

Engineering Cross-Compatible Personalization for Interactive AI-Driven Services

Florian Müller
TU Darmstadt
Darmstadt, Germany
florian.mueller@tu-darmstadt.de

Andrii Matviienko
KTH Royal Institute of Technology
Stockholm, Sweden
matviienko.andrii@gmail.com

Alessandro Bozzon
Knowledge and Intelligence Design
Delft University of Technology
Delft, Netherlands
a.bozzon@tudelft.nl

Carmen Santoro
HIIS Lab
ISTI-CNR
Pisa, Italy
carmen.santoro@isti.cnr.it

Larbi Abdenebaoui
Society
OFFIS Institute for Information
Technology
Oldenburg, Germany
larbi.abdenebaoui@offis.de

Andrea Lavazza
Pegaso University and Centro
universitario internazionale
Naples, Italy
andrea.lavazza@unipv.it

Andreas Komninos
Computer Engineering & Informatics
University of Patras
Patras, Achaea, Greece
akomninos@ceid.upatras.gr

Nicolas Sklavos
Computer Engineering & Informatics
University of Patras
Patras, Greece
ngsklavos@gmail.com

Abstract

AI-driven services continuously infer “models-of-users” to personalize content, automate decisions, and optimize interaction. Yet, most user models remain opaque, fragmented across platforms, and hard to correct, transfer, or revoke. As a result, personalisation in AI-driven services is commonly experienced as a trade-off between privacy and platform lock-in. Users must stay loyal to platforms if a model of adequate utility is to be built around them. This tension is sharpened by emerging regulatory obligations around transparency and user oversight (e.g., GDPR and the EU AI Act). We propose a full-day EICS 2026 workshop that focuses on *engineering cross-compatible personalization*: a vision for future interactive systems in which users can make their digital self-representations *inspectable*, *editable*, and *selectively shareable*. The workshop brings together researchers and practitioners from interactive systems engineering, HCI, AI/ML, and security/privacy to (1) map the design and engineering space of user-sovereign personalization; (2) derive reusable artifacts for the EICS community, such as a reference architecture, protocol sketches, and a pattern language of interaction techniques; and (3) seed a community around building interoperable toolchains for AI personalization that is not only accurate, but also controllable, auditable, and context-adaptive.

Keywords

interactive systems engineering, personalization, human–AI interaction, user modeling, privacy engineering, interoperability, data sovereignty

ACM Reference Format:

Florian Müller, Andrii Matviienko, Alessandro Bozzon, Carmen Santoro, Larbi Abdenebaoui, Andrea Lavazza, Andreas Komninos, and Nicolas Sklavos. 2026. Engineering Cross-Compatible Personalization for Interactive AI-Driven Services. In *Companion Proceedings of the 18th ACM SIGCHI Symposium on Engineering Interactive Computing Systems (EICS Companion '26)*,

June 30–July 03, 2026, Patras, Greece. ACM, New York, NY, USA, 4 pages.
<https://doi.org/10.1145/3807968.3822520>

1 Background and workshop goals

Contemporary interactive systems increasingly embed AI components that personalize feeds, interfaces, and decisions [3, 4]. These systems build and maintain computational user representations (preferences, intents, inferred traits, risk scores) that are consequential yet difficult to access, contest, or move between services. Decades of work on *scrutable* adaptation and user-controllable personalization show both the promise, and the engineering difficulty, of making such representations inspectable and editable [3, 7]. Yet, many deployed transparency features still focus on *post-hoc* explanations rather than interaction mechanisms for *pre-emptive control* and negotiation [5, 9]. This situation raises an EICS-relevant engineering challenge: how do we design architectures, tools, and interaction techniques so that personalization becomes *negotiable*?

We frame negotiable personalization as a system property with three pillars:

(1) User-sovereign representations. Profiles are represented as scoped, composable objects (e.g., “at work”, “commuting”, “anonymous browsing”) with explicit lifetimes, provenance, and revocation semantics.

