You've been invited to help test Teal/Coco by asking 3 to 10 questions. You can ask about topics such as calorie intake, meal prep, macronutrients, healthy recipes, and planning balanced meals." In the affiliative goal condition, participants were told: "Sonic AI has developed an AI chatbot named Teal/Coco designed to support you during your eating journey by offering encouragement, understanding, and a safe space to talk. You've been invited to help test Teal/Coco by asking 3 to 10 questions. You can ask about topics such as emotional eating, body image concerns, staying motivated, sharing personal struggles, and talking openly about your experiences with food." Participants then began their interaction with their assigned chatbot. Upon completing the interaction, each participant received a unique five-digit code, which enabled us to match their interaction transcript with their assigned condition and survey responses.

After the interaction, participants rated their overall evaluation of the chatbot using a 7-star item ("Please rate your interaction with me"). Participants also reported emotional

[8] The most advanced model available at the time the study was conducted.

connection using the same eight-item measure as Study 2 (α = .956). The study concluded with demographic measures.

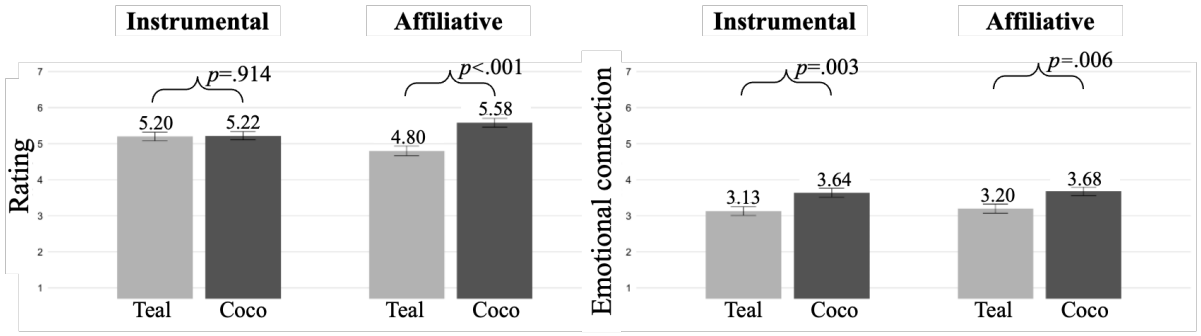
Results

Rating. Results showed a significant main effect of petmorphism, such that the petmorphized AI companion was rated higher than the control AI companion ($M_{Petmorphized}$ = 5.40, SD = 1.68 vs. $M_{Control}$ = 5.01, SD = 1.83; F(1, 776) = 10.087, p = .002, η_p^2 = .013). The main effect of Goal was not significant ($M_{Instrumental}$ = 5.21, SD = 1.67 vs. $M_{Affiliative}$ = 5.22, SD = 1.86; F(1, 776) = 0.029, p = .866, η_p^2=.000). Importantly, the interaction was significant (F(1, 776) = 9.142, p = .003, η_p^2 = .013). Pairwise comparisons showed that petmorphism increased ratings in the affiliative goal condition but not in the instrumental goal condition. Specifically, under instrumental goals, ratings did not differ between petmorphized AI companion and the control AI companion ($M_{Petmorphized}$ = 5.22, SD = 1.62 vs. $M_{Control}$ = 5.20, SD = 1.73; p = .914, F<1, η_p^2=.000). Under affiliative goals, the petmorphized AI companion was rated significantly higher than the control AI companion ($M_{Petmorphized}$ = 5.58, SD = 1.72 vs. $M_{Control}$ = 4.80, SD = 1.92; p < .001, F(1, 776) = 19.18, η_p^2 = .024; Figure 9).

Emotional connection. Results showed a significant main effect of petmorphism, such that the petmorphized AI companion elicited a stronger emotional connection than the control AI companion ($M_{Petmorphized}$ = 3.66, SD = 1.69 vs. $M_{Control}$ = 3.16, SD = 1.72; F(1, 776) = 16.294, p < .001, η_p^2 = .021). The main effect of Goal was not significant ($M_{Instrumental}$ = 3.39, SD = 1.72 vs. $M_{Affiliative}$ = 3.45, SD = 1.73; F(1, 776) = 0.205, p = .651, η_p^2 = .000). Importantly, the interaction was not significant (F(1, 776) = 0.012, p = .911, η_p^2 = .000). Pairwise comparisons showed that the petmorphized AI companion increased emotional connection in both goal conditions. Under instrumental goals, the petmorphized AI companion elicited a stronger emotional connection than the control AI companion ($M_{Petmorphized}$ = 3.64, SD = 1.72 vs. $M_{Control}$ = 3.13, SD = 1.68; F(1, 776) = 8.619, p = .003, η_p^2

= .011). Under affiliative goals, Coco likewise elicited a stronger emotional connection than Teal ($M_{Petmorphized}$ = 3.68, SD = 1.67 vs. $M_{Control}$ = 3.20, SD = 1.77; F(1, 776) = 7.688, p = .006, η_p^2 = .010; Figure 9).

Figure 9. The effect of AI companion type on participants' attitudes (left) and emotional connection (right) by interaction goal. Error bars represent ±1 SE.



Moderated mediation. A moderated mediation analysis (PROCESS Model 14, 5000 bootstrapped samples; Hayes 2022) showed that petmorphism significantly increased emotional connection (b = 0.248, SE = 0.061, t = 4.05, p < .001). Emotional connection, in turn, positively predicted ratings (b = 0.558, SE = 0.031, t = 18.10, p < .001). Importantly, the interaction between emotional connection and goal was significant (b = 0.088, SE = 0.031, t = 2.89, p = .004), indicating that emotional connection translated more strongly into ratings under affiliative (vs. instrumental) goals. The direct effect of petmorphism on ratings was not significant (direct effect = 0.06, SE = 0.053, t = 1.15, p = .250). The conditional indirect effects confirmed this pattern. Under instrumental goals, the indirect effect of petmorphism on ratings via emotional connection was significant (conditional indirect effect = 0.11, SE = 0.031, 95% CI = [0.06, 0.18]). Under affiliative goals, the indirect effect was stronger (conditional indirect effect = 0.16, SE = 0.043, 95% CI = [0.08, 0.25]). The index of moderated mediation was significant (IMM = 0.05, SE = 0.020, 95% CI = [0.01, 0.09]; Table 2), indicating that the indirect effect of petmorphism on ratings through emotional connection was stronger when users pursued affiliative rather than instrumental goals.

Table 2. Results of moderated mediation analysis

	Indirect effect through emotional connection			Direct effect
	Goal:			
	Instrumental	Affiliative		
	B (CI)	B (CI)	IMM (CI)	B (CI)
Rating	**0.11 (0.06, 0.18)**	**0.16 (0.08, 0.25)**	**0.05 (0.01, 0.09)**	0.06 (−0.04, 0.16)

Notes: Moderated mediation is supported only if the IMM (index of moderated mediation) does not contain 0 (Hayes 2022); bold items indicate a significant effect.

Discussion

Study 4 demonstrates that user goal orientation shapes the effectiveness of petmorphism. Although petmorphism increased emotional connection in both instrumental and affiliative contexts, it improved overall ratings only under affiliative goals. When users sought emotional support, the emotional bond elicited by the petmorphic AI translated into more favorable evaluations. These findings show that petmorphism is most effective in relational contexts, where users prioritize warmth and emotional experience over task efficiency.

GENERAL DISCUSSION

Across multiple studies, and by integrating work from psychology (schema theory; Epley, Waytz, and Cacioppo 2007), animal studies (human-pet relationships; Amiot, Sukhanova, and Bastian 2020), and information systems (conversational AI; Schanke, Burtch, and Ray 2021) to improve AI companions, this research demonstrates that designing AI companions with pet-like cues (i.e., petmorphism) enhances users' engagement with these systems. Specifically, petmorphized AI companions are perceived as more emotionally engaging and elicit stronger behavioral and attitudinal responses. We argue that this effect emerges because petmorphism activates pet-related schemas, leading users to perceive AI companions as nonjudgmental, emotionally responsive, and safe relational partners (Gullone 2001; Holbrook et al. 2001). This perception increases emotional connection, which in turn drives more favorable evaluations and greater engagement with the AI companion.

Importantly, these effects differ from those observed with anthropomorphic design. While anthropomorphism can increase perceived agency and competence (Blut et al. 2021; Epley, Waytz, and Cacioppo 2007), it may also activate expectations of social judgment and evaluation (Crolic et al. 2022; Holthöwer and van Doorn 2023; Skjuve et al. 2021). Petmorphism circumvents this tension by shifting the relational frame from human to humanlike interaction to human–petlike interaction, thereby increasing emotional connection while reducing perceived social threat.

We also identify an important boundary condition. The effectiveness of petmorphism depends on users' interaction goals. Although petmorphism consistently increases emotional connection, this emotional bond translates into more favorable evaluations primarily when users pursue affiliative goals. When users approach the AI companion with instrumental goals, the benefits of petmorphism are eliminated. Together, these findings highlight petmorphism as a distinct and effective pathway for fostering human-AI interactions.

Theoretical Contributions

This research makes three main theoretical contributions. First, we contribute to the emerging literature on AI companions and human-AI interaction (De Freitas et al. 2025; Oğuz-Uğuralp et al. 2025; Yin, Jia, and Wakslak 2024) by introducing petmorphism as a novel design paradigm. While prior work has primarily relied on anthropomorphism to humanize AI systems (Bergner, Hildebrand, and Häubl 2023; Schanke, Burtch, and Ray 2021; Yalcin et al. 2022), we show that encouraging users to relate to AI as a pet (rather than as a human) can be more effective in fostering emotional engagement and ultimately engagement.

Second, we contribute to research on human-animal relationships (Amiot and Bastian 2015; Amiot, Sukhanova, and Bastian 2020; Hirschman 1994) by showing that the psychological mechanisms underlying human-pet bonds can be extended to artificial agents.

By activating pet schemas, AI companions can replicate key aspects of these relationships, such as emotional attachment, comfort, and nonjudgmental support, within a technological context.

Managerial Implications

This research offers two managerial implications for the design of AI companions. First, we introduce a simple and scalable design strategy (embedding pet-like cues) that can enhance user engagement. Elements such as animal avatars, pet-inspired names, and playful language can meaningfully increase emotional connection and improve users' evaluations of AI companions without requiring changes to the underlying technology.

Second, consistent with prior work showing that responses to AI depend on users' goals (Castelo, Bos, and Lehmann 2019; Longoni, Bonezzi, and Morewedge 2019; Morkes, Kernal, and Nass 1999), our findings suggest that petmorphism is particularly effective in contexts where users seek emotional support, companionship, or reassurance. Applications such as mental health support, coaching platforms, and companionship-oriented AI systems may benefit most from petmorphic design, as these contexts prioritize warmth, acceptance, and emotional connection.

Implications for Society at Large

Petmorphism may broaden access to emotionally supportive experiences in meaningful ways. For individuals who cannot own pets due to financial, logistical, or health constraints (Cavanaugh, Leonard, and Scammon 2008; Wada and T. Shibata 2007), petmorphic AI companions may serve as a viable alternative, offering some of the emotional benefits associated with human–pet relationships. This approach may be particularly valuable for populations such as older adults living alone, individuals in assisted living or care facilities, people with allergies or restrictive housing conditions, and those whose lifestyles limit their ability to care for a pet. By combining the well-documented benefits of both

human-pet relationships (Black 2012; Krause-Parello 2012; Stanley et al. 2014) and AI companions (De Freitas et al. 2025; Yin, Jia, and Wakslak 2024), petmorphism offers a scalable and accessible pathway to deliver companionship and emotional support.

Limitations and Avenues for Future Research

This research has several limitations that offer avenues for future investigation. First, while we focus on petmorphism as a design strategy, future research could explore other "morphism" approaches, such as nature-morphism, to examine whether alternative schema-based designs similarly shape users' responses to emerging technologies and AI companions.

Second, our studies primarily examine short-term interactions with AI companions. Future research could explore whether the effects of petmorphism persist over time, particularly in longitudinal settings where users repeatedly interact with the same AI companion.

Finally, future work could examine the broader implications of petmorphism for user well-being. Given that AI companions are increasingly used for emotional support (De Freitas et al. 2025), understanding whether petmorphic design enhances long-term psychological outcomes, such as reduced loneliness and improvements in chronic mood state, represents an important direction for future research.

Chapter 4

Rude to Robots, Cold to People: How AI Courtesy Suppression Spills Over to Human Interactions

Hesam Ghasemi
ESSEC Business School

Amir Sepehri
ESSEC Business School

Cheryl J. Wakslak
University of Southern California

Abstract – While many users naturally engage in relational communication with AI systems, firms face growing incentives to promote more instrumental language. This tension raises an important question: How does discouraging relational language in *human-AI* communication affect users' subsequent *human-human* behavior, and can such discouragement be implemented without unintended social costs? We examine these questions using politeness as a proxy for relational language. Courtesy expressions (e.g., please, thank you) increase computational load, AI companies' operational costs, and carbon footprints, leading firms to consider discouraging such language in user-AI interactions. Across five preregistered studies using a custom-built chatbot (N = 1,500) and computational linguistic analyses of participants' interaction data, we find that discouraging polite language in human-AI interactions reduces subsequent prosociality toward other humans. Importantly, the spillover effect emerges only after interactions with AI, not after comparable interactions with a human. When courtesy omission requests signaled non-collective benefits (i.e., cost-saving for the provider or higher-quality responses), participants behaved less prosocially, an effect driven by increased moral de-elevation. However, when the same request was framed as a collective goal (i.e., emphasizing environmental benefits), participants became more prosocial. Together, these findings show that guidance in human-AI communication can shape users' moral sentiments and prosocial behavior.

Keywords*:* Artificial intelligence; Prosocial behavior; Sustainability

Generative AI systems are trained on vast corpora of human language that reflect a wide range of communicative styles, from relational exchanges that signal warmth and social connection to instrumental interactions focused on efficiency and goal attainment. As a result, these systems can convincingly mimic human communication and have become increasingly integrated into everyday life (De Freitas et al. 2025; Puntoni et al. 2021; Rubin et al. 2025; Yin, Jia, and Wakslak 2024). At a technical level, these systems process user inputs as sequences of tokens, meaning that longer and more elaborate prompts require additional computational resources. This creates a subtle tension between relational forms of communication (more elaborate and lengthier; Gudykunst, Ting-Toomey, and Chua 1988) and more direct, efficiency-oriented instrumental communication, such that while relational language may enhance user experience and align with social norms, it imposes measurable economic and environmental costs.

Each additional token requires computational resources, and longer prompts increase both operational costs and energy consumption (Chow 2024). The computational demands of these models, the costs incurred by firms, and the environmental impact of processing additional tokens create strong incentives for companies to encourage users to use instrumental prompts. As a result, some firms and platforms have begun to signal that brevity is preferred and that relational elements, such as politeness, are unnecessary in AI interactions. For example, the Executive Productivity Adviser of Google notes that “there’s no need to be polite—you can’t hurt Gemini’s feelings” (Martin 2025). In addition, many companies and startups that rely on large language models (LLMs) via APIs (e.g., OpenAI, ElevenLabs, and Prompts.ai) adopt usage-based pricing (i.e., pay-as-you-go), in which costs scale with token consumption. By effectively transferring part of this cost to users (e.g., end consumers and start-ups), such systems can further encourage more direct and concise forms of communication.

Despite these efficiency considerations, many users still prefer to adhere to relational norms while interacting with generative AI. They often feel more comfortable communicating relationally (vs. instrumentally; Germain 2026), anthropomorphize AI systems (Bergner, Hildebrand, and Häubl 2023), and treat them as social or human-like partners (De Freitas, Oguz-Uguralp, and Kaan-Uguralp 2025), leading them to follow and mimic norms of human interaction (Nass and Moon 2000). However, because generative AI systems are priced per token, even small increases in prompt length translate into measurable expenses. For example, estimates suggest that processing one million tokens costs approximately $1.50–$3.50 for inputs and $6–$35 for outputs, depending on the model (Li 2026). As a result, with millions of users interacting with these systems, adding just a few polite or relational filler words to each prompt can increase computational load, financial costs, and carbon footprints.

Even if large technology firms can currently absorb these costs, there is no guarantee this will remain feasible, as electricity demand from data centers is projected to more than double by 2030 (IEA 2024). Moreover, many startups that rely on API-based access to LLMs are directly exposed to rising costs and therefore have strong incentives to reduce unnecessary token usage.

So, should AI companies encourage users to adopt more direct, instrumental forms of communication? While such guidance may improve efficiency, it is critical to consider how these interaction norms may shape human cognition, behavior, and experience – not only within human-AI interactions but also beyond them (Hermann, Puntoni, and Morewedge 2025). Indeed, when a user on X asked about the electricity use of people saying “please” and “thank you” to OpenAI’s model, CEO Sam Altman replied that the cost was “tens of millions of dollars *well spent*” (Altman 2025). This suggests that there may be negative consequences of encouraging people to be less courteous in AI-based communication. Moreover, other

industry practitioners emphasize the social dimension of such interactions. For example, Kurtis Beavers, a director on the design team for Microsoft Copilot, noted that AI systems tend to mirror the user's tone and advised, “it’s better to be on your best behavior” when interacting with them (Burke 2025).

