

My research on *Computational Ecosystems* advances integrative approaches to designing, building, and studying socio-technical systems that solve complex human problems and advance core human values. My research is largely multidisciplinary, and bridges across computer science (social computing, human-computer interaction, and artificial intelligence), design, learning sciences, psychology, and philosophy.

Despite the continued development of individual technologies for supporting human endeavors, major leaps in solving complex human problems will require a radical re-thinking of how to combine wedges of human and machine competencies into the design of *integrative* systems [12, 27, 4, 8]. In practice, most socio-technical solutions are built component by component, and layer by layer. Researchers and practitioners focus primarily on advancing individual technologies (i.e., components), and on designing technologies that complement existing ways of working and organizing (i.e., as a layer). While this approach can solve focused problems that are within each technology’s scope, it can fail to address challenges inherent to complex human problems or to advance human values that lie outside of technology’s reach. Taking an alternative approach, my work on computational ecosystems integrates multiple components, and designs entirely new compositions of processes, social structures, and intelligent systems that work well together to enable unique, scalable solutions to complex human problems.

Using this approach, I devised state-of-the-art solutions for community-based problem solving, learning complex skills, and connecting people at distance. To overcome the challenges that limit the efficacy and scale of existing approaches, I introduced new ways to integrate multiple sources of human and machine intelligence; foster sustainable communities and promote synergistic interactions; and shift (but not replace) the roles of human participants. In doing so, my contributions extend beyond the specific application areas to broadly advance general principles and methods for designing integrative human-machine systems.

This statement chronicles my research program, by first sharing the computational ecosystems I have designed, and then sharing my latest work to further advance ecosystem-level thinking and technologies.

1 Designing Computational Ecosystems

My research on computational ecosystems has led to integrative solutions that (1) facilitate complex planning with many knowledgeable participants; (2) solve local problems using crowds and mobility; (3) transform professionally-made artifacts into an authentic learning resource; (4) train a large number of novice researchers to self-direct complex work; and (5) create opportunities for friends and family to connect at distance. I present each contribution below. In doing so, I highlight advances in (1) *integrative computing* to organize and distribute problem solving across people and machines, in ways that account for who can best address them *and* the intrinsic value of human engagement; and (2) *ecological thinking* to create sustainable processes and interactions that support jointly the goals of ecosystem members and ecosystem function.

1.1 Facilitating Community-Based Planning

My research transformed the way we plan large events. I contributed a *community-informed planning (CIP)* approach that produces outcomes that better serve the community as a whole, significantly eases the load on organizers, and engages many participants to communicate their needs and wants.

Planning events involving thousands of people is a challenging organizational and computational problem. Organizers often lack effective processes and tools for gathering and acting on the diverse preferences, constraints, and domain knowledge held by community members that are critical for forming effective plans. Automated systems can resolve easy-to-encode hard constraints, but the plans they output are often incomplete and undesirable because they cannot account for important community goals that remain latent and tacit. As a consequence, the planning task remains arduous, time-consuming, and short-sighted.

To address this challenge, I created a computational ecosystem for community-informed planning called Cobi, that provides a general model for engaging various stakeholders in planning and tools for planning

with community input [17, 2, 7, 3]. As overall project lead, I worked with Rob Miller, Steven Dow, Juho Kim, Lydia Chilton, and others to reinvent how we plan large academic conferences. Our deployments at ACM CHI and CSCW produced three key outcomes: (1) organizers produced better schedules by resolving hundreds of previously hidden conflicts; (2) our tools reduced organizers' planning time from 100 hours to 5 hours; and (3) our process engaged 1.5k+ community members across phases of planning. This research **received a best paper honorable mention at ACM CHI [7] and a notable paper award at HCOMP [2], and for years served as ACM CHI's standard process for planning the schedule for its 3000+ attendees.**

Advancing ecological thinking, CIP contributes new approaches for engaging diverse participants in ways that benefits them and the planning process as a whole [2, 3]. Instead of treating data collection tasks as isolated problems, I designed a socio-technical process called *incentive chaining* to support participants across multiple phases of planning by using the data collected in earlier interactions to promote further data collection in subsequent interactions. Advancing integrative computing, CIP contributes new approaches for integrating multiple sources of human and machine intelligence to support planning. In one direction, Lydia Chilton and I designed a collaborative planning model and tool for flexibly moving between *crowdware* and *groupware* modes of collaboration to support changing needs over the course of a community-based problem solving process [7]. In another direction, to leverage the wealth of information collected from a CIP process, Juho Kim and I developed *community-informed, mixed-initiative interfaces* [17] that empower organizers to make effective decisions using community input, machine intelligence, and their tacit knowledge.

