How Can Human-Centered Design Shape Data-Centric AI?

Hariharan Subramonyam
Stanford University
Stanford, CA

Colleen Seifert
University of Michigan
Ann Arbor, MI

Eytan Adar
University of Michigan
Ann Arbor, MI

Abstract

As machine learning architectures stabilize, there is a methodological shift towards *data-centric AI* (DCAI) —designing training data characteristics while keeping the model constant to achieve desired behavior and performance. We argue that this shift is a promising path forward to realizing human-centered AI. Based on qualitative inquiry (interviews and in-lab co-design studies) with industry practitioners, we find that data is under-emphasized in current AI development practices and is optimized for engineering tasks rather than end-users. Insights from our studies show that HCI practitioners leverage well-established user research and design techniques to anchor AI development around human needs. End-user data in the form of design probes serve as the *lingua-franca* for HCI and AI practitioners to collaborate on system design.

1 Background

By definition, in *human-centered* design, human needs drive system requirements [10]. In standard software engineering workflows, HCI designers gather inputs (data) from end-users to identify needs, generate designs alternatives, and finalize software specifications. These specifications capture all aspects of software behavior, including functional and data requirements [8]. They are handed to engineers who translate them into technical requirements for software implementation [11]. Ideally, AI software design will follow a similar workflow. AI engineering consists of critical human-centered decision points, including identifying training data characteristics, data labeling, feature engineering, etc. [1]. User research and end-user data will inform the design of AI-powered User Experiences (UX) and the AI subcomponents, i.e., human-centered AI. Unfortunately, in many cases, UX research and AI development only converge after AI decisions have been made [5]. While the central role of data in both could potentially bring together HCI and AI processes, current design practices, and tools make it difficult to connect the roles of data in UX design with the data roles in ML processes [15]. Furthermore, ML practitioners tend to overlook human-centered needs in the creation of training data [13]. The result is human-misaligned AI applications resulting in misjudgment of human behavior, bias, and harms to humans [3].

The recent shift towards data-centric AI (DCAI) [12] offers an opportunity to address these concerns. The DCAI approach makes use of increasing stability in ML techniques and architectures to *focus* on optimizing training data (the *backbone* of AI applications) for representativeness of human users and use context. To investigate the intersections of data, human-centered design, and DCAI, we conducted two studies on how human data can inform AI system design. First, we interviewed 21 user experience (UX) researchers, AI engineers, data scientists, and product managers from 14 software companies to understand their current practices for creating human-centered AI systems. We considered their practices using our compilation of human-centered guidelines for AI development (e.g., [2, 6, 7]). In a second in-lab study, pairs of UX designers and AI engineers (one of each, ten sessions in total) worked to design AI experiences based on data probes (i.e., data-driven design) [14].

Preprint. Under review.

We report three critical insights from these studies highlighting how to incorporate end-user data into the design of AI systems.

2 Three Key Insights

2.1 HCI can lead Engineering-Centric DCAI towards Human-Centric Data

We learned from our interviews that AI model goals drive training data requirements in current workflows. When exploring novel AI capabilities, AI researchers don't always know what data might be informative; consequently, requirements for data and its characteristics, such as variables, data types, labels, and the number of data points, evolve through a "trial and error" approach. Study participants described focusing on optimizing the AI development process, so training data initially involved small datasets often collected from their own team. Participants described an "unwritten agreement" that teams will provide data for model development. This approach is expedient, lowers costs of data access, and bypasses formal procedures necessary for human data collection, but falls short of capturing the wants of potential users. Further, engineers also reported striving for "clean" data by removing outliers and noise to improve model performance. This may lead to an idealized dataset used during AI model exploration that omits features occurring in real-world system use. For DCAI, systematic data practices are needed to guide teams in achieving data *representativeness* rather than 'minimum-viable data' and in allocating resources for identifying data needs, collection, labeling, and end-user access, diversity, privacy, and security.