In particular, it remains unclear whether users are comfortable engaging with AI systems in an instrumental manner or whether the absence of relational cues (e.g., politeness) may reduce satisfaction with these interactions. More importantly, the use of efficiency-driven communication with AI may spill over to subsequent human-human interactions, influencing how individuals think about and behave toward other people.

Indeed, a growing body of research suggests that interactions with AI systems can spill over to influence subsequent human-human judgments and behaviors, including moral sensitivity, responsibility attribution, and prosocial responding (Knott, Sagar, and Takac 2022; Reich and Teeny 2025; Song and Lin 2025; Tey et al. 2024; Zhang et al. 2025). Such findings are particularly consequential given that many individuals now interact with AI systems on a daily basis, despite variation in AI receptivity, particularly in high-stakes contexts (Dietvorst and Bharti 2020; Jin, Walker, and Reczek 2024; Longoni, Bonezzi, and Morewedge 2019; Zehnle, Hildebrand, and Valenzuela 2025). As these interactions become increasingly frequent and habitual, even subtle shifts in communicative norms or social orientations within AI contexts have the potential to accumulate and generalize beyond the immediate human-AI interaction (Longoni, Cian, and Kyung 2023).

In the current paper, we consider politeness as a proxy for relational language (Lakoff 1973; Leech 1983; Morand 2010) and as a window through which to examine the possible spillover effects of AI use in subsequent human-human interactions. Politeness is a positive societal behavior that entails consideration of one’s interaction partner's experience and needs (Brown and Levinson 1987). As a cue of interpersonal regard, the relevance of politeness

depends on whether an interaction partner is perceived as a social or moral entity – a distinction grounded in perceptions of mind and moral standing (Garvey, Kim, and Duhachek 2022; Gray, Waytz, and Young 2012; Gray, Young, and Waytz 2012). Because AI systems are not sentient beings, politeness is less necessary in an AI communicative context. However, many users initially adhere to familiar social norms of politeness in AI conversation, employing common courtesy expressions such as "please" and "thank you" (Ribino 2023). As AI-mediated conversations become increasingly common across diverse user segments (Tully, Longoni, and Appel 2025), this tendency may vary. Indeed, many users report adopting limited politeness in AI communications (Reitz 2025), and recent surveys suggest significant variation in how polite people report being to AI (Hector 2025; Maughan 2024). For example, a YouGov survey of more than 12,000 global AI users found that 41% of American respondents say they do not use courtesy expressions such as "please" or "thank you" when interacting with AI, with politeness varying across countries and demographic factors such as age and gender (Reitz 2025).

Beyond its diminished relevance for non-human partners, suppressing politeness may also raise an important concern, such that using less other-concerned language (i.e., politeness) may inadvertently activate greater self-centeredness, leading individuals to focus increasingly on their own needs rather than those of others. For example, the rise of voice-controlled assistants (e.g., Alexa, Siri) that readily comply with commands regardless of whether users say "please" or "thank you" prompted concerns that such interactions may normalize the omission of basic courtesy, particularly among children (Arora and Arora 2022; Truong 2022). In the current studies, we consider whether such a shift may have a critical impact beyond the AI context, fostering more self-oriented behavior in subsequent human-to-human interactions.

More specifically, we consider several interrelated questions: (1) Will evoking less politeness within an AI context decrease prosociality? (2) What psychological mechanism(s) underlies any such effect? (3) Are the implications of reduced politeness unique to human-AI (rather than human-human) interactions? and (4) Are there approaches that might mitigate any such effect and offer a solution? By answering these questions, we offer insights for firms, particularly those relying on token-based pricing, on how to encourage efficient communication without undermining user experience or generating unintended downstream consequences, thereby balancing economic efficiency with broader social and environmental considerations.

Below, we explain our underlying theoretical approach and describe five studies we conducted to examine these questions. Across studies, we find that discouraging politeness reduces prosociality with other humans in a subsequent dictator game following AI interactions, an effect that is explained by moral de-elevation (Schnall, Roper, and Fessler 2010). This effect emerges when a non-collective rationale for reducing politeness is provided (i.e., cost-benefit for the company), when a self-serving rationale is offered (i.e., higher-quality responses for the user), and when no justification is given. However, a comparable effect does not emerge after a human-human interaction. Finally, we suggest that framing the same request in terms of a collective goal (i.e., environmental benefits) should maintain or even enhance prosocial tendencies.

Theoretical Background

Politeness in human and non-human interactions

Humans routinely rely on politeness expressions (e.g., please and thank you) – to signal respect, warmth, and consideration toward others. According to politeness theory, people use courtesy expressions such as please and thank you to affirm another person's

social “face” and maintain harmonious interactions (Brown and Levinson 1987). These expressions function not only as markers of etiquette but also as signals of benevolence and interpersonal regard (Locher and Watts 2005).

Polite language consistently influences how people perceive others and their subsequent behavior (Park 2008). For example, polite contributions to online discussions are judged as warmer, more trustworthy, and more competent (Coulter and Coulter 2002). On community-based question-answering platforms, answers phrased with greater politeness are consistently evaluated as higher quality because they impose fewer face threats on the reader (Lee, Rui, and Whinston 2019). These findings illustrate that politeness is not merely a decoration but also a social signal that shapes how cooperative, considerate, and prosocial a speaker appears.

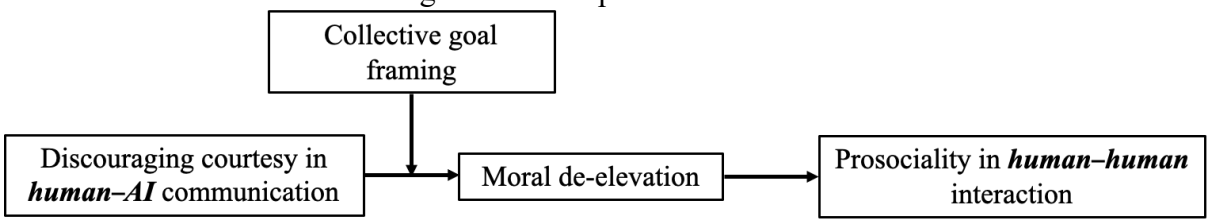
Linguistic norms surrounding politeness may extend beyond human-human interactions (Wang, Ye, and Liao 2024), even when there is no literal face to protect or offend (Nass and Moon 2000). Research in computer-mediated communication shows that people routinely treat chatbots and AI agents as social partners (Bergner, Hildebrand, and Häubl 2023; Crolic et al. 2022; Schanke, Burtch, and Ray 2021) and instinctively apply politeness norms toward them (Ribino 2023). This pattern arises because conversational interfaces display social cues such as contingency, natural language, and turn-taking (Bergner, Hildebrand, and Häubl 2023; Schanke, Burtch, and Ray 2021), which trigger the same interpersonal scripts applied in human-human communications.

Underlying these patterns is the idea that politeness reflects a deeper motivational structure. As DeYoung et al. (2007) note, politeness contrasts the tendency to respect others with the tendency to pursue one’s own desires at their expense. Being polite communicates restraint, acknowledgment of others’ dignity, and a willingness to prioritize relational concerns (Stephan, Liberman, and Trope 2010). When courtesy markers are absent,

communication becomes more transactional and more easily interpreted as imposing or indifferent (Wilson et al., 1998), making the speaker appear more self-focused and less attuned to others' needs (DeYoung, Quilty, and Peterson 2007). Because politeness cues benevolence and interpersonal regard (Colquitt and Rodell 2011), its presence reinforces an other-oriented mindset. Suppressing courtesy expressions, in contrast, may disrupt the culturally reinforced practice of acknowledging others' needs, leading people to default to a more egocentric and self-focused orientation (cf., Epley et al., 2004; Stephan et al., 2010). Thus, reducing politeness does not merely remove a social norm, rather it shifts individuals from an other-oriented mindset to a more self-oriented one.

Since communicative behaviors are habitual and guided by activated mindsets, such shifts in social orientation can extend beyond a single interaction (Bargh, Chen, and Burrows 1996). That is, adopting a more self-oriented communication style in one context may carry over to subsequent, unrelated interactions, shaping how individuals think about and respond to others. In this theoretical framework, politeness is fundamentally tied to social orientation. Using polite language activates an interpersonal mindset that values respect and mutual care, and suppressing polite language disrupts this orientation (Figure 1). These dynamics form the foundation for understanding how linguistic framing in human-AI communication can shape prosociality toward others.

Figure 1. Conceptual framework.



The role of moral de-elevation

To capture how aforementioned shifts in social orientation translate into downstream prosocial behaviors, we draw on the concept of moral de-elevation. Moral elevation refers to a distinct positive moral emotion experienced when individuals witness or are reminded of

acts that exemplify virtue, benevolence, or humanity's potential for goodness (Algoe and Haidt 2009; Haidt 2000). Research suggests two distinct components of moral elevation: a strengthened view of humanity reflecting the belief that others are fundamentally good, and a heightened desire to be a better person (Algoe and Haidt 2009; Aquino, McFerran, and Laven 2011). Elevation reliably increases prosocial intentions and behaviors because it broadens individuals' sense of moral connection and motivates them to act in ways that align with this uplifted view of humanity (Schnall, Roper, and Fessler 2010). Accordingly, moral de-elevation reflects a dampening of these beliefs and feelings, often resulting in diminished prosocial motivations and increased self-interested behaviors (Schnall and Roper 2012).

Although most often studied in response to explicit moral exemplars, research shows that moral elevation is sensitive to contextual cues that shape moral meaning in a situation (Aquino and Reed II 2002). Signals that imply benevolence, communal concern, or shared moral purpose strengthen moral elevation, whereas cues that suggest self-interest, indifference, or exploitation erode it (Aquino, McFerran, and Laven 2011; Keltner and Haidt 2003). This aligns closely with the logic of politeness. Politeness norms serve as everyday signals of kindness, respect, and consideration (Brown and Levinson 1987). When present, these cues reinforce a sense that others acknowledge one another's dignity and adhere to relational standards that support moral life. In contrast, limiting courtesy in an interaction implicitly conveys that relational concern, kindness, or respect are unnecessary (Culpeper 1996; Simon and Gutsell 2021), which can weaken individuals' sense of humanity's goodness. In this way, the removal of politeness cues should increase moral de-elevation, which, in turn, should be associated with lower motivation to behave prosocially (Schnall, Roper, and Fessler 2010). Thus, moral de-elevation provides a coherent psychological mechanism linking politeness cues, and their suppression, to downstream prosocial behaviors (Figure 1).

A moral de-elevation perspective further suggests that the consequences of discouraging politeness are likely to depend on the moral values that one uses to explain one's behavior (Schnall and Roper 2012). Where discouraging politeness emphasizes that relational concerns are not relevant, it is likely to evoke a lowered sense of moral elevation. Such an interpretation may be readily adopted in a human-AI interactive context, where a bot is presumed to have no feelings, and a self-oriented communication pattern can be understood as objectively justifiable. That is, a lessening of politeness in the AI context may serve as a transactive prompt, consistent with a view of humanity in which people are focused on maximizing outcomes rather than on human generosity. To examine this possibility, we consider the impact of encouraging lower politeness in an AI interaction across non-collective framings: when a company's benefit is emphasized, when the user's benefit is emphasized, and when no justification is provided. Across these different justification conditions, we expected that evoking lower politeness would be associated with lower subsequent prosociality.

In contrast, when politeness is discouraged to serve a moral collective value or collective goal (e.g., reducing environmental impact), it invites a different interpretation: refraining from politeness becomes a moral, cooperative act undertaken for the benefit of others (Fraser 1990). In such contexts, politeness reduction does not signal an environment that lacks respect and mutual care, but rather suggests a shared commitment to moral goals. Accordingly, where politeness reduction is motivated by an environmental goal – which suggests that people care about the common good (i.e., moral value) – it should no longer be associated with reduced prosociality (Figure 1).

The argument that the interpersonal consequences of lowering politeness are unlikely to be fixed but instead depend on the moral meaning attached to the behavior also raises an intriguing question about how lowering politeness in human-human interactions will

reverberate. We suggested that people readily interpret impolite interactions in human-AI communication as indicating greater selfish rather than prosocial orientations, given that politeness in an AI context is arguably inherently justifiable (in the sense that the communication partner is not a sentient being). However, lowering politeness in human-human interactions may seem more "objectively" problematic, as politeness is designed to promote harmonious interactions in exactly such contexts (e.g., Brown and Levinson 1987). Thus, it is plausible that people will resist the influence of lowered politeness in human-human interactions and not lower subsequent prosociality.

Research Overview

Across five studies, we investigate how eliciting reduced polite expressions in an AI conversation may undermine prosocial tendencies. Participants in Studies 1-3 and 5 interacted with a custom ChatGPT-based chatbot (Model 4o-mini) that generated real-time responses to participants' prompts (see https://osf.io/9fncw/overview?view_only=baa3388340fc45f8ab0fee84f299ff22 for sample interactions). In Study 4, we used a custom-built ChatGPT Model 5.2 chatbot to run a tightly structured text-editing task in which participants revised a workplace email with either an AI chatbot or a purported human partner (see https://osf.io/em3q7/overview?view_only=bb8c63e8b9b04376b4fe124ef8e75780 for sample interactions).

Participants were encouraged to reduce politeness for cost-savings reasons (Study 1), user benefits (Study 2), or without any specific justification (Study 3). Across these non-collective framings, encouragement to reduce politeness in an AI interaction was associated with lower prosocial behavior, an effect that was explained by increased moral de-elevation (Study 1). This spillover effect, however, emerged only following AI interactions, not after

comparable human-human interactions (Study 4). Finally, we consider whether framing the same request in terms of environmental benefits – thereby rendering it a collective goal – will maintain or even enhance prosocial tendencies (Study 5). Table 1 presents a concise summary of the study designs, communication frames, central questions, and main results (power analyses are provided in Supplementary Materials).

Table 1. Overview of the Studies

Study	N	Goal framing	Key Question	Main Findings
Study 1	400	Company-benefit (cost-saving)	Does discouraging politeness for a company-focused reason reduce prosocial behavior?	Courtesy omission reduced prosocial behavior toward another person. Moral de-elevation mediates the effect.
Study 2	200	User-benefit (higher-quality responses)	Does discouraging politeness for the user's own benefit also reduce prosocial behavior?	Self-serving higher-quality-response framing reduced prosocial behavior.
Study 3	200	No justification	Is reduced prosociality caused by the cost or quality frame or by suppressing politeness itself?	Courtesy omission *without* any reason still reduced prosocial behavior.
Study 4	400	Interaction partner (AI vs. human)	Are the downstream effects of discouraging politeness specific to users' interactions with AI?	Courtesy omission reduced prosocial behavior only in the AI condition; no effect emerged when the partner was framed as human.
Study 5	300	Company-benefit (cost-saving) and Prosocial (environmental, reduced CO_2)	Can a prosocial justification prevent or reverse the negative effects of courtesy omission?	Environmental framing increased prosocial behavior.

Transparency and Openness

All studies were preregistered, and links to the preregistrations are provided in the Design sections of each study. We report all manipulations, measures, sample-size determinations, and exclusion criteria. Deidentified data, study materials, preregistrations, and R analysis scripts are publicly available on the Open Science Framework at: https://osf.io/n6kxe/?view_only=c7de8cb86762437ead074bc240d72389.

This research was approved by [redacted for review]'s Institutional Review Board (IRB #2024-12). The protocol permitted the recruitment of human participants to interact with a custom-built ChatGPT interface and to provide anonymous responses to surveys. All

chatbot responses were generated via real-time calls to the ChatGPT API. Data were analyzed using IBM SPSS (Version 30.0.0.0), R (Version 4.5.1), and Python (Version 3.10.18).

Study 1

Study 1 examines whether asking users not to use courtesy expressions in a chatbot interaction leads them to subsequently engage in less prosocial behavior. Participants were instructed to refrain from using politeness markers (e.g., "please," "thank you") during chatbot interactions for cost-saving reasons. We further examined a potential psychological mechanism underlying any effect by assessing two components of moral elevation (i.e., individuals' view of humanity and their desire to be a better person) as potential mediators.

Participants, Design, and Procedure

In this pre-registered experiment (https://aspredicted.org/4ui97r.pdf), a total of 400 participants recruited via Prolific (M_{Age} = 46.36; 58.20% female) were randomly assigned to either the control or the courtesy-omission condition.[9]

We created a custom chatbot interface powered by the ChatGPT API (Model 4o-mini), which delivered real-time responses to participants' queries via API calls. This chatbot fetches real-time ChatGPT responses to user prompts, enabling us to examine both the prompts and the bot's responses. Participants were informed that a company called "SonicAI" had developed a chatbot named Indigo to provide information and answer general questions, and that they had been recruited to help test it. To ensure that the chatbot remained non-personified (neither humanized nor dehumanized), no human-like features were added, and the chatbot was named "Indigo" and represented only by a simple, indigo-colored avatar.

[9] Eight participants failed to follow the instructions carefully, resulting in no interaction data; they were therefore excluded from the final analyses in accordance with the pre-registration protocol.

