Human Factors in Design, Engineering, and Computing, Vol. 238, 2026, 1962–1972
https://doi.org/10.54941/ahfe1008424

From Inquiry to Handoff: A Reference Architecture and Design Implications for Photorealistic Conversational AI Agents in First-Contact Services

Nick Tugarin and Christian van Husen

Institute for Product and Service Engineering, Furtwangen, 78120, Germany

ABSTRACT

In knowledge-intensive consulting services, the first contact with potential clients helps build mutual understanding and prepares for the next steps. However, this process is often managed through static forms or telephone hotlines that provide limited guidance and lead to repetitive intake tasks. Although conversational AI is established in customer service, its use as an end-to-end first-contact service remains underexplored. This paper presents the design, implementation, and evaluation of a photorealistic conversational AI agent in the real-world consulting environment of a premium kitchen manufacturer. The agent elicits customer requirements in dialogue, generates a structured summary, schedules the consultation, and transfers the result to a human consultant. Following a design science research approach, a reference architecture realizes the service workflow, and the artifact is evaluated with 19 prospective customers through a standardized questionnaire and semi-structured interviews, complemented by interviews with the organization. Results indicate high trust, perceived transparency, and data comfort; a majority would prefer starting with the agent over waiting for a human, whereas social presence and embodiment polarize across interaction styles. The findings are synthesized into five design principles.

Keywords: AI avatar, Conversational agent, Service, AI agents, Human-AI interaction

INTRODUCTION

In knowledge-intensive consulting services, where firms deliver highly customized solutions, the initial contact with prospective clients is decisive: it establishes mutual understanding and prepares the subsequent consultation. Yet this stage is often still handled through static forms or telephone hotlines that offer little guidance and generate repetitive intake work. Conversational agents are an established customer service channel (Gnewuch et al., 2017), and embodied, avatar-based assistants are advancing rapidly. However, their adoption remains concentrated in education, healthcare, and cultural heritage, and many prototypes rely on scripted, menu-driven interactions that constrain user autonomy (Halim et al., 2026). This body of work does not cover knowledge-intensive professional services, where an agent must guide a prospective client end-to-end from initial inquiry to consultation scheduling. The initial stage of consulting—characterized by high customization, minimal

Received August 21, 2026; Revised September 24, 2026; Accepted October 2, 2026; Available online December 1, 2026

© 2026 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

standardization, and reliance on mutual understanding prior to co-creating any value—thus represents a previously unexamined context for embodied conversational agents.

The case scenario originates from a manufacturer of high-quality, personalized kitchens. Most prospective customers asked to speak directly with the CEO, who could not accommodate every request; a service employee was therefore assigned to ask the same standardized questions over a telephone hotline and to forward each inquiry to the CEO through handwritten notes. This process created little value, dissatisfied customers, was expensive, and was misaligned with the employee's skill set.

The study addresses two research questions. RQ1: How can a photorealistic conversational AI agent be designed and architecturally realized to deliver an end-to-end first-contact service in a knowledge-intensive consulting environment? RQ2: How do prospective customers and the service organization evaluate such an agent along the five dimensions of human-AI service interaction? The paper contributes (1) a working artifact and reference architecture for an embodied first-contact service, (2) empirically grounded design principles transferable to comparable professional service contexts, and (3) a construct-based, replicable evaluation approach responding to the inconsistent evaluation practices across avatar studies (Halim et al., 2026).

RELATED WORK

Embodied conversational agents combine natural language interaction with a visually represented body conveying verbal and non-verbal cues (Cassell et al., 2000). People mindlessly apply social rules and persona attributions to computers exhibiting human-like cues (Nass and Moon, 2000), and design knowledge on such cues has been systematized for conversational agents (Diederich et al., 2022). Higher visual realism is associated with trust, empathy, and social presence, but also with uncanny-valley risks; a recent systematic review concludes that AI-powered avatar assistants remain dominated by education, healthcare, and cultural-heritage applications, frequently rely on constrained dialogue designs, and lack standardized evaluation benchmarks (Halim et al., 2026).

