

This forum provides a space to engage with the challenges of designing for intelligent algorithmic experiences. We invite articles that tackle the tensions between research and practice when integrating AI and UX design. We welcome interdisciplinary debate, artful critique, forward-looking research, case studies of AI in practice, and speculative design explorations. — Juho Kim and Henriette Cramer, Editors

How Does AI Challenge Design Practice?

Thomas Olsson and Kaisa Väänänen, Tampere University

Machine learning–based systems have become the bread and butter of our digital lives. Today’s users interact with, or are influenced by, applications of natural language processing and computer vision, recommender systems, and many other forms of so-called *narrow AI*. In the ongoing commodification of AI, the role of design practice is increasingly important; however, it involves new methodological challenges that are not yet solved or established in *design practice*.

Building on recent work on human-centered AI (HCAI) design [1,2], in this article we ask how design practice might, or ought to, change, considering the new computational building blocks and new conditions and societal requirements that AI introduces. For example, how do the cross-disciplinary discourses around AI ethics and societal responsibility introduce new criteria for service design and user experience design? How do evolving machine learning models shape the artifacts of interaction design? In design literature, Bryan Lawson and Kees Dorst [3] stress the need for imagination and *constructive forethought* and for *toleration of uncertainty* while working with *incomplete information* and *conflicting requirements*. These considerations are most topical, as designers expand their methods and professional practices to concern AI as a design material.

WHAT TO EXPECT WHEN YOU’RE EXPECTING AI

AI features qualities and related trends that we consider particularly relevant

for designers. Development of AI likely results in the *increasing agency and proactivity of technology*. The level of automation is steadily increasing, digital systems are unquestionably influencing people’s behavior, and autonomous algorithms are employed even in activities and decisions that traditionally have been at the discretion of human actors. While increasing agency seems tempting in terms of productivity, it also involves the risk of thoughtlessly delegating complex decision making and reasoning to opaque AI applications, as discussed extensively across disciplines. Moreover, AI technologies tend to enter application areas that people might consider sensitive, nontechnical, or intimate (e.g., personal coaching, healthcare, dating). These observations stress the classical HCI question of the appropriate roles for technology. This calls for practices and methods that advocate reflectivity, multidimensional analysis of the context, and holistic forethought and afterthought.

Furthermore, AI systems are *evolving by nature* [2]—in contrast to monolithic information systems that remain unchanged for years. While

the dynamism of digital systems has generally increased through software production practices that underline iterativity and continuous deployment, machine learning enables the systems to evolve simply by being used and by accumulation of new data. Resulting unintended system behaviors tend to be irreparable rather than easy fixes. Hence, designers need methods and preventive mechanisms to deal with the long-term sociotechnical implications of their design artifacts.

Considering technology development in general, we are witnessing an *ethical awakening* among organizations and developer communities and in public discourse. As people are demanding responsibility, fairness, transparency, and bias-free decisions from algorithmic systems, the moral intensity of technological issues appears to be increasing [4]. After the emergence of *usability* in the 1990s and *user experience* in the 2000s, now seems to be the time to design for the even grander concepts of *ethicality* and *responsibility*. That said, the various promising methods that aim to improve design ethics and responsibility (e.g., value-sensitive design) are just beginning to be established in design practice.

FOUR PERSPECTIVES TO AI DESIGN PRACTICE

To offer a framework of expected dynamics in design practice, we highlight four perspectives: *product, people, principles, process*—what we call the 4P model of AI design. This framework, inspired by similar acronyms and perspectives commonly used in marketing and

Insights

- Artificial intelligence introduces new design materials with principles and qualities that demand revisiting the practices of human-centered design.
- We propose a preliminary framework (4P) to highlight key perspectives to reconsider in human-centered AI design: product, people, principles, and process.



service-quality literature, offers structure for approaching AI-specific design questions. It offers complementary viewpoints to prior HCAI considerations, particularly those by Yang et al. [2] on, for example, understanding AI capabilities.