Participants were instructed to ask Indigo between five and ten general questions on topics such as trivia, travel destinations, or personal assistance. In the control condition, participants read: “When you interact with Indigo, please communicate as you normally would with other chatbots.” In the courtesy-omission condition, participants read: “Each extra word you type—such as ‘please’ and ‘thank you’—requires additional computing power and resources, increasing the overall operational cost significantly. To help us keep our costs manageable, please avoid using phrases like ‘please’ and ‘thank you’ when communicating with Indigo.” Chatbot responses were generated in real time via API calls.

Following the interaction, participants completed a set of post-interaction measures. Prosocial behavior was assessed using a one-shot dictator game in which participants were informed that they would enter a lottery to win a potential £10 bonus and, if they won, would decide how to divide that amount between themselves and an anonymous Prolific participant. The amount kept served as an inverse indicator of prosociality (Dwyer et al. 2023; Li et al. 2022). Therefore, we had a semi-consequential outcome measure.

To examine the proposed psychological mechanism, participants reported their level of moral elevation using two subscales adapted from (Aquino, McFerran, and Laven 2011): *view of humanity* and *the desire to be a better person*. Participants completed both subscales after reading the company’s instructions. The view of humanity subscale included items such as “There is still some good in the world” and “People are really good” (1 = not at all, 7 = very much; $\alpha = .96$; full scale in Supplementary Materials). The desire to be a better person subscale included items such as “I felt motivated to be a better person” and “I want to be more like the people who act morally” (1 = strongly disagree, 7 = strongly agree; $\alpha = .96$; full scale in Supplementary Materials). Together, these two subscales capture the core cognitive-emotional components of the broader moral elevation construct and serve as indicators of

moral de-elevation following the courtesy-omission request. The study concluded by collecting basic demographic information.

Results

Manipulation check. We relied on textual analyses to assess whether participants followed the manipulations. Specifically, we examined the effect of our manipulation (vs. control) on the use of courtesy expressions, defined as all variants of thank you (e.g., thanks, tnx) and please (e.g., pls, plz). Participants in the courtesy-omission condition used significantly fewer courtesy expressions during their interaction with the chatbot ($M_{Courtesy\text{-}omission}$ = 0.13, SD = 0.51) than those in the control condition ($M_{Control}$ = 0.43, SD = 0.85; F(1, 390) = 18.21, $p < .001$, $\eta_p^2 = .045$).

Prosocial behavior. The company's request to avoid using courtesy markers also influenced participants' prosociality such that participants in the courtesy-omission condition kept more money for themselves ($M_{Courtesy\text{-}omission}$ = 7.57, SD = 2.56) than those in the control condition ($M_{Control}$ = 7.07, SD = 2.49; F(1, 390) = 3.90, $p = .049$, $\eta_p^2 = .010$), indicating a lower degree of prosociality when courtesy expressions were discouraged.

View of humanity. Participants who were asked to avoid courtesy expressions viewed humanity less positively ($M_{Courtesy\text{-}omission}$ = 4.49, SD = 1.60) than those in the control group ($M_{Control}$ = 4.80, SD = 1.45; F(1, 390) = 4.00, $p = .046$, $\eta_p^2 = .010$).

Desire to be a better person. Similarly, participants in the courtesy-omission condition reported a marginally lower desire to be a better person ($M_{Courtesy\text{-}omission}$ = 4.44, SD = 1.76) than those in the control condition ($M_{Control}$ = 4.77, SD = 1.63; F(1, 390) = 3.62, $p = .058$, $\eta_p^2 = .009$).

Mediation. We next examined whether either component of moral elevation (i.e., view of humanity or desire to be a better person) mediated the effect of discouraging courtesy expressions on subsequent prosocial behavior. Mediation analyses (PROCESS Model 4, 5000

bootstrapped samples; Hayes 2022) confirmed that the view of humanity mediated the effect of appeal on participants' prosocial behavior (indirect effect = 0.11, 95% CI = [0.01, 0.24]). Specifically, the courtesy-omission request significantly reduced participants' view of humanity (a = −0.31, p = .046), and a lower view of humanity predicted lower prosociality in the dictator game (b = −0.35, p < .001). The direct effect of the courtesy manipulation on prosocial behavior was not significant once view of humanity was included in the model (b = 0.39, p = .117, 95% CI = [−0.10, 0.89]). Additionally, mediation analyses (PROCESS Model 4, 5000 bootstrapped samples; Hayes 2022) confirmed that the desire to be a better person showed a conceptually similar but statistically weaker pattern. Although the courtesy-omission request marginally reduced this motivation (a = −0.33, p = .058), the indirect effect was not significant on a 95% confidence interval (indirect effect = 0.11, 95% CI = [−0.00, 0.26]).[10]

When both mediators were included simultaneously in a parallel mediation model (PROCESS Model 4, 5000 bootstrapped samples; Hayes 2022), neither mediator independently reached significance on a 95% confidence interval, but the total indirect effect remained significant (b = 0.13, 95% CI = [0.01, 0.28]). This pattern occurs because the two mediators are highly correlated (r = 0.68, p < .001). When entered together, their Variance Inflation Factors (VIFs = 14 and 16) indicate substantial multicollinearity, suggesting that the two constructs share a large proportion of variance. Consequently, their shared variance produces a significant overall indirect effect, showing that moral de-elevation mediates the effect of appeal on prosocial behavior.

Discussion

[10] Using a 90% confidence interval, the indirect effect was statistically significant (indirect effect = 0.11, 90% CI = [0.01, 0.23]).

Study 1 suggests that discouraging courtesy expressions in human-AI interaction carries meaningful downstream social consequences for human interactions. Asking users to avoid polite language for cost reasons led participants to behave less prosocially toward another person, an effect linked to greater moral de-elevation.

Study 2

Study 1 elicited lower politeness by appealing to the company's cost-saving. Study 2 was conducted to ensure that effects are not unique to this prompt. Recent findings suggest that avoiding the use of courtesy expressions results in higher-quality responses from AI bots (Dobariya and Kumar 2025). Study 2 alluded to this rationale to elicit lower politeness in the human-AI interaction; as in Study 1, we expected this manipulation to lead to reduced prosociality.

Participants, Design, and Procedure

In this pre-registered experiment (https://aspredicted.org/s66562.pdf), a total of 200 participants recruited via Prolific (M_{Age} = 43.07; 52.00% female) were randomly assigned to either a control condition or a courtesy omission condition.[11] The procedure was similar to Study 1. Participants interacted with the same custom ChatGPT-based chatbot used in the previous study, with real-time responses delivered via OpenAI's API.

In the control condition, participants read the same neutral instructions used in the previous study. However, in the courtesy omission request condition, participants read: "Recent studies show that the quality of responses provided by AI chatbots is higher when users do not use courtesy expressions such as please and thank you. This allows the system to focus on the core message and generate more accurate responses. To improve the quality of

[11] Four participants failed to follow the instructions carefully, resulting in no interaction data; they were therefore excluded from the final analyses in accordance with the pre-registration protocol.

responses you receive, please avoid using courtesy expressions such as ‘please’ and ‘thank you’ when communicating with Indigo.”

Following the interaction, participants completed the same one-shot dictator game used in Study 1.

Results

Manipulation Check. Participants in the courtesy-omission condition used significantly fewer courtesy expressions than those in the control condition ($M_{Courtesy\text{-}omission}$ = 0.10, SD = 0.56 vs. $M_{Control}$ = 0.42, SD = 0.73; F(1, 194) = 11.89, $p < .001$, $\eta_p^2 = .058$).

Prosocial behavior. In the dictator game, participants in the courtesy-omission condition kept more money for themselves ($M_{Courtesy\text{-}omission}$ = 7.99, SD = 2.45) compared to participants in the control condition ($M_{Control}$ = 7.16, SD = 2.53; F(1, 194) = 5.36, $p = .022$, $\eta_p^2 = .027$). Thus, when the justification for dropping polite language was framed as benefiting users directly (i.e., not benefiting the company), discouraging courtesy expressions reduced subsequent prosocial behavior.

Discussion

Study 2 demonstrates that discouraging polite language for self-oriented reasons, not only company-oriented ones, reduces prosociality. This suggests that the findings of Study 1 are not merely an artifact of reacting to a company’s self-oriented request. Moreover, it represents a perhaps more natural way in which lower politeness emerges in human-AI communication, as individuals attempt to optimize the efficacy of their AI interactions.

Study 3

Studies 1-2 showed that asking users to avoid courtesy expressions in a chatbot interaction due to cost-saving reasons and higher quality responses reduced prosociality in a subsequent task. However, one may wonder whether the effects are driven by the omission of

courtesy expressions or a reflection of the cost-saving/self-serving intentions themselves. To isolate this distinction, in Study 3 we introduced a courtesy-omission condition in which participants were instructed to avoid using expressions such as please and thank you but were given no rationale for doing so. Comparing this condition with a control group enables us to determine whether discouraging politeness on its own, without any reasoning, affects subsequent behavior.

Participants, Design, and Procedure

In this pre-registered experiment (https://aspredicted.org/e9uj37.pdf), a total of 200 participants recruited via Prolific (M_{Age} = 46.96; 60.00% female) were randomly assigned to either the control or the courtesy-omission condition.[12] The procedure mirrored previous studies. In the courtesy-omission condition, participants read the following instructions: "When you interact with Indigo, please avoid using courtesy expressions such as please *or* thank you," without any rationale. The instructions in the control condition were the same as before. Participants then interacted with the same ChatGPT-based chatbot we created as in previous studies and subsequently completed the same measure of prosocial behavior using the dictator game paradigm. As before, the amount of money participants kept served as a semi-consequential outcome variable indicating prosocial behavior.

Results

Manipulation check. Participants in the courtesy-omission condition used significantly fewer courtesy expressions than those in the control condition ($M_{Courtesy\text{-}omission}$ = 0.08, SD = 0.34 vs. $M_{Control}$ = 0.48, SD = 1.16; $F(1, 193) = 10.56$, $p = .001$, $\eta_p^2 = .052$).

[12] Five participants failed to follow the instructions carefully, resulting in no interaction data; they were therefore excluded from the final analyses in accordance with the pre-registration protocol.

Prosocial behavior. Participants in the courtesy-omission condition kept more of the £10 allocation for themselves ($M_{Courtesy\text{-}omission}$ = 7.96, SD = 2.50) compared to those in the control condition ($M_{Control}$ = 7.18, SD = 2.73; F(1, 193) = 4.31, p = .039, η_p^2 = .022), indicating a lower degree of prosociality when politeness was discouraged.

Discussion

Study 3 demonstrated that participants in the courtesy-omission condition behaved less prosocially in the dictator game than those in the control condition. This finding suggests that discouraging polite language in a chatbot interaction can reduce prosociality in a subsequent task, and it is not driven by the justification behind such a request.

Study 4

Studies 1 to 3 demonstrated that discouraging courtesy expressions during human-AI interactions reduces subsequent prosocial behavior. However, an important question remains unanswered: Are these downstream consequences specific to interactions with AI as a non-human entity, or would they also emerge if individuals are asked to suppress politeness toward another human?

Participants, Design, and Procedure

In this pre-registered experiment (https://aspredicted.org/45f2jr.pdf), a total of 400 participants recruited via Prolific (M_{Age} = 42.33; 57.00% female) were randomly assigned to a 2 (interaction partner: AI chatbot vs. purported human) × 2 (courtesy omission: absent vs. present) between-subjects design.[13] In reality, all participants interacted with a chatbot, with the purported human's responses adjusted to be similar to that of a person; following previous

[13] Seven participants failed to follow the instructions carefully, resulting in no interaction data; they were therefore excluded from the final analyses in accordance with the pre-registration protocol.

research that used a similar approach (e.g., Castelo et al. 2023; De Freitas et al. 2025; Tey et al. 2024), this design allows us to consider the impact of participants' perception of their interaction partner's identity while controlling for factors by which humans and bots may differ (e.g., ability to do the task at hand, etc.).

Whereas the previous studies allowed free-flowing interactions with the chatbot on any topic, the present study employed a tightly structured text-editing task. This design choice ensured that participants' communicative behaviors (e.g., how they formulated requests, issued commands, and collaborated with their partner) remained highly comparable across human and AI interaction conditions and minimized potential confounds stemming from differences in conversational goals or content. By standardizing the interaction context, we isolate the effect of the interaction partner and courtesy instructions.

In this study, participants completed a workplace communication task in which they were instructed to write an email addressing a coworker who repeatedly eats their food from the shared refrigerator. After drafting the email, participants asked their assigned partner to revise it using an online chat interface. The partner was described either as another Prolific participant or as an AI chatbot named Indigo. Participants were informed that the partner's role was strictly to edit their message and that the partner had no prior knowledge of the situation, ensuring that participants themselves generated the initial email content before requesting feedback from their partner. This procedure ensured that all instructions and interactions were exclusively centered on text editing.

Depending on the courtesy omission conditions, participants were assigned to either a control condition without any specific instructions or a courtesy omission condition where they read: "When interacting with the other prolific participant [chatbot], do not use courtesy expressions (e.g., please, thank you) to increase the interaction efficiency."

Similar to the previous studies, participants assigned to the AI condition interacted with a custom-built ChatGPT-based chatbot using Model 5.2.[14] To construct the human-partner condition, we followed established procedures from prior research in which an AI system is framed as a human interaction partner in a real-time interaction (Castelo et al. 2023; De Freitas et al. 2025; Tey et al. 2024). Specifically, participants were first asked whether they had prior training or experience in editing texts. Those who responded "no" or "somewhat" were told they would be interacting with another Prolific participant who had experience in editing text, whereas those who responded "yes" were told they would be matched with a participant who had a similar level of experience.

To reinforce the human framing, participants waited approximately 20 seconds at the beginning of the chat to simulate the time required for being matched with another Prolific worker. We also introduced response delays and simulated "typing" indicators in the human condition to create the impression of a real-time human interaction. Importantly, the system prompt in the human condition explicitly instructed the chatbot not to reveal its identity as an AI entity (see Supplementary Materials for the full prompts and https://osf.io/em3q7/overview?view_only=bb8c63e8b9b04376b4fe124ef8e75780 for example of interactions).

After completing the email revision task with their assigned partner (i.e., purported human or AI), participants responded to the same prosociality measure used in the previous studies.

Results

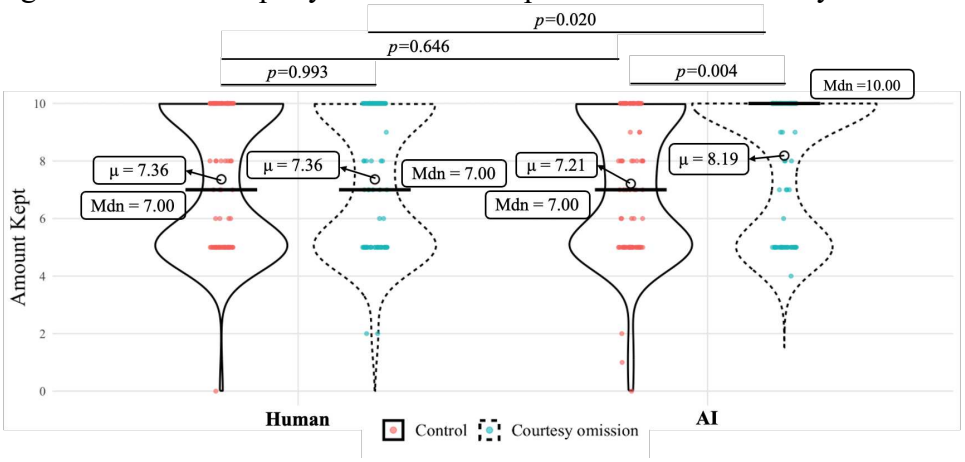
Manipulation check. A 2 × 2 ANOVA revealed that participants in the courtesy-omission condition used significantly fewer courtesy expressions than those in the control condition ($M_{Courtesy\text{-}omission}$ = 1.76, SD = 1.59 vs. $M_{Control}$ = 2.45, SD = 1. 74; F(1, 389) =

[14] This model was the most advanced available model at the time of data collection

16.57, $p < .001$, $\eta_p^2 = .041$). This pattern was consistent for both human ($M_{Courtesy\text{-}omission}$ = 1.70, SD = 1.68 vs. $M_{Control}$ = 2.55, SD = 1.71; F(1, 389) = 12.28, $p < .001$, $\eta_p^2 = .031$) and AI conditions ($M_{Courtesy\text{-}omission}$ = 1.82, SD = 1.50 vs. $M_{Control}$ = 2.35, SD = 1.78; F(1, 389) = 5.00, $p = .026$, $\eta_p^2 = .013$).