Service research conceptualizes AI as fundamentally transforming frontline encounters, in which technology infusion generates automated social presence with consequences for experience and disclosure (van Doorn et al., 2017), and customer reactions to algorithmic providers are task-dependent (Castelo et al., 2019). The framework of Tugarin and van Husen (2025) integrates these strands into a holistic model along five dimensions: (1) AI adoption—factors and features influencing the willingness to use; (2) human-AI dynamics—how interaction, collaboration, and role sharing occur between humans and AI; (3) technologies—enabling a user-centered and enhanced experience; (4) AI implementation—organizational factors including change management and risks; and (5) AI ethics—safeguarding human well-being throughout every process step. Unlike individual-level

adoption models or interaction-level design reviews (Diederich et al., 2022), it spans both customer and organizational perspectives across the service lifecycle; it therefore serves as both a design lens and the structure of the evaluation.

Linking the gaps identified in the avatar literature to the framework dimensions positions this study: trust and social presence have mainly been examined in education, collaboration, and healthcare, not at the front end of professional services; task allocation, role clarity, and the handover between an embodied agent and a human consultant within one service journey remain unexamined; virtual-reality toolchains and game engines dominate, whereas lightweight, web-based agentic architectures are rarely described; most prototypes lack integration with real organizational processes; and most studies show limited ethical readiness regarding disclosure, data protection, and accountability (Halim et al., 2026).

METHODOLOGY

The research followed Peffers et al.'s (2007) design science research methodology. Problem identification and motivation established the need to change how customer requests are handled; solution objectives specified the theoretical contribution and organizational goals; and design and development produced the AI-avatar-based first-contact service. A proof of concept supported demonstration and iterative refinement. Evaluation followed a naturalistic ex-post approach with prospective customers and the participating organization (Venable et al., 2016; Tugarin and van Husen, 2026).

Participants (n = 19) were recruited with the CEO from existing customers and persons interested in purchasing a kitchen and are referred to as P1–P19. Sixteen participants reported basic or no knowledge of generative AI, two reported advanced knowledge, and one reported expert knowledge; all had previously used a customer service chatbot without an avatar. Technical affinity was mixed: ten participants reported high, five moderate, and four low affinity.

A multi-method design combined a quantitative questionnaire with semi-structured interviews and session observations. Participants interacted individually with the deployed system and completed 17 seven-point Likert items and one forced-choice item, followed by 30–45-minute interviews. Social presence was measured with three items adapted from Gefen and Straub (2003, 2004); initial trust was measured with six items covering competence, benevolence, and integrity, adapted from McKnight et al. (2002). Supplementary descriptive anchors addressed embodiment preference, human-AI dynamics, and ethics, with the interviews providing greater depth. On the organizational side, we conducted 60–95-minute interviews with the CEO and two service employees. Questionnaire data were analyzed descriptively; interview and observation notes were analyzed thematically (Braun and Clarke, 2006) and mapped onto the five framework dimensions (Table 1).

Table 1: Primary evidence and analyzed constructs across the five framework dimensions.

Dimension	Primary Evidence	Analyzed Constructs
AI adoption	Customer	Trust, social presence, interaction quality, anthropomorphism effect
Human-AI dynamics	Customer, organization	AI vs. human preference; task allocation, role clarity
Technologies	Customer, artifact	Embodiment preference; system architecture
AI implementation	Organization	Usefulness, risks; role changes; knowledge integration
AI ethics	Customer, organization	Transparency, privacy, data security; accountability

SERVICE WORKFLOW

The first-contact service was designed using service blueprinting (Bitner et al., 2008) and simplified into a process-focused model. From the company homepage, a prospective customer seeking a planning appointment can either visit the open showroom or start a digital consultation with the AI avatar in voice or text mode. The agent introduces itself and its role—organizing the CEO's consultation—and gathers the necessary information; the customer may ask questions at any time. All extracted information stays visible on screen and can be edited before sending, giving the customer control over the transmitted data. The encounter ends with the customer booking a slot on the CEO's synchronized calendar. The system then summarizes the request into a management overview with open questions and suggested focus areas, storing it in the cloud. The CEO is notified to prepare for the consultation. This change shifts the interaction to the agent, eliminating repetitive phone intake from the forefront while keeping the human element in the consultation.

