The Personalized Computer for the 21st Century

Helena Vasconcelos*
hvasconcelos@g.harvard.edu
Harvard University
Massachusetts, USA

Dora Zhao*
dorothyz@stanford.edu
Stanford University
California, USA

Michelle S. Lam
mslam@utexas.edu
University of Texas at Austin
Texas, USA

Dylan Wootton
dwootton@mit.edu
Massachusetts Institute of Technology
Massachusetts, USA

Omar Shaikh
oshaikh@stanford.edu
Stanford University
California, USA

Andy Matuschak
andy@andymatuschak.org
Independent
California, USA

Mitchell Gordon
mlgordon@mit.edu
Massachusetts Institute of Technology & OpenAI
Massachusetts, USA

Michael S. Bernstein
msb@cs.stanford.edu
Stanford University
California, USA

Abstract

We define personalized computing as a class of systems that both *know* about you deeply and can *act* on that information to serve you better. Longstanding work in HCI has built up systems that know about the user, but that knowledge remains largely inert and siloed. For example, our desktop computers, browsers, and web applications minimally adapt to users with ranking and recommendation but generally wait on a user's directive to take action. In contrast, modern AI systems now have the ability to not only consume massive amounts of unstructured data to better know the user, but also *take intelligent, proactive action on this information*. This development has enabled systems that observe users and learn strong user models that can power personalized AI assistance, bespoke tools and interfaces, and personal agents acting on the user's behalf. With these shifts towards a realizable paradigm of personalized computing, we believe our community ought to discuss and envision: what will interaction look like when our computers know more about us than we know about ourselves, and can take meaningful action for us? Will the design of such interactive systems require manual authoring and handcrafted abstractions, or will they succeed only by accruing ever more user data? In an interaction paradigm where information is fundamentally offloaded to AI, how do we ensure that these systems preserve user privacy, agency, control, and understanding? Through this workshop, we aim to explore these questions and advance the development of personalized computing systems.

*Both authors contributed equally.

UIST '26, Woodstock, NY

ACM ISBN 978-1-4503-XXXX-X/2018/06
https://doi.org/XXXXXXX.XXXXXXX

CCS Concepts

• **Human-centered computing** → **Interaction design**; **Interaction paradigms**; • **Computing methodologies** → **Artificial intelligence**.

Keywords

Personalization, HCI, Personal Agents

ACM Reference Format:
Helena Vasconcelos, Dora Zhao, Michelle S. Lam, Dylan Wootton, Omar Shaikh, Andy Matuschak, Mitchell Gordon, and Michael S. Bernstein. 2018. The Personalized Computer for the 21st Century. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (UIST '26)*. ACM, New York, NY, USA, 3 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

A long-standing vision in computing has been to create truly *personalized* computing systems. For example, in 1987, Apple introduced the "Knowledge Navigator," which envisioned a computer that could seek out relevant information based on the user's context. Microsoft's Clippy similarly aimed to anticipate what users needed help with and provide assistance. These early systems share an underlying ethos that while personal computers are owned by a user and hold their data, we ought to pursue *personalized* computers that proactively learn about the user and initiate actions on their behalf. The primary interest of our workshop is this *personalized computing* paradigm: computing systems powered by intelligence that both knows the user and acts on their needs.

Within the human-computer interaction (HCI) community more specifically, creating personalized computing systems has long been a central research goal. Early work explored programming by demonstration to enable end users to create bespoke software through example rather than code [6, 7]. Systems such as Chickenfoot [2] and Greasemonkey [1] allowed end users to customize Internet applications. While these systems required manual customization, others envisioned intelligent agents that could "reduce work and information overload" by acting on users' behalf [13],

prompting debates on whether direct manipulation or personal agents are the best interaction paradigm [17]. Adaptive interface systems such as SUPPLE demonstrated how interfaces could be automatically tailored to users' abilities and preferences [10].

Despite the continued interest in personalized computing, technical limitations stalled many of these visions from realization. Personalization either had to be anticipated in advance through carefully designed abstractions (e.g., computer terminals that are inherently customizable) or predicted with stochastic, error-prone machine learning (ML) algorithms that left users frustrated (e.g., Clippy).