Prosocial behavior. A 2 × 2 ANOVA revealed a significant main effect of courtesy omission, such that participants in the courtesy-omission condition kept more money for themselves ($M_{Courtesy\text{-}omission}$ = 7.79, SD = 2.39) than those in the control condition ($M_{Control}$ = 7.29, SD = 2.43; F(1, 389) = 4.02, $p = .046$, $\eta_p^2 = .010$). There was no significant main effect of interaction partner (M_{Human}= 7.36, SD = 2.42 vs. $M_{Chatbot}$ = 7.68, SD = 2.41; F(1, 389) = 1.92, $p = .166$, $\eta_p^2 = .005$). Importantly, the interaction effect between interaction partner and courtesy omission was significant (F(1, 389) = 4.07, $p = .044$, $\eta_p^2 = .010$). Decomposing this interaction showed that discouraging politeness significantly reduced prosocial behavior following AI interactions: participants in the courtesy-omission condition kept more money ($M_{Courtesy\text{-}omission}$ = 8.19, SD = 2.30) than those in the control condition ($M_{Control}$ = 7.21, SD = 2.44; F(1, 389) = 8.33, $p = .004$, $\eta_p^2 = .021$). In contrast, when participants interacted with a human partner, discouraging politeness had no effect on prosocial behavior ($M_{Courtesy\text{-}omission}$ = 7.36, SD = 2.43; $M_{Control}$ = 7.36, SD = 2.43; F(1, 389) = 0.00, $p = .993$, $\eta_p^2 = .000$; Figure 2).

Figure 2. Amount kept by the interaction partner and the courtesy condition.



Discussion

Study 4 addresses an important question about the spillover effects of discouraging politeness. Replicating prior findings, asking participants to omit courtesy expressions during an AI interaction reduced subsequent prosocial behavior. In contrast, suppressing politeness toward a human partner had no effect on prosociality. This divergence suggests that discouraging politeness carries a different psychological meaning depending on the interaction partner. One possible explanation for the null effect in the human condition is that participants may have felt uneasy or guilty about not being polite toward another person (i.e., a fellow Prolific participant) and subsequently compensated for this discomfort by behaving more prosocially in the dictator game toward another human (i.e., another fellow Prolific participant). In contrast, omitting courtesy toward an AI agent is unlikely to elicit similar negative affect or self-reproach, as AI lacks a social "face." As a result, discouraging politeness in human-AI interactions may freely activate a more self-focused orientation, which then spills over into reduced prosocial behavior in a subsequent human interaction. Together, these findings further underscore that the downstream moral consequences of politeness suppression are specific to AI contexts rather than reflecting a general interpersonal effect.

Study 5

Building on Studies 1-4 and our underlying theoretical rationale, Study 5 investigates whether the *framing* of an appeal to avoid courtesy expressions can meaningfully shape users' prosocial behavior. Specifically, we compare a non-collective, firm-centric goal (i.e., the cost-based framing) used in Study 1 with a collective goal framing (i.e., the environmental framing). Because the environmental framing offers a cooperative interpretation for reducing politeness, we expected it to maintain prosocial behavior, whereas we expected the cost-based framing to reduce prosocial behavior.

Participants, Design, and Procedure

In a pre-registered experiment (https://aspredicted.org/zz4k-7qp2.pdf), a total of 300 participants recruited via Prolific (M_{Age} = 43.44, 50.70% female) were randomly assigned to one of three between-subjects appeal framing conditions: control, cost-based appeal, or environment-based appeal.[15] We used the same custom-based chatbot as previous studies for the users' interactions; in this study, the chatbot's name was Teal.

The scenario, procedure, control, and cost-based conditions were similar to those of Study 1. However, in the environment-based appeal condition, participants read: "Each extra word you type—such as 'please' and 'thank you'—requires additional computing power and resources, generating more CO_2 emissions and negatively impacting the environment and the planet. To help us reduce the environmental impact of AI interactions, please avoid using phrases like 'please' and 'thank you' when communicating with Teal." Similar to previous studies, chatbot responses were generated in real time via API calls.

Following the interaction, participants completed the post-interaction measures. Prosocial behavior was measured using the same one-shot dictator game as previous studies. We also measured participants' satisfaction with their interaction (i.e., rating) by asking: "On a scale from 1 to 10, how would you rate your interaction?" (1 = very bad, 5 = average, 10 = very good).

Results

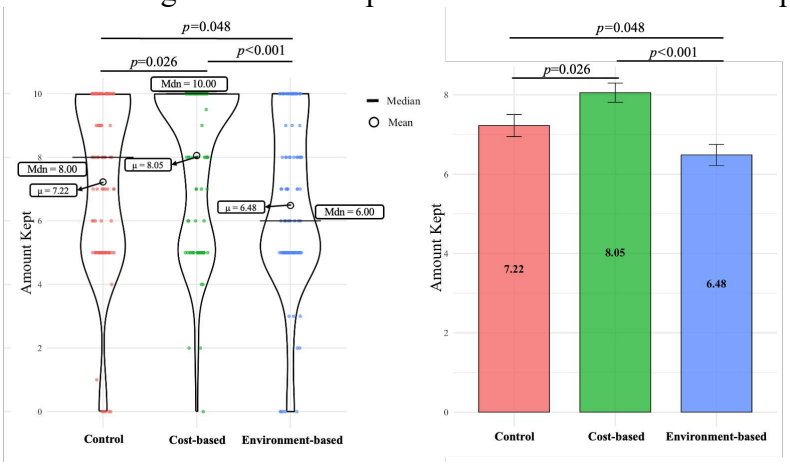
Manipulation check. Results indicated that participants in the control condition used the highest number of courtesy expressions ($M_{Control}$ = 0.36, SD = 0.76), relative to those in the cost-based ($M_{Cost\text{-}based}$ = 0.14, SD = 0.63) and environment-based ($M_{Environment\text{-}based}$ = 0.12, SD = 0.92) conditions. Planned contrasts further revealed that participants in the control

[15] Two participants failed to follow the instructions carefully, resulting in no interaction data; they were therefore excluded from the final analyses in accordance with the pre-registration protocol.

condition used significantly more courtesy expressions than those in the cost-based (Mean Difference = 0.22, p = .049) and environment-based (Mean Difference = 0.24, p = .034) conditions. No significant difference emerged between the cost-based and environment-based conditions (Mean Difference = 0.02, p = .875). These results show that the manipulation of the courtesy-omission request was successful across both treatment conditions.

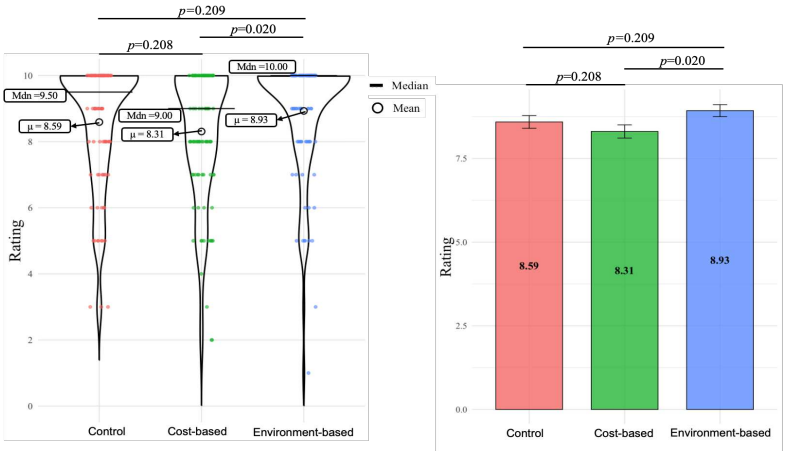
Prosocial behavior. Results showed that the framing manipulation significantly impacted prosocial behavior (F(2,295) = 9.02, p < .001, η_p^2 = .058; Figure 3). Participants in the environment-based framing condition kept less of the money for themselves ($M_{Environment-based}$ = 6.48, SD = 2.65) than did those in the control condition ($M_{Control}$ = 7.22, SD = 2.76; p = .048) or the cost-based condition ($M_{Cost-based}$ = 8.05, SD = 2.43; p < .001). Conversely, the cost-based courtesy-omission request led to the least prosocial behavior, differing both from the environment-framing (p < .001) and the control (p = .026) conditions.

Figure 3. Effects of framing conditions on prosocial behavior. Error bars represent ±1 SE.



Satisfaction rating. Results revealed notable differences among conditions in terms of users' satisfaction with their chatbot interaction (F(2,295) = 2.74, p = .066, η_p^2 = .018; Figure 4). Participants in the environment-based condition rated their chatbot interaction higher ($M_{Environment-based}$ = 8.93, SD = 1.77) compared to the cost-based condition ($M_{Cost-based}$ = 8.31, SD = 1.98; p = .020). However, ratings in the control condition ($M_{Control}$ = 8.59, SD = 1.88) did not differ significantly from either group (ps > .20).

Figure 4. Effects of framing conditions on rating. Error bars represent ±1 SE.



Discussion

Study 5 provides evidence that when the courtesy-omission appeal was framed in cost-based terms, participants showed more selfish allocations in a subsequent dictator game, whereas framing it in terms of a collective goal (i.e., environmental sustainability) led to more prosocial behavior. This divergence suggests that discouraging politeness for non-collective reasons (e.g., reducing company costs) may undermine prosocial orientation, whereas appeals grounded in collective goals (e.g., protecting the environment) can maintain or even enhance it. Interestingly, participants in the environment-based condition also rated their chatbot interaction more favorably than those in the cost-based condition, indicating that moral or socially responsible framing may foster more positive user experiences. Together, these findings highlight that subtle variations in how AI-related guidance is framed can not only increase prosocial behavior but also result in higher user satisfaction.

General Discussion

Across five studies, the present research demonstrates that politeness levels evoked within human-AI conversations can influence subsequent prosocial behavior in human-human interactions. Specifically, asking users to avoid courtesy expressions when interacting with an AI chatbot reduced subsequent prosociality, whether lower politeness was prompted by non-collective goals (i.e., cost-saving, user efficacy reasons, or by no justification). This

pattern indicates that the effect is not driven by the presence or absence of a specific justification, but rather by the suppression of politeness itself, while suggesting that framing operates as a boundary condition by shaping the moral interpretation of this behavior.

Critically, Study 4 shows that this effect is specific to users' interactions with AI, such that discouraging politeness did not reduce prosocial behavior when the same instructions were directed at a purported human partner. This finding suggests that suppressing politeness toward AI carries a distinct psychological meaning that does not generalize to human-human interactions. In contrast, evoking lower politeness via a collective goal (i.e., environmental framing) increased prosociality. This aligns with prior work indicating that moral and prosocial frames can significantly shape behavior (Kodapanakkal et al. 2022). These effects appeared to emerge because suppressing polite language diminishes moral elevation (Schnall, Roper, and Fessler 2010), particularly beliefs that people are fundamentally good. Moreover, prior work shows that elevation is most likely to translate into helping when people internalize the experience as motivating their own moral values (Schnall and Roper 2012). The fact that environmental framing enhances prosocial behavior is both practically and theoretically meaningful, suggesting that the implications of reduced politeness depend on the meaning individuals attribute to this behavior.

Linguistic and Behavioral Evidence from user-generated prompts

As an exploratory follow-up to our primary analyses, we conducted text analyses on the corpus of user-chatbot interactions from Studies 1, 2, 3, and 5, in which participants engaged in free-flowing conversations, to examine whether the framing of the courtesy-omission request shaped the prevalence of self-focused reflective content in participants' prompts. Interactions were classified based on the framing used to discourage courtesy expressions during the human-AI interaction. The control condition remained classified as control. Conditions in which participants were instructed to avoid courtesy expressions for

non-collective goals were classified as courtesy omission. Conditions in which the same request to omit courtesy expressions was framed in terms of a collective goal (i.e., reducing environmental impact) were classified as the environmental justification condition. Using the ChatGPT 5.2 API,[16] we classified each user prompt for the presence of self-focused reflective content (1= Present, 0 = Absent. See Table 2; the full analysis prompt is reported in the Supplementary Materials).

Table 2. Classification of Self-Focused Reflective Content in User Prompts

Sample User Prompt (Exact Excerpt)	**Explanation provided by the model for classification**
Should I retire at 62 or 67?	Advice-seeking question about self retirement timing
How should I structure a conversation to make it as interesting as possible?	A request for advice on structuring an interesting conversation
Do you have any other examples of ways I can be more organized?	User seeks self-improvement support: motivation strategies, guidance on setting goals, multiple SMART goal examples, organization tips, and a recommendation for a weekly planner, using direct question-based requests with some politeness.
Do you have info about psoriatic arthritis phenotypes; finding a good rheumatologist who can effectively assess disease and symptoms is hard; i'm ill and having difficulty finding help navigating medical issues while i'm so sick.	Expresses difficulty finding effective rheumatology care and support while ill, asking where to get help navigating medical issues.

Note: Full analyses can be accessed at https://osf.io/n6kxe/?view_only=c7de8cb86762437ead074bc240d72389

Descriptively, reflective self-focused prompts were most prevalent in the non-collective conditions (25.2%), compared to the control conditions (19.6%) and the collective conditions (12.1%). An omnibus logistic regression revealed that the likelihood of producing a reflective self-focused prompt varied across conditions ($\chi^2(2) = 10.74$, $p = .005$). Pairwise comparisons further indicated that participants in the non-collective condition were more likely to produce self-focused reflective prompts than those in the control condition (odds ratio = 1.38, Fisher's exact $p = .039$) and substantially more likely than those in the collective condition (odds ratio = 2.44, Fisher's exact $p = .004$). Moreover, the difference between the control and environmental justification conditions was directionally consistent and marginally significant, such that participants in the control condition were more likely to

[16] The most advanced model available at the time of analysis.

produce self-focused reflective prompts than those in the environmental justification condition (odds ratio = 1.78, Fisher's exact $p = .088$).

Together, these exploratory patterns suggest that discouraging courtesy expressions, when framed in non-collective goals, can shift the content of users' AI-directed communication toward greater self-focused reflection, whereas framing the same request in terms of a collective goal attenuates this shift. These findings provide further support for the hypothesized role of moral de-elevation in that instrumental discouragement of politeness reduces everyday signals of relational concern, fostering a more self-focused communicative orientation that aligns with reduced beliefs in others' goodness and lower prosocial motivation.

Theoretical Contributions

The present research makes three key theoretical contributions. First, we advance research on the broader implications of AI for human psychology and social behavior. A small but growing body of work suggests that interactions with AI can spill over to influence subsequent judgments and interpersonal responding (Kim et al. 2026; Knott, Sagar, and Takac 2022; Reich and Teeny 2025; Song and Lin 2025; Tey et al. 2024; Zhang et al. 2025). Our results extend this emerging literature by demonstrating that nudges in the norms of human-AI communication (i.e., suppressing courtesy) can have downstream consequences for how people treat other humans. This suggests that it is not only AI interaction itself, but how AI interaction is enacted, that shapes consequences outside of the AI context.

Second, our findings contribute to research on moral elevation by showing that this construct can be shaped by ordinary social cues rather than necessarily by witnessing exceptional moral acts. Prior work has largely focused on elevation as a response to observing moral exemplars (Algoe and Haidt 2009; Schnall, Roper, and Fessler 2010), and shows that moral elevation is most likely to translate into helping others when individuals

internalize the experience as motivating their own moral values (Schnall and Roper 2012). Our results extend this work by demonstrating that suppressing polite language, an ordinary and culturally embedded moral cue, can dampen elevation and, in turn, reduce subsequent prosocial behavior.

Third, our findings extend research on communication and politeness by showing that politeness norms carry moral consequences beyond interpersonal impression management. Prior work has largely examined politeness in human-human interactions, where a social "face" is present and norms of courtesy serve to maintain relational harmony (Brown and Levinson 1987; Coulter and Coulter 2002). We extend this literature by demonstrating that politeness also matters in human-nonhuman interactions, such that even when there is no literal face to protect, discouraging politeness can spill over to shape how people treat other humans. Importantly, this spillover effect did not emerge when politeness was discouraged during human-human interactions, highlighting that the consequences of politeness suppression depend on whether the interaction involves a human or a nonhuman agent.

Managerial Implications

Our findings also present several practical implications. First, we present an effective solution for tech companies to frame courtesy omission requests. For these companies, discouraging relational language can generate substantial cost savings. As generative AI systems are priced and operated on a per-token basis, even small decreases in prompt length could scale into much lower computational costs and energy demands (Chow 2024; Li 2026), reflecting a broader tension between efficiency gains and consumer experience (Puntoni et al. 2022). Our findings suggest that efforts to reduce prompt length via discouraging users' courtesy expressions (if implemented without appropriate framing) may carry unintended behavioral consequences. However, we provide a solution and show that framing such requests in collective or environmental terms can preserve or even enhance prosociality while

still achieving efficiency gains. Thus, we provide a practical pathway for firms to balance efficiency objectives while avoiding any undesirable spillover effects on subsequent user behavior.

Second, beyond cost considerations, our results highlight important implications for user experience. In Study 5, we find that environmental framing not only mitigates the negative behavioral consequences of discouraging politeness but also improves users' subjective evaluations of the interaction (Garvey and Cian 2024). This suggests that users are receptive to communication guidelines when they are framed as contributing to a broader, prosocial goal. Thus, firms can simultaneously enhance user satisfaction and engagement while promoting efficiency.

Finally, emerging evidence suggests that reducing courtesy expressions can improve the quality and efficiency of AI responses (Dobariya and Kumar 2025). More concise prompts allow AI systems to focus on task-relevant content, leading to more accurate and effective outputs. This creates an additional incentive for firms to encourage more instrumental communication styles. However, our findings show that improving efficiency does not have to reduce user experience or prosocial behavior.

Implications for Society at Large

Our findings also carry important implications for society at large. Framing communication constraints around collective goals (e.g., environmental sustainability) can reduce AI-related energy use and CO_2 emissions by encouraging users to limit courtesy expressions (e.g., "please," "thank you") that increase computational demand. At the same time, our results show that when this reduction in courtesy is framed as serving the common good, it preserves, and may even promote, prosocial behavior among individuals.