1.2 Producing Globally Effective Behaviors for Physical Crowdsourcing

My research also advanced how people contribute to local, communal problems. The growth of mobile devices brought about *physical crowdsourcing systems* (e.g., citizen science; ride-sharing) that connect people to tasks and lower barriers to participation. Despite successes, such systems are limited to gathering contributions via one of two approaches: *opportunistic* or *directed*. Opportunistic approaches receive input from users when and where they contribute. This lowers participation barriers, but systems cannot demand for particular tasks to be completed. Directed approaches prompt users to advance system goals, which often require contributions that are outside people's routines and thus require strong incentives. As a consequence, existing models often leave tasks incomplete, or rely on monetary incentives to ensure system outcomes.

As PI of an NSF Cyber-Human Systems award and a Microsoft FUSE Labs award, I worked with Darren Gergle, Yongsung Kim, and others to advance a hybrid approach to physical crowdsourcing that retains the best elements of both opportunistic and directed approaches. Our approach led to computational ecosystems that permit access to a large number of people who when directed to make small, convenient contributions through their existing routines, contribute actions that jointly achieve system goals.

In one direction, we developed interaction techniques that broaden participation by breaking down larger goals into small physical tasks that can be completed through people's existing routines [1, 5], and by tapping into people's physical routines directly [21]. In a complementary direction, we developed new decision-theoretic frameworks for *flexible coordination*, that decide how to opportunistically recruit a well-motivated and minimally disrupted crowd to efficiently achieve system goals. Specifically, we developed (1) the *4X* framework [11] for scaffolding data collection through phases that continually expand data fidelity and spatial coverage; (2) a *hit-or-wait* [19] framework for deciding whether to notify helpers of tasks on their route by reasoning jointly about system needs and a helper's changing patterns of mobility; and (3) a *opportunistic supply management* [18] framework for optimizing and adjusting community-level policies for recruiting helpers. By reasoning about both system needs and the goals, interests, and availability of helpers, these decision-theoretic frameworks enable a flexible approach for connecting helpers to tasks that allows helpers to contribute opportunistically when they desire, in situations when it is convenient to them, valued by the

system, and that ultimately lead to globally effective solutions. More generally, these frameworks advance ecological thinking by offering principled approaches for managing contributions across a computational ecosystem, in ways that support ecosystem members and proper ecosystem function.

1.3 Creating Readily Available Learning Experiences on the Web

While the last two contributions advance community-based problem solving, these next two advance the learning of complex skills. I discuss here our efforts to transform professionally-made artifacts into an authentic learning resource, specifically in the domain of web development.

Aspiring web developers are turning to online resources to teach themselves to code, but existing learning platforms primarily teach syntax or provide practice on constrained tutorial examples. Professional websites can provide a potential resource for authentic learning, as they offer rich details missing from training examples, embed the programming concepts and implementation techniques used by professionals, and are continually updated as new solutions arise. But despite the abundant availability of front-end code, professional examples are difficult for learners to understand. **As PI of an NSF Cyberlearning grant**, I worked with Eleanor O'Rourke, Josh Hibsman, and others to create *Readily Available Learning Experiences (RALE)* that transform the entire web of professional examples into a resource for learning programming concepts and design patterns. Advancing integrative computing, we designed a learning ecosystem that embeds software scaffolds to support self-directed learning, so that learners can more effectively investigate examples, manage their learning process, assess what they have learned, and drive further investigation.

