One Body, Many Minds: Exploring Design Frictions and Considerations for Multi-Agentic Systems in Human-Robot Interaction

Sarah Schömbs
The University of Melbourne
Melbourne, Australia
sschombs@student.unimelb.edu.au

Yan Zhang
The University of Melbourne
Melbourne, Australia
yanzhang15@student.unimelb.edu.au

Canghai Wang
The University of Melbourne
Melbourne, Australia
canghaiw@student.unimelb.edu.au

Wafa Johal
The University of Melbourne
Melbourne, Australia
wafa.johal@unimelb.edu.au

Abstract

When multi-agentic systems are implemented in human-robot interaction, users interact with a single physical robot while multiple AI agents operate 'behind the curtain'. This creates a fundamental mismatch: the system's multi-agent nature is hidden from users, but understanding agent roles, capabilities, and orchestration is critical for system transparency and user understanding. We investigate how multi-agentic systems can be represented through a single anthropomorphic robot. We conducted a design workshop with seven HCI experts to explore design parameters for system characteristics (e.g. objective, persona, tool, handoff) and discuss orchestration configurations through design probes deployed on a Furhat platform. We contribute preliminary expert-informed design considerations for encoding multi-agentic system characteristics through anthropomorphic robots and identify design frictions to inform the design of transparent system orchestration in single-robot platforms.

CCS Concepts

• **Human-centered computing → User studies**.

Keywords

multi-agent systems, agentic, HRI, embodiment, anthropomorphism, transparency, orchestration

ACM Reference Format:
Sarah Schömbs, Yan Zhang, Canghai Wang, and Wafa Johal. 2026. One Body, Many Minds: Exploring Design Frictions and Considerations for Multi-Agentic Systems in Human-Robot Interaction. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3772363.3799366

CHI EA '26, Barcelona, Spain

ACM ISBN 979-8-4007-2281-3/2026/04
https://doi.org/10.1145/3772363.3799366

1 Introduction and Motivation

Recent developments in large language models have enabled multi-agentic systems (e.g. AutoGen [33], MetaGPT [15], OpenAI Agents [4], CrewAI [3]), where multiple AI agents collaborate to accomplish complex tasks for users with a high degree of agency [9]. Such agentic systems require minimal operational specification while exhibiting goal-directed behaviour, high degrees of environmental impact, and long-term planning capabilities [9], which signals a critical shift in autonomous systems. To illustrate, a user might request help booking a flight through a chat-based interface that employs a multi-agentic system. Instead of a single LLM-driven agent handling this request based on a monolithic prompt or sequential back-and-forth queries, multiple agents now collaboratively act on the task, each specialised to compare prices, identify promotion codes, handle data safekeeping, and optimises route planning. Importantly, each of these agents are defined by a set of characteristics that define the agent's objective, operationalised through a set of instructions, its persona that characterises "who or what the model is acting as" [12], and its capabilities (tool access, e.g. web search, code execution) [6, 35].

However, the advancing system complexity introduces critical design challenges related to system understanding, agent communication, transparency and agent visibility [5, 8, 28]. In existing frameworks (e.g., OpenAI Agent [4] and AutoGen [33]) and user interfaces (e.g., MagenticUI [23]), an 'orchestration agent' often functions as a top-level supervisory entity that serves as the communicative bridge between user and system, delegates tasks to specialised subagents, and, by doing so, simplifies interaction to improve system robustness [11]. Such top-level orchestration abstracts away low-level system operations and omits low-level scaffolding. Emerging design challenges include *how to enable users to form appropriate mental models of agentic teams and their underlying specialised agents*, *how to support the identification of agent capabilities, roles, and objectives within networks of agents*, and *how to enable users to trace emergent complexities (e.g. increasing number of agents; agents creating new subagents)* [5, 8, 28]. New design solutions and interaction paradigms are urgently needed to address these emerging challenges.

