ISSN: 2582-5526 www.thevoiceofcreativeresearch.com

The Voice of Creative Research

Vol. 8 & Issue 4 (October 2026)

Peer Reviewed Article

DOI: https://doi.org/10.53032/tvcr/2026.v8n4.05

From Artificial to Humanlike: A Systematic Literature Review of AI Humanisation and Customer Experience

Amit Punia
Assistant Professor,
School of Finance and Commerce,
Galgotias University, Greater Noida
Email: amitpoonia91@gmail.com

Akash Singh
School of Business Management Studies,
Central University of Haryana, Mahendragarh, India
Email: akashsingh5015@gmail.com

Abstract
Artificial intelligence (AI) at the customer interface chatbots, voice assistants, virtual agents and service robots is increasingly designed to look, sound and behave like a human being. This review synthesises the evidence on how that deliberate humanisation shapes customer experience (CX), drawing on 57 sources published between 1994 and 2026 across marketing, information systems, hospitality and psychology. Organising the literature through anthropomorphism theory, the Computers-Are-Social-Actors paradigm and the warmth–competence framework, the review examines the design mechanisms through which AI is humanised, the effects of humanisation on trust, engagement, satisfaction and loyalty-relevant behaviour, the boundary conditions that moderate or reverse those effects, and a dark-side literature documenting uncanny-valley discomfort, privacy inference and misplaced trust. The evidence supports a positive but modest and strongly conditional relationship between humanisation and customer evaluations. Humanisation appears to operate by raising customers' relational expectations of an AI agent: whether its effect is positive or negative depends on whether those expectations can plausibly be met, given the customer's emotional state, the warmth demanded by the service task, the timing of disclosure and the consequences of failure. Beyond a threshold of human-likeness, and in the presence of privacy concern or perceived deception, the relationship turns negative. The review offers an integrative framework, translates the evidence into heuristics for calibrating humanisation, and identifies gaps in longitudinal, generative-AI and emerging-market evidence.

Keywords: Artificial intelligence, Anthropomorphism, Humanisation, Customer Experience, Chatbots, Service Robots

INTRODUCTION

Artificial intelligence (AI) has moved from the back office to the front of the house. Where AI once optimised logistics, pricing or fraud detection largely unseen by the customer, it now routinely occupies the customer-facing role once reserved for a human employee: greeting an online shopper, answering a complaint, recommending a hotel room, or narrating a product demonstration. Chatbots, voice assistants, virtual agents and physically embodied service robots increasingly constitute the first, and sometimes the only, point of contact a customer has with a firm. A deliberate design choice has accompanied this shift. Rather than presenting these systems as neutral tools, many organisations actively humanise them, giving them names, faces, voices, personalities, humour and, increasingly, the conversational fluency of large language models.

This review synthesises the accumulated research evidence on the consequences of that humanisation for customer experience (CX). Humanisation is defined here as the set of visual, verbal, vocal and behavioural design choices that make an AI system appear, sound or act more like a human agent than a machine. The definition extends beyond physical resemblance to the mechanical, thinking and feeling capabilities distinguished by Huang and Rust (2021a), with feeling AI the capacity to recognise and respond to emotion representing the frontier most closely associated with perceived humanness. The question is not whether AI can be humanised, but what humanisation does to the experience customers have.

That question matters because customer experience has become one of the organising constructs of contemporary marketing. Lemon and Verhoef (2016) describe customer experience as the customer's cognitive, emotional, physical, sensorial and social response to a firm across the entire journey rather than at the point of transaction, and argue that firms increasingly compete on the quality of that experience rather than on product attributes alone. As AI has been inserted into more of the touchpoints that constitute this journey, how the design of that AI shapes the experience it produces has become a question with direct managerial consequence.

The theoretical starting point for most of this literature is anthropomorphism theory: the proposition, developed most influentially by Epley, Waytz and Cacioppo (2007), that people are motivated to attribute humanlike mental states to non-human agents when doing so satisfies a need for social connection or helps them explain and predict the agent's behaviour. Applied to AI, the theory predicts a relatively direct route from humanlike design cues to more favourable customer evaluations, mediated by warmth, social presence and trust. The empirical record is considerably less tidy than that prediction implies. Crolic, Thomaz, Hadi and Stephen (2022) show that anthropomorphising a service chatbot backfires when the customer is already angry, because it raises expectations of empathic competence that the chatbot then fails to meet. Work on the uncanny valley (Mathur & Reichling, 2015; Song et al., 2022) shows that the relationship between human-likeness and comfort is not monotonic but curvilinear, turning negative once an agent becomes almost, but not quite, indistinguishable from a person. Even the theoretical bedrock has been challenged directly: Heyselaar's (2023) replication of the original Computers-Are-Social-Actors studies found that the effect no longer holds for desktop computers, suggesting that some of the field's foundational claims may be historically contingent rather than universal. A synthesis of this literature is therefore not a formality; the findings genuinely pull in different directions, and a reader encountering only one or two studies could reasonably come away with an unbalanced picture.

This review does not claim to enter an unexamined field. Husain, Ratna and colleagues (2025) and

Li and colleagues (2025) each published a Theory-Context-Construct-Methods review of anthropomorphism and consumer response within months of one another; Mariani, Hashemi and Wirtz (2023) reviewed conversational agents, and Husain, Ratna and Saxena (2023) reviewed anthropomorphism in hospitality and tourism. The contribution offered here lies elsewhere. Drawing the boundary of the literature around humanisation rather than anthropomorphism admits a growing body of generative-AI work that describes humanlike qualities of large language model agents without using the anthropomorphism label, and which an anthropomorphism-anchored search therefore under-represents. The review also treats the boundary-condition and dark-side evidence as a unifying theme rather than a secondary caveat, and integrates the two into a single account of humanisation as a double-edged design lever rather than three loosely connected research streams.

Three research questions follow. First, through what design mechanisms is AI humanised in customer-facing contexts, and how do these mechanisms differ across chatbots, voice assistants and physically embodied service robots? Second, what effects does humanisation have on customer experience outcomes

specifically trust, engagement, satisfaction and loyalty-related behaviour and under what conditions do these effects strengthen, weaken or reverse? Third, what does the accumulating dark-side literature on uncanny-valley discomfort, privacy concern and misplaced trust imply for how firms should approach the design of humanised AI?

THEORETICAL BACKGROUND

Four bodies of theory recur across the corpus with sufficient frequency and explanatory importance to warrant treatment before the substantive synthesis: anthropomorphism theory, the Computers-Are-Social-Actors (CASA) paradigm, the warmth–competence framework, and the customer experience construct itself. A fifth subsection clarifies how humanisation is used in this review relative to the narrower construct of anthropomorphism. Figure 1 sets out how these bodies of theory relate to one another and to the organising construct of the review.

