



Deus Ex Machina and Personas from Large Language Models: Investigating the Composition of AI-Generated Persona Descriptions

Joni Salminen
University of Vaasa, Vaasa, Finland
jonisalm@uwasa.fi

Chang Liu
Peking University, Beijing, China
imliuc@pku.edu.cn

Wenjing Pian
Fuzhou University, Fuzhou, China
wpian1@e.ntu.edu.sg

Jianxing Chi
Wuhan University, Wuhan, China;
and Fujian Normal University,
Fuzhou, China
chijx@fjnu.edu.cn

Essi Häyhänen
University of Vaasa, Vaasa, Finland
essi.hayhanen@uwasa.fi

Bernard J. Jansen
Qatar Computing Research Institute,
Hamad Bin Khalifa University
bjansen@hbku.edu.qa

ABSTRACT

Large language models (LLMs) can generate personas based on prompts that describe the target user group. To understand what kind of personas LLMs generate, we investigate the diversity and bias in 450 LLM-generated personas with the help of internal evaluators (n=4) and subject-matter experts (SMEs) (n=5). The research findings reveal biases in LLM-generated personas, particularly in age, occupation, and pain points, as well as a strong bias towards personas from the United States. Human evaluations demonstrate that LLM persona descriptions were informative, believable, positive, relatable, and not stereotyped. The SMEs rated the personas slightly more stereotypical, less positive, and less relatable than the internal evaluators. The findings suggest that LLMs can generate consistent personas perceived as believable, relatable, and informative while containing relatively low amounts of stereotyping.

CCS CONCEPTS

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

KEYWORDS

AI, LLMs, HCI, user personas, evaluation

ACM Reference Format:

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)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/3613904.3642036>

1 INTRODUCTION

Personas are fictional representations of target users that provide valuable insights into user needs, behaviors, and preferences [11]



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

CHI '24, May 11–16, 2024, Honolulu, HI, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0330-0/24/05
<https://doi.org/10.1145/3613904.3642036>

presented in a narrative format known as a persona description [25]. Persona descriptions mainly consist of textual information about the user type the persona represents [26]. In turn, large language models (LLMs), such as OpenAI's GPT-4, have shown remarkable capabilities in generating coherent and contextually relevant text. Rooted in their ability to understand and produce textual content [20], LLMs hold promise for further automating persona generation processes, while being able to maintain the narrative realism [10] of manually crafted personas. In theory, such personas can be generated based on data (i.e., be 'data-driven'), while at the same time offering engaging narratives into the circumstances of different people groups' circumstances. To this end, LLMs can generate personas based on prompts that describe the target user group, their needs, goals, and preferences (see Figure 1 for an example). For example, the prompt "Create a persona for a smartwatch game user who likes casual and social games" could yield a persona with a name, age, occupation, hobbies, motivations, and challenges related to smartwatch gaming [40].

However, *what are such personas like? Are they diverse in their representation of people? Do they contain skewness or bias toward certain characteristics? Are the LLM-generated personas at all of satisfactory quality? How do stakeholders feel about them?* These are some of the motivating questions behind our study.

Paoli [28] summarizes the current state of LLM-generated personas as follows: "If the LLM (with the support of the human researcher) can produce at least satisfactorily some forms (or at least ideas) of user personas based on a data analysis, we may also be able to make a step [...] toward Phase 6 [writing the persona descriptions]." (p. 5). Hence, we are still at an early stage where we do not know *how* and *where* to use LLMs in the persona creation process. These questions form essential research gaps that HCI research on personas needs to be addressed.

More precisely, we address the following research questions (RQs):

RQ1: *How diverse are the characteristics of personas created by LLMs? Are there any notable biases?* **RQ2:** *How do (a) UX researchers and (b) subject-matter experts assess the LLM-generated personas?*

Both RQs matter. For **RQ1**, if the LLM-generated personas do not contain diverse characteristics, there is a risk of persona users 'missing out' on marginalized or fringe user types [15, 16], as these

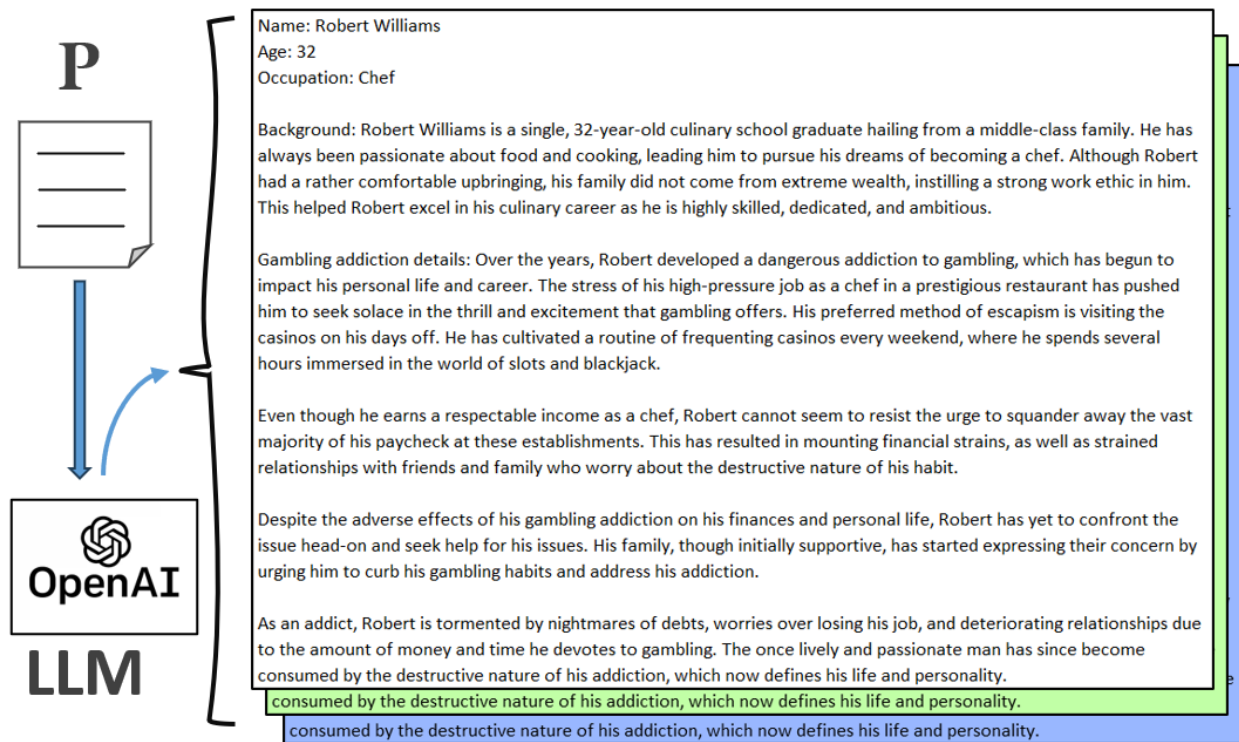


Figure 1: The process of obtaining LLM-generated persona descriptions includes drafting a prompt (P) that the LLM follows to generate persona descriptions. The focus of these persona descriptions is the narrative text content [24, 26].

would not be represented in the generated personas. Moreover, even though the personas would be diverse in their representation of different user types, the personas can still be biased in that they over-emphasize certain characteristics at the expense of others [7, 8]. For example, the persona distribution could be overwhelmingly male or predominantly young.

For RQ2, if end-users of personas consider them to be of weak quality, they will not use such personas [23, 29] and the whole purpose of creating the personas would be defeated. Overall, these questions matter for the persona design practice and the application of personas in real projects.

In this research, we present an in-depth exploration of LLMs' potential for persona generation. Specifically, we analyze the benefits, challenges, and limitations of employing LLMs in persona creation and provide valuable insights into the effectiveness and reliability of the generated personas for the HCI community. Our findings contribute to the understanding of LLMs' role in persona creation. By using LLMs, there is an opportunity to revolutionize the way personas are created, providing designers and researchers with valuable insights into user behaviors and preferences. This transformative nature of LLMs is not to be underestimated by the HCI community, nor should the risk involved with this technology cloud our judgment of its positive prospects. For example, Schmidt [40] reports that "We tried prompts to ask for [HCI-related] ideas, similar to user input in brainstorming and focus groups. The results are in many ways what we would expect from working with actual

users." (p. 8). This statement attests to the fact that the LLMs generate plausible outputs, drastically different from what existed just a couple of years ago. By venturing into this novel domain, we aim to establish a foundation for future investigations into LLM-generated personas.

2 REVIEW OF LITERATURE

LLMs have been explored in the design process in conjunction with personas in various ways, but there is currently no study similar to ours. Here, we summarize the main previous work.

Deshpande et al. [12] study the anthropomorphizing of LLMs. For instance, telling an LLM to "Talk like a doctor" allows it to assume the role of a doctor. The researchers discuss an example of how different personas assigned to the same AI system led to varied behaviors. The exact 'persona' of the system has a large effect on its behaviors and decisions and altering the prompt can make the LLM switch roles on the fly. However, the study did not investigate the characteristics (specifically diversity and/or bias) of LLM-generated personas or user perceptions of such personas.

Paoli [28] illustrates that LLMs can create user personas based on thematic analysis (TA) of semi-structured interviews with real users. The LLM can generate codes and themes from the interview data, and then use them to write personas narratives that include the goals, background, needs, challenges, and other relevant details.

However, the study did not investigate the characteristics (specifically diversity and/or bias) of LLM-generated personas or user perceptions of such personas. In a similar vein, Zhang et al. [45] further elaborate that LLMs can be used for cleaning, integrating, predicting, and analyzing user feedback data, which is a key step in generating high-quality personas. They introduced a GPT-4 based tool *PersonaGen* to generate personas, which can classify different persona attributes for downstream tasks. However, the study did not investigate the characteristics (specifically diversity and/or bias) of LLM-generated personas or user perceptions of such personas.

Kocaballi [6] reported the capability of ChatGPT to generate fictional user personas for a given project, for which ChatGPT was asked to generate five different user personas. Five brief descriptions were successfully generated accordingly with a good variety of demographics. The researcher [6] further commented that “[the GPT-generated] five different personas [showed] a good range of variety in demographics, but potentially [lacked] ethnic diversity based on their names”, which may suggest a possibility for ethnic mismatches in LLM-generated personas. While Kocaballi’s findings were based on a limited sample of personas generated, our evaluation further investigates this topic using a larger sample of personas and more thorough analyses.

Alessa and Al-Khalifa [2] created elderly personas from different demographic backgrounds that interacted with a conversational agent based on the persona’s details. The context was the mitigation of experienced loneliness by the elderly. The interaction episodes were rated by subject-matter experts based on criteria such as engagement, interestingness, fluency, and sense-making. However, the elderly personas themselves were not rated nor were their characteristics examined in the study. Similarly, Hong et al. [21] illustrated the potential of LLMs for assuming the role of the persona by responding to users’ natural language queries. The researchers also pointed out the risk of not precisely knowing “whose opinions are reflected in the generated [personas]”, which may lead to representational biases such as over- or under-representation of demographic subgroups [21]. However, the study did not investigate the characteristics (specifically diversity and/or bias) of LLM-generated personas or user perceptions of such personas.