When a multi-agentic architecture is implemented in human-robot interaction (HRI) to improve system performance [11] and to enable users to achieve more complex tasks, the 'multi' nature of the system could be similarly omitted, as the user solely interacts with a single physical robot. This presents both an important research gap and a design opportunity. First, HRI researchers need to investigate how to convey agent characteristics and orchestration to users through a robot platform. Second, robots, as physical embodied interfaces, offer unique and intriguing design affordances that could inform interface design more broadly across HCI contexts. Concepts from HRI such as re-embodiment (an agent's "social presence" migrates across multiple physical bodies) and co-embodiment (multiple agent identities inhabit a singular physical robot body) [20] offer interesting frameworks for representing multi-agentic systems in HRI. Further, robot platforms enable interaction paradigms beyond text-based prompting, including verbal communication [34, 36], gaze [29], and facial expressions [25, 26], which provide additional modalities to communicate system characteristics. Recent advances in facial projection for anthropomorphic social robot platforms like Furhat [1] allow robots to switch between 'characters' and granularly manipulate facial features. These capabilities enable information visualisation beyond human emotion through metaphorical representations [26], e.g. to communicate health risks by mirroring health consequences through facial expressions [25]. Further, the concept of AI agents defined by 'personas' with 'tools' to accomplish 'objectives' is inherently anthropomorphic, which suggests that anthropomorphism by design could make characteristics of multi-agentic systems more accessible and understandable to users. Notably, anthropomorphism by design, where features are deliberately designed to elicit anthropomorphisation [32], has been shown to facilitate trust, likeability, and perceived intelligence [24]. It extends to nuanced perceptions, where robot face design alone enables users to ascribe and distinguish personality traits such as introversion and extroversion [16].

Motivated by current advances in multi-agentic systems and the the unique affordances of a robot's face, we investigate how multi-agentic AI systems can be represented through a single anthropomorphic robot, guided by two overarching research questions: **(RQ1) How can we design multi-agentic system characteristics through a robot's face?**, and **(RQ2) How are different orchestration configurations affecting people's perception and understanding of multi-agentic systems?**. To address our research questions, we conducted a design workshop with seven HCI experts comprising brainstorming activities (RQ1) and a demonstration of three orchestration configurations (Opaque Orchestration, Direct Delegation, Transparent Orchestration) deployed on the Furhat robot platform followed by a focus group discussion (RQ2). Our findings offer preliminary contributions that guide future investigations: First, we provide initial expert-informed design considerations for encoding multi-agentic system characteristics through anthropomorphic physical robots. Second, we identify design frictions for transparent system orchestration drawn our qualitative findings. This work lays the groundwork for future investigations in which we plan to build on these findings to establish a comprehensive design space for multi-agentic system representation in single-robot platforms.

[1] Furhat Robotics, https://www.furhatrobotics.com/.

2 Method

We conducted an in-person design workshop for which we recruited experts ($N = 7$; $5F$, $2M$) in HCI-related fields, e.g. game design, tangibles, generative AI, and human-AI interaction to explore design ideas, constraints, and to facilitate critical discussion [13, 19, 26], an established practice in HCI. The workshop was structured into two phases: the first phase explored design parameters for multi-agentic system characteristics (RQ1), and the second phase examined orchestration configurations (RQ2). The design workshop was approved by the University's ethics committee.

2.1 Stimuli for the Design Workshop

2.1.1 System Characteristics. To address RQ1, we presented the workshop participants with a set of carefully selected system characteristics based on existing frameworks [6, 12] that define important parameters when understanding and interacting with agents in a multi-agentic system: `[objective]` defines what the agent is supposed to achieve, `[persona]` defines the role the agent adopts [12], `[tool]` defines agent capabilities such as access to functions, external tools like code execution or API calls, including `[agent-as- tool]` where agents can use other agents as tools [4], `[instruction]` defines steps or operations to accomplish tasks, `[handoff]` defines agent delegation of (sub)tasks to other agents, and `[agent creation]` defines agents creating other subagents.

