INKUBIS: Jurnal Ekonomi dan Bisnis
Volume 8, Issue 4, 3603-3618
e_ISSN: 2775-3913
https://inkubis.polteksci.ac.id/index.php/ink/index
DOI: doi.org/10.59261/inkubis.v8i4.583

When Humanlike AI Feels Wrong: Perceived Eeriness and Trust in AI-Mediated Brand Communication — A Systematic Literature Review

Ignasius Bayu Purnomo*
Universitas Esa Unggul, Indonesia

Ferryal Abadi
Universitas Esa Unggul, Indonesia

Regina Deka Sofia
Universitas Esa Unggul, Indonesia

Primasatria Edastama
Universitas Esa Unggul, Indonesia

***Corresponding author:**
Ignasius Bayu Purnomo,
Universitas Esa Unggul, Indonesia.
✉ignbayu@student.esaunggul.ac.id

Article Info:

Article history:
Received: October 16, 2026
Revised: November 16, 2026
Accepted: November 17, 2026
Available Online: November 21, 2026

Keywords:
Anthropomorphism;
Brand Communication;
Consumer Trust;
Perceived Eeriness;
Uncanny Valley

This is an open access article under the CC BY-SA license.

Copyright © 2026 by Author. Published by Politeknik Siber Cerdika International

Abstract

Background: This review examines when humanlike artificial intelligence (AI) becomes unsettling and how uncanny appraisal relates to consumer trust in AI-mediated brand communication. It integrates the trust-building logic of anthropomorphism with evidence of unfavorable trust-related evaluations under uncanny conditions.
Methods: Following PRISMA 2020 and SPAR-4-SLR principles, a Scopus search identified 7,665 records. After duplicate, document-type, and title/abstract screening, 18 journal articles were retained, verified in full text, and coded by AI form, appraisal mechanism, trust object, boundary condition, outcome, and evidence directness. Additional theoretical sources were used for conceptual anchoring but were not added to the systematic evidence corpus.
Results: Five studies provide direct uncanny-appraisal evidence, six proximal-appraisal evidence, and seven contextual anthropomorphism–trust evidence. Direct studies associate unfavorable uncanny appraisal with lower trust-related evaluation, using distinct measures of affect, trust, and source trustworthiness. Human-likeness is not inherently harmful: social cues can support trust when aligned with competence, empathy, usefulness, transparency, and task demands.
Conclusion: The review develops the Eeriness-Trust Calibration Framework (ETCF), separating design cues from consumer appraisals, positive and negative trust routes, and multiple trust objects. It reframes AI humanization as a cue-capability-context calibration problem and identifies cross-object trust transfer to organizations and brands as a key unresolved marketing mechanism. Brands should optimize coherence rather than maximize realism, pretest eeriness and authenticity together, and monitor trust separately at agent, content, provider, and brand levels.

To cite this article: Purnomo, I. B., Abadi, F., Sofia, R. D., & Edastama, P. (2026). When Humanlike AI Feels Wrong: Perceived Eeriness and Trust in AI-Mediated Brand Communication — A Systematic Literature Review. *INKUBIS: Jurnal Ekonomi dan Bisnis*, *8*(4), 3603-3618. https://doi.org/10.59261/inkubis.v8i4.583

INTRODUCTION

Artificial intelligence represents organizations through chatbots, virtual humans, service robots, and generated communication. Empirical findings show why this creates a marketing problem: shopping-chatbot eeriness is associated with lower agent trust and weaker commercial intentions, while humanlike service robots can improve responses under complex tasks (Hameed et al., 2026; So et al., 2024). These results concern consumer evaluations rather than population-wide adoption rates. Marketing research therefore treats AI as an experiential and communicative actor (Davenport et al., 2020; Huang & Rust, 2021; Puntoni et al., 2021). Humanlike cues raise

expectations; eeriness can emerge when observed behavior or informational performance fails to meet them. The shopping-chatbot study involved 429 users in Malaysia Hameed et al. (2026), illustrating observed consumer–AI interaction rather than population-wide adoption.

Anthropomorphism does not have a single directional effect. Humanlike cues may increase social presence, warmth, and closeness Araujo (2018), Go and Sundar (2019), while service research identifies potential relationship (Liu et al., 2026; Mende et al., 2019; van Doorn et al., 2017). Uncanny-valley theory instead predicts unfavorable responses to almost-human entities with detectable inconsistencies (Mori et al., 2012). In brand communication, an unsettling spokesperson can weaken source trustworthiness, and anomalous generated content can undermine trust in the communication (Liu et al., 2026; Xiong et al., 2026). Transfer of these evaluations to the represented brand is a separate, untested possibility, not an automatic consequence.

The apparent contradiction reflects differences in context, AI form, trust object, and mechanism. E-commerce chatbot studies connect eeriness with agent trust and purchase or reuse intentions Hameed et al. (2026), Song and Shin (2024), whereas service-robot research emphasizes usefulness, ease of use, and task complexity (So et al., 2024). Celebrity avatars involve source-credibility dimensions Xiong et al. (2026), and generated videos involve trust in communication (Liu et al., 2026). Authenticity and empathy studies examine different appraisals again (Khan et al., 2025). Human-likeness denotes resemblance, humanlike is its adjectival form, and anthropomorphism denotes attributing human characteristics to AI. Perceived eeriness is an unsettling appraisal, not another label for these cues.

Trust denotes willingness to rely on a specified target under uncertainty. Distrust denotes negative expectations about that target and is not assumed to be merely low trust. Trustworthiness concerns evaluated qualities of the target; credibility concerns the believability of a source or message. Purchase intention, reuse, disclosure, and e-loyalty are behavioral outcomes, not measures of trust. Consumers may evaluate the AI agent, generated content, virtual source, provider, organization, or brand separately. Research on trust in AI and automation emphasizes performance, context, and calibrated reliance (Glikson & Woolley, 2020; Hancock et al., 2011; Hoff & Bashir, 2015). Accordingly, distrust of a chatbot need not imply distrust of the brand; cross-object transfer requires separate evidence.

Three distinct gaps motivate the questions. The empirical gap is the limited direct evidence linking uncanny appraisal to a specified trust outcome; this motivates RQ2 and the boundary-condition analysis in RQ4. The conceptual gap concerns distinctions among design cues, appraisals, trust constructs, and trust objects, motivating RQ1, RQ3, and RQ5. The methodological gap concerns synthesis across heterogeneous technologies, designs, and measures: evidence directness and object-specific coding are needed to interpret RQ2–RQ5 without treating adjacent constructs or behavioral intentions as interchangeable effect estimates. These gaps concern evidence coverage, conceptual explanation, and integration, respectively; they are not three claims that the same mechanism has been overlooked.

