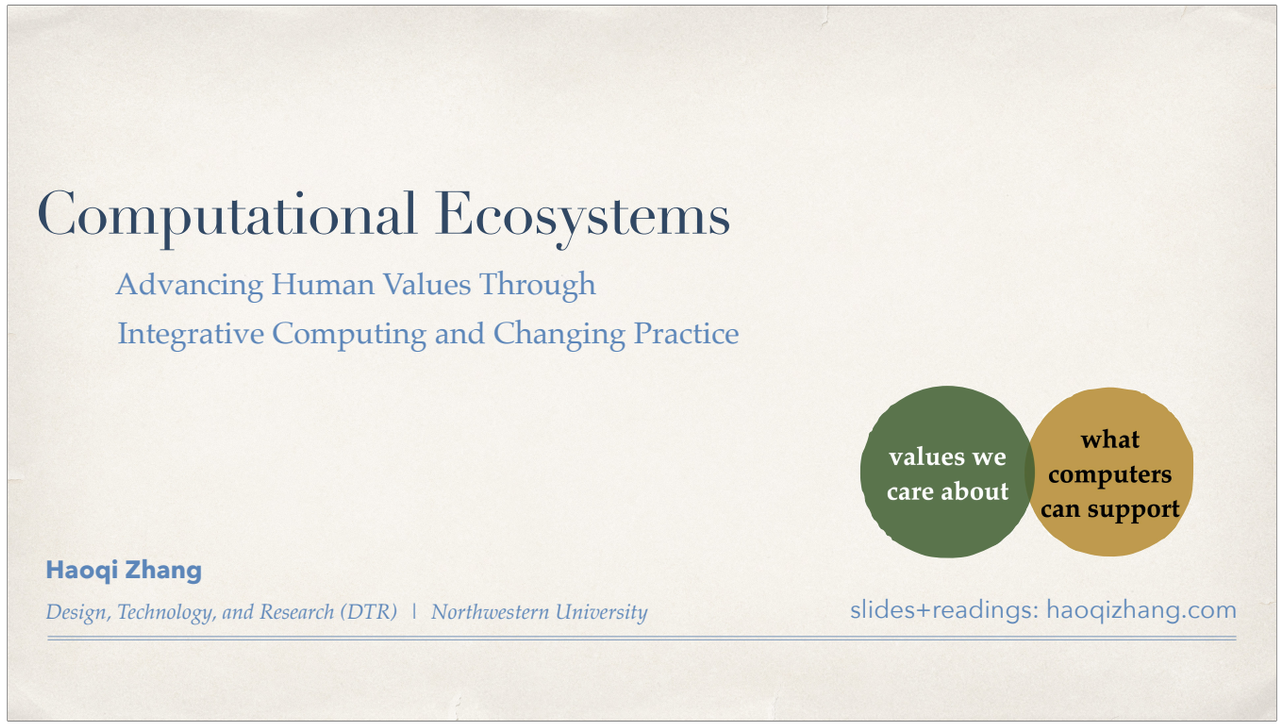

Thank you, and thank you all for coming.

I gave this talk at lots of schools earlier this year, and after an amazing trip, it's really great to be back for homecoming, here with you all.

And today, I want to share with you my journey to align my work with my values.

My hope for this talk, is not that you walk out embracing my values, but instead, that you think about yours. I think that'd be really wonderful.

And with that, let's talk about computational ecosystems, and what they can do for advancing human values.

If impact = winning, computing has won.

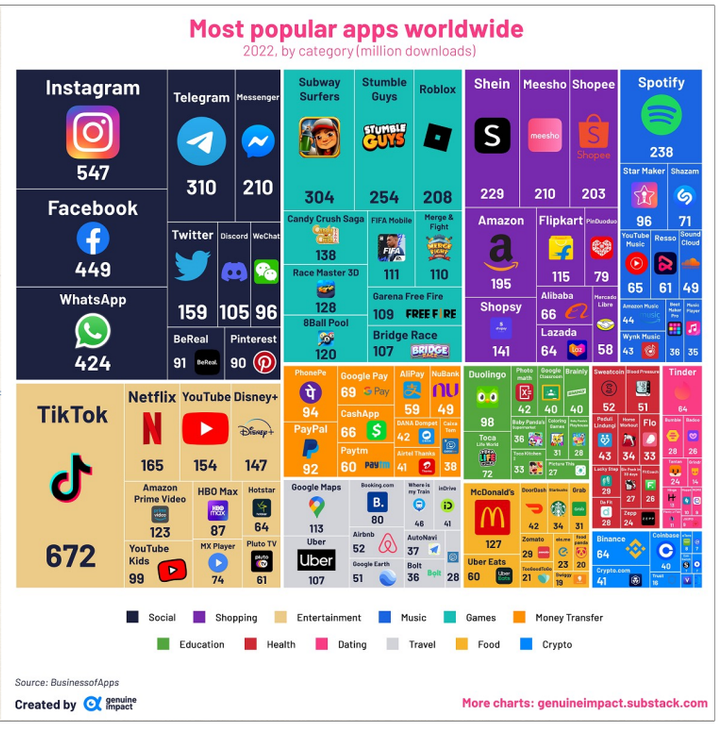

To start my talk, I'd like to pump us up a bit, by assuring you that computing has made a tremendous impact on the world.

If you look at the figure on the right, it certainly looks like computing has touched every facet of human life, and that we are all impacted, if not benefitting, from its advances.

How can we solve human problems and advance human values that remain unaddressed, even after decades of advances in computing?

But when it comes to my research, I keep hitting upon this question: ___

As someone who designs technologies, this is a really uncomfortable question to be asking, especially when computing is totally winning.

But with computing moving so fast…I think it's easy to focus on progressing, than on taking a step back and looking at the areas where despite advances elsewhere, we are no better for it.


Design, Technology, and Research (DTR)

As a small example of this, I want to briefly show you a picture of my students in the Design, Technology, and Research program here at Northwestern, where my students learn to innovate through independent research. I have had many successful students over the years, but my question here is simple:

How come there is no technology that helps me mentor them better? [Pause]

Despite decades of computing advances, why is the best mentoring technology a Notepad? What is a “mentoring technology,” anyway?

world. In this book I would like to contribute to a *critical technical practice* in which rigorous reflection upon technical ideas and practices becomes an integral part of day-to-day technical work itself.

Agre, *Computation and Human Experience*, 1997.

Questions like this got me thinking that maybe the issue isn't a lack of technology, but a lack of reflection on how we design technology, or what Philip Agre calls, a critical technical practice..

And this is the kind of reflection and practice my students and I have been engaged in for over a decade now. Through it, we learned some really important things.

1. There are fundamental limitations in our current approach.
2. We cannot overcome these limitations with better technologies.
3. We need a new approach.
4. Enabling this new approach requires a radical rethinking of the entire sociotechnical configuration.

The first thing we learned, is that for many human problems, there are fundamental limitations in our current approach and ways of doing things…..

1. There are fundamental limitations in our current approach.
2. We cannot overcome these limitations with better technologies.
3. We need a new approach.
4. Enabling this new approach requires a radical rethinking of the entire sociotechnical configuration.

…that cannot be overcome with better technology.

If we could, we would’ve done it already.

1. There are fundamental limitations in our current approach.
2. We cannot overcome these limitations with better technologies.
3. We need a new approach.
4. Enabling this new approach requires a radical rethinking of the entire sociotechnical configuration.

What we need is a new approach, a new way of doing things. That what's broken is not only our technologies for supporting our current way of doing things, but our way of doing things, in and of itself.

1. There are fundamental limitations in our current approach.
2. We cannot overcome these limitations with better technologies.
3. We need a new approach.
4. Enabling this new approach requires a radical reworking of the entire sociotechnical configuration.

Finally, we learned that we cannot enable this new approach without a radical reworking of the sociotechnical configuration. That ultimately, what we need to redesign is not simply a technology, but the entire sociotechnical ecosystem in which we operate.



To illustrate the value of thinking this way, I want to start by showing you five human activities we made better through our approach, starting with,

PLANNING LARGE EVENTS



Paul André
CMU



Anant Bhardwaj
MIT



Lydia Chilton
UW



Juho Kim
MIT



Steven Dow
CMU



David Karger
MIT



Rob Miller
MIT



Haoqi Zhang
Northwestern

Starting with planning large events.

This is work I led over 10 years ago, with some folks that you may recognize, when we were all a lot younger.

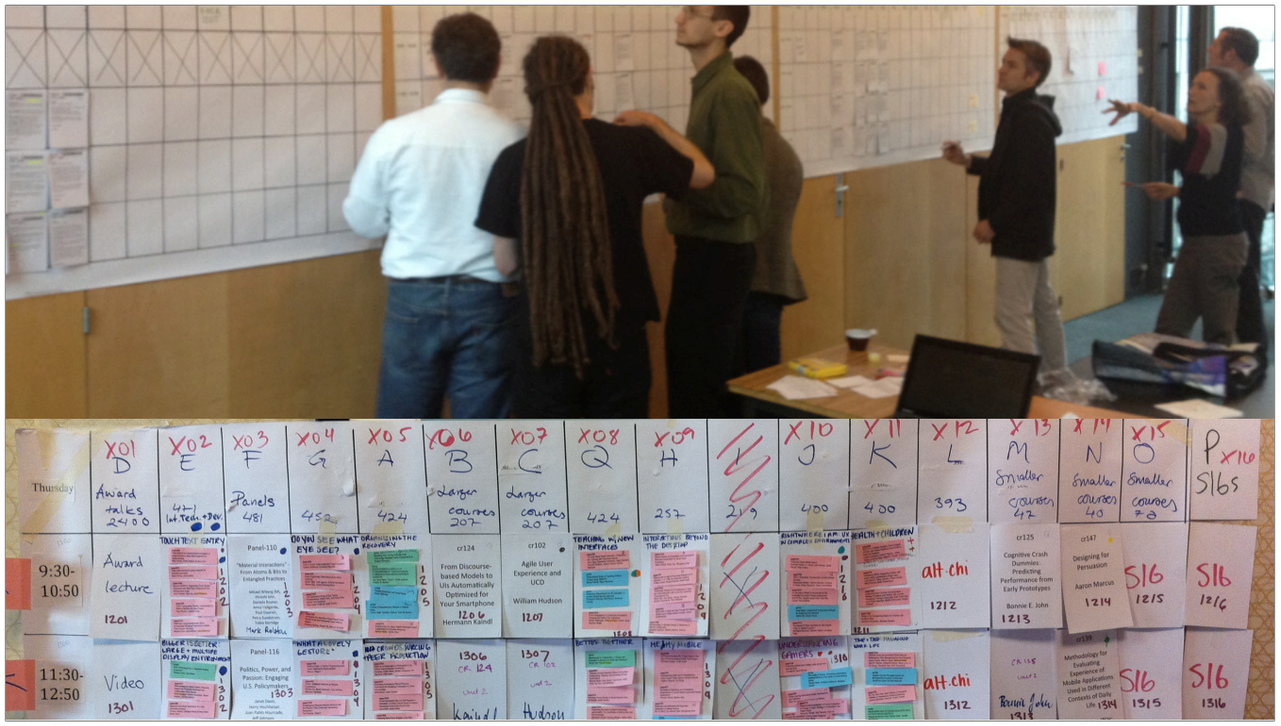

So let’s get into it. What I am showing you here is a handful of ACs working on a preliminary schedule for CHI, where they try to construct hundreds of sessions on paper.

Immediately, you see the limits of this approach.

- The first is that organizers lack information about the preferences, constraints and knowledge held by the larger community.
- The second, is that they lack tools for managing the complexity of the schedule. The problem isn’t that there are no optimizers around, but that without a clear sense of how papers are related and of what preferences and constraints people have, it doesn’t matter how good your optimizer is.