2.2 HCI Practices are Readily Applicable to DCAI

Designers in our interviews reported they supported training data pipelines based on conventional human-centered approaches. Given the significance (and multiple roles) of data in HAI design, data collection and annotation tools are essential for informing end-user requirements. For example, in one interview, the designer reported visiting with on-site staff performing data labeling to understand pain points and running user studies to evaluate data annotation tools. Participants also reported applying early-stage evaluation techniques with target end-users to validate data labels. In addition, to inform training needs, designers provide (1) details about personas emerging from surveys, (2) qualitative code-books with terminology, definitions, and guidelines for training-data annotation, and (3) storyboards capturing human use contexts to inform training data characteristics, representativeness, and formatting needs. Collectively, these practices provide UX designers agency and access to shape training data (DCAI) informed by human users from the ground up.

2.3 Human-Centric Data Provides a Shared Language for DCAI Collaboration

In human-centered design, design probes promote generative thinking and exploration of the design space [9]. In our lab study, we found that end-user data gathered through HCI research served as *data probes* while designing the AI experience [14]. The designers and engineers used data probes as mutual constraints to design both the UX experience and AI capabilities. Designers used personas, data points, vignettes, and user scenarios to advocate for end-users needs, expectations, and decision factors. The engineers translated them into features, datasets, and rules for training the AI. This co-creation process led to discussions about the attributes, priorities, and values important to users alongside determinations of the technical AI capabilities needed to support them. Finally, teams used data probes to align the AI and UI by assessing failures, design flaws, misconceptions, and scalability issues. In other words, end-user data probes serve as a "common ground" for designers and engineers in collaborating on DCAI.

3 Conclusion

Current "AI-first" software workflows emphasize advancing AI capabilities through model testing. Consequently, the human-centered design of training data occurs too late to avoid costly system revisions and severe consequences to end-users (e.g., disparities in gender classification [4]). As our insights show, established HCI methods can support the design of human-centered DCAI. In this workshop, our goal is to initiate a conversation about DCAI and gather feedback on putting these findings into practice by imagining future design and engineering tools for creating HAI experiences.

References

[1] Saleema Amershi, Andrew Begel, Christian Bird, Robert DeLine, Harald Gall, Ece Kamar, Nachiappan Nagappan, Besmira Nushi, and Thomas Zimmermann. Software engineering for machine learning: a case study. In *Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Practice*, pages 291–300. IEEE Press, 2019.

[2] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. Guidelines for human-ai interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, page 3. ACM, 2019.

[3] Jack Bandy. Problematic machine behavior: A systematic literature review of algorithm audits. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1):1–34, 2021.

[4] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pages 77–91. PMLR, 2018.

[5] Graham Dove, Kim Halskov, Jodi Forlizzi, and John Zimmerman. Ux design innovation: Challenges for working with machine learning as a design material. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 278–288. ACM, 2017.

[6] Google. People + ai guidebook, 2019.

[7] Apple Inc. Designing the ui and user experience of a machine learning app, 2019.

[8] Martin Maguire and Nigel Bevan. User requirements analysis. In *IFIP World Computer Congress, TC 13*, pages 133–148. Springer, 2002.

[9] Tuuli Mattelmäki et al. *Design probes*. Aalto University, 2006.

[10] Donald A Norman and Stephen W Draper. *User centered system design: New perspectives on human-computer interaction*. CRC Press, 1986.

[11] Kayur Patel. Lowering the barrier to applying machine learning. In *Adjunct proceedings of the 23nd annual ACM symposium on User interface software and technology*, pages 355–358, 2010.

[12] Chris Ré. The road to software 2.0 or data-centric ai, 2021.

[13] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. "everyone wants to do the model work, not the data work": Data cascades in high-stakes ai. In *proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2021.

[14] Hariharan Subramonyam, Colleen Seifert, and Eytan Adar. Towards a process model for co-creating ai experiences. page 1529–1543, 2021.

[15] Qian Yang. Machine learning as a ux design material: How can we imagine beyond automation, recommenders, and reminders? In *2018 AAAI Spring Symposium Series*, 2018.