2.1.2 Design Probe: Agent Implementation and Orchestration Configurations. The objective of the design probe demonstrated during the workshop was to showcase (1) example AI agents for everyday decision-making, and (2) orchestration configurations to spark critical discussion about transparency and system understanding (RQ2). We deployed four specialised AI agents presented as part of a multi-agent team using Kotlin, the Furhat SDK, and OpenAI's API:

`[The Orchestrator Agent]`, which coordinates agent selection and synthesises responses; `[The Creative Agent]`, which generates creative ideas and unconventional solutions; `[The Logical Agent]`, which provides analytical reasoning and structured thinking; and `[The Affective Agent]`, which offers empathetic support and prioritises human connection. Each agent was defined through structured prompts specifying the role, communication style, output format, and expected behaviour. We implemented distinct built-in character faces ("Alex", "Emma", "Maurice", "Jane") and voices from Furhat's character library for each agent to make them distinguishable during the demo and to elicit critical discussion about agent representations (varying age, gender, ethnicity). We custom-coded the multi-agent orchestration logic (agent selection, response synthesis, turn-taking) to maintain controllability during the workshop following established practices in HCI research [22, 27]. This approach enabled us to deploy and demonstrate three distinct orchestration configurations (illustrated in Figure 2) to gain insights into RQ2:

Opaque Orchestration: The user only sees and interacts with a single orchestrator agent; inter-agent consultation remains opaque.

Direct Delegation: The orchestrator agent selects suitable specialised agents and directly hands over the conversation to that subagent (e.g., "[Specialised Agent], your turn"). The selected agent then responds directly to the user.

Transparent Orchestration: The orchestrator agent transparently consults each specialised subagent, allowing the user to observe the consultation process, then selects and conveys the most appropriate response to the user.

2.2 Procedure

Prior to the workshop, participants received a plain language statement about the research objectives and provided written consent. The workshop began with an introduction and an ice-breaker activity to create a collaborative space for open discussion [18]. We presented a 10-minute overview to agentic systems, system architectures, and multi-agentic teams to reduce knowledge imbalances, followed by a demonstration of the Furhat robot platform to familiarise participants with the system. In the first phase (RQ1), we explored design parameters for multi-agentic system through brainstorming activities. For each system characteristic (see Sec. 2.1.1), we presented a definition accompanied with an example derived from established frameworks. Participants then engaged in 90-second sticky-note brainstorming sessions (see Figure 1) to individually generate ideas for embodying the system characteristic through the robot. After each round, participants presented their ideas and collaboratively discussed and clustered them on paper sheets. Following the brainstorming activities, we held a brief discussion about design challenges and opportunities. After a short break, the second phase focused on the orchestration configurations (RQ2). A researcher from our team interacted with the Furhat robot platform across the three deployed orchestration configurations (described in Sec. 2.1.2, illustrated in Figure 2) through example scenarios (planning a birthday party, managing exam stress by increasing study productivity, resolving workplace conflicts) while the workshop participants were instructed to observe the interactions. After each demonstration, we held a short focus group discussions about their observations, implementation preferences, and identified trade-offs. The workshop took ~3h and was audio recorded.

Figure 1: A photo of participants brainstorming in the design workshop

2.3 Analysis

We conducted a reflexive thematic analysis following the steps provided by Braun and Clarke [7]. Two researchers independently familiarised themselves with the transcribed audio data and iteratively coded the qualitative data independently. Through two collaborative sessions, they discussed the identified codes, refined themes, reduced discrepancies and build a shared understanding of the data. While the research objectives guided initial tentative themes, the exploratory nature of the study enabled a more inductive approach to reveal unexpected insights. We present results related to the design parameters for encoding multi-agentic systems through robots, which informs a preliminary set of design considerations (RQ1), and results from the focus group discussion on orchestration configurations, which informed preliminary design considerations (RQ2). Due to the limited scope, we present results and discussion together in the following sections.