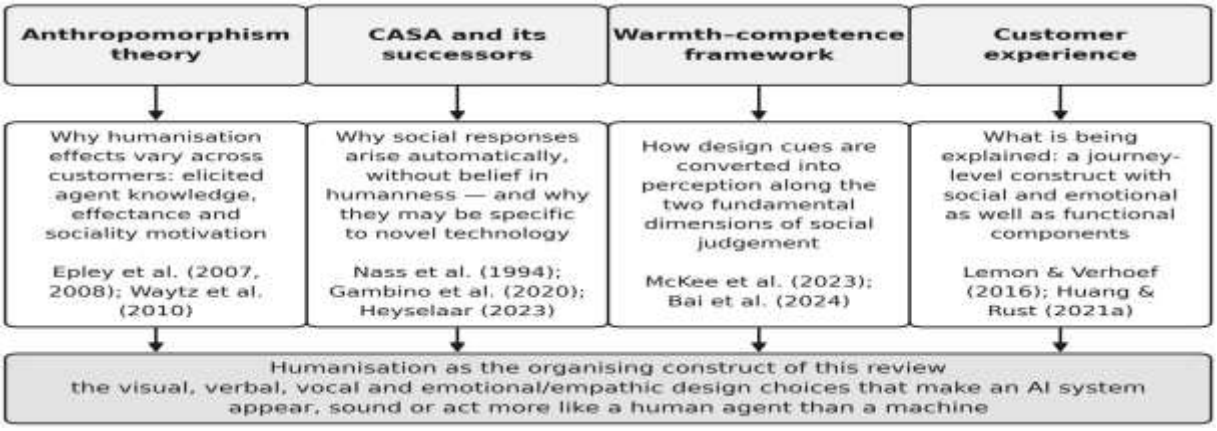


FIGURE 1

Theoretical foundations of the review and their explanatory roles

Anthropomorphism: the psychology of seeing human

The dominant psychological account of why people attribute humanlike qualities to non-human agents is Epley, Waytz and Cacioppo's (2007) three-factor theory. It proposes that anthropomorphism is driven by the interaction of elicited agent knowledge (the accessibility of human-centred concepts as a template for interpreting an unfamiliar agent), effectance motivation (the desire to explain and predict an agent's behaviour, which is easier when the agent is understood in familiar humanlike terms) and sociality motivation (the desire for social connection, heightened for example among people who are lonely or otherwise socially deprived). Epley, Akalis, Waytz and Cacioppo (2008) provide empirical support for the motivational components specifically, showing that dispositional loneliness and situational need for control both increase the tendency to anthropomorphise. Waytz, Cacioppo and Epley (2010) extend the theory by demonstrating that individual differences in anthropomorphism are stable and consequential, predicting the degree of moral concern, trust and social influence a person extends to a non-human agent.

The theory's central implication for this review is that humanisation is not simply a matter of what a firm builds into an AI system, but also of what the customer brings to the interaction. An identically designed chatbot may be strongly anthropomorphised by a customer seeking social connection and only weakly anthropomorphised by one who is not a point that recurs throughout the empirical literature and that cautions against treating humanisation as a fixed property of the technology alone.

Computers as social actors, and their contested present

A second, related but distinct line of theory concerns not why people anthropomorphise but how directly and automatically they apply social rules to machines. The CASA paradigm, originating in Nass, Steuer and Tauber's (1994) experiments, proposes that people apply social scripts politeness norms, reciprocity, in-group favouritism to computers largely mindlessly, without consciously believing the computer to be human or humanlike. This is an important distinction from anthropomorphism theory: where Epley and colleagues emphasise a motivated and sometimes deliberate attribution process, the original CASA account emphasises an overlearned, automatic one.

The paradigm has aged less comfortably than the anthropomorphism literature it partly inspired. Gambino, Fox and Ratan (2020) argue that it requires updating precisely because both people and technology have changed since the 1990s, and propose that users may now be developing distinct human–media social scripts rather than transferring human–human scripts wholesale; Lombard and Xu (2021) extend CASA into a broader Media Are Social Actors framework that explicitly incorporates mindful, and not only mindless, anthropomorphism. Most tellingly, Heyselaar (2023) found through direct replication that the core CASA effect no longer holds for interactions with desktop computers, suggesting the effect may be specific to genuinely novel or emergent technologies rather than a permanent feature of human–computer interaction.

This is a useful corrective for a review of this kind: findings generated using one generation of AI interface should not be assumed to transfer automatically to the generative, large language model agents now becoming the default.

The warmth–competence framework

The third framework, and the one doing the most direct explanatory work across the empirical literature, is the warmth–competence model of social perception, applied to AI most directly by McKee,

Bai and Fiske (2023). Across nine studies and more than three thousand participants they report that people judge AI systems along the same two fundamental dimensions that structure human social perception generally. Systems perceived as pursuing goals aligned with human interests are judged warmer; systems perceived to act more autonomously of human direction are judged more competent; and both judgments independently predict willingness to cooperate with the system.

This two-dimensional structure recurs throughout the applied literature, though its relative weighting shifts with context and with time. Bai and colleagues (2024) find that in a multimodal in-vehicle assistant context competence perceptions matter more early in the usage relationship, when safety and reliability are paramount, while warmth becomes the stronger driver of evaluation as the relationship matures, complicating any simple claim that more warmth is always better. Ng and colleagues (2024), studying a generative-AI hospitality dialogue system with more than 1,300 hotel guests, find that warmth and competence independently predict acceptance and satisfaction, moderated by the guest's readiness to engage with new technology, while Kim, Ryoo and colleagues (2023) show that the two jointly generate a sense of empathy from a chatbot that drives willingness to use it. Warmth and competence are therefore not merely descriptive labels for how humanised a system appears, but the specific psychological mechanism through which humanisation cues are converted into downstream customer experience outcomes.

Customer experience as a construct

Lemon and Verhoef (2016) offer the most widely adopted contemporary definition of the fourth foundation, conceptualising customer experience as a multidimensional construct spanning cognitive, emotional, physical, sensorial and social responses across the entire customer journey, encompassing pre-purchase, purchase and post-purchase stages, and shaped by touchpoints the firm controls only partially. Two implications matter here. First, because customer experience is explicitly social and emotional as well as cognitive and functional, it is a construct for which the humanlike qualities of an interface plausibly matter a great deal — arguably more than for a narrower construct such as task completion time. Second, because customer experience spans the whole journey rather than a single encounter, effects documented in isolated experimental studies, most of which examine one short interaction, may not generalise cleanly to the cumulative experience a customer has with a firm's AI systems over months of repeated use.

From anthropomorphism to humanisation

Much of the literature synthesised below is framed by its own authors as being about anthropomorphism specifically, typically operationalised through visual human-likeness, verbal cues, or, less commonly, vocal cues. This review treats anthropomorphism as the best-studied and most precisely measured component of the broader phenomenon of humanisation, but treats humanisation as the wider construct, encompassing not only static human-likeness but also dynamic, relational qualities such as responsiveness, personalisation and the capacity for something approximating empathy — dimensions foregrounded by Liu-Thompkins, Okazaki and Li (2022) in their treatment of artificial empathy and by Huang and Rust's (2021a) distinction between mechanical, thinking and feeling AI. Generative AI in particular complicates a purely visual or physical definition, since a text-only large language model can feel strikingly humanlike through conversational fluency and apparent emotional attunement alone, without any avatar or embodiment. The review's use of humanisation as the organising term is intended to accommodate that broader reality.

METHOD

Search strategy and selection criteria

The review adopts the general approach of a systematic literature review, synthesising evidence through an explicit and reportable search and selection process rather than a narrative review based on convenience sampling. Searching was conducted iteratively across academic search and citation-indexing tools, covering Scopus-, Web of Science- and PubMed-indexed journals, Semantic Scholar's index, and the SAGE, Springer, Wiley, Taylor & Francis and Elsevier ScienceDirect platforms. Three term families were combined in successive queries: (a) humanisation cues, including anthropomorphism, humanlike and artificial empathy; (b) AI form factors, including chatbot, conversational agent, voice assistant, service robot and large language model; and (c) customer-experience outcomes, including trust, engagement, satisfaction, loyalty and purchase intention. Table A1 reports the full term families, illustrative combined strings and platforms searched. Backward snowballing through the reference lists of the most relevant and most highly cited records recovered foundational theoretical work that a recency-weighted keyword search tends to under-represent (Nass et al., 1994; Epley et al., 2007).

Sources were included where they were peer-reviewed journal articles — together with a small number of advance-online articles from established journals, given the pace of publication in this area — written in English, and substantively concerned with anthropomorphism, humanisation or humanlike design of an AI-based system in a context where a customer, consumer, guest or service user was the subject of study. Empirical, conceptual and review articles were all eligible. Sources were excluded where AI featured only as a back-end capability with no customer-facing or human-likeness dimension, where the agent was studied outside any service or commercial context, or where the source was not peer reviewed. Table A2 states the criteria as applied. No lower date bound was imposed, so that foundational theoretical work could be retained.

