

Searching for the Non-Consequential: Dialectical Activities in HCI and the Limits of Computers

Haoqi Zhang
Northwestern University
Evanston, IL, USA
hq@northwestern.edu

ABSTRACT

This paper examines the pervasiveness of consequentialist thinking in human-computer interaction (HCI), and forefronts the value of non-consequential, dialectical activities in human life. *Dialectical activities* are human endeavors in which the value of the activity is intrinsic to itself, including being a good friend or parent, engaging in art-making or music-making, conducting research, and so on. I argue that computers—the ultimate consequentialist machinery for reliably transforming inputs into outputs—cannot be the be-all and end-all for promoting human values rooted in dialectical activities. I examine how HCI as a field of study might reconcile the consequentialist machines we have with the dialectical activities we value, and propose *computational ecosystems* as a vision for HCI that makes proper space for dialectical activities.

CCS CONCEPTS

• **Human-centered computing** → **HCI theory, concepts and models.**

KEYWORDS

Dialectical activities, consequentialism, philosophy in HCI, human values in HCI

ACM Reference Format:

Haoqi Zhang. 2024. Searching for the Non-Consequential: Dialectical Activities in HCI and the Limits of Computers. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3613904.3641945>

1 INTRODUCTION

With every new technology, we grow increasingly capable of shaping the world to meet our needs and a growing list of desires. Where technology has gone, HCI has largely followed. The growth of HCI into its many subdisciplines—and the ways that our interactions with computational devices, tools, and media shape every aspect of human life and lived experience [11, 114]—is further evidence of technology’s reach into every index of human desire. For any need we cannot yet meet, technological optimism promises that we will find the means to fulfill it sooner or later.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CHI '24, May 11–16, 2024, Honolulu, HI, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0330-0/24/05

<https://doi.org/10.1145/3613904.3641945>

But while the HCI mission of using computational technologies to shape the world to meet our needs and desires rolls on full steam, questions to the very idea of focusing on the production of desired ends remain largely unanswered. As is the case in our culture, much of HCI research and practice is rooted in *consequentialist thinking*: reasoning about actions as means for achieving desired outcomes and ends. But as philosophers have contested across millennia, certain quintessential human values, activities, and ways of being cannot be easily reconciled nor understood through the consequentialist lens. For instance, *dialectical activities* [25], or activities whose values are rooted in the intrinsic nature of the activity itself and that are revealed only through repeated engagement with the activity—such as parenting, being a good friend, engaging in art-making and other creative pursuits, conducting research—do not easily reduce to producing certain desired outcomes. Continued attempts to treat these dialectical activities as a form of production can lead us astray to overfocus on various forms of attainment, than on promoting deep, thoughtful engagement in these valuable activities themselves.

As a field that sits between computer science and the humanities, HCI has a dual mandate to harness the powers of computational technologies, *and* to critically examine how human goals and values can be harmed through our engagement with computation [1, 13, 50, 119, 120, 122]. In this paper, I argue that the allure of consequentialist thinking has not escaped HCI. I argue that computers—the ultimate consequentialist machinery for reliably transforming inputs into desired outputs—amplify the shortcomings of consequentialism. This leads to an impossibility claim: that computers, regardless of their powers of production, can never be the be-all and end-all for promoting human values rooted in dialectical activities. The implications for HCI research and practice are vast, as it requires us to reconcile fundamentally the consequentialist machinery we have with the dialectical activities we value. This reflection is at the heart of HCI’s dual mandate, as it calls on us to rethink how our machines fit in with the most human aspects of our lives.

The paper proceeds as follows. I first outline the key tenets of consequentialism, and share examples of human values and ways of being that cannot be easily captured through the consequentialist lens. Following Talbot Brewer’s arguments in his book *The Retrieval of Ethics* [25], I then broadly introduce the notion of dialectical activities, and argue for their importance. I take an intermission to position our discussion in the context of existing HCI research, and then turn to examine how computers fail to capture dialectical activities that fall out of the purview of consequentialist thought. I discuss the implications of this claim on HCI, and examine the

challenges and ways by which HCI can indirectly or directly support dialectical activities. I close by presenting a vision for HCI research and practice rooted in *computational ecosystems*, that calls for the design and study of socio-technical configurations that jointly advance consequential outcomes *and* promote engagement in dialectical activities.

2 CONSEQUENTIALISM AND ITS LIMITS

I begin by describing and reviewing a view of the world that pervades our practical thinking: that of consequentialist thought. My description of consequentialism and arguments against it largely summarizes arguments made by Talbot Brewer, in *The Retrieval of Ethics* [25].

The core idea behind *consequentialism* is that our actions can be understood in light of the desired *ends* that they are meant to produce. Consequentialist thinking embeds a “world-making prejudice” (p. 12), in that it is concerned primarily with shaping the world in some desired way to achieve some desired effect. Given an end (e.g., a completed task, a desired feeling state), the goal is to find *means* that serve as a mechanism for satisfying that need or desire. Consequentialists may debate which goals and ends to promote, but not the ideal of finding some way to shape the world to reach some desired state of affairs.¹

Many human activities—primarily those concerned with producing a desired state of affairs—are naturally captured in consequentialist terms. A desired end may be to achieve a tangible goal, for example to produce a good, complete a task, or learn a skill. A desired end can also be to have a particular experience or to arrive at a certain state of being, such as engaging playfully, feeling joy or a sense of connection to others, being curious, having a felt experience, and so on. The work of designers and technologists is often to figure out how to help people bring about such ends. While early HCI research and practice focused on achieving a narrower set of ends (e.g., completing cognitive work tasks), much of HCI research and practice today expand this focus to achieving a wide range of desired ends (e.g., in work, play, learning, and relating), for diverse people, in a wide variety of settings, and across devices and interactional modalities [21, 48, 59, 118].

Despite its prevalence, Brewer argues that consequentialist thinking is limiting, in that it locates the value of human activities solely in the objectives that they seek to achieve. As an example, consider parenting. Brewer argues that the meaning of parenting cannot be, and should not be, reduced to the production of children with some desired set of characteristics and capacities. For Brewer, the problem isn't simply one of choosing the *right measure* or the *right end* that best captures the kind of child we wish to raise. Instead, the deeper problem is overfocusing on “the not-yet-present value of the child's possible future traits or achievements,” and “not at all on [honoring] the already-wholly-present value of the child” (p. 93). In other words, with the consequentialist mindset, we fundamentally overlook the *intrinsic value* of relating to one's child, in the here and now. Brewer goes on to argue that the consequentialist account of many important human activities (e.g., parenting, befriending,

art-making, ethical reasoning) in terms of its desired effects can grotesquely distort and obscure the intrinsic good of such activities, in and of themselves.

Within HCI practice, we can see this playing out through various domains of application: dating apps that focus on finding a match; meditation and workout apps that foreground the duration and frequency of engagement; and social media apps that promote building large networks of followers. In these applications, achieving certain outcomes become proxies for the intrinsic good of the activity: of intimately relating to another person, of cultivating healthy living and mindfulness, of being a good friend or colleague. But in actuality, achieving these outcomes do not equate to our grasping the intrinsic good of the underlying human activity. Instead, foregrounding these outcomes can reinforce a consequentialist lens that distorts: we end up trying to achieve them *regardless* of whether we come to see the intrinsic good in our activities, and of how we should best engage with them. Our lives in turn, become a simulacra of the good life [15, 25]. Our achievements give us the sense that we are engaging meaningfully, even as the genuine good in these activities—precisely those that cannot be codified as ends to be achieved—remain obscure to us.

3 DIALECTICAL ACTIVITIES AND THEIR VALUES

As an alternative to the production-oriented picture of the human self, Brewer introduced the concept of *dialectical activities* to highlight human activities that are valuable for their own sake.² Engaging in a dialectical activity—such as parenting, relating to other human beings, making art or music, or conducting research—can form an integral part of a person's life and serve as an important source of meaning. Dialectical activities unfold over time; we engage in them not with a clear picture of an end to achieve, but instead with a vague and imperfect sense of their intrinsic good and their place in our lives. Through successive engagements, we devote ourselves to deepening our comprehension of the intrinsic good of an activity, by striving to embody the ideal form of engagement through our actions. Such an activity in turn has a “self-unveiling character” (p. 37), in that the activity's ideals and internal goods are progressively clarified by means of our ongoing, dialectic engagement in the activity in pursuit of its ideals.

As an example of engaging in a dialectical activity, Brewer describes a blues singer, searching for the best way to perform a particular part of a song. It is through her attempting to find the perfect intonation—exploring musical possibilities and hearing herself as she sings them—that she sharpens her grasp of a “goodness lying untapped in one of its possible renditions” (p. 47). She may very well produce a good performance as a result of singing well, but this is not so much an end she already knows the shape of, but rather the product of the quality of her (improvised) engagement. Were the exact same sounds produced without her engaging with the music as she does—e.g., by copying the motions of her lips, or by using some technology to record her—such a reproduction

¹There is a real sense here that the value of our activities reside in ends. To the extent that certain aspects of the means are undesirable (e.g., it takes too long; has an undesirable side effect), a consequentialist may simply reformulate the desired end state to exclude such undesirable means.

²Brewer's conception of dialectical activities is rooted in the ancient Greek philosophies of Aristotle and Plato, and follows in the tradition of virtue ethics philosophers such as Elizabeth Anscombe, Iris Murdoch, and Alasdair MacIntyre. See *The Retrieval of Ethics* [25] for more.

would not capture the good inherent in her *singing*, in search of the good in singing.