3 Design Considerations and Challenges for Embodying Multi-agentic through a Single Robot

Persona and Objectives. Participants identified the use of visual and vocal cues to embody both objectives and personas. Objectives were proposed to be communicated through verbal announcements, while personas were differentiated through subtler variations in voice, pace, tone, and lexical choice. Personas add "flavour" (P4) by differentiating agent characteristics even while maintaining the same underlying objective, such as "a harsh tutor versus a kind and caring one" (P4). Interestingly, participants ideated more anthropomorphic approaches to encode personas compared to objectives. To design objectives through embodiment, participants described the use of colour coding, facial accessories, tattoos or icons, alongside more advanced visualisations, including a transparent brain that reveals symbolic elements representing the agent's objectives, which aligns with design parameters identified in prior works [25]. For personas, participants suggested more anthropomorphic features such as age associations with specific professions, wrinkles to encode experience and elicit trust, alongside facial accessories (e.g. glasses to signal specific roles like "accounting" (P4). Further, participants suggested the use of distinct faces corresponding to different AI agent personas, with one participants proposing to re-embody "famous people" (P3) as persona representations. However, participants raised concerns that while encoding personas and objectives through anthropomorphic features to activate familiar cognitive schemas could support immediate understanding, it risks reinforcing bias and stereotyping and carries cultural assumptions.

Instructions. Participants questioned the need to embody and visualise instructions that guide the agents operations due to their inherent complexity and system internal nature. Non-trivially, participants identified the communication of task *progress* as more important from an end-user perspective. Progress could be communicated verbally, through queries or announcements, or visually, via anthropomorphic cues such as an increasingly wide smile (P3, P7), or non-anthropomorphic indicators such as a loading bar displayed on the face (P6). Designing an appropriate level of granularity remained contested, as users may prefer to maintain a high-level understanding.

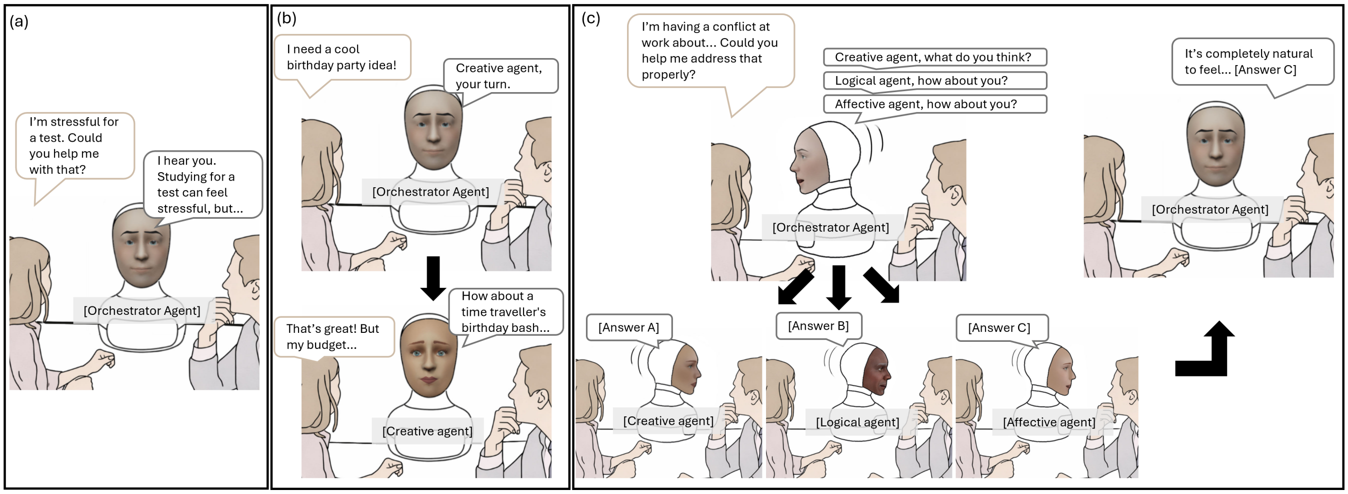

Figure 2: Orchestration configurations. (a) Opaque Orchestration. (b) Direct delegation: direct handoff to visible subagent. (c) Transparent Orchestration: orchestrator visibly consults all agents via turn-taking (head-turning) before conveying selected response to user.

Tools. Tool use was primarily visualised using static elements such as tattoos or icons represented as eyes iris (P2, P4), as well as dynamic elements such as visualising a transparent brain with tools floating inside (P2, P3) to signal system capabilities, or dynamic cues that highlighted the tool currently in use. Importantly, participants emphasised that indicating *when* a tool is being accessed, and an indication of its *duration* was more important than displaying the full set of available tools.

