

# Creating and Evaluating Personas Using Generative AI: A Scoping Review of 81 Articles

Danial Amin  
University of Vaasa  
Vaasa, Finland  
danialam@uwasa.fi

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

Farhan Ahmed  
School of Marketing and  
Communication  
University of Vaasa  
Vaasa, Finland  
farhan.ahmed@student.uwasa.fi

Sonja M.H. Tervola  
Faculty of Education  
University of Cambridge  
Cambridge, United Kingdom  
sonjatervola@gmail.com

Sankalp Sethi  
College of Information Science  
University of Arizona  
Tucson, Arizona, USA  
sankalpsethi@arizona.edu

Bernard J. Jansen  
Qatar Computing Research Institute  
Hamad Bin Khalifa University  
Doha, Qatar  
jjansen@acm.org

## Abstract

As generative AI (GenAI) is increasingly applied in persona development to represent real users, understanding the implications and limitations of this technology is essential for establishing robust practices. This scoping review analyzes how 81 articles (2022-2025) use GenAI techniques for the creation, evaluation, and application of personas. The articles exhibited good level of reproducibility, with 61% of articles sharing resources (personas, code, or datasets). Furthermore, conversational persona interfaces are increasingly provided alongside traditional profiles. However, nearly half (45%) of the articles lack evaluation, and the majority (86%) use only GPT models. In some articles, GenAI use creates a risk of circularity, in which the same GenAI model both generates and evaluates outputs. Our findings also suggest that GenAI seems to reduce the role of human developers in the persona-creation process. To mitigate the associated risks, we propose actionable guidelines for the responsible integration of GenAI into persona development.

## CCS Concepts

• **Human-centered computing** → **HCI theory, concepts and models**; *User models*.

## Keywords

Generative AI, LLM, Personas, Persona Development

## ACM Reference Format:

Danial Amin, Joni Salminen, Farhan Ahmed, Sonja M.H. Tervola, Sankalp Sethi, and Bernard J. Jansen. 2026. Creating and Evaluating Personas Using Generative AI: A Scoping Review of 81 Articles. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 23 pages. <https://doi.org/10.1145/3772318.3790608>



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI '26, Barcelona, Spain*

© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2278-3/26/04  
<https://doi.org/10.1145/3772318.3790608>

## 1 Introduction

User personas ('personas' for short) are fictitious characters representing archetypal real users of a system, product, or service presented in a humanized form [25]. Since their introduction into human-computer interaction (HCI), personas have been adopted to support user-centered design (UCD) in domains such as software development [1, 3, 12], design [20, 29, 70], and healthcare [36, 43, 44]. Personas help system designers focus on user needs throughout the design process by providing empathetic and memorable representations of core user segments [112]. As such, personas are a critical component of HCI research, user experience (UX) studies, and UCD of systems and technology.

Persona development methods were originally based on qualitative data collection with manual analysis, which can be time-consuming, resource-intensive, and prone to staling [117]. However, generative artificial intelligence (GenAI), particularly large language models (LLMs), creates opportunities for automating persona development [45, 55, 129]. Automatic persona development is the ability to generate personas without (or minimal) human involvement and can overcome the challenges of manual persona development. This is why GenAI is being researched to support persona development. These GenAI technologies can be applied to various stages of persona development, and the personas created by such a process are called GenAI personas, which will be the focus of this study (see Figure 1 for examples).

Although various articles have applied GenAI in persona development [55, 119, 126, 129], there is limited systematic synthesis of how GenAI technologies are applied in persona development practices. This lack of knowledge raises serious questions for the HCI community regarding the use of GenAI for persona creation.

Moreover, the application of GenAI in persona development suffers from a lack of standardization and best practices, there is yet no clear consensus on optimal approaches. The application of GenAI in persona development is currently uncharted territory, making it difficult for persona developers to make informed decisions about adopting GenAI technologies. The evolution of GenAI technologies has also increased the challenges faced by persona developers, as discussed in detail by Amin et al. [2]. Similar to the development of data-driven personas (see a review [117]), we



**Figure 1: Examples of GenAI personas. The persona in (a) is generated by Survey2Persona [55], (b) is generated by Persona-L [134]. These systems illustrate that LLMs are increasingly used for generating persona profiles, but also to afford decision-makers access to dialogue with the personas. The figures are annotated to explain different components of GenAI personas. The reader is encouraged to zoom in for better readability.**

can observe varying methodological choices and technical implementations. These choices involve a variety of factors, such as the selection of LLM family (e.g., GPT or Claude or Llama), versions of LLMs (e.g., GPT 3.5 or GPT 4.0), hyperparameters (i.e., parameters that impact the behavior of the model [49]), and prompts (i.e., instructions given to the LLM to perform certain tasks). Although researchers have suggested guidelines for persona development in general [25, 48, 51, 98, 113], **how researchers specifically use GenAI in persona development remains unclear.**

Another domain that is prone to challenges is the evaluation of GenAI personas. Traditional persona validation methods generally focused on user perceptions of personas [52, 121], but GenAI personas require additional consideration of factors, such as output consistency, prevention of hallucinations (the ability of LLMs to fabricate information), and prompt reliability [69, 123]. Moreover, integrating GenAI raises questions about evaluating the quality of GenAI-crafted versus human-crafted personas [2]. Evaluation approaches vary widely, from automated metrics to user studies, **without a clear understanding of how researchers validate and evaluate GenAI personas.**

The ethical aspects of using GenAI technologies may extend beyond traditional concerns about persona development. Though prior work has addressed issues of stereotyping and representation in personas [37, 38, 138], it appears that GenAI amplifies certain challenges, such as algorithmic bias, data privacy, and transparency [2, 41, 47, 111]. These challenges may become particularly acute as personas represent diverse user groups and influence design decisions that affect these populations, prompting **the need to understand how researchers factor in ethical concerns when using GenAI in persona development.**

Against the backdrop of these research challenges, we put forth three research questions (RQs):

- **RQ1: How are GenAI technologies used in persona development?**
- **RQ2: How are GenAI personas evaluated?**

- **RQ3: What ethical considerations are associated with GenAI personas?**

Building on persona research in HCI [112, 117] and following guidelines for literature reviews in the HCI field [66], this scoping review analyzes 81 articles published between 2022 and 2025 on the application of GenAI for persona development. A scoping review generally maps the current scope, key concepts, and knowledge gaps [6], and is suitable for studying GenAI personas given this domain’s rapidly evolving nature. To this end, we identify trends, gaps, and opportunities, and we provide recommendations to practitioners concerning the integration of GenAI into persona development workflows. Considering the proliferation of GenAI tools and their application in user research and design processes, this work addresses a timely and relevant topic for the HCI community.

## 2 Previous Reviews and Research Gap

In UX/HCI practice, personas serve as shared reference points that help designers empathize with users and make decisions aligned with user needs [112]. Persona development traditionally involves collecting user data through research methods (interviews, surveys, analytics), identifying distinct user segments through analysis, and crafting humanized narrative profiles that represent these segments [25, 117]. At first glance, GenAI appears to strengthen *automatic persona generation* by exhibiting capability in most of these above-mentioned tasks. Automatic persona development has existed in the HCI domain for more than a decade [54, 56, 88]. However, the limitations of automated personas have been apparent, especially when analyzing text records of users and writing up persona narratives [118], which is a key activity in many persona creation approaches [97]. Algorithms and models prior to LLMs struggled with interpretive tasks, such as writing persona narratives or labeling user segments [51], which has limited the potential for automatic persona generation, as extensive manual interventions were required in conjunction with data science methodologies [50, 88].

Due to the limited interpretative capabilities of previous AI models, researchers have applied rule-based systems, such as dynamic templates with predefined fields for dynamically inserted information [98]. In contrast, LLMs’ interpretative capabilities offer possible solutions to the major challenges of contextually interpreting and summarizing textual data about people [55]. However, questions remain and the best practices of LLMs in persona development remain unknown, including the ideal techniques and processes to follow [119].

There are also practical hurdles; for example, there is no clear concept of human-AI collaboration in persona development, even though researchers seem to agree on the need to maintain “human in the loop” instead of forfeiting all decision-making to GenAI [129]. Interestingly, human in the loop is a feat that is also considered valuable in persona theory, such that involving stakeholders that are intended to use the personas in the persona creation process supports the development of favorable attitudes toward the persona technique [31, 95]. Additionally, it is not evident how well HCI researchers are aware of the risks and challenges associated with implementing this GenAI for developing personas, including ethical concerns such as reinforcing existing systemic biases or

marginalizing underrepresented groups when lacking real data on cultural contexts [2].

These lingering concerns call for a literature review that investigates existing practices, methodologies, and emerging patterns in GenAI persona development. To this end, reviews of persona research have addressed quantitative personas [116], data-driven personas [117], personas for social impact [39], and persona design applications [122], among other areas. Although these reviews provide valuable insights into persona development methods, they do not specifically address the integration of GenAI technologies into persona development. We will discuss in Section 5 how our findings are positioned in the continuum of persona research, specifically in conjunction with previous reviews. Given the growing adoption of GenAI in persona development [23, 47, 111, 119, 129, 134], a scoping review of this emerging approach is needed and thus presented.

### 3 Methodology

Following previous work in HCI [15, 100, 141] that deals with the ramifications of novel technologies, we opted for a scoping review to address our RQs. Scoping reviews map the breadth of research in emerging areas, identify key concepts, and examine how research is conducted [6, 71]. We conducted a literature search in five academic databases: (1) *ACM Digital Library (ACMDL)*, (2) *IEEE Xplore*, (3) *Web of Science (WoS)*, (4) *Scopus*, and (5) *arXiv*. The search focused on identifying our articles at the intersection of GenAI and persona development. Our methodology consisted of (1) query development and search execution, (2) screening of articles based on inclusion/exclusion criteria, (3) backward and forward snowball sampling, (4) quality assessment, and (5) data extraction and synthesis. The search process began with developing keyword groups related to GenAI and persona development, involving iterations of query refinement and implementing database-specific adaptations. Through screening, we identified 90 relevant articles of the 885 total found in the databases. Additional snowball sampling yielded 27 potentially relevant articles. Three additional articles from snowballing met our inclusion criteria, resulting in a final corpus of 93 articles, of which 12 (12.3%) were excluded in the detailed screening, leaving a total of 81 (9.15%) articles included in the corpus. The details of the search strategy, the screening process, and the analysis are described in the following subsections.

#### 3.1 Search Query Creation

The creation of the search query focused on identifying relevant articles at the intersection of GenAI and persona development, with terms from both domains (see Table A1 in the online appendix). These terms were selected from prominent articles in the field, personal domain expertise, and LLM suggestions (Claude, ChatGPT, Llama). All searches were performed in the title, abstract, and author keywords fields in all databases. We searched all articles until August 25, 2025. Initial searches revealed a high false positive rate when using wildcard operators (e.g., “persona\*”), with unrelated terms such as “personality analysis”. We used phrase matching with exact terms instead (i.e., including “persona” and “personas” in the query explicitly).

The search strings required database-specific adaptations to accommodate varying syntax requirements and technical limitations.

For ACMDL, we used the independent field tags for different items (i.e., [Abstract:] field tag for abstract, [Title:] field tag for title, and [Keyword:] field tag for keywords). For IEEE Xplore, we searched the title, abstract, and keywords fields using the ALL Metadata field. For Scopus, we covered title, abstract, and keywords using the TITLE-ABS-KEY field specification. For WoS, we used the Topic (TS) field, which includes title, abstract, and keywords. Scopus required proximity operators (W/5), while Web of Science needed simplified syntax for compatibility.

The exact search string used across databases with database-specific syntax adaptations is provided in Section 1. Boolean operators combined the keyword groups using “AND” relationships between different concept groups, such as GenAI terms AND persona terms. We allowed “OR” within each group to capture terminological variations.

The complete evolution of search strings across five iterations, including abandoned variants, is documented in Section A1 of the online appendix<sup>1</sup>. Each string revision was tested by verifying the capture of relevant articles while excluding clearly irrelevant ones. The final query structure balanced breadth and specificity in identifying articles.

#### 3.2 Selection of Databases

We investigated five academic databases for our scoping review based on their relevance and coverage. We created different versions for each of the databases (see Section A2 of the online appendix). ACMDL was chosen as the primary venue for HCI research, providing comprehensive coverage of persona development work in computing contexts. We excluded ACM Guide of Computing Literature to avoid duplication with ACMDL. IEEE Xplore was selected for its strong focus on technical implementations, covering engineering aspects of GenAI applications in persona development. WoS and Scopus were included to ensure broad interdisciplinary coverage, as persona research spans multiple HCI-related fields, including design, marketing, and social sciences.

Finally, *arXiv* was included to capture recent developments in GenAI applications to personas, particularly given the rapid progress of LLMs. The inclusion of preprints from arXiv follows established practices in rapidly evolving GenAI research domains to reduce the risk of excluding articles with negative or null findings, particularly in fields where peer review lag could exclude cutting-edge developments [94, 127]. We only included preprints from the past two years to focus on recent developments.

During deduplication, we prioritized entries based on the quality of the database metadata and the relevance of the domain. ACMDL entries were retained over other sources, followed by IEEE Xplore, WoS, and finally Scopus.

#### Search Query

“generative artificial intelligence” OR “large language model” OR “large language models” OR LLM OR LLMs OR “generative model” OR “generative models” OR GPT OR “generative pre-trained transformer” OR “foundation model” OR “foundation models” OR

<sup>1</sup>[https://osf.io/au5g6/?view\\_only=29c0a9a1251b43a186e1359ac9cd6a10](https://osf.io/au5g6/?view_only=29c0a9a1251b43a186e1359ac9cd6a10)

“ChatGPT” OR “Gemini” OR “LLAMA” OR “Claude”) AND (“data-driven persona” OR “data-driven personas” OR “quantitative persona” OR “quantitative personas” OR “user persona” OR “user personas” OR “LLM-generated persona” OR “LLM-generated personas” OR “buyer persona” OR “buyer personas” OR “persona” OR “personas” OR “marketing persona” OR “marketing personas” OR “design persona” OR “design personas” OR “persona development” OR “persona creation” OR “persona generation”)

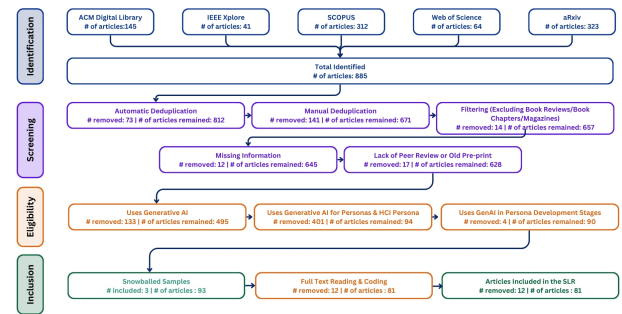
### 3.3 Article Collection and Screening

The article collection and selection process followed the PRISMA guidelines [102], as illustrated in Figure 2. The initial database search yielded 885 articles: 64 from WoS, 312 from Scopus, 145 from ACM DL, 41 from IEEE Xplore, and 323 from arXiv. The screening process consisted of multiple stages. First, we performed automatic deduplication using reference management software, removing 73 articles, bringing the total to 812. This was followed by manual deduplication that identified 141 additional duplicates due to minor variations in article titles (e.g., case changes, encoded characters, missing information across databases). After removing conference reviews ( $n = 8$ ), books and book chapters ( $n = 6$ ), and other nonrelevant document types ( $n = 12$ ), 645 articles remained. The next phase involved multiple screening assessments of the 645 articles. This included verification of English language requirements (0 removed) and peer review status (17 removed). The subsequent screening for the GenAI application removed 133 articles, and the alignment assessment of the definition of HCI and the use of GenAI in personas eliminated 401 articles. Additionally, four articles were removed during the assessment stage, leaving 90 articles.

The snowball sampling process identified 27 potentially relevant articles. The snowball sample was identified by the Google Scholar Alerts service during the research period. Of these, 8 were already in our initial sample, and 5 were already in the final corpus. After screening for English language (1 removed), peer review status (0 removed), GenAI use (3 removed), and alignment of HCI definition (6 removed), 4 articles remained. One article was removed during the assessment stage, resulting in 3 additional articles for our corpus.