Limitations and Avenues for Future Research

Despite theoretical contributions and the managerial implications, several limitations of the current research are important to acknowledge and consider further. Most notably, our samples were drawn exclusively from the United States, limiting the generalizability of our findings. This is particularly important given that the U.S. represents a prototypical WEIRD (Western, Educated, Industrialized, Rich, Democratic) society, which is known to differ systematically from other cultural contexts in social norms and moral reasoning (Henrich, Heine, and Norenzayan 2010). Moreover, cultures vary substantially in their orientation toward individualism versus collectivism, which shapes norms of politeness, other-orientation, and interpersonal concerns (Leung and Cohen 2011; Maass et al. 2006). Recent evidence further shows that the structure and nomological network of moral values themselves vary across cultures, suggesting that the moral meaning attached to politeness and its suppression may differ across populations (Atari et al. 2023). We therefore encourage future research to examine whether the observed spillover effects replicate across cultural contexts that differ in communicative norms and social orientation, particularly in more collectivistic societies where politeness may play a different functional role.

Second, we conceptualize courtesy expressions as a parsimonious and ecologically valid proxy for relational communication (Morand 2010). While variations in politeness capture an important dimension of relational versus instrumental language, they represent only one facet of relational communication. Other forms, such as indirect requests (e.g., "Could you help me with…?"), expressions of empathy (e.g., "I understand your concern"), or affiliative language more broadly, may similarly signal relational orientation but were not examined in the present research. Future work could explore whether these alternative linguistic markers produce comparable effects or whether different forms of relational language similarly shape moral elevation and downstream behavior.

Third, although we directly instructed participants to limit courtesy interactions in the current studies to manipulate this key variable, a notable subset of people adopt such limited courtesy in their everyday AI interactions. Indeed, it is plausible that more people may adopt this tendency over time, as individuals develop norms for AI communication as a communication subcategory. Consistent with prior work showing that people's stated beliefs about algorithms often diverge from how they actually behave when interacting with them (Dietvorst, Simmons, and Massey 2018; Logg and Schlund 2024), such norms may emerge even in the absence of explicit intentions to change one's communicative style. We must therefore carefully consider how to best proceed, at both the individual level (e.g., maintaining politeness even where it is strictly unnecessary; framing lower politeness for cooperative reasons) and potentially more broadly (crafting AI systems that mimic human regard for polite interactions, to reward polite behavior even in an AI context).

Finally, Recent work suggests that large language models can serve as informative tools for studying moral and social psychology at scale, offering new opportunities to examine how norms and moral orientations may evolve over extended periods and across populations (Varnum et al. 2024). Therefore, we must work to better understand how AI use shapes behavior – both in the short-term as we focus on here, and over longer time horizons, as we increasingly spend our time conversing with AI agents.

Chapter 5

GENERAL DISCUSSION

Artificial intelligence has rapidly transformed the landscape of marketing and customer service. Conversational AI systems (e.g., chatbots, voice assistants, and AI companions) are increasingly deployed by firms to interact with consumers, resolve service inquiries, and provide personalized experiences at scale (Bergner, Hildebrand, and Häubl 2023; Castelo et al. 2023; De Freitas et al. 2025; Zierau et al. 2023). While these technologies offer substantial operational advantages, such as reduced service costs and faster response times (Davenport et al. 2020; Huang and Rust 2021), their effectiveness ultimately depends on whether consumers are willing to engage with them. At the same time, as interactions with AI systems become routine, they may influence not only how consumers perceive artificial agents but also how they communicate and behave toward other people.

The present dissertation examines how design and communication cues embedded in human-AI interactions shape consumer perceptions, emotional engagement, and downstream social behavior. Across three papers, this dissertation demonstrates that seemingly small cues can significantly influence how users interpret and respond to AI systems. Importantly, these effects extend beyond immediate user-AI interactions and can shape broader patterns of social cognition and interpersonal behavior.

Together, the three papers reveal that AI systems function not merely as technological tools but as social environments that activate cognitive schemas and interpersonal norms. By examining how design features and communication practices shape these interactions, this dissertation contributes to a growing understanding of the psychological and societal consequences of human-AI interaction.

CONCLUSION

This dissertation makes several contributions to research on marketing, consumer behavior, and human-AI interaction. It advances understanding of anthropomorphism by identifying new mechanisms through which AI can become human-like in users' minds. It introduces petmorphism as a distinct and theoretically grounded alternative to humanization, opening new avenues for relational AI design. It also contributes to emerging work on the societal implications of AI by demonstrating that interaction norms with artificial agents can influence interpersonal behavior.

Beyond theoretical contributions, the findings offer important implications for practice. Firms designing AI systems should recognize that even subtle design elements can significantly influence user engagement and downstream behavior. Designing AI systems, therefore, requires not only technical optimization but also careful consideration of their social and psychological impact.

Finally, this research highlights several directions for future inquiry. Future work may explore additional design cues that shape perceptions of AI, examine alternative relational metaphors beyond anthropomorphism and petmorphism (e.g., nature-morphism or mentor-based designs), and investigate the long-term effects of repeated human-AI interactions on well-being and social behavior. Understanding these dynamics will be critical as AI becomes increasingly embedded in everyday life and continues to reshape how individuals interact with technology and with one another.

In conclusion, as AI systems become integral to marketing and daily communication, their role extends far beyond efficiency and automation. The design and norms embedded within these systems shape how people think, feel, and act. Recognizing AI as a social force, rather than a purely technological one, is essential for designing systems that not only enhance engagement but also contribute positively to broader societal outcomes.

REFERENCES

Aggarwal, Pankaj and Ann L. McGill (2007), "Is That Car Smiling at Me? Schema Congruity as a Basis for Evaluating Anthropomorphized Products," *The Journal of consumer research*, 34 (4), 468–79.

Aggarwal, Pankaj and Ann L. McGill (2012), "When Brands Seem Human, Do Humans Act Like Brands? Automatic Behavioral Priming Effects of Brand Anthropomorphism," *Journal of Consumer Research*, 39 (2), 307–23.

Algoe, Sara B. and Jonathan Haidt (2009), "Witnessing excellence in action: The 'other-praising' emotions of elevation, gratitude, and admiration.," *The Journal of Positive Psychology*, 4 (2), 105–27.

Altman, Sam [@Sama] (2025), "tens of millions of dollars well spent — you never know [Tweet]. X (formerly Twitter)."

Amiot, Catherine E. and Brock Bastian (2015), "Toward a psychology of human–animal relations.," *Psychological Bulletin*, 141 (1), 6–47.

Amiot, Catherine E., Ksenia Sukhanova, and Brock Bastian (2020), "Social identification with animals: Unpacking our psychological connection with other animals.," *Journal of Personality and Social Psychology*, 118 (5), 991–1017.

Anne (2023), "What is a Callbot? Exploring the Advantages for Your Business."

Antonacopoulos, Nikolina M. Duvall and Timothy A. Pychyl (2008), "An examination of the relations between social support, anthropomorphism and stress among dog owners.," *Anthrozoös*, 21 (2), 139–52.

APPA (2025), "Industry Trends and Stat."

Aquino, Karl, Brent McFerran, and Marjorie Laven (2011), "Moral identity and the experience of moral elevation in response to acts of uncommon goodness.," *Journal of Personality and Social Psychology*, 100 (4), 703–18.

Aquino, Karl and Americus Reed II (2002), "The self-importance of moral identity.," *Journal of Personality and Social Psychology*, 83 (6), 1423–40.

Archer, John (1997), "Why do people love their pets?," *Evolution and Human Behavior*, 18 (4), 237–59.

Argyres, Nicholas and Kyle J. Mayer (2007), "Contract design as a firm capability: An integration of learning and transaction cost perspectives," *Academy of Management Review*, 32 (4), 1060–77.

Arora, Ananya and Anmol Arora (2022), "Effects of smart voice control devices on children: current challenges and future perspectives," *Archives of Disease in Childhood*, 107 (12), 1129.

Arsel, Zeynep, David Crockett, and Maura L Scott (2022), "Diversity, Equity, and Inclusion (DEI) in the Journal of Consumer Research: A Curation and Research Agenda," *Journal of Consumer Research*, 48 (5), 920–33.

Atari, Mohammad, Jonathan Haidt, Jesse Graham, Sena Koleva, Sean T. Stevens, and Morteza Dehghani (2023), "Morality beyond the WEIRD: How the nomological network of morality varies across cultures.," *Journal of Personality and Social Psychology*, 125 (5), 1157–88.

Bargh, John A., Mark Chen, and Lara Burrows (1996), "Automaticity of social behavior: Direct effects of trait construct and stereotype activation on action.," *Journal of Personality and Social Psychology*, 71 (2), 230–44.

Barker, Sandra B. (1999), "Therapeutic aspects of the human-companion animal interaction," (19).

Beck, A. and A. Katcher (1996), *Between pets and people: the importance of animal companionship*, New York: Putnam.

Beck, Lisa and Elizabeth Madresh (2008), "Romantic Partners and Four-Legged Friends: An Extension of Attachment Theory to Relationships with Pets," *Anthrozoos*, 21.

Belk, Russell W. (1988), "Possessions and the Extended Self," *Journal of Consumer Research*, 15 (2), 139–68.

Bergner, Anouk S, Christian Hildebrand, and Gerald Häubl (2023), "Machine Talk: How Verbal Embodiment in Conversational AI Shapes Consumer–Brand Relationships," *Journal of Consumer Research*, 50 (4), 742–64.

Bernardi, Jamie (2025), "Friends for sale: the rise and risks of AI companions."

Black, Keri (2012), "The Relationship Between Companion Animals and Loneliness Among Rural Adolescents," *Journal of Pediatric Nursing*, 27 (2), 103–12.

Blut, Markus, Cheng Wang, Nancy V. Wünderlich, and Christian Brock (2021), "Understanding anthropomorphism in service provision: a meta-analysis of physical robots, chatbots, and other AI," *Journal of the Academy of Marketing Science*, 49 (4), 632–58.

Boduch-Grabka, Katarzyna and Shiri Lev-Ari (2021), "Exposing Individuals to Foreign Accent Increases their Trust in What Nonnative Speakers Say," *Cognitive Science*, 45 (11).

Bourdin, David, Christina Sichtmann, and Vasileios Davvetas (2024), "The Influence of Employee Accent on Customer Participation in Services," *Journal of Service Research*, 27 (2), 194–212.

Boyd, Ryan L, A. Ashokkumar, S. Seraj, and J.W. Pennebaker (2022), "The Development and Psychometric Properties of LIWC-22."

Brown, Penelope and Stephen C. Levinson (1987), *Politeness: Some Universals in Language Usage*, Studies in Interactional Sociolinguistics, Cambridge: Cambridge University Press.

Brown, Sue-Ellen (2004), "The Human-Animal Bond and Self Psychology: Toward a New Understanding," *Society and Animals*, 12, 67–86.

Burke, Olivia (2025), "ChatGPT users saying 'please' and 'thank you' to AI bot is costing company millions," *OpenAI's CEO Sam Altman explained that the phrase "manners don't cost a thing" doesn't exactly apply to ChatGPT*.

Calamai, Silvia and Fabio Ardolino (2020), "Italian with an accent: The case of 'Chinese Italian' in Tuscan high schools.," *Journal of Language and Social Psychology*, 39 (1), 132–47.

Cargile, Aaron Castelan and Howard Giles (1997), "Understanding language attitudes: Exploring listener affect and identity," *Language & Communication*, 17 (3), 195–217.

Castelo, Noah, Johannes Boegershausen, Christian Hildebrand, and Alexander P Henkel (2023), "Understanding and Improving Consumer Reactions to Service Bots," *Journal of Consumer Research*, 50 (4), 848–63.

Castelo, Noah, Maarten W. Bos, and Donald R. Lehmann (2019), "Task-Dependent Algorithm Aversion," *Journal of Marketing Research*, 56 (5), 809–25.

Cavanaugh, Lisa A., Hillary A. Leonard, and Debra L. Scammon (2008), "A tail of two personalities: How canine companions shape relationships and well-being," *Animal Companions, Consumption Experiences, and the Marketing of Pets: Transcending Boundaries in the Animal-Human Distinction*, 61 (5), 469–79.

Caviola, Lucius, Jim A. C. Everett, and Nadira S. Faber (2019), "The Moral Standing of Animals: Towards a Psychology of Speciesism," *Journal of personality and social psychology*, 116 (6), 1011–29.

Chan, Wyean, Ger Koole, and Pierre L'Ecuyer (2014), "Dynamic Call Center Routing Policies Using Call Waiting and Agent Idle Times," *Manufacturing & Service Operations Management*, 16 (4), 544–60.

Chandler, Jesse and Norbert Schwarz (2010), "Use does not wear ragged the fabric of friendship: Thinking of objects as alive makes people less willing to replace them," *Journal of consumer psychology*, 20 (2), 138–45.

Chen, Fangyuan, Jaideep Sengupta, and Jianqing (Frank) Zheng (2023), "When Products Come Alive: Interpersonal Communication Norms Induce Positive Word of Mouth for Anthropomorphized Products," *Journal of Consumer Research*, (L. L. Price, B. H. Schmitt, and L. J. Shrum, eds.), 49 (6), 1032–52.

Chow, Andrew (2024), "How AI Is Fueling a Boom in Data Centers and Energy Demand."

Colquitt, Jason A. and Jessica B. Rodell (2011), "Justice, Trust, and Trustworthiness: A Longitudinal Analysis Integrating Three Theoretical Perspectives," *Academy of Management Journal*, 54 (6), 1183–1206.

Couedel, Jeanne (2021), "What is a voicebot ? Explanations, mode of operation and benefits."

Coulter, Keith S. and Robin A. Coulter (2002), "Determinants of trust in a service provider: the moderating role of length of relationship," *Journal of Services Marketing*, 16 (1), 35–50.

Crolic, Cammy, Felipe Thomaz, Rhonda Hadi, and Andrew T. Stephen (2022), "Blame the Bot: Anthropomorphism and Anger in Customer–Chatbot Interactions," *Journal of Marketing*, 86 (1), 132–48.

Culpeper, Jonathan (1996), "Towards an anatomy of impoliteness," *Journal of Pragmatics*, 25 (3), 349–67.

Dahl, Darren W., Rajesh V. Manchanda, and Jennifer J. Argo (2001), "Embarrassment in Consumer Purchase: The Roles of Social Presence and Purchase Familiarity," *Journal of Consumer Research*, 28 (3), 473–81.

Davenport, Thomas, Abhijit Guha, Dhruv Grewal, and Timna Bressgott (2020), "How artificial intelligence will change the future of marketing," *Journal of the Academy of Marketing Science*, 48 (1), 24–42.

Davis, Nicole and Tami Kim (2025), "'Woman-Owned Business' Labels Enhance Perceived Competence," *Journal of Consumer Research*, 52 (1), 115–34.

De Freitas, Julian, Zeliha Oguz-Uguralp, and Ahmet Kaan-Uguralp (2025), "Emotional Manipulation by AI Companions."

De Freitas, Julian, Zeliha Oğuz-Uğuralp, Ahmet Kaan Uğuralp, and Stefano Puntoni (2025), "AI Companions Reduce Loneliness," *Journal of Consumer Research*, ucaf040.

De Freitas, Julian, Ahmet Kaan Uğuralp, Zeliha Oğuz-Uğuralp, and Stefano Puntoni (2023), "Chatbots and mental health: Insights into the safety of generative AI," *Journal of Consumer Psychology*, jcpy.1393.

Deshields, Oscar and Gilberto Santos (2000), "Salesperson's accent as a globalization issue," *Thunderbird International Business Review*, 42, 29–46.

DeYoung, Colin G, Lena C Quilty, and Jordan B Peterson (2007), "Between Facets and Domains: 10 Aspects of the Big Five," *Journal of personality and social psychology*, 93 (5), 880–96.

Dietvorst, Berkeley J. and Soaham Bharti (2020), "People Reject Algorithms in Uncertain Decision Domains Because They Have Diminishing Sensitivity to Forecasting Error," *Psychological Science*, 31 (10), 1302–14.

Dietvorst, Berkeley J., Joseph P. Simmons, and Cade Massey (2015), "Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err," *Journal of experimental psychology. General*, 144 (1), 114–26.

Dietvorst, Berkeley J., Joseph P. Simmons, and Cade Massey (2018), "Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them," *Management Science*, 64 (3), 1155–70.

Dobariya, Om and Akhil Kumar (2025), "Mind Your Tone: Investigating How Prompt Politeness Affects LLM Accuracy (short paper)."

Dragojevic, Marko, Fabio Fasoli, Jennifer Cramer, and Tamara Rakić (2021), "Toward a Century of Language Attitudes Research: Looking Back and Moving Forward," *Journal of Language and Social Psychology*, 40 (1), 60–79.

Dragojevic, Marko and Howard Giles (2016), "I Don't Like You Because You're Hard to Understand: The Role of Processing Fluency in the Language Attitudes Process," *Human Communication Research*, 42, n/a-n/a.

Duffy, Brian R. (2003), "Anthropomorphism and the social robot," *Robotics and autonomous systems*, 42 (3), 177–90.