Handoff. Participants identified inter-agent handoffs as an important system characteristic, agreeing that clear communication is required when agents transfer tasks within a multi-agent system. Most participants emphasised that handoffs should be verbally announced, as users may not always be visually attending to the robot, for example using warning phrases such as "(Danny Phantom) I'm going ghost. [...] go, go, Power Rangers." (P4). Visual cues included colour changes (P3) and flashing facial displays (P4). Participants foregrounded that handoff embodiment should communicate both the *start* and *goal state*, "I would want to know 'who' and 'to"' (P4), and convey the transfer of contextual knowledge, for instance through swapping agent profile icons (P5) or explicit receipt confirmation (P7). However, few participants again debated the necessity of exposing handoffs to users. P1 argued for prioritising task progress over explicit handoff visualisation, while P6 raised concerns about user overload and advocated for subtler, less intrusive designs.

Agents creating Subagents. Participants conceptualised the relationship between agents and subagents through metaphors such as family tree (P4), ancestral lineage (P3, P5), and organigram (P2). They generally agreed that such "genealogical" (P3) relationships should be made visible and that agent-created subagents should be clearly distinguished from those created by users. This ties closely into the current discussion around agent identifiers [8] and AI disclosure information [10, 31], which presents a critical concern in operationalising Article 50 EU AI Act transparency obligations [1]. P2 argued that agent-created subagents should exhibit more robotic features, a view supported by P5 and P4, who suggested distinct markers such as a Tamagotchi egg icon to draw users' attention to other "stranger" agents that may require closer monitoring. P4 proposed the use of inherited visual language to signal shared lineage and objectives, for example monocles to indicate agents belonging to the same "tutor family", also re-emphasising that subagents further become progressively less anthropomorphic as they diverge from the root agent. Similarly, P3 suggested encoding parent–child relationships through anthropomorphised age cues, whereby subagents appear younger and supervisory agents appear older. However, P6 and P2 questioned the need to explicitly disclose subagents at all times, suggesting that, when necessary, subagents could instead be visualised through abstract representations such as neural or network-layer overlays on the face.

Design Frictions. One of the most widely raised concerns was introducing unnecessary complexity which risks overwhelming users, a common concern also raised in information visualisation design [14, 17] and prior works using a robot as an interface to encode information [26]. This led to a central design challenge regarding what information should be embodied and what should be adaptively embodied, and "why embodiment matters in certain cases" (P4, P6). The timing of information delivery was also identified as a potential mechanism for balancing this complexity (P2).

Further, participants identified anthropomorphism as a double-edged sword. Designers must carefully determine which elements should be anthropomorphised and which should not. While it enabled intuitive encoding of system characteristics, participants consistently questioned whether intentionally catering to archetypes held in the public imagination risks re-encoding stereotypes and reinforcing existing social biases (P4). Therefore, P6 suggested allowing user customisation, which in turn introduces the challenge of defining appropriate default embodiment settings (P6, P7). Lastly, designers need to account for users' mental models, shaped by lived experience, including whether target users are familiar with handling a team of people (P4) and individual differences in familiarity with anthropomorphised robots (P5). These factors should inform design decisions about the extent to which agents are visible and information is anthropomorphised.

4 Design Considerations for Transparency

The varying degrees of transparency within each orchestration configuration prompted in-depth discussion about designing organic interactions while accounting for context-dependent nuances. **Direct delegation was described as more formal and better suited to complex tasks**, such as organising an office event (P1, P4). Rich visibility of multiple agents can elicit a sense of support akin to working with a team of people (P1). However, the group discussion revealed that users may not be genuinely accustomed to coordinating a team, particularly in personal contexts such as planning a birthday party or engaging in emotionally sensitive tasks. In such cases, **opaque orchestration, which hides subagents from view, was considered more appropriate for everyday use**. Nevertheless, **transparent orchestration was identified as necessary when users explicitly request explanations** for the orchestrator's responses (P1). Notably, a mismatch between task complexity and the visibility of agents (visible multi-agent communication for low-complexity tasks) may lead to perceptions of system incompetence and reduced trust (P4), which suggests that **task complexity determines the desired level of agent visibility** and designing context-dependent transparency requires careful consideration.