So instead, we reimagined the process of planning itself,

New approach: community-informed planning

committeesourcing

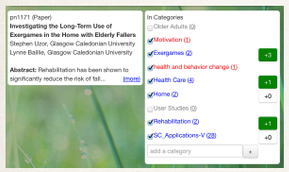

make sessions

[Chilton et al., CHI '14]

expert categories

authorsourcing

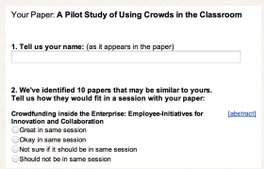

collect affinities

[Andre et al., HCOMP '13]

affinities as seeds

attendeesourcing



collect preferences

[Bhardwaj et al., HCOMP '14]

And came up with a new approach that we call, community informed planning.

Our first core idea, is to engage the entire community in the planning process, whether that's asking committee members to make sessions, collecting paper affinities from authors, or gathering the preferences of attendees.

As we did this, we thought about ways to collect information from certain community members, such as authors, in ways that help attendees, and that promote their contributing even more useful data.

Community-informed, mixed-initiative interface

[Kim et al, UIST 2013]

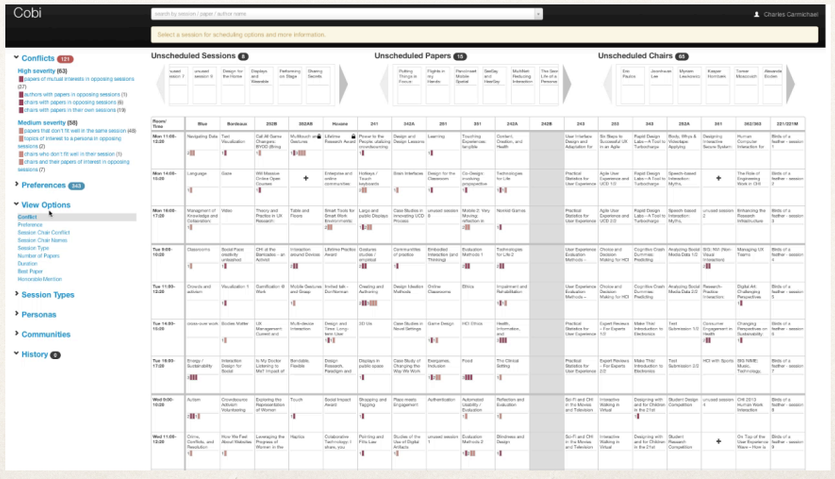

Our second core idea is a community-informed, mixed-initiative interface, that puts organizers in the driver's seat with all the community data we've gathered.

[demo]
This tool gives the organizers a multidimensional view of the schedule, and allows them to make changes that are based on their goals and tacit knowledge, the community's input, and the system's intelligence.

So here you see a session the organizers are trying to move…… based on what the system knows about. This provides organizers with guidance on what may be good options and what not to do, while retaining control of the schedule themselves.

Outcomes

1. engaged 1500 community members
2. reduced organizers' time from 100 hours to 5 hours
3. Resolved 100+ previously hidden conflicts

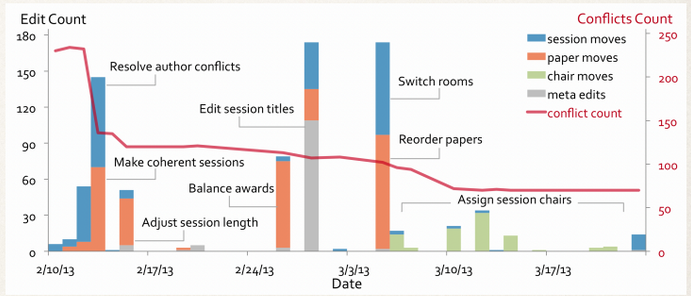

We deployed our system at CHI and CSCW and achieved the following outcomes.

[Outcomes]

But my favorite part is this red line on the right…which shows the number of conflicts going down over time, but what's most interesting to me is the flat part in the middle, where organizers were improving the schedule beyond what the system knows about, but without adding new conflicts.

In this way, organizers can actively exhibit their "care" for creating a "good" schedule and for shaping a "good" conference, beyond just resolving known conflicts. This is a really important idea, that I will come back to later in my talk.

COORDINATING ON LOCAL PROBLEMS



Yongsung Kim

Dissertation: Yongsung Kim. Designing Flexible Coordination Systems to Advance Individual and Collective Goals in Physical Crowdsourcing. 2021.

Funding: National Science Foundation, CHS: Small: Coordination of Opportunistic Actions to Produce Globally Effective Behaviors for Physical Crowdsourcing.

Microsoft FUSE Labs Research Award, Sharing Human-powered Mobility to Improve Societal Efficacy and Efficiency.

Let’s move on to look at the challenge of coordinating on local problems. This is work led by my former PhD student, Yongsung Kim, who is now at Adobe

Mobilizing crowds to solve local problems





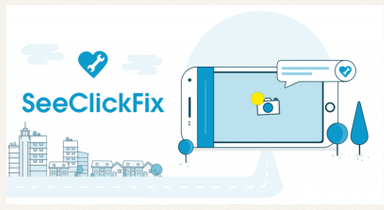



What we are interested in here is mobilizing crowds to solve local problems.

Many of us are used to having crowds help us today, whether that's for rides or food delivery.

Some of us may even serve in citizen crowds, when we use apps to report local problems, or to contribute to citizen science efforts, such as tracking bird migration patterns.

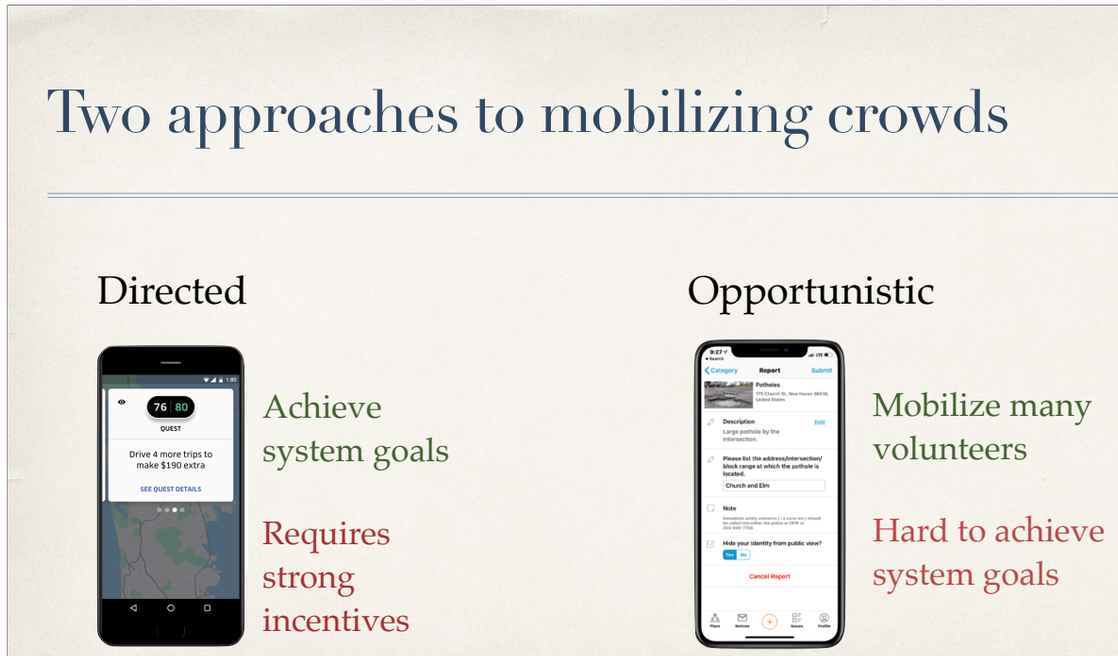

But even after two decades of advances and many successful systems, basically we have just 2 approaches to mobilizing crowds.

The first is a directed approach, where we tell people what to do. This works great for getting things done, but requires strong incentives, which may work in some cases, but not others.

The other approach is to let people opportunistically contribute when they feel like it. This can mobilize many volunteers, but struggle to ensure good system outcomes, whether that's getting good coverage across a region, or getting timely responses.

Instead of settling for either approach, we began to rethink how we can coordinate in an entirely different way….

Flexible coordination [Kim, 2021]

- Ask for contributions from people on-the-go when it is:
  - Convenient and of interest to volunteers
  - Valuable for system goals





which we call, flexible coordination.

The way it works, is that we ask for contributions from people who are on-the-go, and do so only when its convenient and of interest to them, AND valuable to the system.

The key idea here, is that we don’t think of coordination as enforcing rigid ways of working and collaborating that is predetermined. Instead, we are going to LISTEN for opportunities, and ping people when good opportunities arise.

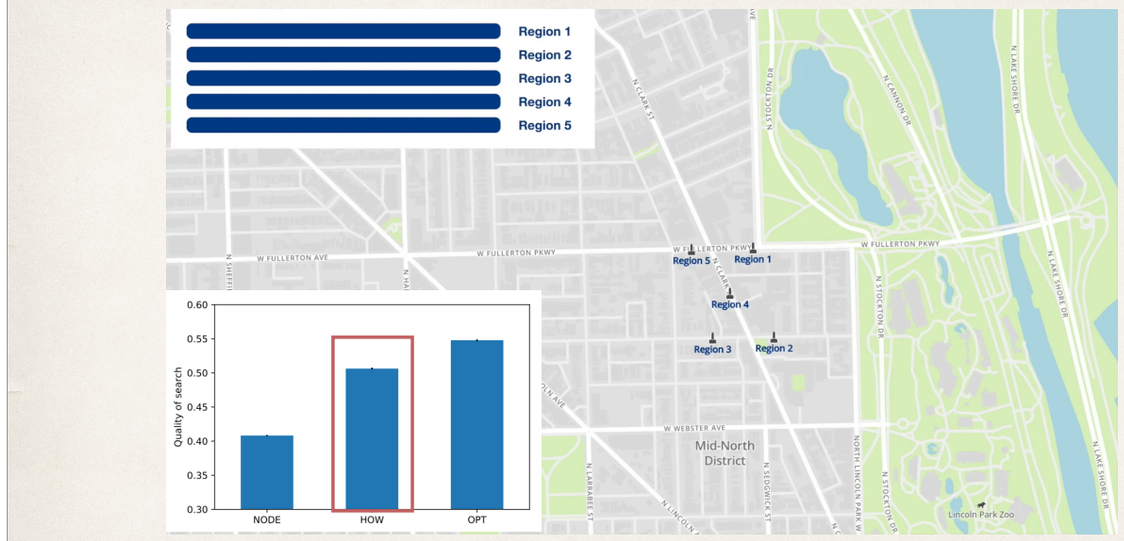

Hit-or-Wait, Yongsung Kim, Darren Gergle, Z., CHI 2018.

Let me show you how this works in practice. What I am showing you here is a lost and found task, where we are trying to search for a lost item across 5 regions, but where everyone who helps may only spend 5-10 seconds looking as they pass by a region.

As people walks past different regions, the system uses decision theory to decide whether to ‘ping’ someone now, or to wait for regions they may come across later, where their help may be more valuable.

Computationally, what’s cool is that this approach achieves outcomes that are almost the same as if we know where people are going ahead of time, which we of course don’t. Practically, what’s really awesome is that by doing this, you get entire new…

new forms of local services and communities



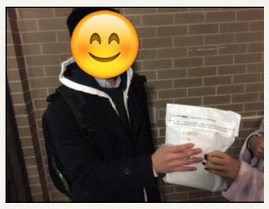



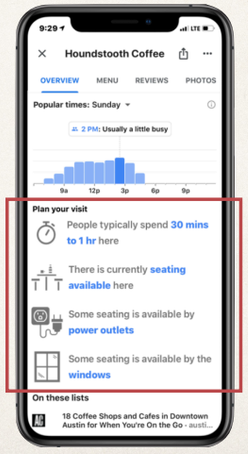

[Kapil Garg et al., CSCW 2019]

forms of local services and communities that are volunteer-based, but that are surprisingly effective, and that are quite different than what we are used to.