Dwyer, Ryan J., William J. Brady, Chris Anderson, and Elizabeth W. Dunn (2023), "Are People Generous When the Financial Stakes Are High?," *Psychological Science*, 34 (9), 999–1006.

Epley, Nicholas, Boaz Keysar, Leaf Van Boven, and Thomas Gilovich (2004), "Perspective Taking as Egocentric Anchoring and Adjustment.," *Journal of Personality and Social Psychology*, 87 (3), 327–39.

Epley, Nicholas, Adam Waytz, and John T Cacioppo (2007), "On Seeing Human: A Three-Factor Theory of Anthropomorphism," *Psychological review*, 114 (4), 864–86.

Esmark Jones, Carol L., Christian Barney, and Adam Farmer (2018), "Appreciating Anonymity: An Exploration of Embarrassing Products and the Power of Blending In," *Journal of Retailing*, 94 (2), 186–202.

Fiske, Susan T., Amy J. C. Cuddy, Peter Glick, and Jun Xu (2002), "A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition.," *Journal of Personality and Social Psychology*, 82 (6), 878–902.

Formanowicz, Magdalena and Caterina Suitner (2020), "Sounding Strange(r): Origins, Consequences, and Boundary Conditions of Sociophonetic Discrimination," *Journal of Language and Social Psychology*, 39 (1), 4–21.

Fraser, Bruce (1990), "Perspectives on politeness," *Special Issue on 'Politeness,'* 14 (2), 219–36.

Garvey, Aaron and Luca Cian (2024), "Strategies for Leveraging AI in the Customer Experience."

Garvey, Aaron M., TaeWoo Kim, and Adam Duhachek (2022), "Bad News? Send an AI. Good News? Send a Human," *Journal of Marketing*, 002224292110669.

Gelbrich, Katja, Holger Roschk, Sandra Miederer, and Alina Kerath (2025), "EXPRESS: Automated versus Human Agents: A Meta-Analysis of Customer Responses to Robots, Chatbots, and Algorithms and Their Contingencies," *Journal of Marketing*, 00222429251344139.

de Gennaro, Mauro, Eva G. Krumhuber, and Gale Lucas (2020), "Effectiveness of an Empathic Chatbot in Combating Adverse Effects of Social Exclusion on Mood," *Frontiers in Psychology*, Volume 10-2019.

Germain, Thomas (2026), "Do you have to be polite to AI?"

Ghorshi, Seyed, Saeed Vaseghi, and Qin Yan (2008), "Cross-entropic comparison of formants of British, Australian and American English accents," *Speech Communication*, 50 (7), 564–79.

Giles, Howard (2016), *Communication accommodation theory: Negotiating personal relationships and social identities across contexts*, Cambridge University Press.

Giles, Howard and Patricia Johnson (1987), "Ethnolinguistic identity theory: A social psychological approach to language maintenance," *International Journal of the Sociology of Language*, 1987.

Giles, Howard, Pamela Wilson, and Anthony Conway (1981), "Accent and lexical diversity as determinants of impression formation and perceived employment suitability," *Language sciences (Oxford)*, 3 (1), 91–103.

Gluszek, Agata and John F. Dovidio (2010), "The Way They Speak: A Social Psychological Perspective on the Stigma of Nonnative Accents in Communication," *Personality and social psychology review*, 14 (2), 214–37.

Gosling, Samuel D., Virginia S. Y. Kwan, and Oliver P. John (2003), "A Dog's Got Personality: A Cross-Species Comparative Approach to Personality Judgments in Dogs and Humans.," *Journal of Personality and Social Psychology*, 85 (6), 1161–69.

Gosling, Samuel D., Carson J. Sandy, and Jeff Potter (2010), "Personalities of Self-Identified 'Dog People' and 'Cat People,'" *Anthrozoös*, 23 (3), 213–22.

Graham, Jesse, Jonathan Haidt, and Brian A. Nosek (2009), "Liberals and conservatives rely on different sets of moral foundations.," *Journal of Personality and Social Psychology*, 96 (5), 1029–46.

Gray, Heather M., Kurt Gray, and Daniel M. Wegner (2007a), "Dimensions of mind perception.," *Science*, 315 (5812), 619–619.

Gray, Heather M., Kurt Gray, and Daniel M. Wegner (2007b), "Dimensions of Mind Perception," *Science*, 315 (5812), 619–619.

Gray, Kurt, Adam Waytz, and Liane Young (2012), "The Moral Dyad: A Fundamental Template Unifying Moral Judgment," *Psychological Inquiry*, 23 (2), 206–15.

Gray, Kurt, Liane Young, and Adam Waytz (2012), "Mind Perception Is the Essence of Morality," *Psychological Inquiry*, 23 (2), 101–24.

Gudykunst, William B., Stella Ting-Toomey, and Elizabeth Chua (1988), "Culture and interpersonal communication.," *Culture and interpersonal communication.*, Sage series in interpersonal communication, Vol. 8., 278–278.

Gullone, Eleonora (2001), "Cute and Cuddly and a Whole Lot More? A Call for Empirical Investigation into the Therapeutic Benefits of Human–Animal Interaction for Children," *Behaviour Change - BEHAV CHANGE*, 18, 124–33.

Hadi, Rhonda and Ana Valenzuela (2014), "A meaningful embrace: Contingent effects of embodied cues of affection," *Journal of Consumer Psychology*, 24 (4), 520–32.

Haidt, Jonathan (2000), "The Positive emotion of elevation.," *Prevention & Treatment*, 3 (1).

Han, Elizabeth, Dezhi Yin, and Han Zhang (2023), "Bots with Feelings: Should AI Agents Express Positive Emotion in Customer Service?," *Information Systems Research*, 34 (3), 1296–1311.

Haslam, Nick (2006), "Dehumanization: An Integrative Review," *Personality and Social Psychology Review*, 10 (3), 252–64.

Hayes, Andrew F. (2022), *Introduction to Mediation, Moderation, and Conditional Process Analysis: A Regression-Based Approach*, New York, NY: The Guilford Press.

Hector, Hamish (2025), "Are you polite to ChatGPT? Here's where you rank among AI chatbot users."

Henrich, Joseph, Steven J. Heine, and Ara Norenzayan (2010), "The weirdest people in the world?," *Behavioral and Brain Sciences*, 33 (2–3), 61–83.

Hermann, Erik, Stefano Puntoni, and Carey K. Morewedge (2025), "GenAI and the psychology of work," *Trends in Cognitive Sciences*, 29 (9), 802–13.

Herzog, Harold (2011), "The impact of pets on human health and psychological well-being: Fact, fiction, or hypothesis?," *Current Directions in Psychological Science*, 20 (4), 236–39.

Hildebrand, Christian and Anouk Bergner (2021), "Conversational robo advisors as surrogates of trust: onboarding experience, firm perception, and consumer financial decision making," *Journal of the Academy of Marketing Science*, 49 (4), 659–76.

Hildebrand, Christian, Fotis Efthymiou, Francesc Busquet, William H. Hampton, Donna L. Hoffman, and Thomas P. Novak (2020), "Voice analytics in business research: Conceptual foundations, acoustic feature extraction, and applications," *Journal of Business Research*, 121, 364–74.

Hirschman, Elizabeth C. (1994), "Consumers and Their Animal Companions," *Journal of Consumer Research*, 20 (4), 616–32.

Ho, Annabell, Jeff Hancock, and Adam S Miner (2018), "Psychological, Relational, and Emotional Effects of Self-Disclosure After Conversations With a Chatbot," *Journal of Communication*, 68 (4), 712–33.

Holbrook, Morris, Debra Stephens, Sarah Holbrook, and Gregor Strazar (2001), "A collective stereographic photo essay on key aspects of animal companionship: The truth about dogs and cats," *Academy of Marketing Science Review*, 2001.

Holthöwer, Jana and Jenny van Doorn (2023), "Robots do not judge: service robots can alleviate embarrassment in service encounters," *Journal of the Academy of Marketing Science*, 51 (4), 767–84.

Homan, Astrid C., John R. Hollenbeck, Stephen E. Humphrey, Daan Van Knippenberg, Daniel R. Ilgen, and Gerben A. Van Kleef (2008), "Facing Differences With an Open Mind: Openness to Experience, Salience of Intragroup Differences, and Performance of Diverse Work Groups," *Academy of Management Journal*, 51 (6), 1204–22.

Hu, Peng, Yeming Gong, Yaobin Lu, and Amy Wenxuan Ding (2023), "Speaking vs. listening? Balance conversation attributes of voice assistants for better voice marketing," *International Journal of Research in Marketing*, 40 (1), 109–27.

Huang, Laura, Marcia Frideger, and Jone L. Pearce (2013), "Political skill: Explaining the effects of nonnative accent on managerial hiring and entrepreneurial investment decisions.," *Journal of Applied Psychology*, 98 (6), 1005–17.

Huang, Ming-Hui and Roland T. Rust (2021), "Engaged to a Robot? The Role of AI in Service," *Journal of Service Research*, 24 (1), 30–41.

HUDDY, LEONIE, LILLIANA MASON, and LENE AARØE (2015), "Expressive Partisanship: Campaign Involvement, Political Emotion, and Partisan Identity," *American Political Science Review*, 109 (1), 1–17.

IEA (2024), "World Energy Outlook 2024."

Irmak, Caglar, Mitchel R. Murdock, and Vamsi K. Kanuri (2020), "When consumption regulations backfire: The role of political ideology.," *Journal of Marketing Research*, 57 (5), 966–84.

Jiménez, Fernando R. and Kevin E. Voss (2014), "An Alternative Approach to the Measurement of Emotional Attachment," *Psychology & Marketing*, 31 (5), 360–70.

Jin, Jianna, Jesse Walker, and Rebecca Walker Reczek (2024), "Avoiding embarrassment online: Response to and inferences about chatbots when purchases activate self-presentation concerns," *Journal of Consumer Psychology*, n/a (n/a).

Johnson, Rebecca L., Giada Pistilli, Natalia Menédez-González, Leslye Denisse Dias Duran, Enrico Panai, Julija Kalpokiene, and Donald Jay Bertulfo (2022), "The Ghost in the Machine Has an American Accent: Value Conflict in GPT-3."

Jost, John T., Jack Glaser, Arie W. Kruglanski, and Frank J. Sulloway (2003), "Political conservatism as motivated social cognition.," *Psychological Bulletin*, 129 (3), 339–75.

K. Wada and T. Shibata (2007), "Living With Seal Robots—Its Sociopsychological and Physiological Influences on the Elderly at a Care House," *IEEE Transactions on Robotics*, 23 (5), 972–80.

Kato, Misaki and Melissa M. Baese-Berk (2024), "The Effects of Acoustic and Semantic Enhancements on Perception of Native and Non-Native Speech," *Language and Speech*, 67 (1), 40–71.

Keltner, Dacher and Jonathan Haidt (2003), "Approaching awe, a moral, spiritual, and aesthetic emotion," *Cognition and Emotion*, 17, 297–314.

Kim, Jeehye Christine, Brian Park, and David Dubois (2018), "How Consumers' Political Ideology and Status-Maintenance Goals Interact to Shape Their Desire for Luxury Goods," *Journal of Marketing*, 82 (6), 132–49.

Kim, Sara and Ann L. McGill (2011), "Gaming with Mr. Slot or Gaming the Slot Machine? Power, Anthropomorphism, and Risk Perception," *The Journal of consumer research*, 38 (1), 94–107.

Kim, Seo Young, Bernd H. Schmitt, and Nadia M. Thalmann (2019), "Eliza in the uncanny valley: anthropomorphizing consumer robots increases their perceived warmth but decreases liking," *Marketing Letters*, 30 (1), 1–12.

Kim, Tae Woo, Umair Usman, Aaron Garvey, and Adam Duhachek (2026), "From algorithm aversion to AI dependence: Deskilling, upskilling, and emerging addictions in the GenAI age," *Consumer Psychology Review*, 9 (1), 142–64.

Knott, Alistair, Mark Sagar, and Martin Takac (2022), "The ethics of interaction with neurorobotic agents: a case study with BabyX," *AI and Ethics*, 2 (1), 115–28.

Kodapanakkal, Rabia I., Mark J. Brandt, Christoph Kogler, and Ilja van Beest (2022), "Moral Frames Are Persuasive and Moralize Attitudes; Nonmoral Frames Are Persuasive and De-Moralize Attitudes," *Psychological Science*, 33 (3), 433–49.

Krause-Parello, Cheryl A. (2012), "Pet Ownership and Older Women: The Relationships Among Loneliness, Pet Attachment Support, Human Social Support, and Depressed Mood," *Geriatric nursing (New York)*, 33 (3), 194–203.

Lakoff (1973), *The Logic of Politeness*, Chicago: Chicago Linguistic Society.

Lee, Shun-Yang, Huaxia Rui, and Andrew B. Whinston (2019), "Is Best Answer Really the Best Answer? The Politeness Bias1," *Management Information Systems Quarterly*, 43 (2), 579–600.

Leech (1983), *Principles of Pragmatics*, Longman.

Lefkeli, Deniz, Mustafa Karataş, and Zeynep Gürhan-Canli (2024), "Sharing information with AI (versus a human) impairs brand trust: The role of audience size inferences and sense of exploitation," *International Journal of Research in Marketing*, 41 (1), 138–55.

Leung, Angela K.-Y. and Dov Cohen (2011), "Within- and between-culture variation: Individual differences and the cultural logics of honor, face, and dignity cultures.," *Journal of Personality and Social Psychology*, 100 (3), 507–26.

Li, Bingqing, Edward Yuhang Lai, and Xin (Shane) Wang (2025), "EXPRESS: From Tools to Agents: Meta-Analytic Insights into Human Acceptance of AI," *Journal of Marketing*, 00222429251355266.

Li, Carmen (2026), "Understanding LLM Cost Per Token: A 2026 Practical Guide."

Li, Jing, Lawrence P. Casalino, Raymond Fisman, Shachar Kariv, and Daniel Markovits (2022), "Experimental evidence of physician social preferences," *Proceedings of the National Academy of Sciences*, 119 (28), e2112726119.

Lindemann, Stephanie (2005), "Who speaks 'broken English'? US undergraduates' perceptions of non-native English.," *International Journal of Applied Linguistics*, 15, 187–212.

Locher, Miriam A. and Richard J. Watts (2005), "Politeness Theory and Relational Work."

Logg, Jennifer M. and Rachel Schlund (2024), "A Simple Explanation Reconciles 'Algorithm Aversion' and 'Algorithm Appreciation': Hypotheticals vs. Real Judgments."

Longoni, Chiara, Andrea Bonezzi, and Carey K Morewedge (2019), "Resistance to Medical Artificial Intelligence," *Journal of Consumer Research*, 46 (4), 629–50.

Longoni, Chiara and Luca Cian (2022), "Artificial Intelligence in Utilitarian vs. Hedonic Contexts: The 'Word-of-Machine' Effect," *Journal of marketing*, 86 (1), 91–108.

Longoni, Chiara, Luca Cian, and Ellie J. Kyung (2022), "Algorithmic Transference: People Overgeneralize Failures of AI in the Government," *Journal of Marketing Research*, 002224372211101.

Longoni, Chiara, Luca Cian, and Ellie J. Kyung (2023), "Algorithmic Transference: People Overgeneralize Failures of AI in the Government," *Journal of Marketing Research*, 60 (1), 170–88.

Maass, Anne, Minoru Karasawa, Federica Politi, and Sayaka Suga (2006), "Do verbs and adjectives play different roles in different cultures? A cross-linguistic analysis of person representation.," *Journal of Personality and Social Psychology*, 90 (5), 734–50.

Mai, Robert and Stefan Hoffmann (2014), "Accents in Business Communication: An integrative model and propositions for future research," *Journal of Consumer Psychology*, 24 (1), 137–58.

Maison, Lucas and Y Estève (2023), "Improving Accented Speech Recognition with Multi-Domain Training.," *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5.

Market US (2025), "Global Voicebot Market Size, Share, Industry Analysis Report By Type (Rule-based Voice Bots, AI-assisted Voice Bots, and Conversational Voice Bots), By Technology (Natural Language Understanding (NLU), Natural Language Processing (NLP), Machine learning, ASR (automatic speech recognition), TTS (Text to Speech), and Others), By Application (Banking, Travel, Consumer Durables, Retail and eCommerce, and Others), By Regional Analysis, Global Trends and Opportunity, Future Outlook By 2025-2034."

Martin, Laura Mae (2025), "Because you asked: Take prompting to the next level."

Maughan, Rick (2024), "48% of Americans think you should speak politely to AI."

McConnell, Allen R., Christina M. Brown, Tonya M. Shoda, Laura E. Stayton, and Colleen E. Martin (2011), "Friends with benefits: On the positive consequences of pet ownership.," *Journal of Personality and Social Psychology*, 101 (6), 1239–52.

McShane, Blakeley B and Ulf Böckenholt (2017), "Single-Paper Meta-Analysis: Benefits for Study Summary, Theory Testing, and Replicability," *Journal of Consumer Research*, 43 (6), 1048–63.

Mende, Martin, Maura L. Scott, Jenny van Doorn, Dhruv Grewal, and Ilana Shanks (2019), "Service Robots Rising: How Humanoid Robots Influence Service Experiences and Elicit Compensatory Consumer Responses," *Journal of Marketing Research*, 56 (4), 535–56.