Screening and extraction

Screening proceeded in two stages. The first assessed titles and abstracts against the inclusion criteria. The second entailed reading the full abstract and, in most cases, the findings and discussion sections of each source that passed the first screen, in order to confirm relevance and to extract theoretical framing, methodology, sample or context and key findings for synthesis. Figure 2 summarises the search and screening flow and the decision rules applied at each stage. Themes were identified deductively, anchored to the theoretical frameworks set out above, and refined inductively as recurring patterns — and recurring points of disagreement — emerged across the corpus. Where a quantitative synthesis of a narrower slice of this literature already exists, most notably Blut, Wang, Wünderlich and Brock's (2021) meta-analysis of anthropomorphism across 108 independent samples, its findings are drawn on directly rather tha re-derid.

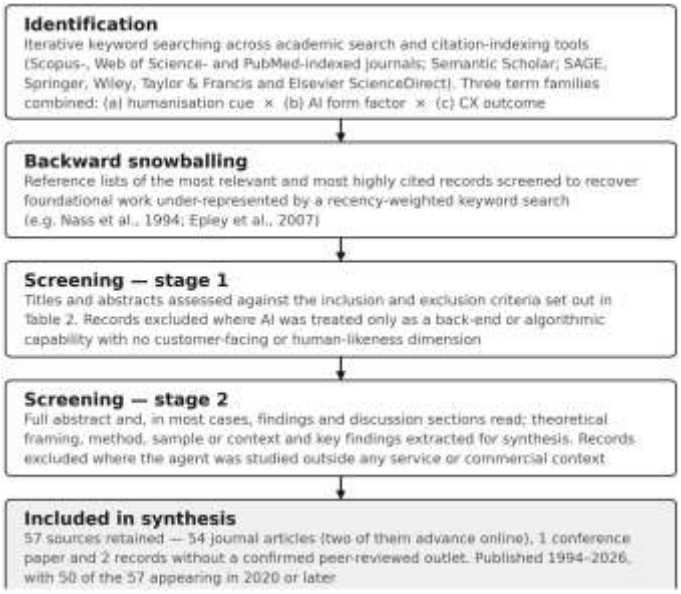


FIGURE 2

Search and screening flow

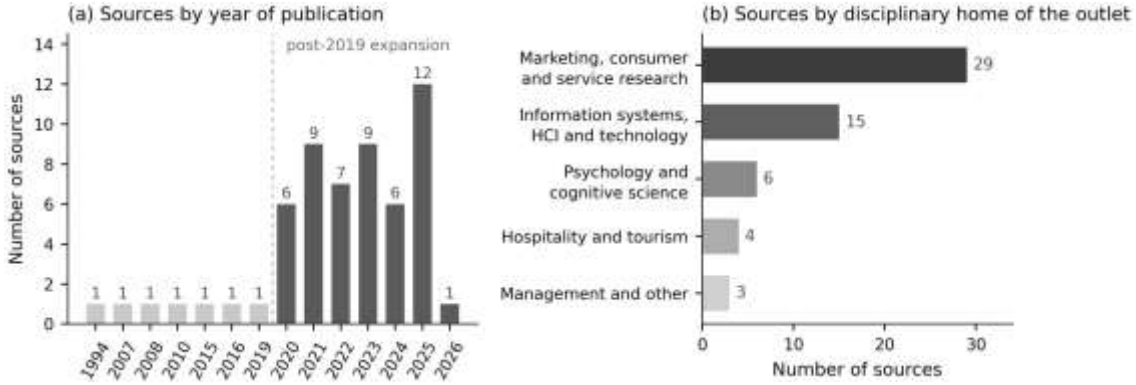
Note. Stage-wise record counts at identification and first-pass screening were not retained, because the search was conducted iteratively across several platforms rather than as a single dated export from one bibliographic database. The figure therefore reports the stages and decision rules rather than a fully enumerated PRISMA flow diagram.

Profile of the corpus

The resulting corpus comprises 57 sources: 54 journal articles, two of which were at advance-online publication at the time of writing, one conference paper, and two records for which a peer-reviewed outlet could not be confirmed. It is weighted toward empirical studies published between 2020 and 2026, which account for 50 of the 57 sources, but is anchored by foundational theoretical work reaching back to 1994 and 2007. The corpus spans marketing, consumer and service research, information systems and human–computer interaction, hospitality and tourism, and psychology. Figure 3 reports its distribution by year of publication and by the disciplinary home of the publishing outlet.

FIGURE 3

Distribution of the corpus by year of publication and by disciplinary home of the outlet



Note. n = 57. Panel (a) shows a marked expansion of the literature from 2020 onward, with pre-2019 sources shaded lighter to distinguish foundational work from the recent empirical wave. Panel (b) classifies each source by the disciplinary home of the outlet in which it appeared; two records without a confirmed peer-reviewed outlet are grouped under management and other outlets.

This disciplinary breadth indicates a genuinely multidisciplinary field rather than one confined to a single research community, but it also introduces heterogeneity the synthesis must handle with care. Constructs vary, in that trust, engagement, satisfaction and loyalty are related but distinct and are frequently operationalised differently across studies; contexts vary, in that retail, hospitality, banking and healthcare settings differ in ways that plausibly moderate effect sizes; and form factors vary, in that a physically embodied hotel robot and a text-based banking chatbot are not interchangeable stimuli.

ANALYSIS AND RESULTS

The synthesis is organised around five themes — the mechanisms through which AI is humanised, the effects of humanisation on trust, engagement and satisfaction, its effects on loyalty and behavioural outcomes, the boundary conditions that moderate or reverse those effects, and the accumulating dark-side literature — followed by a sixth subsection drawing together sector-specific evidence.

Mechanisms of AI humanisation

AI is humanised through at least four distinguishable design mechanisms, which frequently co-occur but are conceptually separable. The first is visual or physical anthropomorphism: giving a chatbot an avatar with a face, or a service robot a body with humanlike proportions, eyes or a head. Belanche, Casaló, Flavián and Schepers (2021) develop and test a Humanness-Value-Loyalty model on this dimension, showing that a robot's perceived physical human-likeness raises customers' expectations of both utilitarian and relational service value, which in turn predict loyalty intentions. Choi and colleagues (2020) manipulate humanoid versus non-humanoid appearance and find that humanoid robots are more strongly associated with perceived warmth specifically, not competence — an important dissociation that recurs across this literature.

The second mechanism is verbal or conversational: first-person pronouns, informal register, humour and a named identity in text-based interaction. Adam, Wessel and Benlian (2020) show in a randomised experiment that verbal anthropomorphic cues, independent of any visual avatar, increase the likelihood that customers comply with a chatbot's requests, mediated by a heightened sense of social presence, while Xu and colleagues (2022) show that humour used during a service failure increases customers' tolerance of that failure, again mediated by perceived warmth and competence. The third is vocal or auditory: a human-sounding voice, natural prosody and conversational timing. Fakhimi, Garry and Biggemann (2023) find through qualitative interviews that a humanised voice interface is central to the development of trust in human–machine relationships, though their informants describe both bright and dark sides to that trust; Oliveira, Lizarelli, Teixeira and Mendes (2023) and Aw, Tan, Cham, Raman and Ooi (2022) both converge on the finding that humanlike attributes function alongside, rather than instead of, conventional technology-acceptance factors such as perceived usefulness and ease of use.

The fourth mechanism, and the most consequential for the current moment, is emotional or empathic humanisation: the capacity to recognise and respond appropriately to a customer's emotional state, increasingly enabled by generative AI. Liu-Thompkins, Okazaki and Li (2022) argue that this dimension, which they term artificial empathy, must become a first-order design consideration precisely because conventional AI agents are so often perceived as cold, while cautioning that it is not always desirable and can feel manipulative or hollow. Truong and colleagues (2025) show that anthropomorphism and responsiveness both increase social presence and perceived empathy, which in turn predict continuance intention, though the effect is dampened among customers with strong privacy concerns. Kumar and colleagues (2025) extend the line into the generative-AI era, showing that anthropomorphic characteristics of generative chatbots raise perceived competence, which combines with perceived warmth to produce a sense of authenticity, and that this authenticity is what ultimately drives engagement, experience and willingness to recommend — positioning perceived authenticity as an important and comparatively under-studied mediating construct. Table 1 summarises the four mechanisms.

