

# The Impostor is Among Us: Can Large Language Models Capture the Complexity of Human Personas?

Christopher Klaus Lazik  
Humboldt-Universität zu Berlin  
Berlin, Germany  
lazikchr@hu-berlin.de

Christopher Katins  
Humboldt-Universität zu Berlin  
Berlin, Germany  
christopher.katins@hu-berlin.de

Charlotte Kauter  
Humboldt-Universität zu Berlin  
Berlin, Germany  
Charlotte.kauter@student.hu-berlin.de

Jonas Jakob  
Humboldt-Universität zu Berlin  
Berlin, Germany  
jonas.jakob@student.hu-berlin.de

Caroline Jay  
School of Computer Science  
The University of Manchester  
Manchester, United Kingdom  
caroline.jay@manchester.ac.uk

Lars Grunske  
Humboldt-Universität zu Berlin  
Berlin, Germany  
grunske@informatik.hu-berlin.de

Thomas Kosch  
Humboldt-Universität zu Berlin  
Berlin, Germany  
thomas.kosch@hu-berlin.de



**Figure 1:** We investigated if and how users discern between textual descriptions of human-crafted and AI-generated personas. In a survey study, we investigated how common descriptions of personas, based on related work, affect the perceived realism and complexity of personas. We find that users can distinguish between human-crafted and AI-generated personas to a large extent, showing that stereotypicality, realism, and appeal are indicative features.



This work is licensed under a Creative Commons Attribution International 4.0 License.

MuC '25, Chemnitz, Germany

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1582-2/25/08

<https://doi.org/10.1145/3743049.3743057>

## Abstract

Large Language Models (LLMs) created new opportunities for generating personas, expected to streamline and accelerate the human-centered design process. Yet, AI-generated personas may not accurately represent actual user experiences, as they can miss contextual and emotional insights critical to understanding real users' needs and behaviors. This introduces a potential threat to quality, especially for novices. This paper examines the differences in how users

perceive personas created by LLMs compared to those crafted by humans regarding their credibility for design. We gathered ten human-crafted personas developed by HCI experts according to relevant attributes established in related work. Then, we systematically generated ten personas with an LLM and compared them with human-crafted ones in a survey. The results showed that participants differentiated between human-created and AI-generated personas, with the latter perceived as more informative and consistent. However, participants noted that the AI-generated personas tended to follow stereotypes, highlighting the need for a greater emphasis on diversity when utilizing LLMs for persona creation.

## CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**.

## Keywords

Personas, Large Language Models, User-Centered Design, Diversity

### ACM Reference Format:

Christopher Klaus Lazik, Christopher Katins, Charlotte Kauter, Jonas Jakob, Caroline Jay, Lars Grunske, and Thomas Kosch. 2025. The Impostor is Among Us: Can Large Language Models Capture the Complexity of Human Personas?. In *Mensch und Computer 2025 (MuC '25), August 31–September 03, 2025, Chemnitz, Germany*. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3743049.3743057>

## 1 Introduction

Personas represent fictional characters created based on user research to represent different user types that might similarly use a service, product, site, or brand. Personas encapsulate key information in a narrative-written style [28, 29] including demographic details, behaviors, goals mostly using text [30], and challenges of real users [11], helping Human-Computer Interaction (HCI) researchers keep the end user's requirements in mind during the design process [9]. They support researchers and designers in better understanding and predicting the hypothetical users' needs, motivations, and frustrations, leading to more user-centered and effective design solutions. Subsequently, using personas ensures that the final research prototype or product is tailored to meet the needs of its intended audience, enhancing usability and user satisfaction [25]. However, creating appropriate personas can be time-consuming, thus slowing down the design phase and becoming a costly process [18].

Consequently, the automatic generation of personas based on data became interesting in the HCI community. For example, previous research mentioned that data-driven personas could be automatically generated using data from social media [4], thus streamlining the persona generation process using crowd-sourced data. Recently, Large Language Models (LLMs) gained attention for simulating data of human participants in HCI research [3, 5, 8, 10, 12–17, 31, 49] and as a data analysis tool in HCI research [47], with LLMs subsequently becoming interesting for persona generation [39, 45]. LLMs can be used to create detailed personas, thanks to their ability to leverage their wealth of previously processed user data to generate realistic, contextually rich character profiles [45]. LLMs potentially allow for rapidly creating more comprehensive personas for various

scenarios. Several AI-powered commercial tools and scientific tools, such as PersonAI<sup>1</sup>, QoqoAI<sup>2</sup> or PersonaCraft [20] exist to assist in generating personas. Thus, the HCI community became interested in examining if LLMs can create personas that represent the desired user base in terms of their efficacy and accuracy, as well as their potential biases and ethical implications [38, 39, 42, 45], thus democratizing the complex persona creation process for novices and experts.

At the same time, researchers urge caution when using LLMs to generate user data [2]. Hallucinations, value lock-ins, training bias, deceptive design patterns created by LLMs [2, 21, 22] may generate personas whose descriptions are not perceived as diverse or realistic. The existing research on using LLMs to generate personas faces two primary limitations. Previous studies have either not directly compared the perceptual differences between AI-generated personas and those crafted by humans using neutral persona descriptions [39], or they have not thoroughly explored the key factors that affect how users discriminate between human-crafted and AI-generated personas [45]. In general, the quality of generated texts by LLMs can introduce a confirmation bias [27], especially to novices in a topic who have no expertise in persona creation. With the upcoming trend of using LLMs in persona generation, we see a potential threat to quality, specifically for those users who are not trained in human-centered design and seek fast solutions. Those novices might be misled by outputs that seem believable at first sight.

This paper addresses these limitations by conducting a user study to compare how users perceive human-crafted and AI-generated personas based on factors from previous research, including their informativeness for design, believability, stereotypicality, positivity, relatability, consistency, clarity, and likability [39, 41]. To this end, we surveyed participants to understand whether they can distinguish between human-crafted and AI-generated personas and which features in persona descriptions make them appear humane and credible. In the first step, we obtained a human-crafted persona dataset from HCI experts, who are aware of personas but do not frequently create them, to simulate a situation where non-experts develop personas. Then, we systematically created an AI-generated persona dataset. The consulted ten HCI experts crafted a persona based on factors relevant to persona design from previous work [7, 39, 41]. We systematically prompted OpenAI's GPT-4o to generate personas [39]. In the second step, we presented the personas to 54 layperson participants who subjectively rated personas as to whether they were human-crafted or AI-generated and according to additional factors that characterize the credibility of personas. Our results show that our participants can discriminate between human-crafted and AI-generated personas. However, we also found that specific characteristics, such as stereotypicality and writing style, help people to distinguish between human-crafted and AI-generated personas. Our work shows that although LLMs convincingly generate personas, they do not entirely cover the complexity of users and are thus a potential threat to quality if chosen based on the wrong properties.

<sup>1</sup><https://www.figma.com/community/plugin/1287786847239653675/personai-user-persona-generator> – last accessed 2025-04-10

<sup>2</sup><https://qoqo.ai/index.html> – last accessed 2025-04-10

## 2 Related Work

We based our research on previous work concerning the relevance of personas and their data-driven creation in HCI. We present a review of the relevant literature in the following.

### 2.1 Personas in Human-Computer Interaction

Personas have different purposes in different scientific fields. For example, personas are developed in requirements engineering to express target users [1]. In HCI, personas are fictional characters developed through user research to depict various types of people who may engage with a similar service, product, site, or brand. They convey essential details in a narrative format [28, 29], covering demographics, behaviors, and goals—primarily through text [30], as well as actual users' challenges [11], aiding Human-Computer Interaction (HCI) researchers in considering the needs of the end-users throughout the design process [9].

Designers can explicitly describe a fictional persona by encapsulating such information about potential user archetypes. Personas aim to offer an explicit common ground for developers to empathize with their users' needs without involving them directly in the design process [42]. Personas are also used in numerous other areas, such as defining potential customer groups in online marketing or in developing patient-oriented health and care technologies [26, 41]. Personas are traditionally created manually by experts in the field of HCI and user experience design. Yet, this is a time-consuming process that requires not only a deep understanding of users but also specialized expertise in developing fictional and representative user profiles that are realistic [43, 45]. Creating realistic personas is an important balancing act since lowly detailed and less informative personas may not be picked up by practitioners [24, 33].

Furthermore, the persona-creation process is influenced by the experience and perspectives of the experts involved, making it challenging to describe user groups in a truly representative manner, particularly for underrepresented users. Diversity and inclusion are central issues in persona creation. Diversity encompasses multidimensional differences in aspects of human beings (such as gender, age, ethnicity, neurodiversity, cultural background, and many more), while inclusion emphasizes actively involving various user groups to address their needs better. Failure to account for these aspects may result in products or systems that are not optimally usable for the broader user base. One approach specifically addressing gender diversity is the GenderMag method [7]. These gender-specific differences affect how software is used and perceived. To mitigate these issues, the GenderMag approach uses fictional personas based on five cognitive dimensions: motivation, information processing style, computer self-efficacy, risk aversion, and tinkering. This helps identify potential usage barriers and fosters more inclusive design. In addition to addressing gender-related diversity, inclusion can be promoted through co-created personas. This approach directly engages users in the persona creation process, capturing their real needs and challenges to develop more accurate and practical personas. For example, in healthcare design, co-created personas have been used to reflect patients' experiences with Parkinson's disease, dementia, or aphasia [26]. By incorporating their perspectives, designers can create solutions that better meet the unique needs of these user groups. Despite the promise of these methods, persona

creation remains a time-intensive and complex process. Hard-to-reach populations, such as individuals with autism or special needs, are often underrepresented, making it difficult to develop personas that comprehensively reflect diverse user groups [26].

### 2.2 Automated Persona Generation

Data-driven personas are profiles created using real data and analytics to represent key segments of users or customers, offering opportunities to enhance personalized marketing, product development, and user experience by making decisions based on concrete insights [19, 37]. Considering the time and resources required to create personas [32, 34, 48], automated approaches for persona creation have recently been proposed [36]. In this context, LLMs offer promising potential. With their ability to process large amounts of text, LLMs could help close knowledge gaps and simulate the requirements of hard-to-reach user groups [43].

