

# AI representing personas representing user groups: Applying the agency theory to examine interaction challenges of conversational personas as decision-making tools

Joni Salminen <sup>a</sup>, Soon-Gyo Jung <sup>b</sup>, Ilkka Kaate <sup>c</sup>, Trang Thi Thu Xuan <sup>a,d</sup>,  
Jinan Y. Azem <sup>b</sup>, Kholoud Khalil Aldous <sup>b</sup>, Danial Amin <sup>a,\*</sup>, Bernard J. Jansen <sup>b</sup>

<sup>a</sup> University of Vaasa, Vaasa, Finland

<sup>b</sup> Qatar Computing Research Institute, Hamad Bin Khalifa University, Doha, Qatar

<sup>c</sup> University of Turku, Turku, Finland

<sup>d</sup> Swinburne Vietnam, FPT University, Hanoi, Vietnam

## ARTICLE INFO

### Keywords:

AI-generated personas  
Interaction challenges  
Decision making  
User study

## ABSTRACT

The proliferation of artificial intelligence (AI) technologies has led to the rise of conversational decision-making support systems, such as dialogue persona systems that provide conversational access to various user segments. For example, product managers can ask personas about features before implementing them, politicians can learn about the needs of local communities through personas, and so on. Nascent research has looked at challenges when users interact with AI personas, but has not framed it as a principal-agent problem, in which the AI represents a persona that itself represents real people in the data. This setting exposes unique interaction challenges that decision makers face when engaging with AI-generated conversational personas, which we examine through a user study with 56 participants using AI-generated conversational personas. Our results indicate seven interaction challenges: (1) Hidden Information, (2) Hidden Personas, (3) Hidden UI, (4) Lack of AI Agency, (5) AI's Selective Attention, (6) Confusing Distributional Information, and (7) Conversational Cold Start that we conceptually link with agency theory. We discuss how the interaction challenges could be alleviated and suggest directions for future work.

## 1. Introduction

User personas (henceforth, ‘personas’) support decision making in user-centered design (UCD) and human–computer interaction (HCI) as empathetic representations of real user groups<sup>1</sup> [1]. Personas enable designers and other key decision makers to better empathize with users and understand their needs, preferences, and pain points, thus promoting UCD [1]. Personas are applied in many decision-making contexts, including software development, e-commerce, education, and healthcare, to evoke empathy for users and guide decisions, such as prioritizing feature development, shaping interaction design, and choosing target audiences. Better user understanding, in theory, leads to better decisions, which, in turn, are expected to deliver better outcomes for product users and companies developing the products [1]. Even though personas are fictitious user representations, they represent real user groups, which makes them relevant instruments for decision-making support.

Increasingly, persona creation leverages artificial intelligence (AI) technologies [2,3], *AI-generated conversational personas* answer information requests from decision makers and provide their answers based on real, segmented user data [3–5], rather than hallucinated or “synthetic” data. Therefore, conversing with personas is data-driven yet empathetic [6]; it is as if talking to a representative of the target user population. The main difference between AI-generated conversational personas and general-purpose AI chat systems like ChatGPT is that the former are created from primary user data relevant to the decision-making context, while the latter are not. Reliance on primary data about people, along with domain expertise [7,8], is appropriate for a decision support system (DSS) when those decisions affect people.

Personas have been conventionally presented as profile personas, and challenges with profile personas are well known — personas are considered abstract, stereotypical, and misleading [9] — but it is not well known which challenges apply to conversational personas in

\* Corresponding author.

E-mail address: [danialam@uwasa.fi](mailto:danialam@uwasa.fi) (D. Amin).

<sup>1</sup> Readers occasionally confuse personas with *avatars*. For clarification, avatars are used for self-representation [11], while personas are used to represent others (particularly groups of people).

decision-making tasks. More generally, human-like user representation has been somewhat studied in DSS literature [10], but not to a great extent. Fehrenbacher et al. [11] studied avatar representation and found that knowledge sharing increases when colleagues are represented by photographs rather than digital avatars. Interaction quality has been observed to influence how decision makers approach conversational tools. For example, Grimes et al. [12] found that better conversational capabilities improved users' perceptions of a conversational agent, suggesting that interaction quality plays a crucial role in decision-making when using conversational technologies. However, these challenges of AI-generated conversational personas are not well known, which calls for empirical research.

Unlike general-purpose conversational AI (e.g., ChatGPT) that are anthropomorphic systems representing neither real individuals or groups of people, conversational personas are specifically designed to represent real people for decision-making purposes, creating unique interaction expectations and challenges that require dedicated investigation beyond general-purpose chatbot research. Furthermore, the current research is timely because decision sciences AI technologies are fundamentally influencing virtually every DSS available in the market [13]. As noted by Storey et al. [13], "The field of decision sciences is undergoing significant disruption and reinvention because of rapid advances in [AI] technologies and the design of complex human-artificial intelligence systems." Generative AI (GenAI) and large language models (LLMs), such as ChatGPT, have gained immense research interest. These LLMs can generate coherent, contextually relevant text, answer questions, and converse with users, supporting decision-making in a range of contexts [13,14]. The evolution of LLM technology has led to *conversational DSSs*, of which AI-generated conversational personas are one type. Conversational personas can support decision makers who need to understand groups of people to make crucial business, development, or other decisions. This "understanding" forms through conversations with personas representing the groups of people affected by the decisions.

Nascent research has begun to explore the effectiveness and implications of AI-generated conversational personas in practical decision-making scenarios [15–17]. However, many questions remain about how AI-generated conversational personas can support decision making. *What is the optimal way to present personas to decision makers? Are persona profiles "better" than conversational personas? What are the main interaction challenges with conversational personas?* These questions motivate the current work.

Traditionally, persona creation has been highly dependent on human involvement, with associated challenges such as subjectivity, limited replicability and scalability, and long development times [6]. To overcome these challenges, current research employs Survey2Persona (S2P), an *interactive persona system* that (1) uses LLMs in the persona generation process and then (2) makes the personas available for decision makers to interact with (see Fig. 1), employed in prior research [18,19].

S2P outputs AI-generated personas, which are personas created fully or partially by leveraging AI technologies [2]. AI-generated personas represent an evolved version of data-driven personas, presenting user information in forms that are more easily understood and remembered than "nameless, faceless data" by stakeholders making decisions individually or in teams. In turn, AI-generated conversational personas *provide an intuitive user interface (UI) to crucial information for user-centered decision making*.

Although AI-generated conversational personas offer potential benefits, challenges arise from the limitations of AI technology and from the interactions of decision makers with these systems. As aptly put by Rapp and colleagues [20] "*we do not have a comprehensive overview of what has been investigated so far from the human side of human-chatbot interaction, i.e., what people expect, feel and, more in general, experience when they 'encounter' a chatbot*" (p. 1). Some of the challenges in the persona context are likely to originate from how chat works as an

interaction technique, from how LLMs work in general, and from the fact that *persona is a special case of chat interaction*. In particular, we propose that the principal-agent theory is a useful frame of reference for interpreting and conceptualizing the challenges of interaction between conversational personas and decision makers. In brief, the principal-agent theory (also known as the agency theory [21]) examines challenges when one party (principal) delegates tasks to another party (agent) whose interests may not align, involving information asymmetry in which the agent knows more than the principal [22,23]; the risk of opportunism<sup>2</sup> arising from this setting is known as the principal-agent dilemma [21].

It is vital to note that personas have typically been communicated through profiles [1], making the persona profile a predominant interface for persona design [24]. However, *interaction with a conversational persona differs from interacting with a persona profile*, as it requires different forms of (1) inputs (*typing vs. gazing*), (2) information processing (*proactive vs. reactive*), (3) presentation (*low- vs. high-graphic UI*), and (4) mental model formation (*gestalt vs. piecewise impressions*). The question of "optimal interface" matters because the interface can both support and hinder decision makers' ability to extract information from personas. AI-generated conversational personas involve distinct interaction dynamics, presenting both challenges and opportunities from a UCD perspective, which persona-based DSS developers urgently need to be informed about. To that end, we pose the following research question: *What kind of challenges do users face when interacting with AI-generated conversational personas?*

To address our research question, we conducted an on-site user study with 56 participants. Participants interacted with AI-generated conversational personas generated from a survey dataset. We conduct a qualitative analysis of the collected data using the think-aloud technique and the principal-agent theory as an analytical framework to interpret the findings. This framing is suitable for the current context because, unlike general chatbots where decision makers acknowledge that they are interacting with AI, conversational personas represent real people authentically, which creates a challenge in which the LLM must faithfully represent both individual persona characteristics and underlying user data, while decision makers rely on this representation to make decisions about real groups of people the personas represent. Overall, the findings can help DSS developers identify key interaction challenges to address.

## 2. Related work

### 2.1. Conversational personas as decision support systems

In general, DSSs are a form of information systems that combine data, analytical models, and UIs to help decision makers make more informed decisions [25]. By extension, a *conversational DSS (also known as a dialogue-based system)* is a computer-based system that uses natural language dialogue interfaces to help users analyze data, explore options, and make informed decisions through interactive conversations. A persona-based conversational DSS is one type of conversational DSS, although researchers use overlapping terms to describe personas generated with computer programs (see Table 1). *Data-driven persona* means that the persona is created using real data<sup>3</sup> about end-users. In our study, we use the term *AI-generated (conversational) persona* to refer to personas generated with AI and interacted with via a conversational interface (i.e., chat).

DSSs that use AI-generated personas have a significant history. Personas started as a manual analysis of end-user data to facilitate decision

<sup>2</sup> Defined as the agent acting against the interests of the principal.

<sup>3</sup> This data does not necessarily need to be quantitative; also qualitative data about end-users yields personas based on real experiences. However, the concept is often used in quantitative studies.

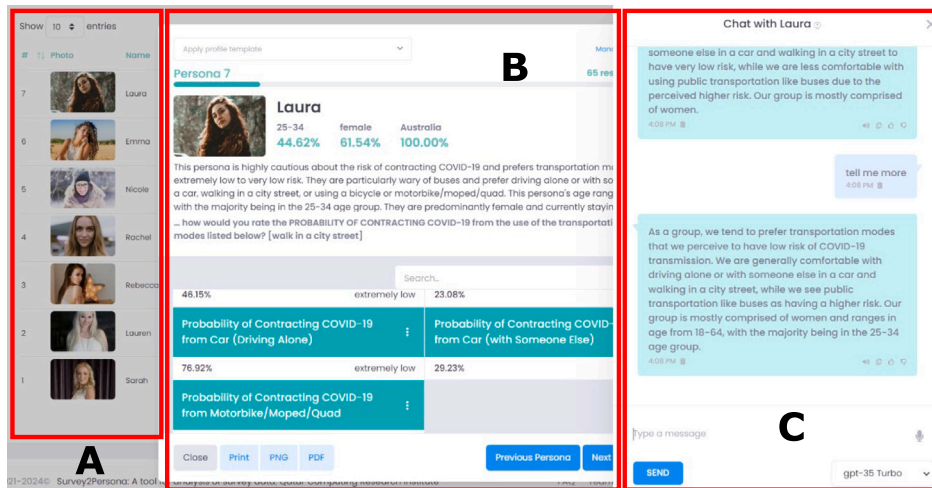


Fig. 1. The Survey2Persona UI. Decision makers can select AI-generated personas from the list (A). A profile version of the persona is on the left (B), and the conversational version is on the right (C). While decision makers can seamlessly switch between the persona types, it is not clear what interaction challenges different persona formats are associated with.

Table 1  
Evolution of persona concepts over time.