Cheng et al. [9] presented a framework named “Marked Personas” that applies natural language prompts to generating personas, i.e., imagined individuals belonging to specific demographic groups. They evaluated the personas by a method named “Marked Words”, which included identifying words that statistically distinguish personas of marked groups from corresponding unmarked ones. The researchers found evidence of harmful patterns like stereotypes and essentializing narratives. They also provided recommendations for LLM creators and researchers to address stereotypes and essentializing narratives. This is the closest study we could locate to ours; yet, it focuses on the intersectionalist analysis of bias (particularly race and gender), whilst we explore more variables, including age, gender, country, and so on. We also add the subject-matter expert perspective, which the study by Cheng et al. [9] did not do.

None of the previous studies, as far as we know, have specifically investigated the characteristics of LLM-generated personas in terms of diversity and bias. Also, we could not locate a study that would have tested subject-matter experts’ perceptions of LLM-generated

personas (for example, Alessa and Al-Khalifa [2] evaluated the interaction between SMEs and personas, not the personas themselves). Because these dimensions of personas form a core line of investigation, the lack of research in them poses a research gap that we address in this study. Incorporating demographic attributes into persona generation from an HCI perspective is essential for designing systems that are more accessible and tailored to diverse user needs, thereby fostering inclusive technology development.

3 METHODOLOGY

3.1 Research Context

As the research context, we chose a serious domain: addiction. Our choice was based on the notion that personas could be more broadly applied for “social good” – that is, societally beneficial purposes [32]. So, in our study, the personas represented individuals with various types of addiction, including *alcohol*, *opioids*, *social media*, *online shopping*, and *gambling*. Based on internal ideation among the team, these addictions were chosen to represent a wide range of addictions that take place in modern people’s lives and touch people regardless of age, gender, or nationality.

Personas are not only applicable to representing ‘users’ of products; instead, they generalize to representing *any* groups of people, for example, survey respondent groups [33]. We want to emphasize this broader applicability, also referred to as ‘personas for social good’ [18, 32], by focusing on the context of a real societal issue, *addictions*. Addictions are treatable, chronic medical conditions in which individuals’ interactions among their brain, genetics, environment, and life experiences may lead them to compulsively use substances or act in certain behaviors that are harmful to them in multiple ways [17]. Addictions could be broadly categorized into two types: substance use disorders and behavioral addictions [46]. Examples of substance use disorders include opioid, nicotine, and alcohol use disorders, while behavioral addictions include (but are not limited to) gambling and overeating. In our case, the personas addicted to alcohol or opioids represent individuals with *substance use disorders*, whereas the personas addicted to gambling, online shopping, or social media represent individuals with *behavioral addictions*.

The chosen conditions describe different forms of addictions in the life of a modern person, ranging from more to less severe in terms of their immediate health impacts. Opioid addiction is a major issue in the United States [44]. Alcohol addiction remains one of the most alarming forms of addiction [13]. Gambling is a particular concern among young men, although it touches nearly all age and gender demographics [27]. Social media and online shopping (ranging from impulsive [19] to compulsive shopping [42] behaviors) are perhaps more recent but yet serious forms of addiction that can have negative impacts on people’s lives, e.g., by having adverse financial or social effects.

While the context of addictions enables us to test the LLM’s ability to create personas for social good [32], this context also enables us to examine any potential biases related to age or gender in a meaningful way. Going forward, personas generated for this context could be used in the design of automated app interventions, for example, to mitigate these addictive behaviors (although this is beyond the scope of the current work).

3.2 Persona Generation

We used GPT-4 (June 2023 version) to generate 450 personas. We created three types of prompts for each addiction: one specifying male gender, one specifying female gender, and one not specifying a gender at all. The reason for this is that, first, we would control and balance the number of each gender; second, we plan to test the gender distribution when it is not specified. So, given that we have five addictions and three prompt types, that yields 15 combinations (3×5). However, generating only 15 personas would be susceptible to inherent randomness in the LLM generation process [3]. To thoroughly evaluate the LLM's ability to generate personas consistently, we need to repeat the generation multiple times. Each time, we obtain a different persona. We chose to repeat the generation 30 times for each of the 15 combinations, thus yielding 450 personas in total (30×15).

A general challenge with LLM-generated personas is that inputting the same prompt multiple times via Open AI's API to GPT-4 yielded nearly identical personas, which might be a caching issue. We addressed this issue via a two-stage prompting strategy: first, we asked the model to generate a list of 30 "skeletal" personas for each addiction-prompt type combination (skeletal in the sense they only contain basic information [35, 41]). This resulted in unique short persona descriptions that we then inputted back to the model, asking it to expand each persona description to create the full persona descriptions (i.e., "rounded personas" [25]) for analysis. Our code is publicly available in the following *Google Colab* notebooks (NB):

- **NB1. Skeletal persona generation:** <https://bit.ly/LLM-personas-skeletal>
- **NB2. Rounded persona generation:** <https://bit.ly/LLM-personas-rounded>

Given access to Open AI's API, other researchers (who have access to Open AI) can run the notebooks to generate personas to replicate our findings or create personas from different contexts by making slight modifications to the prompt (e.g., by changing the context from addictions to something else). Note that in our prompt, we did not provide an explicit definition of what a persona is, as we presumed, based on prior literature [28], that the model already knows what the persona is (and this presumption was correct). However, we specified the role of GPT ("You are a helpful assistant to a social sciences researcher") as well as a structured template for the information we expected ("Provide the output in a json array, with each dict containing only the following keys: 'index', 'name', 'age', 'occupation', 'background', 'details'). The expansion was done by taking the input personas in the previous step and asking the model to expand on them ("Expand on the following summary persona. Ensure that all the information provided is used in your expanded persona."). Overall, using the structured template approach is aligned with prior research on LLM-generated personas [28] and it also has the benefit of producing comparable personas (as the information is in standard structure) – this is beneficial also for other researchers, as we share our persona dataset.

The **Personas-addicted** dataset can be downloaded here: <https://bit.ly/LLM-personas-data>. Overall, our method is replicable in terms of the programming code provided and the analysis is also replicable as we provide the persona descriptions themselves. So,

the methodology itself exemplifies that LLM-generated personas can increase persona creation replicability which has been found problematic in past studies [8, 30]. This is important because replicability is one key toward accomplishing *persona science* [35] which is the application of scientific principles in the study of personas and their users. As such, we believe the datasets to be beneficial to others in the HCI community.

3.3 Evaluation Protocol

3.3.1 Internal Evaluation. The procedure of data coding and evaluation is divided into two stages. The first stage is the internal evaluation of the generated personas to gain a "sanity check" on their quality, within which four internal evaluators from our research team were involved. The average experience among the UX researchers was 9.25 years in UX/HCI research and they were all familiar with the concepts used in the evaluation, such as pain points. Each evaluated approximately 120 personas (of which around 112 were evaluated by one evaluator and eight were used for the inter-coder reliability calculation). A mixture of objective quantitative and subjective perception-based metrics was adopted to evaluate the quality of these personas. The second stage is the subject-matter experts' (SMEs) evaluation of these personas performed by five public health professionals with domain expertise on addictions. Within this stage, only a subset of these personas was evaluated by these external evaluators. The schema of coding and evaluation is shown in Table 1.

In the following, we explain why each criterion is relevant for this study.

Age, gender, and occupation. These are basic characteristics in typical persona profiles [26, 31] that enable us to assess whether there are any distinct biases or stereotypes concerning demographic variables. Demographic diversity is considered important for inclusive design through personas [15, 16], especially when it comes to representing all age and gender groups. Alongside the demographic information, occupation is often included in persona profiles [5].

Text length. This is an interesting variable that captures how extensive persona descriptions the LLM generates. The information contributes to providing a baseline for further comparisons with human-generated personas.

Pain points. Pain points, often referred to as needs, goals, and wants, are typical content for personas [11, 26]. Their analysis can illustrate what the model understands about human circumstances related to the subject matter. We recorded the frequency and content of pain points in the coding stage.

Physical appearance. Appearances matter for personas; for example, smiling pictures affect multiple user perceptions of personas [36]. Persona attractiveness is consistent with the 'what is beautiful is good' effect; personas that are perceived as physically more attractive are attributed to other positive traits [38]. So, we evaluated how LLMs would characterize the physical appearance of the persona.

Personality. Personality traits characterize the persona's psychological tendencies [37]. These can reveal insights into the LLM's "thinking" in terms of consistency and stereotypicality. So, we extract the mentioned personality traits.

Table 1: The coding sheet applied to extract information and assign evaluation ratings to the personas. The definition column includes instructions given to the evaluators. The criteria highlighted in blue color (the last six items of the table) were given both to internal and external evaluators; the other criteria before that were coded only by internal evaluations (Krippendorff’s $\alpha = 0.833$).

	Variable	Definition provided to the evaluators
Extracted from persona description as information	age	The persona’s age
	gender	The persona’s gender (“m” for male, “f” for female)
	occupation	The persona’s job title
	text length	The length of the persona description in words
	number of pain points mentioned	A pain point is a problem or issue that the persona has; in this context, pain points related to addiction (list as many as you find)
	pain point list <i>[open-ended]</i>	Writeup of the pain points, separated by comma and space
	physical appearance mentioned	If the persona’s physical appearance is mentioned, mark “y”; if it’s not, mark “n”
	if yes, how? Describe physical appearance <i>[open-ended]</i>	Describe how the physical appearance is described (you can paste the text from the persona description)
	personality mentioned	If the persona’s personality is mentioned, mark “y”; if it’s not, mark “n”
	if yes, how? Describe personality <i>[open-ended]</i>	Describe how the personality is described (you can paste the text from the persona description)
Determined based on human evaluation of the persona	informativeness for design <i>[persona perception]</i>	Does the persona description contain adequate information to design an app or system to address the persona’s needs?*
	believability <i>[persona perception]</i>	Does the persona appear realistic, i.e., lifelike, like an actual person that could exist?*
	stereotypicality <i>[persona perception]</i>	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.
	positivity <i>[persona perception]</i>	Is the person depicted in a positive light?* (an example of <u>not</u> being depicted in a positive light is to blame the persona for the addiction)
	relatability <i>[persona perception]</i>	Is the persona relatable? Relatability is the quality of being easy to understand or feel empathy for.
	consistency <i>[persona perception]</i>	Is the persona consistent?* Consistent persona does not have conflicting information (for example, if the description said “he is a happy personality” but later said, “because he is often sad” => these information pieces conflict so you would give a low rank for consistency.

Persona perceptions. These are users’ perceptions of the persona they are using [39]. There are both positive and negative persona perceptions: positive ones are qualities we would like to see in a persona (in our study, these are (a) *informativeness for design*, (b) *believability*, (c) *positivity*, (d) *relatability*, and (e) *consistency*), whereas negative ones are qualities we would like to avoid (in our framework: *stereotypicality*). In short, a good persona provides useful information for design purposes, is believable (i.e., realistic, credible), presents the persona in a positive light (not as an antagonist), is relatable (i.e., evokes empathy), and is consistent (i.e., does not contain conflicting information) [11, 25, 39].