TABLE 1

Mechanisms of AI humanisation

Mechanism	Representative design cues	Dimension primarily activated	Representative evidence
Visual / physical	Avatar with facial features; humanoid body shape, eyes and head; expressive gesture	Warmth more than competence	Belanche et al. (2021); Choi et al. (2020)

Mechanism	Representative design cues	Dimension primarily activated	Representative evidence
Verbal / conversational	First-person pronouns; informal register; a named identity; humour; self-disclosure	Warmth, via social presence	Adam et al. (2020); Xu et al. (2022)
Vocal / auditory	Human-sounding voice; natural prosody; conversational timing and turn-taking	Warmth and relational trust	Fakhimi et al. (2023); Oliveira et al. (2023); Aw et al. (2022)
Emotional / empathic	Emotion recognition; affect-congruent response; apology and reassurance; open-ended generative dialogue	Warmth and competence jointly, via perceived authenticity	Liu-Thompkins et al. (2022); Truong et al. (2025); Kumar et al. (2025)

Effects on trust, engagement and satisfaction

The weight of evidence supports a broadly positive relationship between AI humanisation and the immediate relational outcomes of trust, engagement and satisfaction, though almost every study reporting a positive main effect also reports a boundary condition on it. Park and colleagues (2024), studying banking chatbots, find that anthropomorphism and personalisation both significantly increase customer satisfaction, operating through perceived usefulness and ease of use. Gomes and colleagues (2025), with a sample of over thirteen hundred respondents, find that chatbot anthropomorphism raises purchasing decision-making only indirectly, through customer engagement as a mediator — underscoring that anthropomorphism's effect on behaviour is rarely direct, but works by first changing how engaged or emotionally involved the customer becomes. Alalwan and colleagues (2025), applying a flow-experience model to AI chatbots used by five hundred customers in Jordan, similarly find that anthropomorphism predicts online flow experience, which then shapes emotional brand attachment and loyalty.

At the level of quantitative synthesis, Blut and colleagues' (2021) meta-analysis of 108 independent samples and over eleven thousand participants offers the most authoritative statement available. Their model identifies customer traits such as computer anxiety, sociodemographic factors and robot design features as antecedents of anthropomorphism, and finds that anthropomorphism's effect on use intention operates primarily through perceived robot characteristics such as intelligence and functional characteristics such as usefulness, with relational characteristics such as rapport playing a smaller mediating role than anthropomorphism theory alone might predict. Critically, their moderator analysis shows the effect is not uniform: it depends on the type of service, being more strongly helped by anthropomorphism in possession-processing services than in services requiring greater cognitive or emotional engagement, and on robot gender. This meta-analytic evidence tempers the more uniformly positive picture that emerges from any single experimental study: at the level of the accumulated evidence base, anthropomorphism helps, but its effect size is neither large nor constant across contexts.

Effects on loyalty and behavioural outcomes

A smaller but growing set of studies traces the relationship between humanisation and more distal

behavioural outcomes. Li and colleagues (2023), studying friendship AI chatbots such as Replika, find that perceived warmth and perceived competence independently increase continuance-using intention and engagement, mediated by perceived usefulness and moderated by the customer's need to belong and the sensitivity of the information shared. Qi and colleagues (2025) extend this into parasocial-relationship theory, showing that AI agents designed to evoke an assistant script primarily heighten perceived competence while those evoking a friend script primarily heighten perceived warmth, with gendered design traits moderating which pathway dominates — a refinement suggesting firms can deliberately design toward one dimension or the other depending on the relational role the agent is intended to play. Belanche and colleagues' (2021) Humanness-Value-Loyalty model is again relevant, since it explicitly traces the path from perceived human-likeness through service-value expectations to loyalty intentions. Xiao and Kumar (2021) offer a useful corrective to an over-simple reading of this literature, arguing that service robots are for now best understood as a complement to human service employees in most contexts rather than an outright substitute — a framing consistent with Huang and Rust's (2021a) argument that AI is most effective when it augments rather than replaces human judgment and relational labour.

Boundary conditions: when humanisation backfires

The theme that most directly complicates any straightforward reading of this literature concerns the conditions under which humanisation's benefits weaken or reverse. The first is the customer's emotional state. Crolic and colleagues' (2022) study, combining a large real-world dataset from an international telecommunications firm with four controlled experiments, identifies a clear reversal: when a customer enters a chatbot interaction already angry, anthropomorphism reduces satisfaction, firm evaluation and purchase intention, whereas the same cues have no such negative effect, and may help, for customers who are not angry. The mechanism is an expectancy violation, since an anthropomorphised chatbot raises expectations of empathic competence that an angry customer is more likely to have disappointed. The second condition is whether the service task is perceived to require warmth. Peng and colleagues (2022) show that consumers resist assigning AI to high-warmth tasks because they perceive a poor fit between agent and task, a resistance reduced but not eliminated when the AI supports rather than replaces a human employee, while Yang and colleagues (2021) report that anthropomorphic agents are preferred under high perceived customer control but trigger a sense of threat under low control, an effect significant only in explicitly social service scenarios.

The third condition is disclosure and its timing. Luo, Tong, Fang and Qu's (2019) field experiment, involving more than six thousand customers receiving outbound sales calls, found that undisclosed chatbots performed on par with proficient human workers and substantially better than inexperienced ones, but that disclosing the chatbot's identity before the conversation reduced purchase rates by over seventy-nine per cent, an effect traced to customers perceiving the disclosed agent as less knowledgeable and less empathetic. The penalty is smaller when disclosure follows established rapport and among customers with prior AI experience. The fourth is failure and its aftermath: Choi and colleagues (2020) find that because humanoid robots evoke stronger warmth expectations, customers are more dissatisfied when a humanoid robot's service process fails, yet the same robots are better positioned to recover through apology or explanation, because customers extend to them the social scripts CASA predicts. Table 2 sets out these conditions, the direction in which each shifts the relationship, and the design implication that follows. Together they point to a consistent underlying logic: humanisation raises the customer's relational expectations of the system,

and whether that helps or harms depends on whether those expectations can plausibly be met.

Boundary condition	**Effect on the relationship**	**Evidence**	**Design implication**
Customer emotional state	Reverses. Anthropomorphism lowers satisfaction, firm evaluation and purchase intention among customers who are already angry, through violated expectations of empathic competence	Crolic et al. (2022)	De-emphasise humanlike cues on complaint and escalation paths
Task–warmth fit	Weakens. Consumers resist assigning high-warmth tasks to AI because of perceived task–agent misfit; resistance falls, without disappearing, when AI supports rather than replaces an employee	Peng et al. (2022)	Reserve strongly humanised agents for tasks where warmth is genuinely relevant, or pair them with a human
Perceived control over the encounter	Reverses. Under high perceived control, more anthropomorphic agents are preferred; under low perceived control the same cues trigger threat, an effect significant in social service scenarios	Yang et al. (2021)	Give customers visible control, such as an opt-out or a handover, before raising human-likeness
Disclosure and its timing	Reverses when disclosure precedes the conversation: purchase rates fell by more than 79 per cent. The penalty is smaller when disclosure follows established rapport and among customers with prior AI experience	Luo et al. (2019)	Disclose, and design the point at which disclosure occurs; treat honesty as a constraint rather than as a variable
Service failure	Two-sided. Humanoid design	Choi et al.	Invest in recovery scripts

and recovery	deepens dissatisfaction when the service process fails, but also improves recovery through apology and	(2020)	wherever human-likeness is high

The dark side: uncanny valley, privacy and misplaced trust

A parallel and increasingly substantial literature documents the more unambiguous costs of humanisation, most influentially through the uncanny valley hypothesis: comfort and likability rise with human-likeness only up to a point, after which further increases produce a sharp drop into discomfort or eeriness. Mathur and Reichling's (2015) quantitative mapping of this effect across eighty real-world robot faces remains a key empirical anchor. Song and colleagues (2022) provide a direct demonstration in a commercial chatbot context, showing that increasing a chatbot avatar's visual realism and animacy increases feelings of eeriness, which in turn reduce trust, purchase intention and willingness to reuse the chatbot, moderated by how familiar the customer already is with the specific avatar.

