

From Generation to Simulation: Responsible Use of AI Personas in Human-Centered Design and Research

A. Baki Kocaballi
School of Computer Science
University of Technology Sydney
Sydney, Australia
baki.kocaballi@uts.edu.au

Mirjana Prpa
Khoury College of Computer Sciences
Northeastern University
Vancouver, British Columbia, Canada
m.prpa@northeastern.edu

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

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

Bernard J Jansen
Qatar Computing Research Institute
HBKU
Doha, Qatar
bjansen@hbku.edu.qa

Abstract

Generative AI (GenAI) is increasingly transforming human-centered design (HCD) by enabling the generation and simulation of AI-driven personas that can role-play different stakeholders, ideate solutions, and perform usability tests. With their increased efficiency and accessibility, these AI personas offer advantages such as scalability, rapid ideation, coverage of diverse and edge cases, and reduced reliance on costly human subject studies at early stages of design. However, they also pose ethical, representational, and methodological challenges. Although evidence on the use, benefits, and limitations of AI personas is growing, the field still lacks shared validity standards and reporting norms to guide the use of AI personas and their synthetic feedback, and decisions about when to involve actual human participants. This half-day workshop will examine the roles, opportunities, and responsibilities associated with integrating GenAI simulated personas into HCI design practice. Through the presentation of emerging evidence, collaborative discussions, and hands-on activities, we aim to explore responsible methods and develop practical guidelines for the responsible use of AI personas in human-centered design and research.

CCS Concepts

• **Human-centered computing** → Human computer interaction (HCI); HCI design and evaluation methods.

Keywords

Generative AI, Personas, Synthetic Users, User Simulation, UX Evaluation, Ethics & Governance

ACM Reference Format:

A. Baki Kocaballi, Mirjana Prpa, Joni Salminen, Danial Amin, and Bernard J Jansen. 2026. From Generation to Simulation: Responsible Use of AI Personas in Human-Centered Design and Research. In *Extended Abstracts*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CHI EA '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2281-3/26/04

<https://doi.org/10.1145/3772363.3778745>

of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26), April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 5 pages.
<https://doi.org/10.1145/3772363.3778745>

1 Introduction and motivation

Personas, first introduced by Cooper [1], were designed to represent groups of users who share similar characteristics. Over time, they have been adopted across various domains, including education [3], healthcare [14], marketing [9], and technology design [17]; however, they have been most influential in human-centered design (HCD) [7]. Personas address a longstanding challenge in HCD: how to maintain user involvement throughout the design process. While ongoing user participation is central to HCD principles, it is often constrained by budget, accessibility, or scale [20, 24]. Personas emerged as a practical proxy, keeping user perspectives visible when real users cannot be directly involved.

Traditionally, personas have been static artefacts: brief, typically 1-page profiles detailing demographic attributes, motivations, goals, and pain points, often accompanied by a headshot or photos of personal items to make them feel more vivid [23]. These profiles were intended to help design teams “step into the shoes” of their users, empathize with their experiences, and anticipate their responses to design ideas. Yet evidence suggests that these static personas are often under-utilized, treated as paperwork rather than a living part of the process, rarely updated, and sometimes ignored altogether [5, 15, 18].

The advancements in Generative AI (GenAI), particularly large language models (LLMs), is opening up new opportunities and ways to develop and engage with personas. LLMs can create personas from various types of data or enrich existing descriptions with contextual detail, backstories, and lived-experience narratives, providing new ways to generate data-driven personas [19]. More recently, a new line of practice with interactive AI personas have emerged. These are LLM-driven agents that role-play based on persona descriptions and simulate the personas [2, 6, 10, 11, 13, 22, 25]. These interactive AI personas enable practitioners to converse with them, explore different perspectives on design problems, co-create ideas, and receive feedback on early prototypes [2, 6, 22]. A recent study with UX designers [2] demonstrates that these interactive AI personas can provide valuable simulations of user viewpoints across

multiple stages of the design process. These developments promise greater scalability, accessibility, and engagement, but also introduce new challenges. AI-powered personas can misrepresent user groups and introduce bias [8], hallucinate details [11], or exhibit sycophantic agreement [2]. LLM-generated personas' high fidelity can create an illusion of authenticity that risks over-reliance, particularly among less-experienced designers [2]. While static personas have been rarely seen as a *threat* to real user research, these AI-powered versions raise new ethical and methodological concerns: the temptation to substitute simulated engagement for genuine involvement of people [12].