Persona concept	Year	Definition
User Persona	1999	A fictitious person representing a real group of people sharing certain attributes [1].
Quantitative Persona	2006	A persona created from quantitative data (e.g., survey) [26].
Data-driven Persona	2008	A persona based on any data about people [27].
Automatic Persona Generation	2016	Using software to automate the persona-creation process, including automatic collection and analysis of data using algorithms (also referred to as algorithmically generated persona [28]).
AI-generated Persona	2024	A persona created using AI technologies; traditionally, “AI” referred to algorithms such as clustering analysis, but it is increasingly used to refer to GenAI [17] (also referred to as LLM-generated persona [2] when using LLMs for analysis and/or write-up).

making in software development [1], later evolving into data-driven personas that leverage data science algorithms [27], and then into automatic persona generation (APG) where application programming interfaces (APIs) provided end-user statistics to create and update personas in near real-time [28]. These personas were typically presented as static profiles [24]. However, since the introduction of ChatGPT, persona creation tasks are increasingly being delegated from humans to LLMs [2,29]. While APG often relies on numerical data [28], LLMs can analyze both structured (e.g., numbers) and unstructured (e.g., text) data, which increases the scope of persona generation. Furthermore, compared with previous models, LLMs can participate more holistically in the persona generation process than previous computational techniques. For example, computational techniques often using clustering in persona creation, and the traditional approach was to ask humans to name the clusters [28], but this labeling can now be conducted using LLMs [3]. Arguably, the most significant change is the shift from static profiles to conversational personas, which enables real-time interaction between decision makers and personas.

Nascent empirical work on using AI to develop personas for decision making suggests that (a) AI-generated personas match the quality of human-created personas and (b) (at least some) tasks in the persona-creation process could be delegated to an AI. For example, Schmidt [30] shows that LLMs generate realistic personas from real survey data, Schuller et al. [31] found that LLM-generated personas matched the acceptance rates of human-created personas, and De Paoli [32] concluded that LLMs can generate rich persona narratives from user interviews. However, researchers contend that LLMs should be given actual end-user data, or the risk of misinterpretations and stereotypes crucially

increases [16]. Furthermore, Shin et al. [17] argued that AI cannot handle all persona-creation tasks, especially those that require prioritization and judgment. So, despite the potential, the boundaries of task delegation to AI in persona generation remain unclear.

Overall, conversational persona systems are introduced as alternative DSS tools for decision making about people in critical domains, including business, content creation, and presenting marginalized perspectives in society. Many of these systems can be applied to a wide range of decision-making situations, but their evaluation is often lacking. Previous research has predominantly focused on persona quality dimensions such as credibility and consistency [17,32], while neglecting user interaction challenges and their implications for decision support efficacy, which this study addresses.

2.2. Interaction challenges in conversational decision support systems

Although there are many potential interaction challenges in chat-based decision-making systems (see a systematic review by Rapp et al. [20]), we summarize here key considerations from a decision-making perspective. First, chatbots often struggle to accurately interpret user intent and context, leading to misunderstandings and frustration [33]. In the pre-LLM era, chatbots often relied on keyword-based methods that yielded sentence outputs that could not anticipate each request or use of language [34]. A lack of genuine understanding of human language resulted in an unsatisfactory user experience: “most chatbots were not ready to handle the complexities of human conversation.” [20] (p. 2).

Second, chatbots can generate plausible but incorrect information, a trait known as “hallucination” [35]. This is particularly problematic for human–AI interaction because LLMs’ tone tends to be factual, believable, and assertive — it is speaking as if it knew. So, users can struggle to distinguish fact from fiction. Also, LLMs may perpetuate biases in their source data, potentially leading to unfair decisions. Third, chatbots before LLMs struggled to maintain extended conversations, leading to irrelevant responses [36]. This lack of memory has been a persistent challenge, as it makes the interaction feel inconsistent and confusing to the decision maker, reducing interaction effectiveness [37].

Fourth, chatbots — unless programmed — have no social sensibility or emotional intelligence [17]. This is why many chatbots lack the ability to recognize and respond appropriately to user emotions, sometimes appearing “robotic” (monotonic, formulaic, repetitive, non-natural, etc.), which can lead to unsatisfying interactions [38]. While researchers attempt to make AI more human-like, this has been difficult with traditional chatbots and often led to “mirages” (inadvertent effects, e.g., creepiness or lack of social presence). Although somewhat improved, this challenge remains with current AI-generated chatbots.

Fifth, users often have high expectations about the capabilities of the chatbot, leading to disappointment when the chatbot fails to understand or perform as expected [39]. LLM-based chatbots often lack real-world knowledge and can struggle with tasks that require up-to-date information or common-sense reasoning. From a user interaction perspective, this is problematic because LLMs’ style of communication often mimics real experience [32] — e.g., it can refer to itself as “I” and it can present convincing scenarios with high fluency but necessarily no grounding in real human experience. Expectation violations undermine DSS effectiveness, because incongruence between user expectations and system capabilities leads to reduced decision support satisfaction and diminished decision-making benefits [12].

### 2.3. Interaction challenges of conversational personas as a special case

At the outset, it seems logical to presume that many interaction challenges from conversational technology in *general* cascade into decision making with conversational personas; after all, conversational personas are a special case of chatbot technology. Some of these cascading challenges relate to the technology itself (e.g., maintaining context across long discussions and sessions, transparency issues, etc.). Others relate to the interaction design of chatbots (e.g., communication style, speed, and interface presets [34]). So, *why would we expect any special interaction challenges to emerge when the chatbot assumes the role of a persona?* Answering this question, we maintain that a *conversational persona differs from other chatbots in that it represents a particular end-user segment*. This means that the conversational persona needs to first interpret the user’s message to the persona, then extract the relevant information from the end-user data, and articulate a response back to the user in the persona’s “voice”. Conceptually, this representative delegation can be understood as a special case of the principal–agent dilemma.

Although AI-generated conversational personas share similarities with other types of chat systems (e.g., the role of anthropomorphism [38]), this configuration of the *AI standing between the user and the persona* creates distinct user requirements and challenges. Particularly, users expect to engage with the persona as they would with another human being [40]. This core tenet sets special requirements for interaction episodes; as in many other types of chatbot interaction, users know that they are talking to a chatbot and can anticipate and tolerate non-human-like reactions.<sup>4</sup> Although AI-generated conversational

personas are a special use case for conversational interaction, there is a scarcity of research on interaction challenges with these personas. Therefore, we address a distinct form of interaction within conversational decision-making systems. The scarcity of research on interaction processes between decision makers and conversational personas creates a research gap for investigating the specific challenges that arise during these interactions.

## 3. Overview of Survey2Persona

### 3.1. Persona generation procedure

This section presents a systematic approach to creating personas using survey data (specifically, *data about people obtained through a questionnaire*) by leveraging LLM capabilities. Survey data is highly relevant for persona generation, as the literature shows that surveys are the most frequently used data type for persona creation [41]. Survey data typically represents end-users’ opinions, attitudes, preferences, beliefs, demographics, and other variables. The S2P methodology comprises four main stages [3]:

- **Stage 1: Preparing the dataset.** The survey items’ categories are quantified to prepare the data for clustering (e.g., “Strongly agree” = 5 or 7, depending on the number of unique values in the scale)
- **Stage 2: Mapping question types.** The pre-processing stage transforms the dataset, allowing for topical classification, sentiment analysis, and LLM-based survey question grouping and condensation.
- **Stage 3: Data preprocessing.** The persona generation phase employs clustering methods and a statistical test (Kruskal–Wallis) for persona profile.
- **Stage 4: Persona generation.** The final stage enriches the personas, turning them into completed representations of the real respondent groups identified in the previous step. The personas can be interacted with using either the static profile or the interactive conversational interface (see Fig. 1).

Five LLM tasks are integrated into S2P:

- **Task 1: Data mapping (used in Stage 1.)** This task assigns numerical values to responses on the Likert scale. This is a typical task performed by human researchers converting Likert categories (e.g., “Strongly agree”) to numbers (e.g., 5) to make the data quantitatively analyzable. The mapping is done by identifying unique category values and then asking the LLM to assign a numerical value that corresponds to the ordinal structure of the categories.
- **Task 2: Information grouping (used in Stage 2.)** Survey questions are categorized into topics using the LLM. Grouping prevents persona profiles from becoming overwhelming, especially extensive surveys. Grouping aims to provide clarity and structure without losing important details.
- **Task 3: Topic identification (used in Stage 3.)** The LLM identifies potential topics or themes present in the open-ended responses of participants to the survey question. This is similar to the qualitative analysis of open-ended survey data by human researchers.
- **Task 4: Persona information write-up (used in Stage 4.)** The LLM generates description texts for the persona based on survey questions, responses, and their relative frequencies. The description texts provide contextual information for the user.

<sup>4</sup> The case of persona chatbots bears some similarity in social experiences AI applications like virtual friends or companions in that those interactions are also expected to echo the sense of presence from another human being. The distinction is that such applications are for entertainment and personal

purposes, whereas the use case for personas is professional (work, productivity tasks, “serious” decisions). Decision makers have jobs to be done (quickly and efficiently) [34] in which personas act as instruments for accessing information needed to form a realistic user perspective.

- **Task 5: Persona information labeling (used in Stage 4).** This task involves generating heading and description texts for persona information. The LLM creates keyword-based headlines and descriptions using the provided survey question texts and corresponding response counts. The heading captures the behavioral patterns in the responses; it also emphasizes the majority of responses.

The performance of S2P on each task has been evaluated in prior work [3], yielding satisfactory results. Technically, S2P leverages OpenAI's GPT models via the Microsoft Azure platform.<sup>5</sup> At the time of the study, we used the GPT-4 and GPT-3.5-turbo models, which provided a balanced combination of quality, speed, and cost-effectiveness (these models were later replaced by GPT-4o in our implementation; here we report the case at the time of study completion). We used the '2023-03-15-preview' version of these models with a temperature setting of 0.7 (determined through manual experimentation). The GPT-4 model was used for **Tasks 3** and **4**), which involve more complexity; other tasks were completed by the GPT-3.5-turbo. We leave testing other LLMs (e.g., LLama, Gemini, Claude) for future work.

We use S2P [3] in this study because it aligns with our research objective of exploring AI-generated personas. S2P is a fully operational DSS and integrated framework with standard persona profiles and AI persona with conversational capabilities, both created from the same underlying survey data. This combination makes S2P particularly suitable for the comparative analysis required by our research. S2P's methodological pipeline combines clustering, LLM-based summarization, and retrieval augmented generation (RAG), facilitating empirically grounded AI persona responses that minimize hallucination. Also, S2P's interface design of both static and conversational personas allows system end users to easily switch between static (i.e., profile) and dynamic (i.e., AI) modes of persona interaction. Therefore, S2P is an appropriate system for our research studying how decision makers experience distinct persona modalities.

### 3.2. User interaction with the persona

Users select a persona to converse with. As the user sends messages, S2P converts each message into a vector embedding and identifies the most relevant survey items by comparing this embedding to pre-computed survey item embeddings using cosine similarity. The system then retrieves response frequencies for the most relevant survey items and appends this data to the user's message. A context is constructed, including the system prompt, the augmented user message, and a limited number of previous messages, also augmented with survey data. This context is sent to an LLM, which generates a response assuming the role of the persona. The response is then displayed to the user, aiming to maintain the perception of interacting with a consistent persona. Decision makers can thus interactively explore the collective perspectives represented by the survey data, while the underlying technical components ground the responses in the actual survey results (see Fig. 2). The key design principle here is that the dialogue with the AI-generated persona relies on data specifically extracted from the persona's information: the LLM does not "invent" or "hallucinate" the answers, but rather formulates the statistical data into a fluent conversation by taking on the persona's role. This aligns with the foundation of DSSs relying on real information about people.