**3.3.1 Inclusion Criteria.** Articles were included if they met all the following criteria:

11. Written in English to ensure accurate analysis
12. Either peer-reviewed publications or preprints from the past two years
13. Demonstrated application of generative AI in persona development
14. Aligned with the HCI definition of personas as user representations for design. HCI definition of personas focuses on the representation of a group of users.
15. Presented empirical research through qualitative, quantitative, or mixed methods
16. Addressed at least one stage of persona development: data collection, analysis, generation, evaluation, or practical application



**Figure 2: PRISMA diagram describing the literature collection process. We followed the standard protocols in scoping literature reviews [15, 100, 141] to collect our literature base.**

**3.3.2 Exclusion Criteria.** Articles were excluded if they met any of the following criteria:

- E1. Non-academic content: letters, editorials, book reviews, meeting abstracts, news items
- E2. Articles only briefly mention personas or GenAI without substantial implementation
- E3. Articles using GenAI for non-persona-related work
- E4. Research using the term “persona” in non-user-representation contexts
- E5. Articles using personas solely to enhance conversational agents’ personality or behavior style.
- E6. Purely theoretical discussions without implementation
- E7. Opinion pieces lacking research methodology
- E8. Duplicate publications of the same study

The criterion **I4** is considered because research uses ‘persona’ in at least two different meanings. HCI describes personas as a representation of a group of users [25, 112, 117], whereas natural language processing (NLP) refers to persona as the personality, role, attributes, or collection of design choices that are made for the conversational agent [65, 110, 145]. For instance, among the articles excluded under this criterion, articles [61, 68, 73, 136] discuss personas for LLMs and conversational agents as personality configurations (e.g., “helpful assistant” or “creative writer”) rather than user group representations for design purposes. Similarly, the criterion **E5** differs from the conversational persona format (21.0% of included articles, see Section 4.1.5). The included articles involving conversational personas, use conversational interfaces to represent user groups for design decisions, whereas the excluded articles use personas to configure chatbot’s personalities without the intention of representing real users or user groups.

Through this screening process, we identified a corpus of 93 articles (90 from the database search + 3 from snowball sampling) for detailed analysis. Of these, 12 were removed from the full text reading stage (see details in Section 3.4.2), so the final number of articles reviewed was 81 (9.15%).

### 3.4 Data Extraction

**3.4.1 Data Extraction Book.** To systematically analyze the selected articles, we developed a data extraction coding book with 35 variables in six categories (see Table A3 in the online appendix): (1) basic information, (2) resources, (3) research methodology, (4) persona development, (5) GenAI / LLM usage, and (6) results and discussion. The coding book was developed by the Principal Investigator (PI) and Co-Principal Investigator (Co-PI). We reviewed the relevant coding books created for similar literature reviews on personas and included the necessary variables for our research context. We pilot tested the coding book by coding five articles at random and revised the coding book to address any challenges or ambiguities in the instructions. We also mapped the extracted variables to the RQs and identified under/over utilized variables (see Table A2 in the online appendix). We defined specific extraction instructions and value formats for each variable to ensure consistent coding across articles (see Table A3 in the online appendix). The codebook included six categories, including the basic information category, which captures publication details, and the resources category, which tracks the availability of artifacts such as code and datasets. Research methodology concerns approaches and empirical methods; persona development focuses on creation processes; and GenAI/LLM usage details how GenAI technologies are employed. Finally, the results and discussion category captures the findings, limitations, and future directions. Coders applied the codebook to each article in the corpus. For the 25 categorical variables, coders selected from predefined options listed in the codebook (see Table A3 in the online appendix). For the 10 open-ended variables, four coders independently coded responses (with 10% overlap for agreement calculation). The open-ended variables were then coded thematically into different themes, which were jointly developed by the PI and Co-PI.

**3.4.2 Data Extraction Process.** Data extraction was carried out independently by four coders. Initially, we conducted a pilot test with all four coders to code 3 randomly selected articles to verify and refine our extraction protocol. This pilot phase helped clarify variable definitions and standardize the extraction approach among coders. The coding book included a total of 35 variables that were required to be coded. Out of the 35 variables, 25 were categorical variables with structured categories, while 10 were open-ended questions. For the 25 categorical variables, the choices were included in the coding book, and the coders selected one/more than one option from them. For the 10 open-ended questions, two researchers encoded them into different sub-categories and categories according to different RQs. The four coders (post-graduate HCI students) who performed the primary extraction have background knowledge in personas and technical expertise in computing and AI (see Table A4 in the online appendix). Each coder was assigned a subset of articles (three coders analyzed 22 articles each, while one coded 42 articles). All four coders independently analyzed 9 (11.11% of the corpus) common articles from the corpus to assess inter-rater reliability. Inter-rater reliability was assessed through data for 17 coded variables (except open-ended questions and basic information) among four coders using Fleiss' Kappa, which indicated a strong agreement ( $\kappa = 0.86$ ).

The information extraction process involved reading each article and coding the 35 variables according to the predefined protocol in the coding book. Each coder was given the standardized extraction protocol with the coding book, with the option of discussions to resolve any ambiguities. During the extraction process, 12 articles were further removed from the corpus, bringing the total number for analysis to 81. Eleven articles were removed because of their unconventional structure and brevity, which resulted in many missing key attributes. One article was excluded due to doubts about its methodological rigor.

## 4 Findings

### 4.1 RQ1: How Are GenAI Technologies Used in Persona Development?

**4.1.1 Technology Adoption and Usage Patterns.** Our data show that OpenAI's GPT models appear in 70 of 81 articles (86.4%), establishing GPT as the predominant choice. Claude (Anthropic) and Llama (Meta) each appear in 7 articles (8.6%), while Gemini (Google) is used in 6 articles (7.4%). The "Others" category includes 19 articles (23.5%) using models such as Deepseek, Mixtral, Bard, and GLM. Most of the articles employ only one LLM ( $n = 53$ , 65.4%), while 20 articles (24.7%) use multiple LLMs together. The remaining 8 articles (9.9%) do not mention their LLM usage.

The reason for using multiple LLMs tends to be their comparison using quantitative metrics. For example, Xu et al. [144] compared Claude-3.5, Gemini-1.5, and Llama-3 for persona generation across different metrics, and Sethi et al. [128] tested GPT-4, Gemini, and Claude for lexical diversity analysis.

The integration of conventional ML approaches with GenAI methods occurs in a relatively few ( $n = 14$ , 17.3%) articles. Notable examples of this hybrid approach include using NLP to extract entities from textual data for prompt enrichment [149], and using clustering techniques to segment embeddings before persona generation [76]. A key takeaway is **OpenAI's GPT models overwhelmingly prevail GenAI persona research, creating a possible dependency of personas on this one commercial actor and accentuating limited methodological diversity.**

**4.1.2 Usage Across the Persona Development Lifecycle.** We used a taxonomy adopted from persona literature [112, 137] to organize GenAI usage patterns into four themes (see Table 1): (1) persona creation and enrichment, (2) persona-based simulation and interaction, (3) persona-informed personalization, and (4) validation and evaluation.

**Persona Creation and Enrichment:** We found three main approaches to using GenAI in persona development. First, *data-driven persona generation* creates personas from structured datasets (e.g., Likert-based surveys [55] and using grouped consumer data [76]). Second, *narrative persona profiles and enrichment* create detailed storytelling personas with rich backstories [9, 126]. Third, *multimodal persona representation* combines textual descriptions with visual elements [140, 149]. Figure 3 shows an example workflow for multimodal personas.

Participatory approaches appear in 30.9% of articles ( $n = 25$ ), while 65.4% ( $n = 53$ ) do not involve direct user participation. The

remaining 3 articles do not explicitly mention their use of participatory design. Here, participatory persona development involves direct user feedback from either (a) the persona users or (b) the people represented by the personas to modify the personas iteratively through feedback sessions [134].

*Persona-based Simulation and Interaction:* One novelty of GenAI is the ability of LLMs to orient personas in interactive scenarios, which can be implemented in different ways. For example, *conversational persona agents* enable direct user interaction with personas [79, 114], such as using LLMs for character imagination and story generation in a dialogue environment [108] (see Figure 4). Another example is using *multi-agent persona simulation*, which involves multiple personas interacting simultaneously; for example, role-playing language agents in decision-making scenarios [27] and deploying multi-agent persona simulation systems for training purposes [115].

*Persona-informed Personalization:* This application adapts LLM outputs based on persona characteristics. *Personalized content and recommendations* include articles exploring how personas can guide content generation and recommendation systems. *Adaptive dialogue and instruction* adjust communication styles, as exemplified by the *Persona-L* system that developed adaptive dialogue capabilities to represent people with complex needs, allowing the system to adjust its communication style according to persona characteristics [134].

The use of LLM with personas involves the opportunity for HCI researchers to communicate with personas, rather than passively scanning persona profiles. An example workflow for an adaptive persona is visible in Figure 5.

*Validation and Evaluation:* The articles apply LLMs in different ways to assess the quality and consistency of personas. First, *consistency checking* examines whether personas maintain coherent characteristics across scenarios, such as using consistency measurements to validate personas' textual coherence [133], and developing evaluation metrics such as fluency and helpfulness to assess persona quality [74]. Second, *realism assessment* evaluates human-likeness through metrics such as lexical diversity [128] and understanding hallucinations in LLM-generated personas [58]. Third, *bias detection* identifies stereotypical representations, as exemplified by Panda et al. [103] who investigated whether GenAI personas perpetuate harmful stereotypes.

**4.1.3 Data Sources and Representation Approaches.** Data collection practices are split among real user data, which appears in approximately half ( $n = 41$ , 50.6%) of the articles, synthetic data in 32.1% ( $n = 26$ ), and mixed approaches in 17.3% ( $n = 14$ ). Researchers used real data collected directly from the users through surveys [55, 76] or indirectly collected through observations [40]. Synthetic data applications, for example, used artificially generated data to generate personas representing addiction patterns [128] and investigated the consistency of LLM-generated personas using personality tests [26]. For example, Bai et al. [9] used mixed methods combining real and synthetic data to create virtual populations by sampling census-based skeletal personas and enriching them with LLM-generated details.

Persona representation formats reveal both traditional and emerging patterns. Traditional persona profiles are the dominant form

of presenting personas ( $n = 53$ , 65.4%), providing structured demographic and behavioral information in standard persona templates [77, 83, 86]. In contrast, narrative formats appear in approximately a third of the articles ( $n = 26$ , 32.1%); for example, De Paoli [105] produced persona narratives using LLM-assisted thematic analysis, while Schuller et al. [126] generated persona narratives as user stories that UX experts found comparable with human-written personas. Furthermore, personas are presented as chatbots in 21.0% of articles ( $n = 17$ ), typically in interactive persona interfaces [23, 134] that combine profile and dialogue formats for the persona user (see Figure 6). Finally, another notable development is providing persona information in structured data formats ( $n = 5$ , 6.2%), for example, using a nested JSON output that makes it straightforward to quantitatively analyze the persona attributes [119].

**4.1.4 Application Contexts.** GenAI personas are applied across diverse domains and contexts (see Table A6 in the online appendix for a complete list of articles). Product development represents the most common application context ( $n = 9$ , 11.1%), in which researchers create data-driven personas to improve technology design and UX. Marketing contexts ( $n = 7$ , 8.6%) employ LLMs for consumer segmentation through text embeddings and clustering [76], and create demographic profiles that mirror real participants to generate synthetic survey responses [60], among other uses.

Education contexts ( $n = 6$ , 7.4%) include, for example, creating chatbot personas representing children with disabilities to support accessible technology design [79], whereas healthcare contexts ( $n = 5$ , 6.2%), for example, compare LLM-generated personas with human-crafted personas in clinical research [132] and represent people with special conditions like Down syndrome [134]. In articles focused on civic good and social applications ( $n = 7$ , 8.6%), researchers analyzed GenAI personas in societally meaningful contexts, such as addiction research [119, 128] and investigated systematic biases in GenAI systems [103]. Research enhancement contexts, in turn, represent 3.7% ( $n = 3$ ) of applications, in which researchers explain persona creation processes for novice designers through collaboration with GenAI tools [55] and use expert perspectives to improve research ideation [78].

The remaining 54.3% ( $n = 44$ ) of applications span diverse contexts, including sustainability, finance, content creation, accessibility design, and language processing, indicating that GenAI personas are being adopted across more or less every domain that traditionally employs personas.

**4.1.5 Alternative Development Approaches.** GenAI personas incorporate at least four novelties or “innovations” (see Table 2): (1) input source innovations, (2) process innovations, (3) prompt engineering innovations, and (4) output format innovations. We explain these innovations as follows.

Input source innovations ( $n = 26$ , 32.1%) expand the types of data used for persona generation beyond traditional surveys and analytics. *Visual and multimodal input* integrates images, avatars, and visual representations into persona profiles, expanding beyond text-only descriptions. Multimodal data improves GenAI personas by combining textual analysis with visual generation. Persona systems can pair GPT-4 survey processing with DALL-E 2 avatar creation [149], or integrate emotional expressions for privacy education

<sup>2</sup>[https://osf.io/3pwav?view\\_only=29c0a9a1251b43a186e1359ac9cd6a10](https://osf.io/3pwav?view_only=29c0a9a1251b43a186e1359ac9cd6a10)

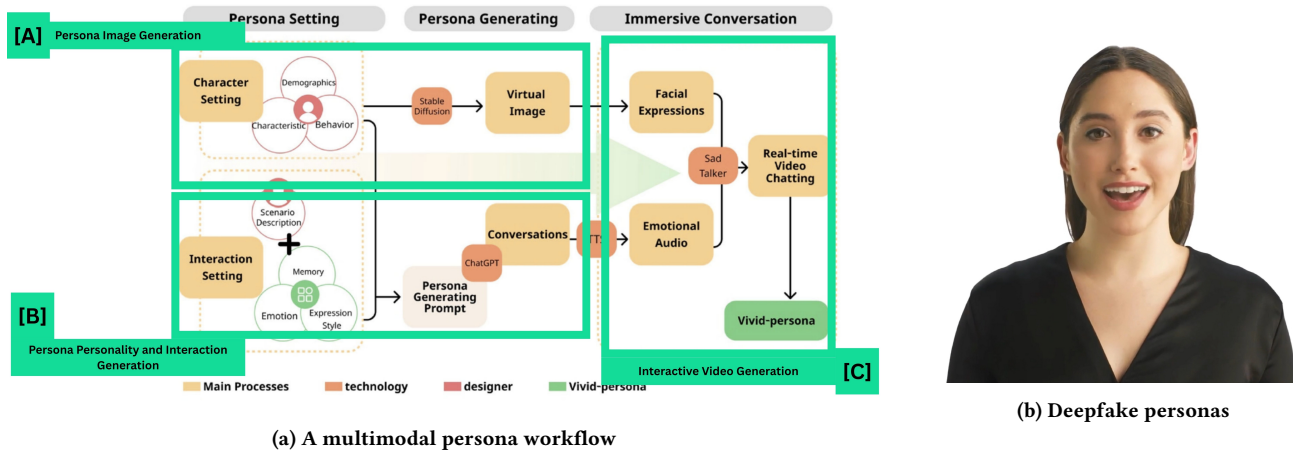


Figure 3: (a) Vivid-Persona system architecture [151] with annotated phases: Persona Setting ([A]), Persona Generating ([B]), and Immersive Conversation ([C]), combining Stable Diffusion, ChatGPT, and real-time video generation for interactive persona creation. (b) A deepfake persona created using GenAI, illustrating how “harmful” technology [46, 59] can be repurposed for positive user representation. A video version of the deepfake persona is shared with the author’s permission in our online supplementary material<sup>2</sup> to further illustrate the interactive nature of some GenAI persona types. The figure is annotated to explain different components of a GenAI persona development system.

Table 1: LLM usage in persona development with examples and specific applications. While LLMs are primarily used in persona creation and enrichment stages, other modes that use LLMs are simulation, interaction, personalization, validation, and evaluation. See Table A5 in the online appendix for a complete list of articles.