There is also a subtler risk, which is that automating persona creation can distance designers from the process of immersing themselves in the user's context and grappling with raw data about users, an act that often yields deeper insights and empathy. By outsourcing this task to AI, we may inadvertently lose that learning opportunity. This concern reflects a broader issue with the increasing automation of AI-generated design artefacts [4]. As Forlizzi and Zimmerman argue [4], the real value often resides in the process of making, not only in the finished design artefact. Applied to persona creation and use, this suggests that what teams need is not "full automation" of personas, but tools and workflows that enable richer human-AI collaboration in crafting, testing, and refining them in order to keep team members active in the loop while leveraging AI's generative capacities [21].

To facilitate the appropriate use of AI personas in HCD practice, we need discipline-wide guardrails and guidance tailored to real workflows: (i) a minimum reporting set that discloses provenance, model version, prompts, intended use, and known limits; (ii) risk audits at key milestones to surface representation harms, failure modes, exposure, and mitigations; (iii) validity methods for coherence and representational faithfulness supported by real user data and checkpoints; and (iv) responsible-use guidelines that address the risks and limitations. This workshop will take stock of this fast-evolving space. We aim to surface emerging evidence, share experiences of AI persona generation and simulation, and work towards a foundation for responsible and transparent use of AI personas in HCD as a complement, not a replacement, to real user research and engagement.

2 Major themes and goals

The workshop will invite HCI/CSCW researchers, UX practitioners, tool builders, methodologists, AI ethics/governance leads, accessibility specialists, and social-computing scholars. We target 25–35 participants with a balance of: (i) AI-personas simulation tool builders; (ii) design researchers who have piloted synthetic users; (iii) critical perspectives (e.g., ethics, marginalized user advocates); and (iv) domain specialists (health, education, civic tech) to share their experiences and insights. We invite short, concrete contributions on (but not limited to):

Studies and Cases

- **Use of AI personas across HCD stages in practice and education:** ideation, scenario generation, counter-scenario creation, early usability "walkthroughs," and lightweight formative evaluation.
- **Team experiences with AI persona workflows:** building persona libraries, curating/retiring personas, orchestrating multi-agent "populated prototypes," and integrating outputs into design sprints.
- **Comparative validity studies:** how AI personas align with or diverge from findings with real participants (means vs. variance, behavioral plausibility, edge cases), and how teams calibrate trust.
- **Sensitive and high-stakes domains:** representing marginalized groups, ability-based use cases, and non-Western contexts; when simulations help, when they may cause harm, and what safeguards are necessary.
- **Designer perceptions and mental models:** empowerment vs. deskilling, over-reliance risks, appropriate reliance, disclosure to stakeholders, and communicating uncertainty.
- **Case reports of failure modes:** stereotype leakage, mode collapse, drift across model versions, hallucinated constraints, and mismatches between simulated and real user behaviors.

Design, Development, and Evaluation

- **Emerging methods for AI persona construction:** data-grounded and synthetic-only personas; retrieval-augmented grounding; memory policies; controllable traits and goals; multi-agent coordination for task flows.
- **Novel UX evaluation techniques:** scripted runs, heuristic + simulated walkthroughs, and mixed pipelines that combine AI personas with human checks; metrics for central tendency, variance/coverage, and ecological plausibility.
- **Context-specific design guidelines:** when to use AI personas (and when not); scoping intended use; counter-persona generation; prompts and constraints that reduce bias; accessibility-aware persona authoring.
- **Standardized measurement and reporting:** practical Persona/Synthetic-User Cards covering provenance, grounding sources, prompts/parameters, model/versioning, known biases, and evaluation plans; logging templates for reproducibility. Design factors in simulated interactions including feedback, feedforward, affordances, and explainability in persona-system exchanges; documenting assumptions embedded in prompts and toolchains.
- **Ethical and regulatory dimensions:** bias and harm audits; records for accountability; guardrails against inappropriate substitution for real participants.
- **Design factors in simulated interactions:** feedback, feedforward, affordances, and explainability in persona-system exchanges; documenting assumptions embedded in prompts and toolchains.
- **Domain-specific practices:** safety-critical and fairness-critical settings (e.g., health, finance, public services); accessibility and personalization. Evolving roles and practices including how designers, researchers, and domain experts collaborate with AI personas; team norms for curation, review, and retirement of persona sets; implications for pedagogy and research training.