ARTIFACT ARCHITECTURE

Component models of avatar-based conversational agents distinguish the embodied avatar—its verbal, non-verbal, and appearance-related elements—from the underlying conversational system of interface, APIs, storage, and language processing (Rashik et al., 2024). We adopt this separation as a structuring lens but extend it toward a process-oriented architecture representing the first-contact workflow: requirement elicitation, inquiry summarization, consultation scheduling, and handover to a human consultant. Architecturally, the artifact is a three-layer web application (Figure 1).

The user interface layer is a Next.js frontend handling session setup, message exchange, guided inputs, the continuously visible request panel with review before submission, and the avatar panel. The orchestration layer is a FastAPI backend exposing a lean REST interface over session state. Each turn traverses a deterministic pipeline rather than an open-ended agent loop—response generation, extraction, speech synthesis, and guided-input

selection composed in code, with extraction and synthesis in parallel—while a CrewAI flow is reserved for the closing summary. A central design decision separates conversational response generation (GPT-5.6 Terra) from structured extraction, decoupling answer quality from data-capture reliability. Speech is transcribed and synthesized with OpenAI models and streamed as paced audio chunks for lip-synchronous playback to the photorealistic avatar, rendered by HeyGen LiveAvatar over a direct WebRTC connection, with the backend acting only as a token broker. The persona ("Clara") is defined in a Markdown system prompt specifying behavior, the six information fields, and service requirements; confirmed inquiries are persisted as structured Markdown through a storage abstraction backed by Supabase.

Two rationales from the avatar literature guided the interaction design: combining voice with text input accommodates differing user preferences, and realistic avatars suit applications requiring trust, empathy, and social presence (Halim et al., 2026)—precisely the constructs selected for the adoption dimension of the evaluation. In voice mode, silence detection governs turn-taking; prolonged silence triggers re-engagement. Latency is further addressed by caching and pre-synthesizing recurring utterances and computing progress deterministically in code. Cal.com, synchronized with the CEO's calendar, is callable from the slot picker and as an agent tool.

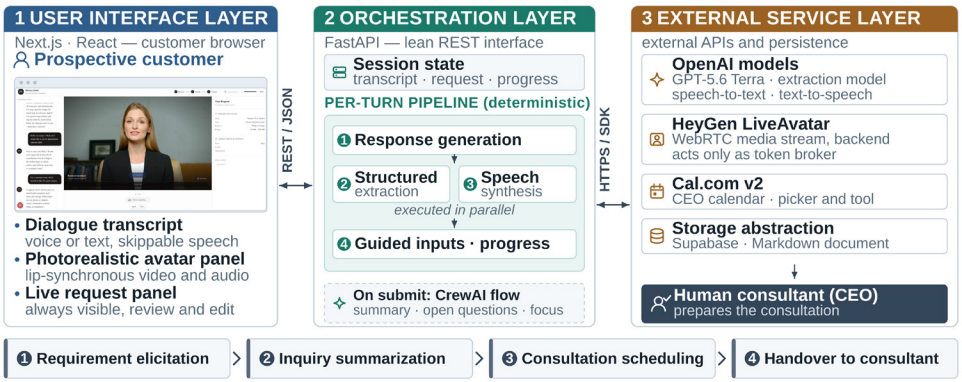


Figure 1: Reference architecture and customer-facing application of the first-contact agent.