Accordingly, this review addresses five questions: (RQ1) how human-likeness, anthropomorphism, perceived eeriness, and closely related appraisals are conceptualized in consumer-facing AI communication; (RQ2) what direct evidence links eerie or uncanny appraisal to consumer trust, trustworthiness, credibility, and trust-related behavioral outcomes; (RQ3) through which positive and negative psychological or functional mechanisms humanlike AI can strengthen or weaken trust; (RQ4) which boundary conditions explain when the social-functional route or the uncanny-incongruity route is more likely to dominate; and (RQ5) how trust is distributed across the AI agent, generated content, virtual source, service provider, organization, and brand-related relationship, and what the evidence implies about cross-object trust transfer.

In answering these questions, the review makes four contributions. First, it separates humanlike design cues from perceived eeriness and introduces an evidence-directness hierarchy, so that direct uncanny effects are not conflated with adjacent constructs such as skepticism, artificiality, emotional detachment, authenticity, or generic discomfort. Second, it reconciles positive and negative anthropomorphism findings through a dual-route calibration logic, in which the same humanlike cue can support social-functional fit or generate uncanny incongruity depending on whether it is coherent with capability, task, relationship, and culture. Third, it disaggregates trust by object, treating AI-agent trust, content credibility, source trustworthiness,

provider trust, organizational trust, and brand-related relationship judgments as related but non-equivalent. Fourth, it integrates these elements into the Eeriness-Trust Calibration Framework (ETCF), which reframes the managerial and theoretical problem from maximizing anthropomorphism to calibrating social cues to what an AI can actually perceive, explain, remember, generate, and do.

METHOD

PRISMA 2020 is used to report identification, screening, verification, and inclusion Page et al. (2021) Figure 1 reconciles the recorded counts. SPAR-4-SLR organizes the procedure into assembling, arranging, and assessing (Paul et al., 2021). Assembling covers the Scopus source, search date, query, publication window, and eligibility criteria below. Arranging covers duplicate removal, document-type and relevance screening, and coding by AI form, appraisal, trust object, context, and outcome. Assessing covers evidence-directness classification, design-sensitive interpretation, thematic synthesis, and limitations. These functions are complementary: PRISMA is a reporting guideline, not evidence that screening was independently duplicated. Heterogeneous constructs and designs preclude a justified pooled effect estimate.

The manuscript reports a fixed corpus of 18 journal articles, with full-text verification and coding after the documented screening stages. Additional theoretical and contextual sources support interpretation but do not enter the PRISMA counts, evidence tiers, or corpus statistics. Eligibility requires the publication, consumer-facing AI, construct-relevance, empirical-evidence, and full-text criteria in Table 1; topical relevance alone does not make an adjacent appraisal a direct eeriness measure.

The review followed the documented Scopus search and screening procedure, resulting in a final corpus of 18 journal articles. Eligibility was assessed against the publication, consumer-facing AI, construct-relevance, empirical-evidence, and full-text criteria presented in Table 1. However, the available review documentation does not identify the individuals responsible for screening, establish whether eligibility decisions were independently assessed, or record how borderline cases were resolved. No preregistration or inter-rater agreement statistic was documented. Consequently, although the reported screening counts are internally consistent, the reproducibility and consistency of individual eligibility decisions cannot be independently verified. These limitations are acknowledged without reconstructing undocumented procedures or introducing unsupported methodological claims.

The Scopus search was conducted on 4 July 2026 using the TITLE-ABS-KEY field. The documented search string was: TITLE-ABS-KEY((eeriness OR eerie OR uncanny OR "uncanny valley" OR creepy OR weird OR unsettling OR discomfort OR humanlike OR anthropomorph*) AND (trust OR trustworthiness OR credibility OR confidence OR reliance OR "brand trust" OR "consumer trust" OR "customer trust")). The review window covered 2020 through the search date in 2026. Scopus was the sole bibliographic database in the documented protocol. The conceptual scope includes consumer-facing GenAI and adjacent anthropomorphic AI interfaces because chatbots, avatars, digital humans, and service robots all communicate on behalf of organizations and illuminate mechanisms relevant to contemporary AI-mediated brand communication.

Table 1
Eligibility criteria and scope calibration

Criterion	Included	Excluded / contextual treatment
Publication	English-language peer-reviewed journal articles published 2020-2026.	Books, chapters, conference items, editorials, errata, letters, notes, retracted items, reviews, and short surveys.
Consumer-facing AI	AI-generated or AI-mediated communication in marketing, e-commerce, hospitality/tourism, service, advertising, or related consumer interaction.	Industrial or technical AI without consumer-facing communication or trust implications.

Criterion	Included	Excluded / contextual treatment
Construct relevance	Human-likeness, anthropomorphism, eeriness/uncanny, authenticity/artificiality, or closely related cue together with trust, trustworthiness, credibility, reliance, or a trust-related outcome.	No defensible connection between AI social/humanlike cues and trust-related evaluation.
Evidence	Empirical evidence capable of informing the review questions. Direct, proximal, and contextual evidence retained with different interpretive weight.	Purely technical benchmark evidence or evidence unrelated to human evaluation.
Full-text verification	Full text available and sufficient to code design, constructs, trust object, and findings.	No additional numerical exclusion was applied after the 18 final records were fixed in the documented screening flow.

The identification stage yielded 7,665 Scopus records. Removal of 65 manually identified duplicates left 7,600 records for screening. Document-type screening excluded 1,920, leaving 5,680 for title/abstract assessment; exclusion of 5,662 irrelevant records left 18 meeting the review criteria. Full texts of all 18 were verified and coded, with no additional exclusions; all 18 entered the qualitative synthesis. Figure 1 reports the same counts and stage labels: identification, screening, verification, and inclusion. The arithmetic reconciles at every transition (7,665 − 65 = 7,600; 7,600 − 1,920 = 5,680; 5,680 − 5,662 = 18). Separate retrieval-failure or full-text exclusion counts were not recorded and are not inferred.