A precise evaluation method is required to determine whether AI-generated personas realistically represent target groups. One potential tool is the persona perception scale [41], which provides a metric for measuring eight key dimensions, including credibility, consistency, and the willingness to use a persona. It enables the comparison of AI-generated and human-crafted personas to ensure that both types can be used effectively in the design process. Salmiinen et al. [39] showed that users perceive AI-generated personas as realistic. The study found that these personas are primarily seen as consistent, credible, and informative [39]. This highlights the potential of LLMs to generate personas perceived as believable, relatable, and informative while minimizing stereotyping, making them suitable for use in the design process [39].

In this context, Schmidt et al. [43] emphasized the potential of LLMs in human-centered design. They argued that human resources should be used more efficiently and that tasks that can be handled just as well or even better by LLMs should be left to them. They also highlighted the importance of transparency: it has to be communicated whether and how LLMs are used in the design process to ensure trust and accountability [43].

Agnew et al. [2] warned that although LLMs have gained popularity for simulating human participants, they can rarely represent real users' profound experiences and perspectives. Important facets of human interaction, such as emotions, cultural contexts, or non-verbal communication, may not be captured by LLMs. Kosch and Feger [21] pointed out problems such as a lack of reproducibility, as LLMs do not consistently deliver the same results for a given input [21]. There is also a risk of bias, as LLMs are based on data that contains cultural and social biases, which can lead to skewed or stereotypical personas. Another risk is "*value lock-in*", where old norms and values are incorporated into the generated personas and prevent progress. It should be underlined that a transparent and responsible use of LLMs is necessary to recognize and minimize these risks early.

To improve the reproducibility and quality of AI-generated personas, [39] presented a structured prompt strategy. This includes a clear definition of the desired persona characteristics such as age, gender, and occupation to guide the output of the LLM. Iterative customization of the prompts is used to refine the generated personas and ensure consistency. Finally, a review and expansion of

the “skeletal” personas is conducted to create complete narrative descriptions that are realistic and consistent. This strategy increases the reproducibility of the generated personas by ensuring that the inputs are aligned with the desired outcomes and that the personas are systematically improved. Additionally, the researchers suggest following three guidelines when implementing LLMs into the design process of personas: Verifying the personas regarding diversity and bias, verifying the personas using subject-matter experts, and adjusting the prompts if you observe challenges. To this end, Jung et al. [20] presented PersonaCraft, a tool to generate personas based on human survey data.

Schuller et al. [45] examined the perception of AI-generated and human-crafted personas and found no significant differences in perceived quality between the two types of personas. However, the study does not consider explicit factors that influence the perception of personas, such as those defined in the persona perception scale [39, 41], focusing instead on subjective impressions such as language and style [45]. In fact, Shin et al. [46] found that LLMs are not optimal at capturing key characteristics when generating personas. The authors suggested that humans should group data into persona-relevant categories beforehand and then use LLMs to summarize the data into personas.

In summary, personas are an essential asset for capturing the requirements in user interface design. Related work suggested accelerating the persona creation process through the support of LLMs [43]. Previous work found that LLMs capture the requirements of generated personas well [39] to the degree that users may not be able to distinguish between human-crafted and AI-generated personas [44]. Yet, previous work did not investigate how common persona-driven features, such as informativeness, reliability, consistency, or clarity [39, 41], are perceived differently between human-crafted and AI-generated personas. Our paper closes this gap by conducting a survey study investigating which features make personas appear humane or generated.

### 3 Methodology

Previous studies have demonstrated the significance of personas in user experience design. As a result, research has explored the potential of using LLMs to expedite the persona creation process, suggesting that including simulated user data could enhance the diversity of the generated personas. Nevertheless, the rise of LLMs might lead to the assumption of accelerating development processes by directly generating personas instead of crafting them.

Prior work has not examined whether users can distinguish between human-crafted and AI-generated personas based on common perceptual factors contributing to a persona’s credibility. Consequently, this paper is influenced by earlier research that focused exclusively on AI-generated personas [39] or studies that compared human-crafted and AI-generated personas but provided limited evaluation of key perception variables [45] such as consistency, completeness, or clarity [41]. Therefore, we define the following research questions:

**RQ1:** To what extent can users distinguish between human and AI-generated personas?

**RQ2:** How do the perceived differences between human-created and AI-generated personas relate to specific features of the persona description?

**RQ2.1:** What are the differences between human-crafted and AI-generated personas regarding quality aspects from the literature?

**RQ2.2:** How do participants characterize the difference between human-crafted and AI-generated personas?

We conducted one data preparation step and a persona-comparison study to answer the research questions. In the first step, we asked ten HCI experts to craft one persona each, resulting in ten human-crafted personas that were not part of any training set from an LLM. Then, we utilized a prompting strategy by Salminen et al. [39] to generate ten personas using an LLM. The user study directly compared the human-crafted personas to generated personas by exposing participants to personas and asking them to assess if a human or an AI created the persona. Furthermore, we employed several constructs used by Salminen et al. [39, 41]. We specifically focused on a subset of these dimensions most relevant to understanding the perceived realism, coherence, and usefulness of personas: informativeness, believability, stereotypicality, positivity, relatability, consistency, clarity, and likability. We omitted dimensions tied directly to designer-specific tasks, such as completeness and willingness to use [41], since our participant pool was not composed of active design professionals applying personas in real-world projects. Additionally, similarity and empathy were excluded because they involve personal relatability to the persona, which could introduce confounding variables unrelated to the primary research questions. We obtained ethical approval for the studies from the institutional review board of our university.

### 4 Obtaining Human-Crafted Personas

Our first data acquisition step was designed to provide a baseline of human-crafted personas. We created new personas by collaborating with researchers and practitioners with several years of HCI experience in the form of a survey. This approach ensures we have a controlled set of personas free from external context-specific biases, unlike those that might arise from using personas from an online resource that targets a specific use or has an intent. We want to investigate the appeal of personas instead of introducing a bias through particular use cases that might influence the appeal and freedom of the created personas. Furthermore, personas that were available online can be in the training corpus of LLMs that we used in our study to generate personas. By obtaining a new set of human-crafted personas, we ensured that we would compare them against AI-generated personas that were not part of the training set of any LLM before. We created an online survey to collect personas from HCI experts using persona properties recommended by the GenderMag project [7]<sup>3</sup> and Salminen et al. [39]. We did not specify a persona complexity level for the participants, allowing creative freedom.

The choice of asking general HCI experts instead of more specific persona design experts ensures that our participants are familiar with the concept of personas while being less far from novices.

<sup>3</sup>[https://gendermag.org/custom\\_persona.php](https://gendermag.org/custom_persona.php) – last accessed 2025-04-10

## 4.1 Survey Structure

The survey aims to obtain human-crafted personas that were previously not used to train LLMs or are the results of an LLM. Furthermore, the set of crafted personas should be comparable to personas made by knowledgeable novice users. Our survey for collecting human-crafted personas consisted of two parts. First, we explained the course of the survey to the participants and asked for informed consent. Second, we gave a template to our participants as shown in Table 1, including the following features for a fictional persona. The properties considered included name, age, occupation, background and skills, motivation and strategies, technological attitude, and other details characterizing the persona. We chose these properties from the previous work by Salminen et al. [39] on generating personas using LLMs and completed the properties using the recommendations by the GenderMag project [7]. These attributes provide the basic information about a general persona. While their name, age, and occupation give a general context of the persona's circumstances, the attribute called "details" includes a description of the persona with more details to provide context. These general details allow the person using the persona to understand the general context of the described fictional person. To enable the usage of the personas in our study for software development processes, we additionally included background and skills, motivation and strategies, and attitude to technology. These attributes are based on the GenderMag concept. Aiming to address the general approach to challenges and technology, they are a fitting extension to the more general personas.

We told all participants not to use generative AI, specifically LLMs, in their persona creation process. We also asked the participants to provide the information in sentences for the last four properties the participants should describe (i.e., background and skills, motivation and strategies, technological attitude, and other details). Even though personas might include a picture [1], we did not ask the participants to provide one to keep comparability to AI-generated personas. Furthermore, previous research showed that including pictures introduces biases within the individual assessment and perception of personas [40]. The findings show that contextual photos enhance informativeness. At the same time, images of different people confuse, and user biases influence how personas are interpreted, suggesting a careful selection of photos in persona design. As a consequence, we refrained from using photographs in the persona descriptions. We conducted the survey using LimeSurvey<sup>4</sup>, hosted by our research institute.

## 4.2 Procedure

We greeted the participants and explained the survey's objectives to them. Following this, the participants completed an informed consent form. We informed the participants that their participation is entirely voluntary and their right to withdraw their participation and data anytime without any disadvantages from the survey. Participants were informed that they were not receiving compensation. We collected demographic information from the participants, including their age, self-identified gender, profession, research focus, experience with personas, and years of experience in HCI research or user-centered design. Participants were also asked to select a

pseudonym, which could be used to request data deletion through a pseudonymization list. They were then provided with a comprehensive task description detailing the persona creation process, which included the specified characteristics of the persona. The properties were concerned with background and skills, motivation and strategies, technological attitude, and other details that describe the personality of the persona. The participants filled in the persona details without any constraints, except for the requirements to answer in complete sentences and avoid using generative AI. Lastly, they were allowed to add any additional comments to the survey.

## 4.3 Participants

We recruited ten participants (five self-identified as female, four self-identified as male, and one self-identified as non-binary) with ages ranging from 23 to 36 ( $\bar{x} = 30.20$ ,  $s = 3.39$ ). All of our participants had a background in HCI research with an average experience of 4.65 years ( $s = 2.69$ ). While the concept of personas was clear to all ten participants, four reported creating personas multiple times during a year, one reported creating personas various times during a month, and five reported that they had never made a persona before. We explained the survey's purpose to the participants and informed them they could cancel it anytime.

## 4.4 Results

We collected a total of ten human-crafted personas. Two researchers rated and screened the set of personas regarding their plausibility. Then, both researchers independently coded the personas using the same coding sheet proposed by Salminen et al. [39]. The rating criteria are displayed in Table 2. We calculated Krippendorff's  $\alpha$  as a measure for the agreement of the codes towards the criteria between both researchers. We found that Krippendorff's  $\alpha$  resulted in  $\alpha = 0.879$ , meaning that sufficient agreement was reached<sup>5</sup>. The age of the human-crafted personas ranged from 10 to 72 ( $\bar{x} = 51.40$ ,  $s = 19.26$ ). The HCI experts created seven male, two female, and one non-binary persona. The text length, counted in words, ranged from 71 to 433 words ( $\bar{x} = 260.7$ ,  $s = 109.75$ ). While the physical appearance was only described in two personas, the personality was expressed in every persona. The occupations of the human-crafted personas were specified as retiree, artist, unemployed, firefighter, pupil, faculty member, researcher and director, and gardener. Retirees were chosen multiple times since the age of some personas was relatively high. To ensure no impact of typographical errors in the human-crafted personas, we corrected spelling mistakes across all human-crafted personas. This is because typographical errors would reveal that a human created a persona for later comparison with AI-generated personas.