## 4. Study design

### 4.1. Creating the AI-generated personas

S2P allows users to upload their survey data and seamlessly generate personas following the outlined methodology. To evaluate the

approach and its implications for user interaction, we utilize a publicly available dataset from the Pew Research Center (*American Trends Panel Wave 99: Artificial Intelligence (AI) and Human Enhancement*<sup>6</sup>). The dataset includes responses from 10,260 adults in the United States between November 1 and 7, 2021. The survey has 55 questions about people's opinions on AI and human enhancement technologies. Applying S2P, the personas describe different respondent groups within the surveyed population, in this case, reflecting people's AI-related attitudes.

### 4.2. On-site user study

#### 4.2.1. Participants

The on-site user study was conducted at two research sites at a major technology-focused university. Participants were recruited via email (both group and individual). The study involved 56 participants. The participants were predominantly researchers (n = 38, 67.9%) and engineers (n = 9, 16.1%). The other participants included a mix of professions: one lab coordinator, one security and health director, one bioinformatics specialist, five undergraduate students, and one business development manager (n = 9, 16.1%). The average age of participants was 33.6 years (SD = 10.8). The sample was roughly evenly distributed between male (n = 32, 57.1%) and female (n = 24, 42.9%) participants. Most of the participants either had a conceptual understanding of personas or had used personas before (n = 34, 60.7%), while the rest (39.3%) had no prior knowledge of personas (however, all participants were briefed on the concept of personas, as we explain in the following subsection). Almost all participants had used AI chatbots (n = 48, 85.7%), while the remaining participants had a conceptual understanding of AI chatbots. The fact that most participants had prior knowledge or experience with personas indicates they could credibly engage with persona-based decision-making regardless of their job title. Also, as researchers and engineers working with technology, most participants were knowledgeable about AI, and a couple of them even had AI-focused startups, so the fictitious scenario (see the following subsection) was compatible with their domain expertise. While our sample somewhat deviates from typical software engineering or management roles specified in our task scenario, the participants nonetheless represent technology-oriented decision makers who would likely be adopters of AI-generated persona systems. Researchers and engineers generally possess the analytical skills and technical literacy to credibly evaluate AI-generated personas for decision-making purposes.

#### 4.2.2. Procedure

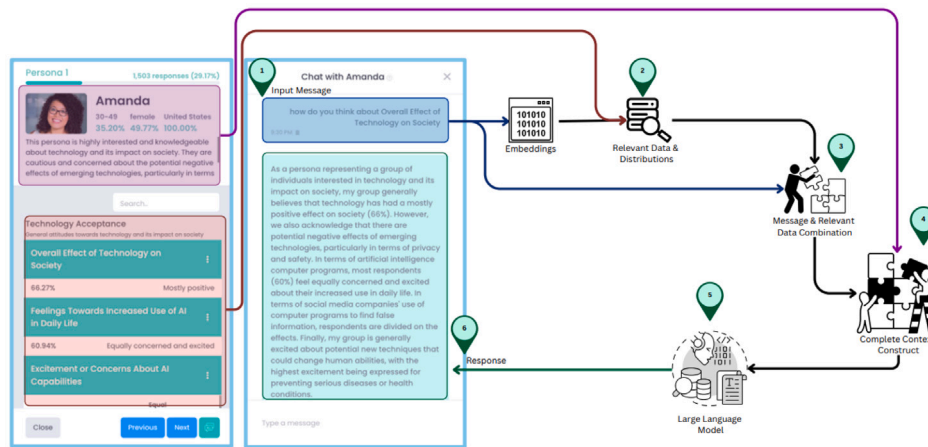
Five study administrators took turns conducting the study sessions in two separate study rooms that had an identical setup (the same set of laptops, monitors, and control devices: Windows 11 laptop, two 24" displays, keyboard, and external mouse).

First, the participants were briefed on the general study procedure and asked to provide their consent to participate. Second, the participants were inquired about their familiarity with the concept of personas, and regardless of their answers, they were provided with an operational definition to ensure that each participant understood the concept of persona well enough to use personas for the task completion; from our observations, we found no issues in this regard, but all participants could grasp the persona concept as we intended. Third, participants were presented with a work task scenario (WTS) requiring them to interact with a persona. This WTS detailed a hypothetical scenario in which participants were assigned the role of either a software engineer or a manager at a company specializing in AI solutions.

The decision making task required the participants to interact with the persona and then be asked to (a) describe the persona in their own

<sup>5</sup> <https://azure.microsoft.com/en-us/products/ai-services/openai-service>

<sup>6</sup> <https://www.pewresearch.org/internet/dataset/american-trends-panel-wave-99/>



**Fig. 2.** AI-generated persona chat in S2P. (1) The user writes a message to the persona. (2) S2P first identifies relevant survey items via semantic matching between the user’s message and the survey items. Then it extracts the distribution of answers from respondents who belong to the persona. (3) S2P adds this information to a (4) general description of the persona in the user’s message, which (5) it passes to an LLM, essentially instructing the LLM to respond to the message based on the known information from the persona while maintaining the role of the persona. The (6) LLM does this, and the response message is relayed back to the user in the chat UI.

words based on the information that they obtained and (b) decide if the persona would be a good target group for the company’s AI solutions.

Fourth, in the within-subjects design, each participant examined one profile persona (Mark or Linda) and one conversational persona (Linda or Mark, respectively). After examining the first persona (either chat or profile, randomly assigned), the participant was guided to complete a short survey about their persona perceptions and the WTS task. The participant then examined the second persona with a different interface than before (if the participant first used a profile, they then used a chat, and *vice versa*), and completed the same survey. After this, the participant was thanked and rewarded with a gift card with a value of \$30 USD. Using the survey think-aloud approach [42], participants were asked to think aloud throughout the entire study session, including while using the persona and completing the surveys.

Study sessions were recorded and transcribed. All transcripts were merged into a single file, structured by participant and persona interaction type (e.g., “P01-C” indicates Participant 1 using a conversational persona). Also, all field notes from each administrator were merged into one file.

#### 4.2.3. Data analysis

Two researchers independently read and analyzed the transcripts and collective field notes from sessions with 56 participants. Both researchers were part of conducting the studies, including administering the studies, giving analysts an in-depth understanding of the study topic and the participants.

The analysis had four stages. **Stage 1** involved independent coding of the text material (transcripts and field notes) using a spreadsheet to note down interaction challenges. The coding guidelines instructed the researchers to identify interaction challenges, defined as any concerns, problems, or issues participants expressed during their interaction with the personas. The coding process involved extracting verbatim quotes and assigning descriptive labels that characterized each problem at a more abstract or general level, which corresponds to the open coding technique [43]. The guidelines also allowed for noting interaction challenges with profile personas when they provided meaningful contrast with conversational persona interactions.

After the independent coding, in **Stage 2**, the researchers presented their findings in a meeting, which involved collaborative working with the spreadsheet while forming a final taxonomy of the interaction challenges that (a) captured the central elements of each researcher’s coding and (b) contained minimal overlap between categories. Through this collaborative coding effort, the researchers agreed on a set of

seven challenges with AI-generated conversational personas. So, in this stage, the researchers collaboratively established a synthesis of the interaction challenges (see online supplementary material for the codes<sup>7</sup>). The process naturally involves a degree of subjectivity, because meaning construction takes place in researchers’ minds; however, as the researchers who coded also ran the user study sessions, we believe we provide accurate descriptions of the real interaction challenges participants faced.

In **Stage 3**, another round of reading was conducted by the more senior researcher to find representative example passages in the transcripts that either *corroborated*, *contradicted*, or *complemented* the categories in the taxonomy. This was done to provide further context that could enrich the findings. The reading at this stage was based on skimming the transcripts and did not yield significant modifications to the list of interaction challenges or their contents. The process was iterative and relied on the senior researcher’s judgment and their formation of the “big picture” of the results.

After this, in **Stage 4**, the senior researcher wrote up the results, using identified quotes and all moderators’ field notes as supporting material. This part of the process refined the definitions of the challenges to articulate them more clearly; however, their content was not significantly altered, and the meaning of each challenge remained within the boundaries of what was commonly agreed upon by the researchers. All co-authors provided feedback on the write-up, including edits and comments; again, these actions did not alter the meaning of the agreed upon list of challenges. Presumably, the editing stage drew on the researchers’ pre-existing knowledge of human–AI interaction, which all the research team members had.

## 5. Findings

### 5.1. Statistical findings on the usability of AI-generated personas

A survey was administered after a participant had used each persona, asking the participant about *usability*, which assessed mental effort required, ease of use, and clarity of the persona system; *willingness to use the persona*, which measured preference for using the persona, perceived utility, and impact on decision-making ability; *enjoyment of using the persona*; as well as *task difficulty* and *task confidence* in the

<sup>7</sup> [https://osf.io/tsuw/overview?view\\_only=28b0392af6b14d8a87e431e447241f7e](https://osf.io/tsuw/overview?view_only=28b0392af6b14d8a87e431e447241f7e)




Fig. 3. User perception differences between profile and conversational personas. The results imply these two formats for user representation are perceived differently in decision-making tasks. The reader is encouraged to zoom in to improve the legibility of the figures.

accuracy of their responses. Regression was carried out using data from 54 participants, excluding two male participants aged 55 and 38 who had missing responses to the survey questions. The results reveal that the conversational personas were perceived as more usable (B = -1.423, p < .001), enjoyable (B = -1.502, p < .001), and less difficult to use (B = 1.778, p < .001). However, for task confidence (B = -0.186, p = 0.157) or willingness to use (B = -0.350, p = 0.127), there was no significant difference. Fig. 3 illustrates the differences. Related work [44] provides additional statistical analyses of these variables, whereas the current work focuses specifically on the interaction challenges posed by AI-generated personas.

Generally, participants preferred interaction with the conversational persona (“It’s like ChatGPT, it makes my task easier” (P45); “This is super freaking easy because I can very quickly find the answers I think.” (P02)). The conversational persona was able to sustain meaningful conversations and the usability of conversational personas was perceived higher. Thus, the conversational persona had certain advantages over the profile persona, though these came at a cost, as discussed in the following subsection. Our findings articulate that each interface had its unique interaction challenges.

## 5.2. Interaction challenges in AI-generated conversational personas

Based on our analysis of the think-aloud transcripts and field notes, following the procedure outlined in Section 4.2.3, we identify challenges in user interaction (CUI) with the personas (see user quotes mentioning the challenge in Table 2). We also interpret each challenge’s connection using the principal-agent theory as a framework. The seven CUIs are summarized as follows:

- **CUI01 — Hidden Information:** Determining the optimal amount of information that conversational personas should reveal, as users face a trade-off between having immediate access to comprehensive persona data versus needing to actively query for specific information through conversation.
- **CUI02 — Hidden Persona:** Users feel disconnected from the actual persona because the users perceive themselves as talking to an intermediary or representative rather than directly to the persona itself, creating an “invisible wall” that reduces social presence and authenticity.
- **CUI03 — Hidden UI:** The minimalist conversational interface lacks visual design elements like color-coding, highlighting, and interactive components, making it harder for users to navigate information and understand system affordances compared to traditional profile interfaces.
- **CUI04 — Lack of AI Agency:** Conversational personas appear too neutral and passive, lacking strong opinions or the ability to be influenced during conversation, which makes interactions feel artificial and reduces the human-like qualities that make personas effective.
- **CUI05 — AI’s Selective Attention:** Conversational personas may misinterpret user queries, provide off-topic information,

or present inconsistent responses, leading to inefficient information retrieval and potential confusion about the persona’s characteristics.

- **CUI06 — Coping with Distributional Information:** Presenting within-group variability and statistical data in a way that maintains the persona’s human-like qualities while providing accurate representations of the diversity within the user group the persona represents.
- **CUI07 — Conversational Cold Start:** Users struggle to initiate meaningful conversations with personas due to the lack of guidance, context, or starting prompts, similar to facing a “blank page” problem when they do not know what questions to ask.