(2) Sandboxed interaction spaces. Building on work on scrutable user models and explainable AI user experiences [3, 5, 10] we propose private and secure safe data spaces, incorporating interfaces that let users inspect what the system “believes”, try changes via what-if previews, and apply edits (e.g., adjusting data scopes, preference weights, feature windows) with predictable downstream effects.

(3) Interoperable protocols and verifiable exchange. External services can request, negotiate, and verify only the necessary

facets of a profile (potentially via verifiable credentials and privacy-enhancing techniques), enabling portability without raw-data over-collection. This pillar interacts with established techniques such as interactive machine learning and explainable and contestable interfaces for user corrections [5] privacy-preserving learning (e.g., federated learning and differential privacy) [2, 6] and emerging work on personal data stores and self-sovereign identity infrastructures [8].

The workshop aims to:

- (1) **Consolidate** a shared vocabulary and design space for negotiable personalization across interaction design, AI engineering, and privacy/security.
- (2) **Produce reusable engineering outputs** for the EICS community: a reference architecture, interface pattern catalog, and “minimum viable” negotiation protocol sketch.
- (3) **Build a cross-community network** spanning multiple established communities in the area of interactive system engineering (e.g. EICS, IUI/CHI, UMAP, MobileHCI), usable security/privacy, and decentralized identity/data infrastructures.

Scope and topics. We solicit contributions on the following (non-exhaustive) list of topics:

- representations for user-controlled personalization (schemas, modular profiles, local adapters);
- interaction techniques for inspecting/editing user models (dashboards, timelines, profile composition tools, conversational negotiation);
- negotiation workflows and consent mechanisms (time-boxing, selective disclosure, revocation, audit trails);
- architectures for user-side AI (on-device/federated learning, privacy-preserving exchange);
- verifiability and accountability (credentials, provenance, logging, evaluation of trust and UX);
- empirical methods and benchmarks for measuring agency, predictability, and utility.

2 Target audience

The workshop targets approximately **20–35 participants** from: (1) EICS researchers working on UI engineering, model-based design, software architectures, and tool support for interactive systems with AI components; (2) HCI and Human–AI Interaction researchers investigating transparency, explainability-as-interaction, and user control; (3) privacy/security and decentralized systems researchers working on personal data stores, identity, and secure computation; and (4) practitioners building personalization stacks, recommender systems, or privacy tooling. We expect a mixed audience, with introductory material in the morning and advanced breakouts in the afternoon. Participants should have familiarity with interactive systems design/engineering; no deep cryptography background is assumed.

3 Pre-workshop plans

Recruiting and community building. Upon acceptance, we will launch a workshop website (hosted by the organizers) featuring the call for participation, program committee, topics, and submission

instructions. We will promote the workshop through: EICS, CHI and IUI community mailing lists; UMAP and RecSys communities; usable security/privacy venues (e.g., SOUPS/PETS); and decentralized identity/data communities (e.g., Solid, W3C VC). We will also directly invite authors of recent papers in the area and participants of related EICS workshops on AI-embedded interactive systems.

Submission types. We invite the following submissions:

- **Position papers** (2–4 pages in ACM format, excluding references) describing research results, open problems, or system visions.
- **Artifact sketches** (1 page) describing an interface concept, protocol/API sketch, or architecture diagram.
- **Optional demo material** (short video or repository link) to support discussion.

All accepted contributions will be presented as *lightning talks* and discussed in breakouts.

Indicative timeline (aligned with EICS workshop schedule).

- **Call for participation released:** mid March 2026.
- **Submission deadline:** 25 April 2026.
- **Notification to authors:** 2 May 2026.
- **Camera-ready for workshop website:** 20 May 2026.

4 Workshop structure and format

We propose a **full-day workshop** combining paper discussion with collaborative engineering activities.