However, advances in AI capabilities are ushering in a new wave of personalized computing systems. These advances are redefining what we can *know* about users. Rather than requiring narrow, hand-tested domains, recent work [8, 15, 16, 21] shows that we can infer general-purpose models of user intent across arbitrary contexts, without needing to be explicitly predefined for any one of them. There are also new ways to *act* upon information about users. For example, natural language prompts and agentic coding allow users to easily build custom tools or adapt interfaces for their tasks [3, 12, 19]. Or, rather than leaving users to act on this understanding themselves, systems can intervene on their behalf, through providing contextual nudges, generating new interfaces, or taking autonomous actions [4, 5, 9, 20].

At this inflection point, our workshop's aim is to chart the driving visions for what personalized computing ought to be, in light of these new technical capabilities. Our workshop centers on the following three questions. First, assuming we have a rich understanding of users, what interactions should we build, and why? Second, *how* should personalized computing systems be built: when do we need pre-specified representations of tasks, contexts, and users, and when does collecting more data about users provide more leverage? Third, what are the ethical considerations for designing these systems? For example, are there inferences about the user that should *not* be made even if they are accurate and useful [11, 14]? And, how do we delineate between a system that is appropriately proactive versus one that infringes on our agency?

This workshop will bring together researchers and practitioners excited about ushering in the future of personalized computing systems. We start with keynote talks and reflect on how we might revive, revise, or replace visions for personalized computing systems. Then, we will host a debate on the methods for constructing these systems. Attendees will develop prototypes of personal computing systems as instantiations of nascent ideas developed in the workshop and provocations for the group. Lastly, we will facilitate an open discussion on the ethical implications of these technologies.

2 Workshop Plan and Goals

Our workshop is a half-day, in-person event aiming to engage 25-30 participants. Prior to the workshop, participants are asked to complete preparatory activities. During the workshop, participants will hear from experts, collaborate to build demos reflecting workshop discussions, and engage in conversation on ethical implications of this field.

2.1 Preparatory Activities

Interested participants will apply via an online form where they will submit a brief description of their research experience, motivation for joining the workshop, and potential ideas for the workshop hackathon. One week before the workshop, we will circulate a slide deck for participants to upload an introduction slide. We will also share a spreadsheet of submitted demo ideas, each with the submitter's name and contact information, to encourage participants with similar interests to discuss prototype ideas ahead of the hackathon.

2.2 Workshop Activities

2.2.1 Opening Remarks and Introduction (5-10 minutes). We will start the workshop with brief opening remarks for the organizing team to communicate the goals and agenda for the day.

2.2.2 Spotlight Talks: Past and Present Visions (30 min). We will have three spotlight talks from experts in academic and industry who have experience building personalized computer systems. The spotlight talks are intended to both revisit earlier visions of personal computing systems and discuss emerging ones. We will ask workshop participants if they are interested in giving a spotlight talk and select three to highlight during the workshop. Talks will run 10 minutes each, and questions will be reserved for during the coffee break.

2.2.3 Debate: How Does the "Bitter Lesson" Apply to Personal Computing? (45-60 min). We will feature a debate on *how* the personalized computing systems ought to be built, covering how far the "bitter lesson"—that general-purpose methods that scale with computation ultimately outperform hand-crafted features—extends to personalized computing [18]. The debate centers on a core tension: is the best path to a deep understanding of users to simply collect more data and learn a general model of them? Or do intentionally designed abstractions and representations of users and tasks remain the right approach? We have invite 1-2 speakers to seed both sides of the debate. Debaters are encouraged to stake out the strongest form of each position.

2.2.4 Interactive Activity: Prototyping Provocative Demos (120 min). In the second half of the workshop, we will host an interactive activity, during which participants can ideate systems that demonstrate their vision for personalized computing. Participants are encouraged to create artifacts to help demonstrate their proposed systems, including paper prototypes, mock-ups, and figures. We will encourage participants to form teams of two to four, and will facilitate grouping by sharing out the ideas participants submitted in their applications ahead of the workshop. Each group will present their proposed system and any accompanying visuals in a lightning round.

2.3 Follow-Up

Following the workshop, we will showcase demos, prototypes, and opinion pieces from participants on our website and social media. We will also prepare a report summarizing the outcomes and future directions that will be shared as a blog post.