As a first step toward RALE, we created new tools that help web developers find relevant source code. We introduced Unravel [13], a tool that helps developers find entry points for understanding UI features on professional websites in 50% less time than when using Chrome Dev Tools. We also introduced Telescope [14], a tool which constructs a composite view of the less than 150 lines of relevant code that produce a UI interaction out of the 10k-100k lines of code on professional websites. But while Unravel, Telescope, and prior tools reduce the complexity of exploring professional code, these solutions target experienced developers and are insufficient for helping novice developers build a conceptual understanding of how code constructs work together to implement a feature. To overcome such difficulties, we designed *learner-centric developer tools* that embed sensemaking scaffolds informed by the learning sciences to explicitly support novices building such conceptual models. We created Isopleth [15], a web-based platform for making sense of complex code constructs and hidden asynchronous relationships in professional web applications. We also created Ply [20], a web inspector that computes visual relevance to highlight the relationship between CSS code and observable effects on professional webpages. **This work received a best paper honorable mention at UIST, and inspired new features that were added to both Firefox and Chrome developer tools.** Results from four studies show that using these tools helped novice developers build significantly more accurate conceptual models of how professional features are implemented, and helped to surface expert design patterns across a diverse range of complex, professional websites.

1.4 Scaling Research Training to Cultivate Self-Directed Learners

My research significantly expanded our ability to train novice researchers to self-direct complex work. My approach cultivates self-directed learners, significantly increases the number of students who receive authentic research training, and lowers the barrier to participating in authentic research experiences.

Apprenticeship, or 1-on-1 mentoring, provides a powerful model for training young researchers, but imposes a heavy orchestration burden on a single faculty researcher who can only mentor a handful of students at a time. Novice researchers lack *regulation skills*, i.e., cognitive, motivational, emotional, metacognitive, and strategic behaviors for reaching desired goals and outcomes. Leading a research project requires students to develop regulation skills to plan research work and to adopt effective help-seeking and collaboration

strategies to overcome challenges. Training and support for developing such skills is necessary but time-consuming for mentors and difficult to support with software alone. As a consequence, most students are relegated to rote tasks and only a few advanced students are trained to direct complex work themselves.

As PI of an NSF Cyberlearning award, I worked with Matt Easterday, Liz Gerber, Leesha Maliakal, and others to design *Agile Research Studios (ARS)* [31, 28], a computational ecosystem for research training that disperses control of the learning process across the entire community. By reconceptualizing the role of students as self-directed learners and distributing support across the community, ARS frees up mentors to respond to challenges they can best address and to teach regulation skills. Applying integrative computing and ecological thinking, ARS provides new socio-technical configurations that address the core computational bottlenecks in scaling by shifting the roles of faculty members, graduate students, and undergraduates to support their respective needs and those of other ecosystem members. Results from a two-year study [31] show that ARS helped to not only engage more students in research, but helped students develop the planning skills and help-seeking dispositions needed to self-direct complex work. For a more complete view of the significant impacts my work on ARS has had, see Section 3.

1.5 Creating Opportunities for Connecting at Distance

My latest computational ecosystem creates opportunities for connecting friends and family at distance. Although there has been extraordinary growth in social technologies to support connection, evidence of social benefits from using such technologies is—at best—mixed. Simple passive engagement on social media (e.g., scrolling feeds) is associated with negative effects while active engagement (e.g., commenting on posts) is associated with positive effects such as decreased loneliness. In other words, human social needs cannot simply be met by having a larger number of connections or passive engagement with others, but instead require meaningful interactions in order to best support core human social needs around connection.

To help address this challenge, together with Darren Gergle and Ryan Louie, I designed a new computational ecosystem called *Opportunistic Collective Experiences (OCEs)* [24], which finds serendipitous moments for engaging in shared activities at a distance. In contrast to existing social technologies which are agnostic to the situation that people are in, people’s interactions in OCEs are made possible by the situations and environments that they find themselves in. By identifying shared situations across distributed contexts, an OCE may include many of the interactional elements available in co-located interactions, such as the ability for people to have similar experiences or attend to shared stimuli. Advancing integrative computing, we contributed (1) interaction structures and theoretically-informed design guidelines that extend how co-located interactions can promote social closeness to people who are not co-located; (2) a programming model that enables experience creators to write concise OCE programs; and (3) an opportunistic execution engine that continually checks for opportunities for social interaction. Through three deployment studies, we found that OCEs are more connecting than posting and commenting on social media, lower barriers to interacting with others, and provide an intimate and engaging way to interact across distance.

2 Addressing Design Limitations by Advancing Ecosystem-level Thinking and Technologies

Beyond designing computational ecosystems, my post-tenure research sought to understand and address the core limitations to designing integrative, socio-technical solutions to support human activities and values. This has led to (1) advances in ecosystem-level thinking and technologies to support building a practice; (2) advances in interfacing between humans and AI systems; and (3) new theory on the limits of computers in supporting intrinsically valuable human activities. I highlight these contributions below.