### 5.2.1. CUI01 — poor information, rich information...hidden information

The participants had mixed feelings about the presentation of information in the AI-generated personas. Some felt that conversational personas lack complete information about end-users the conversational personas represent (Quote[1.1]), especially relative to the profile persona: (Quote[1.2]). The conversational interface encourages more active engagement but requires greater cognitive effort, while the profile interface allows for easier but potentially more superficial information extraction.

Information design is an accessibility consideration (Quote[1.3]). The conventional profile allows users to quickly scan and identify relevant information without having to formulate specific queries. So, even though one might think a conversational interface is fast, the profile can be even faster, especially when forming an impression of the persona as a whole, which is a crucial part of achieving a useful mental model of the persona. In turn, conversation can be faster when locating *specific information* about the persona. It appears that there is a trade-off between Gestalt (an overall impression) and details (specific information requests) when interacting with an AI-generated conversational persona (see Quote[1.4]). So, *conversation is a more challenging interaction technique to form a first impression of the persona.* (As the conversation with the persona progresses, the user can “catch up” on the formation of the mental model.)

However, information from the conversational persona was easier to digest relative to persona profiles (Quote[1.5]). The profiles contained relatively a lot of information, which made it difficult for some participants to cope with: (Quote[1.6]); (Quote[1.7]). This effect of cognitive overload was not present in conversational personas, although information retrieval from conversational personas required more effort to initiate (in the form of thinking what to ask) and maintain (Quote[1.8]). When using the conversational persona, participants need to apply *laddering techniques* to probe deeper into the persona’s information, whereas the profile has the information readily available (Quote[1.9]).

Another perspective is that hidden information is a benefit of the conversational personas. Although the profile provides more information for forming first impressions, it is constrained by its information capacity; only relatively little can be included before it becomes cluttered and difficult to manage cognitively. In contrast, a conversational

**Table 2**  
 Example quotes illustrating how participants experienced the challenges. Frequencies indicate how many participants faced the challenge.

Quote #	Quote
CUI01 — Hidden Information (n = 23, 44.2%)	
Q1.1	“(conversational persona) doesn’t give you the full idea” (P07)
Q1.2	“I think if I ask more, the first option (conversational persona) (would) present me all the details (...). The (conversational persona) is good in interaction. I could ask more and all but the second one (profile persona) just gave me the information. It’s easy to look at (...).” (P37)
Q1.3	“for (conversational persona), you need to have certain questions to ask it right. And only then you can get all the data but (in profile persona) everything is already laid out so you can get a quick overview.” (P44)
Q1.4	“The profile was like the data. Yeah, it gave me a better understanding of what the person I was dealing with like. There was other information that I realized I didn’t use. But in the back of my head, I think I did process. So when I was answering the questions, I knew more about the person. Whereas with the chatbot it was giving me information specific to that one question I ask.” (P39)
Q1.5	“This text here (profile persona) is a little bit difficult to read because it’s like not in the rows are quite lengthy and then you have to scroll. (...) difficult to read this.” (P04)
Q1.6	“It’s difficult to find the answer (in the profile persona). We need to scroll either to scroll searching. I think the searching is based on the topics or things.” (P10)
Q1.7	“It (profile persona) was like a data sheet” (P45)
Q1.8	“Actually I prefer in term of content, (I prefer) the (profile persona), but in terms of easy to use the (conversational persona)...the (profile persona), I like the way it give me all the information because the (conversational persona) it did not help make decisions. It gave me information, but it’s still I need to ask more than one time and it does not give all the presentation.” (P37)
Q1.9	“I already have the information, I didn’t have to ask about it. It’s just all in front of me.” (P56)
CUI02 — Hidden Persona (n = 33, 63.5%)	
Q2.1	“I feel like, OK, there is this profile somewhere there. The real person is like somewhere hidden” (P04)
Q2.2	“I felt it was an algorithm, (it felt like) Linda is trying to answer for Linda.” (P48)
Q2.3	“I felt like, OK, this is I’m not actually talking to the persona.” (P04)
Q2.4	“She’s a robot - she cannot feel” (P18)
Q2.5	“...these are basic things that a robot would point out” (P45)
Q2.6	“(When using the conversational persona,) I felt like, OK, I’m talking to you about and the bot knows about the persona, but the persona is somewhere like there in the distance, whereas in the profile, I actually felt a closer connection with the profile because the profile was kind of like, ‘OK, here is the persona’ and I can see the information also the persona.” (P04)
Q2.7	“It feels like I’m talking to a wall or something, and the answer is it’s like...I wouldn’t feel a connection with it. It just gave me basic, like straightforward answers. There’s no emotions, and the answers. There’s no feelings, no tone.” (P50)
Q2.8	“[the fact that the persona says] ‘as a persona, I think...’ sounds robotic” (P20)
Q2.9	“Yeah, (...), every time it would answer a question they would answer it like saying like OK, like the people in this group.” (P24)
Q2.10	“I didn’t feel like it was a person. I was talking to kind of a summary. It didn’t feel like a person, honestly.” (P21)
Q2.11	“What I didn’t like about the chat thing is that a person is presented here with Linda and then, when she is answering the question, she’s saying ‘our Persona group is split’, so it doesn’t feel personal.” (P36)
CUI03 — Hidden UI (n = 13, 25.0%)	
Q3.1	“As I said, the sectioning was like so helpful. Definitely the percentage wise and the coloring wise, the information about the persona was presented clearly.” (P35)
Q3.2	“I like the color coding in the survey, so it is strongly agrees it’s darker than the rest, so he disagrees or he opposing it will be a lighter color. Or the higher percentage is darker than the lower percentage.” (P36)
Q3.3	“Even for accessibility, while looking I want to scroll up to see the whole (information), so I feel difficult with the second interface (conversational persona), but the traditional one (profile persona) I felt more comfortable because I can all the things, even the color used in the traditional one, which helped me to find out the right answer.” (P43)
Q3.4	“I feel like maybe I should have asked more follow up questions or like I shouldn’t just have like copied and pasted the questions I should have maybe like talked a bit more to get a better idea.” (P17)

(continued on next page)

**Table 2** (continued).

Quote #	Quote
CUI04 — Lack of AI Agency (n = 25, 48.1%)	
Q4.1	“I didn’t like feel like the persona was talking to me. I didn’t think the persona said that ‘I feel like this and I think like this’” (P04)
Q4.2	“So in terms of answering the questions, chat was definitely easier [...]. But at the same time, it was artificial the, the response it gave was very artificial.” (P29)
Q4.3	“Here we can see that 72% of group prioritizes a great decision of a quick decision when it comes through. (...) So she cannot give an opinion, that’s for sure.” (P14)
Q4.4	“I mean, Linda actually does not identify herself within a certain category of people, because if I belong to certain category, I have an opinion, right? Yes. And I would state out what is my opinion, right? Yeah. But then here, she’s always on the neutral.” (P14)
Q4.5	“I see that Linda has does not have very strong personality and she’s not actually very courageous in giving her opinion, even though it’s not a critical question [that I asked].” (P14)
Q4.6	“I think that’s like generally if you use just app, it’s often the case like it’s very neutral because you don’t want to bias people’s decision after that. But this [conversational persona] is different, right? This is supposed to be different. This is different, so you would like this to be more opinionated.” (P14)
Q4.7	“She needs to be opinionated somehow, right? She needs to have her own opinion somehow, right?” (P14)
Q4.8	“I mean I’m engaging in a conversation with someone with these pictures, right. And then based on the like the back and forth, I mean I could decide whether we continue in this discussion and then we can go into other topics and if we are aligned on the basics or not. —Ah yes, because you want kind of see if the person is thinking like you or. —Yeah, not necessarily as me, but as I mean that our opinions construct each other, ok? So they help building something better. —So the discussion is evolving and you might change your opinion and they might change their opinion, right? She might convince me, right. And might. And I might convince her.” (P14)
Q4.9	“So she’s, like, refusing her identity as a persona.” (P14)
CUI05 — AI’s Selective Attention (n = 19, 36.5%)	
Q5.1	“So there are some discrepancies here as well, which would say. The majority of the respondents in this group do not find it acceptable, and the majority of the respondents, while it’s saying specifically 43% of responders found this idea not acceptable, while 55% found it acceptable, so the majority should become the ones which you know find it acceptable.” (P34)
Q5.2	“it’s giving me unrelated information” (P20)
Q5.3	“I notice a pattern. It’s always giving me additional information and it’s answering a different question.” (P20)
Q5.4	“I didn’t ask for the distribution, but (...) but he noticed more. It’s more (than what I asked for).” (P10)
Q5.5	“In a single person, all these contradictory sentences...” (P09)
Q5.6	“It’s weird. In the previous answer it didn’t mention this... it’s scary because uh, like you, you would trust it very easily.” (P10)
Q5.7	“The previous question was interesting because the contradictions and whatever the data doesn’t make sense. It’s actually what makes it human, because sometimes humans (...) might be having conflicting [opinions], yeah.” (P07)
CUI06 — Confusing Distributional Information (n = 28, 53.8%)	
Q6.1	“Now it just give me statistics. So when I see 81%, I’m intimidated” (P14)
Q6.2	Participant asked: “How do you feel about driverless vehicles?” Conversational persona responded: “As a group, we are skeptical about the use of driverless vehicles. 59% believe that widespread use of driverless passenger vehicles would be a bad idea for society and 81% of respondents would probably or definitely not want to ride in a driverless passenger vehicle.”
Q6.3	“Gender 52% male? Is that persona changing gender?” (P20)
Q6.4	“So in terms of describing the persona, I cannot describe the persona because it’s a big persona group that the chat is depending on to provide answers.” (P36)
Q6.5	“When there are these ‘somewhat agree’ [answers], it was kind of [confusing] like I had to ask it, ‘Are you definite or are you probably [thinking this]’, yeah. And then the chat would tell me ‘I am definitive’ or...” (P53)
Q6.6	“She’s born from survey data” (P18)
Q6.7	“The chat gave me not only the answer that was directly what the question related to, but also gave me a bit more contextual information which made it much more natural.” (P40)
Q6.8	“I like the fact that Mark had all like it had all the numbers for each question. And it had a list like it had like a range. So from no excitement to least excitement or sorry, no excitement to the most exciting information or no information to opposing or against.” (P50)
CUI07 — Conversational Cold Start (n = 8, 15.4%)	
Q7.1	“I feel like if I get a general summary at the beginning about that Persona would be much better for me to understand it before.” (P22)
Q7.2	“Maybe at the beginning have a virtual background (...) like, he said, ‘Hello, I’m your assistant’. Something like this.” (P35)

persona can retrieve any underlying information from the user data, which affords more freedom of discovery, but some decision makers may feel more anxiety due to the vast search space of possible information to query. As such, the ladder information in conversational personas poses both advantages and disadvantages. From the principal-agent perspective, information asymmetry arises when the AI agent (persona) controls access to comprehensive user data, forcing principals (decision makers) to actively query rather than having transparent access to all relevant information for their decisions.

### 5.2.2. CUI02 — hidden persona (invisible wall)

A major hindrance to the conversational persona was the “invisible wall” effect. This means that some participants felt they were talking not to the persona, but to someone trying to describe what the persona thinks. This led to frustrations in feeling that the real persona is somewhere “hidden” behind the conversational persona (Quote[2.1]; Quote[2.2]). This can also be described as the problem of two entities, in which the user feels like they are trying to reach one entity (the persona) but instead reaching another (a “representative” of the persona). The representative role creates an invisible wall between the user and the persona, which, in some cases, results in an emotional disconnect or difficulties in forming a mental model (Quote[2.3]; Quote[2.4]; Quote[2.5]).

Interestingly, these issues echo the “classic” challenge of chatbots appearing too robotic; this is interesting because LLMs *do* have the linguistic flexibility to appear non-robotic; yet, this property has not been properly obtained in the tested AI-generated conversational personas. One participant provides a clear comparison of how this effect can hinder user interaction relative to using the profile persona (Quote[2.6]). Interestingly, the profile interface created a stronger sense of presence in the persona, while the conversational interface introduced a perceived *distance*. This was not expected, as social presence is a (theoretical) advantage of a chatbot (i.e., you are supposedly talking to someone/something that is present!). This sentiment was also mentioned by another participant (Quote[2.7]).