EVALUATION RESULTS

Customer Perspective

Table 2 reports the questionnaire results (n = 19; 1 = strongly disagree, 7 = strongly agree). Trust was high (M = 5.86): the agent was perceived as highly capable of the task and of grasping the customer's concern (B2, M = 6.42) and as reliably doing what it announces (B6, M = 6.21)—on this item, 16 of 19 ratings were 6 or 7, two were 5, and one was 4. It was also predominantly perceived as acting in the customer's interest and as unbiased (B3, M = 5.53; B5, M = 5.68). Social presence showed very wide dispersion (person-level means from 2.33 to 6.33), with perceived human

warmth rated lowest (A3, Mdn = 3.0); at the same time, the dialogue was perceived as considerably more human-like and natural than participants' previous chatbot experiences (C2, M = 5.89). The agent's role was rated as very suitable for its purpose (D1, M = 5.84), and only two of 19 participants expressed a desire to switch to a human employee within the dialogue (D3 ≥ 5), whereas opinions on whether a human would perform the task better diverged strongly (D2, SD = 1.84), as did the expected efficiency gain over a hotline (D4, M = 4.53). Notably, in the forced-choice item, a majority (52.6%) would have started directly with the AI agent, only 21.1% would have waited for a human consultant, and 26.3% had no clear preference—with no discernible pattern between technically savvy and less savvy participants. Finally, ethical considerations were the most consistent aspect: participants were nearly universally aware of engaging with an AI (E1, M = 6.74), and nearly all participants expressed no concerns about disclosing the requested information (E2, M = 6.32, min = 4).

Table 2: Descriptive questionnaire results per construct (n = 19; social presence adapted from Gefen and Straub (2003, 2004); trust from McKnight et al. (2002)).

Construct (Items)	M	SD	Mdn
Social presence (A1–A3)	4.14	1.51	4.0
Initial trust (B1–B6)	5.86	1.17	6.0
Embodiment preference (C1–C2)	5.53	1.64	6.0
Role suitability (D1)	5.84	1.07	6.0
Human would perform better (D2)	3.58	1.84	3.0
Desire to switch to human (D3)	3.00	1.70	3.0
Efficiency gain vs. hotline (D4)	4.53	1.47	4.0
Disclosure clarity (E1)	6.74	0.73	7.0
Comfort with data disclosure (E2)	6.32	1.06	7.0

The interviews and observation notes explain this pattern along five themes. First, interaction styles were strongly heterogeneous, and the multimodal design absorbed this heterogeneity. Efficiency-oriented participants (P2, P5, P11) preferred typing, valued the quick-select options, and showed little tolerance for conversational pacing: P5 interrupted the agent mid-sentence and remarked laughingly, "this is annoying me already". P2 questioned the technology as such ("Why do I need AI for this—can't I just use a form?"), and P14 wanted to skip the agent's spoken text to speed up the process. Conversationally oriented participants (P4, P6, P9, P18) used the voice channel almost exclusively, engaged in small talk, and gave detailed, wish-rich answers. Explorative participants (P3, P7, P16) tested the system's limits out of curiosity, probed it for errors, and were more skeptical toward the AI than toward humans. Notably, P11 explicitly praised the guided navigation through the process over a conventional form.

Second, the photorealistic embodiment polarized. For several participants it created genuine social reactions: P6 initially focused on the person rather than the chat, praised the avatar ("she really looks good and likeable, I

like her natural hair") and spontaneously attributed a sales persona to it ("she probably wants to sell something too"); P9 was astonished that "she looks very real"; P7 felt as if already in a real sales process, "almost like human advice, just less annoying and pushy"; and P17 avoided the avatar's gaze yet spoke about the AI as if it were a person. For others, embodiment added no value or caused discomfort: P5 considered the video conversation unnecessary, P1 did not want to be advised by a computer that pretends to be human, P13 articulated a classic uncanny-valley reaction ("It feels uncomfortable, it's really weird—I know it's a computer, but also human"), and P16 reported that knowing the underlying technology suppressed any sense of human-likeness. Appearance also triggered gender attributions (P7: "Why a woman, and not a man?").

