Human-Computer Interaction and AI

What practitioners need to know to design and build effective AI system from a human perspective

Daniel M. Russell*
Principal User Experience Researcher, Google
drussell@google.com

Q. Vera Liao
Principal Researcher, Microsoft Research
veraliao@microsoft.com

Chinmay Kulkarni
Associate Professor, Emory University
Chinmay.kulkarni@emory.edu

Elena Glassman
Assistant Professor, Harvard University.
glassman@seas.harvard.edu

Nikolas Martelaro
Assistant Professor, Carnegie Mellon University
nikmart@cmu.edu

ABSTRACT

AI and ML are now essential parts of many systems that are currently being built. What should CHI practitioners know about the possibilities and potential drawbacks of building AI systems? Understanding the human side of AI/ML based systems requires understanding both how the system-side AI works, but also how people think about, understand, and use AI tools and systems. This course will cover what AI components and systems currently exist, how to design and build usable systems with AI components, along with how the mental models of AI/ML tools operate. These models lead to user expectations of how AI systems function, and ultimately, to design guidelines that avoid disappointing end-users by accidentally creating unintelligible AI tools. We'll also cover the ethics of AI, including data collection, algorithmic and data fairness considerations, along with other risks of AI.

CCS CONCEPTS

• **Human-centered computing**; • **Human computer interaction (HCI)**; • **HCI theory, concepts and models**; • **Artificial Intelligence**; • **Machine Learning**; • **User Interface Design**;

KEYWORDS

HCI, HAI, UI design for AI systems, human-in-the-loop, AI trust, AI fairness

ACM Reference Format:
Daniel M. Russell, Q. Vera Liao, Chinmay Kulkarni, Elena Glassman, and Nikolas Martelaro. 2023. Human-Computer Interaction and AI: What practitioners need to know to design and build effective AI system from a human perspective. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23), April 23–28, 2023, Hamburg, Germany*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3544549.3574170

*Corresponding author.

CHI EA '23, April 23–28, 2023, Hamburg, Germany

ACM ISBN 978-1-4503-9422-2/23/04.
https://doi.org/10.1145/3544549.3574170

1 BENEFITS

As we design and build increasingly sophisticated systems, they're going to have large pieces that will be AI/ML powered. This course is about how we can take on the challenge of building systems that are AI-based, yet work in ways that are understandable by humans. We'll take on issues of capabilities, fairness, interpretability, and accountability. We'll come to learn about what AI systems can (and cannot) do, about what kinds of mental models people have about such systems, and what we can do to design a user experience to make these systems comprehensible. Ultimately, this is a class about the intersection of human intelligence with artificial intelligence—the two don't necessarily fit well together, and each makes unexpected demands on the other. As we design and build out AI-based systems, we will need to have our own deep understanding of the materials of AI, and understand what's possible.

2 INTENDED AUDIENCES

The course is fairly broad, meaning that early-career HCI practitioners (grad students) and seasoned UX designers, researchers, and engineers can all find value in this course.

3 PREREQUISITES

No prerequisites.

4 CONTENT

This team of HCI/AI practitioners will bring together theory and practice from both the academic and industrial perspectives to give students a deep understanding of what AI technology offers to designers and system builders. [1] The instructors have a broad background in building AI systems, bringing their decades of UX experience to the task. Students will learn about the current state of the art for building UIs that use AI, how to evaluate them, ways in which things can go wrong, and how to make them understandable by those that will use them.

This team of HCI/AI practitioners brings together theory and practice from both the academic and industrial perspectives to give students a deep understanding of what AI technology offers to designers and system builders. We will discuss current technical challenges such as AI system hallucinations [3], explanability [2], and trust in using AI tools [4]. The instructors have a broad background in building AI systems, bringing their decades of UX

experience to the task. Students will learn about the current state of the art for building UIs that use AI, how to evaluate them, ways in which things can go wrong, and how to make them understandable by those that will use them.

Our plan for the course is to cover these areas in four 75 minute sections:

* The human aspects of designing and building AI/ML systems - practice and theory
- How people understand AI systems: AI and mental models
- examples/cases of AI system UX
* Designing for AI failures and Feedback to Users
- guardrails and failure modes
- details on what has worked, what hasn't worked, and why
* Data, Knowledge, Fairness, and Ethics
- AI Ethics of Actions, Fairness, Social Acceptability, and Trust
- analysis methods to understand HAI data with respect to fairness
* Interpreting and Explaining AI Algorithms and Systems
- Building AI/ML with humans in the loop / AI in the loop
- understanding spoken language; written language;
- generating language; conversations; large language models
* Computer perception: recognition, classification, uses (and misuses)
* AI & Art: Synthesis systems for creativity, music, imagery
- prompt engineering (what it is and how to do it)
* Where does the future of AI/ML and HCI lead?
- HAI: how to design and build real systems
- Or is AI/ML an existential crisis?

5 PRACTICAL WORK

The course will include 4 different in-class exercises with existing AI tools. Materials for the exercises will be made available in the class via a public class website. Participants should plan on bringing a laptop to the class (or be willing to partner with someone who has one).

6 INSTRUCTOR BACKGROUNDS

Daniel M. Russell is a senior research scientist and 17-year veteran at Google. Within Google, Dan conducts research on understanding how people collect, organize and understand large amounts of information when they search the web. He has been involved in many decisions about the UX for Google search. As an individual contributor, Dan is best known for his studies of the human-computer interaction models of people's sensemaking behavior. He also teaches an online Massive Open Online Course (MOOC) PowerSearchingWithGoogle which has helped more than 4.5 million students become expert searchers.