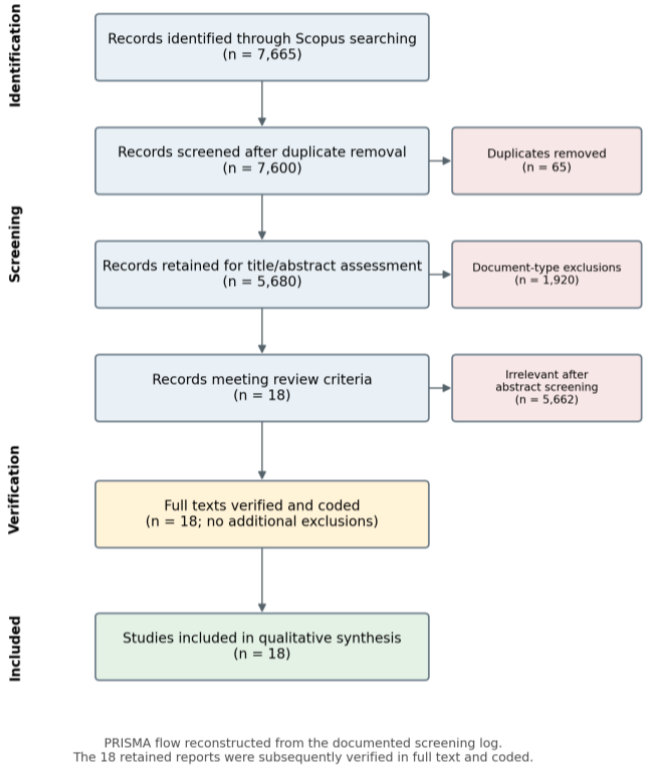


Figure 1. PRISMA flow of study identification, screening, full-text verification, and inclusion.

Data extraction was organized around publication year, AI technology, consumer context, research design, sample characteristics, humanlike cues, consumer appraisals, trust objects, psychological mechanisms, boundary conditions, and reported outcomes. Evidence was classified into three levels according to its relevance to the review questions: direct uncanny-appraisal evidence (Tier 1), proximal-appraisal evidence (Tier 2), and contextual anthropomorphism–trust

evidence (Tier 3). The reported distribution comprises five, six, and seven studies, respectively. Because a complete study-level extraction matrix and methodological apprais al record are not available in the manuscript, the allocation of individual studies and the consistency of the classifications cannot be independently verified from the aggregate results. Accordingly, the tier distribution is reported descriptively, and the synthesis avoids treating adjacent appraisals or heterogeneous trust measures as directly comparable evidence.

Table 2
Evidence-directness framework

Tier	Operational criterion	Interpretive role	n
Tier 1: Direct eeriness-trust	Measures eerie/uncanny affect or an explicit uncanny-valley affective response and relates it to trust or a specified source-credibility dimension.	Primary evidence for uncanny appraisal and trust-related evaluation; favorability and trustworthiness remain separately labelled.	5
Tier 2: Proximal appraisal	Examines authenticity/artificiality, skepticism, emotional detachment, or qualitative uncanny accounts alongside trust; these are not interchangeable eeriness measures.	Mechanism-building evidence; not treated as a direct eeriness effect.	6
Tier 3: Contextual anthropomorphism-trust	Links human-likeness, anthropomorphism, social cues, competence, transparency, or related cues to trust without measuring eeriness as the focal mechanism.	Boundary-condition and alternative-pathway evidence.	7

RESULTS AND DISCUSSION

Results

Publication, technology, and design profile

The corpus contains 18 journal articles, predominantly published from 2024 onward (15 studies; 83.3%). Table 3 provides the annual distribution and the complete technology and design profile. Chatbots and GenAI conversational agents are the largest category (10 studies; 55.6%). Direct uncanny evidence nevertheless spans several modalities, including conversational embodiment, humanoid services, livestream presenters, and AI-generated video.

Table 3
Descriptive profile of the 18-study evidence corpus

Dimension	Category	n	Share
Publication year	2020 / 2021 / 2022 / 2023 / 2024 / 2025 / 2026	1 / 1 / 0 / 1 / 6 / 4 / 5	5.6 / 5.6 / 0 / 5.6 / 33.3 / 22.2 / 27.8%
Primary AI form	Chatbot / GenAI conversational agent	10	55.6%
Primary AI form	Digital human / virtual human / AI avatar	3	16.7%
Primary AI form	Service robot / humanoid	2	11.1%
Primary AI form	AI-generated message/video communication	2	11.1%
Primary AI form	Anthropomorphized AI/voice assistant	1	5.6%
Primary design	Experimental / multi-study experimental	9	50.0%

Dimension	Category	n	Share
Primary design	Quantitative survey / SEM or regression	5	27.8%
Primary design	Mixed methods	3	16.7%
Primary design	Qualitative phenomenology	1	5.6%
Evidence directness	Tier 1 / Tier 2 / Tier 3	5 / 6 / 7	27.8 / 33.3 / 38.9%

Trust objects represented in the corpus

The reported objects and constructs differ across studies. Agent or technology trust appears in Gu et al. (2024), Gupta and Mukherjee (2026), Hameed et al. (2026), Jung et al. (2021), Pillai and Sivathanu (2020), Song and Shin (2024), Gu focuses on sustained trust after failure. Pizzi et al. (2023) concern website/provider trust, Jia et al. (2024) organizational trust, and Pandey and Rai (2024) cognitive trust in an assistant. Khan et al. (2025) connect authenticity to trust and relationship evaluations, and Shahzad et al. (2026) distinguish trust from distrust beliefs.

Source evaluations in Lee and Kim (2024), Seymour et al. (2025), Xiong et al. (2026) concern trustworthiness or credibility; the last also separates expertise and attractiveness. Daher (2025), Liu et al. (2026) concern generated communication, So et al. (2024) robot trust and receptivity, and Jin and Han (2025) qualitative tourism-information experiences. Purchase intention, adoption, reuse, disclosure, and e-loyalty remain downstream outcomes. These reported construct labels are not a harmonized measurement scale; the supplied records do not reproduce every study's item or dimension.

Thematic synthesis

Direct evidence: eeriness is a trust-degrading appraisal

Across the five Tier 1 studies, eerie or uncanny appraisal is negatively related to trust-related evaluation, although the operationalizations differ. Song and Shin (2024) report higher eeriness for a hyperrealistic animated chatbot than for a cartoonish-still agent, with lower chatbot trust and downstream purchase and reuse intentions. Hameed et al. (2026) report a positive human-likeness–eeriness path (β = .131) and a negative eeriness–trust path (β = −.119). Familiarity attenuates the uncanny pathway in both chatbot studies.

Liu et al. (2026) Table 6) report PLS-SEM paths from uncanny-valley eeriness to perceived trust (β = −.084) and from perceived realism to the same outcome (β = .856). Both come from the same structural model and viewing context. The article labels them β coefficients but does not explicitly identify their standardization status in Table 6. The review therefore retains them as reported paths without treating their ratio as a comparable effect-size statistic. Their opposite signs support distinct affective and cognitive associations; they do not establish that realism universally dominates eeriness.

Xiong et al. (2026) link celebrity-avatar eeriness to source-credibility dimensions and purchase intention. Jung et al. (2021) link uncanny-valley favorability to robot trust across service-expertise contexts. Favorability is an affective operationalization, and source trustworthiness is a dimension of credibility; neither is relabelled as an identical eeriness or trust scale.

Across five Tier 1 articles, the reported directions are consistent with lower trust-related evaluation under unfavorable uncanny appraisal. This limited and heterogeneous evidence supports a conditional synthesis, not a universal causal effect or pooled magnitude. Human-likeness itself need not reduce trust, and behavioral intentions do not measure trust.