3 Organizers

A. Baki Kocaballi is a Senior Lecturer at the University of Technology Sydney and an internationally recognized researcher on conversational agents. His work is situated at the intersection of generative AI, technology design, digital health, and education, focusing on collaborative human-AI interaction models that reduce fixation and uncritical adoption, support co-creativity, and preserve human agency. His work has received several academic awards and recognitions, and his research has been published in the top-ranking venues.

Mirjana Prpa is an Assistant Teaching Professor at Northeastern University. Her research interests include leveraging micro-phenomenology in HCI for understanding the complexity of user experiences [16] and extending it to unfold experiences arising from human-AI persona interactions, spanning from LLM use for persona creation to LLM-based agents in social VR.

Joni Salminen is an Associate Professor at the University of Vaasa and an Adjunct Professor at the Turku School of Economics in Finland. His research is concentrated on human-computer interaction (HCI), digital marketing, and AI, with a particular focus on the automatic generation and application of user personas from data sources like social media and web analytics. A former Scientist at the Qatar Computing Research Institute, Dr. Salminen has extensively published on topics including data-driven persona biases, the societal impacts of machine decision-making, and online platforms.

Danial Amin is a doctoral candidate at the University of Vaasa, focusing on ethical AI and persona development for social good, with extensive experience in data science and AI development. His research leverages GenAI technologies to develop ethical and inclusive AI-generated personas that benefit marginalized communities in the Global South.

Dr. Jim Jansen is a Principal Scientist at the Qatar Computing Research Institute working on automatically visualizing user data. He is a West Point graduate with a Ph.D. in computer science from Texas A&M University. Professor Jansen is editor-in-chief of Information Processing & Management, former interim editor-in-chief of the International Journal of Information Management, and former editor-in-chief of Internet Research. He has received several awards and honors, including an ACM Research Award, six application development awards, a product innovation award, and a university-level teaching award, along with other writing, publishing, research, teaching, and leadership honors. Dr. Jansen has authored or co-authored nearly 400 research publications, with articles appearing in a multi-disciplinary range of journals and conferences. He is a co-author of the books *Data-Driven Personas* (Springer Nature) and *Understanding Audiences, Customers, and Users via Analytics* (Springer Nature). For more information, see <https://www.bernardjjansen.com>

4 Pre-workshop plan

We will create a workshop website to share the CFP, updates, and position papers. We will set up a Discord channel for participants to introduce themselves and to conduct a brief survey of participants' tools and goals. The call for participation will be distributed across HCI, AI, design, and health informatics mailing lists. The organizers

will also use social networking, such as LinkedIn, X, and Bluesky, to target a wider research audience. The organizers will also ask for submissions from personal contacts with relevant research expertise, from the organizers' home institutions, from contacts in the relevant industry and other non-academic institutions.

5 Workshop structure

1:00PM – 1:30PM	Introduction and Keynote Brief overview of GenAI personas in HCI: promises and pitfalls. An invited keynote from a senior researcher emphasizing ethical considerations.
1:30PM – 2:30PM	Session 1 – Lightning Talks Participants share short presentations on their experiences using GenAI personas. Focus on lessons learned, challenges faced, and ethical dilemmas encountered.
2:30PM – 3:00PM	Coffee Break
3:00PM – 3.30PM	Session 2 – Anatomy of AI Persona Generation and Simulation: Participants work in small groups to map the end-to-end pipeline of AI persona generation and simulation. Simple templates and colour-coded sticky notes will be used to capture stages, assumptions, risks, and opportunities. This activity will result in a large, annotated visualization of Anatomy of AI persona pipeline.
3:30PM– 4:15 PM	Session 3 – Risk & Bias Audit: Building on the pipeline visualization, participants rotate through breakout tables, each focused on a specific risk lens such as bias, misuse, over-trust, and misalignment with stakeholders. Using short scenario cards and a checklist template, groups rapidly identify concrete failure modes and minimal validation steps for different domains (e.g., healthcare, education). The session culminates in a practical Risk & Bias checklist that any team can apply in less than 15 minutes.
4:15PM – 4:30PM	Wrap up and roadmap Summarizing the day and discussing the next steps with the participants

6 Post-workshop plans

Following the workshop, we will mobilize several initiatives to sustain momentum and build shared resources. With participants'

consent, all accepted position papers will be made available on our workshop website and archived through the **CEUR Workshop Proceedings** (<https://ceur-ws.org>). Together with interested participants, we will establish a **working group** focused on developing reporting standards, guidelines, validity benchmarks, and auditing practices for AI personas and synthetic users. We will also invite workshop participants to co-author a **joint article** synthesizing the workshop discussions and outcomes, with the aim of publishing in ACM Interactions Magazine to reach both academic and practitioner communities. The working group will also release v1 drafts of two community resources: **Anatomy of an AI-Persona Pipeline Visualization** and a **Risk & Bias Audit Checklist**, on our website and in subsequent publications.