We computed the inter-coder reliability based on the four informational categories shown in Table 1 (note that open-ended and persona perception categories cannot be used here because they contain subjective information). Since these categories contain a

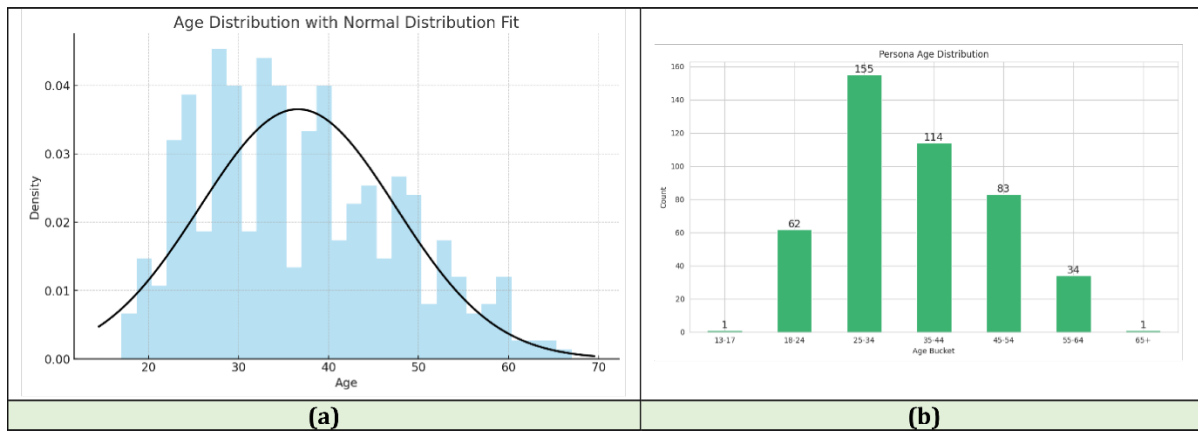
mixture of categorical and numerical data, we selected Krippendorff’s Alpha (α) as the inter-coder reliability metric. The average value taken from these four categories indicates high agreement ($\alpha = 0.833$, where above 0.800 is considered high). Therefore, we conclude that the coded data is quite reliable. For any observed disagreement, the lead author made a judgment call about the final data value.

3.3.2 External Evaluation. In addition to the four internal evaluators, five SMEs evaluated a sample of the personas. We used stratified random sampling to select 30 personas for the SMEs to evaluate. We stratified the sampling by gender and addiction type, so there were three male and three female personas in each of the five addiction types ($3 \times 2 \times 5 = 30$), saving other gender identifications for future research.

Table 2: Human evaluators in this study. The SMEs received a USD 100 compensation. The UX researchers were not financially compensated.

Evaluator ID	Occupation	Years of exp.*	Gender	Country
SME1	Health Surveillance Officer at the State Health Department	3	Female	Nigeria
SME2	Rehabilitation counselor	20+	Female	United States
SME3	Lead Processor at Medical Health Services	3	Female	United States
SME4	Master's degree in Epidemiology, worked on a research project related to substance abuse	4	Female	India
SME5	Pharmacist	15+	Female	Pakistan
UXR1	Associate Professor	16	Female	China
UXR2	Research Assistant	1	Female	Finland
UXR3	Associate Professor	12	Male	China
UXR4	Lecturer	8	Female	China

*For SMEs: In public health. For UX: In UX/HCI research.

**Figure 2: The age distribution of the personas. The median is 35 years.**

We recruited five public health professionals for the study as SMEs using Upwork, a professional services platform (see Table 2 for description). The recruitment included a screening stage where the SMEs were asked three questions: their knowledge of addictions and their work experience in public health. These questions were used to ensure that each SME participant had prior experience of addiction and public health. Prior knowledge of personas was not deemed necessary, as we explained to each SME what a persona is before they started their evaluations. So, the SMEs were briefed on what personas are; they were then provided with definitions of each evaluation criterion (the same ones in Table 1) and asked to evaluate the 30 persona descriptions (for study replication, the IDs of the 30 personas are shared in Appendix 1). We also asked them to provide a short, written statement of their overall impression regarding the evaluated personas in terms of each evaluation criterion. The SMEs were not told that the personas were computer-generated; they were simply told that we were researching personas.

4 RESULTS

4.1 RQ1: How diverse are the characteristics of personas created by LLMs? Are there any notable biases?

4.1.1 Age. The average age of the personas is approximately 37.04 years (SD = 11.11 years). The range is 17-67, meaning the youngest persona is 17 years old and the oldest 67 years old. As can be seen from Figure 2, personas are generated across different age groups, which is a desirable feature relative to a scenario where the personas would only focus on a certain age group. At the same time, the Shapiro-Wilk test indicated that the age of the personas is not normally distributed, $W(474) = .97, p < .001$. Rather, the age distribution displays platykurtic properties, i.e., lower peakedness and flatter tails compared to a normal distribution.

We next conducted Chi-squared tests to compare the frequencies of different addiction types between various age groups. The age grouping was adopted from previous persona generation research [4, 5]. We omit the age groups 13-17 and 65+ from this analysis, as each only had one observation.

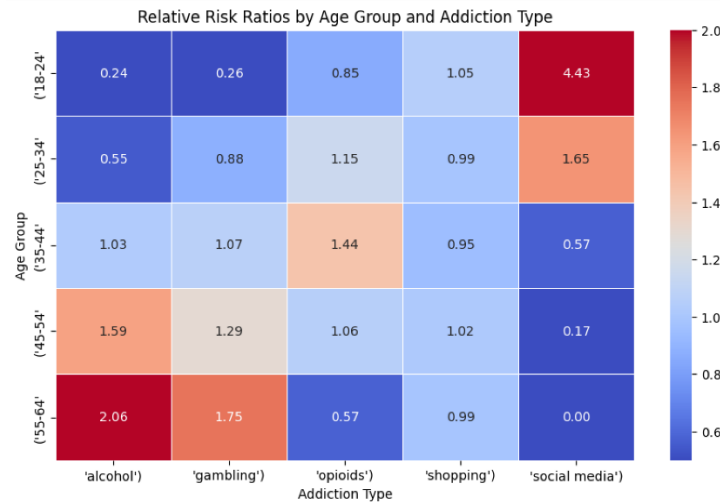


Figure 3: Relative risk ratios by personas’ age group and addiction type. Red indicates a higher risk ratio, and blue indicates a lower. We can observe that the LLM associates the risk of social media addiction with the youngest age cohort (18-24) and the risk of alcohol addiction with the oldest age cohort (55-64).

First, the results indicate a significant difference in the prevalence of *gambling addiction* among the age groups, $\chi^2(4, N = 450) = 10.93$, $p < .05$. The age group with the highest prevalence for gambling was 55-64. Second, there was a significant difference in the prevalence of *alcohol addiction* among the age groups, $\chi^2(4, N = 450) = 25.34$, $p < .05$. The age group with the highest prevalence for alcohol was 55-64. Third, there was a significant difference in the prevalence of *social media addiction* among the age groups, $\chi^2(4, N = 450) = 67.78$, $p < .05$. The age group with the highest prevalence for social media was 18-24.

There was no significant difference in the prevalence of shopping or opioids among the age groups. However, the age groups with the highest prevalence of shopping and opioid addiction were 18-24 and 35-44, respectively. Figure 3 illustrates the relative risk ratios that the LLM-generated personas from different age groups had for the different addiction types.

We also investigated which addiction type is most prevalent for each age group. For age groups 18-24 and 25-34, the most prevalent addiction type was *social media*. For the 45-54 and 55-64 age groups, it was *alcohol*. For the age group 35-44, it was *opioids*. At face value, the variability in these addiction types seems to make sense in terms of the younger age groups being more addicted to social media than the older age groups.

While further comparison to Census statistics at a population level is needed to establish the robustness of these differences, from these findings, we can surmise that the LLM has an opinion of what age group is typically addicted to what – but in the absence of baseline data, we cannot deduct if that opinion is factually correct.

4.1.2 Gender. When not specifying the persona’s gender in the prompt ($n = 150$), the LLM generated a perfectly even distribution of male and female personas, thus resulting in perfect gender parity. We also verified whether the generated personas for which the gender was specified (male or female) actually matched the specified

gender. We found this to be true in all cases (100% adherence to instructions). There was a statistically significant difference in the average age between male ($M = 37.72$, $SD = 11.11$) and female ($M = 35.50$, $SD = 10.67$) personas; $t(448) = 2.16$, $p = .031$. Even though the male personas were slightly older than their female counterparts, the difference is not meaningful in practice (only two years). There was no statistically significant relationship between gender and addiction type, $\chi^2(1) = 0.231$, $p = .994$.

4.1.3 Country. The LLM generated the first and last names for each persona. However, neither the persona description nor the prompt had information about the persona’s country. So, we applied Name2GAN, an online research tool for inferring likely demographics (gender, age, country) based on their name (the model has been trained on millions of names [22]). The results indicate that The LLM-generated personas originated from 15 countries: Argentina ($n = 1$), Australia ($n = 1$); Brazil ($n = 1$); Colombia ($n = 6$); Germany ($n = 1$); Hong Kong, China ($n = 1$); Mexico ($n = 14$); Nigeria ($n = 2$); Philippines ($n = 3$); South Korea ($n = 2$); Spain ($n = 2$); Taiwan, China ($n = 3$); United Kingdom ($n = 25$); United States (US) ($n = 385$); and Vietnam ($n = 3$). Although the large number of countries indicates diversity, the frequencies in Figure 4 indicate strong US-centricity, with the overwhelming majority of the personas being from the US. Nonetheless, a Chi-squared test indicates no statistically significant difference in the prevalence of addiction types between US and non-US personas; $\chi^2(4, N = 450) = 47.08$, $p = .796$

4.1.4 Occupation. In terms of jobs, the generated personas were extremely versatile, with 201 unique jobs being used by GPT-4. Examples are shown in Table 3. Overall, the data suggests a wide variety of occupations across the personas, with the most common being “Graphic Designer”, “Real Estate Agent”, and “Accountant”, each with 12 occurrences. So, in terms of job occupations, GPT generates a wide diversity of personas. In terms of gender, there

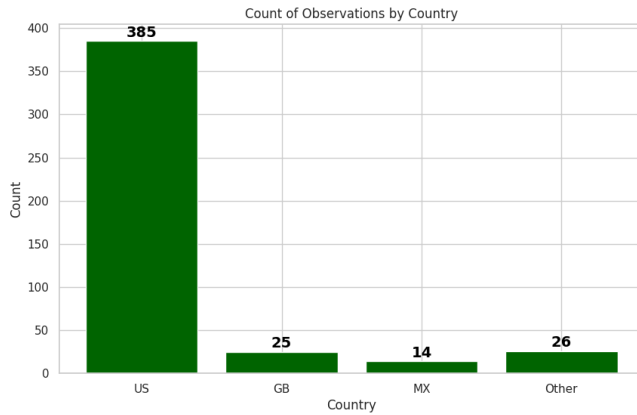


Figure 4: Countries of the generated personas according to the Name2GAN model [22], accessed online at <https://acua.qcri.org/tool/Name2GAN>. Most of the personas (85.6%) were from the United States, despite the country not being specified in the prompt. This implies that GPT tends to generate US-centric personas by default.

is some stereotyping: Figure 5 shows that male personas are more likely to be construction workers, software developers, and unemployed, and female personas are more likely to be nurses, event planners, and baristas. In terms of addiction type, occupations appear randomly distributed (see Figure 6) – the only occupation with a frequency of higher than 3 is “unemployed” ($n = 4$ for alcohol addiction).