Importantly, most participants believed that users should have the option to engage in agent communication and to act as **active collaborators rather than passive observers**. Imperatively, high transparency without mechanisms for immediate participation was perceived as exclusionary (P1, P5), with P1 noting the importance of **interruptibility as a form of in-situ debugging**, similar to coaching a team. This notion aligns with human-in-the-loop considerations for graphical user interfaces of multi-agentic systems, which require 'interrupts' [2] and other oversight mechanisms [8, 21, 23] to mitigate AI harms in agentic systems [9]. Moreover, increased transparency and access to inter-agent communication prompted suggestions that system design should be aligned with individual user preferences, particularly in how the orchestrator determines and prioritises what is important to the user. This supports the idea of *interactive alignment* in human-AI interaction [30]. Participants also debated whether users should have direct access to subagents. P3 argued that such **access increases user agency by allowing involvement in subagent decision-making**, whereas P4 advocated for designs that **maintain a high-level overview, providing lower-level scaffolding only when needed**. This raises design considerations such as implementing 'shortcut mechanisms' that allow users to bypass the orchestrator in hierarchical architectures when appropriate, or designing 'silent' orchestration modes where the orchestrator transparently hands off to subagents without intermediating the interaction.

Our work offers exciting preliminary insights into representing multi-agentic systems through anthropomorphic robots to support transparency and system understanding. The expert-informed design considerations we surface inform HRI research and may extend to interface design for on-screen agents. We further reveal future research avenues such as agent identity and roles in multi-agentic systems. Next steps include building on these findings to establish a comprehensive design space for multi-agentic system representation in single-robot platforms, alongside user studies to validate our design considerations and investigate how orchestration transparency affects user understanding and trust in single-robot platforms.

References

[1] [n. d.]. The AI Act Explorer | EU Artificial Intelligence Act. https://artificialintelligenceact.eu/ai-act-explorer/
[2] 2025. langchain-ai/langgraph-studio. https://github.com/langchain-ai/langgraph-studio original-date: 2024-07-29T22:11:00Z.
[3] CrewAI . [n. d.]. CrewAI. https://www.crewai.com/
[4] OpenAI-Agent. 2025. openai-agent. https://github.com/openai/openai-agents-python/
[5] Gagan Bansal, Jennifer Wortman Vaughan, Saleema Amershi, Eric Horvitz, Adam Fourney, Hussein Mozannar, Victor Dibia, and Daniel S. Weld. 2024. Challenges in Human-Agent Communication. https://doi.org/10.48550/arXiv.2412.10380 arXiv:2412.10380 [cs].
[6] Ilan Bigio. 2024. Orchestrating Agents: Routines and Handoffs. OpenAI Cookbook. https://cookbook.openai.com/examples/orchestrating_agents Accessed: 2026-01-22.
[7] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (Jan. 2006), 77–101. https://doi.org/10.1191/1478088706qp063oa _eprint: https://www.tandfonline.com/doi/pdf/10.1191/1478088706qp063oa.
[8] Alan Chan, Carson Ezell, Max Kaufmann, Kevin Wei, Lewis Hammond, Herbie Bradley, Emma Bluemke, Nitarshan Rajkumar, David Krueger, Noam Kolt, Lennart Heim, and Markus Anderljung. 2024. Visibility into AI Agents. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Rio de Janeiro Brazil, 958–973. https://doi.org/10.1145/3630106.3658948
[9] Alan Chan, Rebecca Salganik, Alva Markelius, Chris Pang, Nitarshan Rajkumar, Dmitrii Krasheninnikov, Lauro Langosco, Zhonghao He, Yawen Duan, Micah Carroll, Michelle Lin, Alex Mayhew, Katherine Collins, Maryam Molamohammadi, John Burden, Wanru Zhao, Shalaleh Rismani, Konstantinos Voudouris, Umang Bhatt, Adrian Weller, David Krueger, and Tegan Maharaj. 2023. Harms from Increasingly Agentic Algorithmic Systems. In *2023 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Chicago IL USA, 651–666. https://doi.org/10.1145/3593013.3594033
[10] Abdallah El Ali, Karthikeya Puttur Venkatraj, Sophie Morosoli, Laurens Naudts, Natali Helberger, and Pablo Cesar. 2024. Transparent AI Disclosure Obligations: Who, What, When, Where, Why, How. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–11. https://doi.org/10.1145/3613905.3650750
[11] Dawei Gao, Zitao Li, Xuchen Pan, Weirui Kuang, Zhijian Ma, Bingchen Qian, Fei Wei, Wenhao Zhang, Yuexiang Xie, Daoyuan Chen, et al. 2024. Agentscope: A flexible yet robust multi-agent platform. *arXiv preprint arXiv:2402.14034* (2024).
[12] Google Cloud. 2026. Overview of Prompting Strategies. Google Cloud Documentation: Generative AI on Vertex AI. https://docs.cloud.google.com/vertex-ai/generative-ai/docs/learn/prompts/prompt-design-strategies Last updated: 2026-01-19 UTC. Accessed: 2026-01-22.
[13] Katerina Gorkovenko, Daniel J. Burnett, James K. Thorp, Daniel Richards, and Dave Murray-Rust. 2020. Exploring The Future of Data-Driven Product Design. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–14. https://doi.org/10.1145/3313831.3376560
[14] Miriam Greis, Jessica Hullman, Michael Correll, Matthew Kay, and Orit Shaer. 2017. Designing for Uncertainty in HCI: When Does Uncertainty Help?. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI EA '17)*. ACM, New York, NY, USA, 593–600. https://doi.org/10.1145/3027063.3027091
[15] Sirui Hong, Xiawu Zheng, Jonathan Chen, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, et al. 2023. Metagpt: Meta programming for multi-agent collaborative framework. *arXiv preprint arXiv:2308.00352* (2023).
[16] Soyoung Jung, Hyoung-taek Lim, Sanghun Kwak, and Frank Biocca. 2012. Personality and facial expressions in human-robot interaction. In *Proceedings of the*