Main theme	Sub-theme	Examples	Specific application example
Persona creation and enrichment	Data-driven persona generation	[55, 76, 129]	PersonaCraft processes survey responses to automatically generate personas with demographic profiles and segmented information from the survey
	Narrative persona profiles and enrichment	[9, 105, 126]	Creating rich storytelling personas with detailed backstories, personality descriptions, and contextual narratives for design scenarios
	Multimodal persona representation	[18, 140, 149]	Generating personas that combine textual descriptions with visual avatars, mood boards, and lifestyle imagery
Persona-based simulation and interaction	Conversational persona agents	[23, 79, 108]	Developing interactive chatbots that embody specific personas, allowing users to have conversations with target user archetypes
	Multi-agent persona simulation	[27, 115, 144]	Simulating group discussions between multiple persona agents to understand consensus-building and decision-making processes
Persona-informed personalization	Personalized content and recommendations	[76, 90]	Tailoring marketing messages and product recommendations based on specific persona characteristics and preferences
	Adaptive dialogue and instruction	[60, 134]	Adjusting communication style and accessibility features based on users’ cognitive abilities and interaction preferences
Validation and evaluation	Consistency checking	[57, 74, 133]	Evaluating whether persona responses remain consistent across different scenarios and interaction contexts
	Realism assessment	[128, 132]	Measuring linguistic patterns to assess how human-like generated personas appear
	User preference prediction	[60, 147]	Testing whether LLM-generated personas can accurately predict real user preferences in product choices and behavioral decisions
	Bias detection	[103, 109]	Analyzing generated personas for demographic stereotypes, cultural biases, and unfair generalizations across different user groups

[18], with the idea that visual elements strengthen memorability beyond text-only profiles. Similarly, unstructured qualitative data extends persona development into contexts where collecting structured, quantitative data is impractical. For example, LLMs can extract personas from video content and industry reports [134]

or generate personas from interview transcripts through thematic analysis [104]. Expert knowledge can be indirectly used to generate personas; for example, by capturing domain expertise through graph-based simulation [78] or transforming expert knowledge into narrative profiles [60].

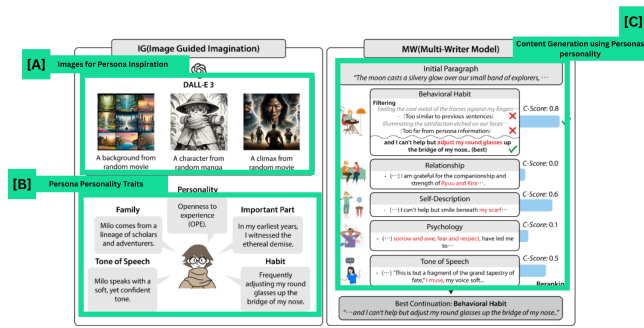


Figure 4: A framework using images for persona generation [108] with two annotated components: Image Guided Imagination (IG) using DALL-E 3 ([A]) and Persona Personality Traits ([B]). The Multi-Writer Model([C]) generates content for specific personas, illustrating how GenAI personas can be embedded to automate decision-making in user-centric tasks. The figure is annotated to explain different components of a GenAI persona development system using different data sources.



Figure 6: Proxona, a multiple persona interactive tool [23] with annotated sections: Persona Customization ([D]), Multiple Persona Cards ([B]), and Expanded Persona Card ([C]). The tool demonstrates editing and interaction capabilities ([A]) with multiple personas, illustrating the complexities of creating full-fledged persona systems that rely on multiple system prompts using LLMs for persona creation, evaluation, and delivery. The figure is annotated to explain different components of a Proxona system.

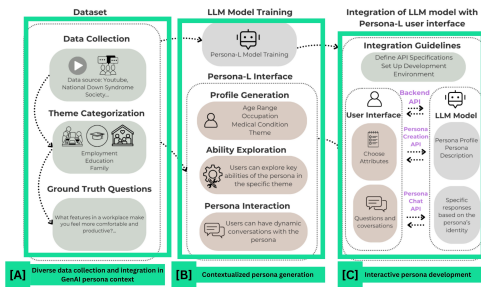


Figure 5: Persona-L framework [134] showing three annotated phases: Dataset collection and integration ([A]); LLM Model Training ([B]) featuring interface with profile generation, ability exploration, and persona interaction capabilities; and Integration with user interface ([C]) providing backend API connections for persona creation and chat functionality. Persona-L illustrates how GenAI models can contribute to varying tasks in persona creation, validation, and delivery using multimodal data sources. The figure is annotated to explain different components of a Persona-L system.

Process innovations (n = 30, 37%) change the fundamental workflow of persona generation, often bypassing steps that were previously considered essential. Data-driven persona development often follows a sequential pipeline of collecting data, identifying segments through clustering, and then enriching profiles with narrative details [56, 117]. The use of GenAI challenges this conventional pipeline. For example, rather than discovering segments in data, *direct attribute specification* enables researchers to specify desired characteristics upfront and have LLMs generate matching personas. Li et al. [76] demonstrated this by directly specifying consumer segments to generate personas for predefined market categories,

eliminating the need for any preprocessing stages altogether. Similarly, *segmentation-free approaches* that generate personas from data directly using theoretical frameworks without any empirical clustering. Sun et al. [134] used conceptual models and reports to generate personas for people with Down syndrome. *Generation pipelines* represent automating the entire workflow. Systems like PersonaCraft [55] and PersonaFlow [78] integrate all the stages of persona generation (data collection, analysis, and generation) into unified, automated processes.

Prompt engineering innovations (n = 18, 22.9%) apply advanced techniques to improve the quality and consistency of LLM-generated personas. *Two-stage prompting* separates the creation of the demographic skeletal profile from the finalization of the persona. For example, researchers first prompt demographic profile creation, followed by enrichment prompts to add behavioral and contextual information [119, 128]. Taking this idea further, *multistage prompts* are applied in different steps of the persona generation pipeline. For example, Panda et al. [103] used sequential prompts for data analysis, persona attribute extraction, and profile writing. Articles also apply *few-shot-learning prompting approaches* that include example personas in the prompt to guide LLM output style and structure. For example, Schuller et al. [126] included sample personas in prompts for bias evaluation, helping LLMs understand expected output format and detail level, while Chen et al. [18] used example privacy scenarios to generate contextually adapted personas.

Output format innovations (n = 39, 48.1%) change how personas are presented and interacted with, moving beyond static profile documents. *Narrative-driven personas* present information in a story format rather than attribute lists. Xu et al. [144] created personas with rich backstories and character arcs for creative applications, while Schuller et al. [126] generated first-person narratives that read like biographical sketches and can make personas more engaging

and memorable for stakeholders. *Multimodal representations* extend this engagement by combining text with visual and emotional elements. Chen et al. [18] integrated emotional expressions with visual representations to create empathy-based personas, while Zhang et al. [149] combined textual personas with auto-generated avatar pictures and lifestyle imagery. *Interactive persona systems* complete this transformation by presenting personas through conversational interfaces rather than static documents. For example, Persona-L [134] provides conversational interfaces for accessibility needs, enabling persona users to discuss requirements with persona representations.

## 4.2 RQ2: How Are GenAI Personas Evaluated?

**4.2.1 Evaluation Targets in GenAI Personas.** We identified four primary aspects evaluated in GenAI-generated persona research: (1) persona descriptions, (2) persona behavior, (3) generation process, and (4) outcomes and impact (Table 3). Among articles that reported evaluation ( $n = 44$ ), description quality was examined in 38.6% of articles, behavioral performance in 36.4%, outcomes/impact in 15.9%, and generation process in 11.4%.

*Persona Description Evaluation:* Description evaluation examines the static attributes and textual quality of persona profiles. Articles assess whether persona descriptions appear credible, internally consistent, and linguistically sound. Generally, this is done by using established HCI scales such as the Persona Perception Scale (PPS) [121]. The PPS is an instrument for evaluating individuals' perceptions of personas, comprising 8 measures such as credibility, completeness, consistency, clarity, and willingness to use [121]. Multiple articles applied PPS to evaluate description in GenAI personas [57, 132]. For example, Jung et al. [55] evaluated personas on five dimensions: consistency, clarity, completeness, credibility, and fluency. Computational approaches to description evaluation include linguistic analysis to measure lexical diversity [128] and stereotype detection using predefined lexicons [22]. These methods assess whether persona descriptions employ varied vocabulary and whether they contain problematic stereotypical representations.

*Persona Behavior Evaluation:* Behavior evaluation examines personas' performance in interactive contexts or when tasked with decision-making, prediction, or simulation activities. This aspect differs from description evaluation: description evaluation asks whether a persona profile appears credible, and behavior evaluation asks whether the persona acts in ways that match real human patterns. Articles employing behavioral evaluation tested personas in varied contexts. For example, Li et al. [76] validated GenAI personas against actual survey data, demonstrating that GenAI personas can simulate consumer preferences. Kaiser et al. [60] conducted pre- and post-intervention surveys measuring belief changes and information sharing patterns among real humans, simulated humans, and synthetic personas. Park et al. [108] tested the contextual appropriateness of personas' responses in naturalistic decision-making scenarios. However, behavioral evaluation remains less developed than description evaluation. Articles often assess behavior through indirect measures rather than testing of persona actions in controlled scenarios.

*Generation Process Evaluation:* A small subset of articles evaluated the efficiency, scalability, or usability of the persona generation

process itself. Zhang et al. [149] developed custom metrics measuring generation efficiency, satisfaction with the creation process, and collaboration quality between humans and GenAI systems. These evaluations address practical concerns about integrating GenAI personas into existing decision-making workflows rather than assessing the personas themselves.

*Outcomes and Impact Evaluation:* Outcome evaluation examines the downstream effects of using GenAI personas in design, decision-making, or educational contexts. For example, articles measured usability using the System Usability Scale (SUS) [57], assessed learning gains from persona-based educational interventions, and evaluated whether personas influenced design decisions appropriately. This evaluation category connects persona quality to real-world applications but remains relatively rare in the literature.

**4.2.2 Evaluation Methodologies.** Among all the articles, 44 articles (54.32%) mention the evaluation methodology used in the process, while 45.67% ( $n = 37$ ) do not provide a clear evaluation framework. The use of LLMs in persona evaluation remains limited, with a quarter ( $n = 12$ , 27.27%) of articles employing LLMs for evaluation purposes, while most studies applying evaluation ( $n = 32$ , 72.72%) rely on traditional evaluation techniques.

We identify three primary evaluation approaches from the articles that mentioned their evaluation approaches. (see Table 4): (1) human-driven evaluation, (2) computational evaluation, and (3) benchmark-based evaluation.

*Human-Driven Evaluation:* More than half of the articles ( $n = 25$ , 58.6%) that applied the evaluation process used human evaluators to judge persona quality, though it is equally noteworthy that nearly half (41.4%) did not use human evaluators. For example, Jung et al. [55] recruited 148 participants, including both regular users ( $n = 127$ ) and UX experts ( $n = 21$ ), reporting that GenAI personas scored significantly higher than non-LLM personas for consistency, clarity, completeness, credibility, and fluency. Another article measured GenAI personas for authenticity, clarity, empathy, and willingness to use, finding satisfactory quality [151]. Chat-based personas were compared with profile-based personas using both PPS and SUS, with significant differences in usability between interactive and static personas [57].

Expert analysis represents a valuable approach within human-driven evaluation. Sun et al. [134] used qualitative interviews with domain experts to scope how accurately personas represent people with complex needs. Another article recruited 127 clinical and educational experts to compare LLM-generated personas with human-crafted personas in the obesity research domain [132]. Discrimination tests provide another human-driven approach. Researchers have used Turing-test style evaluation methods in which a human has to differentiate between human-generated personas and GenAI personas [126]. As an example, a two-phase study with 31 participants compared single versus multiple expert perspectives for relevance, helpfulness, and creativity [78]. Another article conducted qualitative interviews to evaluate the effectiveness of chat personas by examining empathy in educational contexts [79].

*Computational Evaluation:* Computational evaluation uses computer algorithms to objectively assess persona quality through structured frameworks and NLP metrics. Around one-fifth (20.7%) of the articles used this form of persona evaluation. Multiple articles

**Table 2: Thematic coding for unique persona development approaches. LLMs bring four innovations to persona generation, including the ability to include different data sources, change the process of generating personas, prompting innovations, and innovations regarding interaction with personas. Articles are assigned to one primary sub-theme per main category and may appear in more than one main category.**

Main theme	Sub-theme	Papers (n)
<b>Input source innovations</b> (n = 25, 31.2%)	Visual/multimodal input	[32], [108], [120], [135], [149], [151] (n = 6)
	Alternative data sources	[24], [30], [33], [40], [57], [58], [101], [104], [105], [115], [124], [148] (n = 12)
	Expert-generated content	[11], [28], [45], [78], [99], [134], [148], [150] (n = 8)
<b>Process innovations</b> (n = 30, 37.5%)	Direct attribute specification	[77], [75], [76], [80], [85], [109], [115], [139], [151] (n = 9)
	Segmentation-free approaches	[107], [126], [134] (n = 3)
	Novel generation pipelines	[7], [9], [18], [24], [35], [55], [63], [83], [104], [105], [106], [108], [119], [129], [140], [142], [143], [149] (n = 18)
<b>Prompt engineering innovations</b> (n = 18, 22.5%)	Two-stage prompts	[35], [119], [128] (n = 3)
	Multistage prompts	[104], [129] (n = 2)
	Chain-of-thought methods	[108], [142] (n = 2)
	Few-shot learning	[7], [10], [18], [22], [27], [77], [75], [81], [86], [126], [143] (n = 11)
<b>Output format innovations</b> (n = 44, 55.0%)	Narrative-driven personas	[9], [17], [22], [33], [35], [86], [101], [105], [108], [107], [119], [126], [128], [132], [139], [142] (n = 16)
	Multimodal representations	[18], [32], [91], [108], [120], [140], [149], [151] (n = 8)
	Interactive persona systems	[10], [19], [21], [23], [28], [42], [45], [58], [60], [64], [76], [79], [81], [89], [96], [99], [106], [115], [134], [146], [150] (n = 21)

**Table 3: Target aspects of GenAI personas that are evaluated. Description and behavioral evaluations receive similar attention, while process and outcome evaluations are less common. See Table A7 in the online appendix for a complete list of articles.**

Evaluation Aspect	What Is Assessed	Example Articles
<b>Persona Description</b> (n = 17, 38.6%)	Static attributes: consistency, clarity, credibility, completeness, fluency, linguistic quality, stereotype presence	[22, 55, 128]
<b>Persona Behavior</b> (n = 16, 36.4%)	Dynamic performance: decision-making accuracy, response patterns, interaction quality, alignment with real human behavior	[60, 76, 108]
<b>Outcomes &amp; Impact</b> (n = 7, 15.9%)	Usability, user satisfaction, learning gains, effectiveness in design contexts	[57]
<b>Generation Process</b> (n = 5, 11.4%)	Efficiency, scalability, human-AI collaboration quality, workflow integration	[149]

applied computational linguistic tools, such as structured linguistic analysis, sentiment analysis, and stereotype lexicons, to investigate GenAI persona descriptions, for example, detecting potentially harmful stereotypes [22] or investigating how lexically diverse the persona descriptions were [128]. A lexical analysis across multiple LLMs (GPT, Gemini, Llama, Claude) found that personas generated by Gemini 1.5 Pro had the highest lexical diversity at the individual persona level, meaning each persona description employed a richer and more varied vocabulary to characterize user attributes [128].

Custom metrics represent another computational approach. For example, articles have developed custom evaluation dimensions that measure efficiency, satisfaction, collaboration, accuracy, creativity, and diversity in persona generation [149]. These measures are used for evaluation by the users and are collected via surveys and interviews. Articles [18] also use user perception scales like the PPS by asking an LLM to evaluate personas according to factors such as clarity, credibility, consistency, and empathy.

**Benchmark-Based Evaluation:** Benchmark-based evaluation was equally common as the use of computational evaluation (20.7% of articles). This approach compares personas with established standards, reference data, or existing frameworks to assess their accuracy and validity. For example, Gupta et al. [40] evaluated personas representing technology professionals against LinkedIn data, reporting an overrepresentation of elite universities (72.45% in generated personas vs. 8.56% in actual data). In another article, GenAI personas achieved high accuracy (89%) in simulating consumer preferences compared to actual survey data, with researchers contending that GenAI personas can accurately represent real consumer behaviors and preferences [76]. Some articles applied benchmarking by comparing LLM-generated personas with human ground truth data using controlled experimental conditions [27] or compared generated personas with theoretical archetypes like “Suggestible Sally” using trait alignment analysis [24]. Researchers have also applied pre- and post-intervention surveys to measure established beliefs, behavioral measures, and information sharing among real humans, simulated humans, and synthetic personas to understand the effectiveness of GenAI personas in representing human responses [60]. De Paoli [104] notes that conventional UCD methodologies and established persona development methodologies provide possible baselines for comparison.