Brewer's formulation of dialectical activities naturally shifts our focus away from ends to be achieved and toward our present engagement in activity. When we think in terms of ends, we value our activities by their future payoff. But with dialectical activities, Brewer argues that "to dispense with a stretch of life during which one is wholeheartedly engaged in [dialectical] activity would hardly be a boon, even if one is still striving to answer more completely to the activity's constitutive ideals" (p. 90). In other words, dialectical activities represent an important form of human agency that locates their value precisely within our present and perfecting engagement in our activities, which, even in its imperfect form, gives meaning to our doings, and to our lives.³

Brewer's concept of dialectical activities bears some resemblance to Csikszentmihalyi's concept of *autotelic experiences*, or *flow* [34]. Like dialectical activities, autotelic experiences are intrinsically valuable, apart from any desirable ends that might result from an activity. But whereas dialectical activities are intrinsically valuable because of their goodness and virtue, autotelic experiences are intrinsically valuable simply because they are self-motivated and self-rewarding. Autotelic experiences are not concerned with the good [19]; one can have an autotelic experience through morally neutral or even destructive activities [33, 34].⁴ Experientially, autotelic experiences emphasize being in flow, competently working towards clear, proximal goals that are self-endorsed and intrinsically motivating.⁵ With dialectical activities, goals are rarely clear; instead, one strains to bring the good into view through successive, self-deepening engagements, a good that is only dimly visible at first but for whose search, one self-endorses.⁶

To engage in a dialectical activity is to connect with an impersonal good that is shared, but that comes to light through one's personal projects and activities. The experience is personal, but its intrinsic good is hardly arbitrary. One cannot be virtuous by just thinking and stating that one is; one has to be aware of a *genuine* goodness or value [25]. For examples, Brewer argues that making transparently bad arguments and calling it philosophizing, or being impressed by a fallacious mathematical 'proof,' can only indicate

³This is not to say that dialectical forms of engagement are incompatible with achieving ends. In fact, deep engagement in certain activities may be useful for producing certain ends. Rather, the point is that focusing on ends, and achieving them by whatever means necessary, can miss the value of engaging in dialectical activities.

⁴For example, consider a person who enjoys making convincing arguments that are in fact spurious. This act can be self-rewarding and autotelic, but it cannot count as good thinking.

⁵While autotelic experiences are intrinsically valuable to the person, they can be designed for by thinking consequentially. This is because the structure of activity that enables an autotelic experience—namely that of having concrete goals that a person can reach and providing feedback about progress—is well within the realm of consequentialist thought. In contrast, absent any clear, codifiable goals, designing for dialectical activities seems far less straightforward.

⁶Dialectical activities are marked by the difficulty in seeing (and acting on) the good of the activity. For example, in articulating the intrinsic good of thinking like a historian, Sam Wineburg writes: "historical thinking in its deepest form is neither a natural process nor something that springs automatically from psychological development," but it "teaches us what we cannot see, to acquaint us with the congenital blurriness of our vision [137]." One cannot perfect a dialectical activity simply by having the requisite skills for performing a task [26]. Instead, one must come to see things differently, and to grasp the intrinsic good in seeing differently. What is being exercised and developed through a dialectical engagement is not merely one's skills, but one's sensibilities in seeing and in aligning one's actions to one's self-deepening understanding of the good in one's activity.

an incapacity to see the good in philosophizing and in appreciating a mathematical argument (pp. 43, 220). In this way, dialectical activities are personal projects, but they are not merely so; the proper goodness of our engagement (or lack thereof) can in some ways be seen by others, especially those who themselves deeply grasp the good in such activities.

At the same time, a dialectical activity is and must be a first-personal engagement; it cannot be delegated. Another person or machine bringing about a certain state of affairs, or presenting a certain course of action, can never equate to one seeing the good in one's own activity, for oneself. For example, no AI that generates good-sounding music can replace the values intrinsic to one's personal engagement in *music-making*.⁷ And no delegate for parenting—however successful the child becomes—can replace the values intrinsic to one parenting, and relating to one's own child.

With dialectical activities, the good cannot be separated from the person. Both aspects are important for design. From an HCI perspective then, dialectical activities challenge designers and researchers to think about the quality and form of people's engagement in their personal activities, and to consider *how a person comes to see the good in an activity*—to perfect seeing and acting on its ideals, and to search for its value by deepening their engagement. They also challenge designers and researchers to consider people's engagement in long-running activity as life-long pursuits, where the form of engagement can change as one's grasp of an activity's ideals and its place in one's life evolves.

Brewer's concept of dialectical activities also bears some resemblance to Borgmann's concept of *focal things* and *focal practices* [23, 63]. In both conceptions, there is an emphasis on one's direct engagement and presence in unfolding activity over simply achieving ends. For Borgmann, the primary concern is that technological devices can separate *means* and *ends* by providing easy access to (lesser) ends that disengage us from the practice of producing ends ourselves. With dialectical activities, Brewer's concern is with the means-ends way of thinking in and of itself; the concern is less about whether we engage or not, but *how* we engage and come to see the good inherent in the activities in which we are engaged. Specifically, Brewer's concern is that our approaching an intrinsically valuable activity or practice with a production-oriented, consequentialist mindset can miss the good inherent in the activity. This is an important issue for HCI, because it is in our not foregrounding certain forms of human engagement that we can come to advance technologies that do not support it [119], and that can replace such forms of engagement entirely.

Dialectical activities are at once deeply familiar and sadly foreign. On an intuitive level, we understand that engaging, and how we engage, matters: we are more than what we produce or achieve or experience, and we are capable of searching for the good in certain activities of ours, through which we cultivate a deeper sense of *how we can live* and *who we can be*. Yet consequentialism in all its forms (e.g., consumerism, careerism) is deeply ingrained

⁷The point here is that one's first-personal engagement in the dialectical activity of music-making cannot be replaced (without remainder) by another entity bringing about the desirable musical outputs in one's place, as one can (without remainder) if one were a consequentialist, caring only about ends. This does not preclude one from using digital instruments or digital tools, as one dialectically engages in music-making as a musician would.

in modern culture [23, 63]; an emphasis on quantifying, measuring, consuming, and producing drives much of our behavior and the focus of our designs. Until we come to understand the value of dialectical activities in our lives and how to sustain them, we run the continual risk of them being replaced by their consequentialist counterparts.

4 CONNECTIONS TO HCI RESEARCH

Having presented the core ideas behind consequentialism and dialectical activities, I turn to show how consequentialist thinking pervades existing HCI research, and articulate the value of bringing the concept of dialectical activities to HCI.

4.1 Consequentialist HCI

To highlight the various ways that consequentialist thinking pervades HCI research, I present below five familiar types of HCI research contributions, stated in consequentialist terms:

- (1) *Discover new means for reaching desired ends.* HCI systems research is largely focused on inventing and evaluating new technologically-enabled means for reaching desired ends, and learning about the relationship between means and ends. Low-level contributions may focus on creating new devices [72, 132, 136], sensors [145], and interaction techniques [41, 141] to expand interactional capabilities, while application-level contributions may create solutions that directly target the desired needs of end users [27, 28, 54, 105, 144]. Some contributions may focus on demonstrating the initial feasibility of a new technology [85, 86, 88], while others focus on making certain means more accessible and usable [102], for a more diverse set of people [73], and in more contexts and settings [9, 57]. Some contributions focus primarily on the technology itself, while others focus on how people can use and collaborate around a technology as part of a larger socio-technical configuration [18, 95].
- (2) *Understand which ends are desired and advocate for neglected ends.* HCI researchers also contribute by gaining a deeper understanding of people and their desired ends, and of obstacles to reaching them. This may consist of user research aimed at extracting design requirements [53, 67, 97], or providing new frames of reference [113] and vocabulary for describing desired ends that are difficult to conceptualize (e.g., certain experiential states [55, 103]). Such research can contribute to a more inclusive vision of HCI, for instance by highlighting the need to account for differences across people [16] and the needs of particular vulnerable populations [75, 123, 128]. It can also involve an advocacy component, in which researchers bring up problems as matters of concern [36, 121], for instance to sensitize the community to a problem [127], or to argue for the need of a plurality of approaches [32].
- (3) *Understand how new means are being used.* Instead of starting with a desired end, HCI researchers can study what ends are made possible by new technologies, and to raise questions about their desirability and use [78, 96]. Such research can surface unintended consequences from using

new technologies, for example by recognizing the unintended or undesirable outcomes of adopting certain computational technologies [3, 101]. It can also shed light on the reasons and challenges behind technological adoption or non-adoption [17, 77, 133, 140, 142], and their impact on the desired ends that such technology was intended to support [69].

- (4) *Advance the means of research and design.* Instead of focusing on specific means and ends, some HCI contributions focus on improving the way by which HCI research and design is done. This can contribute new understanding of how best to approach devising means to reach certain desired ends, and of knowing when certain ends have actually been reached. Such contributions may include: new design methods and tools [24, 109, 112]; guidelines and perspectives on designing interactions [7, 14, 65]; theories [60, 114]; toolkits that make it easier to create applications [92]; and measurement and analysis techniques and tools for evaluating means-ends relationships [5, 37, 46, 126].
- (5) *Change who can be involved when designing for an end.* HCI research can also provide new ways for stakeholders in a design to participate in the design process. This can change the power dynamics over who can appropriate technology in support of which ends, and for whom [84]. By changing who can design and in what capacity [58], such research can enable more participatory design processes to arrive at solutions that are more likely to meet the actual desires of its stakeholders, and through a process that is more empowering [6, 80, 110].

These common types of HCI research contributions largely showcase the field of HCI as a consequentialist enterprise, one that is focused primarily on producing new knowledge in service of understanding and reaching desired ends. Describing HCI in consequentialist terms, Zimmerman et al. [146] emphasize that the role of designers is precisely to “produce novel integrations of HCI research [to make] a product that transforms the world from its current state to a preferred state.” While such contributions are valuable in their own right, my aim in this paper is to shed light on a different kind of activity, whose value resides not in ends to be produced but in the good that it brings into view.