What you see on the left are images of people helping one another pickup and deliver packages, and finding some really hard to find items. What you see on the right is a service like Google Places, but where you can get additional, community-provided information, like whether there is seating by the window or with outlets at a coffee shop.

What's exciting isn't just that these services work— it's that they empower people to express care and to help others in small and convenient ways, but where their contributions end up being genuinely valuable to their communities.

CONNECTING FRIENDS AT DISTANCE


Ryan Louie


Jennie Werner


Allison Sun

Dissertation: Ryan Louie. Human-AI Interface Layers: Enhancing Communication of Intent for AI-Assisted Creative Pursuits and Social Experiences, 2023.

Funding: Google Faculty Research Award, Computational Tools for Expressing Conceptually Rich Situations to Machines.

Let’s shift gears from solving local problems, to thinking about connecting with friends and loved ones who are at distance.

This is work led by my former PhD student Ryan Louie, who just started his faculty job at Vanderbilt after a postdoc at Stanford, and by Jennie and Allison, who were undergrads in DTR.



So social media has been with us for a while now, but even today, so much of social media is scrolling through people’s past lives.

So you are scrolling, they are scrolling, we are scrolling. But we are just scrolling past each other, not really engaging much with one another.

Somehow, we got locked into this paradigm, and fundamentally, have not reconsidered how we connect with one another at distance for quite some time

Opportunistic Collective Experiences [Louie et al., 2020]

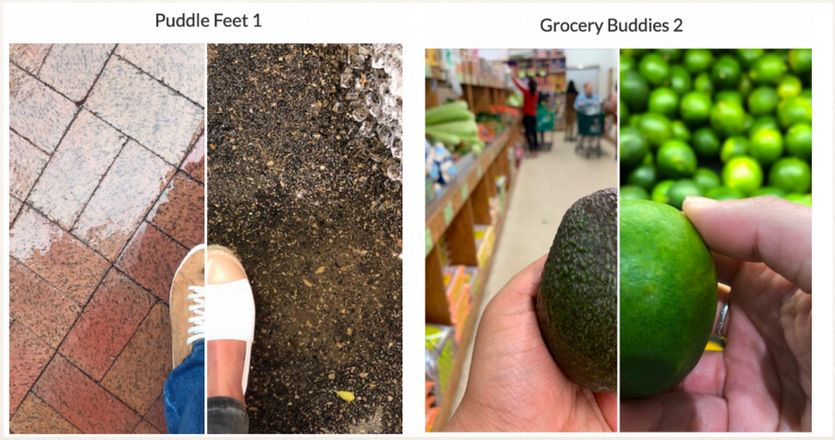

Our research lead us to a totally different approach, which we call, Opportunistic Collective Experiences, or OCEs

OCEs are computer programs that find opportunities for friends and loved ones to engage in shared activities and to connect at distance.

Interactions are grounded in people's physical environments, and only presented when the affordances in people's current places allow them to come together in the desired ways.

Rethinking connecting at distance [Louie et al., 2020, 2022]

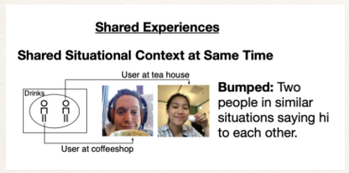

novel interaction model

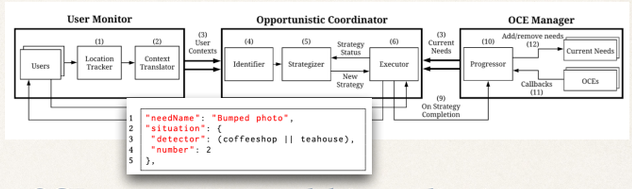

OCE programming model & coordination engine

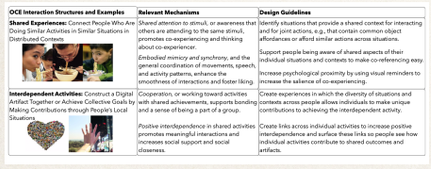

theoretically-informed design principles

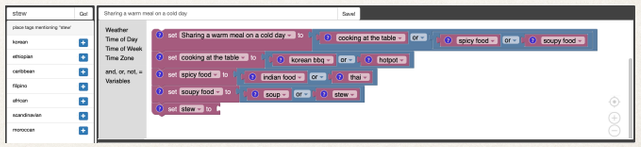

tools for expressing human situations to machines

To do this, we completely rethought what connected at distance looks like, both interactionally and technology wise.

What I am showing on the left is a novel interaction model for connection at distance, for instance when people are in similar situations, or when they can contribute to shared goals.

We read a lot of the literature on social connections, and thought about how what facilitates connecting in-person can be translated to a remote setting, when people for instance, can’t easily attend to the same stimuli.

To make this possible, we built a programming model and coordination engine that recognizes when the right conditions for connection arise, and also tools for expression human situations and experiences to machines, such as where to share a warm meal on a cold day.

Findings: More connecting, fewer barriers, and new experiences

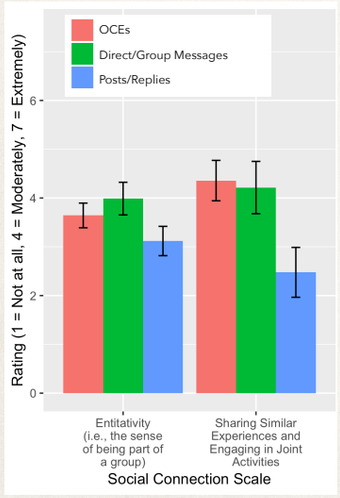

…you're more involved in the other person's life.

Just knowing that what I was doing was explicitly something that [my friend] was doing as well... Not only was it a reminder of them, it's a reminder that I'm around too.

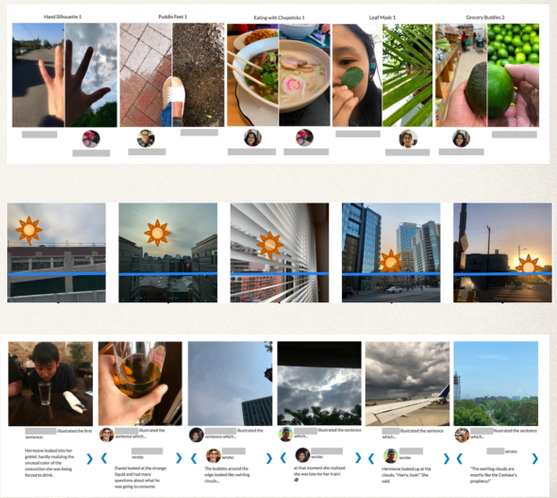

When we do this, we find that we can create a social technology that …

On the left, we see that OCEs are just as connecting as direct messaging, and significantly more connecting than posts and replies on social media. But unlike direct messaging, people find OCEs much easier to engage with throughout their day, whereas with direct messaging, the longer you wait, the harder it is to want to reach out to someone.

Participants felt that with OCEs "…"

In our deployments, people took half-half photos, watched the entire sunset at the same time across places… and told and illustrated stories together.

What's really exciting about all this, is that it demonstrates that WE ARE NOT DOOMED TO SCROLLING PAST each other. If we are willing to rethink how connections may look like, we can have a very different relational experience, even when we are far apart.

LEARNING TO BUILD SOFTWARE


Josh Hibschman


Sarah Lim


Gobi Dasu

Dissertations: Josh Hibschman. Readily Available Learning Experiences in Production Code, 2018.

Gobi Dasu. Process Management for Learning from Professional Source Code Cultivating Diverse Experts in the Age of AI, 2025.

Funding: National Science Foundation, Cyberlearning EXP: Readily Available Learning Experiences: Turning the Entire Web into Progressive Examples to Bridge Conceptual Knowledge Gaps for Novice Web Developers.

Moving on, let’s talk about learning to build software. For about a decade now, my students Josh Hibschman, Sarah Lim, and Gobi Dasu reimagined how novice developers can learn the conceptual knowledge and ways of thinking that professionals have—and what that transition could look like.

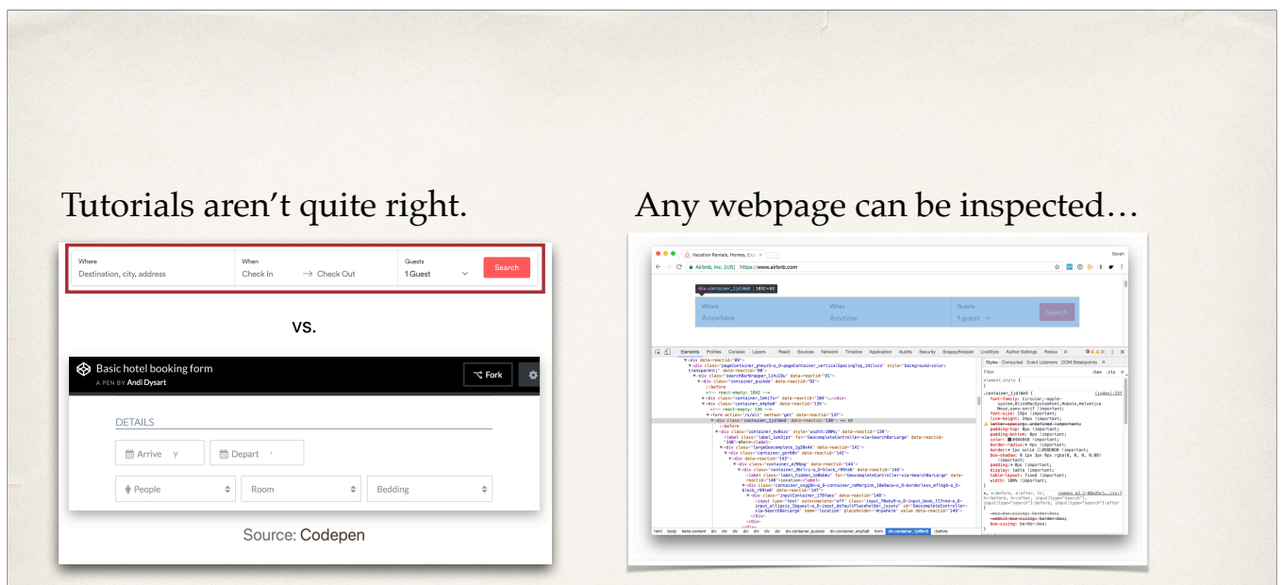

As we looked into this problem, what we saw is that our learning resources just aren't made for this.

On one hand, we have tutorials, but they aren't quite right. On the top left is AirBnB's search bar, and below it is a tutorial for making something similar. But the AirBnB one is just nicer, and knowing how to build the tutorial one on the bottom, doesn't mean you can build the nice one on top.

On the right, I am showing that while you can inspect AirBnB's source code, it is completely overwhelming for a novice, and it's not clear at all, how they build it.

I know there has been better developer tools over the years, but these tools are still largely designed for experts, not novices.

Readily Available Learning Experiences

[Hibschman, 2018]

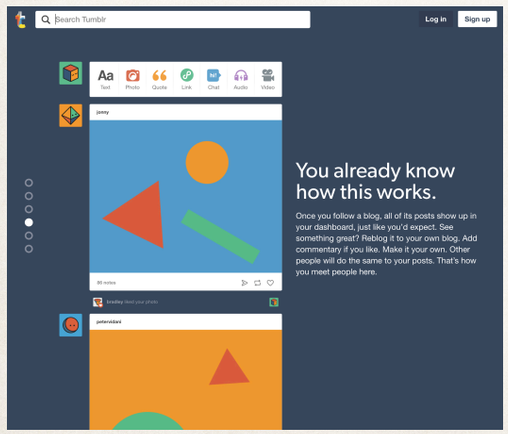

Transform *all* professional websites into a learning resource.