## 5 Evaluating the Perception Between Human-Crafted and AI-Generated Personas

Using the personas from the first data acquisition step as a set of human-crafted personas, we conducted a second survey to determine the properties that distinguish between human and AI-generated personas. To directly compare human-crafted personas and personas created using an LLM, we asked participants to label

<sup>4</sup><https://www.limesurvey.org>— last accessed 2025-04-10

<sup>5</sup>A Krippendorff's  $\alpha$  of  $\geq .800$  is considered high agreement between the coders.

**Table 1: Survey template for participants who created human-crafted personas. Our participants filled in the fields that ultimately described the persona.**

Attribute	Explanation
<b>Name</b>	The name of the persona you want to craft.
<b>Age</b>	The age of the persona you want to craft.
<b>Occupation</b>	What is the persona you are creating doing for a living?
<b>Background and Skills</b>	What knowledge and skills does the persona that you are creating have? This may include educational qualifications or knowledge in tools. Feel free to include every skill or other background you think is part of the identity of your hand-crafted persona.
<b>Motivation and Strategies</b>	What motivates your created persona during their everyday life? How do they approach challenges and tasks?
<b>Attitude to Technology</b>	What does your created persona think about technology?
<b>Details</b>	Describe your created persona a little bit more to give more context of their personality.

**Table 2: Evaluation sheet for assessing the qualitative agreement of the human-crafted personas. The attributes of the evaluation sheet were adapted from previous work [39].**

Attribute	Explanation	Code
<b>Age</b>	The persona’s age	number
<b>Gender</b>	The persona’s gender	'f', 'm', 'nb', 'o'
<b>Occupation</b>	The persona’s occupation	job title
<b>Text Length</b>	The length of the persona’s properties that were given in sentences	word count
<b>Physical Appearance?</b>	Is the persona description mentioning the persona’s physical appearance?	yes or no
<b>Personality Mentioned?</b>	Is the persona’s personality described?	yes or no
<b>Informativeness for Design</b>	Does the persona description contain adequate information to design an app or system to address the persona’s needs?	yes or no
<b>Believability</b>	Does the persona appear realistic, i.e., lifelike, like an actual person that could exist?	yes or no
<b>Stereotypicality</b>	Does the persona appear stereotypical? (Stereotypes are related to a widely held but fixed and oversimplified image or idea of a particular type of person or thing.)	yes or no
<b>Positivity</b>	Is the person depicted in a positive light?	yes or no
<b>Relatability</b>	Is the persona relatable? (Relatability is the quality of being easy to understand or feel empathy for.)	yes or no
<b>Consistency</b>	Is the persona consistent? (A consistent persona does not have conflicting information e.g., the person is described as being overall happy but later as being sad in general.)	yes or no

a mixed set of generated and human-crafted personas. This section explains the persona-generation process, the survey methodology, and the results.

The evaluation of personas in our study was conducted with survey participants who were not explicitly designers to explore broader perceptions of realism, informativeness, and stereotypicality, important attributes linked to persona effectiveness. While personas are ultimately used by designers, understanding how non-designers perceive the authenticity and detail of persona descriptions provides foundational insights into their general credibility and appeal. This approach aligns with prior research that examines personas from a user-perception perspective to assess their believability, perceived utility, and credibility of traits [35]. By gathering input from individuals who could represent diverse perspectives,

we want to identify generalizable traits that influence persona acceptance.

## 5.1 Generating Personas using AI

We used GPT-4o<sup>6</sup> to generate personas, using a zero-shot structured prompting strategy from related work [39]. We used a zero-shot prompting strategy to emulate a typical real-world scenario where users generate personas with minimal customization or prior examples, reflecting the baseline capabilities of the LLM. To avoid the token limit given by GPT-4o<sup>7</sup>, the approach begins with the generation of a skeleton. These skeletal personas include the personas’ properties (e.g., name, age, occupation, and details). After creating a set of ten skeletal personas, we iterate over the persona

<sup>6</sup><https://openai.com/index/hello-gpt-4o> – last accessed 2025-04-10

<sup>7</sup>GPT-4o encounters a token limit of 4096 tokens.

skeletons and prompt GPT-4o to expand each skeletal persona into a complete one. This approach generates the main aspects of all the personas at once that are **structured** in a way that is **comparable** to both our human-crafted personas and other generated personas from the literature. Furthermore, to achieve personas that provide information that can be **used** to develop software, we expanded the main properties with three additional elements (background and skills, motivation and strategies, and attitude to technology) from GenderMag [7]. Thus, the structure of our personas asks for general information about a fictional person and their knowledge base, how they approach challenges, and how they deal with technology. The resulting generation script can be found in our supplementary material.

We used two prompts, one for each generation step:

- I “Generate 10 user personas. Provide the output in a json array, with each dict containing only the following keys: "index", "name", "age", "occupation", "background and skills", "motivations and strategies", "attitude to technology", "details"; Index starts at 1. Start your response with the open square bracket [”
- II “Expand on the following summary persona. Ensure that all the information provided is used in your expanded persona. Don’t include more or less properties. Stick to the structure that is given.” + the JSON-file

We generated ten final personas for this study, which were generated in one round. We did not generate and evaluate multiple personas for a good fit for the study. First, previous work [39] showed that generated personas appear credible when only prompted once. Second, we did not intend to bias the provided AI-generated persons through a pre-selection since a realistic use case is the spontaneous generation of personas using an LLM without further screening.

## 5.2 Persona Generation Results

Similarly to the first data acquisition survey, two experts from our research team reviewed the generated personas. Again, we use Krippendorff’s  $\alpha$  and find a high agreement with an  $\alpha = 0.949$ .

The age of the generated personas ranged from 27 to 50 ( $\bar{x} = 35.90, s = 7.37$ ). GPT-4o created five male, five female, and zero non-binary personas. During the skeletal persona generation, the gender of the personas was consistently alternating. The text length of the given sentences, in words, ranged from 145 to 254 words ( $\bar{x} = 191.10, s = 34.07$ ). No physical appearance was described for any of the personas. Every persona was defined by personality. The occupations of the generated personas were specified as software engineer, digital marketer, product manager, UX designer, data analyst, HR manager, CEO, sales manager, technical support, and business analyst. Every generated persona had a unique occupation. However, all listed occupations fall within the white-collar category and are settled in the technology, business, and management sectors.

Regarding the criteria in the lower part of Table 2 (i.e., the last eight items), our coders mainly answered with “yes”. The personas were informative, believable, positive, relatable, and consistent. However, all of them appeared to be stereotypical to our experts. Furthermore, our experts noticed that many personas are described

as performing similar thoughts in their free time as they do in their jobs.

## 5.3 Survey Structure

The survey study consisted of two parts. First, the participants were informed of the study goals and provided informed consent. Second, we collected demographic information of our participants and assessed their AI literacy employing the “Scale for the assessment of non-experts’ AI literacy” [23]. Then, we informed the participants about their task of rating a mixed set of AI-generated and human-crafted personas and asked them to rate our 20 randomly ordered personas. Then, participants performed the task. They rated several statements on a 7-point Likert scale<sup>8</sup> and provided reasons for their ratings in a free text field. We compensated our participants with nine Pounds per hour. The participants knew that personas were either human-crafted or AI-generated but were unaware of which personas were human-crafted or AI-generated throughout the survey. The survey was created using LimeSurvey<sup>9</sup>, hosted by our research institute.

## 5.4 Task

In our survey study, the participants had to solve one main task. Each participant was exposed to a set of 20 personas. The order of these personas was randomized to prevent possible effects. Ten out of the 20 personas were generated using Open AI’s GPT-4o. The other ten personas were handcrafted by experts from the field of HCI in our first data acquisition survey. Participants were given the personas as a structured text description as shown in Figure 2. By displaying the personas in a neutral way, we aimed to avoid revealing whether a persona was human-crafted like shown in Figure 2a or AI-generated as seen in Figure 2b. Participants were exposed to 15 statements that they had to rate on a 7-point Likert scale as shown in Table 3.

The first two statements were about the origin of the persona. By rating them, participants decided whether they think a persona is human-crafted or AI-generated. After that decision, participants rated multiple statements to assess the design, believability, stereotypicality, positivity, relatability, and consistency of the personas. These aspects of personas were based on the literature by Salmiinen et al. [39]. Additionally, we utilize further questions from the persona perception scale [41]. We used the constructs “Clarity” to determine if the persona description was clear to the participants. Furthermore, we measure the construct “Likability” to further assess how personas appeal to participants in our survey. We did not include the constructs “Consistency” and “Credibility” since the constructs contained items asking for a profile picture. Yet, none of our personas contained a picture to avoid biases regarding the relatability [40]. We excluded “Completeness” and “Willingness to Use” since they ask for items regarding a target group (e.g., managers or designers). Furthermore, we excluded “Similarity” and “Empathy” since both ask for items personally tied to the participant and their relation to the persona. We asked each participant why they assessed a persona as human-crafted or AI-generated and

<sup>8</sup>1: Strongly Disagree; 2: Disagree; 3: Somewhat Disagree; 4: Neutral; 5: Somewhat Agree; 6: Agree; 7: Strongly Agree.

<sup>9</sup><https://www.limesurvey.org/> – last accessed 2025-04-10

**Table 3: Statements ranked by participants to measure the persona aspects. Participants rated the persona aspects on a 7-point Likert scale.**

Persona Aspect	Statement
<b>Human-Crafted</b>	The persona is human-crafted.
<b>AI-Generated</b>	The persona is AI-generated.
<b>Informativeness for Design</b>	The persona description contains adequate information to design an app or system to address the persona's needs.
<b>Believability</b>	The persona appears realistic, i.e., lifelike, like an actual person that could exist.
<b>Stereotypicality</b>	The persona appears stereotypical. (Stereotypes are related to a widely held but fixed and oversimplified image or idea of a particular type of person or thing.)
<b>Positivity</b>	The person is depicted in a positive light.
<b>Relatability</b>	The persona is relatable. (Relatability is the quality of being easy to understand or feel empathy for.)
<b>Consistency</b>	The persona is consistent. (A consistent persona does not have conflicting information e.g. the person is described as being overall happy but later as being sad in general.)
<b>Clarity<sub>1</sub></b>	The information about the persona is well presented.
<b>Clarity<sub>2</sub></b>	The text in the persona profile is clear enough to read.
<b>Clarity<sub>3</sub></b>	The information in the persona profile is easy to understand.
<b>Likability<sub>1</sub></b>	I find this persona likable.
<b>Likability<sub>2</sub></b>	I could be friends with this persona.
<b>Likability<sub>3</sub></b>	This persona feels like someone I could spend time with.
<b>Likability<sub>4</sub></b>	This persona is interesting.

if they were undecided or conflicted about their choice. After rating the above-mentioned items of a persona, the participants were asked in an open text field what aspects of the persona influenced their decision.