4.2 Bringing Dialectical Activities to HCI

I have presented Brewer’s conception of dialectical activities as an articulated picture of a kind of human engagement that is chosen for its own sake than for its effects, and for which successive engagements bring us closer and closer to seeing the good inherent in certain human activities. As it relates to HCI research, dialectical activities can be seen as a response to calls for HCI to more clearly describe who we want to be and become in our inter- and intra-actions with technology [49]. Beyond the need to define a good user experience, I follow prior research that argues for HCI to define its connection to the good life, and to come into touch with the good inherent in certain activities [48, 122]. In other words, by painting a clearer picture of dialectical activities and how it is distinguished from consequentialist thought, I foreground an important form of

human agency and way of approaching one's life and activities, to help us reflect on the state of HCI research and practice [1, 119].

Over the decades, HCI's focus has expanded and shifted far beyond usability to consider situated perspectives [20, 125], and to promote enabling a large variety of human experiences and expression over a narrow focus on completing tasks [11, 98]. This expanded focus is sometimes described as moving across the three waves of HCI research [49, 59, 114], that mark a shift from a cognitive, work-based perspective toward a non-work, non-purposeful perspective. While this expanded focus is important in its own right, I follow Bødker [21] in arguing that this separation of paradigms is not always useful, and in this case, does not provide a sufficient framework for distinguishing between consequential and dialectical forms of engagement. First, dialectical activities are concerned with how one engages and the self-deepening structure of certain activities; this view spans across working, non-working, art, relating, and playing activities. Second, while the focus in the third wave is on issues such as "meaning, complexity, culture, emotion, lived experiences, engagement, motivation, and experience [48]," such concerns can and often are still stated in consequentialist terms, as evoking certain states of being (affective, perceptive, body-based, etc.) as states of affair to be produced or brought about [118]. Third, analogous to Bødker's concern of how emphasizing an art-focus in the third wave can distract from a commitment to the actual users of technology, I am concerned that emphasizing a focus on experiences—many of which are short term—can distract from a commitment to understand and support people in and across their long-term engagement in dialectical activities.

Still, the turn to experience [98] provides a helpful shift from the classical view of man-machine symbiosis for advancing problem solving [89] toward the phenomenological account of making and discovering meaning as humans engage in practical activities [38]. For dialectical activities, one's progressively deepening understanding of meaning in the activity is paramount (i.e., it is phenomenological), as their meaning cannot be brought about by simply taking actions calculated to solve a predetermined problem to bring about a desired state of affairs (i.e., it is not just consequential). But Brewer's account of dialectical activities differs from the general phenomenology account in that engagement in dialectical activities has "a unique phenomenology, and one that at its heart, taps into the human good itself [19]." Whereas the general phenomenological account emphasizes "people's natural attitudes toward the world that lets them easily and unnoticeably make sense of their experience [38]," in contrast, Brewer's account of dialectical activities requires people to engage deeply in experiences as they strive to more fully grasp the intrinsic good in the activity. In other words, dialectical activity is not concerned with just any experience. It is the experience of coming to see the good inherent in our unfolding activity.

Brewer's contrasting account of consequentialism and dialectical activities also provides a useful, unifying perspective across HCI research practice and critique that makes space for non-instrumental and non-production-oriented activities within HCI research, e.g., in ethnographic work [39, 40], humanist HCI [13], critical theory [12] and critical technical practice [1, 119], games research [30], and body-based interactions [66]. In their own ways, these various strands of HCI research can be understood as together promoting

a more dialectical rather than instrumental view of HCI research and research practice. Moreover, the concept of dialectical activities helps to disambiguate between designing for autotelic experiences and promoting dialectical activities. This can help us to reason about how to support dialectical forms of engagement in our activities, beyond making certain activities self-absorbing and enjoyable (autotelic) so that people will continue to engage (e.g., [2, 45, 100]).⁸

Most directly, the concept of dialectical activities helps us to position the work of HCI researchers and designers who themselves engage (and document their engagement) in dialectical activities. For examples, we can understand Sengers and Warner as engaging in a dialectical activity of learning to relate to one another through sympathetic awareness as mediated by a technological medium [22]; and Klefeker and Devendorf as deepening their perception of familiar environments by designing garments as "filters" [79]. In both cases, what's intriguing about the authors' engagement in design is that instead of designing for a known objective, their process brought values and actions into view that were quite distinct from how they had initially thought to engage (e.g., the value and act of relating without codifying emotions; the value and act of attending to the mundane, obnoxious, and man-made aspects of one's environment). What's changing is not simply designs for reaching a desired outcome but instead the authors' evaluative outlook onto (the good in) their activity. Moreover, the authors highlight both the first-personal, idiosyncratic nature of their projects, *and* how the good that came to light through their activities is impersonal, relevant not only to the authors but broadly good in and of itself. Seen this way, these works demonstrate both the form and substance of the authors' dialectical engagement, and in how their own understanding of the inherent good of their activity evolved over time.

My research follows prior efforts to bring moral philosophy to design and HCI. As a focal example, in their paper on value-sensitive design, Friedman and Kahn [50] presented three approaches to developing moral theory: consequentialist, deontological, and virtue-based (dialectical activities belong in this last category). But besides noting that consequentialist and deontological theories are concerned with "what ought I to do?" and virtue-based theories are concerned with "what sort of a person ought I to be?" theirs and subsequent considerations of human values in HCI (e.g., [118]) are largely agnostic to the particular approach to moral theory, and largely provide an undifferentiated menu of human values for designers to consider and trade off. One critical issue with this approach is that our "choice" of moral theory affects our approach to design; in particular, it determines what technology can or cannot do to support certain activities. As we will demonstrate in the next section, virtue-based theories, and dialectical activities in particular, are difficult to design for using computational approaches precisely because such virtues cannot be captured as a mapping of specific situations to specific desired or appropriate actions (as one could potentially do for consequentialist and deontological theories). A computational approach is especially self-defeating for dialectical activities, because the core of the activity is in continually forming and revising one's own conception of the good through first-person

⁸The term autotelic is sometimes used in HCI for both types of aims. For instance, Kreminski and Mateas [82] use the term autotelic, but their positioning of reflection as an intrinsic good in a creative practice appears to me more dialectical in nature.

engagement, rather than one already having a complete and determinate understanding of what is good to encode into a machine a priori. But given our predilection to foreground advances in technology [11, 47], we run a large risk of advancing ways of being that are readily encodable into computational form while neglecting those whose values are not. In other words, we run the serious risk of losing touch with virtue-based ways of being entirely.

5 THE LIMITS OF COMPUTERS

Computing technologies have undoubtedly shaped numerous aspects of modern life. While it is easy to be enamored with the power of computers in light of their impact, it also helps us to see computers for what they are, which are fundamentally, input-output machines.⁹ At their core, computers reliably transform inputs into outputs, which is to say that they reliably produce desired consequential outcomes. But just as a consequentialist, outcome- and production-oriented mindset cannot be the be-all and end-all for engaging in dialectical activities, neither can computers, which effectively encode consequentialist thinking.

Viewing computers as consequentialist machinery can help us to better understand the limits of computers, and to design with and around them. If the intrinsic value of a human activity cannot be captured in consequentialist terms, it cannot be formalized or encoded in a computer. In other words, computers “speak” consequentialism—they accept encodings of consequential ends and the means for achieving them. Any attempt to encode a dialectical activity into a computational system as outputs to be achieved can fail to capture the intrinsic good of the activity. As noted earlier, this can lead people to chase ends that stray from the good inherent in the human activity, e.g., to chase social status and achievement, rather than deepen one’s way of relating and being. Moreover, since the intrinsic value of a dialectical activity comes from one’s direct engagement with the activity, and not from outputs that may be produced, computational procedures that produce certain desired states for us, cannot replace our grasping the intrinsic value of an activity for ourselves. For example, having AI generate the speech to give at our child’s graduation party [99] can never equate to the activity of trying to find the best words to express our feelings for our child ourselves. Understanding the *nondelegability* [25] of dialectical engagement can help us distinguish between the productive, consequentialist outcomes that computers can provide, and the dialectical activities they can (fail to) support.

Like my work here, Agre [1] largely exposes the limitations of the conventional input-output account of computation in favor of an interactionist, situated action [125] account in which actions arise through their continual dependence upon circumstances. A core similarity between my argument and Agre’s is that a pre-defined, input-output view is inherently incompatible with certain forms of human interaction, which are essentially improvised in light of specific situations. But whereas Agre is concerned with routine behaviors and their adaptations to changing situations, with dialectical activities I am concerned not with changing situations per se but with our own evolving conception of the good inherent in these activities. In an important way, dialectical activities are

distinct from routine activities in that our sense of the meaning of these activities evolves as we learn to engage in them. This can change, in a fundamental way, not only what actions we take as situations change, but how we perceive a situation and what constitutes appropriate actions in the first place.

5.1 Understanding contemporary debates in human-AI interaction

Having clarified the limits of computers in consequentialist terms, we can use this understanding to make better sense of various debates within human-AI interaction. As an example, consider the debate over “ethical AI,” in which some researchers and practitioners have argued for the need to encode ethical considerations into AI systems, so that they can produce the desired ethical outcomes [117, 131]. But recent critiques by Johnson and Verdicchio [70] argue that such attempts are doomed to fail, because ethical considerations are not amenable to computational encoding. To advance their claim, Johnson and Verdicchio argue that ethical considerations have too many possible interpretations, and that “no computational model can capture all the possible interpretations.” They also note that some ethical notions may still be contested or in flux, and thus cannot yet be incorporated into a computational system.