Support novices making sense of complex code.

So we took an entirely different approach, which we call, Readily Available Learning Experiences

The key idea, is that we are going to transform all professional websites into a learning resource for novices. So instead of relegating novices to simple tutorials, we are going to support them to make sense of complex code directly.

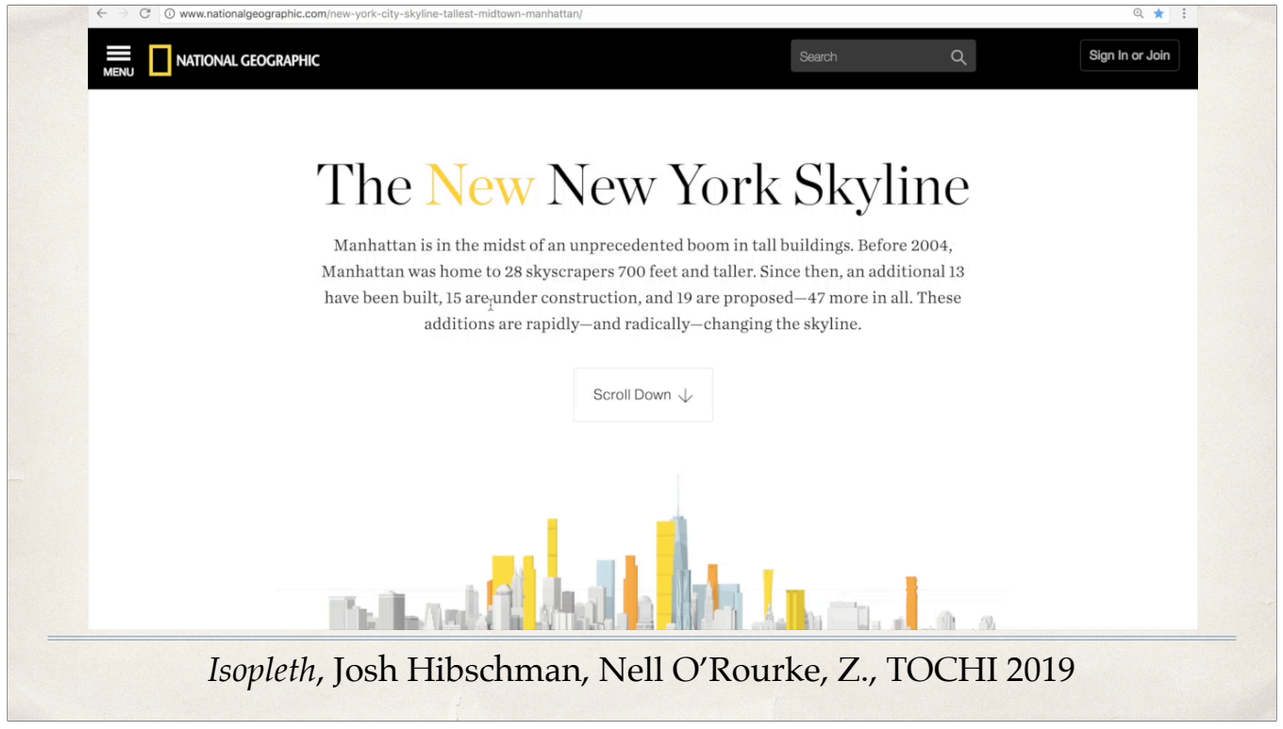

Isopleth, Josh Hibschman, Nell O'Rourke, Z., TOCHI 2019

Let me show you what this looks like, with a tool called Isopleth.

What you are seeing here is an interactive website of the New York City Skyline from National Geographic.

Instead of looking at its source code through an inspector, Isopleth provides visual representations that help a learner see and understand hidden event-driven relationships and function calls.

It gives learners a way to build on their intuitive understanding of visual features, and to have intuitive entry points into the code, while still engaging with the modes of thinking of the discipline, for instance by focusing in on Ajax calls.

Learner-centric developer tools

Isopleth [Hibschman et al., TOCHI 2019]

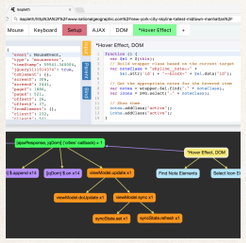

Novices learned complex interactive Javascript features.

31% improvement in conceptual models of code architecture

Ply [Lim et al., UIST 2018]

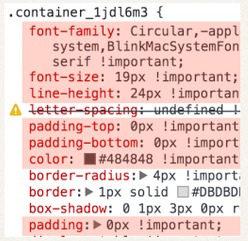

Novices learned unfamiliar CSS concepts.

Replicated complex CSS features 50% faster than CDT.

now in all major browsers

And when we do this and create what we call learner-centric develop tools, we find that they are really effective.

We found significant improvements in both program comprehension and replication than existing approaches for learning Javascript and CSS from professional examples.

And Ply is now in the all major browsers, so even novices can figure out what's going on in the CSS of any website.

Process Management for Learning from Professional Source Code [Dasu, O'Rourke, Z., TOCE 2026]

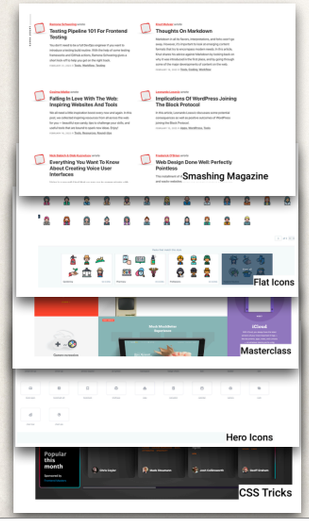

While we could've stopped here, we went further and found that novices needed not only better tools, but better strategies for sensemaking, and for learning.

For example, when given a stack of professional examples of web layouts, we found that novices rarely looked beyond surface differences, and struggled to build and transfer their understanding across examples, just as learning science theory might predict.

Process Management for Learning from Professional Source Code [Dasu, O'Rourke, Z., TOCE 2026]

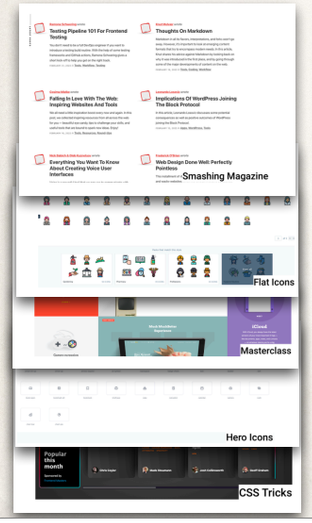

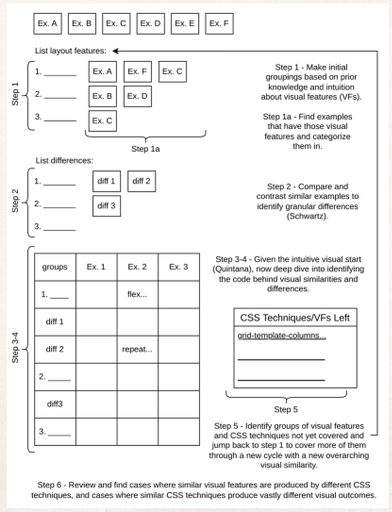

This led us to create better scaffolds and software for managing the sense making process, which completely changed how novices approach professional examples in the first place.

Process Management for Learning from Professional Source Code [Dasu, O'Rourke, Z., TOCE 2026]



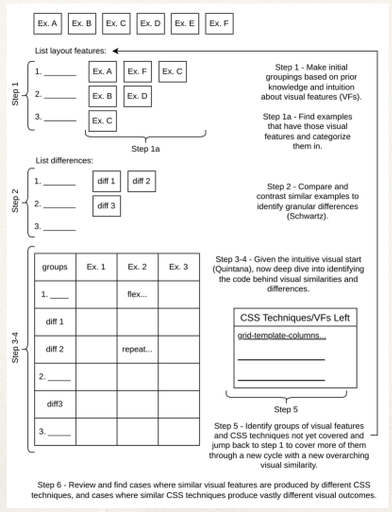

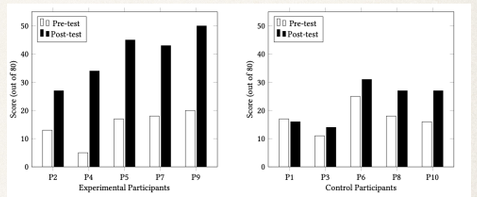

Fig. 28. Comparison of pre-test and post-test scores for experimental (KM) users.

Fig. 29. Comparison of pre-test and post-test scores for control users.

Significant gains in conceptual understanding of professional code across Javascript, CSS, and Python.

Across 3 systems, we found that managing the sense making process to target conceptual learning outcomes led to significant learning gains in modeling Javascript data flows, mastering CSS techniques, and understanding design patterns in Python.

Instead of building a better tutorial or a better code inspector, we focus on changing how novices engage with expert work. With our approach, learning happens through participation in professional practice, and not apart from it.

INNOVATING THROUGH RESEARCH



Leesha Maliakal

Dissertation: Leesha Maliakal. Agile Research Studios: Learning Ecosystems to Scale Effective Research Training, 2023.

Funding: National Science Foundation, Cyberlearning EXP: Agile Research Studios: Scaling Cognitive Apprenticeship to Advance Undergraduate and Graduate Research Training in STEM.

Finally, let's talk about innovating through research, which is work led by my PhD student Leesha Maliakal, who is now faculty at NEIU.


Design, Technology, and Research (DTR)

Going back to the question of training researchers and innovators, I want us to think about why this is so hard, and in particular, what's needed for students to drive research and innovation themselves.

Students need regulation skills

- **Regulation skills**: cognitive, metacognitive, motivational, and emotional skills for reaching a goal *[Jarvela & Hadwin. 2013]*
- Learning to innovate requires regulation skills such as **planning** and **seeking help** to overcome challenges.
- Students lacking these skills are confined to rote tasks, or can struggle to make progress.

To become researchers and innovators, students need to develop regulation skills — or cognitive, metacognitive, motivational, and emotional skills needed to plan and seek help while facing many unknowns. Without them, students can often get stuck doing rote tasks, or they can struggle to make progress, and need to wait for their mentors to rescue them, and get them back on track.

The sad thing is that, as faculty mentors, we don’t always do a good job of teaching these skills. As a result, many students end up banging their head against the wall, and my contention is that there’s got to be a better way.

Agile Research Studio (ARS)

[Z., Easterday, Gerber, Rees Lewis, Maliakal, 2017]

- Model for research training in a learning community
- All students, regardless of seniority, conduct independent research and receive authentic research practice.

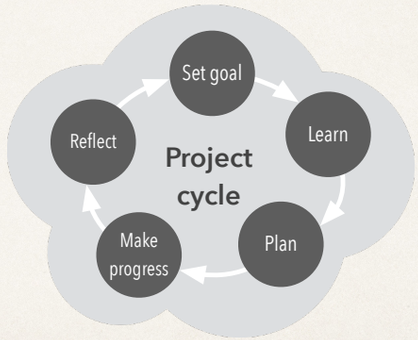

To address this challenge, we created Agile Research Studios, a socio-technical model for creating a learning community for research training.

All students, regardless of how senior they are, receive authentic research practice, by which I mean that they drive the self-directed learning cycle, where they set goals, learn, plan, make progress, and reflect before setting goals again.