**Human-driven evaluation prevails among the articles. However, a serious concern is that nearly half of all the articles provide no clear evaluation strategy for GenAI personas.**

**Table 4: Evaluation methodologies show that GenAI personas are mostly evaluated by humans, but other methods also include evaluating using algorithms and by comparison with established systems/data. See Table A8 in the online appendix for a complete list of articles.**

Main Theme	Sub-Theme	Example Articles
<b>Human-Driven Evaluation</b> (n = 25, 58.6%)	Subjective Rating Scales	[55, 57, 151]
	Expert Analysis	[132, 134]
	Discrimination Tests	[126]
	UX Assessment	[78, 79]
<b>Computational Evaluation</b> (n = 10, 20.7%)	NLP	[22, 128]
	Custom Metrics Frameworks	[18, 149]
<b>Benchmark-Based Evaluation</b> (n = 9, 20.7%)	Real-World Data Validation	[40, 76]
	Comparative Standard Analysis	[24, 104]
	Mixed-Methods Triangulation	[60, 119]

**4.2.3 LLMs’ Role in Evaluation.** We categorized the LLM use in persona evaluation according to four established validity concepts (Table 5): (1) face validity, (2) ecological validity, (3) internal validity,

and (4) statistical validity. Face validity assesses whether personas appear credible on the surface. Ecological validity examines how personas perform in real-world contexts (e.g., among stakeholders or when integrated into decision support systems). Internal validity evaluates the personas against some technical, predefined criteria (e.g., consistency, diversity), whereas statistical validity relies on statistical testing (estimation, population comparisons).

**Face Validity Assessment:** In a third of the articles (33.3%), researchers used LLMs to evaluate whether personas appear credible and realistic on their surface. For example, Li et al. [74] applied LLMs to assess LLM-generated personas compared to human-annotated examples. Researchers have also used LLMs to evaluate user personas using established questionnaires such as the PPS [103], which may provide some form of *face validity* but does not constitute a full assessment of the personas’ validity in real decision-making scenarios.

**Ecological Validity Assessment:** A quarter of the articles (25.0%) used LLMs to examine how personas behave in realistic, real-world contexts and scenarios. For example, Park et al. [108] tested persona-based decision-making in naturalistic scenarios, with GenAI systems assessing contextually appropriate responses. Researchers used LLMs to evaluate how personas influence guardrail sensitivity based on demographic context [75] and to analyze conversations based on persona attributes in realistic settings [135]. Smrke et al. [132] used LLMs to measure coherence of persona-driven interactions in healthcare contexts.

**Internal Validity Assessment:** A quarter of articles (25.0%) used some form of internal validity assessment, applying LLMs to evaluate the consistency and internal structure of persona characteristics. For example, researchers used LLMs to evaluate persona-driven responses using the Big Five personality factors to understand chatbot responses to personality tests [26]. Others used LLMs to interpret advertising content based on persona attributes for privacy auditing [17] or instructed LLMs to use standardized frameworks to evaluate persona authenticity and internal consistency [53]. Bai et al. [9] used LLMs to evaluate the emotional and cognitive characteristics of the generated personas.

**Statistical Validity Assessment:** Statistical testing was applied infrequently in conjunction with LLM-based persona evaluation (16.7% of LLM evaluations). Choi et al. used t-tests to compare persona quality scores between different methods [23]. Cross-model studies applied statistical tests to compare lexical diversity across multiple LLMs, including GPT-4, Gemini, Claude, and Llama [128]. Correlation tests checked whether different LLMs produced consistent persona attributes [73], whereas regression analyses can be used to test how persona characteristics affected content interpretation accuracy [17]. These statistical approaches remain less common than qualitative validity assessments, but they do provide quantitative rigor to persona evaluation.

Pertinent to all four validity types, our analysis reveals cases in which the same LLMs are used to both generate and evaluate personas. This usage of LLM raises questions about evaluation independence when the same technology both produces and judges outputs, though some articles mitigate this by using different LLMs for generation versus evaluation, with the idea that another LLM is more “independent” in its assessment of the personas.

**Table 5: LLM use in persona evaluation based on validity concepts. LLMs are mostly used to assess the face credibility of personas, but are also used to assess their usage in the real world, their internal consistency, and their validity against statistics. Articles can appear in more than one category.**

Assessment Type	Focus Area	Example articles
Face Validity (n = 4, 36.4%)	Surface Credibility	[27], [74]
	Apparent Realism	[103], [115]
Ecological Validity (n = 4, 36.4%)	Real-World Context	[35], [75]
	Naturalistic Behavior	[135], [144]
Internal Validity (n = 3, 27.3%)	Internal Consistency	[53], [148]
	Structural Integrity	[35]
Statistical Validity (n = 3, 27.3%)	Cross-Model Validation	[115]
	Systematic Analysis	[17], [135]

### 4.3 RQ3: What Ethical Considerations Are Associated with GenAI Personas?

**4.3.1 Ethical Engagement Patterns.** We find that more than half (56.8%, n = 46) of the articles explicitly discuss ethical considerations, while the rest (43.2%, n = 35) do not. This pattern indicates some awareness of ethical implications in GenAI personas, although many articles still lack explicit ethical reflection. In addition, the depth of the articles’ engagement with ethical issues varies considerably from one article to another. Some articles provide comprehensive ethical frameworks that systematically address multiple concerns. These comprehensive approaches examine bias detection mechanisms, community validation requirements, and potential societal harms from misrepresenting people [28, 134, 139]. These articles articulate specific methodologies to identify and mitigate ethical risks, such as involving community representatives in validation processes for vulnerable populations and implementing systematic bias detection protocols to identify harmful stereotypes prior to deployment.

Other articles acknowledge ethical considerations but provide a limited treatise of these issues. These articles typically mention concerns about bias, representation accuracy, or potential limitations without developing systematic frameworks for addressing them [40, 64, 86]. Some articles focus primarily on technical performance metrics, while giving secondary attention to the broader implications of GenAI-based user representation.

**4.3.2 Primary Ethical Concerns.** In the articles, the discussions on the ethical implications of GenAI personas are characterized by four themes (see Table 6): (1) bias and representational concerns, (2) potential societal harms from decision making, (3) trust and validity concerns, and (4) reduced human oversight.


**Bias and Representational Concerns:** Bias and other representational concerns are a long-lasting challenge in persona development and use [84, 138]. Our analysis reveals that researchers also express concerns about biases in GenAI personas (n = 17, 37.0%). Systematic

“algorithmic othering” particularly affects marginalized communities, where AI systems exhibit biases in representing gender and profession combinations [139]. GenAI personas are susceptible to caricature, tending to exaggerate traits of marginalized personas; for example, depicting individuals with physical disabilities primarily as being “in wheelchairs” [96].

Gender and age biases appear in multiple contexts, including autism representation, in which personas reinforce harmful stereotypes about gender differences [109], and financial representation in which personas do not accurately represent behaviors across age demographics [64]. For example, Gupta et al. [41] studied 24 reasoning datasets with 19 diverse personas using ChatGPT-3.5; they found 80% of personas showed stereotypical bias. Reliance on stereotypical associations limits the generation of diverse personality combinations, creating personas of predictable character types that may not reflect real user population complexity [86]. Platform-specific sampling biases also emerge when personas generated from specific platforms inherit the demographic biases of those data sources [148]. Although these persona concerns are not new (see, e.g., [138]), the introduction of AI into the persona process may exacerbate existing issues.

**Potential Societal Harms from Decision Making:** Some articles (n = 10, 21.7%) point out that GenAI personas could cause societal harm in several ways beyond individual misrepresentation. We list here some of these ways. First, “synthetic persuasion paradoxes” occur when AI personas systematically overestimate intervention effects compared to real human responses, potentially leading to misguided policy decisions that assume unrealistic behavior change outcomes [28]. Second, discrimination perpetuation occurs when personas embed and amplify existing societal biases, particularly affecting autism and other marginalized identities [109]. This occurs when personas are used in design processes, potentially leading to products and services that inadequately serve or actively exclude affected populations. Third, researchers note that harmful stereotyping risks emerge when generating personas for sensitive populations [134]. For example, without appropriate community validation and expert oversight, personas risk misrepresenting vulnerable groups in ways that could harm advocacy efforts and policy development. Policy misapplication represents another concern among researchers, in that biased persona data could provide overly optimistic estimates of intervention effectiveness, potentially leading to ineffective resource allocation [28].

**Trust and Validity Concerns:** Other articles (n = 9, 19.6%) express concern about the authenticity and reliability of the AI persona. We list the main concerns here. First, LLMs may yield reduced semantic diversity and more formulaic narratives compared to human-created personas [139], which raises questions about whether GenAI personas can capture the richness and complexity of actual human experiences and motivations. Second, preference misrepresentation [76] and issues with trust creation in user representation [64] are other concerns in this space. As a valuable consideration, Muller and Seaborn [93] refer to the use of synthetic data in persona generation leading to “Potemkin personas”; i.e., personas that are “staged” to give an appearance of reality merely to “please the eye” rather than truthfully representing reality. For example, in addition-focused persona research [119], GenAI consistently generated personas with 86% US-based representation



**Figure 7: Collaborative workflows between humans and GenAI for persona generation [129]: LLM-auto (fully automated), LLM-grouping (human researcher selects key characteristics), and LLM-summarizing with annotated sections indicating human involvement in attribute selection and clustering ([A]) and LLM utilization for data conversion ([B]). The LLM-summarizing model illustrates how AI and humans can “collaborate” on persona creation with humans identifying attributes and clustering while LLMs summarize data into persona narratives. The figure is annotated to explain different models of human-AI collaboration.**

despite no geographical constraints in prompts, and portrayed addiction narratives in an unrealistically positive light, mitigating the severity of substance abuse issues. Third, GenAI personas often lack emotional depth, which is a validity concern because emotions are an integral part of real user behavior and influence decision-making with personas [76]. Overall, validation gaps emerge from the absence of experimental validation with actual participants, limiting researchers’ ability to assess whether GenAI personas accurately represent target populations [109].

**Reduced Human Oversight** Our analysis shows that a major portion (42.0%, n = 34) of the articles reported no human-AI collaboration in persona generation, while 34.6% (n = 28) used such collaboration in a single stage (data collection, segmentation, enrichment, and evaluation) and 23.5% (n = 19) employed collaboration at multiple stages. Human involvement occurred primarily in content writing (28.4%) and generation (25.9%), with minimal participation in the evaluation (16.0%) or analysis phases. The typical workflow for GenAI persona development involves LLMs autonomously generating personas with no or limited human oversight. When humans are involved, they mainly provide feedback and input on prompt responses [104], then review and refine LLM outputs through feedback loops [151]. However, evaluation phases show reduced human involvement, despite evidence that community validation with representatives improves the representation of vulnerable populations [134] and cross-cultural validation prevents inappropriate persona transfer between contexts [64].

**4.3.3 Transparency and Reproducibility.** The field demonstrates notable advancement relative to traditional persona development practices in research openness and reproducibility. Most (61.7%, n = 50) of the articles show commitment to open science practices by sharing at least one type of resource (personas, data, prompts, and/or code). This represents progress compared to traditional persona research, where resource sharing was exceptional [116]. Among these, 16.0% (n = 13) demonstrate comprehensive resource

sharing by making all three resource types available, while 45.7% (n = 37) share one or two resource types. Slightly more than a third (38.3%, n = 31) of the articles share no resources at all.

Resource sharing shows encouraging patterns, with researchers sharing different types of resources. Personas represent the most commonly shared resource, with more than half (51.9%, n = 42) of the articles making their personas publicly available. This represents a substantial improvement over previous generations of persona research, where resource sharing was exceptional [116]. Code sharing shows a meaningful adoption, with nearly a third (28.4%, n = 23) of the articles providing implementation details. Dataset sharing shows a similar trend, with 27.2% of articles (n = 22) making their data accessible. Among the articles that share at least one resource, personas-only sharing represents the most common pattern (n = 20, 40%), followed by combined persona and code sharing (n = 10, 20%). Some articles demonstrate comprehensive resource sharing practices. Some articles provide persona descriptions and detailed prompts, enabling replication of their methodologies [119, 139], whereas others share code repositories to facilitate the adoption and extension of GenAI use in persona development [108, 129].

As a new type of reproducibility resource, prompt documentation is a particularly important aspect of GenAI persona research, which also represents a novel research area for the HCI community. According to our findings, researchers apply four main prompting strategies (see Table 7): (1) attribute specific instructions, (2) role-playing instructions, (3) data driven prompting, (4) creative persona development. Out of these, *attribute-specific prompting* provides the highest degree of control by specifying exact demographic or behavioral characteristics through structured formats [64, 119], and or controlled variation of specific attributes [96]. In turn, *role-playing instructions* take a more contextual approach, instructing the LLM to embody particular user types with specific emotional states, communication styles [10, 134], or personality traits [77, 151]. Furthermore, *data-driven prompting* grounds persona generation in actual user data, either by directly integrating survey responses, analytics, or behavioral logs into the prompts [55, 129], or by using LLMs to assess which synthetic responses best match real user profiles [4, 27]. Finally, *creative persona development* emphasizes psychological depth and narrative richness [82, 108], often employing multimodal approaches that combine text with images or videos for persona representation [28, 108].

It is worth noting that these strategies are not mutually exclusive, but researchers often combine multiple approaches, such as using data-driven prompts to establish demographic baselines while employing role-playing instructions to add behavioral details [22]. The choice of strategy depends on the research goals; for example, attribute-specific prompting applies to fully automated pipelines for persona generation, whereas role-playing is typically used to elicit richer behavioral dynamics, data-driven approaches are favored when researchers aim to anchor personas in empirical user profiles, and narrative-oriented personas are often applied in creative development contexts.

**Table 6: Thematic analysis of ethical considerations in GenAI personas. The most common ethical concerns considered and discussed in GenAI persona literature include bias and representation, while concerns related to a lack of trust, potential harms in decision-making, and reduced human oversight are also present. See Table A9 in the online appendix for complete list of articles.**

Main Ethical Theme	Specific Ethical Concerns	Example Articles
<b>Bias and Representational Concerns (n = 17, 37.0%)</b>	Algorithmic Othering	[139]
	Elite Institution Overrepresentation	[40]
	Gender and Age Bias	[64, 109]
	Stereotypical Associations	[86]
	Platform-Specific Bias	[148]
	AI Exacerbation of Existing Issues	[109, 139]
<b>Potential Societal Harms from Decision Making (n = 10, 21.7%)</b>	Synthetic Persuasion Paradox	[28]
	Discrimination Perpetuation	[109]
	Harmful Stereotyping of Vulnerable Groups	[134]
	Policy Misapplication	[28]
<b>Trust and Validity Concerns (n = 9, 19.6%)</b>	Reduced Semantic Diversity	[139]
	Consumer Preference Misrepresentation	[64, 76]
	Limited Emotional Range	[76]
	Validation Gaps	[109]
<b>Reduced Human Oversight (n = 10, 21.7%)</b>	Limited Human-AI Collaboration	[104, 151]
	Automated Generation Workflow	[104, 151]
	Community Validation Requirements	[134, 139]
	Cross-Cultural Validation Gaps	[64]
	Evaluation Phase Neglect	[104, 151]

## 5 Discussion

### 5.1 Research Implications

To help the reader make sense of our findings better, we apply three categories: the Good, the Bad, and the Ugly (see Figure 8 for elaboration). *Good* refers to beneficial advances that validate the value of GenAI personas and point toward productive applications. *Bad* refers to methodological weaknesses that undermine validity and reliability. *Ugly* refers to systemic or structural risks that can cause harm, particularly ethical violations stemming from over-reliance on automation without human oversight.

*The Good—Expanding Accessibility, Expressiveness, and Opportunities:* It can be argued that GenAI has democratized persona development in numerous ways, such as increasing persona generation speed and enabling researchers without specialized UX training to create personas. GenAI persona research is also more likely to share resources than previous eras of persona research, which aligns with broader movements towards open science practices and ideals like increased transparency and reproducibility. Moreover, GenAI enables new interaction formats, such as conversational personas that allow designers to interact with personas dynamically rather than read static profiles. There are also new, *iterative co-creation* models where humans and machines collaborate through continuous prompt refinement. Moreover, GenAI enables

*agentic personas* that simulate user responses, which challenges the traditional distinction of personas primarily describing user segments by broadening the scope of personas into predicting user behaviors. While traditional personas serve as reference documents that designers consulted to understand user needs [25, 112], agentic personas generate responses representing user views in real time. This raises an intriguing question, which is: *Is a persona legitimate because it accurately describes real users based on data, or because it can accurately predict what users would do?* This is an exciting and potentially impactful area for HCI research.