7 Accessibility

Prior to the workshop, we will contact all registered participants to determine the best way to support their participation. A short accessibility questionnaire will be shared via our Discord channel and email, inviting attendees to indicate any access needs. We will work closely with the workshop chairs to address the support needs of our participants. All presenters will use CHI accessible authoring guidelines (e.g., alt text for images, structured headings), and microphones will be used for all speaking turns.

8 Call for participation

Generative AI can rapidly produce simulated personas and synthetic users, empowering design teams to quickly ideate, prototype, and evaluate. However, without careful use, AI-generated personas risk bias, stereotyping, representational harm, and validity gaps. How can we responsibly embrace AI personas while mitigating their risks? We invite HCI researchers, UX practitioners, AI ethics experts, design methodologists, and developers to join our half-day CHI workshop. We aim to collaboratively develop practical guidance and resources for the responsible use of AI personas, focusing on transparency, risk awareness, and validity.

We invite brief contributions that share real experiences with AI personas: where they improved or misled design, how teams integrated them into workflows, and what tools supported reproducibility. We're interested in practical ways to evaluate their plausibility, diversity, and reliability, as well as ethical issues around consent, disclosure, and accessibility.

Interested participants should submit a short position paper (2–4 pages, CHI Extended Abstract format), case study, or artifact abstract, including examples or experiences related to AI personas, synthetic users, or related methods. Submissions will be reviewed by at least two organizers for their relevance and practicality to the workshop goals, clarity of methods, and diversity of perspectives and domains. At least one author of each accepted submission must attend the workshop and that all participants must register for the workshop.

For details and the submission form, visit our workshop website: <https://sites.google.com/view/ai-personas>. Join us to help shape principles and guidelines for the responsible use of AI personas in human-centred design and research.

Acknowledgments

During the preparation of this work the author(s) used ChatGPT 5 to improve readability and language of the final draft. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the published workshop proposal.

