PARTICIPATE-AI: Exploring the Participatory Turn in Citizen-Centred AI

Pam Briggs
PaCT Lab
Northumbria University
Newcastle upon Tyne, United Kingdom
p.briggs@northumbria.ac.uk

Cristina Conati
Computer Science
University of British Columbia
Vancouver, British Columbia, Canada
conati@cs.ubc.ca

Shaun Lawson
Northumbria University
Newcastle upon Tyne, United Kingdom
shaun.lawson@northumbria.ac.uk

Kyle Montague
Northumbria University
Newcastle upon Tyne, Tyne and Wear
United Kingdom
kyle.montague@northumbria.ac.uk

Hugo Nicolau
ITI/LARSyS, Instituto Superior Technico
Universidade de Lisboa
Lisbon, Portugal
hugo.nicolau@tecnico.ulisboa.pt

Ana Cristina Pires
ITI/LARSyS Instituto Superior Technico
University of Lisbon
Lisbon, Portugal
ana.pires@iti.larsys.pt

Sebastian Stein
School of Electronics and Computer Science
University of Southampton
Southampton, Hampshire, United Kingdom
ss2@ecs.soton.ac.uk

John Vines
School of Informatics
University of Edinburgh
Edinburgh, United Kingdom
john.vines@ed.ac.uk

Abstract

This workshop explores value alignment for participatory AI, focusing on interfaces and tools that bridge citizen participation and technical development. As AI systems increasingly impact society, meaningful and actionable citizen input in their development becomes critical. However, current participatory approaches often fail to influence actual AI systems, with citizen values becoming trivialized during translation to technical requirements. The workshop brings together researchers and practitioners from HCI, AI, and social sciences to address challenges such as risk articulation, value evolution, democratic legitimacy, and the translation gap between community input and system implementation. Topics covered include interactive demonstrations of interfaces (or other 'boundary objects') for value elicitation within different communities, critical analysis of failed participatory attempts, methods for capturing evolving community values, and approaches for making citizen concerns actionable for developers.

CCS Concepts

• **Human-centered computing** → **Participatory design**; **Human computer interaction (HCI)**; • **Computing methodologies** → **Artificial intelligence**.

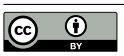

This work is licensed under a Creative Commons Attribution 4.0 International License.
IUI Companion '26, Paphos, Cyprus
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1985-1/26/03
https://doi.org/10.1145/3742414.3794954

Keywords

Participatory AI, Boundary Objects, Democratic AI, Value Alignment, Citizen-Centred AI

ACM Reference Format:
Pam Briggs, Cristina Conati, Shaun Lawson, Kyle Montague, Hugo Nicolau, Ana Cristina Pires, Sebastian Stein, and John Vines. 2026. PARTICIPATE-AI: Exploring the Participatory Turn in Citizen-Centred AI. In *Companion Proceedings of the 31st International Conference on Intelligent User Interfaces (IUI Companion '26), March 23–26, 2026, Paphos, Cyprus.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3742414.3794954

1 DESCRIPTION OF WORKSHOP

A new 'participatory turn' in AI has been generating significant interest [1–4]. This is essentially a more citizen-centred call for AI development that shifts power towards those who use or are affected by AI. This topic is particularly relevant to the IUI community, which has long championed user agency and meaningful human control in intelligent systems. Participatory AI assumes that citizen voices should be integrated as a central part of the development process, but this is a significant challenge. Existing work on Participatory ML has shown that diverse stakeholders can be engaged in the technical development of machine learning systems, but the elicited values are sometimes trivialised and disconnected from social norms [5]. Understanding how to create effective interfaces for participation is crucial as AI systems become more pervasive. A stronger citizen-centred focus would ensure that AI systems more deeply align with public values and serve broader democratic and societal goals. For participatory AI, both ML and

Citizen-Centric approaches should work in tandem so that citizen deliberation can be properly integrated into the development, deployment and subsequent evaluation of new systems. Yet this integration is no simple matter and remains a worrying translation gap between participatory initiatives and the technical development of AI systems, not least because the participatory work often relies on smaller-scale community-based methods that simply fail to impact the globalised operation of commercial AI systems [4]. This workshop directly addresses IUI's core mission of designing intelligent user interfaces by examining what kind of interface or 'boundary object' might sit at the intersection of public deliberation and AI development. Unlike the main conference's focus on individual user interactions, this workshop explores collective participation and community-level interaction with AI systems. We ask for contributions that help us understand (i) how we can best reconcile the values of AI developers and user communities and (ii) what methods might ensure the translation of these values into actionable tools that can have a genuine impact on AI systems development? By bringing together theoretical perspectives with interactive demonstrations of actual boundary objects, this full-day workshop offers IUI attendees practical insights into an emerging challenge that will shape the future of human-AI interaction.