Troshani and colleagues (2020) extend the concept beyond visual appearance to humanness more broadly, including perceived intelligence, finding that consumers experience genuine discomfort as apparent capability begins to approach without quite matching a human's. Kim, Ryoo and Choi (2024) raise a more specifically ethical concern, identifying an uncanny valley of mind: as chatbots become more anthropomorphised, consumers increasingly infer that the agent may possess humanlike autonomous intent to misuse their personal data, a concern a more mechanised chatbot does not evoke to the same degree, and one the authors show can be mitigated through transparent algorithmic disclosure about how data is used. In a healthcare context, Sestino and colleagues (2025) find that avatars with high humanlike similarity trigger stronger negative emotional responses, mediated by perceived creepiness, though the effect is attenuated among users with higher digital literacy. Most directly, Hasan and colleagues (2025) argue that anthropomorphic design can generate a form of misplaced trustworthiness in which the customer's trust outruns what the system warrants, with consumer-protection implications that satisfaction metrics do not capture. Table 3 summarises these risks and the mitigations the evidence supports.

Sector-specific evidence

Hospitality and tourism contexts feature prominently in this corpus: Yang and colleagues (2021) and Xu and colleagues (2022) both use hospitality service-failure scenarios, Ng and colleagues (2024) study a generative-AI hospitality dialogue system directly, and a dedicated review of anthropomorphism within the sector already exists (Husain, Ratna, & Saxena, 2023). Ahmad and colleagues (2026) extend the line into digital intermediary contexts, finding that anthropomorphism in online travel agency chatbots raises both trust and engagement, though perceived risk plays a weaker role than anticipated and trust converts into engagement only once the user has sufficient platform experience. This concentration plausibly reflects the fact that hospitality service quality has long been theorised as depending on interpersonal warmth and emotional labour in a way that transactional retail service does not. Retail and e-commerce are almost as well represented, particularly around voice commerce (Oliveira et al., 2023; Aw et al., 2022; Fakhimi et al., 2023) and chatbot shopping assistance (Song et al., 2022; Gomes et al., 2025), with Blut, Wünderlich and

Brock's (2024) meta-analysis of retail customers' use of virtual assistants complementing their earlier synthesis of service robots. Banking is represented chiefly by Park and colleagues (2024), and healthcare only sparingly, through Sestino and colleagues (2025), even though the stakes of getting humanisation right are plausibly highest there. Table 4 draws this evidence together and indicates where coverage is thinnest relative to the practical stakes involved.

TABLE 3

Dark-side risks of AI humanisation and documented mitigations

Risk	**Mechanism**	**Evidence**	**Mitigation supported by the evidence**
Uncanny-valley eeriness	Visual realism and animacy beyond a threshold raise eeriness, which reduces trust, purchase intention and willingness to Reuse	Mathur & Reichling (2015); Song et al. (2022)	Calibrate realism below the threshold; familiarity with the specific avatar moderates the effect
Discomfort with perceived intelligence	Unease rises as the agent's apparent capability approaches, without matching, a human's	Troshani et al. (2020)	Set capability expectations explicitly within the interface
Uncanny valley of mind	Anthropomorphic cues raise the inference that the agent has humanlike autonomous intent to misuse personal data	Kim et al. (2024)	Transparent algorithmic disclosure of how customer data is used
Creepiness in high-stakes settings	High humanlike similarity in healthcare avatars triggers stronger negative emotional response, mediated by perceived creepiness	Sestino et al. (2025)	Moderate similarity where users are vulnerable; the effect attenuates among users with higher digital literacy
Misplaced trustworthiness	Anthropomorphic design induces trust that outruns what the system actually warrants	Hasan et al. (2025)	Treat as a consumer-protection question; conventional satisfaction metrics will not detect it

TABLE 4

Sector coverage across the corpus

Sector	Representative sources	State of the evidence
Hospitality and tourism	Husain, Ratna & Saxena (2023); Yang et al. (2021); Xu et al. (2022); Ng et al. (2024); Ahmad et al. (2026)	Best developed, with a dedicated systematic review already in place; strongest on service failure and on generative-AI dialogue systems
Retail and e-commerce	Song et al. (2022); Gomes et al. (2025); Blut et al. (2024)	Well represented, and supported by a dedicated meta-analysis of retail customers' use of virtual assistants
Voice commerce	Oliveira et al. (2023); Aw et al. (2022); Fakhimi et al. (2023)	Moderate. Converges on humanlike attributes operating alongside, rather than instead of, conventional technology-acceptance factors
Banking and financial services	Park et al. (2024)	Thin. The proposition that competence and trust outweigh warmth in this setting has not been tested through direct cross-sector comparison
Healthcare	Sestino et al. (2025)	Sparse relative to the stakes, given the sensitivity of the information involved and the vulnerability of many service Users
Advertising and media	Kim et al. (2024); Wu et al. (2024)	Emerging, and concerned principally with privacy inference and uncanny-valley response rather than with service outcomes

DISCUSSION

An integrative account of humanisation and customer experience

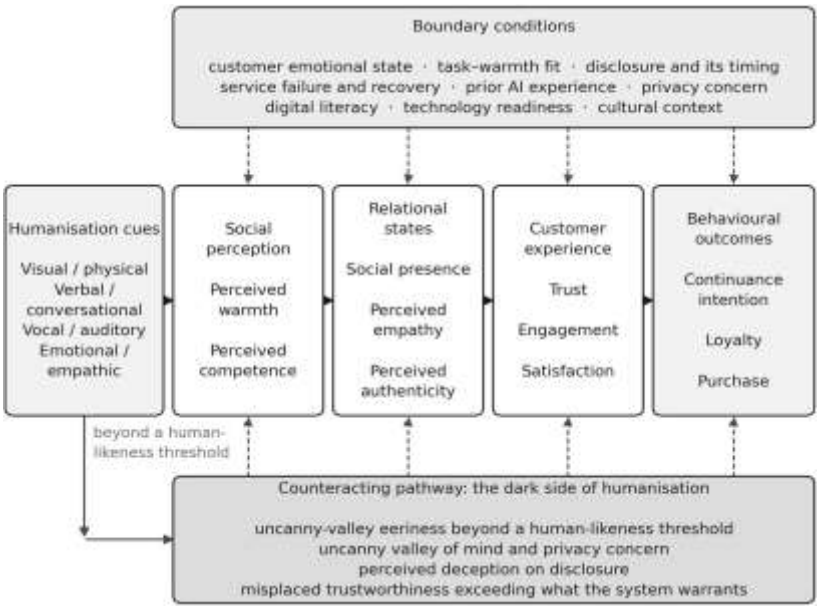
Drawing the themes together, the literature supports an integrative, if necessarily simplified, account of how AI humanisation shapes customer experience. Design-level cues visual, verbal, vocal and emotional or empathic operate on the customer's perception along the two dimensions of the warmth–competence framework (McKee et al., 2023), with visual and physical cues loading more heavily onto warmth (Choi et al., 2020) and cues signalling apparent capability loading more heavily onto competence, and with the

relative importance of each shifting over the course of the relationship (Bai et al., 2024). These perceptions generate the relational states social presence, perceived empathy and perceived authenticity that the artificial-empathy and generative-AI literature identifies as the proximate drivers of engagement (Liu-Thompkins et al., 2022; Kumar et al., 2025), which feed forward into the immediate outcomes of trust and satisfaction (Park et al., 2024; Alalwan et al., 2025) and, more distally, into loyalty-relevant behaviour (Belanche et al., 2021; Li et al., 2023).