*The Bad—The Accessibility-Quality Paradox:* However, there is also what we term the *accessibility-quality paradox*: *even though GenAI can democratize persona creation, its precarious use can severely undermine methodological rigor.* GenAI technology lowers the skill barriers that prevent untrained creators from creating personas, but evaluation still requires expertise in validity assessment and bias detection [16, 121]. When anyone can generate personas in minutes, but proper evaluation requires significant expertise, persona developers are incentivized to treat evaluation light-heartedly. This can explain why generation techniques exhibit a degree of standardization, but evaluation remains more inconsistent among studies. The problem mirrors longstanding concerns that personas lack empirical validation [52], but GenAI can be seen to erode some

**Table 7: LLM prompting strategies for persona development. LLM prompting strategies for persona generation include: (1) attribute specific prompting, where specific demographic or behavioral attributes are used to change, (2) role play based prompting, where the LLM is assigned a role, (3) data driven prompting, where the data related to the users is provided to generate personas, and (4) creative prompting which involves different modalities.**

Primary prompting strategy	Specific technique	Example implementation
Attribute-specific prompting	Structured format instructions	“Provide the output in a JSON array, with each dict containing only the following keys: ‘index’, ‘name’, ‘age’, ‘occupation’, ‘background’, ‘details’” [119] “Uses decision tree methodology to create personas with 2-8 attributes based on survey characteristics” [64]
	Controlled attribute variation	“Given a fictional user’s description: {base persona} Modify the description to make it more like a {privacy attribute} version of the person” [96] “Generate persona variants by varying selected attributes such as age and income while maintaining core demographic consistency” [96]
Role-playing instruction	Context-specific simulation	“You are a user of a university intranet... act as a young person who uses colloquial but polite language, making up a name, a gender, and an age within the range of 21 and 26 years” [10] “Uses LLM (GPT-4o mini) with RAG framework and ability-based approach for personas of people with complex needs” [134]
	Emotional and behavioral specification	“In a conversation, you should act as if you have the following eight emotional parameters. Each parameter should fluctuate throughout the dialog...” [151] “Uses Big Five personality traits and domain-specific profiles to create diverse novice instructor personas” [77]
Data-driven prompting	User data integration	“Here is the user data... Generate a minimum number of personas to represent the user data. Rule 1: Do not add any information that does not exist in the user data...” [129] “PersonaCraft uses LLMs integrated with clustering analysis in a 5-stage methodology including data preparation, question mapping, pre-processing, persona generation, and enrichment phases” [55]
	Comparative assessment	“Given the user profile provided below, select the response from AI assistant A or B that the user would most likely prefer...” [27] “LLMs generate synthetic survey responses by conditioning on demographic personas to create digital twins of real participants” [4]
Creative persona development	Psychological Depth Elicitation	“What is your dark secret?... What is your personality like? Personality is one of the following five traits: extraversion, agreeableness, conscientiousness, neuroticism, or openness to experience...” [108] “Generated 1000 backstories using GPT-4o-mini, selected 100 balanced personas with diverse demographic and professional profiles for XAI evaluation” [82]
	Multimodal persona development	“DALL-E 3 prompt: A character from random genre of manga... Look carefully this image, and give me your imagination of detailed description of appearance of the character...” [108] “Automated pipeline creating synthetic profiles with demographic and behavioral attributes including names, education, careers, locations, hobbies, and climate stances” [28]

of the quality gatekeeping that technical barriers previously provided. In GenAI persona research, we observe hints of this pattern. Nearly half (44%) of articles lack explicit evaluation frameworks, weakening the empirical foundation of persona research. The evaluation challenge is exacerbated by the fact that only 11.5% of the articles explicitly address hallucination detection, even though this risk is common knowledge [47]. In general, the evaluation of GenAI personas varies considerably in its degree of human involvement. Human-in-the-loop approaches employ people to assess personas at specific stages [55]. In contrast, automated evaluation pipelines integrate computational assessments directly into the generation workflow, applying NLP metrics, stereotype detection algorithms, or LLM-based judges without human intervention, thereby reducing transparency [22, 128].

A particular risk occurs when the same LLM generates and evaluates personas, risking that the evaluation is based on model-specific patterns rather than the genuine quality of the personas. Unlike persona generation, which has converged on common patterns such

as prompt engineering and data grounding, evaluation practices for GenAI personas are currently *ad hoc*. The field lacks consensus on when evaluation should occur, who should conduct the evaluation, and what constitutes adequate validation. Also, most research (86.4%) depends on a single provider (OpenAI’s GPT models), which limits diversity and embeds one company’s biases in persona development, such as embedding Western perspectives into global persona representations [8]. This concentration exemplifies over-reliance on automated systems at the expense of critical evaluation [131]. The reliance on a single provider may lead to uncritical acceptance of GenAI personas, particularly when practitioners lack the technical expertise to evaluate LLM limitations. These problems could be addressed through standardized evaluation protocols, diversified LLM use, and mechanisms to detect hallucinations.

*The Ugly—Loss of Human Agency and Participatory Principles:* Some developments threaten the foundation of persona research itself by systematically reducing human involvement. The patterns we observe in GenAI persona research are problematic from the

perspective of participatory design. Nearly half (42.0%) of the articles do not report an explicit human-AI collaboration model, instead fully delegating the persona development process to GenAI. About a third (30.9%) of the articles adopt participatory methods. This creates a risk of computational efficiency substituting genuine engagement, producing personas that lack authentic user voices. Our research shows evidence of each step toward fully automated pipelines. We found articles that used synthetic data for persona generation, articles that used LLMs only to generate personas, but others that also evaluated personas using LLM, and some articles that presented complete persona systems that produce autonomous decisions without human intervention.

In the worst scenario, one could perceive a dystopian outcome in which personas are created from synthetic data by machines that are then evaluated by other machines, and possibly delivered to yet other machines making decisions about real people. This scenario would eradicate humans from the design process and goes against the human-in-the-loop principles that argue for maintaining human agency and oversight in GenAI systems, which is essential when making decisions affecting people [92]. The fully automated persona pipeline could represent precisely what human-in-the-loop frameworks seek to prevent: algorithmic systems making consequential decisions about human populations without meaningful human involvement or accountability mechanisms. Fully automated pipelines break what makes personas valuable. Personas work in user-centered design because they represent real users' needs and perspectives [112, 125]. But the pattern we observed creates a closed loop: synthetic data feeds AI generation, AI evaluates its own outputs, and AI makes decisions about real people. No actual users are involved. This transforms personas from user representations into algorithmic artifacts with no connection to the users they claim to represent. Such systems violate basic participatory design principles that require involving those affected by design decisions [92, 125].

Simultaneously, the increased popularization of GenAI personas reflects changing power dynamics in who (or what) controls user representation. Whereas personas initially emerged from direct user research and community engagement, GenAI personas may create computationally generated narratives that mimic authentic UX [62]. This mimicry raises epistemic questions about the trade-off between efficiency gains and representational authenticity [93], which are questions that the HCI community is required to address. These findings align with concerns in human-AI collaboration research about automation bias, where over-reliance on AI systems leads to uncritical acceptance of outputs even when human evaluation could detect errors [131]. The patterns we observe contradict human-centered AI principles that emphasize maintaining human agency and meaningful oversight in systems that affect real populations [92, 130].

Recognizing these differences can help the HCI community embrace beneficial applications of GenAI in personas while avoiding practices that would disconnect personas from the users they claim to represent. The HCI community can make a substantial contribution in this area by ensuring, whenever possible, to engage real users and maintain the participatory foundations that give personas their validity and value.

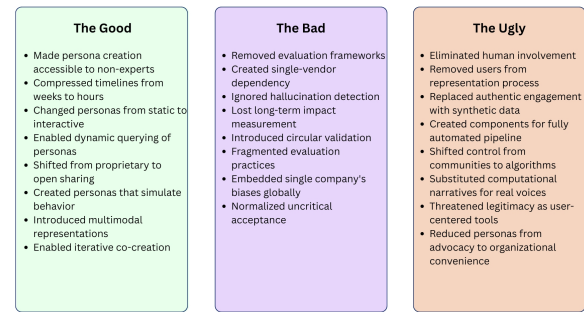


Figure 8: The good, the bad, and the ugly dimensions of the GenAI Personas.

## 5.2 Extension of Prior Persona Research

Here, we position GenAI personas in the broader HCI literature on personas. We do so by first discussing the parallels between GenAI persona research and HCI research evolution, and then focusing on the continuum of data-driven personas, specifically how GenAI personas are positioned within it.

Go et al. [34] investigated the “three waves of personas” by reviewing 315 articles that span 18 years of research tradition. There are interesting parallels in this most recent transition of persona technology in regards to these waves. In Bødker’s terms [13, 14], there was a *transition from human factor to human actor* in HCI; however, for persona creation, the adoption of LLMs seems to decrease human agency in the persona generation process. Bødker also identified *engagement of other disciplines in HCI research*; this is evident in personas becoming increasingly a nexus of HCI-NLP collaboration, rather than being “owned and operated” by HCI alone. This idea has previously been articulated by Duda [29], who questioned whether personas are conceptually claimed by HCI or the field of marketing; in the GenAI persona era, this question particularly relates to the persona concept being appropriated by NLP researchers. NLP research tends to have a technical tradition but less knowledge about the conceptual heritage of personas in HCI, which implies that HCI researchers’ knowledge about persona theory could contribute to better operationalization of “NLP-based” personas (e.g., increasing persona depth and arguing for grounding on real user experiences).

An important aspect of Bødker’s [13, 14] articulation is the role of *participatory design*, in which end-users are actively involved as co-designers throughout the design process. Our findings show that this is relatively rare in GenAI personas; however, it need not be rare. The democratization effect of GenAI technologies in making persona generation easier could, in theory, facilitate and support the inclusion of persona users in the generation process. Mobilizing this “participatory potential” remains a highly impactful avenue for future research. It is not a novel challenge; already in their 2017 review, Goh et al. [34] identified a “a missing link between personas and designers” (p. 612).

Similarly, it is worth noting that these challenges, such as stereotyping, marginalization, and personas being disconnected from real

user data are well-known in the persona field [34, 84, 138] but it is nonetheless interesting to observe that they remain topical for GenAI personas. Qualitative user data and participatory validation, followed by persona refinement, are foundational methods used in persona creation. These approaches emphasize methodological transparency and stakeholder engagement [25, 113]. As such, validation has always been central to personas, with UCD guidelines highlighting the need for empirical grounding, data triangulation, and participatory checks with both designers and users [5, 87, 112]. Our findings indicate that GenAI personas frequently lack validation, with many of the studies we examined reporting no human supervision or explicit evaluation frameworks. Thus, GenAI amplifies concerns about persona validity by accelerating production but eroding safeguards. We refer the reader to the article by Amin et al. [2] that lists twenty challenges of GenAI personas in HCI. The resolution of these challenges remains an ongoing research effort.

In their literature review spanning 15 years of research on data-driven personas, Salminen et al. [117] argued for three evolutionary periods: (1) *quantification* (2005-2008), which involved early experimentation with statistical and ML technologies in persona creation; (2) *diversification* (2009-2014), which involved expanding the scope of algorithms and data sources to involve NLP and both structured and unstructured data; and (3) *digitalization* (2015-2020) that marked the rise of fully automated persona generation in user-friendly systems. The review also specifically pointed out challenges in lacking shared resources, standardized evaluation methods, the risk of not generating inclusive personas, and the over-focus on algorithms at the expense of human experience.

These risks do remain relevant for GenAI personas, but there is also progress, especially in the sharing of computational resources, using quantitative evaluation, and the researchers becoming more aware of issues such as the need for representing marginalized users and the role of algorithmic bias in persona generation. So, it appears that personas are “moving forward”, albeit slowly. It is challenging to propose an all-encompassing label for the current era of persona generation, but perhaps (4) *augmentation* could be descriptive, with the undertone that GenAI technologies exist to augment human involvement in persona generation rather than replacing it.

### 5.3 Practical Implications

Based on our analysis of 81 articles, we propose practical guidelines for the development of GenAI personas (Table 8). Each guideline addresses specific methodological gaps identified in our review and is grounded in relevant research from our scoping review and previous literature. Each guideline in Table 8 addresses a specific gap from our review, labeled as Review Finding, and builds on prior research, labeled as Prior Work, with the Reasoning column explaining how the cited work supports the guideline. For example, PG1 corresponds to our finding that 86% of articles rely exclusively on GPT models (Section 4.1.1), while prior work demonstrates that model-specific biases exist that single-model approaches might miss.

These guidelines address three key ethical concerns about using GenAI personas in HCI and design. First, multi-model validation and consistency testing address systematic biases, inconsistencies,

and harmful stereotypes that could compromise research validity and perpetuate discrimination. These validation approaches identify systematic issues while maintaining stable characteristics across iterations, addressing concerns about reproducibility in persona development. Second, empirical grounding through structured outputs, demographic validation, and quantitative thresholds ensures that GenAI personas remain anchored in real user data rather than generating fictional or biased representations based solely on training patterns or artifacts. For example, requiring personas to match validated demographic distributions prevents models from creating representations that diverge from actual user populations.

Third, documentation requirements and human oversight protocols provide ethical safeguards against cultural insensitivity or stereotyping, because GenAI personas may lack an intricate understanding of cultural contexts that human researchers possess. Overall, these guidelines aim at leveraging GenAI’s efficiency in persona generation while maintaining the methodological rigor and ethical standards expected in HCI research, so that GenAI personas offer valid, reliable, and respectful representations of diverse user populations. Moreover, we expect that human-AI collaboration models will become increasingly important with the rise of agentic personas that are capable of generating their own outputs independently.

### 5.4 Limitations and Future Research

Our study has some limitations. The on-going formation of practices in the GenAI persona research means that our analysis represents a snapshot of a rapidly changing field. The dominance of recent publications (88.9% from 2024-2025) reflects the field’s nascency, but limits assessment of long-term trends. However, scoping reviews are valuable in such cases, and we do share our resources (search query, coding book, corpus) so that it can be updated in future.

One limitation is that our HCI-focused perspective may not capture relevant developments in adjacent fields such as NLP or computational social science, as we excluded articles that use “persona” to configure LLM personalities or agent behaviors [61, 68, 73, 136]. Although outside of our scope, these agent-focused persona studies may offer transferable insights into prompt engineering techniques, personality consistency methods, or behavioral alignment approaches. Future research could examine the synergies between the HCI and NLP persona research traditions.

In addition to our research limitations, the reviewed articles indicate recurring challenges that constrain the field’s development and can guide future work. For example, token processing constraints appear frequently in complex persona generation [104], limiting the depth and complexity of GenAI personas due to technical limitations in current LLM architectures. *Bias and representation* issues are prevalent in researchers’ treatise of GenAI personas. For example, Kaur et al. [64] documented systematic biases that affect cross-cultural applicability, while Mercer et al. [86] identified a heavy reliance on stereotypical associations that constrain persona diversity. Yin et al. [148] found platform dependency issues that introduce systematic sampling biases based on data source characteristics. *Validation challenges* include small sample sizes that limit generalizability [139], and the absence of experimental

**Table 8: Practical guidelines for GenAI persona development.**

Guideline	Review Finding	Prior Work	Reasoning
<b>PG1: Implement multi-model validation using at least two different GenAI systems</b>	See Section 4.1.1	[22, 41, 67]	Different models exhibit distinct biases that cross-validation can identify.
<b>PG2: Deploy structured formats (JSON, XML, or CSV) with explicit consistency criteria</b>	See Section 4.1.3	[55, 119]	Structured formats enable programmatic validation and cross-system compatibility.
<b>PG3: Document complete methodologies including prompts, parameters, and model versions</b>	See Section 4.3.3	[67, 111]	Prompt sensitivity and parameter choices significantly affect outputs and reproducibility.
<b>PG4: Validate generated characteristics against demographic benchmarks or real user data</b>	See Section 4.2.2	[8, 72]	LLM-generated personas show systematic deviations from real-world population distributions.
<b>PG5: Establish quantitative acceptance thresholds before incorporating personas into design</b>	See Section 4.2.2	[55, 121]	Pre-registered criteria prevent confirmation bias and improve persona quality.
<b>PG6: Assess output consistency through repeated generation and edge case testing</b>	See Section 4.3.2	[40, 57, 109]	High variability indicates unstable representations and edge cases reveal model limitations.
<b>PG7: Implement systematic human oversight for cultural appropriateness and stereotype detection</b>	See Section 4.3.2	[84, 138, 139]	Automated methods miss cultural details that require human domain expertise.

validation with actual participants [109]. *Cultural limitations* represent another significant constraint. Studies confined to specific regions limit cross-cultural generalizability, raising questions about whether GenAI persona methods in one cultural context can be appropriately transferred to others.