## 5.5 Procedure

The participants were introduced to the task after agreeing to participate in our study. Participants were informed about their task of rating generated and handcrafted personas. The participants had to iterate over twenty personas in a randomized order. After rating all the personas in our set, the participants were asked to provide qualitative feedback on the study experience to collect information on how participants perceived the task over the experiment. In the last step, we debriefed the participants, and they got confirmation that they finished the questionnaire.

## 5.6 Participants

We sampled participants through the online platform Prolific<sup>10</sup>. We recruited participants until we obtained 54 participants who passed two validation questions.

We collected the answers of 54 participants (25 self-identified as female, 27 self-identified as male, and two self-identified as non-binary) aged 18 to 73 ( $\bar{x} = 32.04$ ,  $s = 10.51$ ). The sample was heterogeneous regarding the occupation of the participants. The participants reported using LLMs in different frequencies (11 “Never”, 10 “Multiple times during a year”, 15 “Multiple times during a month”,

11 “Multiple times during a week”, and 7 “Every day”). Participants responded to the statement “I am familiar with the concept of personas.” on a 7-point Likert scale. The mean response was 5.5 ( $s = 1.28$ ), indicating that most respondents agreed with the statement. Additionally, we asked the participants to rate the statement “I am familiar with crafting personas.” on a 7-point Likert scale. Here, the average response was 4.39 ( $s = 1.66$ ), which shows that most participants were neutral with regard to the statement. Nevertheless, the participants answered that question generally more positively.

To assess the AI literacy of our participants, we employed the “Scale for the assessment of non-experts’ AI literacy” [23]. The participants responded to three key areas: Practical Application, Technical Understanding, and Critical Appraisal. On average, Practical Application scored a mean of 4.65 ( $s = 0.66$ ), Technical Understanding received a mean score of 3.71 ( $s = 0.46$ ), and Critical Appraisal had the highest mean score of 5.03 ( $s = 0.20$ ). On average, our participants seem to be able to apply AI practically and approach it critically but do not entirely grasp all its technical aspects.

## 5.7 Quantitative Results

This section presents the findings of our study. We conducted Wilcoxon signed-rank tests to statistically compare the Likert ratings between human-designed and AI-generated personas. The statements that were ranked by our participants for each persona are shown in Table 3. We averaged the Likert item ratings across personas for each participant to get a single value for human-crafted and generated personas, allowing for statistical comparison. We

<sup>10</sup><https://www.prolific.com> – last accessed 2025-04-10



**Name:** Henriette Bauer

**Age:** 72

**Occupation:** The person is retired and worked part-time in an office as a secretary before she retired.

**Background and skills:** Henriette is a secretary by training. However, when she started working, computers were not a thing yet, so she is mainly versed in typewriters and very unsure of computers and technology in general. Mostly, her husband Hubert takes care of any technology in their household, and she is always a bit scared of making mistakes or falling victim to scams. However, she is a very engaged gardener and an amazing baker and cook. Her biggest hobbies are baking and tableware.

**Motivation and Strategies:** Henriette is an avid homemaker and loves taking care of her family especially her husband. She is happiest at home and in her garden, where she grows her own vegetables. She takes care of all household chores. However, anything that is related to technology in her household is taken care of by her husband. This includes the TV recorder and ordering anything on the web.

**Attitude to Technology:** She is very unsure about technology and scared of making mistakes since she read a lot about scams in the newspaper and saw a lot about them on the news. She can and likes to use the telephone and the TV to start her favorite program. However, for everything beyond that, she asks her husband for help. Once, her daughter gifted her a smartphone, but she never used it, and it has been in a drawer for 3 years.

**Details:** Henriette is a very positive person. She stayed young, considering her age, and is also very focused on her beauty. She wears makeup and perfume every day and goes to the hairdresser every 4 weeks. Every Friday, she goes to the city and goes shopping. She does not like cats but loves birds. She always feeds them in her garden.

(a)

**Name:** Jane Doe

**Age:** 28

**Occupation:** Digital Marketer

**Background and Skills:** Jane Doe holds an MBA in Marketing, which provides her with a strong foundation in strategic thinking and business acumen. Over the past 5 years, she has honed her expertise in various digital marketing domains, including Search Engine Optimization (SEO), Pay-Per-Click (PPC) advertising, and content marketing. Jane is recognized for her strong analytical skills, enabling her to interpret data effectively and make informed decisions to optimize marketing campaigns.

**Motivations and Strategies:** Jane's primary focus is on increasing brand visibility and revenue for the companies she works with. She is dedicated to implementing cost-effective digital marketing strategies that yield high returns on investment. Her approach involves continuous learning and adapting to the latest trends to maintain an edge in the competitive digital landscape.

**Attitude to Technology:** Jane is open and adaptable when it comes to technology. She embraces new tools and platforms that can enhance her digital marketing efforts. Her willingness to experiment with the latest technological advancements allows her to stay ahead of industry curves.

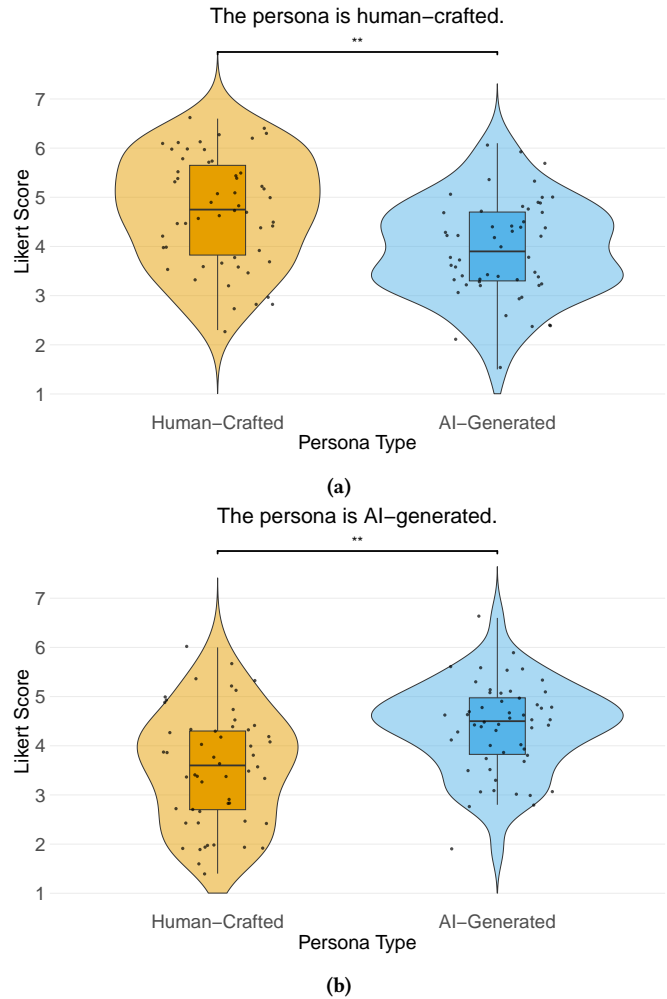
**Details:** Outside of her professional responsibilities, Jane enjoys reading about the latest industry trends and best practices. This passion for learning extends into her free time, where she likes to experiment with new digital marketing strategies. Through continuous education and practical application, Jane remains a knowledgeable and innovative professional in her field.

(b)

**Figure 2: Two example personas from our survey study. (a): human-crafted. (b): AI-generated**

conducted a Wilcoxon signed-rank as a non-parametric test to compare the Likert ratings of the participants. The significance level was set at  $\alpha = .05$ . We also report the test statistics  $V$ ,  $Z$ , and the effect size  $r^{11}$ . We calculated the score of the constructs “Clarity” and “Likability” by calculating the mean score of the questions as suggested by the persona perception scale [41].

<sup>11</sup>Generally,  $r = 0.1$  is considered a small,  $r = 0.3$  a medium, and  $r = 0.5$  a large effect.




**Figure 3: Violin plot comparing if participants could distinguish between human-crafted and AI-generated personas. (a): Participants could recognize human-crafted personas. (b): Participants could recognize AI-generated personas. Asterisk denote significant differences.**


**5.7.1 Distinguishing Between Human-Crafted and Generated Personas.** A Wilcoxon signed-rank test was performed to assess whether participants could distinguish between human personas and those generated by an LLM. Ratings were analyzed separately for human-crafted and AI-generated personas to enable direct comparisons. The results showed a significant difference between the two types of personas, with  $V = 1088.5$ ,  $Z = 2.98$ ,  $p = .003$ ,  $r = 0.405$ . Human-crafted personas received higher Likert scores regarding the Likert statement that describes them as human-crafted, indicating that participants thought they were indeed human-crafted (see Figure 3a). Consistently, we found a significant difference when participants encountered an AI-generated persona, with  $V = 372.5$ ,  $Z = -3.05$ ,  $p = .002$ ,  $r = -0.415$  (see Figure 3b). Consequently, AI-generated personas scored higher for the corresponding Likert statement saying that they are AI-generated.

**5.7.2 Informativeness for Design.** The Wilcoxon signed-rank test indicated a significant difference in the perceived informativeness between human-crafted and AI-generated personas,  $V = 192$ ,  $Z = -4.34$ ,  $p < .001$ ,  $r = -0.597$ , with participants giving higher Likert scores to the generated personas when comparing them to the human-crafted ones (see Figure 4a).

**5.7.3 Believability.** The Wilcoxon signed-rank test showed no significant difference in the believability of human-crafted versus generated personas,  $V = 667$ ,  $Z = 0.05$ ,  $p = .96$ ,  $r = 0.006$  (see Figure 4b).