ARS is a sociotechnical model for developing regulation skills

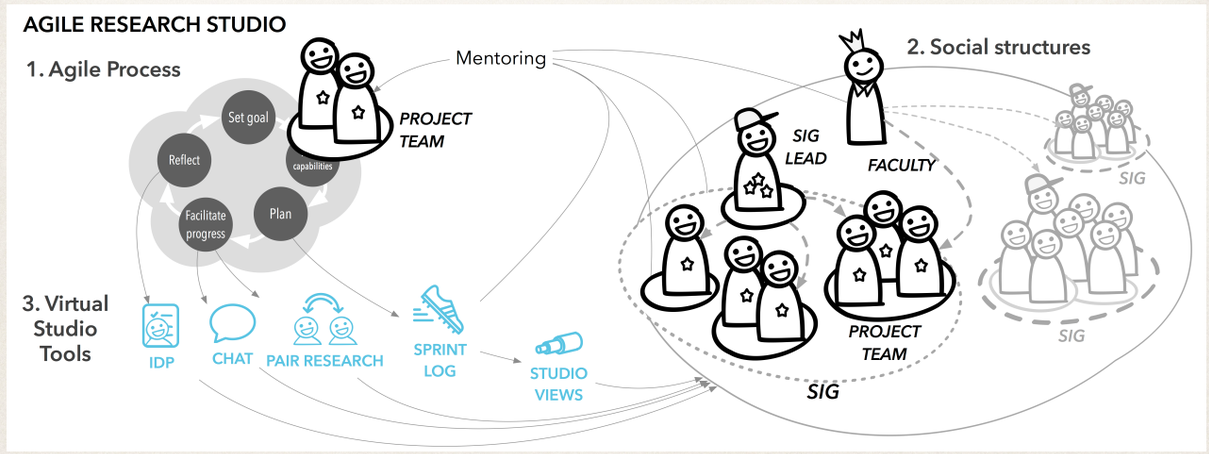

At its core, ARS is a set of processes, social structures, and tools that work together to help students develop regulation skills. It fundamentally reconfigures the roles and relationships of student researchers and their mentors, and creates technologies that facilitate this reconfiguration.

ARS: planning [Maliakal, 2023]

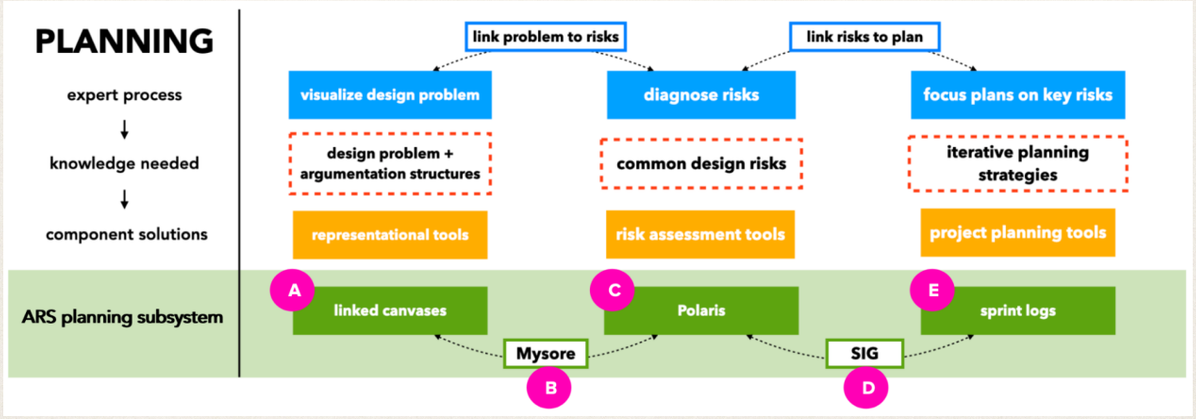

Let me first describe how the model supports students learning to plan. For us, planning isn't simply: figure out what your tasks are for the week.

Instead, it's a process of learning to represent and visualize your existing understanding, to diagnose risks, and to focus plans on those key risks.

We created entire subsystems, or sociotechnical configurations, for learning each of these sub-skills, and for students to connect these skills, so that they can link their problem understanding to their risks, and their risks to their plans.

Together, this provides a way for students to practice driving research themselves.

ARS: help & collaboration

Process:
Distributed help



Social structure:
Studio meeting



Studio tool:
Pair research

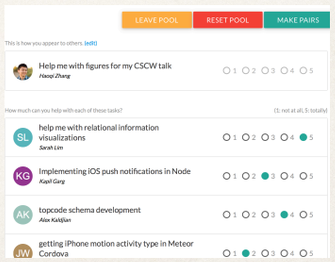

[Miller, Z., Gilbert, Gerber, CSCW '14]

On the help side, instead of relying solely on faculty mentors, ARS distributes help across the community. Every week during studio meetings students sign up for pair research, which pairs students up to help one another on whatever they need help with.

By running a matching algorithm, not only does this allow us to maximally use the expertise within our community, it…

Help & Help-seeking

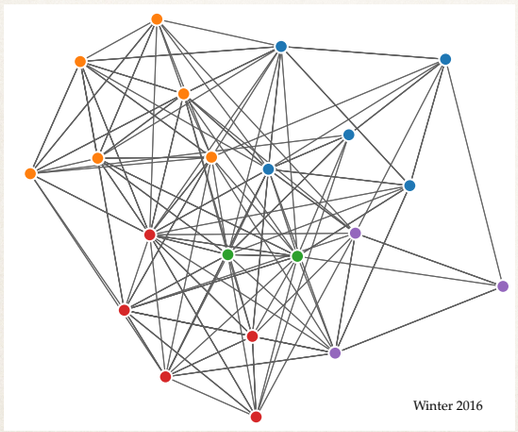

"I can ask for help and that everyone asks for help and it doesn't make them stupid to need help."

Leads to our students getting a ton of help from different students, both within their research area, and across research areas as well.

But what I get really excited about, is seeing students shift their help-seeking dispositions. For instance, one student said [QUOTE]

This is really promising to see, because even our best students can struggle with asking for help.

Outcomes (12 yrs)

- 193 DTR students (169 UG, 11 MA, 15 PhDs) led 70+ research projects.
- 50% women!
- 78 undergraduate research grants
- 30+ papers & extended abstracts; 7 winners at major ACM Student Research Competitions
- Many DTR undergraduates placed at Apple, Google, Microsoft, Meta, Amazon, and OpenAI; others have founded their own companies.



In 12 years, in our home studio here at Northwestern, I am proud to share that I have directly mentored almost 200 DTR students, half of whom are women!

These students received numerous undergraduate research grants, published many papers, and won multiple student research competitions.

A great many of them place at the very top tech companies, and others have founded their own companies as well.

Outcomes (12 yrs)

- 193 DTR students (169 UG, 11 MA, 15 PhDs) led 70+ research projects.
- 50% women!
- 78 undergraduate research grants
- 30+ papers & extended abstracts; 7 winners at major ACM Student Research Competitions
- Many DTR undergraduates placed at Apple, Google, Microsoft, Meta, Amazon, and OpenAI; others have founded their own companies.

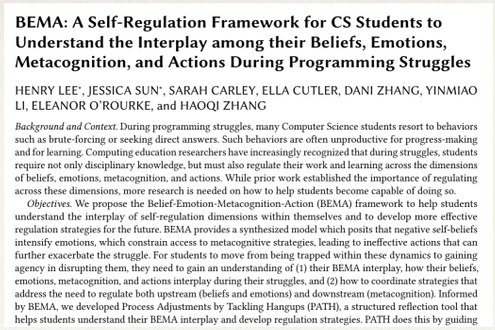

BEMA: A Self-Regulation Framework for CS Students to Understand the Interplay among their Beliefs, Emotions, Metacognition, and Actions During Programming Struggles

HENRY LEE*, JESSICA SUN*, SARAH CARLEY, ELLA CUTLER, DANI ZHANG, YINMIAO LI, ELEANOR O'ROURKE, and HAOQI ZHANG

Background and Context. During programming struggles, many Computer Science students resort to behaviors such as brute-forcing or seeking direct answers. Such behaviors are often unproductive for progress-making and for learning. Computing education researchers have increasingly recognized that during struggles, students require not only disciplinary knowledge, but must also regulate their work and learning across the dimensions of beliefs, emotions, metacognition, and actions. While prior work established the importance of regulating across these dimensions, more research is needed on how to help students become capable of doing so.

Objectives. We propose the Belief-Emotion-Metacognition-Action (BEMA) framework to help students understand the interplay of self-regulation dimensions within themselves and to develop more effective regulation strategies for the future. BEMA provides a synthesized model which posits that negative self-beliefs intensify emotions, which constrain access to metacognitive strategies, leading to ineffective actions that can further exacerbate the struggle. For students to move from being trapped within these dynamics to gaining agency in disrupting them, they need to gain an understanding of (1) their BEMA interplay, how their beliefs, emotions, metacognition, and actions interplay during their struggles, and (2) how to coordinate strategies that address the need to regulate both upstream (beliefs and emotions) and downstream (metacognition). Informed by BEMA, we developed Process Adjustments by Tackling Hangups (PATH), a structured reflection tool that helps students understand their BEMA interplay and develop regulation strategies. PATH does this by guiding




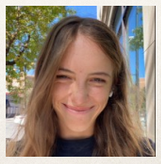


Just accepted at ACM Transactions of Computing Education (TOCE)!

As just one example, two weeks ago an undergraduate team of DTR students had their journal paper accepted at ACM TOCE, which is the top computing education journal!

This is really exciting, and we have been sharing our success…..

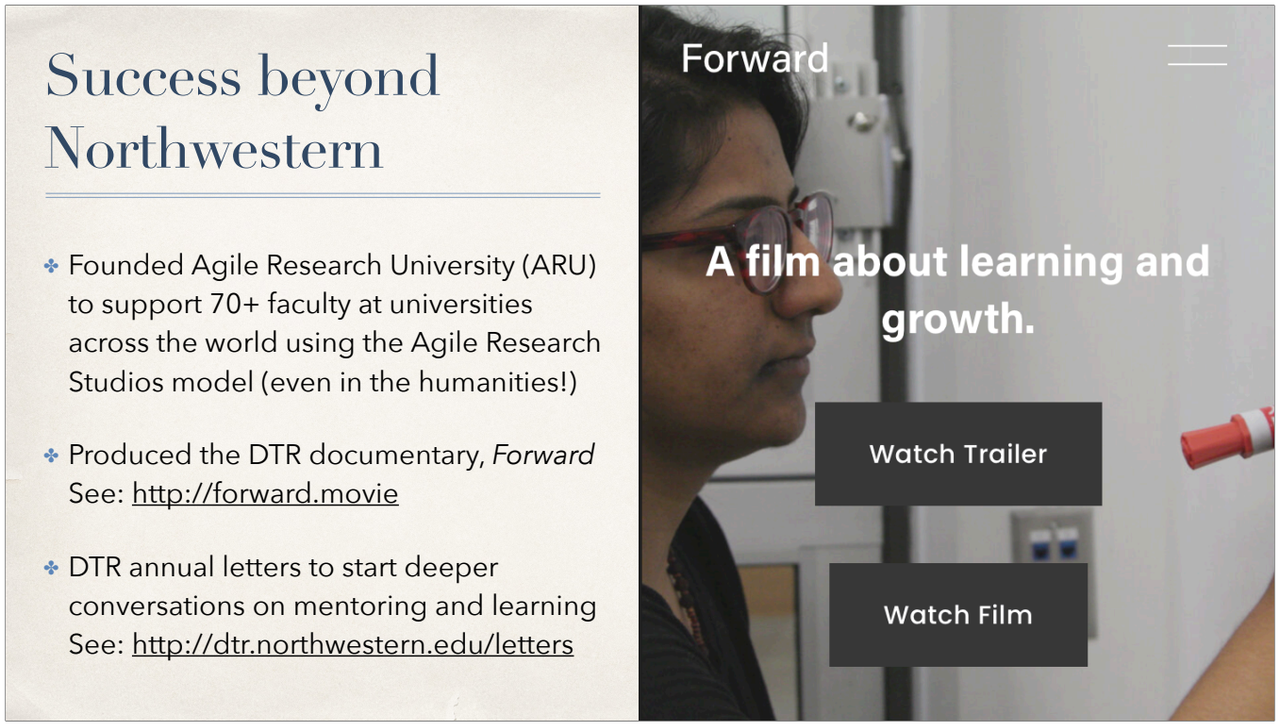

Beyond Northwestern to support 70+ faculty around the world who use our model to support their students.