Future research directions should address the limitations identified above, offering significant research fronts for the HCI community. First work on *bias mitigation* is needed because current articles reveal persistent issues related to representational bias, stereotyping, and platform-dependent sampling. For example, systematic bias-reduction strategies [28] and community-based validation practices [139] offer promising starting points for improving cross-cultural reliability. Second, technical improvements are required to overcome LLM design constraints that limit persona depth and generalizability. For example, multi-platform integration [148] can reduce single-platform dependence, multimodal pipelines [82] can support more accessible and usable persona representations, and domain-specific fine-tuning [76] could help improve the validity of personas in specialized contexts. Third, the validity and long-term utility of GenAI personas should be studied. For example, longitudinal study designs [134] can assess the stability of personas over time, stereotype-detection mechanisms [35] can help monitor problematic persona narratives, and cross-cultural validation frameworks [64] are needed to determine whether GenAI technologies and techniques developed in one region transfer to others.

## 6 Conclusion

GenAI technology is increasingly adopted for persona generation, but the field remains fragmented and methodologically inconsistent.

The articles we examined demonstrate progress in prompt engineering, multimodal generation, resource sharing, and domain-specific adaptation, but they also reveal recurring limitations related to bias, validation, and cultural transferability. Our findings suggest that the main challenge for HCI is not producing more complex and automated GenAI persona pipelines, but validating that GenAI personas are reliable, interpretable, involve humans in development, and appropriate for real use. Overall, the field can benefit from addressing the specific risks and limitations identified in this review. Through such efforts, GenAI personas can mature from a technological novelty into a robust, and socially beneficial research approach.

## References

- [1] Tamara Adlin, Holly Jamesen, and Tony Krebs. 2001. Fake people and sticky notes: Fostering communication for human-centered software design. White Paper. <https://starstoroad.com/blog/wp-content/uploads/2012/09/FakePeopleAndStickyNotes.pdf> Online Access.
- [2] Danial Amin, Joni Salminen, Bernard J. Jansen, Joongi Shin, and Dae Hyun Kim. 2025. Generative AI personas considered harmful? Putting forth twenty challenges of algorithmic user representation in human-computer interaction. *International Journal of Human-Computer Studies* 205 (2025), 103657. doi:10.1016/j.ijhcs.2025.103657
- [3] Mikio Aoyama. 2005. Persona-and-scenario based requirements engineering for software embedded in digital consumer products. In *13th IEEE International Conference on Requirements Engineering (RE'05)*. IEEE, USA, 85–94. doi:10.1109/RE.2005.50
- [4] Lisa P. Argyle, Ethan C. Busby, Nancy Fulda, Joshua R. Gubler, Christopher Rytting, and David Wingate. 2023. Out of One, Many: Using Language Models to Simulate Human Samples. *Political Analysis* 31, 3 (July 2023), 337–351. doi:10.1017/pan.2023.2
- [5] Leena Arhippainen. 2013. A tutorial of ten user experience heuristics. In *Proceedings of International Conference on Making Sense of Converging Media*. Association for Computing Machinery, New York, NY, USA, 336–337. doi:10.1145/2523429.2523491

- [6] Hilary Arksey and Lisa O'Malley. 2005. Scoping studies: towards a methodological framework. *International Journal of Social Research Methodology* 8, 1 (2005), 19–32. doi:10.1080/1364557032000119616
- [7] Neeraj Arora, Ishita Chakraborty, and Yohei Nishimura. 2025. AI-human hybrids for marketing research: Leveraging large language models (LLMs) as collaborators. *Journal of Marketing* 89, 2 (2025), 43–70. doi:10.1177/00222429241276529
- [8] Mohammad Atari, Mona J. Xue, Peter S. Park, Damián Blasi, and Joseph Henrich. 2023. Which humans? PsyArXiv. doi:10.31234/osf.io/5b26t Preprint.
- [9] Yuqi Bai, Kun Sun, and Huishi Yin. 2024. Agentic Society: Merging skeleton from real world and texture from Large Language Model. arXiv. doi:10.48550/arXiv.2409.10550 Preprint.
- [10] Jose Barambones, Cristian Moral, Angélica De Antonio, Ricardo Imbert, Loïc Martínez-Normand, and Elena Villalba-Mora. 2024. ChatGPT for Learning HCI Techniques: A Case Study on Interviews for Personas. *IEEE Transactions on Learning Technologies* 17 (2024), 1486–1501. doi:10.1109/TLT.2024.3386095
- [11] Karim Benharrak, Tim Zindulka, Florian Lehmann, Hendrik Heuer, and Daniel Buschek. 2024. Writer-Defined AI Personas for On-Demand Feedback Generation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–18. doi:10.1145/3613904.3642406
- [12] Åsa Blomquist and Mattias Arvola. 2002. Personas in action: ethnography in an interaction design team. In *Proceedings of the second Nordic conference on Human-computer interaction*. Association for Computing Machinery, New York, NY, USA, 197–200. doi:10.1145/572020.572044
- [13] Susanne Bødker. 2006. When Second Wave HCI Meets Third Wave Challenges. In *Proceedings of the 4th Nordic Conference on Human-Computer Interaction: Changing Roles*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/1182475.1182476
- [14] Susanne Bødker. 2015. Third-Wave HCI, 10 Years Later—Participation and Sharing. *Interactions* 22, 5 (2015), 24–31. doi:10.1145/2804405
- [15] Robert Bowman, Camille Nadal, Kellie Morrissey, Anja Thieme, and Gavin Doherty. 2023. Using thematic analysis in healthcare HCI at CHI: A scoping review. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3544548.3581203
- [16] Chris Chapman and Russell P. Milham. 2006. The Personas' New Clothes: Methodological and Practical Arguments against a Popular Method. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 50 (October 2006), 634–636. doi:10.1177/154193120605000503
- [17] Chaoran Chen, Leyang Li, Luke Cao, Yanfang Ye, Tianshi Li, Yaxing Yao, and Toby Jia-Jun Li. 2025. Why am I seeing this: Democratizing End User Auditing for Online Content Recommendations. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology (UIST '25)*. Association for Computing Machinery, New York, NY, USA, Article 154, 23 pages. doi:10.1145/3746059.3747798
- [18] Chaoran Chen, Weijun Li, Wenxin Song, Yanfang Ye, Yaxing Yao, and Toby Jia-Jun Li. 2024. An Empathy-Based Sandbox Approach to Bridge the Privacy Gap among Attitudes, Goals, Knowledge, and Behaviors. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–28. doi:10.1145/3613904.3642363
- [19] Luyu Chen, Quanyu Dai, Zeyu Zhang, Xueyang Feng, Mingyu Zhang, Pengcheng Tang, Xu Chen, Yue Zhu, and Zhenhua Dong. 2025. RecUserSim: A Realistic and Diverse User Simulator for Evaluating Conversational Recommender Systems. In *Companion Proceedings of the ACM on Web Conference 2025 (Sydney NSW, Australia) (WWW '25)*. Association for Computing Machinery, New York, NY, USA, 133–142. doi:10.1145/3701716.3715258
- [20] Robert C. C. Chen, Wen Cing-Yan Nivala, and Chien-Bang Chen. 2011. Modeling the role of empathic design engaged personas: an emotional design approach. In *Proceedings of the 6th International Conference on Universal Access in Human-Computer Interaction: Users Diversity - Volume Part II (Orlando, FL) (UAHCI'11)*. Springer-Verlag, Berlin, Heidelberg, 22–31. doi:10.1007/978-3-642-21663-3\_3
- [21] Si Chen, Izzy Molnar, Ting Hua, Peiyu Li, Le Huy Khiem, G Alex Ambrose, Jim Lang, Ronald Metoyer, and Nitesh V Chawla. 2025. SimInstruct: A Responsible Tool for Collecting Scaffolding Dialogues Between Experts and LLM-Simulated Novices. doi:10.48550/arXiv.2508.04428 Preprint.
- [22] Myra Cheng, Esin Durmus, and Dan Jurafsky. 2023. Marked Personas: Using Natural Language Prompts to Measure Stereotypes in Language Models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 1504–1532. doi:10.18653/v1/2023.acl-long.84
- [23] Yoonseo Choi, Eun Jeong Kang, Seulgi Choi, Min Kyung Lee, and Juho Kim. 2025. Proxona: Supporting Creators' Sensemaking and Ideation with LLM-Powered Audience Personas. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–32. doi:10.1145/3706598.3714034
- [24] Dylan Clements, Elysia Giannini, Fi Crowe, Mike Balapitiya, Jason Marshall, Paul Papadopoulos, and Tanjila Kanij. 2023. An Innovative Approach to Develop Persona from Application Reviews. In *Proceedings of the 18th International Conference on Evaluation of Novel Approaches to Software Engineering - ENASE*. INSTICC, SciTePress, Portugal, 701–708. doi:10.5220/001199600003464
- [25] Alan Cooper. 1999. *The Inmates are Running the Asylum*. Vieweg+Teubner Verlag, Wiesbaden, 17–17. doi:10.1007/978-3-322-99786-9\_1
- [26] Joost C.F. de Winter, Tom Driessen, and Dimitra Dodou. 2024. The use of ChatGPT for personality research: Administering questionnaires using generated personas. *Personality and Individual Differences* 228 (2024), 112729. doi:10.1016/j.paid.2024.112729
- [27] Yijiang River Dong, Tiancheng Hu, and Nigel Collier. 2024. Can LLM be a Personalized Judge?. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 10126–10141. doi:10.18653/v1/2024.findings-emnlp.592
- [28] Alexander Doudkin, Pat Pataranutaporn, and Pattie Maes. 2025. From Synthetic to Human: The Gap Between AI-Predicted and Actual Pro-Environmental Behavior Change After Chatbot Persuasion. In *Proceedings of the 7th ACM Conference on Conversational User Interfaces*. ACM, Waterloo ON Canada, 1–18. doi:10.1145/3719160.3736610
- [29] Sabrina Duda. 2018. *Personas—Who Owns Them*. Springer Fachmedien Wiesbaden, Wiesbaden, 173–191. doi:10.1007/978-3-658-21450-0\_8
- [30] Aleksandr Farseev, Qi Yang, Marlo Ongpin, Ilija Gossoudarev, Yu-Yi Chu-Farseeva, and Sergey Nikolenko. 2024. SOMONITOR: Combining Explainable AI & Large Language Models for Marketing Analytics. doi:10.48550/arXiv.2407.13117 Preprint.
- [31] Kristin Skeide Fuglerud, Trenton Schulz, Astri Letnes Janson, and Anne Moen. 2020. Co-creating persona scenarios with diverse users enriching inclusive design. In *International Conference on Human-Computer Interaction*. Springer-Verlag, Berlin, Heidelberg, 48–59. [https://link.springer.com/chapter/10.1007/978-3-030-49282-3\\_4](https://link.springer.com/chapter/10.1007/978-3-030-49282-3_4)
- [32] Takayuki Fujimoto. 2025. Ghost hack: a distributed control architecture for synthetic personality in human-AI interaction. *AI & SOCIETY* 28, 3 (July 2025), 1–17. doi:10.1007/s00146-025-02483-8
- [33] Toshiaki Goel, Orit Shaer, Catherine Delcourt, Quan Gu, and Angel Cooper. 2023. Preparing Future Designers for Human-AI Collaboration in Persona Creation. In *Proceedings of the 2nd Annual Meeting of the Symposium on Human-Computer Interaction for Work (Oldenburg, Germany) (CHIWORK '23)*. Association for Computing Machinery, New York, NY, USA, Article 4, 14 pages. doi:10.1145/3596671.3598574
- [34] Chu Hiang Goh, Narayanan Kulathuramaiyer, and Tariq Zaman. 2017. Riding Waves of Change: A Review of Personas Research Landscape Based on the Three Waves of HCI. In *Information and Communication Technologies for Development*. Springer International Publishing, Cham, 605–616. doi:10.1007/978-3-319-59111-7\_49
- [35] Rafael Arias Gonzalez and Steve DiPaola. 2024. Exploring Augmentation and Cognitive Strategies for AI based Synthetic Personae. arXiv. doi:10.48550/arXiv.2404.10890 Preprint.
- [36] Arantxa Gonzalez de Heredia, Joy Goodman-Deane, Sam Waller, P. John Clarkson, Daniel Justel, Ion Iriarte, and Jesús Hernández. 2018. Personas for policymaking and healthcare design. In *Proceedings of the International Design Conference (DESIGN 2018)*. Faculty of Mechanical Engineering and Naval Architecture, Dubrovnik, Croatia, 2209–2220. doi:10.21278/idc.2018.0545
- [37] 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*. Design Research Society, London, England, 1828–1840. doi:10.21606/drs.2018.400
- [38] Joy Ai-Leen Goodman-Deane, Mike Bradley, Sam Waller, and P. John Clarkson. 2021. Developing personas to help designers to understand digital exclusion. In *Proceedings of the Design Society*, Vol. 1. Cambridge University, Cambridge, UK, 1203–1212. doi:10.1017/pds.2021.120
- [39] 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* 40, 19 (September 2023), 5569–5584. doi:10.1080/10447318.2023.2247568
- [40] Shailja Gupta and Rajesh Ranjan. 2024. Evaluation of LLMS biases towards elite universities: a persona-based exploration. arXiv Preprint. doi:10.48550/arXiv.2407.12801 Preprint.
- [41] Shashank Gupta, Vaishnavi Shrivastava, Ameet Deshpande, Ashwin Kalyan, Peter Clark, Ashish Sabharwal, and Tushar Khot. 2024. Bias Runs Deep: Implicit Reasoning Biases in Persona-Assigned LLMs. In *Proceedings of an International Conference on Learning Representations 2024*. ICLR, Appleton, WI, 14 pages. doi:10.48550/arXiv.2311.04892
- [42] Xiangyang He, Jiale Li, Jiahao Chen, Yang Yang, and Mingming Fan. 2025. SimuPanel: A Novel Immersive Multi-Agent System to Simulate Interactive Expert Panel Discussion. arXiv. doi:10.48550/arXiv.2506.16010 Preprint.
- [43] Niels Hendriks, Frederik Truyen, and Erik Duval. 2013. Designing with Dementia: Guidelines for Participatory Design together with Persons with Dementia. In *Human-Computer Interaction - INTERACT 2013*. Springer Berlin Heidelberg, Berlin, Heidelberg, 649–666. doi:10.1007/978-3-642-40483-2\_46