While I agree with Johnson and Verdicchio’s core claim, their argument largely rejects computational encoding on the basis of concerns over tractability and variance. A more fundamental and direct critique is that ethical reasoning is inherently a dialectical activity. The problem isn’t just that there are too many interpretations to capture—but that the act of ethical reasoning calls for a continual engagement in and a deepening of understanding of ethical reasoning itself. In other words, capturing interpretations and producing the desired ethical outcomes would only produce an imitation of the act of ethical reasoning. When we reduce ethical reasoning to the task of figuring out what action to perform in a given circumstance, we lose a valuable way of engaging in ethical reasoning, namely that of interpreting circumstances continually so that we may come to see them more clearly through our thoughtful engagement. Whatever the practical value of codifying the outcomes of ethical reasoning may be, it can never equate to our deliberative engagement in ethical reasoning, which at its best can change our evaluative outlook, in and of itself [25]. Insofar as we think our continual engagement in the dialectical activity of ethical reasoning is valuable in our human societies, we can never replace it with a computer program.¹⁰

My argument here reveals a fundamental limitation of using computational systems to support human activities that goes beyond those identified by Suchman [125]. For Suchman, a key difference between humans and machines is in the human ability to know a situation so as to take appropriate action. Humans can make sense of a situation readily and act to advance our ends, but machines only

⁹At least, as is conventionally defined following Turing [130]. For a more expansive view of computation, see Philip Agre, *Computation and Human Experience* [1].

¹⁰My commentary here is also applicable to critiques of algorithmic systems; see for example Alkhatib [3] and Alkhatib & Bernstein [4]. In these works, the authors highlight how the rigidity of algorithmic decision-making based on their consequentialist encodings can lead to serious problems in deployment and use. This is precisely because these algorithmic systems cannot engage dialectically, so as to continually clarify what proper engagement should look like.

have a limited view into a human situation, and we can never encode a full specification of the situation into a plan for the machine to act out (as Johnson and Verdicchio had argued, in my example). Suchman argues that this difference fundamentally limits the ways that machines can support (consequentialist) human activities. But whereas Suchman's emphasis is on the importance of knowing the situation, my argument (following Brewer) is that the best forms of human activity require us to go beyond our typical sense-making to *coming to see* a situation differently. Even when we know what the situation is, we still have to strive to bring into view the good in our activities through continual, clarifying reinterpretations and actions. It is the human capacity to engage dialectically in this manner that I wish to highlight as a core limitation of computers (i.e., their consequentialist nature), as distinct from their imperfections as consequentialist machines. For dialectical activities, neither situations nor plans are sufficient for good action. It is how we bring the good to light in response to our situations that matters [26].

* * *

As another example, consider the recent interest in generative AI. There is increasing excitement, and concern, that generative AI systems can respond to our prompts with desired artifacts (e.g., music, images, and videos) that rival those produced by human artists and creators [31, 43, 68, 104]. If successful, generative AI represents a significant advance in our ability to produce desired end states; from a consequentialist perspective, it represents a significant advance in machinery.

But if we look past its productive value, it becomes less clear if advances in generative AI constitute an advance of the good inherent in the human activity of *art-making*. As a dialectical activity, art-making comes with its own set of values [26, 56]. Having a computer that can generate the art does not necessarily deepen our own engagement in the art-making process, nor our grasp of such values. A computer can even hinder the process, for example by over-generating complete artifacts to meet a description that leaves little room for people to have finer grain controls over the process of composition [90, 91]. A straightforward implication of this observation is that advancing the productive capacities of generative AI in no way guarantees an advance to the dialectical activity of art-making.¹¹

This is not to say that generative AI cannot be used as a tool when engaging in art-making; it can be. But what I do mean to highlight is that advances in the generative capabilities of AI systems do not imply that people can nor will engage more deeply in art-making. A useful historical parallel is the advent of photography, and the more recent proliferation and mass adoption of smartphone cameras. Hertzmann [61, 62] argues that despite early skepticism of photography as an art-form given that photographs are made by a machine, this view began to change as people recognized that a *photographer-as-artist* uses a camera not to capture visual reality

¹¹My point here is that consequential advances in technology may not lead to advances in the *form* of dialectical engagement. But there may be other concerns as well, for example on whether dialectical forms of engagement may be crowded out professionally if their productive values cannot match that of technology-enabled forms of production to sustain itself. In other words, we have at least two concerns on the impact of technological advances on dialectical activities: (1) on how they may constrain or enable particular forms of dialectical engagement; and (2) on how they may affect production and output in ways that broaden or limit access to engaging dialectically (e.g., as a sustainable profession).

per se but instead to make deliberate choices of depiction, as other artists would.¹² Still, we would not say that anyone who snaps a picture with an iPhone is engaging dialectically, as a photographer-as-artist. Taken together, what this example highlights is the distinction among the productive capacity of a technology (the camera as a machine that makes pictures), its widespread adoption and use (smartphone cameras), and engagement in art-making (making choices in depiction). More broadly, my point is that consequentialist advances in technology do not equate to an advance in the *depth* of our engagement in dialectical activities. We are still left to search for the good in our activities, in the context of the socio-technical configuration in which we presently find ourselves.

In summary, insofar as we adopt a consequentialist view of the world, advances in generative AI are advances in our ability to produce the ends that we desire. But when viewed as a dialectical activity, what is less clear is how the dialectical activity of art-making—and all the human meaning and values embedded in it—will be preserved, diminished, advanced, or transformed as generative AI comes into the mix.

6 IMPLICATIONS FOR HCI

HCI sits between computer science and the humanities—as a field, it has a heightened responsibility to both harness the powers of computational technologies, and to retrieve the values of our human activities as they are mediated through digital technologies [138, 139]. In many ways, reconciling the consequentialist machinery that we have with the dialectical activities that we value is at the heart of HCI itself: it's about making sense of how our machines can fit into the most human aspects of our lives.

In what ways can HCI contribute to enabling and supporting dialectical activities? In this section, I first consider the ways that HCI indirectly supports dialectical activities, and then consider the challenges and opportunities for HCI to support dialectical activities more directly.

6.1 Indirect support for dialectical activities

One way to indirectly support dialectical activities is to simply continue with a consequentialist focus, but with a conscious effort toward recovering the time, space, and attention that is needed for engaging in dialectical activities. Engaging in dialectical activities requires a certain *scholé*, or leisure, through which one can deepen one's engagement into one's pursuits [8]. One way to view the mission of HCI, then, is to make more effective and efficient our meeting the various practical needs that we have, so that more of us can more frequently have the *scholé* needed for engaging in dialectical activities in our lives. In other words, instead of viewing consequentialist advances as the enemy of dialectical activities, we come to value them as an enabler of the dialectical.¹³

¹²Camera manufacturers do pre-make many depiction decisions, which affects which depiction choices are left available to the photographer [62]. Such limitations—like the new capabilities provided by a new technology—do not immediately imply that one can no longer engage in art-making, within and around these restrictions. But how a person engages in art-making—whether they are holding a camera or using generative AI—is not something that the advent nor mass adoption of the technology advances, in and of itself. It is something else altogether.

¹³Of course, it is possible that any free time created by technological advances will be filled with more labor, or forms of passive engagement that would not qualify as dialectical activities. For a critique, see Fallman [48] and Strong & Higgs [124].

In a similar way, HCI can help to meet practical needs that arise in the midst of dialectical engagement. For example, consider my engaging in this research. While my focus is on coming to a clearer understanding of the various ideas presented in this paper, in the midst of my work I found the need to find additional research articles across diverse areas; externalize and review my thoughts across physical and digital media; coordinate and schedule meetings to converse with helpful colleagues; and so on. While none of the computational tools I used to support these activities can ensure or replace my striving to further clarify, deepen, and express my understanding, without them, more of my time and attention would have gone toward meeting these consequentialist needs, than dedicated toward my dialectical engagement in the research itself.

Another direction may be to realize that while computers cannot be the be-all and end-all for supporting dialectical activities, they can make certain dialectical activities more accessible, to more people, and across more diverse media. As a simple example, consider engaging with a conversational partner who is on the other side of the world. Teleconferencing technologies do not ensure deep conversations, but they can still make the opportunities for having them more accessible. Moreover, online communities [81, 107] can create new technologically-mediated spaces for dialectical engagement. Of course, it is possible that in the process of making an activity more accessible, we lose some of the essence of that activity itself: we may not engage in deep conversations when using remote technologies as we would in person, or engage in online communities as thoughtfully as we would in our local communities. But these questions can be examined critically (e.g., [64]), and it is largely our job as HCI researchers and practitioners to attend to such issues.

Yet another direction is to refocus our attention so that we can engage in dialectical activities. In our modern lives, our attention is increasingly fragmented, and distracted by the many activities in which we partake [29]. HCI research on how attention can be recovered (e.g., from consuming social media [129]) or transitioned (e.g., across tasks [76] or at the end of the workday [135]), and popular solutions such as Apple's Screen Time, are examples of technological attempts to help us reclaim our ability to direct our attention as we wish. Future work can more directly examine the challenges to cultivating and sustaining the quality of attention required *during* one's engagement in a dialectical activity, and whether technology may have a role to play there as well.

6.2 Directly supporting dialectical activities

While the aforementioned directions of inquiry are sensible, how might HCI move beyond its focus on producing desired states to directly supporting engagement in a dialectical activity, and the understanding of meaning in a human activity that it affords? What are the challenges in designing HCI systems to promote the kind of quality attention and engagement that we value? On a more personal level, as HCI practitioners and researchers, what must we understand and embody ourselves, in order for us to support other people to engage in dialectical activities?

For starters, it may help to adopt a critical and reflective stance to examine why dialectical activities are sometimes left out, and to imagine how they can be more adequately supported [119]. Part

of this reflection will study what makes engaging in dialectical activities and seeing their value challenging—what gets in the way of proper engagement, and going beyond our routine ways of seeing and doing? Another part will examine how dialectical activities are broadly facilitated—or lost sight of—in our learning institutions, communities, workplaces, and lives (e.g., [23, 25, 56, 63, 87, 137]). Yet another part of the reflection will center on the desired role of technology in dialectical activities. Is the point to make such activities more approachable and sustainable? Is it to help people engage more deeply, or to more quickly grasp an activity's ideals? Reflecting on what it is exactly that we are hoping to support can help us to clarify the value of dialectical engagement in our lives. This reflection may lead us, for example, to conclude that our aim isn't so much to make dialectical engagements so easy or intuitive so as to not require cultivating (as can often be made the focus of design [48, 108, 124]), but to instead make its requisite training, development, and practice more broadly accessible.