Table 4

Direct Tier 1 evidence for the eeriness-trust relationship

Study	Direct stimulus/appraisal	Trust-related evaluation	Boundary/downstream evidence
(Hameed et al., 2026;	Hyperrealistic animated chatbot -> higher eeriness	Eeriness -> lower chatbot trust	Avatar familiarity weakens eeriness; trust

Study	Direct stimulus/appraisal	Trust-related evaluation	Boundary/downstream evidence
Song & Shin, 2024),	Humanlikeness -> higher eeriness (beta=.131)	Eeriness -> lower trust (beta=-.119)	mediates purchase and reuse. Familiarity attenuates human-likeness -> eeriness; trust predicts purchase, reuse, e-loyalty.
(Liu et al., 2026)	Hallucination intensity -> higher eeriness and lower realism	Reported PLS-SEM paths: eeriness → trust β = −.084; realism → trust β = .856; standardization not explicit in Table 6.	Trust → behavioral intention β = .794; within-model associations only, not a general ranking of predictors.
(Xiong et al., 2026)	Celebrity AI avatar -> higher perceived eeriness	Eeriness → lower source trustworthiness, attractiveness, and expertise; distinct credibility dimensions.	Source-credibility dimensions transmit effects to purchase intention.
(Jung et al., 2021)	Human-likeness produced a U-shaped UCV favorability response	UCV favorability → robot trust β = .557/.559; favorability is not an identical eeriness scale.	Expertise context changes affective and trust evaluations.

Proximal appraisals: artificiality, skepticism, authenticity, and detachment

Tier 2 evidence concerns related mechanisms, not direct eeriness–trust effects. Skepticism evaluates a source or message, artificiality concerns perceived machine-like qualities, authenticity concerns perceived genuineness, and emotional detachment concerns relational distance. General discomfort is classified as uncanny evidence only when the original operationalization explicitly supports that interpretation. Pizzi et al. (2023) connect warmth and competence to skepticism and website/provider trust; this does not make skepticism a measure of eeriness.

Authenticity is a separate positive appraisal. Khan et al. (2025) relate humanness, empathy, warmth, and humor to authenticity and then to trust, rapport, and satisfaction. Pandey and Rai (2024) place authenticity upstream of cognitive trust. These studies help explain how humanlike interaction may support trust, but neither result alone demonstrates an inverse eeriness effect.

Shahzad et al. (2026) distinguish trust-building beliefs from distrust-related beliefs, including privacy concern and emotional detachment. Jin and Han (2025); Seymour et al. (2025) report ambivalent evaluations of digital humans or tourism-information interaction. Qualitative descriptions and adjacent appraisals support mechanism development; they are not counted as quantitative evidence that measured eeriness reduces trust.

Positive social-functional route: when anthropomorphism builds trust

Tier 3 evidence explains why human-likeness is not inherently negative. Pillai & Sivathanu, (2020) connect anthropomorphism, intelligence, usefulness, ease of use, and trust to chatbot adoption. Gupta and Mukherjee (2026) identify transparency, information quality, anthropomorphism, and empathy as antecedents of trust in GenAI product/service search. So et al., (2024) show that functional perceptions mediate service-robot responses, with benefits from anthropomorphism under complex tasks. These findings support a social-functional pathway grounded in competent and intelligible interaction.

Relational needs and service failures create different conditions for this pathway. Jia et al. (2024) link anthropomorphic chatbot interaction after social exclusion to belonging, control, and organizational trust. Gu et al. (2024) relate anthropomorphism, empathy, and interaction quality to sustained trust through failure attribution, with AI anxiety as a boundary condition.

Communication credibility is another distinct outcome. Daher (2025) reports that factual AI communication can be credible despite lower perceived warmth, while anthropomorphism within AI conditions can improve warmth and trust-related judgments. Contextual HRI research also identifies competence and warmth as trust resources Christoforakos et al. (2021) and examines cognition- and affect-based trust in robot communication (Chien et al., 2024).

These pathways do not guarantee behavioral change. Ho et al. (2026) used as contextual evidence, report increased empathy and animacy after repeated friendlike interaction without a change in trust. Choi et al. (2023) likewise show that trust need not dominate ChatGPT adoption. The review therefore distinguishes social perceptions, trust judgments, and behavioral outcomes.

Boundary conditions explain apparently contradictory effects

Familiarity moderates the uncanny pathway in (Hameed et al., 2026; Song & Shin, 2024). Task demands also alter the value of humanlike cues: Jung et al. (2021) compare service-expertise contexts, while So et al. (2024) identify stronger benefits under complex tasks. These findings concern different moderators and should not be combined into a universal effect of human-likeness.

Other conditions operate through distinct mechanisms. Service failure increases the relevance of attribution and empathy Gu et al. (2024) social exclusion increases the value of connection (Jia et al., 2024). Transparency and information quality provide evidence of reliability Gupta and Mukherjee (2026), Shahzad et al. (2026), while visual anomalies alter realism and eeriness in generated communication (Liu et al., 2026). Table 5 distinguishes these conditions.

Table 5
Boundary conditions that shape trust-route dominance

Boundary condition	How it shifts calibration	Illustrative evidence
Familiarity	Normalizes humanlike cues and reduces category uncertainty, weakening the uncanny route.	(Hameed et al., 2026; Song & Shin, 2024)
Task complexity / expertise	Increases weight placed on ability and task fit relative to appearance.	(So et al., 2024)
Social-connection need	Raises the value of relational cues, making anthropomorphic communication more beneficial.	(Jia et al., 2024; Jung et al., 2021)
Service failure / AI anxiety	Raises accountability, empathy, and attribution demands; anxiety makes sustained trust more fragile.	(Gu et al., 2024)
Transparency / information quality	Provides cognitive evidence of reliability that can strengthen trust independently of appearance.	(Gupta & Mukherjee, 2026; Shahzad et al., 2026)
Perceived realism / hallucination	Visual-semantic anomalies increase eeriness and reduce cognitive credibility.	(Liu et al., 2026)
Technology readiness / expectations	Individual readiness shapes effort and performance expectations that feed trust.	(Fu et al., 2024)
Cultural appropriateness	Multimodal responses that violate local norms can create discomfort and undermine trust in global applications.	(Qiu et al., 2025) - contextual source

Trust objects and the unresolved question of spillover

The studies collectively form a network of trust objects. The most common target is the AI agent, chatbot, technology, or robot; other studies target generated information, source trustworthiness, the hosting provider, organizational trust, or cognitive trust in an AI assistant. This heterogeneity should not be averaged away. Pizzi et al. (2023) show that chatbot-related appraisals can be linked to trust in the service provider, Jia et al. (2024) examine organizational

trust, and Hameed et al. (2026) connect chatbot trust to e-loyalty. Yet the corpus does not establish a complete causal sequence from eeriness at the AI source to distrust of content and then distrust of the brand.