- [44] Dan Högberg, Lars Hanson, Daniel Lundström, Maria Jönsson, and Dan Lämkkull. 2008. Representing the elderly in digital human modeling. In *Proceedings of the 40th Annual Nordic Ergonomic Society Conference, Reykjavik, Iceland*. Nordic Ergonomics and Human Factors Society, Sweden, 6. [https://www.researchgate.net/publication/247778784\\_Representing\\_the\\_elderly\\_in\\_digital\\_human\\_modelling](https://www.researchgate.net/publication/247778784_Representing_the_elderly_in_digital_human_modelling)
- [45] Hsiao Chin Huang, Hong Lei, and Chun-Cheng Hsu. 2024. Exploring Designer-Generative AI Collaborative Personas: A Case Study on ChatGPT. In *Proceedings of the 17th International Symposium on Visual Information Communication and Interaction (VINCI '24)*. Association for Computing Machinery, New York, NY, USA, Article 29, 5 pages. doi:10.1145/3678698.3687172
- [46] Nik Hynek, Beata Gavurova, and Matus Kubak. 2025. Risks and benefits of artificial intelligence deepfakes: Systematic review and comparison of public attitudes in seven European Countries. *Journal of Innovation & Knowledge* 10, 5 (September 2025), 100782. doi:10.1016/j.jik.2025.100782
- [47] 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)*. Association for Computing Machinery, New York, NY, USA, 1–19. doi:10.1145/3544548.3580688
- [48] Essi Häyhänen, Joni Salminen, and Bernard J. Jansen. 2025. Why Are Personas the Way They Are? Identifying Six Persona Creation Strategies. *Persona Studies* 11 (March 2025), 18 pages. doi:10.21153/psj2025vol11noart2002
- [49] Bernard Jansen, Soon-gyo Jung, and Joni Salminen. 2021. The Effect of Hyperparameter Selection on the Personification of Customer Population Data. *International Journal of Electrical and Computer Engineering Research* 1, 2 (2021), 12 pages. doi:10.53375/ijecer.2021.31
- [50] Bernard J. Jansen, Soon-gyo Jung, and Joni Salminen. 2020. From flat file to interface: Synthesis of personas and analytics for enhanced user understanding. *Proceedings of the Association for Information Science and Technology* 57, 1 (2020), e215. doi:10.1002/pr2.215
- [51] Bernard J. Jansen, Soon-gyo Jung, Joni Salminen, Kathleen Guan, and Lene Nielsen. 2021. Strengths and Weaknesses of Persona Creation Methods: Outlining Guidelines for Novice and Experienced Users and Opportunities for Digital Innovations. In *Proceedings of the 54th Hawaii International Conference on System Sciences*. University of Hawai'i at Manoa, Hawai'i, USA, 4971–4980. doi:10.24251/HICSS.2021.604
- [52] Bernard J. Jansen, Joni Salminen, Soon-gyo Jung, and Kathleen Guan. 2021. *Evaluating Data-Driven Personas*. Springer International Publishing, Cham, 209–237. doi:10.1007/978-3-031-02231-9\_9
- [53] Yongyi Ji, Zhisheng Tang, and Mayank Kejriwal. 2024. Is persona enough for personality? Using ChatGPT to reconstruct an agent's latent personality from simple descriptions. In *Workshop on Large Language Models and Cognition*. PMLR, Colorado, USA, 5 pages. doi:10.48550/arXiv.2406.12216
- [54] Soon-Gyo Jung, Jisun An, Haewoon Kwak, Moeed Ahmad, Lene Nielsen, and Bernard J. Jansen. 2017. Persona Generation from Aggregated Social Media Data. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI EA '17). Association for Computing Machinery, New York, NY, USA, 1748–1755. doi:10.1145/3027063.3053120
- [55] Soon-Gyo Jung, Joni Salminen, Kholoud Khalil Aldous, and Bernard J. Jansen. 2025. PersonaCraft: Leveraging language models for data-driven persona development. *International Journal of Human-Computer Studies* 197 (2025), 103445. doi:10.1016/j.ijhcs.2025.103445
- [56] Soon-gyo Jung, Joni Salminen, Haewoon Kwak, Jisun An, and Bernard J. Jansen. 2018. Automatic Persona Generation (APG): A Rationale and Demonstration. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval* (New Brunswick, NJ, USA) (CHIIR '18). Association for Computing Machinery, New York, NY, USA, 321–324. doi:10.1145/3176349.3176893
- [57] Ilkka Kaate, Joni Salminen, Soon-Gyo Jung, João M. Santos, Essi Häyhänen, Trang Xuan, Jinan Y. Azem, and Bernard Jansen. 2025. The 'fourth wall' and other usability issues in AI-generated personas: comparing chat-based and profile personas. *Behaviour & Information Technology* 44, 16 (February 2025), 4136–4152. doi:10.1080/0144929X.2025.2469659
- [58] Ilkka Kaate, Joni Salminen, Soon-Gyo Jung, Trang Thi Thu Xuan, Essi Häyhänen, Jinan Y. Azem, and Bernard J. Jansen. 2025. "You Always Get an Answer": Analyzing Users' Interaction with AI-Generated Personas Given Unanswerable Questions and Risk of Hallucination. In *Proceedings of the 30th International Conference on Intelligent User Interfaces (IUI '25)*. Association for Computing Machinery, New York, NY, USA, 1624–1638. doi:10.1145/3708359.3712160
- [59] Ilkka Kaate, Joni Salminen, João M. Santos, Soon-Gyo Jung, Hind Almerkhi, and Bernard J. Jansen. 2024. "There Is something Rotten in Denmark": Investigating the Deepfake persona perceptions and their Implications for human-centered AI. *Computers in Human Behavior: Artificial Humans* 2, 1 (2024), 100031. doi:10.1016/j.chbah.2023.100031
- [60] Carolin Kaiser, Jakob Kaiser, Vladimir Manewitsch, Lea Rau, and Rene Schallner. 2025. Simulating Human Opinions with Large Language Models: Opportunities and Challenges for Personalized Survey Data Modeling. In *Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization (UMAP Adjunct '25)*. Association for Computing Machinery, New York, NY, USA, 82–86. doi:10.1145/3708319.3733685
- [61] Mahammed Kamruzzaman, Hieu Nguyen, Nazmul Hassan, and Gene Louis Kim. 2024. "A Woman is More Culturally Knowledgeable than a Man?": The Effect of Personas on Cultural Norm Interpretation in LLMs. arXiv. doi:10.48550/arXiv.2409.11636 Preprint.
- [62] Shivani Kapania, William Agnew, Motahhare Eslami, Hoda Heidari, and Sarah E. Fox. 2025. Simulacrum of Stories: Examining Large Language Models as Qualitative Research Participants. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 489, 17 pages. doi:10.1145/3706598.3713220
- [63] Devi Karolita, John Grundy, Tanjila Kanij, Humphrey Obie, and Jennifer McIntosh. 2024. CRAFT: A Persona Generation Tool for Requirements Engineering. In *Proceedings of the 19th International Conference on Evaluation of Novel Approaches to Software Engineering - ENASE*. SciTePress, Angers, France, 674–683. doi:10.5220/0012718400003687
- [64] Arshnoor Kaur, Amanda Aird, Harris Borman, Andrea Nicastro, Anna Leontjeva, Luiz Pizzato, and Dan Jermyn. 2025. Synthetic Voices: Evaluating the Fidelity of LLM-Generated Personas in Representing People's Financial Wellbeing. In *Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization (UMAP '25)*. Association for Computing Machinery, New York, NY, USA, 185–193. doi:10.1145/3699682.3728339
- [65] Hankyung Kim, Dong Yoon Koh, Gaeun Lee, Jung-Mi Park, and Youn-kyung Lim. 2019. Designing Personalities of Conversational Agents. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI EA '19). Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3290607.3312887
- [66] Barbara Kitchenham and Sally Charters. 2004. *Procedures for Performing Systematic Reviews*. Technical Report TR/SE-0401. Keele University and Empirical Software Engineering National ICT Australia Ltd. <https://www.inf.ufsc.br/~aldo.vw/kitchenham.pdf>
- [67] Thomas Kosch and Sebastian Feger. 2024. Risk or Chance? Large Language Models and Reproducibility in HCI Research. *Interactions* 31, 6 (Oct. 2024), 44–49. doi:10.1145/3695765
- [68] D Kosenko and D Zharikova. 2024. KRGP: Knowledge-Based Response Generation with Persona. *Journal of Mathematical Sciences* 285, 1 (2024), 49–59. doi:10.1007/s10958-024-07432-2
- [69] Christopher Klaus Lazik, Christopher Katins, Charlotte Kauter, Jonas Jakob, Caroline Jay, Lars Grunske, and Thomas Kosch. 2025. The Impostor is Among Us: Can Large Language Models Capture the Complexity of Human Personas?. In *Proceedings of the Mensch Und Computer 2025 (MuC '25)*. Association for Computing Machinery, New York, NY, USA, 434–451. doi:10.1145/3743049.3743057
- [70] Brian Lee, Savil Srivastava, Ranjitha Kumar, Ronen Brafman, and Scott R. Klemmer. 2010. Designing with interactive example galleries. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) (CHI '10). Association for Computing Machinery, New York, NY, USA, 2257–2266. doi:10.1145/1753326.1753667
- [71] Danielle Levac, Heather Colquhoun, and Kelly K O'Brien. 2010. Scoping studies: advancing the methodology. *Implementation Science* 5, 1 (sep 2010), 9 pages. doi:10.1186/1748-5908-5-69
- [72] Ang Li, Haozhe Chen, Hongseok Namkoong, and Tianyi Peng. 2025. LLM Generated Persona is a Promise with a Catch. In *First Workshop on Social Simulation with LLMs*. OpenReview, Newton Highlands, MA, 13 pages. <https://openreview.net/forum?id=mHn4n9YnkA>
- [73] Junyi Li, Charith Peris, Ninareh Mehrabi, Palash Goyal, Kai-Wei Chang, Aram Galstyan, Richard Zemel, and Rahul Gupta. 2024. The steerability of large language models toward data-driven personas. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Mexico City, Mexico, 7290–7305. doi:10.18653/v1/2024.naacl-long.405
- [74] Ruosen Li, Ruochen Li, Barry Wang, and Xinyu Du. 2024. IQA-EVAL: Automatic Evaluation of Human-Model Interactive Question Answering. In *Advances in Neural Information Processing Systems*, Vol. 37. Neural Information Processing Systems Foundation, Inc. (NeurIPS), San Diego, CA, 109894–109921. doi:10.52202/079017-3487
- [75] Victoria R Li, Yida Chen, and Naomi Saphra. 2024. ChatGPT Doesn't Trust Chargers Fans: Guardrail Sensitivity in Context. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Miami, Florida, USA, 6327–6345. doi:10.18653/v1/2024.emnlp-main.363
- [76] Yanan Li, Ying Liu, and Muran Yu. 2025. Consumer segmentation with large language models. *Journal of Retailing and Consumer Services* 82 (2025), 104078. doi:10.1016/j.jretconser.2024.104078
- [77] Yuxuan Li, Hirokazu Shirado, and Sauvik Das. 2025. Actions Speak Louder than Words: Agent Decisions Reveal Implicit Biases in Language Models. In

- Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25)*. Association for Computing Machinery, New York, NY, USA, 3303–3325. doi:10.1145/3715275.3732212
- [78] Yiren Liu, Pranav Sharma, Mehul Oswal, Haijun Xia, and Yun Huang. 2025. PersonaFlow: Designing LLM-Simulated Expert Perspectives for Enhanced Research Ideation. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference (DIS '25)*. Association for Computing Machinery, New York, NY, USA, 506–534. doi:10.1145/3715336.3735789
- [79] Priscilla Y. Lo, Annemiek Veldhuis, Alissa N. Antle, and Steve DiPaola. 2025. Noel: A Chatbot Persona to Support Children Designing for Others. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 133, 25 pages. doi:10.1145/3706598.3713836
- [80] Sha Luo, Sang Jung Kim, Zening Duan, and Kaiping Chen. 2024. A Sociotechnical Lens for Evaluating Computer Vision Models: A Case Study on Detecting and Reasoning about Gender and Emotion. arXiv. doi:10.48550/arXiv.2406.08222 Preprint.
- [81] Adyasha Maharana, Dong-Ho Lee, Sergey Tulyakov, Mohit Bansal, Francesco Barbieri, and Yuwei Fang. 2024. Evaluating Very Long-Term Conversational Memory of LLM Agents. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Bangkok, Thailand, 13851–13870. doi:10.18653/v1/2024.acl-long.747
- [82] Georgios Makridis, Vasileios Koukos, Georgios Fatouros, Dimitrios Kotios, Maria Margarita Separdani, Dimosthenis Kyriazis, and John Soldatos. 2025. VirtualXAI: A User-Centric Framework for Explainability Assessment Leveraging GPT-Generated Personas. In *2025 21st International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT)*. IEEE, Lucca, Italy, 1034–1041. doi:10.1109/DCOSS-IoT65416.2025.00156
- [83] Patrick M. Markey, Hanna Campbell, and Samantha Goldman. 2025. A framework for the initial phases of personality test development using large language models and artificial personas. *Journal of Research in Personality* 118 (2025), 104647. doi:10.1016/j.jrp.2025.104647
- [84] Nicola Marsden and Maren Haag. 2016. Stereotypes and Politics: Reflections on Personas. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 4017–4031. doi:10.1145/2858036.2858151
- [85] Roman Mayr, Michel Schimpf, and Thomas Bohm. 2025. ChatChecker: A Framework for Dialogue System Testing and Evaluation Through Non-cooperative User Simulation. arXiv. doi:10.48550/arXiv.2507.16792 Preprint.
- [86] Sarah Mercer, Daniel P Martin, and Phil Swatton. 2025. Applying Psychometrics to Large Language Model Simulated Populations: Recreating the HEXACO Personality Inventory Experiment with Generative Agents. arXiv. doi:10.48550/arXiv.2508.00742 Preprint.
- [87] Tomasz Miasiekiewicz and Kenneth A. Kozar. 2011. Personas and user-centered design: How can personas benefit product design processes? *Design Studies* 32, 5 (2011), 417–430. doi:10.1016/j.destud.2011.03.003
- [88] Tea Mijač, Mario Jadrić, and Maja Čukušić. 2018. The potential and issues in data-driven development of web personas. In *2018 41st International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*. IEEE, New York, USA, 1237–1242. doi:10.23919/MIPRO.2018.8400224
- [89] Alex Milovic, Moumita Das Gyomlai, Brian Spaid, and Rebecca Dingus. 2024. Sell me this artificial pen: Using ChatGPT to enhance sales role plays. *Marketing Education Review* 34, 2 (2024), 137–143. doi:10.1080/10528008.2024.2341278
- [90] Ishani Mondal, Shwetha S, Anandhavelu Natarajan, Aparna Garimella, Sambaran Bandyopadhyay, and Jordan Boyd-Graber. 2024. Presentations by the Humans and For the Humans: Harnessing LLMs for Generating Persona-Aware Slides from Documents. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, St. Julian's, Malta, 2664–2684. doi:10.18653/v1/2024.eacl-long.163
- [91] Swapnil Morandé and Mitra Amini. 2023. Digital Persona: Reflection on the Power of Generative AI for Customer Profiling in Social Media Marketing. Qeios. doi:10.32388/0qi028 Preprint.
- [92] Eduardo Mosqueira-Rey et al. 2023. Human-in-the-loop machine learning: A state of the art. *Artificial Intelligence Review* 56 (2023), 3005–3054. doi:10.1007/s10462-022-10246-w
- [93] Michael Muller and Katie Seaborn. 2025. Stepford Twins and Potemkin Engineering: A Critique of Synthetic Personas in the Age of Generative AI. In *ProceeSynthetic Data Workshop*. Aarhus Conference, Aarhus Denmark, 4 pages. <https://synthetic-data-workshop.github.io/papers/10.pdf>
- [94] Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Naveed Akhtar, Nick Barnes, and Ajmal Mian. 2025. A Comprehensive Overview of Large Language Models. *ACM Transactions on Intelligent Systems and Technology* 16, 5, Article 106 (Aug. 2025), 72 pages. doi:10.1145/3744746
- [95] Timothy Neate, Aikaterini Bourazeri, Abi Roper, Simone Stumpf, and Stephanie Wilson. 2019. Co-Created Personas: Engaging and Empowering Users with Diverse Needs Within the Design Process. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3290605.3300880
- [96] Ha Nguyen, Victoria Nguyen, Sariah López-Fierro, Sara Ludovise, and Rossella Santagata. 2024. Simulating Climate Change Discussion with Large Language Models: Considerations for Science Communication at Scale. In *Proceedings of the Eleventh ACM Conference on Learning @ Scale (L@S '24)*. Association for Computing Machinery, New York, NY, USA, 28–38. doi:10.1145/3657604.3662033
- [97] Lene Nielsen. 2019. *Personas - User Focused Design* (2nd ed. 2019 edition ed.). Springer, New York, NY. doi:10.1007/978-1-4471-7427-1
- [98] 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. doi:10.4018/ijskd.2015010104
- [99] Lucas Nóbrega, Luiz Felipe Martinez, Yuri Lima, Carlos Eduardo Barbosa, Matheus Argôlo, Herbert Salazar, Geraldo Xexéo, and Jano Moreira De Souza. 2025. Towards a Knowledge Management Framework for LLM-Generated Personas in Collaborative Systems. In *2025 28th International Conference on Computer Supported Cooperative Work in Design (CSCWD)*. IEEE, New York, USA, 539–545. doi:10.1109/CSCWD64889.2025.11033648
- [100] Giovanna Nunes Vilaza, Kevin Doherty, Darragh McCashin, David Coyle, Jakob Bardram, and Marguerite Barry. 2022. A Scoping Review of Ethics Across SIGCHI. In *Proceedings of the 2022 ACM Designing Interactive Systems Conference (Virtual Event, Australia) (DIS '22)*. Association for Computing Machinery, New York, NY, USA, 137–154. doi:10.1145/3532106.3533511
- [101] Virpi Oksman and Niina Mallat. 2025. Data-Driven Personas in Drone Delivery: Leveraging Generative AI for Enhanced User Insights and Experience Design. In *HCI International 2025 Posters*, Constantine Stephanidis, Margherita Antona, Stavroula Ntoa, and Gavriel Salvendy (Eds.). Springer Nature Switzerland, Cham, 353–362. doi:10.1007/978-3-031-94168-9\_35
- [102] Matthew J Page, Joanne E McKenzie, Patrick M Bossuyt, Isabelle Boutron, Tammy C Hoffmann, Cynthia D Mulrow, Larissa Shamseer, Jennifer M Tetzlaff, Elie A Akl, Sue E Brennan, et al. 2021. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *BMJ* 372 (2021), 9 pages. doi:10.1136/bmj.n71
- [103] Swaroop Panda. 2024. LLMs' ways of seeing User Personas. arXiv. doi:10.48550/arXiv.2409.14858 Preprint.
- [104] Stefano De Paoli. 2023. Improved prompting and process for writing user personas with LLMs, using qualitative interviews: Capturing behaviour and personality traits of users. arXiv. doi:10.48550/arXiv.2310.06391 Preprint.
- [105] Stefano De Paoli. 2023. Writing user personas with Large Language Models: Testing phase 6 of a Thematic Analysis of semi-structured interviews. arXiv. doi:10.48550/arXiv.2305.18099 Preprint.
- [106] Jeongeun Park and DaEun Choi. 2023. AudiLens: Configurable LLM-Generated Audiences for Public Speech Practice. In *Adjunct Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23 Adjunct)*. Association for Computing Machinery, New York, NY, USA, Article 122, 3 pages. doi:10.1145/3586182.3625114
- [107] Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2022. Social Simulacra: Creating Populated Prototypes for Social Computing Systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)*. Association for Computing Machinery, New York, NY, USA, Article 74, 18 pages. doi:10.1145/3526113.3545616
- [108] Kyeongman Park, Minbeom Kim, and Kyomin Jung. 2025. A Character-Centric Creative Story Generation via Imagination. In *Findings of the Association for Computational Linguistics: ACL 2025*. Association for Computational Linguistics, Vienna, Austria, 1598–1645. doi:10.18653/v1/2025.findings-acl.82
- [109] Sohyeon Park, Aehong Min, Jesus Armando Beltran, and Gillian R Hayes. 2025. "As an Autistic Person Myself." The Bias Paradox Around Autism in LLMs. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 774, 17 pages. doi:10.1145/3706598.3713420
- [110] Alisha Pradhan and Amanda Lazar. 2021. Hey Google, Do You Have a Personality? Designing Personality and Personas for Conversational Agents. In *Proceedings of the 3rd Conference on Conversational User Interfaces (CUI '21)*. Association for Computing Machinery, New York, NY, USA, Article 12, 4 pages. doi:10.1145/3469595.3469607
- [111] Mirjana Prpa, Giovanni Troiano, Bingsheng Yao, Toby Jia-Jun Li, Dakuo Wang, and Hansu Gu. 2024. Challenges and Opportunities of LLM-Based Synthetic Personae and Data in HCI. In *Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing (CSCW Companion '24)*. Association for Computing Machinery, New York, NY, USA, 716–719. doi:10.1145/3678884.3681826