References

[1] Greasemonkey. https://addons.mozilla.org/en-US/firefox/addon/greasemonkey/.

[2] Bolin, M., Webber, M., Rha, P., Wilson, T., and Miller, R. C. Automation and customization of rendered web pages. In *Proceedings of the 18th annual ACM symposium on User interface software and technology* (2005), pp. 163–172.

[3] Cao, Y., Jiang, P., and Xia, H. Generative and malleable user interfaces with generative and evolving task-driven data model. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (2025), pp. 1–20.

[4] Chen, J., Zhang, Y., Zhang, Y., Shao, Y., and Yang, D. Generative interfaces for language models. In *Findings of the Association for Computational Linguistics: ACL 2026* (2026), pp. 1499–1519.

[5] Cheng, R., Barik, T., Leung, A., Hohman, F., and Nichols, J. Biscuit: Scaffolding llm-generated code with ephemeral uis in computational notebooks. In *2024 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)* (2024), IEEE, pp. 13–23.

[6] Cypher, A. Eager: Programming repetitive tasks by example. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (1991), pp. 33–39.

[7] Cypher, A., and Halbert, D. C. *Watch what I do: programming by demonstration*. MIT press, 1993.

[8] Danry, V., Billa, J. G., Samaradivakara, Y., Liang, P. P., and Maes, P. Mind mapper: Modeling and predicting behavioral patterns from everyday conversations with wearable ai systems and llms. In *Proceedings of the 31st International Conference on Intelligent User Interfaces* (2026), pp. 2059–2083.

[9] Fang, C. M., Samaradivakara, Y., Maes, P., and Nanayakkara, S. Mirai: A wearable proactive ai" inner-voice" for contextual nudging. In *Proceedings of the extended abstracts of the CHI conference on human factors in computing systems* (2025), pp. 1–9.

[10] Gajos, K., and Weld, D. S. Supple: automatically generating user interfaces. In *Proceedings of the 9th international conference on Intelligent user interfaces* (2004), pp. 93–100.

[11] Klein, E. I saw something new in san francisco. *The New York Times* (Mar. 2026).

[12] Lam, M. S., Shaikh, O., Xu, H., Guo, A., Yang, D., Heer, J., Landay, J. A., and Bernstein, M. S. Just-in-time objectives: A general approach for specialized ai interactions. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems* (2026), pp. 1–26.

[13] Maes, P. Modeling adaptive autonomous agents. *Artificial life 1*, 1_2 (1993), 135–162.

[14] Monteiro, K., Park, M., Ioffrida, A., Sanna, A., Mireshghallah, N., Wang, Y., Das, S., et al. When are llm inferences acceptable? user reactions and control preferences for inferred personal information. *arXiv preprint arXiv:2605.10013* (2026).

[15] Sapkota, S., Jörke, M., Sabbagh, Z., Shaikh, O., Wang, G., and Landay, J. A. " what are you really trying to do?": Co-creating life goals from everyday computer use. *arXiv preprint arXiv:2605.00497* (2026).

[16] Shaikh, O., Sapkota, S., Rizvi, S., Horvitz, E., Park, J. S., Yang, D., and Bernstein, M. S. Creating general user models from computer use. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology* (2025), pp. 1–23.

[17] Shneiderman, B., and Maes, P. Direct manipulation vs. interface agents. *interactions 4*, 6 (1997), 42–61.

[18] Sutton, R. S. The bitter lesson, 3 2019.

[19] Vaithilingam, P., Glassman, E. L., Inala, J. P., and Wang, C. Dynavis: Dynamically synthesized ui widgets for visualization editing. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (2024), pp. 1–17.

[20] Yan, L., Head, A., Milne, K., Le, V., Gulwani, S., Parnin, C., and Murphy-Hill, E. The invisible mentor: Inferring user actions from screen recordings to recommend better workflows. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems* (2026), pp. 1–17.

[21] Zhao, D., Lam, M. S., Yang, D., and Bernstein, M. S. Behavior latticing: Inferring user motivations from unstructured interactions. *arXiv preprint arXiv:2604.07629* (2026).