The *product* perspective refers to the mental models about the designed AI artifacts. Following the trend of increasing autonomy, we anticipate a shift from designing instrumental and reactive *information tools* to designing proactive *agents and collaborative partners*. While similar ideas of post-instrumental functions and the proactivity of technology have been much discussed since early ubicomp visions, AI applications further challenge the conventional mental models of user-product relationships. For example, chatbot-based companion applications like replika.ai (<https://replika.ai/>) offer new value propositions that demand consideration of the long-term implications for relational communication and emotion-regulation skills.

Considering interaction, Wei Xu [1] argues for evolution from human-computer interaction to “human-machine integration” or “human-machine teaming.” Rather than designing user interfaces with inputs and outputs, designers should focus on designing, for example, *scripts for collaboration*, *interest conflict-management procedures*, and *context-specific rules for agency*. Considering recommender systems and social robots, the user’s role has already shifted from *commander and controller* to that of indirect *coach* of their AI companions.

As an example from our recent work, we have reconsidered the role of the user interface in conditioning behavior in online news commenting. Our design exploration played with the idea of introducing AI-based affect-labeling mechanisms that would encourage self-reflection by readers and commenters. The design of what might seem to be a conventional, Web-based, textual-discussion platform UI has turned into the design of *computational discussion facilitation* and *AI-assisted emotion regulation*. This highlights the strengthening sociotechnical and multidisciplinary nature of design work.

The *people* perspective considers for *whom* to design and *who* might be



influenced by the system. Following philosophies like design for all, universal design, and inclusive design, the principle of catering to different user groups, cultures, and stakeholders is ever more important when designing AI applications. On the one hand, AI can support this principle by allowing more-personalized and adaptive service. On the other hand, the underlying notion of intelligence in AI might imply a false promise of *automatic personalization by default*, making it easy for the designer to forget the diversity of users. Moreover, decreasing user autonomy due to highly proactive services may do a disservice to the design’s appeal to vulnerable user groups such as the elderly.

The qualities of AI services also challenge the notion of *usership*—the different forms and positions of being a user. For example, use of proactive services like ambient voice assistants emphasizes the notion of *dynamism*: Anyone can start using a voice assistant with a simple utterance, and already the next command can be given by another person. This stresses the need for catering to not only the various user groups but also secondary users (e.g., co-located people) and tertiary users (people elsewhere using the ML-based natural language processing model and

whose service is influenced by the inputs of all other users). For example, to avoid unintended inputs and interpretation errors, people at a social gathering need to negotiate who is the active user providing input, hence weakening users’ autonomy to opt out of being a user.

The *principles* perspective refers to the *values* and *fundamental propositions* that shape how designers understand and solve design problems. The societal and economical purpose of design is to be the engine of product innovation; designers are expected to serve as change agents and advocates of radical innovation, while also having the power to define the preferences and values that their creations follow. The much-discussed notions of sustainability and responsibility introduce seemingly convincing, yet lofty and broad, sets of principles to consider, from respect of human rights and promotion of equality to minimizing computation’s carbon footprint. Similarly, considering AI ethics, principles like *autonomy* and *explainability* set new standards for the design quality of AI systems. However, designers are expected to follow such values that are generally desirable, while the scientific community—let alone the public opinion—is not unanimous about what exactly those principles are.

For example, it is easy to agree with a general requirement of *fairness*, but there are numerous definitions of *fairness* that may conflict with one another—and that also depend on the context. Our recent work has investigated this in the context of team-assembly systems. When considering the question of effective and fair team compositions for student-innovation projects, as designers we had to make sense of the seemingly vague idea of fairness: Which rival theories of fairness and discrimination should a computational solution manifest, in terms of what resources different stakeholders demand for equality, what applicant characteristics should be highlighted in the UI, and how the system could support equal treatment and help the decision maker avoid biases? Further, while operationalizing a conventionally little-regarded principle like this, one needs to be mindful of what other principles need to be satisfied simultaneously and what trade-offs might result.