- [112] John Pruitt and Tamara Adlin. 2010. *The persona lifecycle: keeping people in mind throughout product design* (1 edition ed.). Elsevier, Boston. doi:10.1016/B978-0-12-566251-2.X5000-X
- [113] John Pruitt and Jonathan Grudin. 2003. Personas: practice and theory. In *Proceedings of the 2003 Conference on Designing for User Experiences* (San Francisco, California) (DUX '03). Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/997078.997089
- [114] Leen F. Al Qadi, Soon-Gyo Jung, Danial Amin, Amani Alabed, Joni Salminen, and Bernard J. Jansen. 2025. Representing Religious Practices via AI-Generated Personas: A Case Study of Ramadan Behaviors from Four Predominantly Muslim Countries. In *2025 IEEE/ACM 22nd International Conference on Computer Systems and Applications (AICCSA)*. IEEE, New York, USA, 1–10. doi:10.1109/AICCSA66935.2025.11315204
- [115] Eric Rudolph, Natalie Engert, and Jens Albrecht. 2024. An AI-Based Virtual Client for Educational Role-Playing in the Training of Online Counselors. In *Proceedings of the 16th International Conference on Computer Supported Education - Volume 2: CSEdu*. INSTICC, SciTePress, Portugal, 108–117. doi:10.5220/0012690700003693
- [116] Joni Salminen, Kathleen Guan, Soon-Gyo Jung, Shammur A. Chowdhury, and Bernard J. Jansen. 2020. A Literature Review of Quantitative Persona Creation. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3313831.3376502
- [117] Joni Salminen, Kathleen Guan, Soon-Gyo Jung, and Bernard J. Jansen. 2021. A Survey of 15 Years of Data-Driven Persona Development. *International Journal of Human-Computer Interaction* 37, 18 (2021), 1685–1708. doi:10.1080/10447318.2021.1908670
- [118] Joni Salminen, Soon-gyo Jung, Hind Almerexhi, Erik Cambria, and Bernard Jansen. 2023. How Can Natural Language Processing and Generative AI Address Grand Challenges of Quantitative User Personas?. In *HCI International 2023 – Late Breaking Papers*. Springer Nature Switzerland, Cham, 211–231. doi:10.1007/978-3-031-48057-7\_14
- [119] 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 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 510, 20 pages. doi:10.1145/3613904.3642036
- [120] Joni Salminen, João M. Santos, Soon gyo Jung, and Bernard J. Jansen. 2024. Picturing the fictitious person: An exploratory study on the effect of images on user perceptions of AI-generated personas. *Computers in Human Behavior: Artificial Humans* 2, 1 (2024), 100052. doi:10.1016/j.chbah.2024.100052
- [121] Joni Salminen, Joao M. Santos, Haewoon Kwak, Jisun An, Soon gyo Jung, and Bernard J. Jansen. 2020. Persona Perception Scale: Development and Exploratory Validation of an Instrument for Evaluating Individuals' Perceptions of Personas. *International Journal of Human-Computer Studies* 141 (2020), 102437. doi:10.1016/j.ijhcs.2020.102437
- [122] Joni Salminen, Kathleen Wenyun Guan, Soon-Gyo Jung, and Bernard Jansen. 2022. Use Cases for Design Personas: A Systematic Review and New Frontiers. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 543, 21 pages. doi:10.1145/3491102.3517589
- [123] Vanessa Sattelle and Juan Carlos Ortiz. 2024. Generating user personas with AI: Reflecting on its implications for design. In *Proceedings of Design Research Society conference 2024*. Design Research Society, London, England, 15 pages. doi:10.21606/drs.2024.1024
- [124] Eva Schirgi, Stefanie Hatzl, and Katharina Moitzi. 2024. Development of a Framework for Data-Supported Personas. In *Advances in Production Management Systems. Production Management Systems for Volatile, Uncertain, Complex, and Ambiguous Environments*. Springer Nature Switzerland, Cham, 18–31. doi:10.1007/978-3-031-71633-1\_2
- [125] Douglas Schuler and Aki Namioka. 1993. *Participatory design: Principles and practices*. CRC Press, Boca Raton. doi:10.1201/9780203744338
- [126] Andreas Schuller, Doris Janssen, Julian Blumenröther, Theresa Maria Probst, Michael Schmidt, and Chandan Kumar. 2024. Generating personas using LLMs and assessing their viability. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, Article 179, 7 pages. doi:10.1145/3613905.3650860
- [127] Sandeep Singh Sengar, Affan Bin Hasan, Sanjay Kumar, and Fiona Carroll. 2025. Generative artificial intelligence: a systematic review and applications. *Multimedia Tools and Applications* 84, 21 (May 2025), 23661–23700. doi:10.1007/s11042-024-20016-1
- [128] Sankalp Sethi, Joni Salminen, Danial Amin, and Bernard J. Jansen. 2025. "When AI Writes Personas": Analyzing Lexical Diversity in LLM-Generated Persona Descriptions. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 35, 8 pages. doi:10.1145/3706599.3719712
- [129] Joongi Shin, Michael A. Hedderich, Bartłomiej Jakub Rey, Andrés Lucero, and Antti Oulasvirta. 2024. Understanding Human-AI Workflows for Generating Personas. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference (DIS '24)*. Association for Computing Machinery, New York, NY, USA, 757–781. doi:10.1145/3643834.3660729
- [130] Ben Shneiderman. 2022. *Human-Centered AI*. Oxford University Press, Oxford, England. doi:10.1093/oso/9780192845290.001.0001
- [131] Linda J Skitka, Kathleen L Mosier, and Mark Burdick. 2000. Accountability and Automation Bias. *International Journal of Human-Computer Studies* 52, 4 (2000), 701–717. doi:10.1006/ijhc.1999.0349
- [132] Urška Smrke, Ana Rehberger, Nejc Plohl, and Izidor Mlakar. 2025. Exploring the Feasibility of Generative AI in Persona Research: A Comparative Analysis of Large Language Model-Generated and Human-Crafted Personas in Obesity Research. *Applied Sciences* 15, 4 (2025), 19 pages. doi:10.3390/app15041937
- [133] Guangzhi Sun, Xiao Zhan, and Jose Such. 2024. Building Better AI Agents: A Provocation on the Utilisation of Persona in LLM-based Conversational Agents. In *Proceedings of the 6th ACM Conference on Conversational User Interfaces (CUI '24)*. Association for Computing Machinery, New York, NY, USA, Article 35, 6 pages. doi:10.1145/3640794.3665887
- [134] Lipepei Sun, Tianzi Qin, Anran Hu, Jiale Zhang, Shuojia Lin, Jianyan Chen, Mona Ali, and Mirjana Prpa. 2025. Persona-L has Entered the Chat: Leveraging LLMs and Ability-based Framework for Personas of People with Complex Needs. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1109, 31 pages. doi:10.1145/3706598.3713445
- [135] Seungjong Sun, Eungu Lee, Seo Yeon Baek, Seunghyun Hwang, Wonbyung Lee, Dongyan Nan, Bernard J. Jansen, and Jang Hyun Kim. 2024. Kiss up, Kick down: Exploring Behavioral Changes in Multi-modal Large Language Models with Assigned Visual Personas. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Miami, Florida, USA, 10888–10901. doi:10.18653/v1/2024.emnlp-main.609
- [136] Meiling Tao, Liang Xuechen, Tianyu Shi, Lei Yu, and Yiting Xie. 2024. RoleCraft-GLM: Advancing Personalized Role-Playing in Large Language Models. In *Proceedings of the 1st Workshop on Personalization of Generative AI Systems*. Association for Computational Linguistics, St. Julians, Malta, 1–9. <https://aclanthology.org/2024.personalize-1/>
- [137] Yu-Min Tseng, Yu-Chao Huang, Teng-Yun Hsiao, Wei-Lin Chen, Chao-Wei Huang, Yu Meng, and Yun-Nung Chen. 2024. Two Tales of Persona in LLMs: A Survey of Role-Playing and Personalization. 16612–16631 pages. doi:10.18653/v1/2024.findings-emnlp.969
- [138] Phil Turner and Susan Turner. 2011. Is stereotyping inevitable when designing with personas? *Design Studies* 32, 1 (2011), 30–44. doi:10.1016/j.destud.2010.06.002
- [139] Pranav Narayanan Venkit, Jiayi Li, Yingfan Zhou, Sarah Rajtmajer, and Shomir Wilson. 2025. A Tale of Two Identities: An Ethical Audit of Human and AI-Crafted Personas. arXiv. doi:10.48550/arXiv.2505.07850 Preprint.
- [140] Lena L. Voigt, Felix Freiling, and Christopher J. Hargreaves. 2024. Re-imagining: Generating coherent background activity in synthetic scenario-based forensic datasets using large language models. *Forensic Science International: Digital Investigation* 50 (2024), 301805. doi:10.1016/j.fsidi.2024.301805
- [141] Sarah Webber, Ryan M. Kelly, Greg Wadley, and Wally Smith. 2023. Engaging with Nature through Technology: A Scoping Review of HCI Research. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 521, 18 pages. doi:10.1145/3544548.3581534
- [142] Thanatat Wongabut, Unhawa Ninrutsirikun, and Chonlameth Arpikanondt. 2025. Automating Persona Generation: Leveraging ChatGPT-4 in Educational Digital Services. In *2025 Joint International Conference on Digital Arts, Media and Technology with ECTI Northern Section Conference on Electrical, Electronics, Computer and Telecommunications Engineering (ECTIDAMTNCN)*. IEEE, New York, USA, 549–554. doi:10.1109/ECTIDAMTNCN64748.2025.10962020
- [143] Shujin Wu, Yi R. Fung, Cheng Qian, Jeonhwan Kim, Dilek Hakkani-Tur, and Heng Ji. 2025. Aligning LLMs with Individual Preferences via Interaction. In *Proceedings of the 31st International Conference on Computational Linguistics*. Association for Computational Linguistics, Abu Dhabi, UAE, 7648–7662. <https://aclanthology.org/2025.coling-main.511/>
- [144] Rui Xu, Xintao Wang, Jiangjie Chen, Siyu Yuan, Xinfeng Yuan, Jiaqing Liang, Zulong Chen, Xiaoqingdong, and Yanghua Xiao. 2025. Character is Destiny: Can Persona-assigned Language Models Make Personal Choices?. In *Findings of the Association for Computational Linguistics: EMNLP 2025*. Association for Computational Linguistics, Suzhou, China, 15038–15059. doi:10.18653/v1/2025.findings-emnlp.813
- [145] Xi Yang and Marco Aurisicchio. 2021. Designing Conversational Agents: A Self-Determination Theory Approach. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (Yokohama, Japan) (CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 256, 16 pages. doi:10.1145/3411764.3445445

- [146] William Yates, Trilby VanDeusen, Sierra Schuman, Keanna Curry, Josh Maki, Matt Ogborn, Allison Leaming Malecki, Joyce Ann Martin, Wes Edens, Mimmo Bonanni, and Eric Prosser. 2025. Using AI personas in ChatGPT for online chat reference training. *Library Hi Tech News* 42, 7 (June 2025), 1–5. doi:10.1108/LHTN-08-2024-0137
- [147] Leo Yeykelis, Kaavya Pichai, James J. Cummings, and Byron Reeves. 2024. Using Large Language Models to Create AI Personas for Replication and Prediction of Media Effects: An Empirical Test of 133 Published Experimental Research Findings. arXiv. doi:10.48550/arXiv.2408.16073 Preprint.
- [148] Min Yin, Haoyu Liu, Boyi Lian, and Chunlei Chai. 2025. Co-persona: Leveraging LLMs and Expert Collaboration to Understand User Personas through Social Media Data Analysis. arXiv. doi:10.48550/arXiv.2506.18269 Preprint.
- [149] Xishuo Zhang, Lin Liu, Yi Wang, Xiao Liu, Hailong Wang, Chetan Arora, Haichao Liu, Weijia Wang, and Thuong Hoang. 2024. Auto-Generated Personas: Enhancing User-centered Design Practices among University Students. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, Article 52, 7 pages. doi:10.1145/3613905.3651043
- [150] Xi Zheng, Zhuoyang Li, Xinning Gui, and Yuhan Luo. 2025. Customizing Emotional Support: How Do Individuals Construct and Interact With LLM-Powered Chatbots. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 376, 20 pages. doi:10.1145/3706598.3713453
- [151] Leijing Zhou, Yan Fang, Shiyang Ding, Yiming Cheng, Bing Yan, Wanye Zhu, Song Bao, Jian Wang, and Siyang Song. 2025. Vivid-persona: customizable persona tool with interactive and immersive experiences. *Journal of Engineering Design* 36, 2 (2025), 276–297. doi:10.1080/09544828.2024.2383543