Running alongside this broadly positive pathway is a second, counteracting pathway that becomes active once human-likeness crosses a threshold, or once expectations are violated by disclosure, failure or an emotionally charged encounter. Here the uncanny-valley and dark-side literature (Mathur & Reichling, 2015; Song et al., 2022; Kim et al., 2024; Hasan et al., 2025) and the boundary-condition literature (Crolic et al., 2022; Yang et al., 2021; Luo et al., 2019) converge on a shared conclusion, even though they are rarely cited together: humanisation does not act on customer experience directly, but by raising the customer's relational expectations of the system, and the sign and size of its effect depend on whether those expectations are subsequently met. A moderately humanised, warm but not deceptively humanlike system, used for a task where warmth is genuinely relevant, with honest disclosure of its non-human status and a graceful failure-recovery design, appears from this literature to be a considerably safer design target than a system optimised purely to maximise apparent human-likeness. Figure 4 represents this account in schematic form.

FIGURE 4

An integrative framework of AI humanisation and customer experience



Note. Solid arrows denote the principal positive pathway; dashed arrows denote moderating and counteracting influences. The framework is derived from the synthesis presented here and is offered as a structure for future empirical testing rather than as an estimated model.

Theoretical implications

The synthesis clarifies why the field's two most-cited theoretical frameworks anthropomorphism theory (Epley et al., 2007) and the CASA paradigm (Nass et al., 1994) though often invoked almost interchangeably in the applied literature, are better treated as complementary than as identical accounts. Anthropomorphism theory explains why some customers, and some situations, generate stronger humanisation effects than others, through sociality and effectance motivation; CASA, particularly in its updated forms (Gambino et al., 2020; Lombard & Xu, 2021), explains the more automatic, script-based social responses that occur even absent any deliberate belief that the agent is humanlike. Heyselaar's (2023) finding that the original CASA effect does not replicate for desktop computers is a reminder that this automatic pathway may be specific to novel technology, with an important implication for generative AI: the current strength of CASA-type effects around conversational agents may partly reflect the genuine novelty of large language model interfaces, and could plausibly weaken as those interfaces become as familiar and unremarkable as the desktop computer once was.

A second theoretical contribution is to draw the warmth–competence framework, the boundary-condition literature and the dark-side literature into a single account of humanisation as a double-edged design lever, rather than treating them, as much of the primary literature does, as three loosely connected streams. The account proposes that expectation formation around an agent's true nature and capabilities is the thread connecting an otherwise disparate set of boundary conditions emotional state, task type, disclosure timing and failure a proposition that is itself directly testable.

MANAGERIAL IMPLICATIONS

The pattern of evidence reviewed here suggests that firms should treat the degree of AI humanisation as a design variable to be calibrated to context rather than maximised by default. Where a service task genuinely calls for warmth, and where customers are unlikely to arrive in a heightened emotional state, moderate humanisation appears to carry real relational benefit; where a task is more purely functional, or where failure is likely and its consequences serious, overtly emphasising human-likeness may raise expectations the system cannot meet. Honest and well-timed disclosure of an agent's AI status, combined with transparent communication about how customer data is used, appears to reduce rather than eliminate the risk of a negative reaction, and is difficult to justify withholding on customer-experience grounds alone once the ethical dimension raised by Hasan and colleagues (2025) is taken into account.

Three further implications follow. Brands whose agents sit low on perceived human-likeness should invest in the mechanisms that raise warmth without raising visual realism — conversational register, responsiveness and empathic acknowledgement — since these carry the relational benefit with less uncanny-valley exposure. Because humanisation raises expectations of recovery as much as of service, investment in failure-recovery scripts should scale with the degree of human-likeness rather than be treated as a separate programme. And since privacy concern, digital literacy and technology readiness all moderate the response, a single optimal level of humanisation across a customer base is unlikely to exist. Table 5 translates the boundary-condition evidence into a set of calibration heuristics.

TABLE 5

Calibrating the degree of humanisation: design heuristics derived from the review

Design decision	Lower humanisation is indicated when…	Higher humanisation is indicated when…
Task type	The task is transactional, informational or possession-processing	The task calls for warmth, reassurance or relationship-building
Customer emotional state	Customers are likely to arrive frustrated, as on complaint, outage or billing-dispute paths	Customers are likely to be in a neutral or positive state, as in browsing, onboarding or recommendation
Consequence of failure	Failure is likely and its consequences are serious	Failure is unlikely or easily recovered, and recovery scripts are already in place
Data sensitivity	The interaction involves sensitive personal, health or financial data	The interaction involves low-sensitivity data
Relationship stage	The relationship is new, and competence signals carry most weight	The relationship is established, and warmth has become the stronger driver of evaluation
Customer segment	The segment shows high privacy concern or low technology readiness	The segment has prior AI experience and higher digital literacy

Note. Disclosure is not a variable in this table. The evidence reviewed here supports disclosing an agent's AI status in all conditions and treating the timing and framing of that disclosure, rather than the fact of it, as the design choice.

LIMITATIONS AND FUTURE RESEARCH DIRECTIONS

Four limitations should be weighed alongside these conclusions. First, the search process, though iterative and wide-ranging across academic search tools and publisher platforms, was not a single formal export from one bibliographic database following a fully specified PRISMA protocol, and screening was conducted by a single reviewer rather than by independent coders reconciling disagreements; the review is therefore best treated as a structured map of the field rather than as a substitute for a fully protocolled bibliometric search. Second, bibliographic details were compiled from academic search and indexing tools, and while every cited source was verified to exist and reviewed for its substantive content, page-level details for a small number of entries should be cross-checked against the publisher record. Third, the corpus is weighted toward English-language, internationally indexed journals, which may under-represent relevant work published in regional or non-English outlets. Fourth, because generative AI has changed rapidly even within the period covered, some earlier findings particularly those generated using rule-based chatbots or narrow-AI voice assistants may already be dated, which argues for treating this synthesis as a snapshot to be periodically refreshed.

Several research gaps follow from the synthesis itself. Most of the evidence base is cross-sectional, using single-encounter surveys or short laboratory experiments; Bai and colleagues' (2024) finding that the relative importance of warmth and competence shifts as a usage relationship matures is a rare exception whose temporal design is worth emulating. Longitudinal or field-experimental evidence tracing how

humanisation effects evolve over months of repeated interaction, closer to how Lemon and Verhoef (2016) conceptualise the customer journey, remains scarce. The literature is also weighted toward North American, Western European and East Asian samples, alongside a smaller cluster of Middle Eastern evidence (Alalwan et al., 2025); South Asian and specifically Indian evidence is largely absent, despite India's scale as a services and digital-commerce market and despite cultural factors, such as differing norms around service formality, deference and interpersonal warmth, that plausibly moderate the effects documented here.

Generative AI is addressed by only a handful of the most recent sources (Kumar et al., 2025; Grewal et al., 2025; Hermann & Puntoni, 2025), while the great majority of the evidence base, including the foundational warmth–competence and uncanny-valley studies, was generated using earlier-generation rule-based or narrow-AI systems; whether the same mechanisms apply with the same strength to agents capable of open-ended humanlike dialogue is an assumption the literature has not yet tested as thoroughly as its practical importance warrants. Healthcare and financial services remain understudied relative to hospitality and retail, despite arguably higher stakes for trust and vulnerability, and almost no study in this corpus examines humanisation from the perspective of the frontline employees whose roles increasingly sit alongside these systems. Table 6 states these gaps as research questions, each paired with a design capable of addressing it.