In 2022 we released the DTR documentary, Forward, and I began writing the DTR annual letters, all to start a conversation about learning and growing, and the challenges we all face, as students and educators.

If you are interested, you can follow the links here, and, I brought hard copies of this year's annual letter with me today, and you are welcome to grab a copy on your way out.



Alright, so so far, I have shown you 5 examples of problems and entirely new ways of approaching them.

Before I summarize what we did across the five, any questions about any particular system?

PAUSE FOR QUESTIONS

- There were fundamental limitations with existing approaches.
- These limitations could not be overcome with technology alone.
- We took a different approach.
- We created entirely new sociotechnical configurations to support it.

So, in each of these activities we looked at, there were fundamental limitations with existing approaches, that could not be overcome with technology alone.

What we did is to take a new approach.

In doing so, created an entirely new sociotechnical configurations to support it, which we call,

Computational Ecosystems

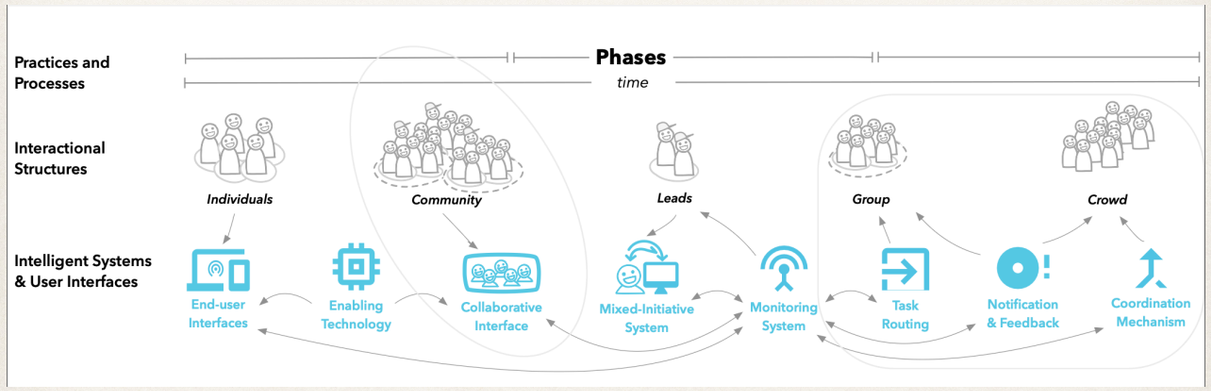

Computational ecosystems.

Computational ecosystems are systems, but that are designed not piecemeal, but as integrative solutions — where practices, interactional structures, and technologies are conceived of together as one coherent whole.

Typically…

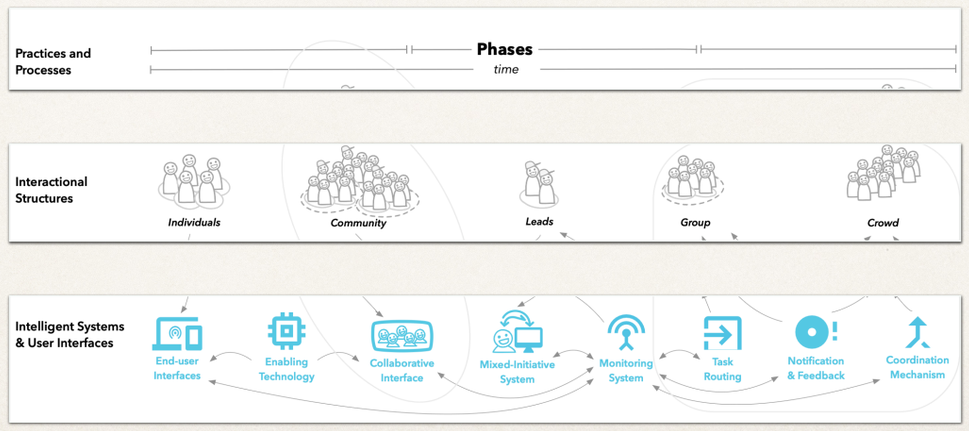

Typically, we design sociotechnical systems layer by layer:
- we define a process,
- then shape the social structures around it,
- and finally build tools to support both.

And this is a thoughtful approach — especially if existing practices can’t easily change.

Instead…

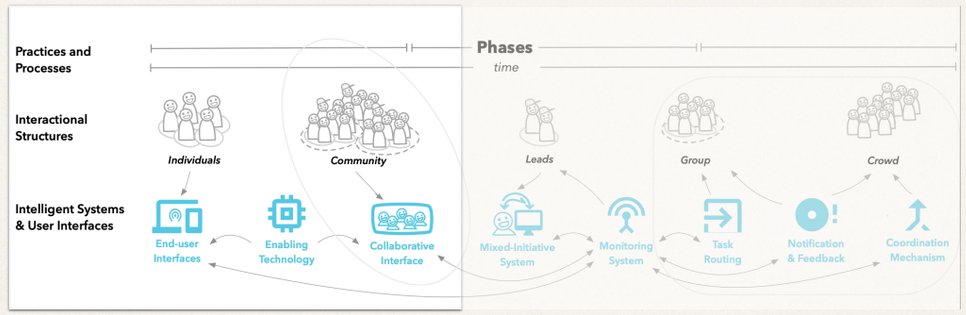

But if we *can* change the practice, why should we assume that it has to stay fixed? You can get much better solutions, as we have demonstrated across the five examples, if you are willing to consider changing all three layers at the same time.

In other words, computational ecosystems free us to explore more creative solutions, that are more likely to succeed, when our existing ways of doing things are fundamentally limited or flawed.

A call for systems: having great components is not enough.



Atul Gawande

"…We've been obsessed in medicine with components. We want the best drugs, the best technologies, the best specialists, but we don't think too much about how it all comes together. It's a terrible design strategy actually."

TED 2012

Our FOCUS on a system's approach to design echoes the calls from Atul Gawande from the Harvard Medical School, who argues that in essence, having great components is not enough.

He says…

A call for systems: having great components is not enough.



Atul Gawande

"Making systems work - whether in health care, education, climate change, and making a pathway out of poverty - is the great task of our generation as a whole."

[Quote]

I think in computing, we are still not doing our part in thinking systematically enough.

But when we do, I think we can not only have the impact that we have had, but can start to address some of the fundamental human problems and values that we have largely neglected, if not created.



But before we get too excited, I do want to note that doing research this way is really brutal.

- My students and I often start by trying many existing approaches, only to get no where.
- Only after YEARS of intense study in and outside of computer science do we arrive at a new approach.
- It then takes years to develop increasingly sophisticated systems, and to gather evidence across multiple real-world deployments
- And after all that work, and perhaps because of it, other human values can become legible that challenge us to rethink and reconfigure our ecosystem yet again.

So while this work is important, it takes MANY YEARS of effort to design, build, and study even a single computational ecosystem. So it's understandable why not everyone wants to work this way, or to go after the problems we've gone after.

- There were fundamental limitations with existing approaches.
- These limitations could not be overcome with technology alone.
- We took a different approach.
- We created entirely new sociotechnical configurations to support it.

Before I move onto the second half of my talk, I want to pause again here.

Any questions?
BRIEF PAUSE HERE

RESTART HERE:
So far, we've looked at examples of computational ecosystems that open up entirely new ways of doing things.

What are computational ecosystems good for?

[So far, we've looked at examples of computational ecosystems that open up entirely new ways of doing things.]

In the next part of this talk, I want to ask a deeper question: what are computational ecosystems actually good for?

Are they just better ways to solve problems, or do they contribute something else — something intrinsically valuable that individual technologies cannot?

DEVELOPING A SELF/PRACTICE



Kapil Garg

Dissertation: Situated Practice Systems: Developing Worker's Capabilities for Complex Work in Networked Workplaces, 2024.

Funding: National Science Foundation, RITEL: Collaborative Research: Situated Practice Systems: Supporting Coaches and Students to Develop Regulation Skills for Design, Research and STEM Innovation.

I actually think they do, and I want to give you two examples.

The first is that, computational ecosystems can create spaces for people to grow, by supporting their developing a self and a practice.

This is work led by my PhD student Kapil Garg, who is now at UC Irvine.



Design, Technology, and Research (DTR)

Let’s go back to DTR for a minute. While we do train students to produce research, that’s actually not our main focus in DTR.

I apologize for the slight bait and switch. But in reality, our main focus is to use research as a vehicle for students to learn about themselves, and to grow as people.

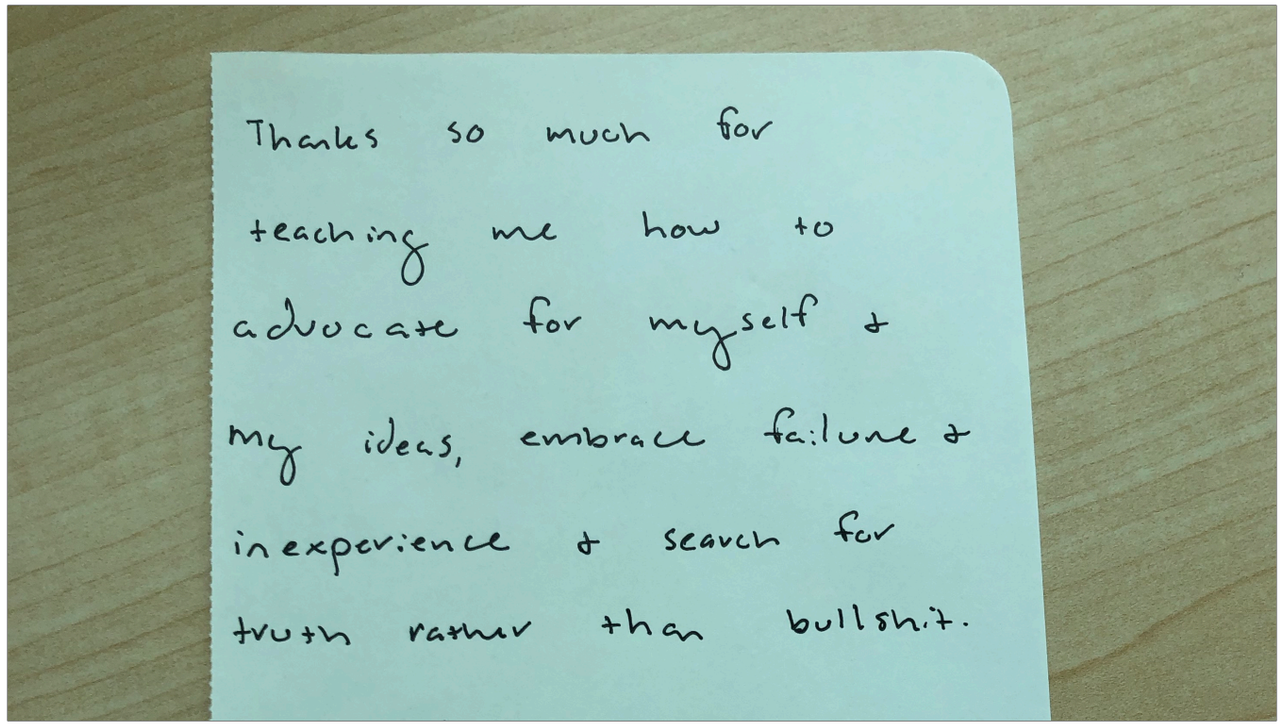

I want to read you a note from a DTR alum, who said: [quote]

Notice that what's most important here is not just producing research, but the student seeing herself differently

And the question I have is: how does a space for research training, provide this kind of learning and growth? And what does this kind of personal growth have to do with doing research anyway?