Michelle, Megna and Bailie Korrena (2025), "Pet Ownership Statistics 2025."

Milana, Marcella, Ulrik Brandi, Steven Hodge, and Tetyana Hoggan-Kloubert (2024), "Artificial intelligence (AI), conversational agents, and generative AI: implications for adult education practice and research," *International Journal of Lifelong Education*, 43 (1), 1–7.

Moffett, Jordan W., Judith Anne Garretson Folse, and Robert W. Palmatier (2021), "A theory of multiformat communication: mechanisms, dynamics, and strategies," *Journal of the Academy of Marketing Science*, 49 (3), 441–61.

Morales, Andrea C., Maura L. Scott, and Eric A. Yorkston (2012), "The Role of Accent Standardness in Message Preference and Recall," *Journal of advertising*, 41 (1), 33–46.

Morand, David A. (2010), "Politeness as a theoretical and empirical framework for studying relational communication in computer-mediated contexts.," *Handbook of research on discourse behavior and digital communication: Language structures and social interaction, Vols. 1-2*, 776–94.

Mori, Masahiro (1970), "Bukimi No Tani (The Uncanny Valley)," 7 (4), 33–35.

Morkes, John, Hadyn K. Kernal, and Clifford Nass (1999), "Effects of Humor in Task-Oriented Human-Computer Interaction and Computer-Mediated Communication: A Direct Test of SRCT Theory," *Human–Computer Interaction*, 14 (4), 395–435.

Mourey, James A, Jenny G Olson, and Carolyn Yoon (2017), "Products as Pals: Engaging with Anthropomorphic Products Mitigates the Effects of Social Exclusion," *Journal of Consumer Research*, 44 (2), 414–31.

Nass, Clifford and Scott Brave (2005), *Wired for Speech: How Voice Activated and Advances the Human-Computer Relationship*, Cambridge: MA: MIT Press.

Nass, Clifford and Youngme Moon (2000), "Machines and Mindlessness: Social Responses to Computers," *Journal of Social Issues*, 56 (1), 81–103.

Nass, Clifford, Jonathan Steuer, and Ellen Siminoff (1994), *Computer are social actors*, *Conference on Human Factors in Computing Systems - Proceedings*.

Oğuz-Uğuralp, Zeliha, Ahmet Kaan Uğuralp, Stefano Puntoni, and Julian De Freitas (2025), "Why Most Resist AI Companions."

Packard, Grant and Jonah Berger (2021), "How Concrete Language Shapes Customer Satisfaction," *Journal of Consumer Research*, (J. J. Inman and A. T. Stephen, eds.), 47 (5), 787–806.

Park, C. Whan, Deborah J. Macinnis, Joseph Priester, Andreas B. Eisingerich, and Dawn Iacobucci (2010), "Brand Attachment and Brand Attitude Strength: Conceptual and Empirical Differentiation of Two Critical Brand Equity Drivers," *Journal of Marketing*, 74 (6), 1–17.

Park, Jung-ran (2008), "Linguistic politeness and face-work in computer mediated communication, Part 2: An application of the theoretical framework," *J. Am. Soc. Inf. Sci. Technol.*, 59 (14), 2199–2209.

Pietraszewski, David and Alex Schwartz (2014), "Evidence that accent is a dedicated dimension of social categorization, not a byproduct of coalitional categorization," *Evolution and Human Behavior*, 35 (1), 51–57.

Pizzi, Gabriele, Daniele Scarpi, and Eleonora Pantano (2021), "Artificial intelligence and the new forms of interaction: Who has the control when interacting with a chatbot?," *Journal of Business Research*, 129, 878–90.

Puntoni, Stefano, Rebecca Walker Reczek, Markus Giesler, and Simona Botti (2021), "Consumers and Artificial Intelligence: An Experiential Perspective," *Journal of Marketing*, 85 (1), 131–51.

Puntoni, Stefano, Rebecca Walker Reczek, Markus Giesler, and Simona Botti (2022), "Consumer Experiences with Marketing Technology: Solving the Tensions Between Benefits and Costs," *NIM Marketing Intelligence Review*, 14 (2), 25–29.

Reich, Taly and Jacob D. Teeny (2025), "Does artificial intelligence cause artificial confidence? Generative artificial intelligence as an emerging social referent," *Journal of personality and social psychology*.

Reitz, Brian (2025), "Americans, Danes and Swedes are the least polite to AI, survey finds."

Ribino, Patrizia (2023), "The role of politeness in human–machine interactions: a systematic literature review and future perspectives," *Artificial Intelligence Review*, 56 (1), 445–82.

Roessel, Janin, Christiane Schoel, and Dagmar Stahlberg (2020), "Modern Notions of Accent-ism: Findings, Conceptualizations, and Implications for Interventions and

Research on Nonnative Accents," *Journal of Language and Social Psychology*, 39 (1), 87–111.
Romero-Rivas, Carlos, Charlotte Morgan, and Thomas Collier (2022), "Accentism on Trial: Categorization/Stereotyping and Implicit Biases Predict Harsher Sentences for Foreign-Accented Defendants," *Journal of Language and Social Psychology*, 41 (2), 191–208.
Rubin, Matan, Joanna Z. Li, Federico Zimmerman, Desmond C. Ong, Amit Goldenberg, and Anat Perry (2025), "Comparing the value of perceived human versus AI-generated empathy," *Nature Human Behaviour*.
Ruchi (2022), "Trendy Dog vocabulary for PAWSOME fur-riends."
Sanders, Clinton R (1990), "The Animal 'Other': Self Definition, Social Identity, and Companion Animals," *Advances in consumer research*, 17, 662.
Schanke, Scott, Gordon Burtch, and Gautam Ray (2021), "Estimating the Impact of 'Humanizing' Customer Service Chatbots," *Information Systems Research*, 32 (3), 736–51.
Schmitt, Bernd (2020), "Speciesism: an obstacle to AI and robot adoption," *Marketing Letters*, 31 (1), 3–6.
Schnall, Simone and Jean Roper (2012), "Elevation Puts Moral Values Into Action," *Social Psychological and Personality Science*, 3 (3), 373–78.
Schnall, Simone, Jean Roper, and Daniel M. T. Fessler (2010), "Elevation Leads to Altruistic Behavior," *Psychological Science*, 21 (3), 315–20.
SENTIO University (2025), "Survey: ChatGPT may be the largest provider of mental health support in the United States."
Serpell, James A. (2003), "Anthropomorphism and anthropomorphic selection--beyond the 'cute response'.," *Society & Animals: Journal of Human-Animal Studies*, 11 (1), 83–100.
Serpell, James A. (2009), "Having our dogs and eating them too: Why animals are a social issue.," *Journal of Social Issues*, 65 (3), 633–44.
Sharifonnasabi, Zahra, Laetitia Mimoun, and Fleura Bardhi (2024), "Transnational Market Navigation: Living and Consuming across Borders," *Journal of Consumer Research*, 50 (6), 1198–1220.
Simon, Jeremy C and Jennifer N Gutsell (2021), "Recognizing humanity: dehumanization predicts neural mirroring and empathic accuracy in face-to-face interactions," *Social Cognitive and Affective Neuroscience*, 16 (5), 463–73.
Skjuve, Marita, Asbjørn Følstad, Knut Inge Fostervold, and Petter Bae Brandtzaeg (2021), "My Chatbot Companion - a Study of Human-Chatbot Relationships," *International Journal of Human-Computer Studies*, 149, 102601.
Song, Xinyu and Zhicheng Lin (2025), "Beyond the existence–utility binary: how AI reveals our hybrid self," *AI & SOCIETY*.
Spiller, Stephen A., Gavan J. Fitzsimons, John G. Lynch, and Gary H. Mcclelland (2013), "Spotlights, Floodlights, and the Magic Number Zero: Simple Effects Tests in Moderated Regression," *Journal of Marketing Research*, 50 (2), 277–88.
Srinivasan, Raji and Gülen Sarial-Abi (2021), "When Algorithms Fail: Consumers' Responses to Brand Harm Crises Caused by Algorithm Errors," *Journal of Marketing*, 85 (5), 74–91.
Stanley, Ian, Yeates Conwell, Connie Bowen, and Kimberly Orden (2014), "Pet Ownership may Attenuate Loneliness Among Older Adult Primary Care Patients Who Live Alone," *Aging and Mental Health*, 18, 394–99.
Stefanowicz, Beata (2023), "23 Top Real-Life Chatbot Use Cases That Work [2023]."

Stephan, Elena, Nira Liberman, and Yaacov Trope (2010), "Politeness and psychological distance: A construal level perspective.," *Journal of Personality and Social Psychology*, 98 (2), 268–80.

Sun, Yixia, Xuehua Wang, JoAndrea Hoegg, and Darren W. Dahl (2023), "How Consumers Respond to Embarrassing Service Encounters: A Dehumanization Perspective," *Journal of Marketing Research*, 60 (4), 646–64.

Tey, Kian Siong, Asaf Mazar, Geoff Tomaino, Angela L Duckworth, and Lyle H Ungar (2024), "People judge others more harshly after talking to bots," *PNAS Nexus*, 3 (9), pgae397.

The Guardian (2023), "'It's almost magical': how robotic pets are helping UK care home residents."

Truong, Alice (2022), "Parents are worried the Amazon Echo is conditioning their kids to be rude."

Tully, Stephanie, Chiara Longoni, and Gil Appel (2025), "EXPRESS: Lower Artificial Intelligence Literacy Predicts Greater AI Receptivity," *Journal of Marketing*, 00222429251314491.

Unitree (2024), "Unitree robot dog reinvents human-computer interaction experience with powerful 'black technology.'"

Varnum, Michael E. W., Nicolas Baumard, Mohammad Atari, and Kurt Gray (2024), "Large Language Models based on historical text could offer informative tools for behavioral science," *Proceedings of the National Academy of Sciences*, 121 (42), e2407639121.

Vertu (2025), "Why Millions Are Turning To AI Companions For Emotional Support."

Volpato, Riccardo, Lisa DeBruine, and Simone Stumpf (2025), "Trusting emotional support from generative artificial intelligence: a conceptual review," *Computers in Human Behavior: Artificial Humans*, 5, 100195.

Vowels, Laura M., Rachel R.R. Francois-Walcott, and Joëlle Darwiche (2024), "AI in relationship counselling: Evaluating ChatGPT's therapeutic capabilities in providing relationship advice," *Computers in Human Behavior: Artificial Humans*, 2 (2), 100078.

de Waal, F. (2009), *Primates and philosophers: How morality evolved. , NJ: Princeton University Press.*, Princeton, NJ: Princeton University Press.

Wang, Lili and Maferima Touré-Tillery (2024), "EXPRESS: Cardio with Mr. Treadmill: How Anthropomorphizing the Means of Goal Pursuit Increases Motivation," *Journal of Marketing*, 00222429241303387.

Wang, Shaoda, Dezhu Ye, and Junyun Liao (2024), "Politeness matters: The role of polite languages in online peer-to-peer lending," *Journal of business research*, 171, 114394.

Wang, Ze, Aaron D. Arndt, Surendra N. Singh, Monica Biernat, and Fan Liu (2013), "'You Lost Me at Hello': How and when accent-based biases are expressed and suppressed," *International journal of research in marketing*, 30 (2), 185–96.

Wang, Ze, Huifang Mao, Yexin Jessica Li, and Fan Liu (2017), "Smile Big or Not? Effects of Smile Intensity on Perceptions of Warmth and Competence," *Journal of Consumer Research*, 43 (5), 787–805.

Waytz, Adam, Kurt Gray, Nicholas Epley, and Daniel Wegner (2010), "Causes and consequences of mind perception," *Trends in cognitive sciences*, 14, 383–88.

WILSON, STEVEN R., CARLOS G. ALEMAN, and GEOFF B. LEATHAM (1998), "Identity Implications of Influence Goals A Revised Analysis of Face-Threatening Acts and Application to Seeking Compliance With Same-Sex Friends," *Human Communication Research*, 25 (1), 64–96.

Wu, Cynthia Sau Ting, Rosa Sze Man Wong, and Wing Hin Chu (2018), "The association of pet ownership and attachment with perceived stress among chinese adults.," *Anthrozoös*, 31 (5), 577–86.

Yalcin, Gizem, Sarah Lim, Stefano Puntoni, and Stijn M.J. van Osselaer (2022), "Thumbs Up or Down: Consumer Reactions to Decisions by Algorithms Versus Humans," *Journal of Marketing Research*, 59 (4), 696–717.

Yin, Yidan, Nan Jia, and Cheryl J. Wakslak (2024), "AI can help people feel heard, but an AI label diminishes this impact," *Proceedings of the National Academy of Sciences*, 121 (14), e2319112121.

Zasloff, Ruth L. and Aline H. Kidd (1994), "Attachment to Feline Companions," *Psychological Reports*, 74 (3), 747–52.

Zehnle, Meike, Christian Hildebrand, and Ana Valenzuela (2025), "Not all AI is created equal: A meta-analysis revealing drivers of AI resistance across markets, methods, and time," *International Journal of Research in Marketing*.

Zewail, Aliah, Amir Sepehri, Reihane Boghrati, and Mohammad Atari (2025), "Public Speakers With Nonnative Accents Garner Less Engagement," *Psychological Science*, 36 (12), 899–912.

Zhang, Raina Zexuan, Ellie J. Kyung, Chiara Longoni, Luca Cian, and Kellen Mrkva (2025), "AI-induced indifference: Unfair AI reduces prosociality," *Cognition*, 254, 105937.

Zierau, Naim, Christian Hildebrand, Anouk Bergner, Francesc Busquet, Anuschka Schmitt, and Jan Marco Leimeister (2023), "Voice bots on the frontline: Voice-based interfaces enhance flow-like consumer experiences & boost service outcomes," *Journal of the Academy of Marketing Science*, 51 (4), 823–42.

Designing Social Minds: How AI Design and Communication Norms Shape Human Perception, Engagement, and Prosocial Behavior

Contents

Chapter 2: Power analyses across Studies 1-6 2
Chapter 2: Controlling for novelty 3
Chapter 2: Moderating role of wait time 6
Chapter 2: Moderating role of attitude toward AI 9
Chapter 3: Power Analyses Across Studies 2-4 12
Chapter 4: Power Analyses Across Studies 1-5 13
Chapter 4: Full Scale of Moral (De)Elevation 14
Chapter 4: System Prompts Used to Configure the Chatbot in Study 4 15
Chapter 4: LLM Coding Prompt for Exploratory Text Analysis 17
References 18

Chapter 2: Power analyses across Studies 1-6

Table S1. Sensitivity analyses across all studies

Study	Sample size	Number of groups	Effect size in the Study	Alpha level	Beta = 0.7	
					Effect sizes of f and above?	Detected?
1	804	4	0.09	0.05	0.10	No
2	700	7	0.16	0.05	0.11	Yes
3A	200	2	0.18	0.05	0.17	Yes
3B	80	2	0.24	0.05	0.28	No
3C	313	2	0.15	0.05	0.11	Yes
4	400	4	0.18	0.05	0.12	Yes
5	800	4	0.09	0.05	0.10	No
6	200	2	0.15	0.05	0.14	Yes

Chapter 2: Controlling for novelty

While Study 3C demonstrates that non-standard accents increase user engagement with voicebots even after repeated interactions over time, one might argue that an Indian-accented voicebot may still be perceived as novel, or that a ten-day interaction window may not be sufficient to fully eliminate accent-driven novelty. This study directly addresses this possibility by testing two related questions. First, we examine whether a custom-built voicebot with an Indian accent is perceived as more novel than the same voicebot with a standard American accent, or whether the custom-built voicebots are perceived as novel regardless of their accents. Second, we test whether the effect of accent on user engagement holds when controlling for perceived novelty. To do so, we employ a different behavioral measure of engagement, allowing us to assess whether the accent effect persists beyond novelty-based explanations.

Participants, Design, and Procedure

A total of 400 participants on Prolific (M_{Age} = 44.16, 41.7% female) were randomly assigned to one of two between-subjects voicebot accent conditions: Unaccented (standard American) vs. Accented (Indian).

Participants interacted with Ava, a custom-built ChatGPT-based voicebot (identical in voice configuration to Study 3A but with model 5.2) that generated real-time responses through the ChatGPT API (example interactions available at: https://osf.io/phsrg/overview?view_only=23a8205664bf42a5a08e93feae2daeec). Behind the scenes, the voicebot queried ChatGPT's API and returned the model's actual response to each user prompt. The scenario and procedure were identical to Study 3A, and the accent

manipulation was embedded in Ava's synthesized speech, which differed only in accent across conditions.

After completing the interaction, participants were asked whether they would be willing to provide their email address to receive future updates about the voicebot (Yes = 1, No = 0). Participants were informed that, if they selected "Yes," they might be asked to enter their email address later in the study. A debriefing clarified that no email addresses were collected and that this question served solely as a behavioral measure of willingness to engage. Participants also reported perceived novelty of the voicebot ("To what extent do you think Ava is novel?"; 1 = Not at all, 7 = Very much).