TABLE 6

Future research agenda

Gap	Research question	Suggested design
Temporal dynamics	Does the relative weight of warmth and competence shift as a customer–AI relationship matures, and do humanisation effects decay or compound over repeated encounters?	Longitudinal panel or field experiment across repeated encounters, extending the temporal design used by Bai et al. (2024)
Emerging-market and Indian evidence	Do norms of service formality, deference and interpersonal warmth in South Asian markets moderate effects documented in predominantly Western and East Asian samples?	Cross-cultural experiment using matched stimuli in at least two national markets, with measured rather than assumed cultural moderators
Generative-AI mechanisms	Do warmth–competence and uncanny-valley mechanisms hold with the same strength for open-ended large language model agents as for rule-based chatbots and embodied robots?	Direct replication of canonical studies substituting a generative agent as the stimulus, following the replication logic of Heyselaar (2023)
Under-researched sectors	Does the relative weight of warmth and competence invert in high-stakes settings such as healthcare and financial services?	Cross-sector experimental comparison holding agent design constant while varying the service context

Frontline-employee perspective	How does customer-facing AI humanisation affect the frontline employees who work alongside these systems?	Dyadic or multi-level design linking employee-level and customer-level data within the same service encounters
Disclosure as a design variable	Can the disclosure penalty be reduced through timing, framing or repetition without compromising the honesty of the disclosure?	Field experiment varying the point and framing of disclosure, with purchase and trust as outcomes
Authenticity as a mediator	Is perceived authenticity empirically distinct from warmth, competence and social presence, or a relabelling of them?	

CONCLUSION

AI tools are humanised deliberately and increasingly, and that humanisation has real, measurable, but far from unconditional consequences for customer experience. Across 57 sources spanning marketing, information systems, hospitality and psychology, the weight of evidence supports the view that humanlike design cues, operating through the two fundamental dimensions of perceived warmth and competence, generally increase customer trust, engagement and satisfaction and, more distally, loyalty-relevant behaviour. The same evidence base is equally clear that these benefits are contingent: on the customer's emotional state, on whether the service task genuinely calls for warmth, on honest and well-timed disclosure of the agent's non-human status, and on what happens when the system fails. Beyond a certain threshold of human-likeness, and in the presence of privacy concern or perceived deception, the relationship between humanisation and customer experience turns negative, a pattern the uncanny-valley and dark-side literatures document with increasing precision. The practical implication is that humanisation is a lever to be calibrated rather than a quality to be maximised, and the theoretical implication is that its effects are best understood through the expectations it creates rather than through human-likeness itself. The clearest opportunities for future work lie in generative-AI-era evidence, in longitudinal designs that track the customer–AI relationship over time rather than at a single encounter, and in emerging-market contexts that remain largely absent from the literature assembled here. Measurement work establishing discriminant validity, followed by competitive mediation testing.

TABLE 6

Future research agenda

Gap	Research question	Suggested design
Temporal dynamics	Does the relative weight of warmth and competence shift as a customer–AI relationship matures, and do humanisation effects decay or compound over repeated encounters?	Longitudinal panel or field experiment across repeated encounters, extending the temporal design used by Bai et al. (2024)
Emerging-market and Indian evidence	Do norms of service formality, deference and interpersonal warmth in South Asian markets moderate effects documented in predominantly Western and East Asian samples?	Cross-cultural experiment using matched stimuli in at least two national markets, with measured rather than assumed cultural moderators
Generative-AI mechanisms	Do warmth–competence and uncanny-valley mechanisms hold with the same strength for open-ended large language model agents as for rule-based chatbots and embodied robots?	Direct replication of canonical studies substituting a generative agent as the stimulus, following the replication logic of Heyselaar (2023)
Under-researched sectors	Does the relative weight of warmth and competence invert in high-stakes settings such as healthcare and financial services?	Cross-sector experimental comparison holding agent design constant while varying the service context
Frontline-employee perspective	How does customer-facing AI humanisation affect the frontline employees who work alongside these systems?	Dyadic or multi-level design linking employee-level and customer-level data within the same service encounters
Disclosure as a design variable	Can the disclosure penalty be reduced through timing, framing or repetition without compromising the honesty of the disclosure?	Field experiment varying the point and framing of disclosure, with purchase and trust as outcomes
Authenticity as a mediator	Is perceived authenticity empirically distinct from warmth, competence and social presence, or a relabelling of them?	Measurement work establishing discriminant validity, followed by competitive mediation testing

REFERENCES

Adam, M., Wessel, M., & Benlian, A. (2020). AI-based chatbots in customer service and their effects on user compliance. Electronic Markets, 30(3), 427–445.

Ahmad, K., et al. (2026). Humanizing AI: How anthropomorphic chatbots drive customer engagement in online travel agencies. International Journal of Tourism Research. Advance online publication.

Alalwan, A. A., et al. (2025). Examining the impact of anthropomorphism and AI-chatbots service quality on online customer flow experience: Exploring the moderating role of telepresence. Journal of Consumer Marketing.

Aw, E. C.-X., Tan, G. W.-H., Cham, T.-H., Raman, R., & Ooi, K.-B. (2022). Alexa, what's on my shopping list? Transforming customer experience with digital voice assistants. Technological Forecasting and Social Change, 180.

Bai, S., et al. (2024). Warmth trumps competence? Uncovering the influence of multimodal AI anthropomorphic interaction experience on intelligent service evaluation. Technological Forecasting and Social Change.

Belanche, D., Casaló, L. V., Flavián, C., & Schepers, J. (2021). Examining the effects of robots' physical appearance, warmth, and competence in frontline services: The Humanness-Value-Loyalty model. Psychology & Marketing, 38(12).

Blut, M., Wang, C., Wünderlich, N. V., & Brock, C. (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), 632–658.

Blut, M., Wünderlich, N. V., & Brock, C. (2024). Facilitating retail customers' use of AI-based virtual assistants: A meta-analysis. Journal of Retailing, 100(2), 293–315.

Cheng, X., Zhang, X., Cohen, J., & Mou, J. (2022). Human vs. AI: Understanding the impact of anthropomorphism on consumer response to chatbots from the perspective of trust and relationship norms. Information Processing & Management, 59(3), 102940.

Chiang, A.-H., et al. (2022). Emotion and service quality of anthropomorphic robots. Technological Forecasting and Social Change.

Choi, S., et al. (2020). To err is human(-oid): How do consumers react to robot service failure and recovery? Journal of Service Research.

Crolic, C., Thomaz, F., Hadi, R., & Stephen, A. T. (2022). Blame the bot: Anthropomorphism and anger in customer–chatbot interactions. Journal of Marketing, 86(1).

Davenport, T., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. Journal of the Academy of Marketing Science, 48, 24–42.

Epley, N., Akalis, S., Waytz, A., & Cacioppo, J. T. (2008). When we need a human: Motivational determinants of anthropomorphism. Social Cognition, 26(2), 143–155.

Epley, N., Waytz, A., & Cacioppo, J. T. (2007). On seeing human: A three-factor theory of anthropomorphism. Psychological Review, 114(4), 864–886.

Fakhimi, A., Garry, T., & Biggemann, S. (2023). The effects of anthropomorphised virtual conversational assistants on consumer engagement and trust during service encounters. Australasian Marketing Journal.

Gambino, A., Fox, J., & Ratan, R. A. (2020). Building a stronger CASA: Extending the Computers Are Social Actors paradigm. Human-Machine Communication, 1, 71–85.

Gomes, S., et al. (2025). Anthropomorphism in artificial intelligence: A game-changer for brand marketing. Future Business Journal, 11.

Grewal, D., Satornino, C. B., Davenport, T., & Guha, A. (2025). How generative AI is shaping the future of marketing. Journal of the Academy of Marketing Science, 53(3), 702–722.

Han, E., Yin, D., & Zhang, H. (2023). Bots with feelings: Should AI agents express positive emotion in customer service? Information Systems Research, 34(3), 1296–1311.

Hasan, R., et al. (2025). The dark side of AI anthropomorphism: A case of misplaced trustworthiness in service provisions [Manuscript].

Hermann, E. (2021). Anthropomorphized artificial intelligence, attachment, and consumer behavior. Marketing Letters, 32(4).

Hermann, E., & Puntoni, S. (2025). Generative AI in marketing and principles for ethical design and deployment. Journal of Public Policy & Marketing, 44(3), 332–349.

Heyselaar, E. (2023). The CASA theory no longer applies to desktop computers. Scientific Reports, 13, 19664.

Huang, M.-H., & Rust, R. T. (2021a). A strategic framework for artificial intelligence in marketing. Journal of the Academy of Marketing Science, 49(1), 30–50.