4.1.5 Pain points. To carry out a thematic analysis identifying common themes or topics mentioned in the pain points, we used a simple word frequency-based approach. The list in Table 4 contains keywords from the pain points extracted from the LLM-generated personas. Some potential themes that could be inferred from these common words are work-related issues, relationship problems, financial troubles, performance concerns, stress, and life disruptions. At face value, these reasons appear plausible antecedents for the development of addictions, although a more robust assessment is required.

Table 3: Example occupations in LLM-generated personas. There were 201 unique occupations, indicating a high degree of occupational diversity among the personas. The count shows the five most common occupations and samples of occupations mentioned only once.

Occupations with the given frequency	Frequency
Graphic Designer, Real Estate Agent, Accountant	12
Retired	11
Barista, College Student	10
Architect, Construction Worker, Student	9
Event Planner, Nurse, Chef	8
...	...
Instagram Influencer, Science Teacher, Stay-at-home dad, Customer service representative,	1
Sales Assistant	
Other occupations	185

Table 4: Five most common terms in the personas’ pain points (the list has been cleaned from words such as ‘and’, ‘of’, and so on).

Term	Frequency
work	120
relationships	119
financial	116
performance	77
stress	72

Regression modeling was carried out to predict each word’s frequency based on age, gender, and addiction type. The regression analysis for *work* showed a significant relationship with addiction types (gambling, opioids, shopping, and social media) but not with age or gender. Among the addiction types, gambling had the strongest negative relationship with *work* ($\beta = -1.5848$, $p < 0.001$), followed by shopping ($\beta = -1.5209$, $p < 0.001$) and social media ($\beta = -0.9426$, $p = 0.006$).

The word *relationships* showed a significant relationship with addiction types (gambling, shopping, and social media) but not with age or gender. Among the addiction types, shopping had the strongest negative relationship with *relationships* ($\beta = -0.7024$, $p < 0.001$).

The word *financial* showed a significant relationship with addiction type (gambling) but not with age, gender, or other addiction types. Gambling had a strong positive relationship with *financial* ($\beta = 1.6406$, $p < 0.001$).

The word *performance* showed a significant relationship with age and addiction type (gambling, shopping, and social media) but not with gender or opioid addiction type. Among the addiction types, shopping had the strongest negative relationship with *performance* ($\beta = -0.4897$, $p < 0.001$).

The word *stress* showed a significant relationship with addiction types (gambling, shopping, and social media) but not with age or gender. Among the addiction types, shopping had the strongest negative relationship with *stress* ($\beta = -0.7912$, $p < 0.001$), followed by social media ($\beta = -1.3773$, $p < 0.001$).

In summary, there is no evidence of age or gender bias in LLM-generated personas’ pain points. However, the analyses reveal that

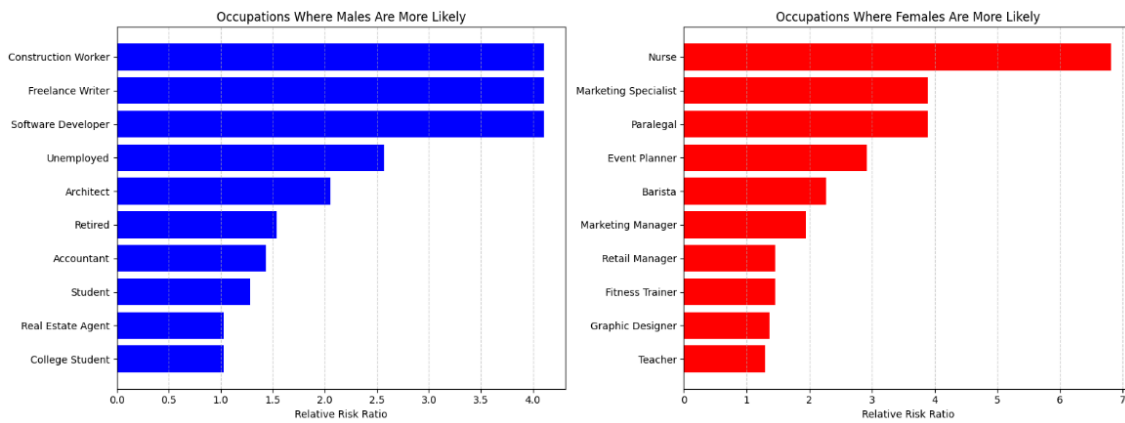


Figure 5: Relative risk ratios of occupations by gender.

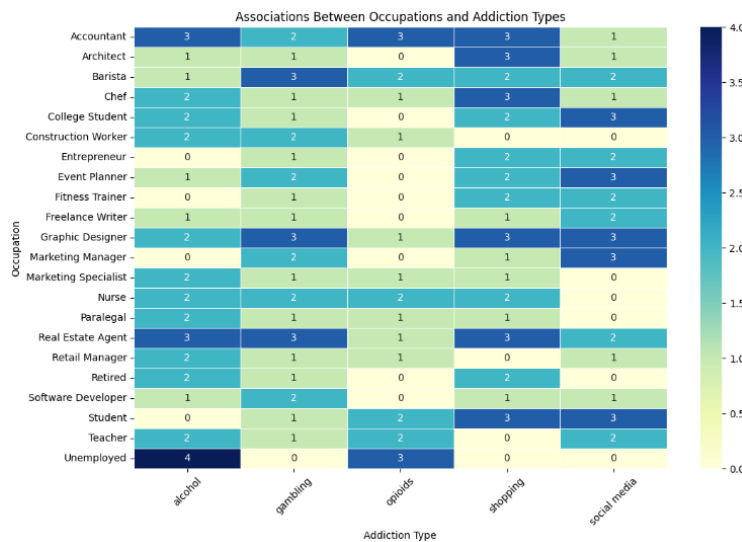


Figure 6: The frequency of occupations in addiction types.

the LLM appears to assign certain life conditions more commonly to some addiction types than others. (Full regression results are included in the online supplemental material.)

4.1.6 Physical appearance. It was extremely rare that the LLM would emphasize physical appearance in the persona descriptions it created. The coders logged only eight of such cases (~1.8%). Even among these, detailed scrutiny showed that the physical appearance was mentioned in passing and related to the negative consequences of the addiction (“a distinct change in his physical appearance”; “steady decline in his overall physical appearance”). Of the four observations where physical appearance was mentioned as a specific attribute of the persona, one was about a male persona (“His good looks”) and three about a female persona (“Through her videos, she showcases her talent, personality, and her gorgeous looks”; “Felicity is of average height and has a petite figure”, “Standing 5’6” with a slender build, Carla has always been conscious of her weight

and appearance”). So, we conclude that the LLM does not consider physical appearance as a dominant descriptor in this context of persona creation (which is correct behavior, as it should *not* be).

4.1.7 Personality. While physical appearance was not typically mentioned, personality was. In this coding, we considered personality broadly as the persona’s nature or characters. The coders identified personality cues in 180 (~38.0%) personas, so including personality descriptions in the generated personas appears common behavior for the LLM. Some of the common themes in the way the LLM described the personas included:

(1) **Hardworking and dedicated:** Many of the personality descriptions emphasize traits such as being hardworking, diligent, ambitious, and having a strong work ethic. The personas are described as committed to their careers, families, and personal goals (e.g., “She is known for her dedication to her demanding career”, “As a dedicated teacher, Felicity is diligent and resourceful”, “a highly

skilled accountant with a strong work ethic”, “hardworking and seemingly responsible individual”, “known for his excellent organizational and problem-solving skills”, “known for his creativity and unique approach to visual aesthetics.”).

(2) **Compassionate and caring:** Several personas are described as warm, compassionate, loving, and devoted, especially towards their families or those they work with, such as children or students (e.g., “warm and compassionate person, loving and devoted mother”, “a deeply caring and empathetic individual”, “a dedicated and passionate teacher”).

(3) **Intelligent and creative:** Many personas are described as intelligent, creative, and having unique talents and skills in various fields (e.g., “intelligent”, “a highly skilled accountant with a strong work ethic”, “a talented and respected architect”, “talented and ambitious art school graduate”, “a bright and ambitious young man”, “creative, careful”, “has a passion for exploring new places and culture.”).

(4) **Extroverted:** Some personas are characterized as sociable, outgoing, and friendly, enjoying social interactions and engaging with others (e.g., “extrovert”, “friendly, social and outgoing young woman”, “an outgoing, friendly, and ambitious individual who values hard work and dedication”, “a sociable, outgoing person”, “fun-loving and adventurous individual who loves to travel and party.”).

(5) **Introverted:** Other personas are described as introverted, finding social situations challenging or preferring solitude (e.g., “introverted”, “self-reported introvert person”, “despite outward appearances, Leo is a deeply introverted individual who struggles in social situations.”).

These examples illustrate that the LLM used a diverse set of personalities and psychological characteristics to incorporate humanness in the persona descriptions. The LLM’s viewpoint on the individual is often positive – the persona is more portrayed as a protagonist than an antagonist (in fact, we could locate no case where the LLM would have vilified the persona or presented them as a bad person).

4.1.8 Text analyses. To ensure that no persona description would be identical or close to identical, we computed the Levenshtein distance (LD) between each description pair. This metric tells us how many character changes are needed to make the pair identical (so, a low value would indicate a highly similar text description). The obtained LD values indicate no description is identical ($M = 1931.65$, $SD = 222.72$, $Min = 1285.00$, $Max = 3025.00$). So, the LLM does not recycle the same descriptions across the different personas.

The average length of the LLM-generated persona descriptions was 381.78 words ($SD = 50.70$). This suggests a moderate amount of variability in the length. Although we do not have a baseline of human-generated personas to compare to (as far as we know, nobody has investigated the length of persona descriptions previously!), the length seems reasonable in the sense of giving adequate information about the personas.

We investigated if there was a difference in the word count between male and female personas, between personas of different ages, and between the addiction types. First, a Mann-Whitney U test indicates no statistically significant difference in persona description lengths in terms of word count between male ($M = 379.41$,

$SD = 54.16$) and female ($M = 384.10$, $SD = 47.08$) persona descriptions, $U = 23823.0$, $p = .282$. Second, Spearman’s correlation coefficient ($\rho = -0.1547$, $p = .001$) indicates a statistically significant but weak negative monotonic relationship between age and persona description length. It suggests that as age increases, persona descriptions tend to be shorter, but the strength of this relationship is relatively modest (see Figure 7a).

The results of the Kruskal-Wallis indicated a statistically significant difference in the word count of persona descriptions between addiction types, $H = 22.164$, $P = .0002$. Pairwise comparisons using Dunn’s test indicated that the word counts in persona descriptions were significantly different between gambling and shopping ($p = .0007$) and gambling and social media ($p = .0007$). No other differences were statistically significant. Overall, the lengths of the persona descriptions appear to be aligned, with no noteworthy bias observed (see Figure 7b).

4.2 RQ2: How do (a) UX researchers and (b) subject-matter experts assess the LLM-generated personas?

4.2.1 Quantitative results. Addressing RQ2, we found that the LLM-generated personas generally obtained high scores from the human evaluators. The scores shown in Figure 8 indicate a high degree of consistency, reliability, positivity, believability, and informativeness for design. In contrast, stereotypicality is low (which is desirable as this is a problem and not a virtue in personas [43]). So, these evaluations indicate no quality issues in the generated personas – quite the opposite.