Third, trust was grounded in task competence and neutrality rather than human-likeness: the agent made no errors and rendered statements correctly (P1: "Wow, she even wrote my name correctly"), was perceived as neutral, focused, and precise (P7), and withstood active consistency checks (P3). P15 summarized the quality perception in contrast to conventional chatbots: "It felt high-quality and provided a feeling of competence and that real value will come out." Trust was accordingly high even among participants with low social presence ratings.

Fourth, participants drew a clear role boundary between structured intake and interpretive consultation. For capturing requirements and preparing an appointment, the agent was considered sufficient or even preferable due to immediacy, convenience, and the absence of social pressure: P15 noted that the process "took not even 5 minutes and I didn't have to drive anywhere," and P19, initially critical with high expectations, "would prefer it over a human live call or standardized predefined forms." For open-ended consultation, humans were expected to interpret ambiguous wishes better, and several participants expected deeper consulting, although the intake purpose was clearly communicated: "There was a lack of follow-up questions, which made the conversation feel stiff and less natural. A human would adapt more and provide a more personalized service" (P1). Perceived pressure to answer was the strongest friction—P8 was annoyed when the agent urged him to answer questions he could not yet answer. Balancing this is difficult, since other participants explicitly wanted a fast, efficient process.

Fifth, transparency and data handling were unproblematic, but the medium shaped disclosure behavior. No participant felt deceived, and comfort held as long as the requested data remained non-sensitive (P9: "I could have given any name"); P12 nevertheless argued that the agent should communicate even more clearly that it is an AI, so that customers do not feel the service is imitating human support. Several participants adapted their effort to the medium: P5 deliberately entered only short words because it is an AI, P2 would have elaborated more toward a human, and P3 monitored the AI more attentively than a human employee. Frequent improvement wishes included visual support such as example images of styles and objects, interruptible or skippable agent speech, lower latency, which made the conversation feel artificial despite the realistic avatar, and calmer avatar movement: "for me, it is too much movement" (P16).

Organization Perspective

In terms of usefulness and risk, the CEO aimed to balance efficiency with service quality: gathering necessary information before consultations while providing an on-demand service that surpassed basic forms or traditional chatbots. Positioning was equally pivotal — the agent was conceived as a distinctive differentiator not yet accessible from competitors, functioning as an intermediary between personal appointment scheduling and complete automation. This prevented customers from perceiving the initial contact as merely filling out a form. The value achieved manifested in three ways: (1) reduced preparation time as the agent organized inquiries and highlighted missing details; (2) reduced consultation times by using structured summaries instead of handwritten notes; and (3) interaction logs that became a valuable analytical tool, enabling the CEO to review contact times and recurring terms, and to refine the service accordingly.

Analyzing the role changes showed that the service employee's responsibilities expanded from appointment scheduling to include customer support tasks. The telephone channel was still maintained: about one in five participants preferred initial contact with a human, so disabling it was seen as a potential revenue risk. Therefore, both channels continue to operate simultaneously, and the employee must stay available for incoming calls. At the time of review, it was not yet clear whether the freed capacity would lead to a lasting redistribution of work. As part of the knowledge integration, the employee was involved from the beginning and contributed tacit knowledge about the previous intake process—such as common customer questions, relevant information for consultations, and how the task was handled previously. This input influenced the agent's elicitation process, including few-shot examples for common questions and the selection of data to ensure current pricing and availability. Both employees tested the application, and their feedback confirmed the accuracy of the agent's responses before deployment.

Derivation of Design Principles

Synthesizing customer and organizational perspectives yields five design principles (DPs), structured according to the anatomy from Gregor et al. (2020).

In each of the five cases, the implementers are designers of embodied first-contact agents for knowledge-intensive consulting services, and the users are prospective customers unless otherwise specified.