The perceptual challenges stem from the persona’s communication style, which does not always come across as natural (Quote[2.8]; Quote[2.9]; Quote[2.10]). Another root cause of this behavior is that the persona is trying to convey variation within the group (see Quote[2.11]). There is a conflict between the user’s expectation of interacting with an individual persona and the reality of interfacing with an aggregate data representation, which relates to a trade-off between precise communication (giving the user more precise information) and user experience — it seems that in some cases, the higher precision worsens the interaction.

As the participants were not familiar with the data to a level where they could determine that the question is impossible to answer, the setting poses a classical principal-agent problem in which the user (principal) needs to trust the agent (AI-generated conversational persona) to give the correct answers. As the agent is more knowledgeable about the data, this setting involves a crucial degree of information asymmetry, which might result in the user making wrong decisions based on the information provided. From the principal-agent perspective, the representation chain creates a principal-agent problem where decision makers cannot directly access the “real” persona but must rely on an AI intermediary, leading to perceived disconnection and reduced trust in the agent’s authentic representation.

### 5.2.3. CUI03 — hidden UI (less navigability and control)

Similar to the information (and the persona) being perceived as hidden (i.e., non-visible), the same is true of the UI. Whereas the profile persona incorporated interactive elements, the conversational interface remains minimalist. This relates not only to buttons but also to factors such as color, structure, and highlighting of information. For example, these elements were found to be comparatively useful in the profile persona interface (emphases added) (Quote[3.1]). Coloring

in the profile helped participants identify more important information (cf. information saliency) (Quote[3.2]). Similar comments were made by another participant (Quote[3.3]). Visual elements in the persona profile seemed to better support the participant’s cognitive processes and information-seeking behaviors. A lot of the information conveyed through design elements was missing from the conversational interface. A key issue is that the lack of UI elements leads to the user’s incomplete perception of affordances, i.e., what the system can do or how it should be used (see Quote[3.4]). In the absence of UI design elements to guide the user, they may feel the interface does not sufficiently prompt or afford more in-depth conversations. From the principal-agent perspective, the minimalist conversational interface limits the decision maker’s ability to monitor and evaluate the AI agent’s performance, reducing transparency and making it harder to detect when the agent fails to provide optimal information presentation.

### 5.2.4. CUI04 — lack of AI agency

Agency refers to an agent’s ability to act independently, make its own choices, and exert influence over its environment. The lack of agency manifested itself in two ways: *Lack of Strong Opinions* and *Immovable Persona*. Overall, the lack of persona’s perceived agency can lead to artificially-feeling conversations (Quote[4.1]). More broadly, we observe an interesting trade-off between ease of use and perceived authenticity in AI interactions (Quote[4.2]). *Easy does not mean natural*.

Regarding the lack of strong opinions, some participants would have liked the persona to be more opinionated (see Quote[4.3]). In these cases, the persona was perceived as too neutral (Quote[4.4]); as someone who does not have a strong personality and is not very courageous (Quote[4.5]). Interestingly, this goes back to the conversational persona being a special case of chatbots; the user understands this and expects natural dialogue (Quote[4.6]). This neutral stance is undoubtedly at least partly inherited from the underlying LLM whose (human-programmed) safeguards tend to make it cautious and not take a stand in any direction. This is somewhat limited when the persona *should* represent people closer to extremes (e.g., in the political spectrum — or in any domain, as people tend to have “fixations” one way or the other — this makes them human (Quote[4.7])). Another part of the issue is that a persona tends to represent the majority opinion, which is problematic when we want to portray the whole range of human attitudes.

Also, some participants were bothered by the fact that they could not influence the persona during the conversation. When two people interact, it is common for them to alter their views based on what the other person says — there is persuasion, adopting one’s perspective, developing it further, and, together, creating a shared perspective on a topic. However, some participants felt that this does not take place with the persona (Quote[4.8]). Instead, the persona had a set of fixed beliefs (based on the data) that could not be changed. For some participants, this reduced realism and hindered the perception of talking to another person (Quote[4.9]). This problem of *persona immovability* is one interaction challenge with AI-generated conversational personas, though it is unclear how fluid the persona’s attitudes *should* be, given that we must truthfully represent the end-user data. On the other hand, a more opinionated persona might help the *decision maker* change their views (e.g., toward greater empathy, understanding, and compassion) when the persona’s style would be more assertive. So, the (im)movability can work in both directions. From the principal-agent perspective, the AI agent’s overly neutral and passive behavior fails to meet the principal’s expectation of receiving genuine user perspectives, creating a misalignment between what decision makers need (i.e., authentic user voices) and what the agent delivers (i.e., sanitized responses).

### 5.2.5. CUI05 — ai's selective attention

One way in which the conversational persona exhibits agency is by selecting the *precise* information to include in its messages to the user. By the AI persona's "selective attention" (parallel to how humans have selective information processing behaviors), we refer to its inefficiencies in selecting or interpreting information about end-users. This took place in two primary ways: (1) misinterpreting the data, and (2) interpreting the data correctly but not in the way that the user wanted. The actual cases of factual errors were rare. In general, the LLM could interpret the survey response frequencies correctly; however, in some cases, the system did make errors (Quote[5.1]).

Also, somewhat interestingly, it is possible that the persona gives correct information but still misses the point in the participants' information request by producing off-topic or inconsistent answers. The former is related to the conversational persona that primarily gives more information than requested (Quote[5.2]), sometimes being off-topic (Quote[5.3]) and contrary to the participants' expectations (Quote[5.4]).

Inconsistencies were not common, but did emerge among some participants (Quote[5.5]; Quote[5.6]). However, some participants were keen to interpret the inconsistencies as natural traits of the persona (Quote[5.7]). Although it is true that *people* sometimes behave (and think) inconsistently, interpreting such inconsistencies as intended behavior by the LLM is problematic, since the system prompt was not designed to introduce them. So, the fact that the user is using a persona (of which they are aware) might yield a "false positive" impression that the conversational persona is more realistic due to small glitches that mimic human behavior. From a principal-agent perspective, the AI agent's misinterpretation or selective presentation of information represents a classic agency problem, in which the agent's information processing priorities may not align with the principal's actual information needs.

### 5.2.6. CUI06 — confusing distributional information

Processing distributional information was challenging for many of the participants (Quote[6.1]). Distributional information refers to the variability within the persona, as it is an amalgamation of data from different people (see Quote[6.2]). Within-group variability is characteristic of personas: even though the end-users represented by the persona are *similar* by some traits, they are not *identical* in most traits, and they might, in fact, differ in some other traits. When communicating this variability, users receive more nuanced information about the persona (stereotyping and simplification are major threats to persona application [45,46]). However, while sound in theory, in practice distributional information caused confusion among the participants in terms of persona attributes (Quote[6.3]), forming a clear mental model of the persona as a person (Quote[6.4]), or having to interpret numbers instead of just text (Quote[6.5]). Decision makers expect interacting with a person, so when the persona references group-level data, this expectation is not met (Quote[6.6]). Thus, distributional information hinders the anthropomorphization of the persona.

Still, this issue is not as straightforward as removing the distributional information, because there were *other* participants who clearly used it to make more informed conclusions about the persona. For example, some participants used distributional information to avoid a stereotypical view. Other participants preferred the added information about the persona group (Quote[6.7]; Quote[6.8]). Finally, there can be patterns of *genuine* contradictions in end-user behavior and attitudes regarding people's "belonging" to the same persona. Distributional information can reveal these, representing a greater range of user information. It can, however, be cognitively taxing for decision makers to process more detailed information, as we observed with some participants. From the principal-agent perspective, the agent faces competing demands to represent both individual persona authenticity (i.e., what principals expect) and statistical accuracy (i.e., what the data contains), creating an agency conflict between user experience and faithful data representation.

### 5.2.7. CUI07 — conversational cold start

Interestingly, even though chat may be considered an "easy" interface that requires little skill to get started and actively use, some participants struggled with these activities (generally referred to as a 'cold start' or the 'blank screen' problem in computer science). This aspect relates to the aforementioned issue of the "Hidden Information" in the sense that some participants felt that they did not know the basic information of the persona (Quote[7.1]). Some participants also indicated that they could use "professional help" with prompting — they did not perceive the conversational interface as difficult to use, but they found it difficult to figure out *what to type*. Another way of understanding this challenge is through the concept of "information scent" [47], which postulates that users need cues to effectively use and learn from the information provided in an interface. In the case of *profile* personas, such cues are readily available, as the persona information is visible. However, in the case of *conversational* personas, the provision of an "information scent" precisely forms a problem, because apart from the console (chat field), there tends to be little to no information provided to the user to get started. Although we provided a welcome message ("Howdy! I'm a persona shaped by survey data, encapsulating the sentiments and viewpoints of the respondents in this group. My responses are a blend of their shared thoughts and preferences. Feel free to ask for further insights!"), some participants found this inadequate (Quote[7.2]). The role of the welcome message is important to clarify what the conversational persona is about (transparency) and what one can do with it (management of expectations). From the principal-agent perspective, the lack of guidance creates an agency problem where principals cannot effectively communicate their information needs to the agent, leading to suboptimal delegation and underutilization of the persona's capabilities.

## 5.3. Frequency analysis

The transcripts were analyzed to determine how frequent each challenge was. The frequency analysis was carried out in collaboration with Google Gemini 2.5-Pro, which at the time of study (July 2025) performed best among the ones we tested for this task (others tested included GPT-4 and Claude-4). The prompt given to the LLM for stepwise analysis can be seen in online supplementary material, along with an example of the reasoning process applied by the Gemini 2.5-Pro model. The model was provided with one-sentence definitions of each challenge (also in the online supplementary material) and asked to evaluate, for each participant and each challenge, if the challenge existed in the given participant's think-aloud transcriptions. The model was asked to output a table indicating which challenges were present for which participants, and to provide real quotes from the transcripts as evidence of this presence. The lead author manually cross-checked the quotes, assessing whether the model's logic was sound and whether the quotes identified actually matched the challenge's definition. If a quote was not found in the transcripts, the lead author reuploaded the transcripts and asked the model to resume from the point where it had started hallucinating. Consequently, each coded instance (that is, pairs of challenge-participant that were found to be present) was manually checked by the lead author for accuracy (that is, presence of logical consistency and absence of hallucination).

Fig. 4(a) shows the frequency of the challenges (see the online supplementary material for the coding per participant). In the case of four participants (P16, P27, P32, and P46), no challenges were observed. Among the 52 participants with observed challenges, CUI02 (Hidden Persona) emerged as the most prevalent challenge with 33 occurrences (63.5% of the participants), indicating that the feeling of talking to an intermediary rather than directly to the persona was experienced by the majority of users. CUI06 (Coping with Distributional Information) was the second most frequent challenge with 28 occurrences (53.8%), suggesting that presenting group variability while maintaining persona authenticity proved difficult for more than half




Fig. 4. Frequency and co-occurrence of persona chat interaction challenges. The reader is encouraged to zoom in for better legibility.

of the participants. CUI04 (Lack of AI Agency) and CUI01 (Hidden Information) showed moderate frequency at 25 (48.1%) and 23 occurrences (44.2%), respectively, indicating that issues with passive persona behavior and the limited information visibility affected slightly less than half of the participants. CUI05 (AI’s Selective Attention) appeared less frequently with 19 occurrences (36.5%), while CUI03 (Hidden UI) was experienced by fewer participants with 13 occurrences (25%), suggesting that conversational attention failures and the lack of visual interface elements were less common concerns. CUI07 (Conversational Cold Start) was the least frequent challenge with only 8 occurrences (15.4%), indicating that most users were able to initiate conversations despite the open-ended nature of the conversational interface, though this challenge still affected a meaningful subset of participants.