Huang, M.-H., & Rust, R. T. (2021b). Engaged to a robot? The role of AI in service. Journal of Service Research, 24(1), 30–41.

Husain, R., et al. (2025). Anthropomorphism and consumer experience: Review and future research agenda. International Journal of Consumer Studies.

Husain, R., Ratna, V. V., & Saxena, A. (2023). Past, present, and future of anthropomorphism in hospitality & tourism: Conceptualization and systematic review. Journal of Hospitality Marketing & Management, 32(8), 1077–1125.

Kim, W., et al. (2023). What makes people feel empathy for AI chatbots? Assessing the role of competence and warmth. International Journal of Human–Computer Interaction.

Kim, W., Ryoo, Y., & Choi, Y. K. (2024). That uncanny valley of mind: When anthropomorphic AI agents disrupt personalized advertising. International Journal of Advertising. Advance online publication.

Kumar, A., et al. (2025). Anthropomorphic generative AI chatbots for enhancing customer engagement, experience and recommendation. Journal of Consumer Marketing.

Lemon, K. N., & Verhoef, P. C. (2016). Understanding customer experience throughout the customer journey. Journal of Marketing, 80(6), 69–96.

Li, B., et al. (2023). How do friendship artificial intelligence chatbots (FAIC) benefit the continuance using intention and customer engagement? Journal of Consumer Behaviour.

Li, et al. (2025). Anthropomorphism of artificial intelligence service agent and consumer responses: A systematic literature review and future research agenda. International Journal of Consumer Studies.

Liu-Thompkins, Y., Okazaki, S., & Li, H. (2022). Artificial empathy in marketing interactions: Bridging the human-AI gap in affective and social customer experience. Journal of the Academy of Marketing Science, 50(6).

Lombard, M., & Xu, K. (2021). Social responses to media technologies in the 21st century: The Media

are Social Actors paradigm. Human-Machine Communication, 2, 29–55.

Luo, X., Tong, S., Fang, Z., & Qu, Z. (2019). Frontiers: Machines vs. humans: The impact of artificial intelligence chatbot disclosure on customer purchases. Marketing Science, 38(6), 937–947.

Mariani, M. M., Hashemi, N., & Wirtz, J. (2023). Artificial intelligence empowered conversational agents: A systematic literature review and research agenda. Journal of Business Research, 161.

Mathur, M. B., & Reichling, D. B. (2015). Navigating a social world with robot partners: A quantitative cartography of the Uncanny Valley. Cognition.

McKee, K. R., Bai, X., & Fiske, S. T. (2023). Humans perceive warmth and competence in artificial intelligence. iScience, 26(8), 107256.

Nass, C., Steuer, J., & Tauber, E. R. (1994). Computers are social actors. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 72–78). ACM.

Ng, W., et al. (2024). From function to relation: Exploring the dual influences of warmth and competence on generative artificial intelligence services in the hospitality industry. Journal of Hospitality & Tourism Research.

Oliveira, G. G. de, Lizarelli, F. L., Teixeira, J. G., & Mendes, G. H. de S. (2023). Curb your enthusiasm: Examining the customer experience with Alexa and its marketing outcomes. Journal of Retailing and Consumer Services, 71.

Park, Y., et al. (2024). Impact of artificial intelligence (AI) chatbot characteristics on customer experience and customer satisfaction. Journal of Global Scholars of Marketing Science.

Peng, C.-Y., et al. (2022). The effect of required warmth on consumer acceptance of artificial intelligence in service: The moderating role of AI-human collaboration. International Journal of Information Management.

Puntoni, S., Walker Reczek, R., Giesler, M., & Botti, S. (2020). Consumers and artificial intelligence: An experiential perspective. Journal of Marketing, 84(6).

Qi, T., et al. (2025). An assistant or a friend? The role of parasocial relationship of human-computer interaction. Computers in Human Behavior.

Sestino, A., et al. (2025). Designing AI-based interactive avatars in healthcare services: The impact of human-like similarity on emotional responses. Management Decision.

Song, S., Shim, M., et al. (2022). Uncanny valley effects on chatbot trust, purchase intention, and adoption intention in the context of e-commerce: The moderating role of avatar familiarity. International Journal of Human–Computer Interaction.

Troshani, I., et al. (2020). Do we trust in AI? Role of anthropomorphism and intelligence. Journal of Computer Information Systems.

Truong, T. T. H., et al. (2025). When empathy is enhanced by human–AI interaction: An investigation of anthropomorphism and responsiveness on customer experience with AI chatbots. Asia Pacific Journal of Marketing and Logistics.

Vlačić, B., Corbo, L., Costa e Silva, S., & Dabić, M. (2021). The evolving role of artificial intelligence in marketing: A review and research agenda. Journal of Business Research, 128, 187–203.

Waytz, A., Cacioppo, J., & Epley, N. (2010). Who sees human? The stability and importance of individual differences in anthropomorphism. Perspectives on Psychological Science, 5(3), 219–232.

Wu, H., et al. (2024). Research on the uncanny valley effect in artificial intelligence news anchors. Multimedia Tools and Applications.

Xiao, L., & Kumar, V. (2021). Robotics for customer service: A useful complement or an ultimate substitute? Journal of Service Research, 24(1), 9–29.

Yang, Y., et al. (2021). Anthropomorphism and customers' willingness to use artificial intelligence service agents. Journal of Hospitality Marketing & Management, 30(2).

Yatawara, K., Sampath, T., Kalupahana, P. L., Rathnayake, S., Jayasuriya, N., & Rathnayake, N. (2025). A systematic review on consumer adoption of AI-driven chatbots.

APPENDIX: SEARCH PROTOCOL

TABLE A1

Search term families, illustrative search strings and platforms searched

Term family	Search terms	Purpose
(a) Humanisation cue anthropomorphism;	humanisation / humanization; anthropomorphic; humanlike; human-likeness; artificial empathy	Captures the design-side construct, including recent work that describes humanlike qualities of generative agents without using the term anthropomorphism
(b) AI form factor	artificial intelligence; chatbot; conversational agent; voice assistant; virtual agent; service robot; generative AI; large language model	Spans text, voice and physically embodied interfaces so that the corpus is not confined to a single modality
(c) Customer-experience outcome	customer experience; customer engagement; trust; satisfaction; loyalty; continuance intention; purchase intention	Anchors the search on customer-facing outcomes and screens out purely technical work
Illustrative combined strings	"AI anthropomorphism customer experience"; "warmth competence chatbot"; "uncanny valley service robot"; "CASA paradigm"; "artificial empathy marketing"; "chatbot disclosure purchase intention"	Terms from each family were combined iteratively, with successive queries refined in the light of earlier results
Platforms searched	Scopus-, Web of Science- and PubMed-indexed journals; Semantic Scholar; SAGE, Springer, Wiley, Taylor & Francis and Elsevier ScienceDirect; supplemented by backward snowballing	Breadth of platform coverage was prioritised over the reproducibility of a single dated database export

ISSN: 2582-5526 www.thevoiceofcreativeresearch.com

The Voice of Creative Research

Vol. 8 & Issue 4 (October 2026)

TABLE A2

Inclusion and exclusion criteria

Criterion	Included	Excluded
Publication type	Peer-reviewed journal articles, together with a small number of advance-online-publication articles from established journals	Non-peer-reviewed blogs, trade press and vendor white papers, other than a small number of authoritative practitioner sources cited only for factual, non-evaluative context
Language	English	Languages other than English
Subject matter	Anthropomorphism, humanisation or humanlike design of an AI-based system	AI treated only as a back-end or algorithmic capability with no customer-facing or human-likeness dimension
Study context	A customer, consumer, guest or service user is the subject of study	Physical or social robots studied outside any service or commercial context
Study type	Empirical (survey, experimental, qualitative and meta-analytic) as well as conceptual and review articles	—
Period	1994–2026; no lower date bound was applied, so that foundational theoretical work could be retained	—