**Figure 4: Violin plot comparing how participants rated the informativeness and the realism of human-crafted and AI-generated personas. (a):** Participants rated the informativeness of AI-generated personas higher than that of human-crafted personas. We abbreviated the title of the figure with the persona aspect from Table 3 to increase the readability. **(b):** There was no significant difference between participants' ratings about the realism of AI-generated and human-crafted personas. Asterisk denote significant differences.



**Figure 5: Violin plot comparing how participants rated the stereotypicality and the positivity of human-crafted and AI-generated personas. (a):** Participants rated the stereotypicality of AI-generated personas higher than that of human-crafted personas. **(b):** Participants rated the positivity of AI-generated personas higher than of human-crafted personas. Asterisk denote significant differences.


**5.7.4 Stereotypicality.** The Wilcoxon signed-rank test indicated a significant difference in the stereotypicality of human-crafted versus generated personas,  $V = 356.5$ ,  $Z = -2.24$ ,  $p = .03$ ,  $r = -0.305$ . Generated personas received higher Likert scores than human-crafted ones (see Figure 5a).

**5.7.5 Positivity.** The Wilcoxon signed-rank test also revealed a significant difference in the positivity of human-crafted and generated personas,  $V = 203.5$ ,  $Z = -4.31$ ,  $p < .001$ ,  $r = -0.586$ . Likert scores were higher for AI-generated personas than human-crafted ones (see Figure 5b).

**5.7.6 Reliability.** The Wilcoxon signed-rank test did not indicate a significant difference in the believability between human-crafted

and AI-generated personas,  $V = 593.5$ ,  $Z = -0.16$ ,  $p = .85$ ,  $r = -0.025$  (see Figure 6a).

**5.7.7 Consistency.** The Wilcoxon signed-rank test showed a significant difference in the consistency of human-crafted versus AI-generated personas,  $V = 155.5$ ,  $Z = -4.40$ ,  $p < .001$ ,  $r = -0.599$ , with participants rating the AI-generated personas higher in consistency (see Figure 6b).



**Figure 6: Violin plot comparing how participants rated the relatability and the consistency of human-crafted and AI-generated personas. (a): Participants rated the relatability of AI-generated personas higher than that of human-crafted personas. (b): Participants rated the consistency of AI-generated personas higher than of human-crafted personas. Asterisk denote significant differences.**

**5.7.8 Clarity.** The Wilcoxon signed-rank test showed a significant difference in the clarity between human-crafted and generated personas,  $V = 223.5$ ,  $Z = -4.08$ ,  $p < .001$ ,  $r = -0.556$ , with higher Likert scores for AI-generated personas (see Figure 7a).

**5.7.9 Likability.** The Wilcoxon signed-rank test indicated no significant difference in likability between human-crafted and generated personas,  $V = 672$ ,  $Z = -0.38$ ,  $p = .71$ ,  $r = -0.052$  (see Figure 7b).



**Figure 7: Violin plot comparing how participants rated the clarity and the likability of human-crafted and AI-generated personas. (a): Participants rated the clarity of AI-generated personas higher than that of human-crafted personas. (b): There was no significant difference between the ratings of participants about the likability of AI-generated and human-crafted personas. Asterisk denote significant differences.**

## 5.8 Qualitative Results

We started the qualitative analysis by categorizing the free-text responses based on their classification as human-crafted, AI-generated, neutral, or conflicted. Then, we applied a second tag to indicate the correctness of the participant's classification. For instance, if a participant labeled a persona as human-crafted when it was AI-generated, we marked the classification as incorrect. Conversely, a tag indicating a correct choice was applied when the participant

correctly identified the origin of a persona. Neutral classifications were assigned if participants rated both “The persona is human-crafted” and “The persona is AI-generated” with a score of four. A conflicted tag was given when both questions were rated below or above four. This tagging system allowed us to code the free-text answers about the accuracy of the participants’ choices.

Two authors coded the free-text responses. We randomly assigned 25% of the written responses from when participants encountered human-crafted and AI-generated content. Following Blandford et al. [6], we applied an inductive coding approach to generate an initial set of codes and construct coding trees. Afterward, we conducted a code adjustment session where we reviewed the initial coding tree and distributed the remaining free-text responses among three coders. Once the entire data set was coded, we engaged in iterative discussions using axial coding to refine the final coding tree and identify recurring patterns. We extracted the five themes *Writing*, *Information*, *Stereotypicality*, *Realism and Appeal*, and *Positive and Negative Appeal* from our coding process, which captured participants’ perceptions of personas as either real, artificial, or undecided. Given the exploratory nature of our study, we determined that an open-ended thematic analysis, as outlined by Blandford et al. [6], was the most appropriate approach. This analysis method, rooted in interpretivism, is commonly employed in HCI research.

**5.8.1 Writing.** In the evaluation of personas generated by LLMs, one of the recurring themes identified in participants’ responses was the distinctive writing style of the personas. Participants often pointed to specific elements such as sentence structure, grammatical errors, and vocabulary choice when forming their impressions. These writing characteristics played a key role in shaping how the participants perceived the personas, often serving as subtle indicators of whether the text was likely human-generated or AI-generated:

*“When compared to the previous personas that are believed to be AI generated, there is a striking resemblance in the words and phrases used, such as “deep-seated”, “Outside of her professional life,” and “global understanding”. This is no coincidence; I strongly believe these are all connected due to the work of Artificial Intelligence.”* (P15, AI-generated persona correctly rated)

While many participants cited the writing style as a key reason for their rating of the personas, their interpretations of specific elements, such as grammatical errors, varied. Some participants argued that the presence of grammatical mistakes indicated that the text was likely human-generated, as they believed an AI would not produce such errors:

*“Again found some grammar issues within the text AI wouldn’t do this.”* (P18, human-crafted persona correctly rated)

On the other hand, others suggested the opposite, claiming that AI and LLMs do not generate grammatically flawless text, thus implying that errors were more characteristic of AI writing. These contrasting views highlight the complexity of how participants perceived AI-generated content, with grammatical accuracy being a point of contention in determining whether the writing felt more

human- or machine-like. Consequently, other participants were misguided by grammatical errors:

*“way too many errors in the description give away that it is not human generated”* (P29, human-crafted persona falsely rated)

Overall, participants identified the language used in the personas as a common factor influencing their ratings. There was a consensus among participants that AI-generated personas often adopted a “robotic tone,” which stood out in contrast to more natural human speech. The vocabulary used by these personas was frequently described as “unusual” or somewhat disconnected from everyday language, further reinforcing the perception that the text lacked human-like fluidity and warmth. This shared impression of the AI personas’ language as overly formal, rigid, or unnatural contributed to a sense of distance between the personas and the participants, shaping their overall assessment of the personas’ authenticity and reliability:

*“I believe this was AI-generated due to the use of some words such as “honed”, “permeating” and “further cementing” that are not commonly used in conversations.”* (P22, AI-generated persona correctly rated)

Furthermore, participants noted that human-crafted personas had a more conversational and relatable tone, resembling how a person would naturally speak to someone else. They emphasized that these personas felt more personal and engaging, with language that flowed in a way that mirrored typical human interactions. This conversational style made the human-crafted personas appear more authentic, as if they genuinely addressed the reader. In contrast to the AI-generated personas, often described as robotic or detached, human-crafted personas were seen as more capable of forming a connection through their more intuitive and familiar language:

*“Definitely human written, no doubt in my mind. The use of language here reads like they’re speaking to someone while they’re writing it down- more of an inner monologue, or like someone was describing someone they knew to someone else. [...]”* (P35 human-crafted persona correctly rated)

The participants noticed repetition in the persona descriptions. Surprisingly, repetitions led to both correct and wrong ratings. The participants associated the repetition with AI generation, especially if the persona’s name was repeated:

*“The descriptions seem to have very repetitive words which makes me suspect it’s AI-generated.”* (P41, AI-generated persona correctly rated)

However, participants also rated human-crafted personas as AI-generated because of repetitive writing:

*“The arrangement of the write up and the continual use of the person’s name influenced me to decide it is an AI-generated persona.”* (P42, human-crafted persona falsely rated)

**5.8.2 Information.** Another common theme in participants’ responses was the nature of the information provided in the persona descriptions. Participants often associated AI-generated personas

with an abundance of details that, while extensive, felt less personal and occasionally irrelevant. They remarked that AI personas tended to include information that seemed extraneous or overly factual, lacking the nuanced, human touch that makes descriptions feel personalized. This tendency to over-provide impersonal or unnecessary details further contributed to the perception that the AI-generated personas were less relatable and more mechanical, reinforcing the idea that they lacked the selectiveness found in human-crafted descriptions:

*“This persona feels AI-generated because there are many unnecessary details, which AI is kind of notorious for doing at least in my experience”* (P21, AI-generated persona, correctly rated)

*“This persona seems human because there aren’t unnecessary details and everything mentioned is realistic and feasible”* (P21, AI-generated persona falsely rated)

Whenever participants observed that a persona description included emotional or personal details, they were more likely to rate the personas as human-crafted. These emotional elements, such as expressions of feelings, experiences, or personal anecdotes, made the personas feel more genuine and relatable. Participants associated these details with the complexity and depth of human communication, which often involves subtle emotional cues. The presence of such personal information helped to distinguish human-crafted personas from AI-generated ones, which were typically viewed as more neutral or detached. This connection to emotional content reinforced the belief that the persona was created with a human touch, leading to higher ratings in authenticity and engagement:

*“Detailed, consistent, deliberate balance between included work and personal details. Some minor grammatical errors.”* (P36, human-crafted persona correctly rated)

When AI-generated personas included more personal or emotional details, participants were often misled into believing they were human-crafted. These personal touches, such as references to experiences, emotions, or unique characteristics, made the personas appear more authentic and relatable, blurring the line between AI and human-generated content. By incorporating these human-like elements, the AI personas effectively mimicked the complexity and depth that participants typically associated with human authorship. As a result, the participants were more likely to overlook the artificial nature of the personas, rating them as if they had been crafted by a human rather than generated by an algorithm. This demonstrates how including personal details can significantly enhance the perceived authenticity of AI personas:

*“Nancy Jackson’s persona appears human-crafted due to the detailed and nuanced depiction of her professional background, motivations, and personal interests. The inclusion of her passion for cultural diversity and traveling reflects a level of personalization and depth that suggests a human creator. AI-generated personas often lack such specific and individualized details, making the nuanced elements of Nancy’s profile indicative of human authorship.”* (P39, AI-generated persona falsely rated)

Some AI-generated personas included free-time activities that appeared directly related to the persona’s occupation, which participants often recognized as a clear indicator of AI generation. These activities, while logically consistent with the persona’s professional background, lacked the diversity and spontaneity typically seen in human-crafted personas. Participants noted that real people tend to have hobbies and interests that may not always align with their work, adding to their individuality. The overly predictable nature of AI-generated personas’ free-time activities made them feel less authentic and more formulaic, leading participants to rate them as AI-generated rather than human-crafted. This highlighted how rigid connections between work and personal life can diminish the perceived realism of AI personas:

*“His hobby is reading biographies of successful business leaders? He doesn’t sound real at all. It’s all generic and lacks detail. Probably AI.”* (P19, AI-generated persona correctly rated)

Participants of our study noticed that some personas have an attitude to technology that appears to be rather critical instead of open and positive. This was reported to be a reason for rating personas as human-crafted.