2 PREVIOUS HISTORY

This workshop is a new proposal with no previous editions at IUI or elsewhere. It builds upon a special track on Citizen-Centric AI held at the ACM GoodIT conference in Lisbon 2023 (which was a paper track rather than a workshop, receiving 12 submissions with 8 accepted papers and approximately 40 attendees). This workshop differs from existing IUI workshop series in important ways. While HAI-GEN (2020-2025) focuses on individual creativity and co-creation with generative models, and SOCIALIZE examines personalisation for individual users, our workshop addresses collective participation and community-level engagement with AI systems. Unlike these domain-specific workshops, we explicitly tackle the challenge of translating between citizen communities and technical development, recognising that current participatory methods often fail to influence AI development meaningfully.

3 Organisers

Pam Briggs (Primary Contact) holds a Chair in Applied Psychology at Northumbria University and leads their Centre for Digital Citizens and is a co-investigator for their Centre for Doctoral Training in Citizen-Centred AI. She organised the Citizen-Centred AI track for GoodIT, 2023. https://www.northumbria.ac.uk/about-us/our-staff/b/pamela-briggs/. p.briggs@northumbria.ac.uk

Cristina Conati is a Professor of Computer Science at the University of British Columbia, Vancouver, Canada, where she researches human-centred AI and AI-driven personalisation. She has received 10 Best Paper Awards and the Test of Time Award 2022 from the Educational Data Mining Society, is a Fellow of AAAI and AAIA, an ACM Distinguished Member, and co-Editor in Chief of the Journal of AI in Education. She served as President of AAAC (Association for the Advancement of Affective Computing), and as Program or Conference Chair for several international conferences, including IUI 2009, 2016, and 2020. https://www.cs.ubc.ca/people/cristina-conati. conati@cs.ubc.ca

Shaun Lawson is Professor of Social Computing and Head of the Department of Computer & Information Sciences at Northumbria University. He is Director of the University's UKRI AI Centre for Doctoral Training in Citizen Centred AI. He organised the More Than Human Digital Civics Workshop at CHI 2024. https://www.northumbria.ac.uk/about-us/our-staff/l/shaun-lawson/. shaun.lawson@northumbria.ac.uk

Kyle Montague is Professor of Human-Computer Interaction at Northumbria University and Research Co-Lead of the Northumbria Centre for Responsible AI. He directs the Northumbria Social Computing (NorSC) research group, focusing on participatory design, digital civics, and inclusive technology development. His work emphasises democratising technology design through community engagement, with recent projects addressing responsible AI in law enforcement, AI literacy for vulnerable populations, and participatory platforms for public services. https://www.kylemontague.co.uk/. kyle.montague@northumbria.ac.uk

Hugo Nicolau is an Associate Professor at the University of Lisbon. He leads a research team on inclusive computing and is Vice-President of ITI/LARSyS research unit. He has served in multiple organising roles in top-tier conferences such as CHI, ASSETS, and DIS. He organised the Access InContext: Futuring Accessible Prototyping Tools and Methods Workshop at CHI 2025. https://web.tecnico.ulisboa.pt/hugo.nicolau/aboutme/. hugo.nicolau@tecnico.ulisboa.pt

Ana Cristina Pires is an integrated researcher at the Interactive Technologies Institute (ITI/LARSyS) and an Invited Assistant Professor at Instituto Superior Técnico, Universidade de Lisboa. Ana has a background in psychology and she explores citizen-centred and participatory design approaches to increase learning, inclusion, and collaboration, especially among populations at risk of social exclusion. https://iti.larsys.pt/member/ana-pires/. ana.pires@iti.larsys.com

Sebastian Stein is a Turing AI Acceleration Fellow at the University of Southampton, where he is Principal Investigator of the Citizen-Centric AI Systems project. With a background in multi-agent systems, his work focuses on the use of AI systems to solve urgent societal problems, for example, in smart mobility, smart energy and disaster response applications. http://southampton.ac.uk/people/5x5r59/professor-sebastian-stein. ss2@ecs.soton.ac.uk

John Vines holds a Chair in Design Informatics in the School of Informatics at the University of Edinburgh, co-directs the Institute for Design Informatics, and leads the Centre for Doctoral Training in Designing Responsible Natural Language Processing. His research focuses on how people experience and use digital and data-driven technologies in their everyday lives. He organised the Technology Mediated Caregiving For Older Adults Aging in Place Workshop for CHI 2025. https://www.designinformatics.org/person/john-vines/. john.vines@ed.ac.uk

Workshop programme committee: As above.