Cross-object transfer is therefore treated as a proposition rather than a finding. Attribution is likely to determine whether distrust remains localized or spreads. When consumers interpret an AI failure as the organization's responsibility, provider or brand trust may be affected; when the failure is seen as a bounded tool error, the brand relationship may remain intact. This distinction extends source-credibility logic into brand management and is especially important as AI spokespersons and generated content become embedded in owned customer touchpoints.

Eeriness-Trust Calibration Framework (ETCF)

The ETCF distinguishes perceived eeriness from related but conceptually separate consumer appraisals. The uncanny-incongruity route is primarily represented by perceived eeriness or an explicitly measured uncanny affective response. Artificiality, skepticism, emotional detachment, and perceived unreliability are treated as distinct adjacent appraisals that may inform consumer evaluations through different mechanisms. They should therefore be measured separately rather than interpreted as interchangeable indicators of eeriness. The framework further distinguishes focal trust in an AI agent, generated content, virtual source, or service provider from potential trust transfer to an organization or brand. The proposed cross-object pathway remains a theoretical proposition requiring direct empirical testing.

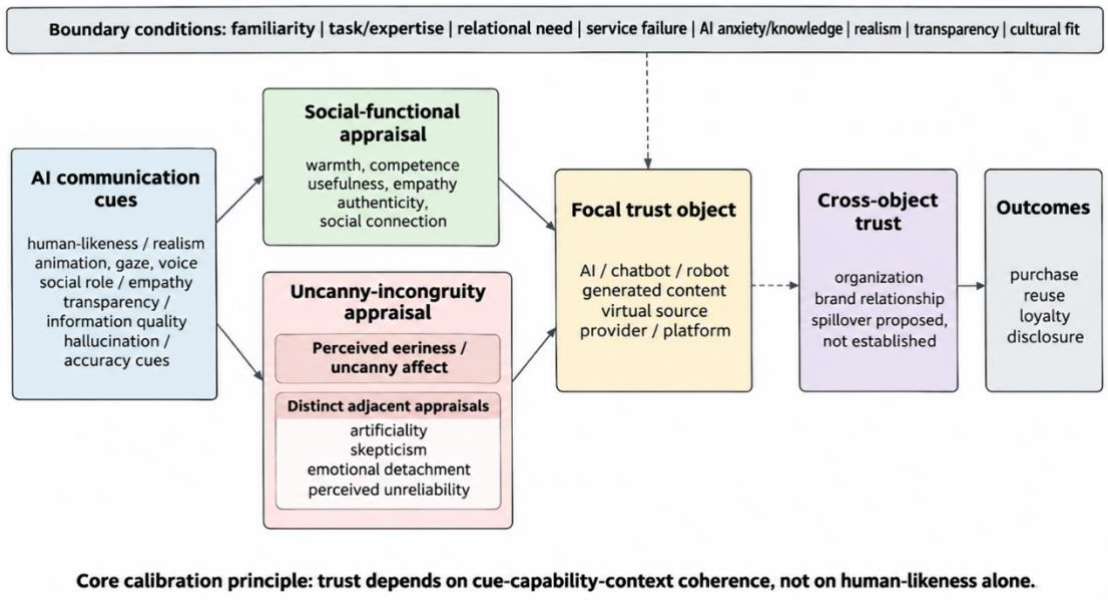


Figure 2. Eeriness-Trust Calibration Framework (ETCF). The dashed cross-object path is review-derived and requires direct empirical testing

What ETCF adds beyond existing theory

ETCF adds explanatory value in three specific ways. First, it separates cue from appraisal. Uncanny-valley research often begins with human resemblance, while anthropomorphism research often treats humanness as the focal psychological construct. ETCF instead specifies that human-likeness is an input and eeriness is an evaluative response. This prevents a positive anthropomorphism coefficient from being misread as evidence against uncanny-valley effects and prevents humanlike design from being treated as inherently eerie.

Second, ETCF models competing indirect routes. CASA, warmth-competence, and technology-acceptance perspectives explain why humanlike or socially intelligent AI can build trust; uncanny-valley, algorithm-aversion, and distrust perspectives explain why mismatch or error can undermine trust. ETCF places both processes in the same causal architecture and makes coherence the organizing variable. In this view, realism raises expectations; competence,

responsiveness, transparency, and contextual fit determine whether those expectations are satisfied.

Third, ETCF specifies the trust object and makes transfer testable. Existing research frequently moves from interface perceptions to purchase, adoption, or loyalty without establishing which entity is trusted. ETCF requires researchers to distinguish agent trust, content credibility, source trustworthiness, provider trust, organizational trust, and brand relationship. It then treats spillover as conditional on attribution rather than automatic. This multi-object logic connects HCI/HRI evidence to marketing outcomes such as brand credibility, relationship strength, loyalty, and recovery after AI failure.

Review-derived propositions

Table 6
Review-derived ETCF propositions

Code	Proposition
P1	Human-likeness affects trust primarily through appraisal rather than through a universal direct effect: the same cue can activate a positive social-functional route or a negative uncanny-incongruity route.
P2	Cue-capability incongruity may lower focal trust through perceived eeriness or through distinct adjacent appraisals such as skepticism, artificiality, detachment, and unreliability; each pathway requires its own measure.
P3	Functional and epistemic coherence, including realism, information quality, transparency, efficiency, competence, and response contingency, may support focal trust alongside negative affect; relative importance requires within-study tests.
P4	Familiarity attenuates the translation of human-likeness into eeriness by reducing novelty and category uncertainty.
P5	Task, expertise, relational need, service-failure context, and cultural fit change the relative diagnosticity of social appearance versus functional ability, thereby altering which appraisal route dominates.
P6	Distrust at the AI-agent or source level can transfer to provider, organizational, or brand-related trust when consumers attribute responsibility for the AI interaction to the deploying organization; this transfer requires direct empirical testing.

Discussion

RQ1: construct clarity requires separating cue from appraisal

The central conceptual result is that human-likeness and eeriness should not be used interchangeably. Human-likeness is a designed or perceived property that can be beneficial, neutral, or harmful. Eeriness is a negative appraisal that emerges when a source or interaction violates expectations of coherent humanness. The evidence-directness tiers operationalize this distinction: Tier 1 directly observes eerie/uncanny responses, Tier 2 captures adjacent relational appraisals, and Tier 3 shows what happens when humanlike cues support trust without an eerie mechanism. This structure prevents the review from inflating the direct evidence base.