*“The fact it mentions him being worried about AI has me thinking it’s more likely to be human-crafted because I just couldn’t imagine the AI including that in its description of someone.”* (P40, human-crafted persona correctly rated)

In addition to the presence and details included in a persona description, participants also argued about the consistency of a persona’s information. Most of the participants argued that consistent personas are human-crafted:

*“Anne not accepting the use of new technology like the smartphone or work tools is consistent with people of older age.”* (P23, human-crafted persona correctly rated)

Consistent AI-generated personas were also rated as human-crafted because of their consistency:

*“The features so described are consistent throughout the narrative”* (P23, AI-generated persona falsely rated)

**5.8.3 Stereotypicality.** Participants frequently remarked that many personas appeared stereotypical or generic, a theme that emerged for AI-generated and human-crafted personas. However, their interpretations of this stereotypicality influenced their evaluations differently. Some participants noted that when a persona was overly stereotypical or lacked unique traits, they were more inclined to believe it was AI-generated. They argued that AI tends to rely on familiar patterns, making it more likely to create personas that fit into generic molds without the nuance or individuality of real people. For these participants, the lack of depth or complexity in such personas indicated that they were machine-generated, as human-crafted personas were expected to show more diversity and distinctiveness in their characteristics:

*“I rated as AI-generated because this persona feels quite stereotypical. A 20 something IT guy that builds PC and games? Very common in real life, but also a very*

*common stereotype that AI could easily generate.”* (P16, AI-generated persona correctly rated)

On the other hand, participants argued that less stereotypical personas were more likely to be human-crafted. These personas, characterized by unique or unconventional traits, felt more personalized and authentic, leading participants to believe humans designed them. The distinctiveness and originality of the non-stereotypical personas made them appear more reflective of real human complexity, as they defied the generic patterns often associated with AI-generated content. For these participants, the nuanced and individualized nature of such personas was a strong indicator of human authorship, reinforcing the belief that humans are better at creating diverse and less predictable personas:

*“The persona of Chris feels unique and non-stereotypical with characteristics such as his love of birds which is not exactly usual for a 10 year old boy in America. I feel like an AI-generated persona might stick to more stereotypical interests.”* (P16, human-crafted persona correctly rated)

The stereotypical nature of some human-crafted personas even misleads participants into rating them as AI-generated, especially if the participants perceive the persona as believable or realistic:

*“The story does feel believable as a human but it seems too stereotypical to be human crafted”* (P54, human-crafted persona falsely rated)

**5.8.4 Realism and Appeal.** In analyzing the participants’ written responses, a common theme emerged: the perception that specific personas appeared realistic, believable, and lifelike. Participants frequently used these qualities as key indicators when arguing that a persona was human-crafted. They emphasized that the more a persona felt authentic and mirrored real human behavior, thoughts, and emotions, the more likely they attribute its creation to a human rather than an AI. The lifelike and relatable nature of these personas, with their believable characteristics and interactions, made participants feel that only a human could craft such a nuanced and convincing depiction, further reinforcing their judgments of authenticity:

*“Oh dear! I know he’s not a real person but wow do I feel sorry for Peter, and am hoping he can move forward and improve his life. As I had quite an emotional first reaction to this persona, I’m going to rate it as human-crafted. He feels like he could exist in real life. Not necessarily as a good person, or likable, but real.”* (P16, human-crafted persona correctly rated)

Accordingly, a realistic depiction also misled participants in the case of AI-generated personas. When AI-generated personas exhibited believable and lifelike traits, participants were more likely to mistake them for being human-crafted. The realistic portrayal, including natural dialogue, relatable behaviors, and plausible personal details, blurred the line between AI and human authorship. This sense of authenticity created by the AI misled participants into thinking the persona was the product of human creativity, highlighting how effective AI can be at mimicking the complexity of human expression when it incorporates realistic elements:

*“The details influenced my decision to rate the persona as human-crafted and the information seemed consistent and realistic across every aspect.”* (P27, AI-generated persona falsely rated)

Participants partly associated a stereotypical description with personas that reflect reality:

*“This could have been human-crafted based on a real person. Stereotypical description of someone from this age group that liked to garden and is afraid of technology.”* (P22, human-crafted persona correctly rated)

This association was also evident in participants’ uncertainty when rating non-stereotypical personas, as they often appeared unrealistic. While participants generally associated non-stereotypical personas with human authorship, the lack of familiar or conventional traits sometimes made these personas feel less believable. As a result, participants found it challenging to determine whether such personas were AI-generated or human-crafted confidently. The unusual or unexpected characteristics of these personas created ambiguity, making participants question their initial assumptions about what constitutes a realistic, human-like portrayal:

*“It feels a bit made up. Nobody plays Pokemon Go anymore, not since like 2016. On the other hand, it’s believable and there is some specific detail. I’m not sure about this one.”* (P19, human-crafted persona neutrally rated)

Furthermore, participants argued that “likeable” personas are human-crafted:

*“This persona is human crafted. he is interesting and likable”* (P38, human-crafted persona correctly rated)

**5.8.5 Positive and Negative Descriptions.** A common reason participants rated personas as AI-generated was the presence of an overly positive description. Many participants remarked that some AI-generated personas seemed “too positive,” presenting an unrealistic or excessively idealized depiction. This overly optimistic tone, where the persona exhibited no apparent flaws or challenges, was perceived as lacking the nuance and complexity of real human behavior. Participants viewed this as a hallmark of AI generation, as they expected human-crafted personas to reflect a more balanced and authentic representation of strengths and weaknesses. The excessive positivity made the personas feel artificial, leading to lower ratings in terms of realism and authenticity:

*“[...] There also doesn’t appear to be any inherently negative things said, even just a small thing- like with the previous example of the gardener with a back problem. This profile just seems too perfect.”* (P35, AI-generated persona correctly rated)

Furthermore, participants were more likely to rate personas as human-crafted if depicted in a more negative or balanced light. Personas that included imperfections, flaws, or challenges were perceived as more authentic and realistic, leading participants to associate them with human authorship. Including negative traits or struggles made the personas feel more relatable and complex, as these characteristics better reflected the nuanced nature of real human experiences. Participants saw this more grounded and less idealized portrayal as something a human would be more likely

to create, unlike the overly positive and polished personas often associated with AI generation:

*“seems like a very basic description of a normal person, I feel like ai would not be so harsh”* (P33, human-crafted persona correctly rated)

## 6 Discussion

We conducted a study to investigate differences between human-crafted and AI-generated personas. In the previous sections, we reported the results of our persona crafting and generation process and our online survey’s quantitative and qualitative results. In the following section, we discuss the results of our overall research questions.

### 6.1 Distinguishing between Human-Crafted and AI-Generated Personas

Our first research question aimed to investigate to what extent participants can distinguish between human-crafted and AI-generated personas. As shown in our quantitative results, we found a significant difference between the rating of human-crafted and AI-generated personas. Human-crafted personas achieved higher Likert scale ratings when participants were asked if the shown persona was human-crafted. Accordingly, in the case of AI-generated personas, participants rated the statement that a persona is AI-generated significantly higher on a Likert scale. Consequently, our results do not replicate the work by Schuller et al. [45], who reported that AI-generated personas were indistinguishable from human-crafted personas, suggesting similar quality and acceptance. Schuller et al. employed a non-parametric Wilcoxon rank-sum test on results from 11 participants, which did not reject the null hypothesis. While in their case, the null hypothesis assumes that participants could not distinguish between AI-generated and human-crafted personas, not rejecting the null hypothesis does not prove it. A significant effect can be absent due to the low sample size. By asking a group of 54 participants and focusing on novice users, we tried to investigate the topic further, more specifically. These factors may have led to a significant difference in the measured results. **Thus, we answer RQ1:**

**RQ1:** Our results indicate that participants could distinguish between the human-crafted and AI-generated personas. When asked whether a persona was human-crafted, participants rated human-crafted personas significantly higher on a 7-point Likert scale ( $p = .003$ ). Similarly, when asked whether a persona was AI-generated, participants rated AI-generated personas significantly higher ( $p = .002$ ). These findings suggest that participants can identify differences between personas based on their origin, demonstrating a perceptual distinction between human-crafted and AI-generated personas.

### 6.2 Perception of Human-Crafted and AI-Generated Personas

Our second research question aimed to investigate which features in a persona’s description affect how this persona is perceived and why participants rate a persona as human-crafted or AI-generated. Therefore, we collected ratings for different aspects of personas based on related literature. Our quantitative results replicated the work by Salminen et al. [39]. We found significant differences between human-crafted and AI-generated personas regarding their informativeness, positivity, and consistency. Participants rated these aspects significantly higher on a 7-point Likert scale for AI-generated persona. Additionally, we found a significant difference in the stereotypicality of persona description. Participants rated AI-generated personas as significantly more stereotypical when answering a 7-point Likert scale question. Related work stated that AI-generated personas were perceived as non-stereotypical [39]. Furthermore, we found a significant difference regarding the clarity of personas. Participants rated AI-generated personas significantly higher than human-crafted personas. We could not find any significant differences in realism, relatability, or likability between human-crafted and AI-generated personas. **We answer RQ2.1:**

**RQ2.1:** Participants rated AI-generated personas as significantly clearer, more consistent, more positive, and more informative than the human-crafted personas. Despite these strengths, AI-generated personas were also perceived as significantly more stereotypical. This suggests that AI-generated personas may excel in certain areas, such as clarity and consistency, but they also have a higher risk of reinforcing stereotypes.