2.1 Ecosystem-Level Supports for Building a Practice

The first track of my post-tenure research advances ecosystem-level thinking and technologies to support building a practice. Specifically, I contribute a *process management framework* to support learning a practice

across multiple support subsystems, and computational abstractions of a situated practice that enable the creation of *situated co-pilots* to support complex work.

A complex practice has many components and facets. For example, when planning research work, an expert researcher doesn't just plan; they represent and visualize the problem, they diagnose risks, and they focus plans on addressing key risks. To scaffold novice research learners building and integrating these component skills to develop an effective planning practice, Leesha Maliakal and I developed multiple planning support subsystems into the Agile Research Studios learning ecosystem. But as the number of supports for scaffolding a complex practice grew in complexity, we observed breakdowns in how learners practice across subsystems of supports [29, 9]. To help learners strategically adapt and execute effective planning processes, we introduced a *process management framework* and tool called Compass to help learners track and update their plans across their interactions in a learning ecosystem [29]. An 8-week deployment study comparing student planning processes with and without Compass showed that students using Compass revised their plans more following support opportunities, and that mentor assessed their plans to be more structurally aligned across their practice. In other work, we found that process management scaffolds are also effective for learners recognizing and addressing their metacognitive risks [26].

Complex practices are necessarily situated in rich physical and digital contexts, and often involve different people and contexts over time. Supporting a situated practice with technology has historically been challenging, largely because we lack *ecosystem-level supports* that can monitor, coordinate, and restructure interactions within a socio-technical system to support building a practice over time. My work on designing computational ecosystems does not address this challenge, as it largely focuses on constructing assemblages of component supports that work well together, but not on supporting one's practice within the complex ecosystems that are constructed [9]. To address this gap, Kapil Garg and I developed computational abstractions of an organization's ways of working, that allow for modeling, tracking, and suggesting situated practices in software [10]. We used these abstractions to create Orchestration Scripts, a system for encoding and tracking signals of emerging work situations and the strategies for addressing them in relevant situations. In a field study, we found that Orchestration Scripts created opportunities to practice individual and collaborative strategies for addressing emerging work situations that previously would have gone unnoticed.

In ongoing work, we are continuing to develop richer computational models and co-pilots to support people's situated practice within a computational ecosystem. A key idea is to leverage both the automated tracking and facilitation of interactions afforded by computational agents, and to allow people to refine their own understanding as they develop a personal practice. Broadly, this line of work demonstrates how computational abstractions that model the interactions within a computational ecosystem allows us to provide in-situ supports for a situated practice, in ways that designing a computational ecosystem does not.

2.2 Human-AI Interface Layers

The second track of my post-tenure research advances the interface between humans and AI systems, particularly for supporting what are inherently human activities and experiences. While early uses of artificial intelligence (AI) aimed to automate repetitive and burdensome tasks, modern AI systems show great potential in assisting people in human activities imbued with personal meaning and importance, such as creative pursuits and social experiences. For instance, Opportunistic Collective Experiences use context-aware AI technologies to surface opportunities for friends and family to connect at distance. Such AI capabilities can lower the barriers for people to engage in creative and social activities, but using AI in these human domains requires the AI to have a deeper understanding and sensitivity to a person's own ideas. As **PI of a Google Faculty Research Award**, I worked with Darren Gergle and Ryan Louie to contribute a crucial layer in an AI-powered application's stack called the *Human-AI Interface Layer*, which provides intuitive

constructs and interactive tools that empower people to more effectively communicate their intentions to AI systems. Realizing this approach, Ryan’s dissertation developed (1) a *steering* interface layer that partitions and constrains AI-generated outputs to support creators incrementally guiding outputs [23]; (2) an *expression* interface layer equipped with cognitive bridging tools to help designers flesh out their high-level concepts and represent them using AI constructs [25]; and (3) an *execution* interface layer for stating and surfacing implicit expectations so an AI agent can appropriately adjust after handoff [22]. Largely, our research demonstrates how human-AI interface layers enable the design of new socio-technical configurations that effectively use and integrate AI capabilities to support personally meaningful human activities.