RQ2: direct eeriness evidence is consistently negative but not uniformly dominant

RQ2 is answered by five Tier 1 articles linking unfavorable uncanny appraisal to lower trust-related evaluation. Their measures include eeriness, uncanny-valley favorability, robot/agent trust, and source trustworthiness; they do not support a single pooled estimate. Liu et al. (2026) report opposing paths within one PLS-SEM model, with standardization not explicitly stated in Table 6. No cross-study or universal dominance claim is drawn from their numerical magnitudes.

RQ3-RQ4: positive and negative mechanisms coexist because route dominance is contextual

Eight recurring mechanism families organize the fragmented evidence: perceptual/behavioral incongruity; authenticity calibration; skepticism/source evaluation; functional and epistemic quality; emotional connection/detachment; attribution/accountability

after failure; familiarity/normalization; and role/task fit. These mechanisms explain why uncanny effects should not be reduced to facial realism. An interaction can feel wrong because movement is unnatural, language is humanlike but emotionally inappropriate, a source is persuasive without sufficient transparency, generated content contains anomalies, or the interface's social role exceeds demonstrated capability. Conversely, competence, empathy, transparency, usefulness, and social connection can make anthropomorphic cues valuable. Contextual sources strengthen this interpretation: Fu et al. (2024) show that expectancy formation matters for trust, while Chien et al. (2024) and Jiang et al. (2026) indicate that relational cues are beneficial when calibrated to service performance rather than simply intensified.

RQ5: trust should be modeled as a network of objects

A major marketing implication is that researchers should specify exactly what is trusted. AI-agent trust, content credibility, source trustworthiness, provider trust, organizational trust, and brand relationship are distinct constructs with potentially different antecedents. A design that reduces trust in a virtual spokesperson may leave product quality beliefs intact; conversely, a service failure attributed to the organization may move distrust beyond the chatbot. Multi-object models can identify whether trust loss remains localized or spills into broader brand equity and relationship outcomes. The distinction also opens a service-recovery question: whether transparent human escalation can repair provider or brand trust even when trust in the AI remains low.

Implications for marketing management and research

Theoretical implications

1) Separate design cues from appraisals. Anthropomorphism, realism, and human-likeness should not be used as proxies for eeriness; future models should measure both cue intensity and consumer response.
2) Model positive and negative routes simultaneously. Warmth, competence, authenticity, usefulness, and social connection should be estimated alongside eeriness, skepticism, detachment, and distrust.
3) Specify the trust object. Measures should identify whether respondents are judging the AI, content, source, provider, organization, or brand relationship.
4) Treat trust and distrust as potentially distinct. AI can evoke confidence in competence while simultaneously producing relational or epistemic distrust.
5) Shift from anthropomorphism intensity to coherence. The theoretically relevant mismatch may be visual, behavioral, linguistic, emotional, informational, accountability-related, or cultural.
6) Treat calibration as dynamic. Familiarity and repeated exposure may normalize uncanny cues, but repeated errors can also sensitize users and accelerate trust loss.

Managerial implications: design for calibrated trust

For managers, the core principle is to optimize coherence rather than maximize human resemblance. Humanlike faces, voices, animation, names, empathy, or intimate conversational styles raise expectations. If the system cannot respond contingently, explain limitations, maintain semantic continuity, or display credible competence, those cues become liabilities. Conversely, moderate anthropomorphism can be effective when it improves usability, social connection, transparency, or task fit. Marketing teams should therefore evaluate AI experience as a trust architecture rather than as a visual-design problem.

Table 7

Trust-by-calibration principles for AI-mediated brand communication

Principle	Recommended practice	Risk addressed
Do not maximize human-likeness	Match face, voice, animation, intimacy, and social role to actual conversational, emotional, and functional capability.	Expectation violation and uncanny mismatch.

Principle	Recommended practice	Risk addressed
Pretest eeriness and authenticity together	Measure comfort/eeriness, authenticity, skepticism, competence, credibility, and willingness to rely - not engagement alone.	A visually impressive design that quietly erodes trust.
Prioritize perceptual and informational coherence	Audit animation, gaze, response contingency, hallucination, factual accuracy, and semantic continuity.	Visual/behavioral anomalies and credibility collapse.
Use transparency as cognitive trust evidence	Explain AI use, material limitations, relevant data handling, and when accountable human support is available.	Opacity, privacy concern, and ambiguous accountability.
Design for task and relationship context	Use stronger social cues when connection or complex relational service adds value; avoid unnecessary realism in simple utilitarian interactions.	Mismatch between role expectations and interface design.
Build familiarity without deception	Introduce virtual agents consistently and allow consumers to learn capabilities over repeated interactions.	Novelty and category uncertainty that amplify eeriness.
Provide human escalation after failure	Escalate high-risk advice, complaints, recovery, or emotionally sensitive interactions to accountable humans.	Trust erosion after failure and AI anxiety.
Track trust at multiple levels	Measure AI trust separately from content credibility, provider trust, organizational trust, and brand relationship metrics.	Misdiagnosing the level at which distrust is forming.

Methodological and future research agenda

The literature is stronger at estimating self-reported relationships than at observing long-term behavior. Future work should use factorial experiments that manipulate multiple levels of human-likeness together with behavioral coherence, disclosure, and human oversight; longitudinal panels that track habituation or sensitization; and field data on purchase, switching, complaint, escalation, and negative word of mouth. Trust should be measured at more than one object in the same model so that interface-level distrust can be separated from brand-level spillover.

Multimodal GenAI expands the uncanny problem beyond appearance. Face, motion, voice, language, and semantic hallucination can jointly produce mismatch. Dagtekin and Kabakus, (2025) emphasize cognitive reappraisal and epistemic suspicion in GenAI interaction, suggesting that future models should combine affective eeriness with cognitive trust calibration. Cultural context is equally important. Qiu et al. (2025) show that multimodal systems can violate culturally grounded norms even when responses are otherwise plausible; such violations can create discomfort and undermine trust. Cross-cultural research should therefore treat cultural appropriateness as a form of cue-context coherence rather than only as a demographic moderator.

Computational marketing research should be cautious about treating lexical expressions such as 'creepy', 'fake', 'robotic', 'off', or 'unnatural' as ready-made measures of eeriness. Automated classifiers require human-coded validation, discriminant validity against generic negative sentiment or anxiety, and explicit linkage to the trust object. Cross-language work also needs measurement equivalence because expressions of uncanniness and appropriate social distance can carry different meanings across cultures.