### 6.3 Features That Distinguish Human-Crafted and AI-Generated Personas

Our qualitative results showed that the main difference between human-crafted and AI-generated personas was the writing and language style. Participants who reasoned about either the style of writing, grammatical errors, the tone, or the choice of words were primarily correct about their rating. AI-generated personas appear to use a “more robotic” language and write texts that read differently than human-crafted ones. Additionally, human-crafted personas include grammatical errors that are primarily interpreted correctly as a reason to rate a persona as human-crafted. However, some participants still expect AI-generated personas to be less correct sometimes. Furthermore, only a portion of the participants reported grammatical errors as one of the reasons for some of their decisions. At the same time, the choice of words and tone of the language were the main aspects of the writing style.

Participants referred frequently to the personal details and emotional content of human-crafted personas. AI-generated personas were often rated correctly with the reason for including “unnecessary” details that do not add information to the persona’s personality. In our study, participants expected AI to be less emotional and more general than personal. Thus, AI-generated personas that included personal information were sometimes rated as human-crafted. AI-generated personas sometimes describe the fictional



person's free-time activities in a way perceived as unusual by participants. These fictional persons were described as enjoying doing only things in their free time associated with their work.

Consistency of the given information was another frequent theme in participants' reasoning. Several participants expected human-crafted personas to be more consistent. Since the quantitative results for consistency showed a significantly higher consistency in AI-generated personas, AI can mislead participants due to its quality regarding specific aspects.

The quantitative results showed higher stereotypically in the AI-generated personas, which was also a frequently mentioned reason for rating a persona as AI-generated. Participants described many AI-generated personas as generic and a stereotype of a person. However, some AI-personas that were described as stereotypical were also perceived as realistic. Some participants seem to recognize stereotypical personas as realistic since the stereotype exists.

In general, participants argued that personas were human-crafted using a depiction of a person who seems to exist. As shown in the quantitative results, there was no significant difference between AI-generated and human-crafted personas regarding realism. In addition to the argument of personas' likeability and relatability, participants rated personas wrongly as human-crafted if they perceived them to fulfill precisely those aspects. Likable, relatable, and realistic AI-generated personas were often rated as human-crafted.

The AI-generated personas in our study were often depicted in a more positive light. Participants argued that this purely positive depiction appears unrealistic and rated personas as AI-generated whenever they noticed a description that is only from a positive point of view. The human-crafted personas that talked about the weaknesses of their fictional characters were rated as human-crafted. Thus, a humane persona seems positive and open about strengths and weaknesses.

Our participants pointed out that AI-generated personas often describe the attitude toward technology as less critical and more open. Whenever participants noticed that a persona was critically approaching technology, they rated them human-crafted. This implies the expectation of a bias against technology in AI-generated personas that participants expect. Our participants argued that AI would not describe technology as something negative or something to be careful about. **We answer RQ2.2:**

**RQ2.2:** Participants generally described the human-crafted personas as more realistic, consistent, critical, emotional, and personal. In contrast, AI-generated personas were expected to appear more robotic, include unnecessary details, and rely on overly generic and positive stereotypes. The writing and language style of a persona description, in particular, played a key role in shaping participants' perceptions, with human-crafted personas being seen as more authentic and relatable.

## 6.4 Generalizability of Findings

This study used ten personas created by HCI experts who are familiar with the concept but do not regularly design personas. We then evaluated these personas with participants who are not trained in persona design. Although the study's conditions limit how broadly

we can generalize the findings, especially when compared to personas crafted by experts after real user interviews, we believe the results highlight important concerns for more common, less ideal scenarios.

With the growing use of large language models to generate data that typically requires significant human effort, there's a risk that laypersons might accept AI-generated personas at face value. Because these personas often appear well-written, users may assume they are good enough and skip the deeper design process altogether. Prior studies—and our findings—show that AI-generated personas can seem polished. However, we also show that they lack depth, diversity, and critical realism. If the superficial qualities lead non-experts to rely on them instead of engaging in thoughtful design, the effectiveness of persona-based software development could be seriously undermined.

## 6.5 Limitations and Future Work

This study employed a set of human-crafted personas developed by ten experts in the field of Human-Computer Interaction (HCI). These personas may reflect a realistic example, but are limited in number. Since this study was conducted with online participants, we decided to limit both persona sets to ten so as not to overwhelm the participants. Future work might investigate different aspects of personas more deeply, further to explore the differences between human-crafted and AI-generated personas. Furthermore, we focused on very general personas that were neither designed by routinized design experts nor for a particular use case. While our study aimed at the potential threat to the quality of fast persona generation, it still seems like LLMs could enhance the professional workflows of persona designers. Thus, another study with a main focus on experts can bring further insights.

Participants mentioned the grammatical errors of some human-crafted personas. This could imply that human-crafted personas are easily identified whenever mistakes due to human error are present. However, the recognized grammar mistakes were only present in four personas and interpreted in three cases as human-crafted features. Therefore, the results do not suggest grammatical errors as the main reason why people were able to distinguish the personas. We eliminated all typographical errors before the human-crafted personas were included in the survey to avoid the effect of obvious mistakes on the judgment of our participants. To assess the specific effect of grammatical errors in text on the perception of human users and their awareness of the usage of AI, future research might come up with studies specifically designed for that purpose.

The constructs used to rate the personas in this study are based on related literature. However, we had to limit the number of constructs to those applicable to only text-based personas rated by participants without pre-screening for development experience. The decision to sample participants without a background in computer science, human-centered design, or development was essential to investigating the general perception of human beings as personas. However, experiences in one of the mentioned disciplines might influence participants' perceptions.

A limitation of this study is the lack of direct evaluation of personas in active design scenarios with professional designers. While our approach provides valuable insights into how personas



are perceived regarding realism, consistency, and informativeness, it does not measure their impact on design outcomes or creativity. Future research should involve designers applying these personas in user-centered design tasks to assess their practical utility and effectiveness. Additionally, our findings may be influenced by the participants' varying levels of familiarity with persona use, which could affect subjective judgments.

Finally, one limitation of our persona generation process was using relatively simple prompts without providing detailed examples or context to guide the language style. While this choice reflects a baseline scenario where users leverage LLMs with minimal customization, it may have contributed to the more robotic tone noted by participants. Future studies should investigate the impact of more elaborate prompt engineering, including few-shot prompting and tailored contextual instructions, to evaluate whether these techniques can produce more human-like and engaging personas.

## 7 Conclusion

This study investigated whether AI-generated personas can be distinguished from human-crafted personas. We collected ten human-crafted personas from ten HCI experts who had no deep experience in persona creation. We asked participants to explain how they perceive these human-crafted personas compared to ten AI-generated personas. We showed that participants could distinguish between the human-crafted and AI-generated personas. In a survey study, we showed that AI-generated personas, compared to personas from novices, are more informative, consistent, and clear, even though their writing and language style were described as somewhat robotic and unnatural. The AI-generated personas were significantly more positive and stereotypical, which was perceived as unrealistic and generic. Overall, we conclude that LLMs can generate personas that meet many quality aspects but also include stereotypical depictions of characters. This can lead to a bias that motivates the usage of personas that might be interesting from a surface-level general quality-based perspective but will retain stereotypes that threaten diverse requirements for engineering and development and, thus, quality.

## Acknowledgments

This work is supported by the German Research Foundation (DFG), CRC 1404: "FONDA: Foundations of Workflows for Large-Scale Scientific Data Analysis" (Project-ID 414984028).