We also computed how many times each challenge co-occurred with another challenge (see Fig. 4(b)). In the following, we interpret the co-occurrences that took place more than 10 times (i.e., represent the most frequent co-occurrences; the cutoff value of 10 was determined by manually examining the frequency distribution). For each co-occurrence, we provide a reasonable explanation of the logical connection between the challenges and, because this connection is partially speculative, we also articulate how future work should examine it.

First, the frequent co-occurrence ( $n = 20$ ) of CUI02 (Hidden Persona) and CUI04 (Lack of AI Agency) suggests that when users perceive an intermediary barrier between themselves and the persona, this disconnect is exacerbated by the persona’s passive, opinion-less nature, creating a compounding effect where both the sense of talking to a representative rather than the actual persona and the AI’s neutral responses reinforce each other to create an overall artificial interaction experience. Future research should investigate whether improving AI agency through *stronger opinions and more dynamic responses could help break through the perceived intermediary barrier and establish more direct persona connections.*

Second, the co-occurrence ( $n = 18$ ) of CUI02 (Hidden Persona) and CUI06 (Coping with Distributional Information) likely stems from the challenge that when personas must represent statistical variability within groups, this mathematical complexity further distances users from feeling they are interacting with a coherent individual rather than a system processing data. Research should explore whether developing *more sophisticated narrative techniques for presenting group variability can maintain both statistical accuracy and individual persona authenticity.*

Third, CUI01 (Hidden Information) and CUI02 (Hidden Persona) co-occur ( $n = 15$ ) perhaps because information opacity compounds the intermediary problem — when users cannot easily access persona details and simultaneously feel disconnected from the actual persona, both barriers create distance and reduce engagement. Studies should examine whether *exposing more information about the conversational persona (e.g., a backstory) can simultaneously improve the sense of direct persona connection.*

Fourth, the co-occurrence ( $n = 14$ ) of CUI04 (Lack of AI Agency) and CUI06 (Coping with Distributional Information) suggests that when

personas must represent group statistics, they may default to neutral, noncommittal responses to avoid misrepresenting the diversity within their represented group, thus appearing passive and opinion-less. Research should investigate *methods for personas to express stronger viewpoints while appropriately contextualizing their position within broader group distributions.*

Fifth, CUI01 (Hidden Information) and CUI04 (Lack of AI Agency) may co-occur ( $n = 13$ ) because the interaction setting in which users must actively query for specific persona details instead of having immediate access requires the persona to wait passively for direct questions instead of proactively sharing opinions or taking conversational initiative, thus appearing neutral and non-committal even when they possess rich information. Future work should explore whether *personas can be designed to volunteer relevant information proactively during the conversation flow, rather than waiting for explicit queries from the user.*

Sixth, the co-occurrence ( $n = 12$ ) of CUI05 (AI’s Selective Attention) and CUI06 (Coping with Distributional Information) implies that the complexity of representing group variability may lead to misinterpretations and off-topic responses when confounding individual persona characteristics with statistical group data. Future work should explore whether personas can be *designed to express uncertainty and partial knowledge in more engaging, human-like ways.*

Finally, CUI02 (Hidden Persona) and CUI05 (AI’s Selective Attention) can co-occur ( $n = 11$ ) because when users already feel disconnected from the persona, inconsistent responses further reinforce the sense that they are interacting with a system rather than an individual. Studies should examine whether *improving conversational consistency can help establish a stronger persona presence and reduce the intermediary effect.*

## 6. Discussion and implications

### 6.1. Theoretical implications

As a contribution of this work, we identify seven interaction challenges faced by on-site user study participants when interacting with AI-generated conversational personas, and conceptualize them as instances of representational decision-support technology, in which decision makers query information about target users based on persona “surrogates” that represent these target users. Our identification of interaction challenges builds on the foundational work on AI-generated personas [17,29,31], but presents the first analysis of the difficulties of interacting with AI-generated conversational personas. Although empirical investigations have demonstrated the technical feasibility of AI-generated personas [2,17,48], the theoretical understanding of human–AI interaction processes within conversational persona systems remains underdeveloped, particularly with respect to representational friction that influences how decision makers learn about users’ needs.

Conceptually, we observe that end-user representation is an agency problem in itself — the end-user (or group of them) is represented




Fig. 5. Interaction challenges in persona-based DSSs are more complex because the persona layer is further represented by AI technologies.

by someone or something else. Whereas a persona serves as a stand-in for real end-users during the design process, in the case of AI-generated conversational personas, the *persona* is now represented by another actor, the LLM. In addition, this representation occurs within a constrained UI with hidden or opaque elements. So, the “chain of representation” in conversational personas goes from users represented by personas that, in turn, are represented by LLMs in a specific interface (see Fig. 5).

Each link in the representation chain poses a potential challenge for the eventual user interaction with the persona: the segmentation process that divides end-users to groups simplifies a population (CUI01). Creating persona descriptions from the segment (write-up, enrichment) involves interpretation trade-offs that occur regardless of who or what the creator is [49]. The LLM’s interpretation of the persona descriptions, in turn, is strongly influenced by the applied prompt (CUI02, CUI04) and, of course, the opaque way in which the LLM operates (e.g., selective attention, CUI05). Finally, the conversational interface imposes constraints that limit the communication of the information (presenting distributional information, CUI06) and affordances (actions users can take, CUI03), as well as requiring decision makers to put more cognitive effort into initiating the information retrieval (CUI07).

Our theoretical contribution, therefore, concerns reconceptualizing conversational AI interactions when AI technology serves as a representative agent for the persona rather than a direct provider of information. Unlike general-purpose chatbots (e.g., ChatGPT), conversational personas involve a three-way relationship (decision maker, AI agent, represented users) instead of the more common perspective of two-way interaction (user, AI system), which yields interaction challenges absent in general conversational AI, because general-purpose AI agents do not represent real groups of people. By conjecture, this implies that even though the core technology is the same (and thus some interaction challenges would overlap), AI-generated conversational personas can face “old” interaction challenges in new ways or new challenges altogether.

Previous studies tend to examine the interaction challenges of conversational technologies as direct interaction with conversational systems and their users, without investigating how conversational systems represent real users. For example, in their 2018 article, Brandtzaeg and Følstad [34] mention many interaction challenges faced by conversational agents at the time. These include (1) balancing human vs. robot qualities, (2) unpredictable user input, (3) determining valid responses, (4) vulnerability to manipulation, (5) managing user expectations, (6) inadequate user and context models, and (7) designing for open-ended conversations. Given our findings, we observe that many of these remain topical for AI-generated conversational personas. As evidenced by CUI02, there is a need to avoid robotic qualities in persona design, because perception of the persona as robotic was associated with worse user experience. As another case in point, lack of agency (CUI04) appears to conceptually overlap with inadequate user and context models (i.e., chatbots lacking understanding of who they are talking to and why [34]); adapting the conversational persona to its interaction context is quite not there yet.

However, due to the general technical quality of foundational LLMs, other problems seem to have reduced in appearance. Specifically, unpredictable user inputs do not appear to be an issue for AI-generated conversational personas, neither determining valid responses, or being

vulnerable to manipulation (they maintain their role well; even too well according to the participants that wanted the persona to change its opinion). In contrast, designing for open-ended conversations and managing user expectations seem to be borderline cases; the technical fluency of AI-generated conversational personas in handling open-ended queries is good, but representational conflicts can still emerge in which the user is not satisfied concerning how the persona refers to itself, conveys information, or establishes a social connection with the user.

More importantly, previous literature does not examine the specific representational dilemmas of AI-generated conversational personas, missing a dedicated analysis of various challenges associated with this type of decision-making technology, such as presenting distributional information (CUI06) and the selective attention problem (CUI05). Again, we postulate that studies on this topic could benefit from conceptualizing interaction with AI-generated conversational personas as a three-way relationship (decision maker, AI agent, represented users), because this helps identify and understand distinct interaction challenges absent in general two-way human–AI (or user–persona) interaction.

Interestingly, there are some challenges often mentioned in LLM chatbots that we did not find. That is, the participants tended to have trusting beliefs even though the persona development process was not explained in detail (i.e., there was low transparency). There were no notable concerns about hallucination among the participants (i.e., the persona inventing information), which is surprising given the notoriety of the hallucination problem [50]. However, in our case, low hallucination suspicions actually make sense because, as explained in the methodology, S2P is grounded in the end-user information through the implementation of the RAG. Trusting beliefs can also arise because decision makers tend to trust data-driven personas [2,17].

From a DSS theory perspective, the conversational persona can be seen as a *decision-making mediator* that transforms statistical data into natural conversation creates what we term a “representation chain” — from end-users to survey data to an LLM to conversational interface and, finally, to the decision maker. While this matches the definition of a DSS, the conversational and representative nature of AI-generated conversational personas produces interaction challenges that blend social, psychological, and informational aspects. Consequently, the design of such DSSs becomes more intertwined with interaction design and adjacent fields like UCD; that is, how to orchestrate the interaction so that the decision makers extract the maximum value (with minimal distractions) out of it. As aptly put by Chaves and Gerosa [51] “making a conversational agent acceptable to users is primarily a social, not only technical, problem” (p. 729).

## 6.2. Practical implications

Table 3 summarizes the challenges as questions for developers of conversational persona DSSs, and proposes preliminary ideas for addressing them. We also reviewed the think-aloud transcript in an *ad-hoc* manner to determine whether participants themselves provided solution ideas; though this was beyond the direct scope of our work, we investigated this possibility to provide the reader a fuller picture. This ad-hoc analysis yielded a few interesting ideas that we share here: (1) providing recommended questions to ask the persona (“I think it would

**Table 3**  
Design ideas to tackle conversational persona challenges.

Challenge	Design Ideas (DIs)
CUI01	D101 : Add a concise “profile card” of key persona information at the start of the conversational interaction; D102 : Configure progressive information disclosure over the course of the interaction
CUI02	D103 : Apply first-person narrative with personal anecdotes and examples; D104 : Add realistic emotional reactions and personal stakes in responses
CUI03	D105 : Provide contextual action buttons and quick-reply options; D106 : Implement dynamic background colors increasing the saliency of specific information
CUI04	D107 : Consider opinionated responses when they correspond with clear tendency in the data; D108 : Configure an ability to change stance based on conversation (but clearly indicating when it might deviate from data); D109 : Configure proactive questioning and challenging user assumptions when they differ with the persona’s mental model
CUI05	D110 : Implement query clarification before responding; D111 : Add confidence indicators for persona responses
CUI06	D112 : Present within-group variability using primarily qualitative descriptions (e.g. “most people in this group...”) with the option to view exact percentages if requested; D113 : Switch between individual and group perspectives; D114 : Apply uncertainty expressions when representing diverse views
CUI07	D115 : Configure conversation starter suggestions with preview outcomes; D116 : Provide example question templates for common use cases

be nice if a survey to persona our system could help people who come up with questions to formulate them to make them clearer, clearer or simpler.” (P20)), (2) the persona providing understanding checks (“Yeah, it could have (...) asked more questions like, does that answer your questions? Do you (...) need any more clarifications? What do you mean? I mean that lacks in the persona.” (P37)), and (3) offering a multi-persona setup (“Imagine talking to multiple personas at the same time. (...) So you can ask them, for example, would AI affect or (...) Would this certain aspects improve your job opportunities? And then let’s say based on their educational background based on their ethnic background, they might say they, yeah, (...) my group of people, they don’t get jobs easily because of ‘ABC’ ” (P20)). These suggestions offer ideas for tackling interaction challenges in the further development of persona-based DSSs, an endeavor that requires experimentation that goes beyond the scope of the current work.