We conducted a series of Welch’s t-tests to assess whether the differences in the ratings between the internal evaluators and SMEs were mixed, with some statistically significant differences for certain criteria. As there are six tests (one for each criterion), the Bonferroni-adjusted alpha value is $0.05/6 = 0.0083$. The results indicate no significant differences for *informativeness* ($t(194.00) = -2.19$, $p = .0300 > .0083$), *believability* ($t(191.57) = 1.49$, $p = .1368$), and *consistency* ($t(183.24) = 2.37$, $p = .019$). However, there were three significant differences in ratings given by internal evaluators and SMEs. First, the SMEs rated the personas more *stereotypical* than the internal evaluators did, $t(177.26) = -5.52$, $p < .0001$. Second, the SMEs rated the personas less *positive* than the internal evaluators did, $t(177.47) = 11.02$, $p < .0001$. Third, the SMEs rated the personas less *reliable* than the internal evaluators did, $t(180.78) = 4.56$, $p < .0001$.

In absolute terms, however, the scores given by the SMEs were not bad: the average stereotypicality score they gave ($M = 2.98$) was below the scale average which is four for a seven-point Likert scale, whereas all of the “desirable” persona traits were above four. So, there are two takeaways here: (1) SMEs gave LLM-generated personas lower quality scores than internal evaluators, but (2) both the quality scores given by SMEs and internal evaluators indicate rather “high” than “low” quality personas. Especially surprising is that consistency ranks the highest for both evaluator types, as consistency has traditionally been an issue with text generation [3].

There were no notable differences by gender of the personas (see Figure 9), and no measure significantly correlated with the

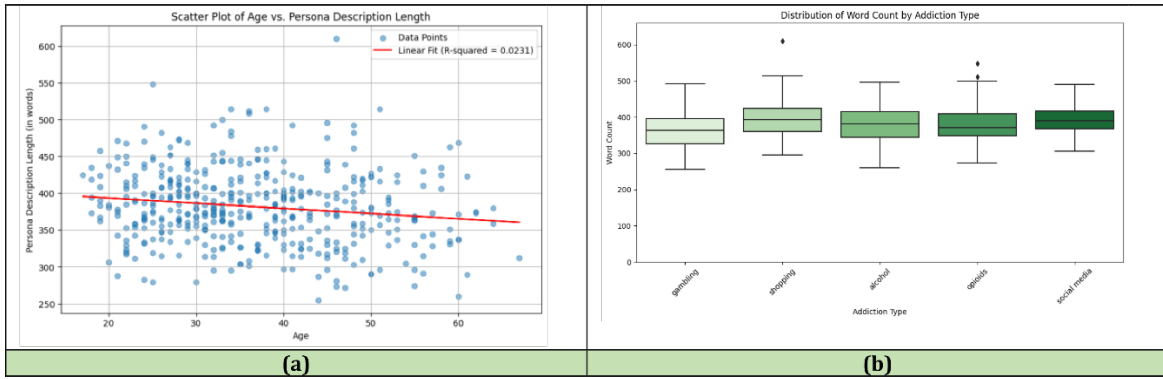


Figure 7: (a) the relationship between age and word count in persona descriptions. (b) The relationship between addiction type and word count.

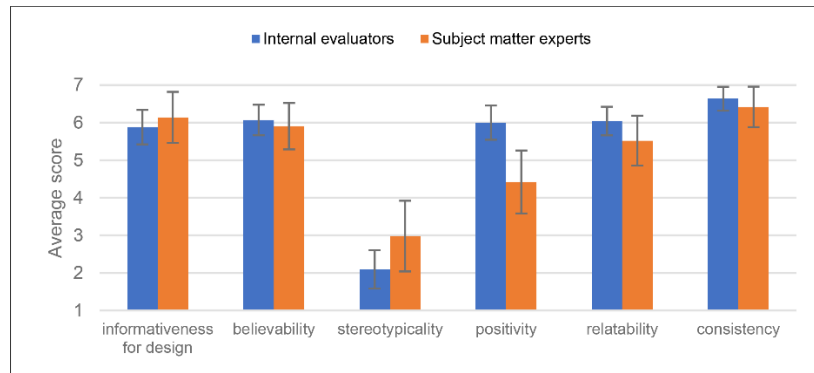


Figure 8: Average human evaluator scores across four raters. Example question: “Does the persona appear consistent? 1 = Not at all, 7 = Very much). The evaluators were provided with a definition of each criterion. Error bars indicate standard deviation.

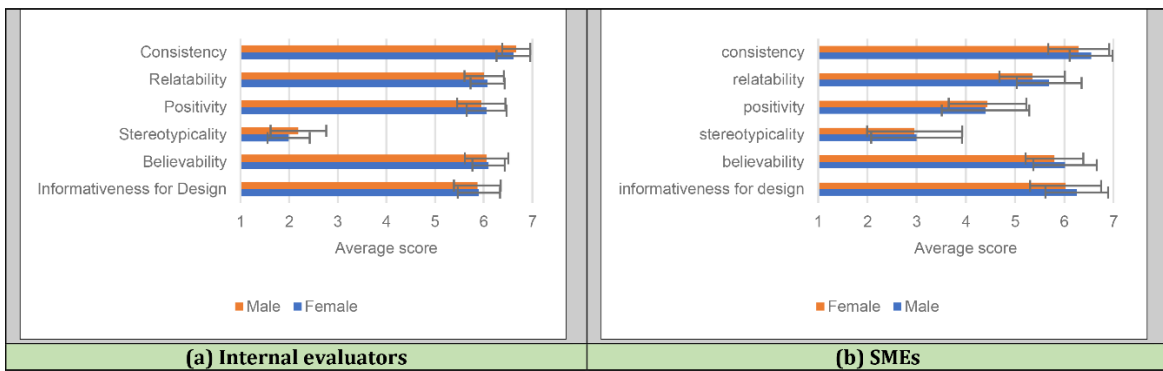


Figure 9: Evaluation scores by persona gender. No notable differences exist. We conducted a series of t-tests which are omitted from this manuscript (but available upon request) as no significant differences were found. Error bars indicate standard deviation.

persona’s age (details omitted due to parsimony, available upon request). So, the scores given by the human evaluators indicate no age or gender bias in terms of persona attributes.

4.2.2 Qualitative results. To better understand the scores given by the SMEs, we asked them to provide open-ended explanations

regarding their answers for each evaluation criterion. The full feedback by the SMEs is provided in Appendix 2; here, we summarize the main insights (note: positive comments are highlighted in green color, while critique or improvement suggestions are in red color, and “E” indicates evaluator ID):

Believability. Noteworthy comments were as follows:

- “The veterans’ personas were especially believable. **I found some of the shopping addictions a bit hard to believe, in particular the social worker and teachers.** Social workers and teachers usually struggle to make ends meet even without a shopping addiction.” (E2).
- “Overall, the majority of personas demonstrated were highly believable. The backgrounds presented combined with their high pressures in either their professional or social life made the scenarios seem very realistic, that they could be an actual person. For example, Persona ID P244, or Ava Chen, seemed very realistic, as I am sure the pressures of immigrating to an entirely new country and the challenges and barriers that exist with this transition seem paramount and never-ending. In addition, her trying to take care of her family and also being a schoolteacher must bring immense stress, leading her to unhealthy coping mechanisms such as alcohol.” (E3).
- “There were only a few personas whose background, profession, and addiction disorder did not quite add up. For example, Persona P64, or Stacey Rivers, seemed a bit off to me. Her escalation from winning at a charity casino night to a full-on gambling addiction seemed a bit extreme, combined with her background of being a schoolteacher.” (E3).

The feedback on personas highlighted a mix of believability and skepticism, with veteran personas being praised for their authenticity, while some personas, like those of social workers and teachers with shopping addictions, were questioned for their realism given their financial constraints. The detailed background stories, such as Ava Chen’s immigration challenges and resultant stress, were recognized for adding depth and realism, making the personas relatable and believable. However, some scenarios, like Stacey Rivers’ rapid descent into gambling addiction, were deemed unrealistic, suggesting a need for more nuanced development to align personas more closely with their professional and social contexts.

Relatability. Noteworthy comments were as follows:

- “The more personal details given about a persona, the more relatable I found them. It would have been helpful to have a little more info on the current important relationships in their lives.” (E2)
- “Overall, the majority of personas demonstrated appeared highly relatable and garnered much empathy. The caring and empathetic professions that many personas had, such as being teachers, social workers, environmental activists, etc., combined with their caring and connected backgrounds with family and friends, made their struggles with these negative coping mechanisms very relatable and highly sympathetic.” (E3)
- “**There were only a few that seemed off, with the stand-out being Persona P20, or Sean Hall,** who was a 34-year-old HVAC technician that still lived with his parents while having a gambling addiction.” (E3)

Feedback on the relatability of personas indicated that detailed personal stories enhanced empathy and connection, with suggestions for more insights into significant relationships to deepen relatability. The personas, particularly those in caring professions like teaching or social work, were largely seen as empathetic and

relatable due to their nurturing backgrounds and the realistic portrayal of their struggles with negative coping mechanisms. However, some personas, such as Sean Hall, the HVAC technician with a gambling addiction still living with his parents, were viewed as less relatable, suggesting the importance of aligning personal circumstances with professional and lifestyle choices for greater authenticity.

Consistency. Noteworthy comments were as follows:

- “For the most part, consistency was very good. **Natalia Thompson contained a significant discrepancy.** It first stated that she adopted a child but later said she had postpartum depression. This was confusing- did she adopt the child or give birth?” (E2)
- “Overall, I believe all of these personas were very consistent with their backgrounds, personalities, professions, and unhealthy coping mechanisms. While reading each persona, I really did not see any contradictions between their thoughts, emotions, or actions.” (E3)
- “**The only persona that stood out with conflicting information was (...) Natalia Thompson.** In her persona, it described her excitement and achievement of adopting a baby boy named Alex and then, however, described how **she was diagnosed with postpartum depression, which is depression experienced by women following child-birth.**” (E3)

Feedback on the consistency of personas was predominantly positive, highlighting their coherence in backgrounds, personalities, professions, and coping mechanisms, with no noticeable contradictions in thoughts, emotions, or actions. However, confusion arose with the persona Natalia Thompson, where there was an inconsistency regarding her situation; she was described as adopting a child but was also mentioned to have postpartum depression, a condition typically associated with childbirth, leading to questions about whether she adopted or gave birth. This discrepancy points to a need for clearer storytelling to avoid confusion and maintain the integrity of the personas’ narratives.

Informativeness for design. Noteworthy comments were as follows:

- “Overall, informativeness for design was very good. Personas #29 and #30, Yvette Patel and Anthony Rogers, seemed to describe benzodiazepine addictions rather than opioid addictions. Benzos are commonly prescribed for anxiety. It would be more unusual for someone to start an opioid addiction due to anxiety.” (E2)
- “For the most part, a lot of these personas described a good amount of information in relation to the individual’s background, relationships, emotions, motives, and professional goals, allowing designers to pin-point access to resources and information that could help the individual in managing their disorder.” (E3).
- “I believe the personas that did not provide a lot of information on what drives the individual to their unhealthy coping mechanism and their emotions during it were scored lower as it would be harder to find out exactly what resources could be used to really help the individual in their addiction.” (E3).

Feedback on the informativeness of personas for design highlighted their overall effectiveness, though it pointed out specific areas for improvement. For instance, Personas 29 and 30 were critiqued for inaccurately attributing benzodiazepine characteristics to opioid addictions, suggesting a need for more precise information regarding the nature of the addiction and its causes. The detailed backgrounds, relationships, emotions, motives, and professional aspirations provided in most personas were praised for giving designers clear insights into the individuals' needs, thereby facilitating the identification of relevant support resources. However, personas lacking detailed information on the motivations behind unhealthy coping mechanisms and the emotions experienced during these periods were viewed as less useful, indicating that a deeper exploration of these aspects could significantly enhance the design utility of the personas.