I present a few challenges from my own reflections on directly supporting dialectical activities in HCI. One core challenge is that design and HCI methods are often grounded in consequentialist language and modes of thinking. For example, design methods typically emphasize coming up with a problem, identifying user needs, goals, and obstacles, and coming up with characteristics of a design or system (means) that overcome the obstacles to achieve the desired goals (ends) [44, 111]. But this language—which I and others use to do our design thinking and teaching—is more often than not oriented consequentially: embedded within it is a world-making impulse [115].¹⁴

To address this challenge, we will need language and methods that are oriented toward understanding dialectical activities, that describe the ways we'd like to be as we engage in an activity—and ways to operationalize that in practice, as HCI designers and researchers. For instance, instead of focusing on what users want to produce or achieve, we will want to understand ideal forms of engagement in an activity, and what it is actually like to engage more deeply in it. For example, in the context of dating apps, instead of focusing on what users wish to achieve (e.g., find a match), we might focus on understanding what relating to another person might look like in the process of dating someone, and how people develop their capacities for relating more intimately [10]. Broadly, we will need ways to understand how a dialectical practice is built, what gives shape to it, and what gets in the way of it. We will need to understand how people bridge the gap between their aspirational form of activity to concrete actions [25], and how that in turn, can affect their understanding of the activity and its good in their lives.

* * *

A second core challenge in supporting dialectical activities is considering how a person's grasp of the meaning of an activity evolves and deepens over the course of their long-term engagement

¹⁴This problem can be exacerbated in HCI practice. For instance, to create a new venture, startup teams are asked to state their value proposition in measurable, consequentialist terms, in ways that show the impact of their creations on shaping the world. A team may very well have a vision for enabling a way of being that is dialectical in nature. But when it comes to instantiating that vision into a product, they are often forced to think in terms of the problem the product solves, the end state it brings out, and the measurable success it will have—than in terms of the modes and quality of engagement that it affords.

with the activity. Viewed this way, dialectical activities are *learning activities*. What helps us to engage with the activity changes over time, not only as our skills develop, but as our outlook on its meaning and what proper engagement should look like matures.

But HCI research programs—in particular HCI systems research—are not typically set up for the design, study, and “living with” systems and engaging in activities over longer time frames [22]. For example, exciting recent HCI research on new experiences that embody valuable ways of being (e.g., in caring interactions [93]; in experiencing one’s own body [66, 103]) beg the question of how such experiences may progress, evolve, and deepen through dialectical engagements over longer time frames. Ultimately, we will need methods for designing for and analyzing dialectical activities that view them as long-running practices that unfold across temporal trajectories. But whereas practices are generally seen as relatively stable performances [83], a dialectical activity is marked by our changing understanding of them, which can fundamentally change what we do. For example, as our understanding of what it means to be a good friend evolves, our friendship activities and how we engage in them can change substantially. This will require new methods for design and study.

* * *

A third core challenge to directly supporting dialectical activities is that the widely accepted notion of “know thy user” is likely not enough. The value of a dialectical activity is revealed over repeated, first-person engagement over a long period of time. As a designer or HCI researcher, one will likely need the personal experience of engaging and grasping at the good of a dialectical activity to know what engaging in such an activity is truly about. When we don’t do this, we are likely to attempt to solve some user problem through a consequentialist lens, based on our impoverished understanding of what ideal engagement in the activity may actually entail.

Here I follow calls for the need for first-person research in HCI [35, 66, 79, 94]. In this context, I call for designers and researchers to engage deeply in dialectical activities and to grasp at their meaning ourselves, with which to go about doing the design and technology building work. But even this is likely insufficient, as this does not necessarily give us a clear picture of how we would support another person coming to grasp the good in their activity for themselves. Successful design attempts will likely need to triangulate across the designer’s understanding of the activity and its ideals; the challenges particular people face in their engagement with the activity; and perhaps most importantly, of effective ways to promote dialectical engagement in another person.

7 A VISION OF HCI THAT EMBRACES DIALECTICAL ACTIVITIES: COMPUTATIONAL ECOSYSTEMS

If computational technologies are to have a role in advancing human values rooted in dialectical activities, we need to more deeply understand how technologies designed to achieve consequential ends can be effectively embedded in the larger socio-technical systems in which they operate [51], folded within the fabric of our dialectical activities. Are there ways of embedding technology that are particularly helpful or harmful for engaging in dialectical activities? In

what ways do our consequentialist pursuits and the technologies that support them interface with our dialectical activities?

An underlying challenge in answering these questions is uniting critical thinking [13] and system design [106]. Ideally, critical thinking about how certain human values can be advanced should go hand-in-hand with the design of HCI systems. But this is easier said than done, as critical thinkers and system designers are often concerned with different challenges. A critical thinker may focus on the values inherent in the human activity (dialectical), whereas a system designer may focus on solving a practical user problem to achieve a desired state of affairs (consequential). Focusing on one but not the other can lead to impractical solutions, or practical ones that compromise some human value. The challenge is in finding a way to work with both concerns, so that certain human values can be advanced within a practical design.

To illustrate the value of thinking dialectically *and* consequentially, let us briefly revisit Wetmore’s account of how the Amish use technology, in their quest to reinforce their values and community [134]. In Amish communities, technology is regulated to minimize distractions and its potential harmful effects on their way of life (i.e., their dialectical activities). In adopting technology, they choose ones that they believe better promote the values that they practice (e.g., humility, equality, and simplicity). In this way, the Amish engage in a critical reflective practice to preserve their culture and the dialectical activities that support it. At the same time, they are attentive to the needs of the community to sustain itself through production (consequentialist). This involves, for example, adopting some tools and increases in technology to meet the economic need to produce and sell affordable products. Taken together, by considering both consequentialist concerns and dialectical values, the Amish continually reason about how to sustain their dialectical activities in an evolving modernity.

To generally support our accounting for both the consequentialist and the dialectical perspectives in HCI, I propose a systems approach [11, 42, 51, 71, 118] to design and study, that I call *computational ecosystems*. To take a computational ecosystems approach is to consider a socio-technical system in light of both the consequential and the dialectical; the ecosystem simultaneously produces desired goods and services, *and* promotes people’s engagement in dialectical activities. Thinking computationally, we consider how activities and problem solving may be organized and distributed across people and machines, in ways that account for who can best address them (consequential) *and* for the intrinsic value of human engagement (dialectical). Thinking ecologically, we consider how ecosystem interactions may sustain people’s deepening engagement in their activities (dialectical) *and* proper ecosystem function (consequential). Taken together, this dual-view supports our thinking about technology’s role in improving the productive capacity of a system, and its role in a larger ecology that promotes dialectical engagement.

To illustrate the computational ecosystems approach to HCI research, consider my work on Agile Research Studios (ARS) [116, 143], a learning ecosystem for advancing research training. Viewed consequentially, ARS provides new socio-technical configurations of work processes, support structures, and computational tools that enable a learning community to train more students to effectively self-direct research projects given limited mentor time [143].

Viewed dialectically, ARS engages students to critically examine their ways of working and to come to see themselves differently, as they come to approach research as a way of being—as a good in its own right [116]. Thinking computationally, ARS takes a dispersed control approach to orchestrate learning and support across a network of student researchers and mentors. Doing so scales mentor time and leads to students regularly getting help (consequential), but it also reconfigures the roles of mentors and students so that every student can practice leading authentic research inquiry (dialectical). Thinking ecologically, ARS structures activities and social interactions that help to build a supportive community in which students' self-deepening engagement in research is valued (dialectical), and which sustains students making consistent research progress and learning relevant skills (consequential). In these ways, ARS is designed to promote thoughtful engagement in research as its own good, while also addressing the practical needs of students and mentors to learn and produce.

The computational ecosystems view brings into clearer focus the value of dialectical engagement, by giving them a proper face and place in our design and critical reflection. Moreover, it brings to light any potential conflicts (or synergies) across the consequential and the dialectical. For example, we can examine whether ecosystem-supported interactions are generally conducive to engaging in dialectical activities, or if they crowd them out due to an overfocus on production. By bringing into focus both a system's productive values and its non-consequential, dialectical values, we can better make design decisions that are aligned with our vision for humanity, and our needs.

The use of the ecosystem metaphor may also help us with thinking across multiple interactions, each of which may be facilitated by their respective component technologies. Instead of focusing solely on the design of components, it opens us to consider designs that support the workings of the entire ecosystem. This may include, for example, the design of “connectors” that facilitate practice across multiple ecosystem interactions [52, 74]. Finally, the ecosystem view challenges us to see ourselves and the systems in which we are embedded as living entities that can evolve, adapt, and reconfigure across time. This may help us to think more flexibly about how we engage in our activities, and the people and technologies that support them.

8 CONCLUSION

This paper examines the limits of consequentialism (and by extension, computers) and the implications for HCI. I foreground the value of dialectical activities in human life, and discuss the challenges in supporting dialectical activities within existing HCI research and practice. I reflect on the consequentialist nature of computers, and on how we can still support dialectical activities within HCI. I note how even when computers play a largely supporting and facilitating role, there are important ways in which they can support the functioning of a computational ecosystem, to enable the ecosystem to not only produce, but to sustain engagement in dialectical activities. By placing our consequentialist machines and our dialectical activities in their proper place, HCI can play a critical role in advancing a more meaningful and sustainable vision of human life.