Before joining Google, Dan held research positions at IBM's Research Almaden Research Center (San Jose, CA), Apple's Advanced Technology Group (ATG) and Xerox PARC. Dan has also been an adjunct lecturer in computer science at the University of Santa Clara, Stanford University, the University of Zürich, and is currently an adjunct faculty member at Stanford, the University of Maryland, College Park, and the University of Colorado, Boulder.

Dan is a Fellow of the ACM Computer Human-Interaction Society. He is also an international speaker, giving keynotes and invited lectures at conferences and universities around the world. His latest book is The Joy of Search: An Insider's Guide to Going Beyond the Basics, available from MIT Press and fine bookstores everywhere.

Dan is teaching the "HCI & AI/ML" (CS339Hsyllabus) course at Stanford (fall term, 2022).

Q. Vera Liao is a Principal Researcher at Microsoft Research Montréal, where she is part of the FATE (Fairness, Accountability, Transparency, and Ethics of AI) group. Her current research interests are in human-AI interaction, explainable AI, and responsible AI. She has taught several courses at HCI and AI conferences, including Intro to Explainable AI (CHI2021), Human-Centered Evaluation of AI Explanations (NAACL 2022), and Algorithmic Discrimination at the Intersection (NeurIPS 2022). She currently serves as the Co-Editor-in-Chief for Springer HCI Book Series, and in the Editors team for ACM CSCW conferences. Prior to joining MSR, she worked at IBM T.J. Watson Research Center, and studied at the University of Illinois at Urbana-Champaign.

Elena L. Glassman is an Assistant Professor at Harvard University's John A. Paulson School Of Engineering And Applied Sciences. She specializes in human-computer interaction and human-AI interaction, teaching graduate and undergraduate CS classes in introductory and advanced HCI as well as leading a research group focused on developing interfaces that augment human sensemaking and decision-making about piles of complex data. At Harvard she has specifically taught Research Topics in HCI (with a strong focus on human-AI interaction–syllabus), Design of Useful and Usable Interactive Systems, and Engineering Usable Interactive Systems. Prior to that, at MIT during her final year as a PhD student, she was co-instructor for a large undergraduate UI design class. Her undergraduate degree is in EE and her Master's is in EECS, also both from MIT; her postdoctoral training was at the Berkeley Institute of Design (BiD). She has served for many years on the CHI PC's Engineering Interactive Systems & Tools subcommittee, as well as on the UIST and CSCW PCs in years past. Her publications and talks are available on her labwebsite.

Chinmay Kulkarni is an Assistant Professor of Human Computer Interaction at Emory University, where he directs the Expertise@Scale lab. His work is at the intersection of large-data and automation, learning, and the future of work. In his research, he investigates how the new affordances of scale and automation can create new opportunities for people to learn and work more effectively. Typically, this research introduces novel computer systems that demonstrate benefits at large scale that are otherwise not achievable. Research from his group has resulted in systems that have been used by more than 100,000 in 150 countries; and has been adopted by companies such as Coursera and edX.

Before joining Emory University, Chinmay earned a PhD in Computer Science from Stanford University, where his thesis won the Arthur P Samuel Award. He has also previously worked at Microsoft Research, the Barcelona Supercomputing Center, and Carnegie-Mellon University.

Nikolas Martelaro is an Assistant Professor of Human-Computer Interaction at Carnegie Mellon University where he directs the Augmented Design Capability Studio. His work focuses on understanding how designers work and looks to develop new computational tools to support them in doing their work better.

Most recently, Nik and his students have been exploring how designers learn to collaborate with generative AI systems for mechanical and industrial design, what AI image generators know about architecture and product design, and how people learn to do prompt engineering use generative image models to support their work practices. Nik also explores how people interact with autonomous physical systems, such as sidewalk robots and autonomous vehicles, with the goal of helping interaction designers rapidly prototype new system behaviors by using Wizard-of-Oz techniques in contextually situated scenarios. Finally, the research group also explores how to help designers and engineers proactively think about the risks and hazards of AI and autonomy through though-provoking and playful activities such as card game named "What Could Go Wrong?"

Before joining Carnegie Mellon University, Nik was a Technology R&D Principal at Accenture Technology Lab's Digital Experience Group. He earned his Ph.D. in Mechanical Engineering from Stanford University's Center for Design Research.

7 RESOURCES

A web site will be created to host the presentation materials, notes, and reference lists for the class.

8 ACCESSIBILITY

All of the materials will be available to the students with appropriate accessibility features (e.g., ALT tags).

REFERENCES

[1] Shneiderman, B. *Human-Centered AI.* Oxford University Press, 2022.
[2] Ehsan, Upol, Liao, Q. V., Muller, M., Riedl, M. O., & Weisz, J. D. "Expanding explainability: Towards social transparency in ai systems." Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. 2021. James W. Demmel, Yozo Hida, William Kahan, Xiaoye S. Li, Soni Mukherjee, and Jason Riedy. 2005. Error Bounds from Extra Precise Iterative Refinement. Technical Report No. UCB/CSD-04-1344. University of California, Berkeley.
[3] Wu, Tongshuang, Michael Terry, and Carrie Jun Cai. "Ai chains: Transparent and controllable human-ai interaction by chaining large language model prompts." CHI Conference on Human Factors in Computing Systems. 2022.
[4] Cai, C. J., Winter, S., Steiner, D., Wilcox, L., & Terry, M. (2019). " Hello AI": uncovering the onboarding needs of medical practitioners for human-AI collaborative decision-making. Proceedings of the ACM on Human-computer Interaction, 3(CSCW), 1-24.