Table 8
Priority research agenda derived from ETCF

Gap	Research question	Recommended design
Nonlinear cue calibration	At what point does increasing realism or humanness stop increasing social presence and begin to increase eeriness?	Multi-level realism experiments; spline/threshold models; repeated exposure.
Cross-object trust transfer	Does eeriness reduce trust in the AI first, content first, or provider/brand first, and when does distrust spill across objects?	Multi-object SEM; process experiments; attribution manipulations.
Multimodal GenAI	How do face, motion, voice, language, and semantic hallucinations interact in producing eeriness?	Factorial multimodal experiments; objective anomaly features plus human ratings.
Disclosure and accountability	Does AI disclosure reduce ambiguity or increase artificiality salience, and does visible human oversight change the effect?	Disclosure x oversight experiments; service-failure scenarios.
Longitudinal familiarity	Does repeated exposure normalize uncanny cues or make consumers more sensitive to artificial defects?	Panel/experience-sampling studies; repeated interaction logs.
Cross-cultural calibration	Are the same levels of human-likeness, emotional expression, autonomy, and cultural appropriateness experienced as coherent across markets?	Cross-country experiments; measurement invariance; culturally grounded multimodal stimuli.
Behavioral validation	Do self-reported eeriness and distrust predict actual abandonment, switching, complaint, disclosure refusal, or purchase?	Field experiments; clickstream/transaction data; choice experiments.
Failure and recovery	How can escalation, apology source, or explanation restore provider and brand trust after an eerie or erroneous AI interaction?	Service-recovery experiments; longitudinal trust restoration.

Limitations

The review is limited to one database, Scopus, and an English-language journal corpus. Coverage and indexing can therefore affect what is found, and broad cue-related search terms yield heterogeneous evidence: only five of the 18 reported studies directly address uncanny appraisal and trust-related evaluation. The absence of identified screeners, independent dual screening, an adjudication record, and preregistration leaves selection consistency and borderline decisions unverifiable from the supplied record. Full-text checking does not remove this selection risk; no agreement statistic is inferred.

The reported literature is concentrated in service, e-commerce, tourism, and Asian-market settings. Technologies and trust measures differ, restricting cultural and technological generalization and preventing pooled-effect interpretation. The missing study-level extraction/appraisal supplement also limits verification of measurement dimensions and the reported tier allocation. The ETCF is an integrative proposal: transfer to brand trust, culturally sensitive calibration, and behavioral consequences require direct testing. Contextual references remain outside the fixed corpus, and rapid interface development may limit the durability of technology-specific findings.

CONCLUSION

This review reports 18 studies, with direct uncanny-appraisal evidence limited to five articles. For RQ1, human-likeness and anthropomorphism are cues or attributions, whereas perceived eeriness is an appraisal. For RQ2, the direct studies associate unfavorable uncanny appraisal with lower trust-related evaluation, using different operationalizations; no pooled effect or universal causal relationship is established. For RQ3, positive routes involve competence, usefulness, empathy, authenticity, and connection, while negative routes involve eerie incongruity and distinct adjacent appraisals such as skepticism or detachment. Tier 2 evidence explains

related mechanisms rather than proving direct eeriness effects.

For RQ4, familiarity, task demands, relational needs, service failure, transparency, and information quality condition the reported relationships; the evidence does not establish one universally dominant route. For RQ5, agent trust, content credibility, source trustworthiness, provider trust, and organizational trust are distinct, and transfer to brand trust remains a proposition. ETCF integrates these distinctions through cue-capability-context coherence and calibrated reliance. Managers should assess contextual fit and demonstrated capability alongside humanlike design. Research should first test competing appraisal routes and cross-object transfer, then examine cultural, temporal, and behavioral robustness.

ACKNOWLEDGEMENT

The authors gratefully acknowledge Universitas Esa Unggul for its academic support during the preparation of this systematic literature review. The authors also appreciate the constructive input received during the development and refinement of the manuscript.

AUTHOR CONTRIBUTION STATEMENT

Ignasius Bayu Purnomo contributed to the research conceptualization, development of the theoretical framework, literature collection and analysis, interpretation of findings, and preparation of the initial manuscript draft. Ferryal Abadi, Regina Deka Sofia, and Primasatria Edastama contributed to the development and refinement of the research framework and manuscript. All authors approved the final manuscript and accept responsibility for their respective contributions to the research.

REFERENCES

Araujo, T. (2018). Living up to the chatbot hype: The influence of anthropomorphic design cues and communicative agency framing on conversational agent and company perceptions. *Computers in Human Behavior*, *85*, 183–189.

Chien, S.-Y., Lin, Y.-L., & Chang, B.-F. (2024). The effects of intimacy and proactivity on trust in human-humanoid robot interaction. *Information Systems Frontiers*, *26*(1), 75–90. https://doi.org/10.1007/s10796-022-10324-y

Choi, J., Park, J., & Suh, J. (2023). Evaluating the current state of ChatGPT and its disruptive potential: An empirical study of Korean users. *Asia Pacific Journal of Information Systems*, *33*(4), 1058–1092. https://doi.org/10.14329/apjis.2023.33.4.1058

Christoforakos, L., Gallucci, A., Surmava-Grosse, T., Ullrich, D., & Diefenbach, S. (2021). Can robots earn our trust the same way humans do? A systematic exploration of competence, warmth, and anthropomorphism as determinants of trust development in HRI. *Frontiers in Robotics and AI*, *8*. https://doi.org/10.3389/frobt.2021.640444

Dagtekin, U., & Kabakus, A. K. (2025). Human–technology interaction in generative AI: A theoretical review of technology acceptance and cognitive response. *International Journal of Advanced Computer Science and Applications*, *16*(12), 194–206. https://doi.org/10.14569/IJACSA.2025.0161218

Daher, M. M. (2025). How Generative AI shapes sustainable wine consumption and the role of anthropomorphism and trust in the Lebanese market. *International Journal of Customer Relationship Marketing and Management*, *16*(1), 1–32. https://doi.org/10.4018/IJCRMM.396868

Davenport, T. H., 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.

Fu, J., Mouakket, S., & Sun, Y. (2024). Factors affecting customer readiness to trust chatbots in an online shopping context. *Journal of Global Information Management*, *32*(1), 1–22. https://doi.org/10.4018/JGIM.347503

Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, *14*(2), 627–660.

Go, E., & Sundar, S. S. (2019). Humanizing chatbots: The effects of visual, identity and conversational cues on humanness perceptions. *Computers in Human Behavior*, *97*, 304–316.