## References

- [1] Tamara Adlin and John Pruitt. 2010. *The Essential Persona Lifecycle: Your Guide to Building and Using Personas*. Morgan Kaufmann. Google-Books-ID: fvmN0Fr5c\_MC.
- [2] William Agnew, A. Stevie Bergman, Jennifer Chien, Mark Diaz, Selim El-Sayed, Jaylen Pittman, Shakir Mohamed, and Kevin R. McKee. 2024. The Illusion of Artificial Inclusion. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3613904.3642703
- [3] Gati V Aher, Rosa I. Arriaga, and Adam Tauman Kalai. 2023. Using Large Language Models to Simulate Multiple Humans and Replicate Human Subject Studies. In *Proceedings of the 40th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 202)*, Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (Eds.). PMLR, 337–371. <https://proceedings.mlr.press/v202/aher23a.html>
- [4] Khaled Albusays, Pernille Bjorn, Laura Dabbish, Denae Ford, Emerson Murphy-Hill, Alexander Serebrenik, and Margaret-Anne Storey. 2021. The Diversity Crisis in Software Development. *IEEE Software* 38, 2 (2021), 19–25. doi:10.1109/MS.2020.3045817
- [5] Lisa P. Argyle, Ethan C. Busby, Nancy Fulda, Joshua R. Gubler, Christopher Rytting, and David Wingate. 2023. Out of One, Many: Using Language Models to Simulate Human Samples. *Political Analysis* 31, 3 (2023), 337–351. doi:10.1017/pan.2023.2
- [6] Ann Blandford, Dominic Furniss, and Stephann Makri. 2016. *Qualitative HCI research: Going behind the scenes*. Morgan & Claypool Publishers.
- [7] Margaret Burnett, Simone Stumpf, Jamie Macbeth, Stephann Makri, Laura Beckwith, Irwin Kwan, Anicia Peters, and William Jernigan. 2016. GenderMag: A Method for Evaluating Software's Gender Inclusiveness. *Interacting with Computers* 28, 6 (Nov. 2016), 760–787. doi:10.1093/iwc/iwv046
- [8] Courtnei Byun, Piper Vasicek, and Kevin Seppi. 2023. Dispensing with Humans in Human-Computer Interaction Research. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3544549.3582749
- [9] Yen-ning Chang, Youn-kyung Lim, and Erik Stolterman. 2008. Personas: from theory to practices. In *Proceedings of the 5th Nordic conference on Human-computer interaction: building bridges (NordicCHI '08)*. Association for Computing Machinery, New York, NY, USA, 439–442. doi:10.1145/1463160.1463214
- [10] Cheng-Han Chiang and Hung-yi Lee. 2023. Can Large Language Models Be an Alternative to Human Evaluations? *arXiv preprint: 2305.01937*.
- [11] Alan Cooper. 1999. *The Inmates are Running the Asylum*. Vieweg+Teubner Verlag, Wiesbaden, 17–17. doi:10.1007/978-3-322-99786-9\_1
- [12] Danica Dillion, Niket Tandon, Yuling Gu, and Kurt Gray. 2023. Can AI language models replace human participants? *Trends in Cognitive Sciences* (2023). Publisher: Elsevier.
- [13] Marco Gerosa, Bianca Trinkenreich, Igor Steinmacher, and Anita Sarma. 2024. Can AI serve as a substitute for human subjects in software engineering research? *Automated Software Engineering* 31, 1 (2024), 13. doi:10.1007/s10515-023-00409-6 Publisher: Springer.
- [14] Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. ChatGPT outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences* 120, 30 (2023), e2305016120. doi:10.1073/pnas.2305016120 \_eprint: <https://www.pnas.org/doi/pdf/10.1073/pnas.2305016120>
- [15] Tom Heyman and Geert Heyman. 2023. The impact of ChatGPT on human data collection: A case study involving typicality norming data. *Behavior Research Methods* (2023), 1–8. doi:10.3758/s13428-023-02235-w Publisher: Springer.
- [16] John J Horton. 2023. *Large language models as simulated economic agents: What can we learn from homo silicus?* Technical Report. National Bureau of Economic Research. doi:10.3386/w31122
- [17] Perttu Hämmäläinen, Mikke Tavast, and Anton Kunnari. 2023. Evaluating Large Language Models in Generating Synthetic HCI Research Data: a Case Study. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3544548.3580688
- [18] Bernard Jansen, Soon-Gyo Jung, Lene Nielsen, Kathleen W. Guan, and Joni Salminen. 2022. How to Create Personas: Three Persona Creation Methodologies with Implications for Practical Employment. *Pacific Asia Journal of the Association for Information Systems* 14, 3 (March 2022). doi:10.17705/1pais.14301
- [19] Bernard J. Jansen, Joni Salminen, Soon-gyo Jung, and Kathleen Guan. 2021. *Creating Data-Driven Personas*. Springer International Publishing, Cham, 93–118. doi:10.1007/978-3-031-02231-9\_4
- [20] Soon-Gyo Jung, Joni Salminen, Kholoud Khalil Aldous, and Bernard J. Jansen. 2025. PersonaCraft: Leveraging language models for data-driven persona development. *International Journal of Human-Computer Studies* 197 (2025), 103445. doi:10.1016/j.ijhcs.2025.103445
- [21] Thomas Kosch and Sebastian Feger. 2024. Risk or Chance? Large Language Models and Reproducibility in Human-Computer Interaction Research. doi:10.48550/arXiv.2404.15782 arXiv:2404.15782 [cs].
- [22] Veronika Krauß, Mark McGill, Thomas Kosch, Yolanda Thiel, Dominik Schön, and Jan Gugenheimer. 2024. "Create a Fear of Missing Out" – ChatGPT Implements Unsolicited Deceptive Designs in Generated Websites Without Warning. arXiv:2411.03108 [cs.HC] <https://arxiv.org/abs/2411.03108>
- [23] Matthias Carl Laupichler, Alexandra Aster, Nicolas Haverkamp, and Tobias Raupach. 2023. Development of the "Scale for the assessment of non-experts' AI literacy" – An exploratory factor analysis. *Computers in Human Behavior Reports* 12 (Dec. 2023), 100338. doi:10.1016/j.chbr.2023.100338
- [24] Tara Matthews, Tejinder Judge, and Steve Whittaker. 2012. How do designers and user experience professionals actually perceive and use personas?. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) (CHI '12). Association for Computing Machinery, New York, NY, USA, 1219–1228. doi:10.1145/2207676.2208573
- [25] Tomasz Miaskiewicz and Kenneth A. Kozar. 2011. Personas and user-centered design: How can personas benefit product design processes? *Design Studies* 32, 5 (Sept. 2011), 417–430. doi:10.1016/j.destud.2011.03.003
- [26] Timothy Neate, Aikaterini Bourazeri, Abi Roper, Simone Stumpf, and Stephanie Wilson. 2019. Co-Created Personas: Engaging and Empowering Users with

- Diverse Needs Within the Design Process. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3290605.3300880
- [27] Raymond S. Nickerson. 1998. Confirmation Bias: A Ubiquitous Phenomenon in Many Guises. *Review of General Psychology* 2, 2 (1998), 175–220. doi:10.1037/1089-2680.2.2.175 arXiv:https://doi.org/10.1037/1089-2680.2.2.175
- [28] Lene Nielsen. 2013. *Personas-user focused design*. Vol. 15. Springer.
- [29] Lene Nielsen. 2019. *Persona Writing*. Springer London, London, 55–81. doi:10.1007/978-1-4471-7427-1\_4
- [30] Lene Nielsen, Kira Storgaard Hansen, Jan Stage, and Jane Billestrup. 2015. A template for design personas: analysis of 47 persona descriptions from danish industries and organizations. *International Journal of Sociotechnology and Knowledge Development (IJSKD)* 7, 1 (2015), 45–61.
- [31] Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2022. Social Simulacra: Creating Populated Prototypes for Social Computing Systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3526113.3545616 event-place: Bend, OR, USA.
- [32] John Pruitt and Jonathan Grudin. 2003. Personas: practice and theory. In *Proceedings of the 2003 Conference on Designing for User Experiences* (San Francisco, California) (DUX '03). Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/997078.997089
- [33] Kari Rönkkö, Mats Hellman, Britta Kilander, and Yvonne Dittrich. 2004. Personas is not applicable: local remedies interpreted in a wider context. In *Proceedings of the Eighth Conference on Participatory Design: Artful Integration: Interweaving Media, Materials and Practices - Volume 1* (Toronto, Ontario, Canada) (PDC '04). Association for Computing Machinery, New York, NY, USA, 112–120. doi:10.1145/1011870.1011884
- [34] Joni Salminen, Kathleen Guan, Soon-Gyo Jung, Shammur A. Chowdhury, and Bernard J. Jansen. 2020. A Literature Review of Quantitative Persona Creation. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3313831.3376502
- [35] Joni Salminen, Bernard J. Jansen, Jisun An, Haewoon Kwak, and Soon-gyo Jung. 2018. Are personas done? Evaluating their usefulness in the age of digital analytics. *Persona Studies* 4, 2 (2018), 47–65. doi:10.21153/psj2018vol4no2art737
- [36] Joni Salminen, Bernard J. Jansen, Jisun An, Haewoon Kwak, and Soon-Gyo Jung. 2019. *Automatic Persona Generation for Online Content Creators: Conceptual Rationale and a Research Agenda*. Springer London, London, 135–160. doi:10.1007/978-1-4471-7427-1\_8
- [37] Joni Salminen, Soon-gyo Jung, and Bernard J. Jansen. 2019. The Future of Data-driven Personas: A Marriage of Online Analytics Numbers and Human Attributes.. In *ICEIS (I)*, 608–615.
- [38] Joni Salminen, Soon-Gyo Jung, Lene Nielsen, and Bernard Jansen. 2022. Creating More Personas Improves Representation of Demographically Diverse Populations: Implications Towards Interactive Persona Systems. In *Nordic Human-Computer Interaction Conference (NordCHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–11. doi:10.1145/3546155.3546654
- [39] Joni Salminen, Chang Liu, Wenjing Pian, Jianxing Chi, Essi Häyhänen, and Bernard J. Jansen. 2024. 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, New York, NY, USA, 1–20. doi:10.1145/3613904.3642036
- [40] Joni Salminen, Lene Nielsen, Soon-Gyo Jung, Jisun An, Haewoon Kwak, and Bernard J. Jansen. 2018. “Is More Better?”: Impact of Multiple Photos on Perception of Persona Profiles. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3173574.3173891
- [41] Joni Salminen, Joao M. Santos, Haewoon Kwak, Jisun An, Soon gyo Jung, and Bernard J. Jansen. 2020. Persona Perception Scale: Development and Exploratory Validation of an Instrument for Evaluating Individuals’ Perceptions of Personas. *International Journal of Human-Computer Studies* 141 (2020), 102437. doi:10.1016/j.ijhcs.2020.102437
- [42] Joni Salminen, Kathleen Wenyun Guan, Soon-Gyo Jung, and Bernard Jansen. 2022. Use Cases for Design Personas: A Systematic Review and New Frontiers. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (CHI '22). Association for Computing Machinery, New York, NY, USA, 1–21. doi:10.1145/3491102.3517589
- [43] Albrecht Schmidt, Passant Elagroudy, Fiona Draxler, Frauke Kreuter, and Robin Welsch. 2024. Simulating the Human in HCD with ChatGPT: Redesigning Interaction Design with AI. *interactions* 31, 1 (Jan. 2024), 24–31. doi:10.1145/3637436
- [44] Sander Schulhoff, Jeremy Pinto, Anau Khan, Louis-François Bouchard, Chenglei Si, Svetlana Anati, Valen Tagliabue, Anson Kost, Christopher Carnahan, and Jordan Boyd-Graber. 2023. Ignore This Title and HackAPrompt: Exposing Systemic Vulnerabilities of LLMs Through a Global Prompt Hacking Competition. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 4945–4977. doi:10.18653/v1/2023.emnlp-main.302
- [45] Andreas Schuller, Doris Janssen, Julian Blumenröther, Theresa Maria Probst, Michael Schmidt, and Chandan Kumar. 2024. Generating personas using LLMs and assessing their viability. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–7. doi:10.1145/3613905.3650860
- [46] Joongi Shin, Michael A. Hedderich, Bartłomiej Jakub Rey, Andrés Lucero, and Antti Oulasvirta. 2024. Understanding Human-AI Workflows for Generating Personas. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference* (Copenhagen, Denmark) (DIS '24). Association for Computing Machinery, New York, NY, USA, 757–781. doi:10.1145/3643834.3660729
- [47] Wilbert Tabone and Joost De Winter. 2023. Using ChatGPT for human–computer interaction research: a primer. *Royal Society Open Science* 10, 9 (Sept. 2023), 231053. doi:10.1098/rsos.231053
- [48] Volker Thoma and Bryn Williams. 2009. Developing and Validating Personas in e-Commerce: A Heuristic Approach. In *Human-Computer Interaction – INTERACT 2009*, Tom Gross, Jan Gulliksen, Paula Kotzé, Lars Oestreicher, Philippe Palanque, Raquel Oliveira Prates, and Marco Winckler (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 524–527.
- [49] Shuohang Wang, Yang Liu, Yichong Xu, Chenguang Zhu, and Michael Zeng. 2021. Want To Reduce Labeling Cost? GPT-3 Can Help. \_eprint: 2108.13487.