Rethinking HCI Education for the Era of AI

Kasper Hornbæk, University of Copenhagen
Per Ola Kristensson, University of Cambridge
Antti Oulasvirta, Aalto University and ELLIS Institute Finland

Insights

- HCI education has a key position in efforts to develop AI systems that are usable, useful, and safe.
- What we teach in HCI, and how we teach it, must evolve to address the distinctive challenges of human-AI interaction.
- We outline four challenges to renewing HCI education to meet this need.

Jonathan Grudin famously described how the fields of HCI and AI are bound together in seasonal cycles [1]. Competing for the same limited economic and intellectual resources, HCI tends to experience hardship when AI flourishes. When HCI enjoys its summer, AI braces for winter.

We believe many readers will agree that HCI has been enduring a winter. The impact of HCI research articles is declining [2], and there is rampant *LLM-ification* [3]. It is also unclear whether HCI is succeeding outside the ivory towers of academia. Users feel widespread frustration with computing systems [4]. Some products with AI-infused features, seemingly well-grounded in HCI, have flopped in the market (e.g., Apple Vision Pro, Metaverse, AI Pin). Meanwhile, AI as a field is experiencing explosive growth across bibliometrics, industry funding, and beyond [5].

We argue that renewing HCI education is the pathway to a new summer. By training better cohorts of students, those students can go on to create the AI systems people actually like and use. To this end, these cohorts must be able to deeply understand the novel character of human-AI interaction, build on insights about users' changing needs,

conceptualize exciting and effective interactions, design and engineer them with rigor, and evaluate them with a critical eye. They need to achieve this in a world that is moving fast, where problems are harder than ever, where facts must prevail over opinions, and where collaboration is key to success.

In this article, we present four challenges to achieving this by rethinking how HCI is taught. These four draw from our experiences writing an HCI textbook for students in computer science, engineering, and electrical engineering. Two of the four challenges concern the content of education and how it should be renewed; the other two address teaching the interdisciplinary stance.

CHALLENGE 1: RESPOND TO THE CHANGING FUNDAMENTS OF INTERACTION

Interaction is the defining concept of the field and needs to remain the centerpiece of education. The prevailing way of teaching HCI, however, draws from concepts of interaction formed in the era of the desktop computer [6], raising the question of how well they capture the profound changes wrought by AI at the frontier of computing. Where conventional computing systems yield predictable outputs to user inputs, AI now offers unpredictable responses and can take actions autonomously. Another point of distinction is the bidirectionality of knowledge formation: It is not just the user forming goals, beliefs, and understandings—the AI is continuously doing the same.

Although contemporary research on human-AI interaction is mostly not theory-obsessed, several emerging theoretical concepts can be identified in recent literature. Some address the challenges of AI directly, such as human-AI teaming and multi-agent systems, while others augment and extend broader, long-standing HCI concepts, such as dialogue and tool use. Some of these long-standing concepts are being actively developed for AI interaction. For example, Hari Subramonyam and colleagues expanded Norman's *gulfs of evaluation and execution* with one specific to interaction with LLMs: the *gulf of envisioning* [7]. In the absence of clear affordances and signifiers, the user's challenge is to gauge which intentions are supported by an LLM.

The renewal of educational materials should cover emerging accounts like these. Together, they offer a rich overview of interaction with AI. They span different scopes and explain key empirical phenomena, such as trust and alignment, in different ways.

CHALLENGE 2: FOCUS TEACHING ON TIME-ENDURING PRINCIPLES

The most valuable things that can be taught in university classes, whether in HCI or elsewhere, are foundational principles. With AI developing at a fast pace, this is even more, not less, important.

HCI's principles reveal something fundamental about interaction: how people understand it, how interfaces should work, and how to design for it. *Direct manipulation* is a well-known example from the pre-AI era. But the AI era has prompted both renewed interest in "old" principles and the development of new ones, many of which are not conveyed to students. These address key *phenomena* in human-AI interaction, such as alignment, theory of mind, explainability, ironies of automation, and shared control. They are increasingly applied in empirical studies, the results of which are highly relevant for practitioners, too. For example, J.D. Zamfirescu-Pereira and colleagues studied why users struggle to design effective prompts for LLMs, explaining that this happens because their expectations are drawn from human-human interactions [8].