References

- [1] Alan Cooper. 1999. The Inmates are Running the Asylum. In U. Arend, E. Eberle, and K. Pitschke, eds., *Software-Ergonomie '99: Design von Informationswelten*. Vieweg+Teubner Verlag, Wiesbaden, 17–17.
- [2] Paluck Deep, Monica Bharadhidasan, and A. Baki Kocaballi. 2025. "She was useful, but a bit too optimistic": Augmenting Design with Interactive Virtual Personas. *International Journal of Human-Computer Studies*: 103646.
- [3] Ali Farooq, Amani Alabed, Pilira Stella Msefula, et al. 2025. Representing groups of students as personas: A systematic review of persona creation, application, and trends in the educational domain. *Computers and Education Open* 8: 100242.
- [4] Jodi Forlizzi and John Zimmermann. 2025. Design Envoy. Medium. Retrieved October 3, 2025 from <https://medium.com/@forlizzi/design-envoy-bd8c2c342cba>.
- [5] Erin Friess. 2012. Personas and decision making in the design process: an ethnographic case study. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Association for Computing Machinery, 1209–1218.
- [6] (Eric) Heng Gu, Senthil Chandrasegaran, and Peter Lloyd. 2025. Synthetic users: insights from designers' interactions with persona-based chatbots. *AI EDAM* 39: e2.
- [7] Kathleen W. Guan, Joni Salminen, Lene Nielsen, Soon-Gyo Jung, and Bernard J. Jansen. 2021. Information Design for Personas in Four Professional Domains of User Experience Design, Healthcare, Market Research, and Social Media Strategy.
- [8] Shashank Gupta, Vaishnavi Shrivastava, Ameet Deshpande, et al. 2024. Bias Runs Deep: Implicit Reasoning Biases in Persona-Assigned LLMs. Retrieved February 14, 2025 from <http://arxiv.org/abs/2311.04892>.
- [9] Pei-Fang Hsu, Yu-Han Lu, Shih-Chu Chen, and Patricia Pei-Yi Kuo. 2024. Creating and validating predictive personas for target marketing. *International Journal of Human-Computer Studies* 181: 103147.
- [10] Hang Jiang, Xiajie Zhang, Xubo Cao, Cynthia Breazeal, Deb Roy, and Jad Kabbara. 2024. PersonalLLM: Investigating the Ability of Large Language Models to Express Personality Traits. Findings of the Association for Computational Linguistics: NAACL 2024, Association for Computational Linguistics, 3605–3627.
- [11] Ilkka Kaate, Joni Salminen, Soon-Gyo Jung, et al. 2025. "You Always Get an Answer": Analyzing Users' Interaction with AI-Generated Personas Given Unanswerable Questions and Risk of Hallucination. Proceedings of the 30th International Conference on Intelligent User Interfaces, ACM, 1624–1638.
- [12] Shivani Kapania, William Agnew, Motahhare Eslami, Hoda Heidari, and Sarah Fox. 2024. "Simulacrum of Stories": Examining Large Language Models as Qualitative Research Participants. Retrieved February 4, 2025 from <http://arxiv.org/abs/2409.19430>.
- [13] A. Baki Kocaballi. 2023. Conversational AI-Powered Design: ChatGPT as Designer, User, and Product. Retrieved January 21, 2025 from <http://arxiv.org/abs/2302.07406>.
- [14] Cynthia LeRouge, Jiao Ma, Sweta Sneha, and Kristin Tolle. 2013. User profiles and personas in the design and development of consumer health technologies. *International Journal of Medical Informatics* 82, 11: e251–e268.
- [15] Tara Matthews, Tejinder Judge, and Steve Whittaker. 2012. How do designers and user experience professionals actually perceive and use personas? Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Association for Computing Machinery, 1219–1228.
- [16] Mirjana Prpa, Sarah Fdili-Alaoui, Thecla Schiphorst, and Philippe Pasquier. 2020. Articulating Experience: Reflections from Experts Applying Micro-Phenomenology to Design Research in HCI. Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, Association for Computing Machinery, 1–14.
- [17] John Pruitt and Tamara Adlin. 2005. The Persona Lifecycle: Keeping People in Mind Throughout Product Design. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- [18] Kari Rönkkö, Mats Hellman, Britta Kilander, and Yvonne Dittrich. 2004. Personas is not applicable: local remedies interpreted in a wider context. Proceedings of the eighth conference on Participatory design: Artful integration: interweaving media, materials and practices - Volume 1, Association for Computing Machinery, 112–120.
- [19] Joni Salminen, Guan, Kathleen, Jung, Soon-Gyo, and Bernard J. and Jansen. 2021. A Survey of 15 Years of Data-Driven Persona Development. *International Journal of Human-Computer Interaction* 37, 18: 1685–1708.

- [20] Syed Ghulam Sarwar Shah and Ian Robinson. 2007. Benefits of and barriers to involving users in medical device technology development and evaluation. *International Journal of Technology Assessment in Health Care* 23, 1: 131–137.
- [21] Joongi Shin, Michael A. Hedderich, Bartłomiej Jakub Rey, Andrés Lucero, and Antti Oulasvirta. 2024. Understanding Human-AI Workflows for Generating Personas. *Proceedings of the 2024 ACM Designing Interactive Systems Conference, Association for Computing Machinery*, 757–781.
- [22] Lipepei Sun, Tianzi Qin, Anran Hu, *et al.* 2025. Persona-L has Entered the Chat: Leveraging LLMs and Ability-based Framework for Personas of People with Complex Needs. *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, Association for Computing Machinery*, 1–31.
- [23] Martin Tomitsch, Madeleine Borthwick, Naseem Ahmadvour, *et al.* 2021. *Design. Think. Make. Break. Repeat.: Revised edition.* BIS, Amsterdam.
- [24] Øivind Klungseth Zahlens, Elena Parmiggiani, and Yngve Dahl. 2023. Challenges of Scaling Participatory Design: A Systematic Literature Review. *Proceedings of the 34th Australian Conference on Human-Computer Interaction, Association for Computing Machinery*, 143–159.
- [25] Leijing Zhou, Yan Fang, Shiyong Ding, *et al.* 2024. Vivid-persona: customizable persona tool with interactive and immersive experiences. *Journal of Engineering Design* 0, 0: 1–22.