4.1 Overview

Morning: shared framing and rapid exchange (keynote-style provocations + lightning talks). **Afternoon:** hands-on breakouts to create actionable outputs (patterns, reference architecture, protocol sketch), with a synthesis session.

4.2 Detailed schedule (draft)

Breakout deliverables. Each breakout produces a tangible output:

- **Interaction pattern cards** (problem, context, UI components, engineering hooks);
- **Protocol sketch** for negotiating a profile facet (request, consent, expiry, audit);
- **Reference architecture** showing user-side components (profile store, local model, mirror UI, connectors).

All outputs will be consolidated into a public repository.

5 Operational requirements

We request: (1) a room with movable tables to support three parallel breakout groups; (2) projector + audio; (3) whiteboards/flip charts and sticky notes; (4) reliable Wi-Fi and sufficient power outlets for laptops. No specialized hardware is required.

6 Advertising plan

The workshop website will include accepted papers, schedule, organizers, and participation instructions. We will advertise through: EICS channels; relevant mailing lists (UMAP, RecSys, CHI, IUI,

| Time | Activity |
|-------------|--|
| 09:00–09:15 | Welcome, goals, and warm-up poll (participants' domains and "pain points"). |
| 09:15–09:45 | Invited provocation #1: <i>From explanation to negotiation</i> . |
| 09:45–10:30 | Lightning talks I (5 min each) + structured Q&A. |
| 10:30–11:00 | Coffee break + poster-style mingling. |
| 11:00–11:30 | Invited provocation #2: <i>Interoperability and verifiable exchange for user models</i> . |
| 11:30–12:00 | Lightning talks II + discussion. |
| 12:00–12:30 | Panel (legal/ethical perspectives): <i>Rights, accountability, and governance constraints for negotiable personalization</i> . |
| 12:30–13:00 | Lunch. |
| 13:00–14:15 | Breakout formation: three tracks with canvases (Interaction Patterns / Protocols / Architectures). |
| 14:15–15:30 | Breakout work session I (produce first drafts of artifacts). |
| 15:30–16:00 | Coffee break. |
| 16:00–16:45 | Breakout work session II (refine and cross-review artifacts). |
| 16:45–17:30 | Plenary synthesis: merge outputs into a shared roadmap; define next steps. |
| 17:30–17:45 | Closing: commitments, publication plan, and community channel. |

Table 1: Draft schedule for the full-day workshop.

SOUPS/PETS); social media (LinkedIn/Mastodon); and targeted invitations to authors of recent work in personal data stores, decentralized identity, and human-centered AI. We will also cross-advertise with prior EICS workshops on AI-embedded interactive systems to foster continuity while offering a distinct focus on user-sovereign personalization.

7 Post-workshop plans

Proceedings and dissemination. In line with EICS workshop plans, we will invite authors to revise their submissions after the workshop discussion for inclusion in the joint LNCS workshop volume. We will additionally publish: (1) a workshop report summarizing insights and open challenges; (2) the artifact repository; and (3) a short "EICS poster" summarizing results to be presented at the main conference if the option is available.

Sustaining the community. We will set up a lightweight community channel (e.g., mailing list and/or Slack/Discord) and plan follow-up activities: a special session at a related venue (EICS/CHI/UMAP), and an interoperability "challenge" where participants implement a minimal connector between a mock service and a user-side profile store.

8 Organizers and expertise

The organizing team combines expertise in engineering interactive systems, human-centered AI/personalization, and socio-technical and ethical perspectives on user agency.

Florian Müller (TU Darmstadt). Assistant Professor for Mobile HCI at TU Darmstadt; works on mobile/XR interaction and human-AI collaboration in real-world contexts; served in key conference roles and has experience in organizing workshops at EICS [1].

Andrii Matviienko (KTH Royal Institute of Technology). Assistant Professor in HCI at KTH leading immersive technologies lab. He researches interaction techniques and prototyping for XR and context-aware systems; served in key conference roles including MUX 2022, CHI 2023-2024, MUM 2024 and more.