DP1 (Scope framing): To achieve calibrated expectations about the service scope and agent capabilities, define the agent's role, purpose, and boundaries at the beginning of the session. Reinforce this framing when requests go beyond the defined scope. The agent should communicate this naturally during the conversation, backed by the organizational context. This approach leverages photorealistic embodiment and fluent dialogue to align user expectations with human consultations, since simply revealing the AI identity does not effectively calibrate those expectations.

DP2 (Customer-governed pacing): To achieve complete elicitation without making the customer feel pressured, guide the conversation smoothly by keeping unanswered fields pending, yielding instantly if interrupted, allowing speech to be omitted, and transferring these deferred fields into the handoff artifact as open questions. The customer should set the pace, while the agent accepts deferral, since in a photorealistic embodiment, insistence is seen as interpersonal pressure rather than a neutral system feature.

DP3 (Configurable modality and representation): To achieve interaction fit across diverse customer styles, enable interchangeability of text, speech, and guided input at every stage. Allow visual representations to be simplified without sacrificing functionality. Keep idle motion calm and stable, and support decision-making with example images. Customers should choose their preferred modality, while the agent maintains dialogue continuity across switches. Since embodiment functions as a segmentation variable rather than a universal enhancer, fixed configurations are ineffective for certain user segments.

DP4 (Demonstrable attentiveness): To achieve competence-based trust and prevent strategic effort reduction, ensure reliability-critical functions are deterministic instead of generative, repeat captured statements exactly and instantly, keep latency below the level where turn-taking feels unnatural, and ask follow-up questions that clearly relate to what was just said. The agent should intermittently extract information and probe during customer verification, as trust depends on observable task performance rather than human-like qualities. Customers also adjust their willingness to share based on how attentive they perceive the recipient to be.

DP5 (Organizationally defined elicitation scope): To achieve a handoff artifact ready for consultation without needing post-hoc review, the consulting organization should predefine the relevant fields and their meanings, guide, test, and restrict the agent to this scope, display the collected data continuously, and include a review and editing step before sending. The agent should gather information within these parameters, while the customer verifies the captured data. A clearly defined and tested scope limits the potential for the agent to drift into unnecessary or incorrect content, and continuous visibility allows the party best suited to assess accuracy to perform verification.

DISCUSSION

The findings carry three theoretical implications. First, they decouple trust from social presence for transactional service episodes: trust was high while social presence was widely dispersed and perceived warmth low, and a majority preferred starting with the AI over waiting for a human—independent of technical affinity. For structured intake tasks, competence-based trust thus dominates presence-based trust, nuancing service research that positions automated social presence as a central experiential lever (van Doorn et al., 2017) and consistent with task-dependent acceptance of algorithmic providers (Castelo et al., 2019). Second, embodiment operated as a segmentation variable rather than a universal enhancer: reactions ranged

from genuine social engagement to uncanny-valley discomfort (P13) and knowledge-based suppression of human-likeness (P16)—offering a candidate framework extension: interaction-style heterogeneity as an explicit design factor. Third, transparent disclosure did not harm the interaction, in contrast to concerns that AI disclosure deteriorates outcomes (Luo et al., 2019); instead, it shifted behavior toward shorter, monitored input—a trade-off between honesty-enabled comfort and information richness that first-contact designs must actively manage.

Several limitations qualify these results. The sample (n = 19) is modest and self-selected, and all results are descriptive. The study covers a single case, language, cultural context, and a single female avatar persona; single-session use implies possible novelty effects. Future work should vary embodiment and persona systematically and examine downstream effects on conversion and consultation quality.

CONCLUSION

This paper designed, implemented, and evaluated a photorealistic conversational AI agent delivering an end-to-end first-contact service in a knowledge-intensive consulting environment, and demonstrates that such a service can be realized with a lightweight, web-based agentic stack. Customers trusted and accepted the agent for structured intake without requiring social presence, while embodiment polarized along interaction styles—a segmentation the design must accommodate rather than resolve. The five design principles are transferable to professional service contexts in which the first contact prepares, rather than replaces, human consultation.