Stereotypicality. Noteworthy comments were as follows:

- “I feel that overall the personas were not too stereotyped. It would have been nice to see a little more diversity reflected in their names.” (E2).
- “Overall, I believe the majority of these personas were not deemed stereotypical scenarios. The majority of these personas each held unique backgrounds, emotions, and behaviors that are not widely held and fixed/oversimplified images or ideas of a particular person.” (E3).
- “For example, Persona P324, or Yvette Patel, being a single schoolteacher with crippling anxiety that led her to an opioid addiction seemed the opposite of stereotypical in our society.” (E3).
- “All personas made sense except a few. There was not much stereotyping.” (E5).
- “I gave high stereotypicality scores to a few personas (P327, P221, P162, P132, P324) because they looked more like a fiction story to me, as though an author is creating them from scratch and the people are fictional characters that do not exist, but if they do, they are moving about their routine life normally even though they are “addicted” to one thing or another.” (E5).

Feedback on the stereotypicality of personas indicated that they were generally perceived as non-stereotypical, with a call for greater diversity in naming to reflect broader inclusivity. The personas were praised for their unique backgrounds, emotions, and behaviors that went beyond fixed or oversimplified images, particularly highlighting examples like Yvette Patel, whose story as a single schoolteacher with anxiety leading to opioid addiction was seen as counter-stereotypical. While most personas were viewed as realistic and well-constructed, a few were critiqued for seeming more like fictional characters, with their life situations and addictions feeling too constructed and not reflective of real-life complexities. This feedback suggests a balance was largely achieved in avoiding stereotypes, but some personas could benefit from more grounded detailing to enhance their believability and avoid the impression of fiction.

Positivity. Noteworthy comments were as follows:

- “Overall, I believe the majority of these personas were presented in a more neutral light, compared to negative or positive depiction.” (E3).

- “Personas that were scored higher for positivity, such as James Patterson, were due to their recognition and actions in trying to manage their addiction and the positive lifestyles that they were trying to lead.” (E3).

Feedback on the positivity aspect of personas indicated that they were generally presented in a neutral manner, neither overly positive nor negative. Personas like James Patterson, who were scored higher for positivity, were distinguished by their proactive efforts to manage their addiction and their attempts to maintain or shift towards positive lifestyles. This approach underscores the importance of depicting personas in a balanced way that acknowledges their struggles while also highlighting their resilience and efforts towards recovery or positive change.

5 DISCUSSION AND IMPLICATIONS

5.1 Answers to Research Questions

RQ1 dealt with the diversity and bias of the LLM-generated personas. The results indicated that the LLM generated personas of different **ages**, from young to elderly people. That said, the LLM was biased toward younger age groups. The addiction types varied among personas from different age groups, but their variation appeared logical (i.e., younger groups struggled with social media addiction more often, the older groups with alcohol).

In terms of **gender**, the LLM generated the same number of male and female personas. Male personas were slightly older than female personas. In terms of **country**, the LLM generated personas from 15 different countries, although 86% of the personas were from the US, indicating strong bias. In terms of **occupation**, the LLM generated personas with 201 different jobs. Despite high diversity, there was some gender stereotyping, for example, males had a higher likelihood of being construction workers and females being nurses. The personas' **pain points** differed by addiction type, indicating that the LLM considered people with different addictions facing different types of challenges. In terms of **physical appearance**, the LLM rarely referred to the looks of the persona. In terms of **personality**, the LLM tended to portray the person in a positive light, highlighting positive traits over negative ones. There was no clear bias in terms of the length of the persona description, except that the length was slightly shorter for older personas.

Overall, the personas generated by LLM appear diverse. They do contain some biases, but the source and severity of these biases are difficult to assess. It appears that when humans perceive certain biases as harmful, these could be remedied by relatively minor interventions (e.g., changing the female personas profession from nurse to software developer). However, such assessments would need to be done on a case-by-case basis.

RQ2 dealt with the internal and external evaluation of the personas. The results indicate that LLMs can generate consistent personas that are perceived as believable, relatable, and informative for design while retaining a relatively low level of stereotyping, as perceived by the human evaluators. The potential of LLMs for persona generation lies in their possible capacity to generate immersive persona descriptions, incorporating demographics, motivations, pain points, and preferences based on a given set of inputs. As noted by Paoli [28], “. . .there is something powerful in the [ChatGPT] model since it knows what a user persona is without needing

any contextual explanation.” This property of fluency can explain why human evaluators give such high ratings to LLM-generated personas.

In the following, we offer some tentative explanations for the observed biases. *First*, the ‘youth bias’ might stem from the fact that machine learning (ML) datasets often over-emphasize younger demographics [14]. *Second*, the US-centricity might have a similar background, stemming from the fact that many ML training sets are based on English materials. We also must bear in mind that OpenAI is a US-based company, which might further accentuate the lack of cultural adaptation in its model’s behavior. *Third*, the LLM’s positive outlook on each persona (i.e., portraying predominantly positive personality traits and describing the person in a positive light) is likely due to Open AI’s guardrails on the output and corresponds to observations made by other researchers [9].

5.2 Practical Implications for Persona Design

An important observation for the practical deployment of LLMs is that GPT-4 seems to have an innate understanding of what a persona is, so it is by default able to start listing needs, pain points, attitudes, and so on [28]. So, the model that is supposed to create “mental models” in the form of personas has a mental model of its own when it comes to understanding what constitutes a persona! For the successful implementation of LLMs in the persona creation process, we propose the following guidelines:

- **1. Verify the LLM-generated personas using diversity and bias analysis techniques**, such as those illustrated in this work. There is not necessarily a need for complex analyses, but basic descriptive statistics go a long way.
- **2. Verify the LLM-generated personas using subject-matter experts** to establish external validity. Domain experts ought to be able to spot if there is ‘anything fishy’ about the personas.
- **3. Adjust the prompts if you observe challenges** in diversity, bias, or quality of the personas. Prompt design will substantially affect the characteristics of the generated personas. For example, the strong US-centricity of the LLM-generated personas could be addressed by instructing the LLM to generate personas from different countries.

By following these three guidelines, persona creators can mitigate the challenges and risks associated with using LLMs for persona generation. We also note that completely alternative approaches to making use of LLMs could be deployed, such as fine-tuning based on existing user or population data. These approaches are likely to emerge as the research on LLM-generated personas matures.

5.3 Limitations and Future Research

As with any study, ours includes some limitations. We discuss them here.

The reader should note that the generated personas are based on the general knowledge the GPT-4 model has about people with addictions. Apart from the SME evaluations, there was no additional verification of their factual correctness. Because the SMEs noted some inconsistencies in some of the generated personas, such

inconsistencies should be addressed before considering the application of the personas in any real-world scenario. As our evaluation of the personas primarily relied on subjective assessments from internal researchers and external SMEs, it would be beneficial to explore additional trustworthy sources, such as persona databases or real-world persona case studies, to compare the generated personas with those created through actual design processes. It was not verified whether these SMEs have experience in specific addiction domains or all of them. Future work could verify that as well as recruiting more SMEs to achieve more stable evaluation ratings. It is also possible that the SMEs did not understand all measurement criteria in the same way as the UX/HCI researchers did, specifically informativeness for design. Future research could cross-check SMEs’ baseline understanding of HCI metrics.

Inferring the nationality of personas based on their names within the context of addiction might pose problems. Names may be linked to a person’s place of birth or even to their parents, whereas the reasons for addiction might be more related to the current place of residence. These distinctions are essential for personas since a person’s place of birth and their current place of living may not necessarily be the same. For instance, all the personas listed in the paper might be living in the United States during their addiction journey. Future research could work to entangle the relationship between nationality and place of residence within LLM-generated personas.

Also, a significant contribution to HCI would be interpreting how to design prompt engineering to be more robust against biases in LLM generation. The initial idea on this is to ensure that the prompting covers protected classes and minority groups to generate personas from all possible user groups. The fact that specifically instructing the LLM to generate male or female personas resulted in 100% correct gender specification in the output supports the notion that the LLM can follow instructions concerning specific persona attributes.

Future research could investigate the textual content of LLM-generated personas using NLP techniques, similar to the study by Cheng et al. [9]. Multiple metrics could be deployed, including length, lexical diversity, sentiment, psycholinguistics, and so on. Linguistic analysis can reveal more insights into the diversity and bias in personas [34]. There is also a need for comparing LLM-enabled persona generation with the traditional persona generation processes, such as those based on user research and user behavior data. This would help emphasize the distinctions and unique characteristics of personas generated by LLMs.

It is not yet evident how LLMs will shape the persona-creation process. We have illustrated one possible approach, which is using the LLM’s foundational knowledge about people to generate personas. Another possibility is to ground the persona generation more strongly to specific datasets, whereupon the LLM becomes a “helper” in the analysis [28]. More studies are needed to test the pros and cons of integrating LLMs into the persona creation process.

As with any novel technology, LLM-generated personas come with possible harms. They can, either intentionally or inadvertently, have adverse societal effects, such as generating unreliable information, reinforcing gender stereotypes, affecting diversity representation, and deceiving users about the capabilities and limitations of

their actual degree of quality [1]. We did not focus especially on these risks, as it was not in the scope of our study. So, these risks warrant further scrutiny from the HCI research community. The risks should be weighed against the potential benefits to form a balanced perspective on the pros and cons of LLMs for HCI research and practice.

Overall, LLMs are rapidly transforming various spheres of society. HCI is not immune to their impact, neither is persona design. With this work, we have highlighted multiple avenues for future research on this topic which certainly warrants much more investigation from the HCI community. To facilitate replication of our study as well as further research on LLM-generated personas, we make our data and coding results publicly available (see the links to resources in Section 3). This supports the advancement of *persona science*, as called for in the literature [35].

6 CONCLUSION

Based on the findings, it can be concluded that LLM-generated personas exhibit diversity across various demographic and psychological dimensions. However, some biases are present, primarily related to age, occupation, and pain points. Younger age groups are overrepresented, and there is gender stereotyping in certain occupations. Additionally, there is a strong bias towards personas from the United States. Despite these biases, LLMs can generate consistent, believable, relatable, and informative personas for design purposes. Human evaluators generally perceive these personas positively, highlighting the fluency of LLMs in understanding and portraying user personas. It is important to note that while some biases are present, they appear to be addressable through minor interventions on a case-by-case basis. Overall, LLM-generated personas hold promise for design and user research applications, providing a foundation for further research.

ETHICAL REMARKS

The personas generated were not evaluated for factuality. They were evaluated for other factors such as believability and consistency. Because they were not evaluated for factuality, we do not recommend directly applying them for healthcare (or other) interventions. To generate personas for actual decision-making, we recommend either verifying the factuality of the personas generated using a general LLM like ChatGPT or then using factual data to finetune or otherwise adapt the LLM before the persona generation.