We have also begun to explore how AI systems can help humans form and clarify their own ideas for human activities and experiences. As **PI of a Center for Advancing Safety of Machine Intelligence (CASMI) grant**, I am working with Darren Gergle and students to advance new human-AI tools that: (1) help designers form rich and accurate conceptions of human situations to encode into machines systems; (2) provide support for refining concept expressions; and (3) identify and address issues of difference across settings and contexts of use. By supporting forming an integrated human-machine understanding of a human situation, our approach broadly advances socio-technical solutions that reflect the entirety of the human-machine system, and not just the technical and algorithmic components.

2.3 The Limits of Computers for Supporting Intrinsically Valuable Human Activities

The third track of my post-tenure research builds new theory to expose the limits of computers in supporting intrinsically valuable human activities. It also clarifies the important role that computational ecosystems play in not only solving complex human problems, but in promoting and sustaining core human values.

As computing deepens its reach into every aspect of society, there is a real opportunity for it to advance the richness of human life, but it would be naive to think that technology alone can advance the human values we hold dearest. My CHI 2024 paper [30], *Searching for the Non-Consequential: Dialectical Activities in HCI and the Limits of Computers*, provides a philosophical argument for why computers—consequentialist machines that reliably transform inputs into desired outputs—cannot be the be-all and end-all for advancing intrinsically valuable human activities whose values do not reside solely in ends to be produced. The paper provides theoretical explanations for how the consequentialist nature of computers limit how they can be used to support intrinsically valuable human activities that require continual, self-deepening reinterpretations and actions in search of the good of the activity (e.g., in being a good friend, art-making, ethical reasoning, conducting research). The paper further argues for the value of computational ecosystems in promoting people’s engagement in intrinsically valuable activities, than simply reaching consequential outcomes. Whereas my earlier thinking on computational ecosystems had focused primarily on its role in advancing (consequentialist) problem solving, **this paper and my post-tenure work largely advance our thinking on how computational ecosystems play an important in promoting and sustaining human values and ways of being that cannot be, and should not be, replaced by computing technologies.**

3 Research Impact and Future Directions

My research, to date, has resulted in 40 major publications with 1800+ citations, leading to an h-index of 22 (using Google Scholar). My work is published at premier venues in HCI (CHI, TOCHI, UIST), AI (AAAI, IJCAI), Social and Crowd Computing (CSCW, HCOMP), and the Web (WWW), and has been recognized with 4 notable paper awards or best paper honorable mentions. It is supported by 5 awarded NSF grants (4 on which I serve as PI), a Center for Advancing Safety of Machine Intelligence award, a grant from the Buffett Institute of Global Affairs, a Google Faculty Research Award, a Microsoft FUSE Labs award, eight Murphy Society Awards, and an Office of the Provost Award for Digital Learning. To broaden the reach of my research, I have given invited talks at 12 major research institutions and at the AAAS annual meeting.

Since tenure review I have published eight significant papers at major conferences and journals (TOCHI, CHI, CSCW), written two letters on mentoring and learning, and produced the documentary film “Forward: A Story about Learning and Growth.” I entered tenure as PI on three active, on-going NSF projects; since I have acquired an additional \$1.45 million of funding. I also graduated 3 PhD students: Yongsung Kim (postdoc at CMU; Northwestern Communication Studies Dissertation Award); Ryan Louie (postdoc at Stanford; Google PhD Fellowship); and Leesha Maliakal (Assistant Professor of CS, Northeastern Illinois University).

My research impact is best demonstrated through my work on the Agile Research Studios (ARS) model for research training, which I use to direct the Design, Technology, and Research (DTR) program at Northwestern. Through DTR I have mentored 160 students to self-direct independent research projects (76 since tenure). 73 of my mentees are female (46%), far surpassing the national average of 22% as reported in the latest CRA Taulbee Survey. My students published 28 papers at major computing research conferences (13 since tenure) and placed at 7 major student research competitions, including 1st and 2nd at the 2022 ACM CHI Student Research Competition. My students won 65 undergraduate research grants (26 since tenure) from Northwestern, the most of any research lab or group on campus. 40% of DTR undergraduate students (44 out of 110) placed at Microsoft (15), Apple (12), Google (10), Meta (5), and Amazon (2). Since tenure, **my efforts to advance undergraduate research mentoring have been recognized with the Office of Undergraduate Research Faculty Honor Roll, two departmental nominations for the CRA-E Undergraduate Research Mentoring Award, and an invitation to join the CRA-E board of directors.**