Developing regulation skills for building a self-directed research and innovation practice

Cognitive skills

- representing problem and solution spaces *[Carlson et al., 2020]*
- assessing risks *[Carlson et al., 2018]*
- critical thinking and argumentation *[Maliakal, 2023]*
- core design, research, and STEM methods

Metacognitive skills and dispositions

- planning: forming feasible plans and planning effective iterations *[Z. et al., 2017; Rees Lewis et al., 2018]*
- help-seeking: leveraging resources; seeking help; communication skills *[Garg et al., 2022]*
- reflection: awareness of one's own skills, abilities, and metacognitive blockers *[Pribble et al., 2022]*

Emotional regulation and disposition toward self and learning

- emotional regulation: understanding one's fears and anxieties
- disposition: dealing with failure, embracing challenges, embracing self-direction

To begin answering these questions, I want to start by highlighting how many regulation skills are involved in self-directing research, and in learning to innovate.

The truth is, it's not easy for students to be good at any of these skills, let alone all of them.

(You can see that planning and help-seeking are just two of these skills).

All kinds of scaffolds can help, and our lab has created many of them, but good coaching really matters, because ultimately, it's about the person.

Challenge: understanding and addressing the gaps in students' practice and regulation is really complicated.

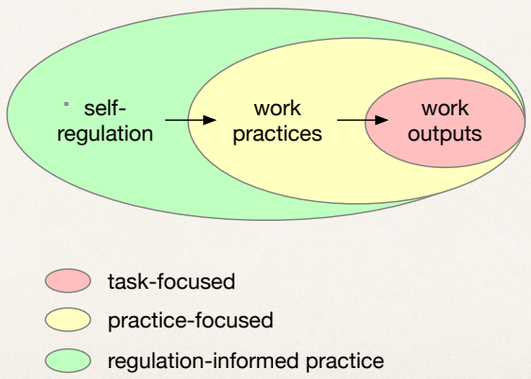

[Wang, Lee, Ly, Z., in preparation]

A core challenge is that understanding and addressing the gaps in students' practice and regulation is really complicated.

We understand that students' work outputs aren't everything — and that if we want them to grow, students need to learn better work practices, and even more so, to have better ways of regulating their work and learning.

Challenge: understanding and addressing the gaps in students' practice and regulation is really complicated.

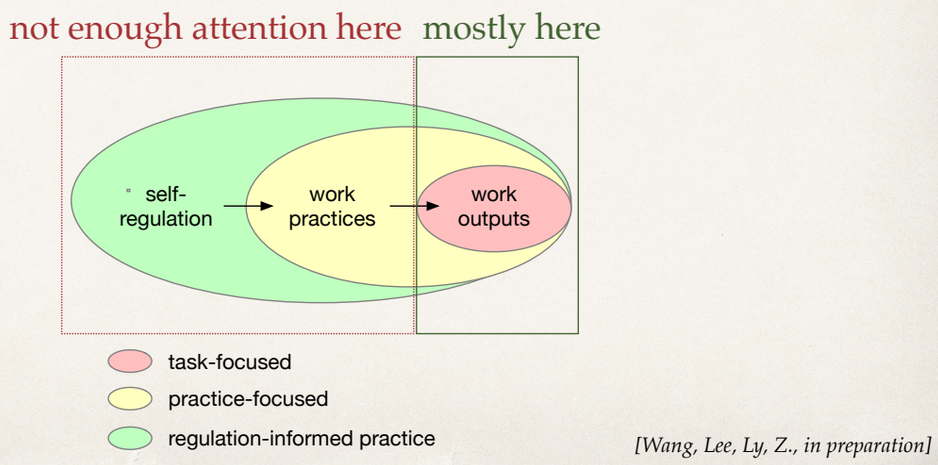

[Wang, Lee, Ly, Z., in preparation]

However, this is often not what we focus on in project-based learning classes. A lot of effort is spent troubleshooting students' work output, and there is often little focus on understanding a student's practice and regulation, even when those are actually the desired learning outcomes.

Interactive CAP Notes

[Garg, Gergle, Z., CSCW 2026]

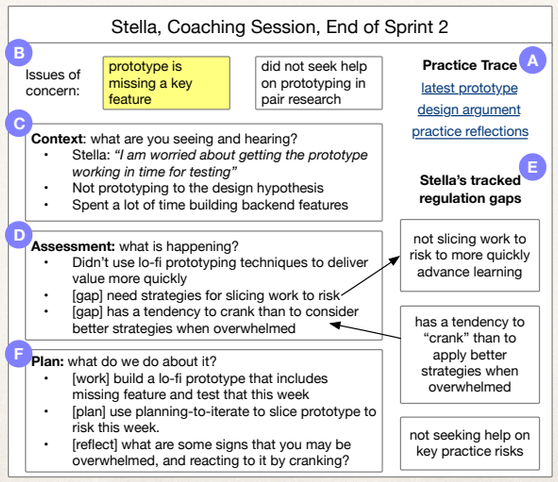

To address this, we developed a mentoring practice and tool called CAP notes, where coaching focuses on addressing issues in students practice and regulation than just their work output.

During a coaching session, a coach still notes practical issues of concern, such as that Stella's prototype is missing a key feature. But instead of just telling her to fix it, we try to understand why this is happening in the first place.

Interactive CAP Notes

[Garg, Gergle, Z., CSCW 2026]

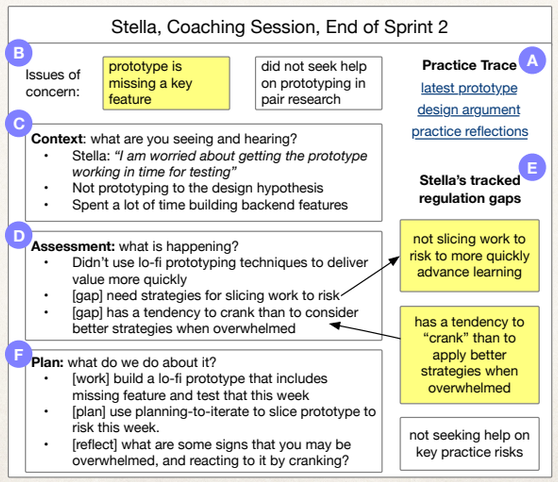

And what we may find is that Stella lacks strategies for slicing work to risk, in that she prioritizes delivering software each week over new understanding. We find also that she tends to overwork when overwhelmed, which often leads to ineffective work, despite all the hours she puts in.

Practice Objects for tracking practice and regulation [Garg, Gergle, Z., CSCW 2026]

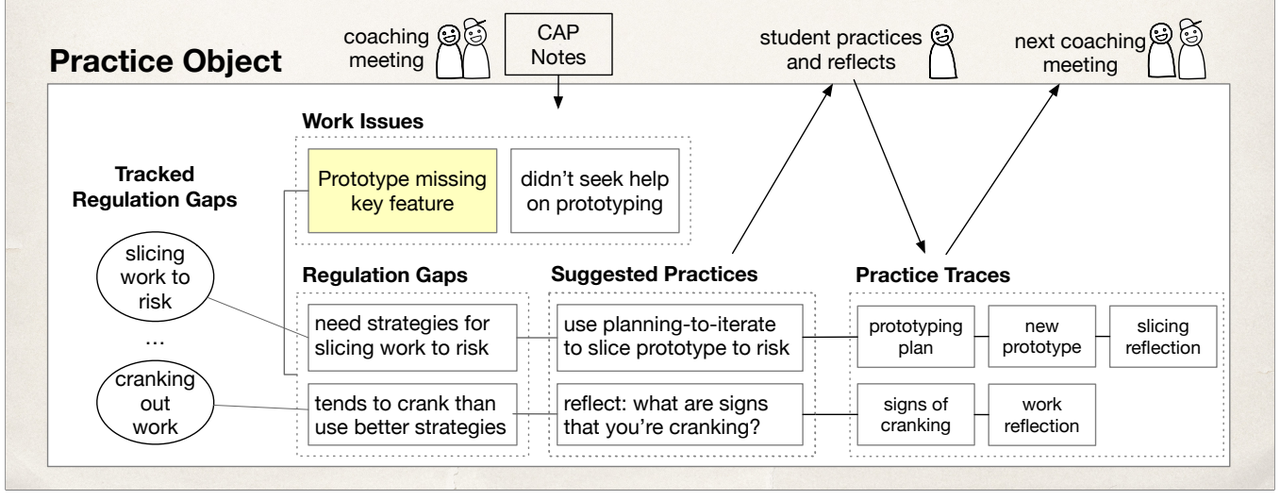

The key idea is that, instead of seeing student's work issues each week as isolated problems, we recognize how ineffective ways of working and regulating are patterns that trigger across situations.

What this allows us to do as coaches is to suggest practices for working directly with their gaps in their regulation and practice, than just helping students get over the surface problems they face each week.

To make this possible, we created what we call practice objects, which are computational representations that captures how a student's regulation and practice evolves over time. These objects then allow us to create new genres of tools and agents that support learning and mentoring, at the level of practice, and at the level of regulation.

Creating space for learning about one's self

- Build strong relationships
- Hold students for who they are *and* are becoming
- Recognize that patterns + beliefs are quite sticky
- Value developing regulation skills over producing research
- Problematize how students approach problems and look at themselves

From *the DTR annual letters*:
- A College Education (2025)
- Interruption, Discord, and Drama (2025)
- Holding Space (2026)
- The Research Project (2026)

Beyond computation, what really shifted is that we no longer see self-regulation as just a skill to be built. Instead, we see it as a doorway into seeing students more clearly, and for students to see themselves more clearly.

To do that, we make an incredible effort to..

[SLIDE — don't say too much…]

Challenge #2: seeing and acting on the *good* of engaging in research

- Dancing with not knowing
- Re-examining phenomenon
- Re-examining beliefs and worldviews
- …

From *the DTR annual letters*:
- The Good (2024)
- Beyond Production (2025)
- Sea Otters (2026)

These goods are largely not about production,
but about how we *engage* in research.

In addition to learning regulation skills, we focused on helping students engage more deeply in research.

What I value most about research is that it offers an opportunity to learn to dance with not knowing, to intensely re-examine phenomena, and to see into our deeply held beliefs and worldviews, that may limit our ability to see the world clearly.

These goods connect us to the scientific tradition, and they are largely not about producing more papers, but are valuable ways to engage in research in their own right, and also in being a human.

If you are interested in learning more, please refer to these essays from the DTR annual letters.