The most valuable things that can be taught in university classes, whether in HCI or elsewhere, are foundational principles.

CHALLENGE 3: INTRODUCE HCI AS AN APPLIED SCIENCE

We hold that HCI should be taught as an applied science, not just a craft. While practical skills are important, framing HCI as merely a practice—or worse, as a mindset—may fail to equip students with the tools required to tackle the hard, complex, systemic problems that human-AI interaction incurs. Relying on seat-of-the-pants approaches may erode the quality of outcomes as well as a practitioner's ability to deliver impact and persuade in real-world projects.

What does this mean in practice? Human-AI interaction needs to be approached as an *evidence-based discipline*. Students need to learn how to back up their insights with empirical, formal, and demonstrative evidence. They need to be able to draw reliable conclusions from the flood of empirical studies published on interaction with AI. They need to form evaluations whose results can be trusted even when a new large language model is released. And they need to justify their design choices by reference to relevant theories.

We envision that, equipped with an applied science ethos, students can approach AI with results they can trust and avoid the temptation to conclude that *everything is messy*.

CHALLENGE 4: EMBRACE INTEGRATIVE CONCEPTS FOR INTERDISCIPLINARITY

HCI benefits from many concepts such as theories, methods, and frameworks. However, HCI courses should teach *integrative* concepts—those that combine knowledge from relevant disciplines in actionable ways. Science, design, and engineering each offer a fundamental way of knowing in HCI, but it is at their intersection where the action lies. This calls for increasing emphasis on integrative forms of knowledge.

What do they look like? HCI theories are a prime example of integrative knowledge. Take, for instance, *self-determination theory*, a popular theory for understanding motivation, such as why people play computer games, now extended to

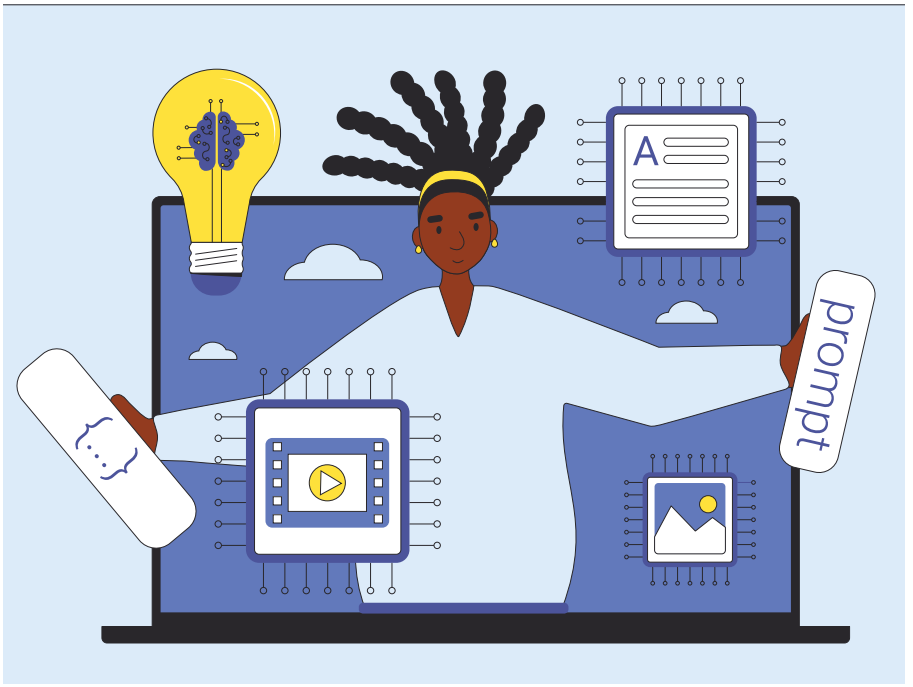

explain interaction with AI. It includes theoretical propositions, measurement instruments such as questionnaires, and empirical findings. Or consider *information theory*, which rests on a few core assumptions about interaction that remain relevant when interacting with AI. Information transfer is capacity-limited, which can be expressed formally, implemented in code to drive intelligent interfaces, or used to measure user behavior.