Alessandro Bozzon (TU Delft). Professor of Human-Centered AI at TU Delft; expertise in user modelling, recommender systems, and scalable personalization pipelines bridging AI and interactive systems engineering.

Carmen Santoro (CNR-ISTI). Researcher at CNR-ISTI with long-standing expertise in end-user development and interactive system engineering, including tools and methods that make complex functionality configurable, inspectable, and adaptable by non-experts.

Larbi Abdenebaoui (OFFIS). Researcher at OFFIS working at the intersection of human-computer interaction and artificial intelligence, with a focus on explainable AI and value-driven design. His work explores how interactive intelligent systems can be engineered to be transparent, accountable, and aligned with human values in real-world applications.

Andrea Lavazza (Pegaso University). Associate Professor of Moral Philosophy and Neuroethics; researches AI ethics and neuroethics; supports the workshop's normative framing and responsible-design outcomes.

Andreas Komninos (University of Patras). Assistant Professor at the University of Patras focusing on mobile and pervasive interactive systems (incl. context-aware services and IoT); active in SIGCHI community-building through Global Chapters and MedCHI service.

Nicolas Sklavos (University of Patras). Professor at the University of Patras; expertise in cybersecurity, cryptography, hardware security, embedded systems, IOT; active as ACM Distinguished Speaker of Security and Privacy.

References

- [1] Damiano Distante, Marco Winckler, Regina Bernhaupt, Judy Bowen, José Creissac Campos, Florian Müller, Philippe Palanque, Jan Van den Bergh, Benjamin Weyers, and Alexandra Voit. 2019. Trends on engineering interactive systems: an overview of works presented in workshops at EICS 2019. In *Proceedings of the ACM SIGCHI Symposium on Engineering Interactive Computing Systems* (Valencia, Spain) (EICS '19). Association for Computing Machinery, New York, NY, USA, Article 22, 6 pages. doi:10.1145/3319499.3335655
- [2] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Found. Trends Theor. Comput. Sci.* 9, 3–4 (Aug. 2014), 211–407. doi:10.1561/04000000042
- [3] Judy Kay. 2006. Scrutable adaptation: because we can and must. In *Proceedings of the 4th International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*. Springer-Verlag, Berlin, Heidelberg, 11–19. doi:10.1007/11768012_2
- [4] Alfred Kobsa. 2007. Privacy-enhanced personalization. *Commun. ACM* 50, 8 (Aug. 2007), 24–33. doi:10.1145/1278201.1278202

- [5] Q. Vera Liao, Daniel Gruen, and Sarah Miller. 2020. Questioning the AI: Informing Design Practices for Explainable AI User Experiences. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3313831.3376590
- [6] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*. PMLR, 1273–1282.
- [7] Filip Radlinski, Krisztián Balog, Fernando Diaz, Lucas Dixon, and Ben Wedin. 2022. On Natural Language User Profiles for Transparent and Scrutable Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Madrid, Spain) (SIGIR '22). Association for Computing Machinery, New York, NY, USA, 2863–2874. doi:10.1145/3477495.3531873
- [8] Frederico Schardong and Ricardo Custódio. 2022. Self-Sovereign Identity: A Systematic Review, Mapping and Taxonomy. *Sensors* 22, 15 (2022). doi:10.3390/s22155641
- [9] Nava Tintarev and Judith Masthoff. 2007. A Survey of Explanations in Recommender Systems. (2007), 801–810. doi:10.1109/ICDEW.2007.4401070
- [10] Huan Wang and Andrii Matviienko. 2025. Experiencing Art Museum with a Generative Artificial Intelligence Chatbot. In *Proceedings of the 2025 ACM International Conference on Interactive Media Experiences (IMX '25)*. Association for Computing Machinery, New York, NY, USA, 430–436. doi:10.1145/3706370.3731650