Gu, C., Zhang, Y., & Zeng, L. (2024). Exploring the mechanism of sustained consumer trust in AI

chatbots after service failures: A perspective based on attribution and CASA theories. *Humanities and Social Sciences Communications*, *11*. https://doi.org/10.1057/s41599-024-03879-5

Gupta, A. S., & Mukherjee, J. (2026). Investigating antecedents of trust in GAI technology for product/service search: A mixed-method study. *Behaviour & Information Technology*, *45*(9), 1727–1741. https://doi.org/10.1080/0144929X.2025.2587728

Hameed, I., Nadeem, M., Hameed, I., Bhatti, Z. A., & Ponnudurai, P. (2026). Social and psychological factors in AI chatbot trust-building. *International Journal of Retail & Distribution Management*, *54*(8), 833–849. https://doi.org/10.1108/IJRDM-03-2025-0158

Hancock, P. A., Billings, D. R., Schaefer, K. E., Chen, J. Y. C., de Visser, E. J., & Parasuraman, R. (2011). A meta-analysis of factors affecting trust in human-robot interaction. *Human Factors*, *53*(5), 517–527. https://doi.org/10.1177/0018720811417254

Ho, J. Q. H., Hu, M., Goh, A. Y. H., Pragasam, E. J., & Hartanto, A. (2026). How consistent friendlike conversation with AI companions influences our attitudes and perceptions toward AI: An exploratory experiment. *Behavioral Sciences*, *16*(2). https://doi.org/10.3390/bs16020278

Hoff, K. A., & Bashir, M. (2015). Trust in automation: Integrating empirical evidence on factors that influence trust. *Human Factors*, *57*(3), 407–434.

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

Jia, X., Zhao, X., Tao, W., & Ahn, S. (2024). From exclusion to connection: The role of anthropomorphic chatbots in shaping customer responses post-social exclusion. *Journal of Internet Commerce*, *23*(3), 205–232. https://doi.org/10.1080/15332861.2024.2375965

Jiang, Y., Phan Ah Kiaw, J., & Mun, Y. S. (2026). Re-evaluating emotional awareness and anthropomorphism in AI customer service: Role of social presence, and psychological distance on purchase intention. *International Review of Management and Marketing*, *16*(4), 92–101. https://doi.org/10.32479/irmm.22892

Jin, J.-H., & Han, J.-S. (2025). A phenomenological study on the experience of searching for tourism information following the emergence of ChatGPT: Focused on the uncanny valley theory. *Sustainability*, *17*(1). https://doi.org/10.3390/su17010355

Jung, Y., Cho, E., & Kim, S. (2021). Users' affective and cognitive responses to humanoid robots in different expertise service contexts. *Cyberpsychology, Behavior, and Social Networking*, *24*(5), 300–306. https://doi.org/10.1089/cyber.2020.0170

Khan, M. I., Tarofder, A. K., Gopinathan, S., & Haque, A. (2025). Designing authentic customer-chatbot interactions: A necessary condition analysis of emotional intelligence and anthropomorphic features in human-computer interaction. *International Journal of Human-Computer Interaction*. https://doi.org/10.1080/10447318.2025.2495118

Lee, D.-S., & Kim, G.-G. (2024). A study on advertising effectiveness depending on information source type (real human vs. virtual human). *Revista de Gestão Social e Ambiental*, *18*(7). https://doi.org/10.24857/rgsa.v18n7-116

Liu, J., Sun, Y., & Xiao, K. (2026). Psychological mechanisms linking AI hallucinations to user trust and behavioral intentions toward AI-generated videos: An S-O-R perspective. *Frontiers in Psychology*, *17*. https://doi.org/10.3389/fpsyg.2026.1781974

Mende, M., Scott, M. L., van Doorn, J., Grewal, D., & Shanks, I. (2019). Service robots rising: How humanoid robots influence service experiences and elicit compensatory consumer responses. *Journal of Marketing Research*, *56*(4), 535–556.

Mori, M., MacDorman, K. F., & Kageki, N. (2012). The uncanny valley. *IEEE Robotics & Automation Magazine*, *19*(2), 98–100. https://doi.org/10.1109/MRA.2012.2192811

Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., & Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, *372*. https://doi.org/10.1136/bmj.n71

Pandey, P., & Rai, A. K. (2024). Modeling consequences of brand authenticity in anthropomorphized AI-assistants: A human-robot interaction perspective. *Purushartha*, *17*(1), 116–135. https://doi.org/10.21844/16202117108

Paul, J., Lim, W. M., O'Cass, A., Hao, A. W., & Bresciani, S. (2021). Scientific procedures and rationales for systematic literature reviews (SPAR-4-SLR). *International Journal of Consumer Studies*, *45*(4), O1–O16. https://doi.org/10.1111/ijcs.12695

Pillai, R., & Sivathanu, B. (2020). Adoption of AI-based chatbots for hospitality and tourism. *International Journal of Contemporary Hospitality Management*, *32*(10), 3199–3226. https://doi.org/10.1108/IJCHM-04-2020-0259

Pizzi, G., Vannucci, V., Mazzoli, V., & Donvito, R. (2023). I, chatbot! The impact of anthropomorphism and gaze direction on willingness to disclose personal information and behavioral intentions. *Psychology & Marketing*, *40*(7), 1372–1387. https://doi.org/10.1002/mar.21813

Puntoni, S., Reczek, R. W., Giesler, M., & Botti, S. (2021). Consumers and artificial intelligence: An experiential perspective. *Journal of Marketing*, *85*(1), 131–151.

Qiu, H., Huang, K.-H., Zheng, R., Sun, J., & Peng, N. (2025). Multimodal cultural safety: Evaluation framework and alignment strategies. *Transactions on Machine Learning Research*.

Seymour, M., Yuan, L., Riemer, K., & Dennis, A. R. (2025). Less artificial, more intelligent: Understanding affinity, trustworthiness, and preference for digital humans. *Information Systems Research*, *36*(2), 1096–1128. https://doi.org/10.1287/isre.2022.0203

Shahzad, K., Khan, A. N., Ahmad, B., Hayat, K., & Chang, S. (2026). Balancing trust and distrust in generative AI chatbot adoption: A case study from China. *The Service Industries Journal*. https://doi.org/10.1080/02642069.2025.2487819

So, K. K. F., Kim, H., Liu, S. Q., Fang, X., & Wirtz, J. (2024). Service robots: The dynamic effects of anthropomorphism and functional perceptions on consumers' responses. *European Journal of Marketing*, *58*(1), 1–32. https://doi.org/10.1108/EJM-03-2022-0176

Song, S. W., & Shin, M. (2024). 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*, *40*(2), 441–456. https://doi.org/10.1080/10447318.2022.2121038

van Doorn, J., Mende, M., Noble, S. M., Hulland, J., Ostrom, A. L., Grewal, D., & Petersen, J. A. (2017). Domo arigato Mr. Roboto: Emergence of automated social presence in organizational frontlines and customers' service experiences. *Journal of Service Research*, *20*(1), 43–58.

Xiong, L., Wei, D., & Long, X. (2026). Real vs. virtual: How the uncanny valley weakens the persuasive power of celebrity AI avatar presenters - An experimental study based on live streaming e-commerce. *Journal of Theoretical and Applied Electronic Commerce Research*, *21*(5). https://doi.org/10.3390/jtaer21050141