Integrative knowledge, however, is not limited to theories. Models express, in formal terms, how theories can be applied to real-world problems. Systems, and artifacts more generally, can embody cross-disciplinary, counterfactual knowledge ("What if we designed it this way instead?"). In short, they embody visions of the future, something more relevant than ever in the era of AI.

Importantly, HCI educators should try to reach beyond traditional HCI boundaries in search of such concepts. For example, recognizing that intelligent systems introduce new types of system-level challenges, concepts such as systems engineering, safety, and human error are critical—topics traditionally covered in human factors courses but less frequently in HCI. We believe that concepts like these will outlast any technological fad.

Finally, critical thinking will be a central integrative requirement for the new HCI cohorts. AI is not a panacea, and neither are HCI methods. That is why media literacy, the ability to debate and explore alternative perspectives, and a continual reading of HCI literature are key goals for education.

OPTIMISM FOR THE NEW ERA

We believe the new HCI summer is within reach. There is tremendous untapped potential in HCI research. We need to figure out a way to transfer that to our students. This call is timely. When the development of AI technology stagnates, the focus will shift from computational models to applications. As Ben Shneiderman noted, HCI has profoundly shaped some of the most influential technologies we use today, from mobile devices to computer games [9]. AI must not be an exception.

As the first step in our call to renew HCI education, we wrote a textbook for introductory HCI classes that was published by Oxford University Press. This open-access book is available as a free PDF on the publisher's platform.[a]

a Kasper Hornbæk, Per Ola Kristensson, and Antti Oulasvirta, *Introduction to Human-Computer Interaction*, Oxford University Press, 2025, https://bit.ly/3XlzBjf

References

1. Grudin, J. AI and HCI: Two fields divided by a common focus. *AI Magazine 30*, 4 (2009), 48–57.
2. Chen, Z. and Li, Y. The sharply decreasing disruptiveness of HCI. *Proc. of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, 2025, Article 454, 1–24.
3. Pang, R.Y. et al. Understanding the LLM-ification of CHI: Unpacking the impact of LLMs at CHI through a systematic literature review. *Proc. of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, 2025, Article 456, 1–20.
4. Hertzum, M. and Hornbæk, K. Frustration: Still a common user experience. *ACM Transactions on Computer-Human Interaction 30*, 3 (2023), Article 42, 1–26.
5. Ahmed, N., Wahed, M., and Thompson, N.C. The growing influence of industry in AI research. *Science 379*, 6635 (2023), 884–886.
6. Hornbæk, K. and Oulasvirta, A. What is interaction? *Proc. of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 2017, 5040–5052.
7. Subramonyam, H., Pea, R., Pondoc, C., Agrawala, M., and Seifert, C. Bridging the gulf of envisioning: Cognitive challenges in prompt based interactions with LLMs. *Proc. of the 2024 CHI Conference on Human Factors in Computing Systems*. ACM, 2024, Article 1039, 1–19.
8. Zamfirescu-Pereira, J.D., Wong, R.Y., Hartmann, B., and Yang, Q. Why Johnny can't prompt: How non-AI experts try (and fail) to design LLM prompts. *Proc. of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, 2023, Article 437, 1–21.
9. Shneiderman, B. Revisiting the astonishing growth of human–computer interaction research. *IEEE Computer 50*, 10 (2017), 8–11.

Kasper Hornbæk is a professor of computer science at the University of Copenhagen. He works on theories of HCI. He is also interested in the methodology of HCI, including the role of replications, measures of usability, and solid experimental work.
→ kash@di.ku.dk

Per Ola Kristensson is a professor of interactive systems engineering in the Department of Engineering at the University of Cambridge and a fellow at Trinity College Cambridge. He is a cofounder and codirector of the Centre for Human-Inspired Artificial Intelligence at the University of Cambridge.
→ pok21@cam.ac.uk

Antti Oulasvirta is a professor at Aalto University, where he leads the Computational Behavior Lab (cbl.aalto.fi), and a principal investigator at the ELLIS Institute Finland.
→ antti.oulasvirta@aalto.fi

DOI: 10.1145/3777517