Beyond directing DTR and developing the ARS model at Northwestern, I founded the Agile Research University program (<http://agileresearch.io>), which to date, has supported 15 faculty across institutions on site visits and 70+ faculty who use the research mentoring tools, resources, and starter kits that I has developed. My research mentoring methods and best practices in DTR has since been produced and disseminated through a documentary film [28] (see <http://forward.movie>), which has engaged hundreds of faculty and students across 30+ institutions. They have even been adapted for use in the humanities [16]. I also write an annual letter to foster dialogue in academic communities and innovation on topics related to mentoring and learning (<http://dtr.northwestern.edu/letters>), and facilitate a cross-institutional support group for junior faculty in computing (<http://haoqizhang.com/group>) to support faculty mentors across 12 institutions. Senior faculty readers of my annual letters describe them as ‘thought-provoking,’ ‘enlightening,’ ‘open,’ and ‘honest;’ junior faculty attendees of my support group note that it “was a source of comfort for me at times when I felt overwhelmed,” and provided “an enriching and safe space.”

Broadly, my study of computational ecosystems charts a course through which we can better align the systems we build with the values we hold. It calls on us to envision solutions in which people and technology are part of a larger ecology, and that bring to bear all of our human and machine abilities to deal with the complexities of the world, and to advance fundamental human values and ways of being. To complement my technical work, I have continued to seek a deeper understanding of core human values, and how they can be better supported and advanced. Specifically, post-tenure I have assembled two multi-disciplinary teams to (1) study how people learn to lead a self-directed life, through leading a Buffett Institute Idea Dialogue and a Spencer Foundation proposal on "Fostering Self-Direction in Human Living," and (2) advance women’s leadership development at universities globally, through co-leading a Buffett Institute of Global Affairs working group with Jennifer Tackett in Psychology. These teams bring together significant expertise across computer science, design, embodiment and cultural theory, engineering, entrepreneurship, health equity, learning sciences, philosophy, and psychology. My expectation is that these collaborations will provide in-depth understanding of two core human activities across multiple contexts and perspectives that sets the groundwork for developing integrative, socio-technical solutions that improve how diverse people learn, across settings, to self-direct their lives and become leaders in their communities.

References

- [1] Zachary Allen. Gaze: Using mobile devices to promote discovery and data collection. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, pages 97–102. ACM, 2015.
- [2] Paul André, Haoqi Zhang, Juho Kim, Lydia Chilton, Steven Dow, and Rob Miller. Community clustering: Leveraging an academic crowd to form coherent conference session. In *Proc. HCOMP 2013*, 2013.
- [3] Anant Bhardwaj, Juho Kim, Steven P. Dow, David Karger, Samuel Madden, Robert C. Miller, and Haoqi Zhang. Attendee-sourcing: Exploring the design space of community-informed conference scheduling. In *Proc. HCOMP 2014*, 2014.
- [4] Anthony S Bryk, Louis M Gomez, Alicia Grunow, and Paul G LeMahieu. *Learning to improve: How America’s schools can get better at getting better*. Harvard Education Press, 2015.
- [5] Stephen Chan, Nicole Zhu, and Haoqi Zhang. Tapshare: Participatory communitysensing through physical gestures and mobile sensors. Technical report, Northwestern University, 2015.
- [6] Kevin Chen and Haoqi Zhang. Remote paper prototype testing. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 77–80. ACM, 2015.
- [7] Lydia Chilton, Juho Kim, Paul Andre, Felicia Cordeiro, James Landay, Dan Weld, Steven Dow, Rob Miller, and Haoqi Zhang. Frenzy: Collaborative data organization for creating conference sessions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2014.
- [8] George W Furnas. Future design mindful of the MoRAS. *Human-Computer Interaction*, 15(2):205–261, 2000.
- [9] Kapil Garg, Darren Gergle, and Haoqi Zhang. Understanding the practices and challenges of networked orchestration in research communities of practice. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2):1–28, 2022.
- [10] Kapil Garg, Darren Gergle, and Haoqi Zhang. 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*, pages 1–17, 2023.
- [11] Kapil Garg, Yongsung Kim, Darren Gergle, and Haoqi Zhang. 4x: A hybrid approach for scaffolding data collection and interest in low-effort participatory sensing. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–28, 2019.
- [12] Atul Gawande. Atul Gawande: How do we heal medicine. *TED video*, 19:19, 2012.
- [13] Josh Hibsichman and Haoqi Zhang. Unravel: Rapid web application reverse engineering via interaction recording, library detection, and source tracing. In *Proceedings of the 28th annual ACM symposium on user interface software and technology*, UIST ’15. ACM, 2015.
- [14] Josh Hibsichman and Haoqi Zhang. Telescope: Fine-tuned discovery of interactive web ui feature implementation. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, pages 233–245. ACM, 2016.