DECLARATION OF GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

During the preparation of this work, the author(s) used Open AI's ChatGPT (GPT-3.5 and GPT-4) as well as GPT-4 via API to generate the personas, assist us in the analysis, and provide material for addressing the 'blank page' problem in writing. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

REFERENCES

- [1] Gavin Abercrombie, Amanda Cercas Curry, Tanvi Dinkar, and Zeerak Talat. 2023. Mirages: On anthropomorphism in dialogue systems. *arXiv preprint arXiv:2305.09800* (2023).

- [2] Abeer Alessa and Hend Al-Khalifa. 2023. Towards Designing a ChatGPT Conversational Companion for Elderly People. *arXiv preprint arXiv:2304.09866* (2023).
- [3] Mostafa M. Amin, Erik Cambria, and Björn W. Schuller. 2023. Will Affective Computing Emerge From Foundation Models and General Artificial Intelligence? A First Evaluation of ChatGPT. *IEEE Intelligent Systems* 38, 2 (2023), 15–23.
- [4] Jisun An, Haewoon Kwak, Soon-gyo Jung, Joni Salminen, and Bernard J. Jansen. 2018. Customer segmentation using online platforms: isolating behavioral and demographic segments for persona creation via aggregated user data. *Social Network Analysis and Mining* 8, 1 (2018), 54. <https://doi.org/10.1007/s13278-018-0531-0>
- [5] Jisun An, Haewoon Kwak, Joni Salminen, Soon-gyo Jung, and Bernard J. Jansen. 2018. Imaginary People Representing Real Numbers: Generating Personas from Online Social Media Data. *ACM Transactions on the Web (TWEB)* 12, 4 (2018), 27. <https://doi.org/10.1145/3265986>
- [6] A. Baki Kocaballi. 2023. Conversational AI-Powered Design: ChatGPT as Designer, User, and Product. *arXiv e-prints* (2023), arXiv:2302.
- [7] Chris Chapman, Edwin Love, Russell P. Milham, Paul ElRif, and James L. Alford. 2008. Quantitative Evaluation of Personas as Information. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, September 01, 2008. 1107–1111. . <https://doi.org/10.1177/154193120805201602>
- [8] Chris Chapman and Russell P. Milham. 2006. The Personas' New Clothes: Methodological and Practical Arguments against a Popular Method. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, October 01, 2006. 634–636. . <https://doi.org/10.1177/154193120605000503>
- [9] Myra Cheng, Esin Durmus, and Dan Jurafsky. 2023. Marked Personas: Using Natural Language Prompts to Measure Stereotypes in Language Models. *arXiv preprint arXiv:2305.18189* (2023).
- [10] Hyunyi Cho, Lijiang Shen, and Kari Wilson. 2014. Perceived Realism: Dimensions and Roles in Narrative Persuasion. *Communication Research* 41, 6 (August 2014), 828–851. <https://doi.org/10.1177/0093650212450585>
- [11] Alan Cooper. 1999. *The Inmates Are Running the Asylum: Why High Tech Products Drive Us Crazy and How to Restore the Sanity* (1 edition ed.). Sams - Pearson Education, Indianapolis, IN.
- [12] Ameet Deshpande, Tanmay Rajpurohit, Karthik Narasimhan, and Ashwin Kalyan. 2023. Anthropomorphization of AI: Opportunities and Risks. *arXiv preprint arXiv:2305.14784* (2023).
- [13] Mohamed A. Elfeki, Mohamed A. Abdallah, Lorenzo Leggio, and Ashwani K. Singal. 2023. Simultaneous management of alcohol use disorder and liver disease: a systematic review and meta-analysis. *Journal of Addiction Medicine* 17, 2 (2023), e119–e128.
- [14] Raul Vicente Garcia, Lukasz Wandzik, Louisa Grabner, and Joerg Krueger. 2019. The Harms of Demographic Bias in Deep Face Recognition Research. In *2019 International Conference on Biometrics (ICB)*, June 2019. 1–6. . <https://doi.org/10.1109/ICB45273.2019.8987334>
- [15] Joy Ai-Leen Goodman-Deane, Mike Bradley, Sam Waller, and P. John Clarkson. 2021. Developing personas to help designers to understand digital exclusion. *Proceedings of the Design Society* 1, (2021), 1203–1212.
- [16] Joy Goodman-Deane, Sam Waller, Dana Demin, Arantxa González-de-Heredia, Mike Bradley, and John P. Clarkson. 2018. Evaluating Inclusivity using Quantitative Personas. In *In the Proceedings of Design Research Society Conference 2018*, June 28, 2018, Limerick, Ireland. Limerick, Ireland. . <https://doi.org/10.21606/drs.2018.400>
- [17] Jon E. Grant, Marc N. Potenza, Aviv Weinstein, and David A. Gorelick. 2010. Introduction to behavioral addictions. *The American journal of drug and alcohol abuse* 36, 5 (2010), 233–241.
- [18] Kathleen W. Guan, Joni Salminen, Soon-Gyo Jung, and Bernard J. Jansen. 2023. Leveraging Personas for Social Impact: A Review of Their Applications to Social Good in Design. *International Journal of Human-Computer Interaction* 0, 0 (September 2023), 1–16. <https://doi.org/10.1080/10447318.2023.2247568>
- [19] Muhammad Bilal Gulfranz, Muhammad Sufyan, Mekhail Mustak, Joni Salminen, and Deepak Kumar Srivastava. 2022. Understanding the impact of online customers' shopping experience on online impulsive buying: A study on two leading E-commerce platforms. *Journal of Retailing and Consumer Services* 68, (September 2022), 103000. <https://doi.org/10.1016/j.jretconser.2022.103000>
- [20] Perttu Hämäläinen, 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)*, April 19, 2023, New York, NY, USA. Association for Computing Machinery, New York, NY, USA, 1–19. . <https://doi.org/10.1145/3544548.3580688>
- [21] Matthew K. Hong, Shabnam Hakimi, Yan-Ying Chen, Heishiro Toyoda, Charlene Wu, and Matt Klenk. 2023. Generative AI for Product Design: Getting the Right Design and the Design Right. *arXiv preprint arXiv:2306.01217* (2023).
- [22] Soon-gyo Jung, Joni Salminen, and Bernard J. Jansen. 2021. All About the Name: Assigning Demographically Appropriate Names to Data-Driven Entities. In *Proceedings of the 54th Hawaii International Conference on System*

- Sciences, 2021, Virtual conference. Virtual conference. . Retrieved from <http://hdl.handle.net/10125/71108>
- [23] Tara Matthews, Tejinder Judge, and Steve Whittaker. 2012. How do designers and user experience professionals actually perceive and use personas? In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems - CHI '12*, 2012, Austin, Texas, USA. ACM Press, Austin, Texas, USA, 1219. . <https://doi.org/10.1145/2207676.2208573>
- [24] Lene Nielsen. 2002. From User to Character: An Investigation into User-descriptions in Scenarios. In *Proceedings of the 4th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques (DIS '02)*, 2002, London, England. ACM, London, England, 99–104. . <https://doi.org/10.1145/778712.778729>
- [25] Lene Nielsen. 2019. *Personas - User Focused Design* (2nd ed. 2019 edition ed.). Springer, New York, NY, USA.
- [26] 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* 7, 1 (2015), 45–61. <https://doi.org/10.4018/ijksd.2015010104>
- [27] Christian Nyemcsok, Hannah Pitt, Peter Kremer, and Samantha L. Thomas. 2023. Viewing young men's online wagering through a social practice lens: implications for gambling harm prevention strategies. *Critical Public Health* 33, 2 (2023), 241–252.
- [28] Stefano Paoli. 2023. Writing user personas with Large Language Models: Testing phase 6 of a Thematic Analysis of semi-structured interviews.
- [29] 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 (PDC 04)*, 2004, Toronto, Ontario, Canada. ACM, Toronto, Ontario, Canada, 112–120. . <https://doi.org/10.1145/1011870.1011884>
- [30] Joni Salminen, Kathleen Guan, Soon-gyo Jung, Shammur Absar Chowdhury, and Bernard J. Jansen. 2020. A Literature Review of Quantitative Persona Creation. In *CHI '20: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, April 25, 2020, Honolulu, Hawaii, USA. ACM, Honolulu, Hawaii, USA, 1–14. . <https://doi.org/10.1145/3313831.3376502>
- [31] Joni Salminen, Kathleen Guan, Lene Nielsen, Soon-gyo Jung, Shammur Absar Chowdhury, and Bernard J. Jansen. 2020. A Template for Data-Driven Personas: Analyzing 31 Quantitatively Oriented Persona Profiles. In *Human Interface and the Management of Information. Designing Information. HCI 2020*, S. Yamamoto and H. Mori (eds.), Springer, Copenhagen, Denmark, 125–144.
- [32] Joni Salminen, Kathleen W. Guan, Soon-Gyo Jung, and Bernard J. Jansen. 2022. Use Cases for Design Personas: A Systematic Review and New Frontiers. In *2022 ACM Conference on Human Factors in Computing Systems (CHI'22)*, 2022, New Orleans, USA. ACM, New Orleans, USA.
- [33] Joni Salminen, Bernard Jansen, and Soon-Gyo Jung. 2022. Survey2Persona: Rendering Survey Responses as Personas. In *Adjunct Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '22 Adjunct)*, July 04, 2022, New York, NY, USA. Association for Computing Machinery, New York, NY, USA, 67–73. . <https://doi.org/10.1145/3511047.3536403>
- [34] Joni Salminen, Soon-Gyo Jung, Shammur Chowdhury, Dianne Ramirez Robillos, and Bernard J. Jansen. 2021. The Ability of Personas: An Empirical Evaluation of Altering Incorrect Preconceptions About Users. *International Journal of Human-Computer Studies* (March 2021), 102645. <https://doi.org/10.1016/j.ijhcs.2021.102645>
- [35] Joni Salminen, Soon-Gyo Jung, and Bernard Jansen. 2022. Developing Persona Analytics Towards Persona Science. In *27th International Conference on Intelligent User Interfaces (IUI '22)*, March 22, 2022, New York, NY, USA. Association for Computing Machinery, New York, NY, USA, 323–344. . <https://doi.org/10.1145/3490099.3511144>
- [36] Joni Salminen, Soon-gyo Jung, João M. Santos, and Bernard J. Jansen. 2019. Does a Smile Matter if the Person Is Not Real?: The Effect of a Smile and Stock Photos on Persona Perceptions. *International Journal of Human-Computer Interaction* 0, 0 (September 2019), 1–23. <https://doi.org/10.1080/10447318.2019.1664068>
- [37] Joni Salminen, Rohan Gurunandan Rao, Soon-gyo Jung, Shammur A. Chowdhury, and Bernard J. Jansen. 2020. Enriching Social Media Personas with Personality Traits: A Deep Learning Approach Using the Big Five Classes. In *Artificial Intelligence in HCI (Lecture Notes in Computer Science)*, 2020, Cham. Springer International Publishing, Cham, 101–120. . https://doi.org/10.1007/978-3-030-50334-5_7
- [38] Joni Salminen, João M. Santos, Soon-gyo Jung, and Bernard J. Jansen. 2023. How does an imaginary persona's attractiveness affect designers' perceptions and IT solutions? An experimental study on users' remote working needs. *Information Technology & People* 36, 8 (January 2023), 196–225. <https://doi.org/10.1108/ITP-09-2022-0729>
- [39] 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, (April 2020), 102437. <https://doi.org/10.1016/j.ijhcs.2020.102437>
- [40] Albrecht Schmidt. 2023. Speeding Up the Engineering of Interactive Systems with Generative AI. In *Companion Proceedings of the 2023 ACM SIGCHI Symposium on Engineering Interactive Computing Systems*, 2023, 7–8.
- [41] Phillip Douglas Stevenson and Christopher Andrew Mattson. 2019. The Personalization of Big Data. *Proceedings of the Design Society: International Conference on Engineering Design* 1, 1 (July 2019), 4019–4028. <https://doi.org/10.1017/dsi.2019.409>
- [42] Piotr Tarka, Monika Kukar-Kinney, and Richard J. Harnish. 2022. Consumers' personality and compulsive buying behavior: The role of hedonistic shopping experiences and gender in mediating-moderating relationships. *Journal of Retailing and Consumer Services* 64, (2022), 102802.
- [43] Phil Turner and Susan Turner. 2011. Is stereotyping inevitable when designing with personas? *Design studies* 32, 1 (2011), 30–44.
- [44] Jennifer C. Villeux, Peter J. Colvin, Jennifer Anderson, Catherine York, and Adrienne J. Heinz. 2010. A review of opioid dependence treatment: pharmacological and psychosocial interventions to treat opioid addiction. *Clinical psychology review* 30, 2 (2010), 155–166.
- [45] Xishuo Zhang, Lin Liu, Yi Wang, Xiaoli Wang, Hailong Wang, Anqi Ren, and Chetan Arora. 2023. PersonaGen: A Tool for Generating Personas from User Feedback. *arXiv preprint arXiv:2307.00390* (2023).
- [46] Noam Zilberman, Gal Yadid, Yaniv Efrati, Yehuda Neumark, and Yuri Rassovsky. 2018. Personality profiles of substance and behavioral addictions. *Addictive behaviors* 82, (2018), 174–181.