ACKNOWLEDGMENTS

I thank Darren Gergle, Alex Allain, Elizabeth Lenaghan, Maalvika Bhat, Linh Ly, Arya Bulusu, Laura Tom, and anonymous reviewers for their comments and suggestions on drafts, and Desiree Foerster, Pedro Lopes, Aravindan Vijayaraghavan, and Talbot Brewer for helpful conversations. Attending a faculty summer writing retreat organized by Northwestern's Office of the Provost helped me to prepare an initial draft. Funding for this research was provided in part by UL Research Institutes through the Center for Advancing Safety of Machine Intelligence.

REFERENCES

- [1] Philip Agre. 1997. *Computation and human experience*. Cambridge University Press.
- [2] Morteza Akbari, Mozghan Danesh, Azadeh Rezvani, Nazanin Javadi, Seyyed Kazem Banihashem, and Omid Noroozi. 2023. The role of students' relational identity and autotelic experience for their innovative and continuous use of e-learning. *Education and Information Technologies* 28, 2 (2023), 1911–1934.
- [3] Ali Alkhatib. 2021. To live in their utopia: Why algorithmic systems create absurd outcomes. In *Proceedings of the 2021 CHI conference on human factors in computing systems*.
- [4] Ali Alkhatib and Michael Bernstein. 2019. Street-level algorithms: A theory at the gaps between policy and decisions. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*.
- [5] Ayman Alzayat, Mark Hancock, and Miguel A Nacenta. 2019. Quantitative measurement of tool embodiment for virtual reality input alternatives. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*.
- [6] Aloha Hufana Ambe, Margot Breerton, Alessandro Soro, Min Zhen Chai, Laurie Buys, and Paul Roe. 2019. Older people inventing their personal internet of things with the IoT un-kit experience. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*.
- [7] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 CHI conference on human factors in computing systems*.
- [8] Gavin Ardley. 1967. The role of play in the philosophy of Plato. *Philosophy* 42, 161 (1967), 226–244.
- [9] Nivedita Arora, Steven L Zhang, Fereshteh Shahmiri, Diego Osorio, Yi-Cheng Wang, Mohit Gupta, Zhengjun Wang, Thad Starner, Zhong Lin Wang, and Gregory D Abowd. 2018. SATURN: A thin and flexible self-powered microphone leveraging triboelectric nanogenerator. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 2 (2018).
- [10] Carolina Bandinelli and Arturo Bandinelli. 2021. What does the app want? A psychoanalytic interpretation of dating apps' libidinal economy. *Psychoanalysis, Culture & Society* 26 (2021), 181–198.
- [11] Liam Bannon. 2011. Reimagining HCI: toward a more human-centered perspective. *interactions* 18, 4 (2011), 50–57.
- [12] Jeffrey Bardzell. 2009. Interaction criticism and aesthetics. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 2357–2366.
- [13] Jeffrey Bardzell and Shaowen Bardzell. 2016. Humanistic HCI. *Interactions* 23, 2 (2016), 20–29.
- [14] Shaowen Bardzell. 2010. Feminist HCI: taking stock and outlining an agenda for design. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1301–1310.
- [15] Jean Baudrillard. 1994. *Simulacra and simulation*. University of Michigan press.
- [16] Amanda Baughan, Nigini Oliveira, Tal August, Naomi Yamashita, and Katharina Reinecke. 2021. Do cross-cultural differences in visual attention patterns affect search efficiency on websites?. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.
- [17] Eric PS Baumer, Phil Adams, Vera D Khovanskaya, Tony C Liao, Madeline E Smith, Victoria Schwanda Sosik, and Kaiton Williams. 2013. Limiting, leaving, and (re)lapsing: an exploration of facebook non-use practices and experiences. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 3257–3266.
- [18] Emma Beede, Elizabeth Baylor, Fred Hersch, Anna Iurchenko, Lauren Wilcox, Paisan Ruamviboonsuk, and Laura M Vardoulakis. 2020. A human-centered evaluation of a deep learning system deployed in clinics for the detection of diabetic retinopathy. In *Proceedings of the 2020 CHI conference on human factors in computing systems*.
- [19] Lorraine Besser-Jones. 2011. Drawn to the Good? Brewer on Dialectical Activity. *Journal of Moral Philosophy* 8, 4 (2011), 621–631.
- [20] Susanne Bødker. 1998. Understanding representation in design. *Human-Computer Interaction* 13, 2 (1998), 107–125.

- [21] Susanne Bødker. 2006. When second wave HCI meets third wave challenges. In *Proceedings of the 4th Nordic conference on Human-computer interaction: changing roles*.
- [22] Kirsten Boehner, Phoebe Sengers, and Simeon Warner. 2008. Interfaces with the ineffable: Meeting aesthetic experience on its own terms. *ACM Transactions on Computer-Human Interaction (TOCHI)* 15, 3 (2008).
- [23] Albert Borgmann. 1984. *Technology and the character of contemporary life: A philosophical inquiry*. University of Chicago Press.
- [24] Anne E Bowser, Derek L Hansen, Jocelyn Raphael, Matthew Reid, Ryan J Gamett, Yurong R He, Dana Rotman, and Jenny J Preece. 2013. Prototyping in PLACE: a scalable approach to developing location-based apps and games. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1519–1528.
- [25] Talbot Brewer. 2009. *The retrieval of ethics*. Oxford University Press, USA.
- [26] Talbot Brewer. 2021. The Aesthetic Dimension of Practical Wisdom. In *Neglected Virtues*. Routledge, Chapter 8, 163–178.
- [27] Emeline Brule, Gilles Bailly, Anke Brock, Frédéric Valentin, Grégoire Denis, and Christophe Jouffrais. 2016. MapSense: multi-sensory interactive maps for children living with visual impairments. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 445–457.
- [28] Carrie J Cai, Emily Reif, Narayan Hegde, Jason Hipp, Been Kim, Daniel Smilkov, Martin Wattanberg, Fernanda Viegas, Greg S Corrado, Martin C Stumpe, et al. 2019. Human-centered tools for coping with imperfect algorithms during medical decision-making. In *Proceedings of the 2019 CHI conference on human factors in computing systems*.
- [29] Nicholas Carr. 2020. *The shallows: What the Internet is doing to our brains*. WW Norton & Company.
- [30] Marcus Carter, John Downs, Bjorn Nansen, Mitchell Harrop, and Martin Gibbs. 2014. Paradigms of games research in HCI: a review of 10 years of research at CHI. In *Proceedings of the first ACM SIGCHI annual symposium on Computer-human interaction in play*. 27–36.
- [31] Michael Chui, Eric Hazan, Roger Roberts, Alex Singla, and Kate Smaje. 2023. The economic potential of generative AI. (2023). McKinsey & Company.
- [32] Marianela Ciolfi Felice, Marie Louise Juul Søndergaard, and Madeline Balaam. 2021. Resisting the medicalisation of menopause: Reclaiming the body through design. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.
- [33] Mihaly Csikszentmihalyi and Reed Larson. 1978. Intrinsic rewards in school crime. *Crime & Delinquency* 24, 3 (1978), 322–335.
- [34] Mihaly Csikszentmihalyi and Jeanne Nakamura. 2002. The Concept of Flow. In *Handbook of positive psychology*, Charles Richard Snyder and Shane J Lopez (Eds.). Oxford university press, Chapter 7, 89–105.
- [35] Audrey Desjardins, Oscar Tomico, Andrés Lucero, Marta E Cecchinato, and Carman Neustaedter. 2021. Introduction to the Special Issue on First-Person Methods in HCI. *ACM Transactions on Computer-Human Interaction (TOCHI)* 28, 6, Article 37 (2021).
- [36] Carl DiSalvo, Jonathan Lukens, Thomas Lodato, Tom Jenkins, and Tanyoung Kim. 2014. Making public things: how HCI design can express matters of concern. In *Proceedings of the SIGCHI Conference on Human factors in Computing Systems*. 2397–2406.
- [37] Seungwon Do, Minsuk Chang, and Byungjoo Lee. 2021. A simulation model of intermittently controlled point-and-click behaviour. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.
- [38] Paul Dourish. 2001. Seeking a foundation for context-aware computing. *Human-Computer Interaction* 16, 2-4 (2001), 229–241.
- [39] Paul Dourish. 2006. Implications for design. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*. 541–550.
- [40] Paul Dourish. 2007. Responsibilities and implications: further thoughts on ethnography and design. In *Proceedings of the 2007 conference on Designing for User eXperiences*.
- [41] Pierre Dragicevic, Gonzalo Ramos, Jacobo Bibliowicz, Derek Nowrouzezahrai, Ravin Balakrishnan, and Karan Singh. 2008. Video browsing by direct manipulation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 237–246.
- [42] Val Dusek. 2006. What is Technology? Defining or Characterizing Technology. In *Philosophy of technology: An introduction*. Blackwell, Chapter 2, 26–37.
- [43] Yogesh K Dwivedi, Nir Kshetri, Laurie Hughes, Emma Louise Slade, Anand Jeyaraj, Arpan Kumar Kar, Abdullah M. Baabdullah, Alex Koohang, Vishnupriya Raghavan, and Manju Ahuja. 2023. “So what if ChatGPT wrote it?” Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management* 71 (2023).
- [44] Matthew Wayne Easterday, Daniel G Rees Lewis, and Elizabeth M Gerber. 2016. The logic of the theoretical and practical products of design research. *Australasian Journal of Educational Technology* 32, 4 (2016).
- [45] Ziv Epstein, Céane Boulais, Skylar Gordon, and Matt Groh. 2020. Interpolating GANs to scaffold autotelic creativity. *arXiv preprint arXiv:2007.11119* (2020).
- [46] João Marcelo Evangelista Belo, Anna Maria Feit, Tiare Feuchtnr, and Kaj Grnbæk. 2021. XRgonomics: facilitating the creation of ergonomic 3D interfaces. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.
- [47] Daniel Fallman. 2010. A different way of seeing: Albert Borgmann’s philosophy of technology and human-computer interaction. *Ai & Society* 25 (2010), 53–60.
- [48] Daniel Fallman. 2011. The new good: exploring the potential of philosophy of technology to contribute to human-computer interaction. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1051–1060.
- [49] Christopher Frauenberger. 2019. Entanglement HCI the next wave? *ACM Transactions on Computer-Human Interaction (TOCHI)* 27, 1 (2019).
- [50] Batya Friedman and Peter H Kahn Jr. 2007. Human values, ethics, and design. In *The human-computer interaction handbook*. CRC press, 1267–1292.
- [51] George W Furnas. 2000. Future design mindful of the MoRAS. *Human-Computer Interaction* 15, 2-3 (2000), 205–261.
- [52] Kapil Garg, Darren Gergle, and Haoqi Zhang. 2023. Orchestration Scripts: A System for Encoding an Organization’s Ways of Working to Support Situated Work. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*.
- [53] Ryan Colin Gibson, Mark D Dunlop, Matt-Mouley Bouamrane, and Revathy Nayar. 2020. Designing clinical AAC tablet applications with adults who have mild intellectual disabilities. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*.
- [54] Elena L Glassman, Juho Kim, Andrés Monroy-Hernández, and Meredith Ringel Morris. 2015. Mudslide: A spatially anchored census of student confusion for online lecture videos. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 1555–1564.
- [55] Erik Grönvall, Sofie Kinch, Marianne Graves Petersen, and Majken K Rasmussen. 2014. Causing commotion with a shape-changing bench: experiencing shape-changing interfaces in use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2559–2568.
- [56] Olivia Gude. 2013. New school art styles: The project of art education. *Art Education* 66, 1 (2013), 6–15.
- [57] Luke Haliburton, Natalia Bartłomiejczyk, Albrecht Schmidt, Pawel W Woźniak, and Jasmin Niess. 2023. The Walking Talking Stick: Understanding Automated Note-Taking in Walking Meetings. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*.
- [58] Christina Harrington and Tawanna R Dillahunt. 2021. Eliciting tech futures among Black young adults: A case study of remote speculative co-design. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.
- [59] Steve Harrison, Deborah Tatar, and Phoebe Sengers. 2007. The three paradigms of HCI. In *Alt. Chi. Session at the SIGCHI Conference on human factors in computing systems San Jose, California, USA*.
- [60] Eric B Hekler, Predrag Klasnja, Jon E Froehlich, and Matthew P Buman. 2013. Mind the theoretical gap: interpreting, using, and developing behavioral theory in HCI research. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 3307–3316.
- [61] Aaron Hertzmann. 2018. Can computers create art? *Arts* 7, 2 (2018).
- [62] Aaron Hertzmann. 2022. The choices hidden in photography. *Journal of Vision* 22, 11 (2022).
- [63] Eric Higgs, Andrew Light, and David Strong. 2010. *Technology and the good life?* University of Chicago Press.
- [64] Jim Hollan and Scott Stornetta. 1992. Beyond being there. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 119–125.
- [65] Kristina Höök, Martin P Jonsson, Anna Ståhl, and Johanna Mercurio. 2016. Somaesthetic appreciation design. In *Proceedings of the 2016 chi conference on human factors in computing systems*. 3131–3142.
- [66] Kristina Höök, Anna Ståhl, Martin Jonsson, Johanna Mercurio, Anna Karlsson, and Eva-Carin Banka Johnson. 2015. Somaesthetic design. *interactions* 22, 4 (2015), 26–33.
- [67] Dominik Hornung, Claudia Müller, Irina Shklovski, Timo Jakobi, and Volker Wulf. 2017. Navigating relationships and boundaries: Concerns around ICT-uptake for elderly people. In *Proceedings of the 2017 CHI conference on human factors in computing systems*. 7057–7069.
- [68] Nanna Inie, Jeanette Falk, and Steve Tanimoto. 2023. Designing Participatory AI: Creative Professionals’ Worries and Expectations about Generative AI. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*.
- [69] Eunkyung Jo, Daniel A Epstein, Hyunhoon Jung, and Young-Ho Kim. 2023. Understanding the benefits and challenges of deploying conversational AI leveraging large language models for public health intervention. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*.
- [70] Deborah G Johnson and Mario Verdicchio. 2023. Ethical AI is not about AI. *Commun. ACM* 66, 2 (2023), 32–34.
- [71] Deborah G Johnson and Jameson M Wetmore. 2021. *Technology and society: Building our sociotechnical future*. MIT press.
- [72] Brett R Jones, Hrvoje Benko, Eyal Ofek, and Andrew D Wilson. 2013. IllumiRoom: peripheral projected illusions for interactive experiences. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 869–878.
- [73] Hernisa Kacorri, Kris M Kitani, Jeffrey P Bigam, and Chieko Asakawa. 2017. People with visual impairment training personal object recognizers: Feasibility