- [15] Joshua Hibsichman, Darren Gergle, Eleanor O’Rourke, and Haoqi Zhang. Isopleth: Supporting sense-making of professional web applications to create readily available learning experiences. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 26(3):1–42, 2019.
- [16] EB Hunter. Structuring courses with agile research studio: Five components and four pedagogical values. *Theatre Topics*, 31(1):31–42, 2021.
- [17] Juho Kim, Haoqi Zhang, Paul André, Lydia B. Chilton, Wendy Mackay, Michel Beaudouin-Lafon, Robert C. Miller, and Steven P. Dow. Cobi: A community-informed conference scheduling tool. In *Proc. UIST 2013*, 2013.
- [18] Yongsung Kim. *Designing Flexible Coordination Systems to Advance Individual and Collective Goals in Physical Crowdsourcing*. PhD thesis, Northwestern University, 2020.
- [19] Yongsung Kim, Darren Gergle, and Haoqi Zhang. Hit-or-wait: Coordinating opportunistic low-effort contributions to achieve global outcomes in on-the-go crowdsourcing. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, page 96. ACM, 2018.
- [20] Sarah Lim, Joshua Hibsichman, Haoqi Zhang, and Eleanor O’Rourke. Ply: A visual web inspector for learning from professional webpages. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, pages 991–1002, 2018.
- [21] Katherine Lin, Henry Spindell, Scott Cambo, Yongsung Kim, and Haoqi Zhang. Habitsourcing: Sensing the environment through immersive, habit-building experiences. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, pages 639–650. ACM, 2016.
- [22] Ryan Louie. *Human-AI Interface Layers: Enhancing Communication of Intent for AI-Assisted Creative Pursuits and Social Experiences*. PhD thesis, Northwestern University, 2023.
- [23] Ryan Louie, Andy Coenen, Cheng Zhi Huang, Michael Terry, and Carrie J Cai. Novice-ai music creation via ai-steering tools for deep generative models. In *Proceedings of the 2020 CHI conference on human factors in computing systems*, pages 1–13, 2020.
- [24] Ryan Louie, Kapil Garg, Jennie Werner, Allison Sun, Darren Gergle, and Haoqi Zhang. Opportunistic collective experiences: Identifying shared situations and structuring shared activities at distance. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW3):1–32, 2021.
- [25] Ryan Louie, Darren Gergle, and Haoqi Zhang. Affinder: Expressing concepts of situations that afford activities using context-detectors. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–18, 2022.
- [26] Molly Pribble, Neha Sharma, Haoqi Zhang, and Leesha Maliakal Shah. Mindyoga: Scaffolding the metacognitive reflection process within learning ecosystems. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, pages 1–8, 2022.
- [27] Hope Reese. Mastery of AI has been ‘harder than expected’ and ‘future is uncertain,’ says Microsoft’s AI chief. *TechRepublic*, 2015.
- [28] Sergio Salgado, Sarah Hanson, and Haoqi Zhang. *Forward: A Story about Learning and Growth*. FurnaceFPS, 2020.

- [29] Leesha Maliakal Shah. *Agile Research Studios: Learning Ecosystems to Scale Effective Research Training*. PhD thesis, Northwestern University, 2023.
- [30] Haoqi Zhang. Searching for the non-consequential: Dialectical activities in HCI and the limits of computers. In *Proceedings of the 2024 CHI conference on human factors in computing systems*, pages 1–13, 2020.
- [31] Haoqi Zhang, Matthew W Easterday, Elizabeth M Gerber, Daniel Rees Lewis, and Leesha Maliakal. Agile research studios: Orchestrating communities of practice to advance research training. In *Companion of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, pages 45–48. ACM, 2017.