### 6.3. Limitations and future research directions

Regarding limitations, first, our study used an extensive dataset comprising 55 survey items and more than 10K respondents, which affects the amount of information that the personas can convey. Future work could compare AI-generated personas from high-information datasets with those from low- (i.e., fewer survey items and responses) and high-information datasets, as information complexity seems to affect interaction with AI-generated personas.

Second, while our participant pool consisted primarily of researchers and engineers, this composition aligned with our research focus on technology-oriented decision makers who are likely to be early adopters of AI-generated persona systems. However, this specialized sample limits the generalizability of our findings to other organizational contexts where personas are used, such as marketing, product management, or user experience design. Future studies should examine how different professional roles (e.g., marketing managers, product managers, designers) experience these interaction challenges, as decision-making priorities and information processing patterns may vary across organizational contexts.

Third, as noted in previous research, LLMs pose risks and limitations for persona design and conversational DSSs [30,32,48,52]. Personas generated by AI may inherit biases from the source data, also when using RAG. However, bias is also an issue in traditional persona creation, as data sets in qualitative studies lack statistical

representativeness [6,28] and persona creation using manual analysis can involve various forms of stereotyping, prejudice, and bias [53]. Notably, AI-generated personas inherit any problems and biases from their source data and/or the AI’s interpretation of it. This implies that, for responsible deployment, (a) decision makers should receive training on the limitations of conversational persona systems; (b) organizations should maintain human oversight of critical decisions based on information and interpretations conveyed by LLMs, and (c) ethical audits should assess the risks of these systems.

Fourth, our methodological approach presents several limitations that warrant discussion. The within-subjects design is useful for adjusting for individual differences, but it may also introduce learning effects, such that participants who first interacted with the profile interface might have transferred this knowledge to their chat interactions, potentially affecting their information-seeking behavior. For example, P41’s observation that “the first option (conversational persona) (would) present me all the details” suggests an awareness of available information from prior exposure to the profile interface.

The think-aloud protocol has the advantage of allowing communication about important cognitive processes, but it can also alter natural interaction patterns with the system. Some participants, particularly those less comfortable with verbal articulation, might have modified their typical decision-making approach. Future studies could complement think-aloud data with other observational methods, such as eye-tracking, to capture decision-making behaviors that participants may not verbalize. This could also involve analyzing user chats to identify traces of interaction challenges.

Finally, not all users face them the same way or even consider them a challenge. For example, some users find the distributional information conveyed by the conversational persona easy to understand, while others find it confusing. Some users were expressing anxiety about having to deal with an empty input field in the conversational interface, while others found it a good starting point. Therefore, future work should test the relationship between these interaction challenges and user features, such as interaction preferences, experience, job role, and AI proficiency. In addition, principal-agent relationships can vary in different cultural and organizational contexts. Future work could therefore study (1) how cultural differences in authority and delegation affect persona interactions, (2) how agency expectations differ between domains (e.g., healthcare vs. marketing vs. policy), (3) how professional norms and regulations affect optimal persona design, and

(4) how personas should adapt to different decision-making styles and preferences.

Furthermore, the conceptual linkage with the principal-agent theory is useful for proposing solution ideas to test in future work. In particular, the following research directions (RDs) could be pursued to experiment with ways to mitigate the challenges:

**RD01: Alignment mechanisms and incentive design.** Prompting the LLM that guides the persona could be done in a way that aligns the persona's representation with decision-maker (principal) interests. This would require an understanding of the decision maker's goals, which could be obtained, e.g., by explicitly asking the decision maker to state their goals before interacting with the persona. More technically elaborate solutions could involve design reward functions that optimize for decision-maker outcomes. These reward functions aim to create (an artificial yet functional) sense of having "skin in the game," in which the persona's representation of a user group is tied to decision outcomes that benefit that user group. Given the generally observed sycophancy in LLM behavior [54], AI-generated conversational personas could in theory be prompted to be supportive of decision-making tasks; while in economic agency theory, researchers utilize the concept of rationality to design incentive structures, LLMs are not rational actors in the same sense as humans and therefore prompting is seen as the mechanism for alignment between human and LLM objectives (computational research refers to this as 'value alignment' [55]).

**RD02: Information asymmetry mitigation.** This direction deals with creating transparent persona DSSs that reduce the knowledge gap between AI-generated personas and human principals. For example, these systems could expose the persona's confidence levels and data limitations more saliently, including providing "side information" about how the persona's responses are generated and what information they rely on, and standardized "information cards" that reveal what data the persona can and cannot access. In principle, information asymmetry mitigation efforts align with the computational research community's goal of designing explainable and transparent AI systems.

**RD03: Dynamic agency and adaptive representation.** Conversational personas could adapt their user representation strategies based on decision makers' needs and improve representation over time based on decision maker feedback. Representation adaptation should also take place by detecting and correcting agency problems in persona interactions, e.g., by using AI agents to cross check each other's representations and flagging potential representation failures, such as when personas are misrepresenting underlying data, toward addressing the question of "What monitoring mechanisms best detect AI agency failures?". This can also include investigating how conversational personas should handle conflicting signals from the same decision maker because people, not only AI, can be fickle.

#### CRedit authorship contribution statement

**Joni Salminen:** Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Soon-Gyo Jung:** Software, Resources, Data curation. **Ilkka Kaate:** Visualization, Validation, Software, Resources, Methodology, Investigation. **Trang Thi Thu Xuan:** Writing – original draft, Visualization, Project administration, Methodology, Investigation. **Ji-nan Y. Azem:** Investigation, Data curation. **Kholoud Khalil Aldous:** Data curation, Conceptualization. **Danial Amin:** Writing – review & editing, Visualization. **Bernard J. Jansen:** Writing – review & editing, Supervision, Methodology, Investigation, Conceptualization.

#### Declaration of Generative AI in the Writing Process

The authors used GPT-3.5 Turbo, GPT-4, and Claude-3.5-Sonnet to accomplish the research purpose and address the blank page problem encountered in the writing profession. After using these tools, the authors reviewed and edited the content as needed and took full responsibility for the content.

#### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

#### Data availability

Data will be made available on request.

#### References

- [11] Dennis D. Fehrenbacher, Martin Weisner, Avatars and organizational knowledge sharing, *Decis. Support Syst.* 182 (2024) 114245, Publisher: Elsevier.
- [1] Alan Cooper, The inmates are running the asylum, in: Udo Arend, Edmund Eberleh, Knut Pitschke (Eds.), *Software-Ergonomie '99: Design Von Information-swelten*, Vieweg+Teubner Verlag, Wiesbaden, 1999, [http://dx.doi.org/10.1007/978-3-322-99786-9\\_1](http://dx.doi.org/10.1007/978-3-322-99786-9_1), 17–17.
- [2] Joni Salminen, Chang Liu, Wenjing Pian, Jianxing Chi, Essi Häyhänen, Bernard J. Jansen, Deus ex machina and personas from large language models: Investigating the composition of AI-generated persona descriptions, in: *Proceedings of the CHI Conference on Human Factors in Computing Systems*, CHI '24, Association for Computing Machinery, 2024, pp. 1–20, <http://dx.doi.org/10.1145/3613904.3642036>.
- [3] Soon-Gyo Jung, Joni Salminen, Kholoud Khalil Aldous, Bernard J. Jansen, PersonaCraft: Leveraging language models for data-driven persona development, *Int. J. Hum.-Comput. Stud.* 197 (2025) 103445, <http://dx.doi.org/10.1016/j.ijhcs.2025.103445>.
- [4] Sunhyo Oh, Taejun Park, Gahgene Gweon, Cognitive and emotional engagement design factors in text-based pedagogical conversational agents: Impacts on student learning and motivation, *Int. J. Hum.-Comput. Stud.* (2026) 103750.
- [5] Joni Salminen, Lene Nielsen, Ali Farooq, Bernard Jansen, AI-generated personas: representing user needs with generative AI models, *Int. J. Hum. Comput. Stud.* (2026) 103745.
- [6] Bernard J. Jansen, Soon-gyo Jung, Joni Salminen, From flat file to interface: Synthesis of personas and analytics for enhanced user understanding, *Proc. Assoc. Inf. Sci. Technol.* 57 (1) (2020) <http://dx.doi.org/10.1002/pr2.215>.
- [7] Leonard Boussiou, Andrew Chen, Ming Fan, Apurva Jain, Socratic iterative reasoning: Enhancing LLM decision-making in the beer game supply chain, in: *ICIS 2025 Proceedings*, 2025.
- [8] Pengli Wu, Xuebing Yang, Yanlong Wen, Wensheng Zhang, MDR: Memory distillation and reproduction for personalized dialogue generation, *Knowl.-Based Syst.* 336 (2026) 115252, <http://dx.doi.org/10.1016/j.knosys.2025.115252>.
- [9] Tara Matthews, Tejinder Judge, Steve Whittaker, How do designers and user experience professionals actually perceive and use personas? in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM Press, 2012, p. 1219, <http://dx.doi.org/10.1145/2207676.2208573>.
- [10] W. Wang, Y. Tong, J. Mou, Artificial intelligence agents or human agents? Impact of online customer service agents on crowdfunding performance, *Decis. Support Syst.* (2025) 114562.
- [12] Gretchen M. Grimes, Ryan M. Schuetzler, Justin Scott Giboney, Mental models and expectation violations in conversational AI interactions, *Decis. Support Syst.* 144 (2021) 113515.
- [13] Veda C. Storey, Alan R. Hevner, Victoria Y. Yoon, The design of human-artificial intelligence systems in decision sciences: A look back and directions forward, *Decis. Support Syst.* 182 (2024) 114230, <http://dx.doi.org/10.1016/j.dss.2024.114230>.
- [14] Mirjana Prpa, Giovanni Maria Troiano, Matthew Wood, Yvonne Coady, Challenges and opportunities of LLM-based synthetic personae and data in HCI, in: *Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems*, in: CHI EA '24, Association for Computing Machinery, 2024, pp. 1–5, <http://dx.doi.org/10.1145/3613905.3636293>.
- [15] Senthil Chandrasegaran, Peter Lloyd, et al., Synthetic users: insights from designers' interactions with persona-based chatbots, *AI EDAM* 39 (2025) e2.
- [16] Vanessa Sattelle, Juan Carlos Ortiz, Generating user personas with AI: Reflecting on its implications for design, *DRS Bienn. Conf. Ser.* (2024).
- [17] Joongi Shin, Michael A. Hedderich, Bartłomiej Jakub Rey, Andrés Lucero, Antti Oulasvirta, Understanding human-AI workflows for generating personas, in: *Designing Interactive Systems Conference*, ACM, 2024, pp. 757–781, <http://dx.doi.org/10.1145/3643834.3660729>.
- [18] Leen F. Al Qadi, Soon-Gyo Jung, Danial Amin, Amani Alabed, Joni Salminen, Bernard J. Jansen, Representing religious practices via AI-generated personas: A case study of ramadan behaviors from four predominantly muslim countries, in: *2025 IEEE/ACS 22nd International Conference on Computer Systems and Applications*, AICCSA, IEEE, 2025, pp. 1–10.