ACKNOWLEDGMENT

The authors acknowledge all participants who contributed to the study.

REFERENCES

Bitner, M. J., Ostrom, A. L., Morgan, F. N. (2008). Service blueprinting: A practical technique for service innovation. California Management Review, Volume 50, No. 3, pp. 66–94.

Braun, V., Clarke, V. (2006). Using thematic analysis in psychology. Qualitative Research in Psychology, Volume 3, No. 2, pp. 77–101.

Cassell, J., Sullivan, J., Prevost, S., Churchill, E. F. (Eds.) (2000). Embodied Conversational Agents. Cambridge, MA: MIT Press.

Castelo, N., Bos, M. W., Lehmann, D. R. (2019). Task-dependent algorithm aversion. Journal of Marketing Research, Volume 56, No. 5, pp. 809–825.

Diederich, S., Brendel, A. B., Morana, S., Kolbe, L. (2022). On the design of and interaction with conversational agents: An organizing and assessing review of human-computer interaction research. Journal of the Association for Information Systems, Volume 23, No. 1, pp. 96–138.

Gefen, D., Straub, D. W. (2003). Managing user trust in B2C e-services. e-Service Journal, Volume 2, No. 2, pp. 7–24.

Gefen, D., Straub, D. W. (2004). Consumer trust in B2C e-commerce and the importance of social presence: Experiments in e-products and e-services. Omega, Volume 32, No. 6, pp. 407–424.

Gnewuch, U., Morana, S., Maedche, A. (2017). Towards designing cooperative and social conversational agents for customer service, in: Proceedings of the 38th International Conference on Information Systems (ICIS 2017), Seoul.

Gregor, S., Chandra Kruse, L., Seidel, S. (2020). Research perspectives: The anatomy of a design principle. Journal of the Association for Information Systems, Volume 21, No. 6, pp. 1622–1652.

Halim, N. A. A., Suaib, N. M., Fadzli, F. E. (2026). AI-powered avatar assistants: A systematic literature review of tools, application domains and challenges. Journal of Ambient Intelligence and Humanized Computing, Volume 17, No. 2, pp. 633–670.

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, Volume 38, No. 6, pp. 937–947.

McKnight, D. H., Choudhury, V., Kacmar, C. (2002). Developing and validating trust measures for e-commerce: An integrative typology. Information Systems Research, Volume 13, No. 3, pp. 334–359.

Nass, C., Moon, Y. (2000). Machines and mindlessness: Social responses to computers. Journal of Social Issues, Volume 56, No. 1, pp. 81–103.

Peffers, K., Tuunanen, T., Rothenberger, M. A., Chatterjee, S. (2007). A design science research methodology for information systems research. Journal of Management Information Systems, Volume 24, No. 3, pp. 45–77.

Rashik, M., Jasim, M., Kucher, K., Sarvghad, A., Mahyar, N. (2024). Beyond text and speech in conversational agents: Mapping the design space of avatars, in: Proceedings of the 2024 ACM Designing Interactive Systems Conference (DIS 2024), Copenhagen, pp. 1875–1894.

Tugarin, N., van Husen, C. (2025). Development and integration of human-AI interactions in service applications: Conceptual framework and review. International Journal of Information Management Data Insights, Volume 5, No. 2, 100357.

Tugarin, N., van Husen, C. (2026). Designing and evaluating human–AI interaction service systems: An industrial case study. International Journal of Information Management Data Insights, Volume 6, No. 2, 100438.

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, Volume 20, No. 1, pp. 43–58.

Venable, J., Pries-Heje, J., Baskerville, R. (2016). FEDS: A framework for evaluation in design science research. European Journal of Information Systems, Volume 25, No. 1, pp. 77–89.