- and challenges. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 5839–5849.
- [74] Victor Kaptelinin and Liam J Bannon. 2012. Interaction design beyond the product: Creating technology-enhanced activity spaces. *Human-Computer Interaction* 27, 3 (2012), 277–309.
- [75] Naveena Karusala, David Odhiambo Seeh, Cyrus Mugo, Brandon Guthrie, Megan A Moreno, Grace John-Stewart, Irene Inwani, Richard Anderson, and Keshet Ronen. 2021. “That courage to encourage”: Participation and Aspirations in Chat-based Peer Support for Youth Living with HIV. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.
- [76] Harmanpreet Kaur, Alex C Williams, Daniel McDuff, Mary Czerwinski, Jaime Teevan, and Shamsi T Iqbal. 2020. Optimizing for happiness and productivity: Modeling opportune moments for transitions and breaks at work. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*.
- [77] Elizabeth Kaziunas, Michael S Klinkman, and Mark S Ackerman. 2019. Precarious interventions: Designing for ecologies of care. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019).
- [78] Ryan M Kelly, Yueyang Cheng, Dana McKay, Greg Wadley, and George Buchanan. 2021. “it’s about missing much more than the people”: how students use digital technologies to alleviate homesickness. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.
- [79] Josephine Klefeker, Libi Striegl, and Laura Devendorf. 2020. What HCI can learn from ASMR: Becoming enchanted with the mundane. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*.
- [80] Yasmine Kotturi, Herman T Johnson, Michael Skirpan, Sarah E Fox, Jeffrey P Bigham, and Amy Pavel. 2022. Tech Help Desk: Support for Local Entrepreneurs Addressing the Long Tail of Computing Challenges. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*.
- [81] Robert E Kraut and Paul Resnick. 2012. *Building successful online communities: Evidence-based social design*. MIT Press.
- [82] Max Kreminski and Michael Mateas. 2021. Reflective Creators. In *Proceedings of the 12th International Conference on Computational Creativity (ICCC '21)*. 309–318.
- [83] Kari Kuutti and Liam J Bannon. 2014. The turn to practice in HCI: towards a research agenda. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 3543–3552.
- [84] Stacey Kuznetsov, Carrie Doonan, Nathan Wilson, Swarna Mohan, Scott E Hudson, and Eric Paulos. 2015. DIYbio things: open source biology tools as platforms for hybrid knowledge production and scientific participation. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 4065–4068.
- [85] Ben Lafreniere, Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2014. Investigating the feasibility of extracting tool demonstrations from in-situ video content. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 4007–4016.
- [86] Gierad Laput, Chouchang Yang, Robert Xiao, Alanson Sample, and Chris Harrison. 2015. Em-sense: Touch recognition of uninstrumented, electrical and electromechanical objects. In *Proceedings of the 28th annual ACM symposium on user interface software & technology*. 157–166.
- [87] Megan J Laverty. 2015. “There is no substitute for a sense of reality”: Humanizing the humanities. *Educational Theory* 65, 6 (2015), 635–654.
- [88] Hanchuan Li, Eric Brockmeyer, Elizabeth J Carter, Josh Fromm, Scott E Hudson, Shwetak N Patel, and Alanson Sample. 2016. Paperid: A technique for drawing functional battery-free wireless interfaces on paper. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 5885–5896.
- [89] JCR Licklider. 1960. Man-machine symbiosis. *IRE Transactions on Human Factors in Electronics* 1 (1960), 4–11.
- [90] Ryan Louie, Andy Coenen, Cheng Zhi Huang, Michael Terry, and Carrie J Cai. 2020. Novice-AI music co-creation via AI-steering tools for deep generative models. In *Proceedings of the 2020 CHI conference on human factors in computing systems*.
- [91] Ryan Louie, Jesse Engel, and Cheng-Zhi Anna Huang. 2022. Expressive Communication: Evaluating Developments in Generative Models and Steering Interfaces for Music Creation. In *27th International Conference on Intelligent User Interfaces*. 405–417.
- [92] Hao Lü, James A Fogarty, and Yang Li. 2014. Gesture script: recognizing gestures and their structure using rendering scripts and interactively trained parts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1685–1694.
- [93] Jasmine Lu and Pedro Lopes. 2022. Integrating Living Organisms in Devices to Implement Care-based Interactions. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*.
- [94] Andrés Lucero, Audrey Desjardins, Carman Neustaedter, Kristina Höök, Marc Hassenzahl, and Marta E Cecchinato. 2019. A sample of one: First-person research methods in HCI. In *Companion Publication of the 2019 on Designing Interactive Systems Conference 2019 Companion*. 385–388.
- [95] Yuhan Luo, Peiyi Liu, and Eun Kyoung Choe. 2019. Co-Designing food trackers with dietitians: Identifying design opportunities for food tracker customization. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*.
- [96] Anne-Marie Mann, Uta Hinrichs, Janet C Read, and Aaron Quigley. 2016. Facilitator, functionary, friend or foe? Studying the role of iPads within learning activities across a school year. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 1833–1845.
- [97] Laura Maye, Sarah Robinson, Nadia Pantidi, Liana Ganea, Oana Ganea, Conor Linehan, and John McCarthy. 2020. Considerations for implementing technology to support community radio in rural communities. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*.
- [98] John McCarthy and Peter Wright. 2004. Technology as experience. *interactions* 11, 5 (2004), 42–43.
- [99] Microsoft. 2023. Introducing Microsoft 365 Copilot with Outlook, PowerPoint, Excel, and OneNote. <https://www.youtube.com/watch?v=ebls5x-gb0s>
- [100] Juan D Millan Cifuentes, Ayse Göker, and Andrew MacFarlane. 2014. Designing autotelic searching experience for casual-leisure by using the user’s context. In *Proceedings of the 5th Information Interaction in Context Symposium*. 348–350.
- [101] Nusrat Jahan Mim. 2021. Gospels of modernity: Digital cattle markets, urban religiosity, and secular computing in the global South. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.
- [102] Martez E Mott, Radu-Daniel Vatavu, Shaun K Kane, and Jacob O Wobbrock. 2016. Smart touch: Improving touch accuracy for people with motor impairments with template matching. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 1934–1946.
- [103] Florian ‘Floyd’ Mueller, Richard Byrne, Josh Andres, and Rakesh Patibanda. 2018. Experiencing the body as play. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*.
- [104] Michael Muller, Lydia B. Chilton, Anna Kantosalo, Charles Patrick Martin, and Greg Walsh. 2022. GenAICHI: generative AI and HCI. In *CHI conference on human factors in computing systems extended abstracts*.
- [105] Elizabeth L Murnane, Xin Jiang, Anna Kong, Michelle Park, Weili Shi, Connor Soohoo, Luke Vink, Iris Xia, Xin Yu, John Yang-Sammarato, et al. 2020. Designing ambient narrative-based interfaces to reflect and motivate physical activity. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*.
- [106] Brad Myers. 1994. Challenges of HCI design and implementation. *interactions* 1, 1 (1994), 73–83.
- [107] Elizabeth D Mynatt, Annette Adler, Mizuko Ito, Charlotte Linde, and Vicki L O’Day. 2002. The network communities of SeniorNet. In *ECSW’99: Proceedings of the Sixth European Conference on Computer Supported Cooperative Work 12–16 September 1999, Copenhagen, Denmark*. Springer, 219–238.
- [108] Don Norman. 2009. *The design of future things*. Basic books.
- [109] William Odom, Ron Wakkary, Youn-kyung Lim, Audrey Desjardins, Bart Hengeveld, and Richard Banks. 2016. From research prototype to research product. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 2549–2561.
- [110] Alisha Pradhan, Ben Jelen, Katie A Siek, Joel Chan, and Amanda Lazar. 2020. Understanding older adults’ participation in design workshops. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*.
- [111] Jenny Preece, Helen Sharp, and Yvonne Rogers. 2015. *Interaction Design: Beyond Human-Computer Interaction*. Wiley.
- [112] Marc Rettig. 1994. Prototyping for tiny fingers. *Commun. ACM* 37, 4 (1994), 21–27.
- [113] Gisela Reyes-Cruz, Joel E Fischer, and Stuart Reeves. 2020. Reframing disability as competency: Unpacking everyday technology practices of people with visual impairments. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*.
- [114] Yvonne Rogers. 2022. *HCI theory: classical, modern, and contemporary*. Springer Nature.
- [115] Daniela K Rosner. 2018. *Critical fabulations: Reworking the methods and margins of design*. MIT Press.
- [116] Sergio Salgado, Sarah Hanson, and Haoqi Zhang. 2022. Forward: A Story about Learning and Growth. <http://forward.movie>
- [117] Jeffrey Saltz, Michael Skirpan, Casey Fiesler, Micha Gorelick, Tom Yeh, Robert Heckman, Neil Dewar, and Nathan Beard. 2019. Integrating ethics within machine learning courses. *ACM Transactions on Computing Education (TOCE)* 19, 4 (2019).
- [118] Abigail Sellen, Yvonne Rogers, Richard Harper, and Tom Rodden. 2009. Reflecting human values in the digital age. *Commun. ACM* 52, 3 (2009), 58–66.
- [119] Phoebe Sengers, Kirsten Boehner, Shay David, and Joseph ‘Jofish’ Kaye. 2005. Reflective design. In *Proceedings of the 4th decennial conference on Critical computing: between sense and sensibility*. 49–58.
- [120] Phoebe Sengers, John McCarthy, and Paul Dourish. 2006. Reflective HCI: articulating an agenda for critical practice. In *CHI’06 extended abstracts on Human factors in computing systems*. 1683–1686.
- [121] Lucy A. Sparrow, Martin Gibbs, and Michael Arnold. 2021. The ethics of multiplayer game design and community management: Industry perspectives and challenges. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.