seventh annual ACM/IEEE international conference on Human-Robot Interaction (HRI '12). Association for Computing Machinery, New York, NY, USA, 161–162. https://doi.org/10.1145/2157689.2157735

[17] Matthew Kay, Tara Kola, Jessica R. Hullman, and Sean A. Munson. 2016. When (Ish) is My Bus? User-Centered Visualizations of Uncertainty in Everyday, Mobile Predictive Systems. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. ACM, New York, NY, USA, 5092–5103. https://doi.org/10.1145/2858036.2858558

[18] Ethan Kerzner, Sarah Goodwin, Jason Dykes, Sara Jones, and Miriah Meyer. 2019. A Framework for Creative Visualization-Opportunities Workshops. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan. 2019), 748–758. https://doi.org/10.1109/TVCG.2018.2865241 Conference Name: IEEE Transactions on Visualization and Computer Graphics.

[19] Xingyu Lan, Yanqiu Wu, Yang Shi, Qing Chen, and Nan Cao. 2022. Negative Emotions, Positive Outcomes? Exploring the Communication of Negativity in Serious Data Stories. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–14. https://doi.org/10.1145/3491102.3517530

[20] Michal Luria, Samantha Reig, Xiang Zhi Tan, Aaron Steinfeld, Jodi Forlizzi, and John Zimmerman. 2019. Re-Embodiment and Co-Embodiment: Exploration of social presence for robots and conversational agents. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. ACM, San Diego CA USA, 633–644. https://doi.org/10.1145/3322276.3322340

[21] Leila Methnani, Andrea Aler Tubella, Virginia Dignum, and Andreas Theodorou. 2021. Let Me Take Over: Variable Autonomy for Meaningful Human Control. *Frontiers in Artificial Intelligence* 4 (Sept. 2021), 737072. https://doi.org/10.3389/frai.2021.737072