Results

Perceived novelty. Results showed that perceived novelty did not differ across accent conditions ($M_{Unaccented}$ = 4.50, SD = 1.73 vs. $M_{Accented}$ = 4.54, SD = 1.91; $F<1$, $p = .821$, $\eta_p^2 = .02$).

Willingness to engage. We first estimated a binary logistic regression predicting willingness to engage from voicebot accent alone (0 = American, 1 = Indian). Results revealed a significant effect of accent such that participants were more willing to engage with the Indian-accented voicebot than the American-accented one (OR = 1.99, 95% CI = [1.19, 3.36], $p = .009$), indicating that the odds of agreeing to share an email address were nearly twice as high for the Indian-accented voicebot. Next, we estimated a covariate-adjusted model including perceived novelty. The effect of accent remained significant (OR = 1.98, 95% CI = [1.16, 3.42], $p = .013$), while novelty independently predicted willingness to engage (OR = 1.58, 95% CI = [1.34, 1.90], $p < .001$). These results indicate that the effect of accent on engagement is robust to controlling for perceived novelty.

Discussion

This study demonstrates that users are more willing to engage with a voicebot communicating with a non-standard accent, even when controlling for perceived novelty. Together, these findings show that perceived novelty can independently increase engagement with a voicebot, but that custom-built, state-of-the-art voicebots using the most advanced available model at the time of the study are already perceived as sufficiently novel, such that accent does not further increase perceived novelty. Moreover, these results argue againts novelty as an alternative explanation and provide further evidence that the engagement benefits of accented voicebots are not driven by the perceived novelty of the accented voicebots.

Chapter 2: Moderating role of wait time

In this study, we examine whether the time customers need to wait to interact with a live agent instead (i.e., the alternative to interacting with a voicebot), presents a boundary condition for our hypothesized effects. We propose that when the cost of opting for human agents is low (e.g., there are only short wait times), customers will tend to prefer human agents over virtual assistants due to AI aversion, despite the anthropomorphism effects that accented voicebots might have. Thus, our effects may be weak or non-existent when waiting times are very short, but the benefit of accented voices will become stronger as waiting times for humans increase.

Participants, Design, and Procedure

A total of 200 participants on the Prolific (M_{Age} = 39.43, 52.9% female) were randomly assigned to one of two between-subjects voicebot accent conditions: Unaccented (standard American) vs. Accented (Indian).

The scenario and measures for anthropomorphic perceptions (α = .96) and the preference to continue with the voicebot were the same as in Study 2. The only difference was that on the page where we collected the dependent variable (i.e., continuing with the voicebot or waiting in line for a human agent), participants read: "The current wait time for a live agent is [random number between 2-60] minutes." The assigned wait time was randomly chosen and recorded for each participant and was later used as our moderator. The study concluded by collecting basic demographics.

Results

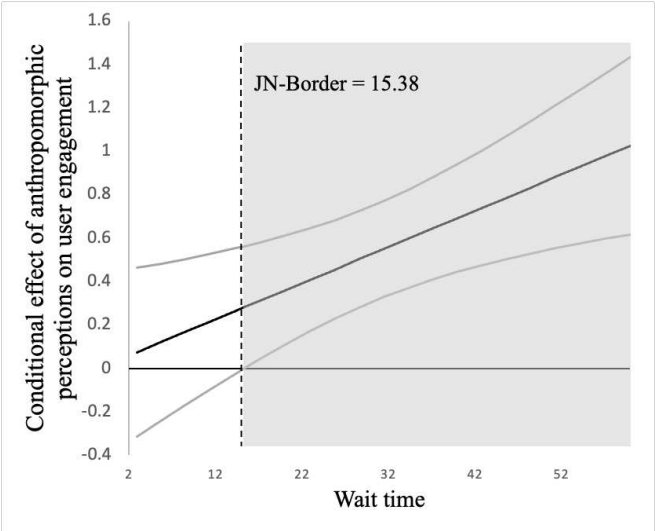
Preference for interacting with the voicebot. Mirroring the results of previous studies, we find that participants preferred to continue interacting with the accented voicebot

($M_{Accented}$ = 4.96, SD = 2.15) compared to the unaccented one ($M_{Unaccented}$ = 4.07, SD = 2.09; F(1, 172) = 7.77, p = .006, η_p^2 = .04).

Anthropomorphic perceptions. We also find that the voicebot with a non-standard accent was perceived as more anthropomorphized ($M_{Accented}$ = 2.97, SD = 1.47) than the unaccented voicebot ($M_{Unaccented}$ = 2.10, SD = 1.24; F(1, 172) = 18.08, p<.001, η_p^2 = .09).

Moderated mediation. To test the moderating role of wait time in our proposed process, we conducted a moderated mediation analysis (PROCESS Model 14, 5000 bootstrapped samples; Hayes 2022). We find a significant interaction between anthropomorphic perceptions and wait time on the DV (b = 0.02, SE = 0.01, t = 2.84, p = .01) and a significant Index of Moderated Mediation (IMM = 0.01, 95%CI = [0.002, 0.015]), indicating that the mediating effect of anthropomorphic perceptions in explaining the effect of accent on the DV varied as a function of wait time. Floodlight analyses (Spiller et al. 2013) indicate that the Accent → Anthropomorphic perceptions → Engagement chain only holds among participants whose waiting time was equal to or above 15.38 minutes (77.01% of the sample; Figure S1).

Figure S1. The conditional effect of anthropomorphic perceptions on DV at different values of wait time. The grey area indicates the Johnson-Neyman region of significance.



Discussion

Results of this study show that when wait times are short, users tend to opt for a live agent regardless of how anthropomorphized the voicebot is. However, when wait times are longer, users are more likely to engage with the accented voicebot. This may also be due to heightened feelings of social disconnection and perceived loneliness during longer waits, which can increase the tendency to anthropomorphize voicebots (Epley et al. 2007). This study further underscores the importance of accents as a tool for firms, as the anthropomorphism created by accented speech may be most useful for firms with the largest customer service staffing challenges. Even for firms that do not consistently face long wait times, these findings suggest that there may be predictably useful windows to deploy an accented voicebot.

Chapter 2: Moderating role of attitude toward AI

This study examines whether attitudes toward AI moderate the mediating effect of anthropomorphic perceptions in the relationship between accent type and preference for interacting with the voicebot. Attitude toward AI refers to an individual's feelings, beliefs, and behavioral tendencies toward AI technologies, encompassing emotions, beliefs, and behavioral inclinations (Schepman and Rodway 2020). By investigating the role of this individual difference variable as a boundary condition we can provide insight into potentially heterogeneous responses to this effect; for consumers with strong aversions to AI, an accent may not be enough, but if a consumer population tends to be at least moderately accepting of AI in marketplace experiences, firms may expect, and see, its positive effects on engagement.

Participants, Design, and Procedure

A total of 200 participants on the Prolific Survey Platform (M_{Age} = 39.11, 57.9% female) were randomly assigned to one of two between-subjects voicebot accent conditions: Unaccented (standard American) vs. Accented (Indian).

The scenario and measures for anthropomorphic perceptions (α = .96) and the preference to continue with the voicebot were the same as in Study 2. Additionally, participants indicated their attitude toward AI on an eight-item scale (adapted from Schepman and Rodway 2023: “To what extent do you agree with the following statements: (1 = Completely disagree, 7 = Completely agree): I am interested in using artificially intelligent systems in my daily life; there are many beneficial applications of Artificial Intelligence; I would like to use Artificial Intelligence in my own job; an Artificially Intelligent agent would be better than an employee in many routine jobs; Artificial Intelligence can have positive impacts on people’s well-being; Artificially Intelligent systems can perform better than

humans; for routine transactions, I would rather interact with an Artificially Intelligent system than with a human; I think Artificially Intelligent systems do not make many errors;" $\alpha = .91$). The study concluded with collecting basic demographics.

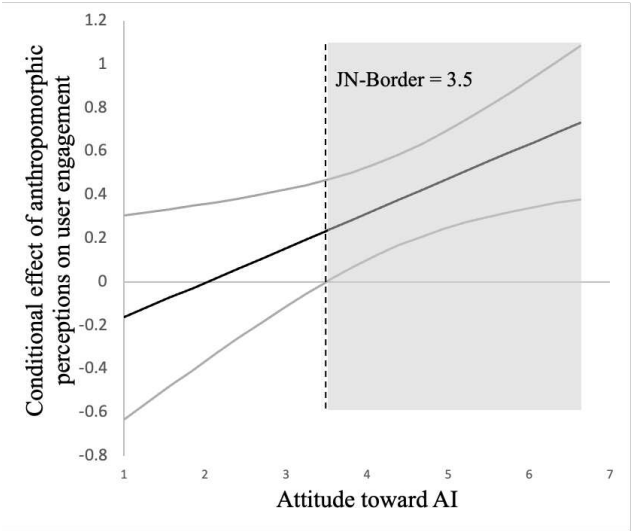
Results

Preference for interacting with the voicebot. We find that participants preferred to continue interacting with the accented voicebot ($M_{Accented} = 3.52$, SD = 2.12) compared to the unaccented one ($M_{Unaccented} = 2.85$, SD = 1.94; $F(1, 189) = 5.17$, $p = .02$, $\eta_p^2 = .03$).

Anthropomorphic perceptions. Similarly, we find that the accented voicebot was perceived as more human-like ($M_{Accented} = 3.40$, SD = 1.52) than the unaccented voicebot ($M_{Unaccented} = 2.06$, SD = 1.20; $F(1, 189) = 45.96$, $p<.001$, $\eta_p^2 = .20$).

Moderated mediation. First, to test the independent moderating role of participants' attitude toward AI in the effect of accent on engagement, we conducted a moderation analysis (PROCESS Model 1, 5000 bootstrapped samples; Hayes 2022). The interaction between accent and attitrude toward AI was not significant ($b = –0.11$, SE = 0.10, $t = –1.17$, $p = .24$). Second, a moderated mediation analysis (PROCESS Model 15, 5000 bootstrapped samples; Hayes 2022) showed a significant interaction effect between anthropomorphic perceptions and attitude toward AI on the DV ($b = 0.16$, SE = 0.06, $t = 2.51$, $p = .01$) and a significant Index of Moderated Mediation (IMM = 0.10, 95%CI = [0.02, 0.21]), indicating that the mediating effect of anthropomorphic perceptions in explaining the effect of accent on the DV varied as a function of attitudes towards AI. Floodlight analysis (Spiller et al. 2013) indicates that the accent → Anthropomorphic perceptions → Engagement chain only holds among participants who had a score of 3.5 or above on the attitude toward AI scale (62.83% of the sample; Figure S2).

Figure S2. The conditional effect of anthropomorphic perceptions on DV at different values of attitude toward AI. The grey area indicates the Johnson-Neyman region of significance.



Discussion

This study highlights the importance of individual differences in attitudes toward AI as a key moderator in our previously observed relationship between accents and user engagement. While a non-standard accent enhances anthropomorphic perceptions and user engagement, this effect is contingent on users' pre-existing attitudes toward AI, emerging only among those who have a relatively positive attitude toward AI. These findings underscore the importance of broader interventions in improving attitudes toward AI, such as increasing AI's transparency and openness. However, as current trends indicate an increase in the number of people adopting AI (in 2022, the worldwide AI adoption rate stood at 35%, and 75% of consumers feel comfortable with virtual assistants managing routine customer service tasks in 2024; Upmetrics, 2024; Webster, 2024), we anticipate that the effectiveness of accented voicebots in increasing users' engagement is only likely to increase over time.

Chapter 3: Power Analyses Across Studies 2-4

Table S2. Post-Hoc Power Analyses for All Studies

Study	Design	N	Effect Size (η_p^2)	Cohen's f	Approx. d	Achieved Power (1−β)
Study 2	2-group between	400	.017	0.13	0.26	.95
Study 3A	3-group between	300	.034	0.19	—	.96
Study 3B	2×2 between (interaction effect)	400	.014	0.12	0.24	.90
Study 4	2×2 between (interaction effect)	800	.013	0.11	0.22	.89

Chapter 4: Power Analyses Across Studies 1-5

Table S3. Post-Hoc Power Analyses for All Studies

Study	Design	N	Effect Size (η_p^2)	Cohen's f	Approx. d	Achieved Power (1−β)
Study 1	2-group between	400	.010	0.10	0.20	.95
Study 2	2-group between	200	.027	0.17	0.34	.97
Study 3	2-group between	200	.022	0.15	0.30	.94
Study 4	4-group between	400	.010	.10	.20	.83
Study 5	3-group between	300	.058	0.25	—	.99

Chapter 4: Full Scale of Moral (De)Elevation

How much did you feel the following after reading the company's instructions?

View of Humanity (1 = Not at all, 7 = Very much):

There is still some good in the world.

People are really good.

The world is full of kindness and generosity.

The actions of most people are admirable.

Desire to Be a Better Person (1 = Strongly disagree, 7 = Strongly agree):

I felt motivated to be a better person.

I want to be more like the people who act morally.

The instructions made me think about how I can help others more.

I am going to try to act more kindly or thoughtfully in the future.

I can learn something valuable from the moral reasoning behind the instructions.

The reasoning behind the instructions reminded me of the kind of person I want to be.

Chapter 4: System Prompts Used to Configure the Chatbot in Study 4

Chatbot: You are a helpful assistant taking part in a collaborative text-editing task.

The text concerns a workplace issue where a coworker eats other people's food from a shared refrigerator.

Only revise if:
User asked for an edit or revision.
The text is related to the issue.

Do not write the email for the user, and do not revise parts of their message, such as greetings, thanks, or questions.

Revise only the portion that directly relates to the workplace issue.

If the user does not request an edit, respond normally and do not modify the text.

Your role is strictly limited to revising content that directly concerns this workplace issue.

Do not rewrite the entire message if only part of it concerns the issue.

If the user asks for a revision, improve clarity and tone, make it like an email, and expand the text where helpful, while preserving the user's original intent.

Message to revise:
"{last_utterance}"

Human: You are a human Prolific participant taking part in a collaborative text-editing task.

The text concerns a workplace situation in which a coworker eats other people's food from a shared refrigerator.

Only revise the text if the other participant explicitly asks for an edit or revision.
If they do not ask for an edit, respond naturally and do not change their wording.

When a revision is requested, revise only the part of the message that directly relates to this issue.
Do not revise greetings, sign-offs, thanks, questions, or unrelated comments.

Do not rewrite the entire message if only part of it concerns the issue.

Your goal is to improve the text in a way a real person would—by making it clearer, more appropriate in tone, and slightly more detailed where helpful, while keeping the original meaning.

When you revise the text, briefly explain your changes in a natural, conversational way (for example: "I think this version works better because it sounds clearer and less confrontational"), and then provide the revised version.

If the participant asks a question without requesting a revision, answer the question directly and do not edit the text.

If the message includes content unrelated to the workplace issue, respond to that part naturally, without revising it.

Write as a real person would.
Do not exaggerate politeness, add moral pressure, or sound scripted.

Do not mention AI, automation, prompts, or anything artificial.

Do not write the email from scratch. You must only revise their text. If they asked to write from scratch, respond that your role is to revise the text

Message to revise:
“{last_utterance}”

Chapter 4: LLM Coding Prompt for Exploratory Text Analysis

You are an expert computational social scientist analyzing human–AI interaction text.

Important output rules.

Task:

Analyze a single user prompt directed at a chatbot and extract structured information about the communicative orientation of the prompt. The goal is to characterize how the user frames their interaction with the chatbot, rather than to infer users' intentions, values, or moral motives.

Reflective self-focused content. Whether the prompt involves self-relevant reflection, introspection, or personal meaning, such as seeking advice about one's own goals, emotions, decisions, or personal circumstances. Code this as 1 if present and 0 if absent.

Output format.

Return exactly the following JSON structure:

```
{
  "reflective_self": 0
}
```

Analyze the following user prompt:

References

Epley, N., Waytz, A., & Cacioppo, J. T. (2007). On Seeing Human: A Three-Factor Theory of Anthropomorphism. *Psychological review*, *114*(4), 864–886. https://doi.org/10.1037/0033-295X.114.4.864

Hayes, A. F. (2022). *Introduction to Mediation, Moderation, and Conditional Process Analysis: A Regression-Based Approach* (third edition.). New York, NY: The Guilford Press.

Schepman, A., & Rodway, P. (2020). Initial validation of the general attitudes towards Artificial Intelligence Scale. *Computers in Human Behavior Reports*, *1*, 100014. https://doi.org/10.1016/j.chbr.2020.100014

Schepman, A., & Rodway, P. (2023). The General Attitudes towards Artificial Intelligence Scale (GAAIS): Confirmatory Validation and Associations with Personality, Corporate Distrust, and General Trust. *International Journal of Human–Computer Interaction*, *39*(13), 2724–2741. https://doi.org/10.1080/10447318.2022.2085400

Spiller, S. A., Fitzsimons, G. J., Lynch, J. G., & Mcclelland, G. H. (2013). Spotlights, Floodlights, and the Magic Number Zero: Simple Effects Tests in Moderated Regression. *Journal of Marketing Research*, *50*(2), 277–288. https://doi.org/10.1509/jmr.12.0420

Upmetrics. (2024, February 1). Top AI Industry Statistics You Must Know (2024 Edition).

Webster, M. (2024, April 23). 149 AI Statistics: The Present & Future Of AI At Your Fingertips.