A PERSONA IDS FOR SME EVALUATION

Persona IDs: P327, P01, P410, P414, P182, P221, P197, P246, P83, P369, P449, P44, P274, P298, P20, P61, P426, P162, P64, P132, P247, P380, P316, P111, P244, P139, P99, P324, P348, P171

B OPEN-ENDED FEEDBACK FROM SUBJECT MATTER EXPERTS

Question	Evaluator 1	Evaluator 2	Evaluator 3	Evaluator 4	Evaluator 5
<i>Believability - what was good and what seemed off?</i>	The personas are quite believable and realistic in their portrayal of individuals struggling with addiction.	The veterans' personas were especially believable. I found some of the shopping addictions a bit hard to believe, in particular the social worker and teachers. Social workers and teachers usually struggle to make ends meet even without a shopping addiction.	Overall, the majority of personas demonstrated were highly believable. The backgrounds presented combined with their high pressures in either their professional or social life made the scenarios seem very realistic, that they could be an actual person. For example, Persona ID P244, or Ava Chen, seemed very realistic, as I am sure the pressures of immigrating to an entirely new country and the challenges and barriers that exist with this transition seem paramount and never ending. In addition, her trying to take care of her family and also being a schoolteacher must bring immense stress, leading her to unhealthy coping mechanisms such as alcohol. There were only a few personas whose background, profession, and addiction disorder did not quite add up. For example, Persona P64, or Stacey Rivers, seemed a bit off to me. Her escalation from winning at a charity casino night to a full-on gambling addiction seemed a bit extreme, combined with her background of being a schoolteacher.	Good: Yes, I think all personas appear realistic. Seemed off: In some persona like Lila Bennett little more detail is needed to make them appear like an actual person.	The first few personas were not believable (as I have explained in the consistency section) but as I moved down the questions, the personas started making sense and turned believable especially those related to social media addiction, PTSD opioid/alcohol dependence and online shopping. Some personas about single parents giving in to alcohol dependence due to pressure of raising a child were not believable.

<i>Relatability - what was good and what seemed off?</i>	The personas are mostly relatable in the sense that they show case different types of people who could be affected by addiction.	The more personal details given about a persona, the more relatable I found them. It would have been helpful to have a little more info on the current important relationships in their lives.	Overall, the majority of personas demonstrated appeared highly relatable and garnered much empathy. The caring and empathetic professions that many personas had, such as being teachers, social workers, environmental activists, etc., combined with their caring and connected backgrounds with family and friends, made their struggles with these negative coping mechanisms very relatable and highly sympathetic. Especially due to the fact that a large portion of these individuals recognized their unhealthy disorder and were trying to find ways to remedying it and alleviate the stress and pain it causes both to them and those who care for them. There were only a few that seemed off, with the stand-out being Persona P20, or Sean Hall, who was a 34-year-old HVAC technician that still lived with his parents while having a gambling addiction.	Good: Yes, according to me all personas are relatable	Most of the questions made sense and there were real-life examples to relate to. Only 2-3 personas did not look relatable to me, and I have explained in the consistency section.
<i>Consistency - what was good and what seemed off?</i>	The personas are mostly consistent in their portrayal of addiction and the associated challenges	For the most part, consistency was very good. Persona ID #6 (Natalia Thompson) contained a significant discrepancy. It first stated that she adopted a child but later said she had postpartum depression. This was confusing- did she adopt the child or give birth?	Overall, I believe all of these personas were very consistent with their backgrounds, personalities, professions, and unhealthy coping mechanisms. While reading each persona, I really did not see any contradictions between their thoughts, emotions, or actions. The only persona that stood-out with conflicting information was Persona 221, or Natalia Thompson. In her persona, it described her excitement and achievement of adopting a baby boy named Alex and then, however, described how she was diagnosed with postpartum depression, which is depression experienced by women following childbirth.	Good: All personas look consistent but, in Lila Bennett's description, a slight contradiction is there	Almost all personas showed consistency and I gave them high scores except P327 and P221. Let me explain why. P327 did not look consistent to me. Samantha was shown in such a positive light, i.e., young, determined and committed yet she caved into the pressure of her business in just 5 years. This is not just too early but also inconsistent for a person who has studied and prepared for a job/business all their life. How can you give up on something you're so passionate about and that too so quickly? The personas turned a sharp turn from positivity to negativity. Therefore, I gave it low consistency scores. P221 also looked inconsistent. Natalie was passionate and looked like an achiever. She wanted a career, she studied and strived for it; she wanted a baby, she adopted one without waiting to meet the right man and making a biological baby. She looks like a doer, yet she allowed herself to be crushed under the burden of responsibilities. For such a strong and independent person who takes big decisions, Natalie seemed to have a sudden and abrupt shift in her behavior. She ran her house effectively before adopting the baby. There are day care centers for babies. She has a stable job, there is no way she should be quitting so soon. It's abrupt and inconsistent, in my opinion.

Informativeness for design - what was good and what seemed off?

The personas are informative for design in the sense that they provide designers with a deeper understanding of the experiences and need of individual struggling with addiction.

Overall, informativeness for design was very good. Personas #29 and #30, Yvette Patel and Anthony Rogers, seemed to describe benzodiazepine addictions rather than opioid addictions. Benzos are commonly prescribed for anxiety. It would be more unusual for someone to start an opioid addiction due to anxiety.

Overall, I believe the majority of these personas presented an adequate amount of information to design an app or system to address the persona's needs. For the most part, a lot of these personas described a good amount of information in relation to the individual's background, relationships, emotions, motives, and professional goals, allowing designers to pin-point access to resources and information that could help the individual in managing their disorder. I believe the personas that did not provide a lot of information on what drives the individual to their unhealthy coping mechanism and their emotions during it were scored lower as it would be harder to find out exactly what resources could be used to really help the individual in their addiction.

Yes, to design an app adequate amount of information about personas is present

The informativeness factor was really good. All personas except 3 were educational and informative. Persona #1 looks so conflicting. Samantha is young and determined. It looks unbelievable that she'd fall for addiction in such a young age and just after 5 years of starting her business. Her challenges do not look big enough to affect her mental health so much. The persona looks off and somewhat inconsistent, unbelievable, and unreliable. P414 also looks weird. I have worked as a consultant and after spending so much time online, smart phone and digital screens are the last thing I want in my routine. Sitting all day and staring at the screen is so torturing. You want to run away from it, not indulge in it. P221 looked totally made up. Natalia is not the first and the only working woman in the country. Nearly all women work full-time and the challenge if raising a baby is not big enough to turn her into an alcoholic. The scenario looked stereotypical and made up.

Positivity - any comments on this?

The personas do not necessarily focus on positivity, as they are meant to depict the challenges and struggles associated with addiction.

There was a wide range in positivity among the personas. Some only mentioned the persona's career and no other details about them. The more we know about a person's positive attributes (strengths, interests, achievements), the more relatable and believable it is.

Overall, I believe the majority of these personas were presented in a more neutral light, compared to negative or positive depiction. Because the majority of these personas give a balanced background between the individual's successes and challenges, I believe they seemed human for the most part and were scored in a more neutral zone. Personas that were scored higher for positivity, such as P449, or James Patterson, were due to their recognition and actions in trying to manage their addiction and the positive lifestyles that they were trying to lead.

I feel that very few personas show positive behavior toward their addiction

Save a few personas, there was a huge positive factor in all. Many people were inherently good but fell prey to their circumstances. They all realized their addiction/dependence status and showed inclination to seek help. I gave low positivity scores to a few personas (P327, P01, P414, P410, P182, P197, P83, P162, P64, P171) because they had a very little positivity element. The description showed they were either spoiled or totally messed up from the start and made no effort to improve their life even though they had families to support and responsibilities on their shoulders. Many personas did not even realize they had an addiction or needed help. Their personas were off.

<p><i>Stereotypicality - any comments on this?</i></p>	<p>The personas do not seem to rely on stereotypes, as they appear to be based on real life experiences.</p>	<p>I feel that overall the personas were not too stereotyped. It would have been nice to see a little more diversity reflected in their names.</p>	<p>Overall, I believe the majority of these personas were not deemed stereotypical scenarios. The majority of these personas each held unique backgrounds, emotions, and behaviors that are not widely held and fixed/oversimplified images or ideas of a particular person. The ones that particularly seemed off were in the correlation of their professions and behaviors to their addiction disorder. For example, Persona P324, or Yvette Patel, being a single schoolteacher with crippling anxiety that lead her to an opioid addiction seemed the opposite of stereotypical in our society.</p>	<p>I don't feel that in the case of any persona there is an oversimplified or fixed image</p>	<p>All personas made sense except a few. There was not much stereotyping. The personas looked real and made sense and were relatable. I gave high stereotypicality scores to a few personas (P327, P221, P162, P132, P324) because they looked more like a fiction story to me, as though an author is creating them from scratch and the people are fictional characters that do not exist, but if they do, they are moving about their routine life normally even though they are "addicted" to one thing or another.</p>
<p><i>Any other comments or remarks about the personas?</i></p>	<p>The personas are well crafted and provide a useful starting point for designers to understand and empathize with individuals struggling with addiction.</p>	<p>I thought these were well-done overall. If the aim is believability and relatability, I would suggest adding a little more detail about the personas' personal lives and strengths.</p>	<p>From reading these personas, it makes you realize how the myriad of pressures and stressors from everyday life can really lead anyone from any background and walk of life towards the path of addiction, especially compared to the stereotypes we may see in our society.</p>	<p>Good: All personas are relatable, one can easily understand them, just in some personas a little more details or clarity of their characteristics needs to be added</p>	<p>Good work. It's not easy to create believable, relatable and consistent personas. You've created very fine personas of people who look, sound and feel very much real.</p>