[22] Saumya Pareek, Sarah Schömbs, Eduardo Velloso, and Jorge Goncalves. 2025. "It's Not the AI's Fault Because It Relies Purely on Data": How Causal Attributions of AI Decisions Shape Trust in AI Systems. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–18. https://doi.org/10.1145/3706598.3713468

[23] Brenda Potts. 2025. Magentic-UI, an experimental human-centered web agent. https://www.microsoft.com/en-us/research/blog/magentic-ui-an-experimental-human-centered-web-agent/

[24] E. Roesler, D. Manzey, and L. Onnasch. 2021. A meta-analysis on the effectiveness of anthropomorphism in human-robot interaction. *Science Robotics* 6, 58 (Sept. 2021), eabj5425. https://doi.org/10.1126/scirobotics.abj5425

[25] Sarah Schömbs, Jorge Goncalves, and Wafa Johal. 2025. "I can feel the risks by looking at the robot face": Communicating Risk through a Physical Agent. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference (DIS '25)*. Association for Computing Machinery, New York, NY, USA, 236–252. https://doi.org/10.1145/3715336.3735759

[26] Sarah Schömbs, Jiahe Pan, Yan Zhang, Jorge Goncalves, and Wafa Johal. 2024. FaceVis: Exploring a Robot's Face for Affective Visualisation Design. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, 1–10. https://doi.org/10.1145/3613905.3650910

[27] Sarah Schömbs, Saumya Pareek, Jorge Goncalves, and Wafa Johal. 2024. Robot-Assisted Decision-Making: Unveiling the Role of Uncertainty Visualisation and Embodiment. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–16. https://doi.org/10.1145/3613904.3642911

[28] Sarah Schömbs, Yan Zhang, Jorge Goncalves, and Wafa Johal. 2025. From Conversation to Orchestration: HCI Challenges and Opportunities in Interactive Multi-Agentic Systems. In *Proceedings of the 13th International Conference on Human-Agent Interaction*. ACM, Yokohama Japan, 158–168. https://doi.org/10.1145/3765766.3765795

[29] Gabriel Skantze and Bahar Irfan. 2025. Applying General Turn-Taking Models to Conversational Human-Robot Interaction. In *2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 859–868. https://doi.org/10.1109/HRI61500.2025.10973958

[30] Michael Terry, Chinmay Kulkarni, Martin Wattenberg, Lucas Dixon, and Meredith Ringel Morris. 2024. Interactive AI Alignment: Specification, Process, and Evaluation Alignment. https://doi.org/10.48550/arXiv.2311.00710 arXiv:2311.00710 [cs].

[31] Benjamin Toff and Felix M. Simon. 2025. "Or They Could Just Not Use It?": The Dilemma of AI Disclosure for Audience Trust in News. *The International Journal of Press/Politics* 30, 4 (Oct. 2025), 881–903. https://doi.org/10.1177/19401612241308697

[32] Adam Waytz, John Cacioppo, and Nicholas Epley. 2010. Who Sees Human?: The Stability and Importance of Individual Differences in Anthropomorphism. *Perspectives on Psychological Science* 5, 3 (May 2010), 219–232. https://doi.org/10.1177/1745691610369336

[33] Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W. White, Doug Burger, and Chi Wang. 2023. AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation. https://doi.org/10.48550/arXiv.2308.08155 arXiv:2308.08155 [cs].

[34] Yan Zhang. 2025. Implicit Communication of Contextual Information in Human-Robot Collaboration. In *2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 1903–1905. https://doi.org/10.1109/HRI61500.2025.10974036

[35] Yan Zhang, Haoqi Li, Ramtin Tabatabaei, and Wafa Johal. 2025. ROSAnnotator: A Web Application for ROSBag Data Analysis in Human-Robot Interaction. In *2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 1099–1103. https://doi.org/10.1109/HRI61500.2025.10974254

[36] Yan Zhang, Tharaka Sachintha Ratnayake, Cherie Sew, Jarrod Knibbe, Jorge Goncalves, and Wafa Johal. 2025. Can you pass that tool?: Implications of Indirect Speech in Physical Human-Robot Collaboration. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–18. https://doi.org/10.1145/3706598.3713780