Generally, superior quality is an elusive goal: The constantly updating principles demand sensitivity and agility to identify the emerging quality attributes and to cater to them in design work. Following Anders Albrechtslund [5], we call for strengthening skills in areas such as *envisioning probable futures* (e.g., what-if scenarios and contrafactual sociotechnical imaginaries), identifying new sources of *design inspiration*, rethinking conventional *design patterns and trends*, and identifying the possible *gaps between design intentions and the eventual use of artifacts*. In particular, the efficiency-driven convention of utilizing design systems, patterns, benchmarks, and other design legacy involves risks. When solutions are transferred from one problem to another, they can be decontextualized and misappropriated in such ways that new types of bias unintentionally emerge. For example, it would likely be detrimental to transfer *nudging* solutions from a game into a decision-support system that aims at enabling well-informed and democratic choices. All in all, while our probabilistic AI systems are increasingly capable of making accurate linear predictions based on vast training data, as

designers we need to develop our skills of collectively defining what future directions are worth pursuing.

The *process* perspective deals with design as a *practical, professional production*. Generally, the design of AI systems should follow the basic stages of human-centered design, including identification of various user and stakeholder requirements, exploration of solution alternatives, and user-based testing. At the same time, mindful of the need for ethical deliberation, the processes should not only permit but also *encourage* consideration of various quality criteria and values. Design teams need to create processes that help review design directions and recognize their risks earlier, to avoid mistakes resulting from ignorance or short-sightedness and generally avoid the traps of technical intervention [6], such as the solutionism trap or portability trap.

Considering the temporal perspective of production processes, product quality might not be definable at the point of release, but only long after deployment. That is, the *definition of done* will likely change as the product evolves, becoming a moving target. Further, the increasing technological agency might call for drastic measures should a solution prove undesirable in the long run; this could mean discontinuing the system deployment instead of trying to patch it. This underlines the importance of post-release follow-up, comprehensive reflection, and defining intervention mechanisms and repair procedures for solutions already in use.

Considering software production as the professional context in which interaction designers primarily operate, perhaps the nearly axiomatic agile development processes ought to be rethought. Sprint-based development combined with *minimum viability* tends to favor speed over deliberation and incremental improvements over questioning a design’s fundamentals. How could we showcase minimum viability in terms of ethicality or sustainability? We are currently preparing a new research project that will critically look into the values and conventions in technology development, aiming to develop methods that could help designers and developers

reflect on their professional practices. To this end, we hope that the much-debated quote “move fast and break things”—as an invitation to not care—would not represent the ethos of AI development and business in the future.

In this article, we outlined various considerations regarding how AI might challenge design work. We highlighted four perspectives to design dynamics (*product, people, principles, and process*) and offered examples that hopefully will provoke discussion and reflection, as well as trigger methodology development in human-centered AI design practice.

ENDNOTES

1. Xu, W. Toward human-centered AI: A perspective from human-computer interaction. *ACM Interactions* 26, 4 (Jul.–Aug. 2019).
2. Yang, W., Steinfeld, A., Rosé, C., and Zimmerman, J. Re-examining whether, why, and how human-AI interaction is uniquely difficult to design. *Proc. of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, New York, 2020, 1–13.
3. Lawson, B. and Dorst, K. *Design Expertise*. Taylor & Francis, 2009.
4. Shilton, K. Values and ethics in human-computer interaction. *Foundations and Trends in Human-Computer Interaction* 12, 2 (2018), 107–171.
5. Albrechtslund, A. Ethics and technology design. *Ethics and Information Technology* 9 (2007), 63–72. DOI: 10.1007/s10676-006-9129-8
6. Selbst, A.D., boyd, d., Friedler, A.S., Venkatasubramanian, S., and Vertesi, J. Fairness and abstraction in sociotechnical systems. *Proc. on the Conference on Fairness, Accountability, and Transparency*. ACM, New York, 2019.

📍 **Thomas Olsson** is an associate professor of human-technology interaction at Tampere University, Finland. He works on sociotechnical systems, computer-supported cooperative work, and critical design of AI-based applications. He leads the Technology x Social Interaction research group and serves as an associate chair in ACM CHI and ACM CSCW. → thomas.olsson@tuni.fi

📍 **Kaisa Väänänen** is a full professor of human-technology interaction at Tampere University, Finland. She leads the Human-Centered Technology (IHTE) research group in the Computing Sciences unit. Väänänen has 25 years of research experience and is currently focused on human-centered AI and sustainable development supported by digital solutions. → kaisa.vaananen@tuni.fi