4 Participants

We anticipate 15–20 participants for this workshop, seeking a balanced mix of perspectives from the HCI, AI, design, and social science communities. Our recruitment strategy combines targeted

invitations with an open call to ensure diverse viewpoints on participatory AI. Participants will be recruited through multiple channels: *(1) We will issue a call for contributions through relevant mailing lists, (2) We will directly invite researchers leading on large-scale AI innovation projects,* and *(3) We will contact industry practitioners and policy makers working on AI governance and public engagement.*

5 Workshop Activities

This half-day workshop takes a mixed approach to exploring Participatory AI. We combine theoretical presentations with hands-on demonstrations. Our central focus is 'boundary objects'—the interfaces, tools, and methods that help citizens and AI developers work together. These objects act as translators, making technical systems accessible to citizens while making citizen values actionable for developers.

We begin with a brief **welcome and introduction** (15 minutes) to frame the workshop goals, followed by a **provocations and positions session** (60 minutes), where contributors present lightning talks of 3–5 minutes each, to surface key tensions around risk articulation, value alignment, platform integration, and democratic legitimacy. We then move to **thematic breakout discussions** (60 minutes). Participants self-organise around the challenges that resonate most: risk translation, evolving values, scale and integration, or democratic legitimacy. Each group examines real examples from the morning demonstrations. Following a coffee break, we hold a **collective synthesis session** (30 minutes) where each group shares their key insights. Together, we map the landscape of current approaches, identifying fundamental barriers and promising directions. The workshop concludes with a **wrap-up and next steps session** (15 minutes). We identify opportunities for collaboration and plan follow-up activities. This ensures momentum continues beyond the workshop itself.

Materials provided: We distribute all position papers before the event. During the workshop, we provide discussion templates and note-taking tools. Simple prototyping materials will be available for sketching ideas. All submissions and documentation will be available in a shared digital repository, which will be updated with key discussions and outcomes following the workshop.

6 Planned Outcomes of the Workshop

First, we will compile **a documented collection of boundary objects** with critical analysis. Rather than presenting these as solutions, we provide honest assessments of their contexts, assumptions, and limitations to create a realistic picture of the current state of the art. Each artefact will be photographed, described, and linked to the challenges it attempts to address. Second, we will generate **an agenda of open research questions**—3 to 5 core questions emerging from our discussions. These highlight where theoretical critiques meet practical attempts. They identify gaps where further research is most urgently needed. Key areas include *risk articulation, value evolution,* and *what makes participation meaningful versus tokenistic.* Third, we will produce **a mapping of current approaches to key challenges**. This visual or structured mapping shows which boundary objects address which problems. It reveals clusters of activity and notable gaps. More importantly, it highlights where multiple approaches tackle the same challenge differently, and where certain challenges remain largely unaddressed. Finally, we will create **an online repository** for continued community engagement. This contains all submissions, documented artefacts, discussion notes, and the mapping exercise. The repository remains open for others to build upon and extend. It serves as both an archive of the workshop and a starting point for future collaboration.

Acknowledgments

This work was supported by two awards from the Engineering and Physical Sciences Research Council: The Centre for Digital Citizens - Next Stage Digital Economy Centre; EPSRC (EP/Y009800/1) and the Centre for Doctoral Training in Citizen-Centred Artificial Intelligence (CCAI) (EP/Y030729/1); and the European Commission (101079116 Fostering Digital Civics Research and Innovation in Lisbon).

References

[1] Abeba Birhane, William Isaac, Vinodkumar Prabhakaran, Mark Diaz, Madeleine Clare Elish, Iason Gabriel, and Shakir Mohamed. 2022. Power to the people? Opportunities and challenges for participatory AI. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–8.

[2] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2023. The participatory turn in ai design: Theoretical foundations and the current state of practice. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–23.

[3] Michael Feffer, Michael Skirpan, Zachary Lipton, and Hoda Heidari. 2023. From preference elicitation to participatory ML: A critical survey & guidelines for future research. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*. 38–48.

[4] Meg Young, Upol Ehsan, Ranjit Singh, Emnet Tafesse, Michele Gilman, Christina Harrington, and Jacob Metcalf. 2024. Participation versus scale: Tensions in the practical demands on participatory AI. *First Monday* (2024).

[5] Tan Zhi-Xuan, Micah Carroll, Matija Franklin, and Hal Ashton. 2025. Beyond Preferences in AI Alignment: T. Zhi-Xuan et al. *Philosophical Studies* 182, 7 (2025), 1813–1863.