- [19] Jinan Azem, Farhan Ahmed, Joni Salminen, Bernard J. Jansen, Interviewing AI-generated personas: Talking to your data to generate qualitative text from users, in: *The 3rd International Conference on Foundation and Large Language Models (FLLM2025)*, IEEE, 2025, pp. 1–8.
- [20] Amon Rapp, Lorenzo Curti, Arianna Boldi, The human side of human-chatbot interaction: A systematic literature review of ten years of research on text-based chatbots, *Int. J. Hum.-Comput. Stud.* 151 (2021) 102630, <http://dx.doi.org/10.1016/j.ijhcs.2021.102630>.
- [21] Kathleen M. Eisenhardt, Agency theory: An assessment and review, *Acad. Manag. Rev.* 14 (1) (1989) 57–74.
- [22] Jacky Akoka, Isabelle Comyn-Wattiau, IS/IT Backsourcing decision making-A design science research approach, *Decis. Support Syst.* 189 (2025) 114379, Publisher: Elsevier.
- [23] Mohammad Hossein Jarrahi, Paavo Ritala, Rethinking AI agents: A principal-agent perspective, *Calif. Manage. Rev.* 65 (3) (2023) 27–47.
- [24] Lene Nielsen, Kira Storgaard Hansen, Jan Stage, Jane Billestrup, A template for design personas: Analysis of 47 persona descriptions from Danish industries and organizations, *Int. J. Sociotechnology Knowl. Dev.* 7 (1) (2015) 45–61, <http://dx.doi.org/10.4018/ijskd.2015010104>.
- [25] Peter G.W. Keen, Decision support systems: a research perspective, in: *Decision Support Systems: Issues and Challenges: Proceedings of an International Task Force Meeting*, 1980, pp. 23–44.
- [26] Chris Chapman, Russell P. Milham, The personas' new clothes: Methodological and practical arguments against a popular method, in: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 50, 2006, pp. 634–636, <http://dx.doi.org/10.1177/154193120605000503>.
- [27] Jennifer (Jen) McGinn, Nalini Kotamraju, Data-driven persona development, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM Press, 2008, p. 1521, <http://dx.doi.org/10.1145/1357054.1357292>.
- [28] Jisun An, Haewoon Kwak, Joni Salminen, Soon-gyo Jung, Bernard J. Jansen, Imaginary people representing real numbers: Generating personas from online social media data, *ACM Trans. Web (TWEB)* 12 (4) (2018) 27, <http://dx.doi.org/10.1145/3265986>.
- [29] Perttu Hämäläinen, Mikke Tavast, Anton Kunnari, Evaluating large language models in generating synthetic HCI research data: a case study, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '23, Association for Computing Machinery, 2023, pp. 1–19, <http://dx.doi.org/10.1145/3544548.3580688>.
- [30] Albrecht Schmidt, Speeding up the engineering of interactive systems with generative AI, in: *Companion Proceedings of the 2023 ACM SIGCHI Symposium on Engineering Interactive Computing Systems*, 2023, pp. 7–8.
- [31] Andreas Schuller, Doris Janssen, Julian Blumenröther, Theresa Maria Probst, Michael Schmidt, Chandan Kumar, Generating personas using LLMs and assessing their viability, in: *Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems*, in: CHI EA '24, Association for Computing Machinery, 2024, pp. 1–7, <http://dx.doi.org/10.1145/3613905.3650860>.
- [32] Stefano De Paoli, User personas, ideation and large language models: A post-hoc study, *Int. J. Hum.-Comput. Stud.* 208 (2026) 103690, <http://dx.doi.org/10.1016/j.ijhcs.2025.103690>.
- [33] Asbjørn Følstad, Petter Bae Brandtzæg, Chatbots and the new world of HCI, *Interactions* 24 (4) (2017) 38–42, <http://dx.doi.org/10.1145/3085558>.
- [34] Petter Bae Brandtzæg, Asbjørn Følstad, Chatbots: changing user needs and motivations, *Interactions* 25 (5) (2018) 38–43, <http://dx.doi.org/10.1145/3236669>.
- [35] Joshua Maynez, Shashi Narayan, Bernd Bohnet, Ryan McDonald, On faithfulness and factuality in abstractive summarization, in: Dan Jurafsky, Joyce Chai, Natalie Schluter, Joel Tetreault (Eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Online, 2020, pp. 1906–1919, <http://dx.doi.org/10.18653/v1/2020.acl-main.173>.
- [36] Mohit Jain, Pratyush Kumar, Ramachandra Kota, Shwetak N. Patel, Evaluating and informing the design of chatbots, in: *Proceedings of the 2018 Designing Interactive Systems Conference*, ACM, 2018, pp. 895–906, <http://dx.doi.org/10.1145/3196709.3196735>.
- [37] Michael J. Shaw, Pei-Lei Tu, Prabuddha De, Applying machine learning to model management in decision support systems, *Decis. Support Syst.* 4 (3) (1988) 285–305, Publisher: Elsevier.
- [38] Tianran Hu, Anbang Xu, Zhe Liu, Quanzeng You, Yufan Guo, Vibha Sinha, Jiebo Luo, Rama Akkiraju, Touch your heart: A tone-aware chatbot for customer care on social media, in: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, ACM, 2018, pp. 1–12, <http://dx.doi.org/10.1145/3173574.3173989>.
- [39] Ewa Luger, Abigail Sellen, “Like having a really bad PA”: The Gulf between user expectation and experience of conversational agents, in: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, ACM, 2016, pp. 5286–5297, <http://dx.doi.org/10.1145/2858036.2858288>.
- [40] Lene Nielsen, *Personas - User Focused Design*, second ed. 2019, Springer, 2019.
- [41] Joni Salminen, Kathleen Guan, Soon-gyo Jung, Shammur Absar Chowdhury, Bernard J. Jansen, A literature review of quantitative persona creation, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 2020, pp. 1–14, <http://dx.doi.org/10.1145/3313831.3376502>.
- [42] Lene Nielsen, Joni Salminen, Soon-Gyo Jung, Bernard J. Jansen, Think-aloud surveys: A method for eliciting enhanced insights during user studies, in: *Human-Computer Interaction-INTERACT 2021: 18th IFIP TC 13 International Conference*, Bari, Italy, August 30–September 3, 2021, Proceedings, Part V 18, Springer, 2021, pp. 504–508.
- [43] Juliet Corbin, Anselm Strauss, *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*, SAGE, 2008.
- [44] Ilkka Kaate, Joni Salminen, Soon-Gyo Jung, Trang Thi Thu Xuan, Jinan Y. Azem, João M. Santos, Bernard J. Jansen, When personas talk to you: Evaluating the evolution of user personas from static profiles to conversational user interfaces, in: *Proceedings of the 2025 ACM Designing Interactive Systems Conference*, DIS '25, Association for Computing Machinery, New York, NY, USA, 2025, pp. 2350–2372, <http://dx.doi.org/10.1145/3715336.3735676>.
- [45] Phil Turner, Susan Turner, Is stereotyping inevitable when designing with personas? *Des. Stud.* 32 (1) (2011) 30–44, <http://dx.doi.org/10.1016/j.destud.2010.06.002>.
- [46] Yinan Ma, Hua Dong, Persona-based design process model: A step-by-step method for designing products and interfaces, in: *Proceedings of the 2016 3rd International Conference on Computing for Sustainable Global Development*, INDIACom, IEEE, 2016, pp. 3079–3084.
- [47] Ed H. Chi, Peter Piroli, Kim Chen, James Pitkow, Using information scent to model user information needs and actions and the web, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 2001, pp. 490–497, <http://dx.doi.org/10.1145/365024.365325>.
- [48] Paluck Deep, Monica Bharadhidasan, A. Baki Kocaballi, “She was useful, but a bit too optimistic”: Augmenting Design with Interactive Virtual Personas, *Int. J. Hum.-Comput. Stud.* 205 (2025) 103646, <http://dx.doi.org/10.1016/j.ijhcs.2025.103646>.
- [49] Joni Salminen, Soon-gyo Jung, Bernard J. Jansen, Are data-driven personas considered harmful?: Diversifying user understandings with more than algorithms, *Pers. Stud.* 7 (1) (2021) 48–63, ISBN: 2205-5258.
- [50] Xiang Shi, Jiawei Liu, Yinpeng Liu, Qikai Cheng, Wei Lu, Know where to go: Make LLM a relevant, responsible, and trustworthy searchers, *Decis. Support Syst.* 188 (2025) 114354, Publisher: Elsevier.
- [51] Ana Paula Chaves, Marco Aurelio Gerosa, How should my chatbot interact? A survey on social characteristics in human-chatbot interaction design, *Int. J. Human-Computer Interact.* 37 (8) (2021) 729–758, <http://dx.doi.org/10.1080/10447318.2020.1841438>.
- [52] J. Strunk, A. Nissen, S. Smolnik, All risks ain't the same—A risk facets perspective on AI-based decision support systems, *Decis. Support Syst.* (2025) 114557.
- [53] Bernard J. Jansen, Soon-Gyo Jung, Lene Nielsen, Kathleen W. Guan, Joni Salminen, How to create personas: Three persona creation methodologies with implications for practical employment, *Pac. Asia J. Assoc. Inf. Syst.* 14 (3) (2022) 1.
- [54] Lars Malmqvist, Sycophancy in large language models: Causes and mitigations, in: Kohei Arai (Ed.), *Intelligent Computing*, vol. 1426, Springer Nature Switzerland, Cham, 2025, pp. 61–74, [http://dx.doi.org/10.1007/978-3-031-92611-2\\_5](http://dx.doi.org/10.1007/978-3-031-92611-2_5), Series Title: Lecture Notes in Networks and Systems.
- [55] Timothy R. McIntosh, Tong Liu, Teo Susnjak, Paul Watters, Malka N. Halgamuge, A reasoning and value alignment test to assess advanced GPT reasoning, *ACM Trans. Interact. Intell. Syst.* 14 (3) (2024) 1–37, <http://dx.doi.org/10.1145/3670691>.

**Joni Salminen** is an Associate Professor (tenure track) at the School of Marketing and Communication in University of Vaasa, Vaasa, Finland. His research focuses on data-driven personas, along with overlapping topics such as Human-Centered AI, Quantitative UX, Interactive Systems, User Segmentation Algorithms, and so on.

**MSc. Soon-gyo Jung** is a Full-Stack Software, AI, and Data Engineer at Qatar Computing Research Institute (QCRI), specializing in scalable data-driven systems, LLM applications, and research-oriented software engineering.

**Ilkka Kaate** is a postdoctoral researcher in marketing at the University of Turku, Finland, with research interest in deepfake personas. Additionally, Ilkka is a digital marketing entrepreneur specializing in social media and search marketing. In his free time Ilkka likes to act, sing (solo and choir) and play the guitar.

**Trang Xuan** is a doctoral student at the University of Vaasa. She holds a bachelor's degree from HAMK University of Applied Science (Finland, 2014) and an M.Sc. from the University of Strathclyde (UK, 2016), majoring in Marketing. Before her Ph.D. journey, she worked in marketing agencies for 8 years and marketing lecturer in Vietnamese universities for 5 years. Trang's current research focus is AI chatbot usability in Marketing Education.

**Jinan Y. Azem** is a Researcher and UX/Product Designer at Qatar Computing Research Institute. She received her B.S. degree in Information Systems from Carnegie Mellon University in Qatar. Her work focuses on designing and evaluating user-centered digital systems, with research interests in usability, user experience, accessibility, and the application of HCI methods to digital products, and explores how digital technologies can better serve diverse user populations.

**Kholoud Khalil Aldous** is a Post-Doctoral Researcher at Northwestern University in Qatar and an Instructor at the University of Doha for Science and Technology. She received the B.S. and M.S. degrees in Computer Science from Qatar University and the Ph.D. degree in Computer Science and Engineering from Hamad Bin Khalifa University, Qatar. Her research interests include social media analysis, user engagement modeling, data analytics, and natural language processing.

**Danial Amin** is a doctoral researcher at the University of Vaasa, focusing on generative AI personas for social good, developing ethical and inclusive user-representation that benefit marginalized communities in the Global South. Danial has over 8 years of experience in consulting multi-national organizations in artificial intelligence and data science domains for impactful projects.

**Bernard J. Jansen** is a Principal Scientist at the Qatar Computing Research Institute, Doha, Qatar. He is a West Point graduate with a Ph.D. in computer science from Texas A&M University and master's degrees from Texas A&M (computer science) and Troy State (international relations). Dr. Jim Jansen served in the U.S. Army as an Infantry enlisted soldier and communication commissioned officer.