- [122] Eric Stolterman and Anna Croon Fors. 2008. Human Computer Interaction (HCI): Towards a Critical Research Position. *Design Philosophy Papers* 6, 1 (2008), 17–40.
- [123] Elizabeth Stowell, Mercedes C Lyson, Herman Saksono, Renée C Wurth, Holly Jimison, Misha Pavel, and Andrea G Parker. 2018. Designing and evaluating mHealth interventions for vulnerable populations: A systematic review. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*.
- [124] David Strong and Eric Higgs. 2010. Borgmann's Philosophy of Technology. In *Technology and the good life?*, Eric Higgs, Andrew Light, and David Strong (Eds.). University of Chicago Press, Chapter 1, 19–37.
- [125] Lucy Suchman. 1987. *Plans and situated actions: The problem of human-machine communication*. Cambridge university press.
- [126] Hyewon Suh, Nina Shahriaree, Eric B Hekler, and Julie A Kientz. 2016. Developing and validating the user burden scale: A tool for assessing user burden in computing systems. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 3988–3999.
- [127] Sharifa Sultana, François Guimbretière, Phoebe Sengers, and Nicola Dell. 2018. Design within a patriarchal society: Opportunities and challenges in designing for rural women in bangladesh. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*.
- [128] Reem Talhouk, Sandra Mesmar, Anja Thieme, Madeline Balaam, Patrick Olivier, Chaza Akik, and Hala Ghattas. 2016. Syrian refugees and digital health in Lebanon: Opportunities for improving antenatal health. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 331–342.
- [129] Jaime Teevan. 2019. Attending to What Matters. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*.
- [130] Alan M. Turing. 1937. On Computable Numbers, with an Application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society* 2nd series, 42 (1937), 230–265.
- [131] Ibo Van de Poel. 2020. Embedding values in artificial intelligence (AI) systems. *Minds and Machines* 30, 3 (2020), 385–409.
- [132] Keith Vertanen, Haythem Memmi, Justin Emge, Shyam Reyaj, and Per Ola Kristensson. 2015. VelociTap: Investigating fast mobile text entry using sentence-based decoding of touchscreen keyboard input. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 659–668.
- [133] Jenny Waycott, Frank Vetere, Sonja Pedell, Ameer Morgans, Elizabeth Ozanne, and Lars Kulik. 2016. Not for me: Older adults choosing not to participate in a social isolation intervention. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 745–757.
- [134] Jameson M Wetmore. 2007. Amish technology: Reinforcing values and building community. *IEEE Technology and Society Magazine* 26, 2 (2007), 10–21.
- [135] Alex C Williams, Harmanpreet Kaur, Gloria Mark, Anne Loomis Thompson, Shamsi T Iqbal, and Jaime Teevan. 2018. Supporting workplace detachment and reattachment with conversational intelligence. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*.
- [136] Karl Willis, Eric Brockmeyer, Scott Hudson, and Ivan Poupyrev. 2012. Printed optics: 3D printing of embedded optical elements for interactive devices. In *Proceedings of the 25th annual ACM symposium on User interface software and technology*. 589–598.
- [137] Sam Wineburg. 2010. Historical thinking and other unnatural acts. *Phi delta kappan* 92, 4 (2010), 81–94.
- [138] Terry Winograd. 1997. From computing machinery to interaction design. In *Beyond Calculation: The Next Fifty Years of Computing*. Springer-Verlag, 149–162.
- [139] Terry Winograd and Fernando Flores. 1986. *Understanding computers and cognition: A new foundation for design*. Addison-Wesley.
- [140] Alan Yusheng Wu and Cosmin Munteanu. 2018. Understanding older users' acceptance of wearable interfaces for sensor-based fall risk assessment. In *Proceedings of the 2018 CHI conference on human factors in computing systems*.
- [141] Haijun Xia, Bruno Araujo, Toví Grossman, and Daniel Wigdor. 2016. Object-oriented drawing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 4610–4621.
- [142] Qian Yang, Aaron Steinfeld, and John Zimmerman. 2019. Unremarkable AI: Fitting intelligent decision support into critical, clinical decision-making processes. In *Proceedings of the 2019 CHI conference on human factors in computing systems*.
- [143] Haoqi Zhang, Matthew W Easterday, Elizabeth M Gerber, Daniel Rees Lewis, and Leesha Maliakal. 2017. Agile research studios: Orchestrating communities of practice to advance research training. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 220–232.
- [144] Xiaoyi Zhang, Lilian de Greef, Amanda Swearngin, Samuel White, Kyle Murray, Lisa Yu, Qi Shan, Jeffrey Nichols, Jason Wu, Chris Fleizach, et al. 2021. Screen recognition: Creating accessibility metadata for mobile applications from pixels. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*.
- [145] Yang Zhang, Chouchang Yang, Scott E Hudson, Chris Harrison, and Alanson Sample. 2018. Wall++ room-scale interactive and context-aware sensing. In *Proceedings of the 2018 CHI conference on human factors in computing systems*.
- [146] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research through design as a method for interaction design research in HCI. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 493–502.