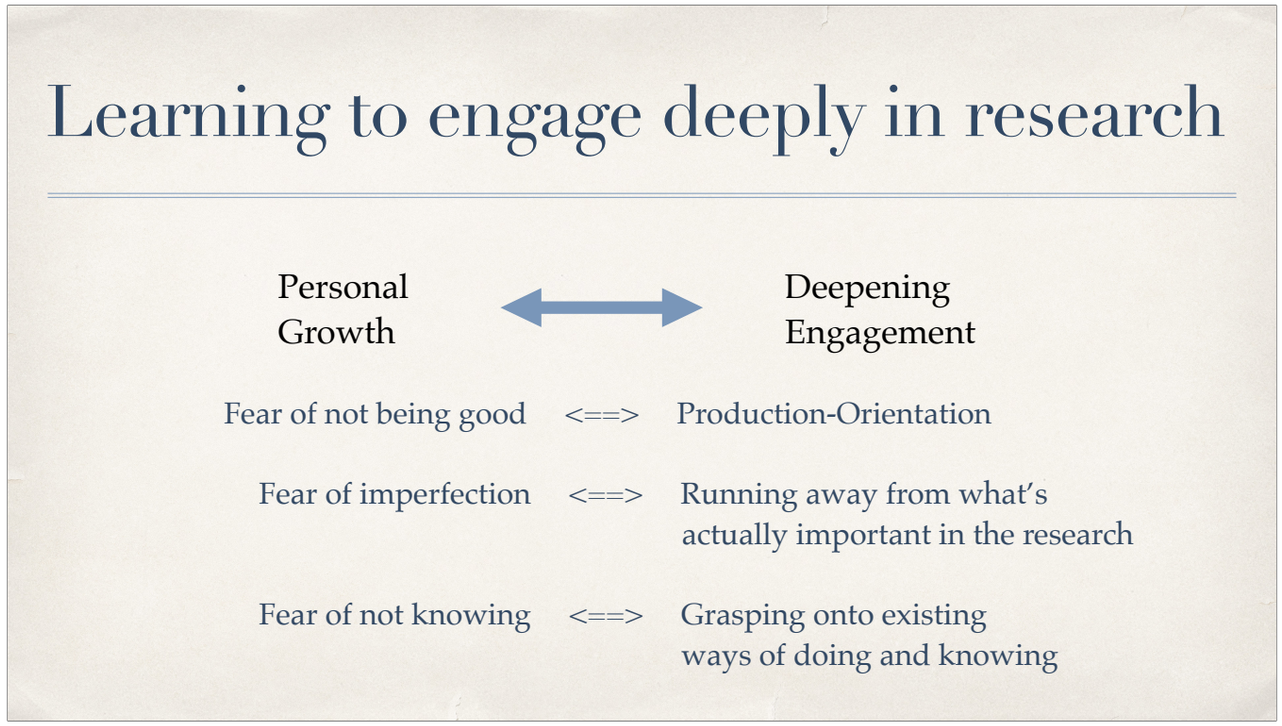

So how does one learn to engage more deeply? This is a complicated question, but a partial answer is that we need to understand what gets in the way, and to work with that.

And here we see that personal growth is intricately tied to our ability to engage deeply, because for example, when we fear we are not good, it's easy to slip into a production-oriented mindset, as a way to prove our own worth.

Our fear of imperfections can lead us to run away from what's actually important in the research, and our fear of not knowing can lead us to grasp onto existing ways of doing things than to explore new ways of investigating that may be more effective.

What we are doing in DTR is to use our learning ecosystem to support students in having this kind of growth, both because it's valuable for their persons, and because it allows them to engage more deeply in the research.

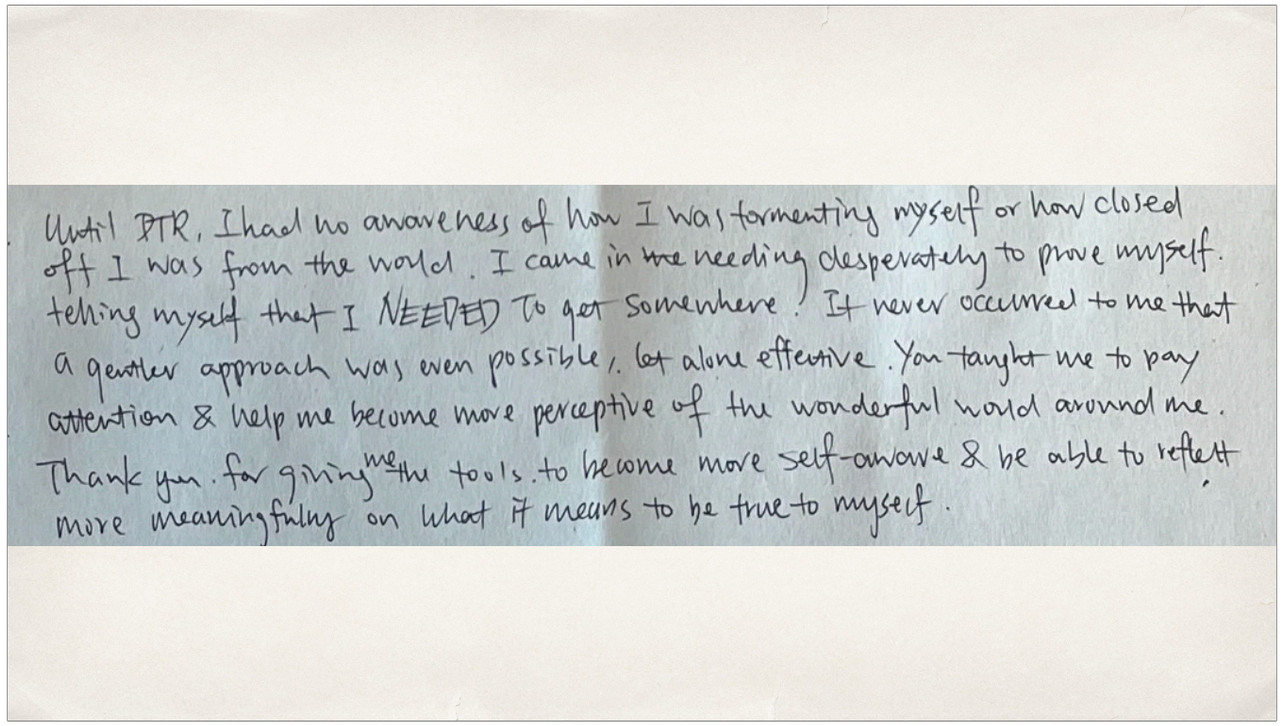
Until DTR, I had no awareness of how I was tormenting myself or how closed off I was from the world. I came in needing desperately to prove myself. telling myself that I NEEDED To get somewhere. It never occurred to me that a gentler approach was even possible, let alone effective. You taught me to pay attention & help me become more perceptive of the wonderful world around me. Thank you for giving me the tools to become more self-aware & be able to reflect more meaningfully on what it means to be true to myself.

I want to share this note with you from a student, who said that until DTR,… [quote]

This note really strikes me, because how can any of us engage deeply in research, if we can't see the beautiful world around us? How can we seek new knowledge with open eyes?

We all deserve to feel okay about ourselves, and to have the space and community behind us that make valuable ways of learning and growing not only possible, but commonplace. In DTR, what we have a computational ecosystem that supports not only our work, but our growing into the persons we are becoming.

A focus on becoming



Kalina Silverman



Meg Grasse



Katherine Lin

And this focus on BECOMING has fundamentally changed my view on what research learning is about, and on what success looks like.

In DTR, success isn't just sending someone to grad school. It's students develop the courage and skills needed to pursue a life that is fully their own, and to devote themselves to problems, practices, and relationships that are meaningful to their persons.

As examples, Kalina started Big Talk, a social movement to foster more meaningful conversations and combat loneliness (her book just came out this year!)
Meg went from leading a team at Apple to studying environmental engineering to improve water supply systems in the UK
Katherine was the head of engineering at a sustainable tech startup, but recently spent a year at the whale and dolphin conservation in Massachusetts.

I am really excited for these students of mine personally, and broadly, for shifting the dialogue on what research experiences can provide for our students, and the meaning it can have in their lives. [ANY QUESTIONS????]

SUPPORTING HUMAN EXPERIENCES



Jackie He



Shirley Zhang

Dissertation: [yes please!]

Funding: Center for Advancing Safety of Machine Intelligence, Human-AI Tools for Expressing Human Situations and Contexts to Machines.

Okay, so my previous example highlighted how an ecosystem can support our inner growth. This example is about how they can help us to engage with the outer world.

This is work led by former DTR students Jackie and Shirley, and its an area we are actively investigating today.

The Internet Is Very Flat, And We're Used To It

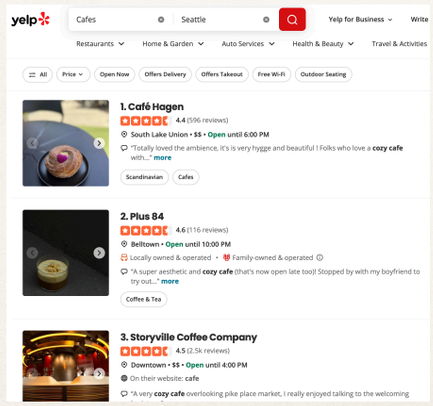

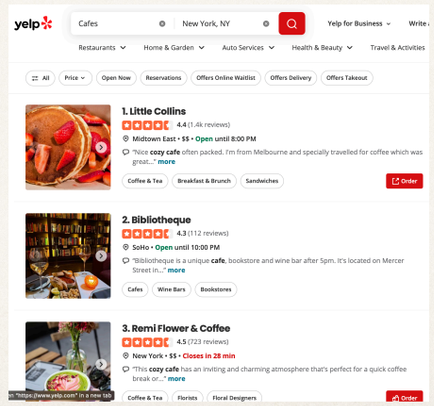

When we look at the world through the internet, it's just very flat..

You can search for Cafes anywhere, …. And you will get a list of the best cafes no matter where you search.

And if you look at the terms are bolded in the reviews, it's cozy cafe, cozy cafe, cozy cafe. Because according to Yelp, what makes a cafe good is just that they are cozy, and this is what it highlights, and this is very flat

The Cost of Ubiquity

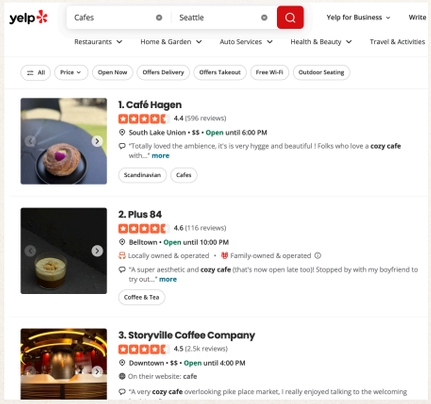

What makes a place special?

What experiences does it afford, and for whom?

How does that differ across user populations, geographies, cultures?

And this is the COST of ubiquity, right? We built systems that work beautifully everywhere — and this is what we prize as Computer Scientists.

But in doing so, we've replaced these questions you see on the right

with:

The Cost of Ubiquity

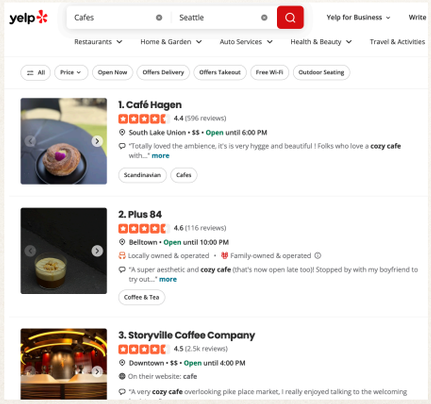

What is the best?

What is the best place to go?

And in many ways, we have accepted this as not only okay, but even, desired!

Our own explorations of the world can often guided and consumed by searching for what’s best, and valuing what’s best over anything else.

Computational Understanding of Human Experiences

"an increasing understanding on the part of system developers that human activities are enmeshed in a variety of [cultural and social] practices and relations that make them meaningful…"

Seeking a Foundation for Context-Aware Computing
Paul Dourish, **2001 (!!)**

But when we look at this Dourish quote from more than 20 years, we see that even back then, we had an increasing understanding on the part of ….
[quote]

So even 25 years ago, we understood that it's important to more meaningful capture our experiences in computing systems. So why are we where we are now?

Conceptualizing a human experience across many contexts is a really complex problem

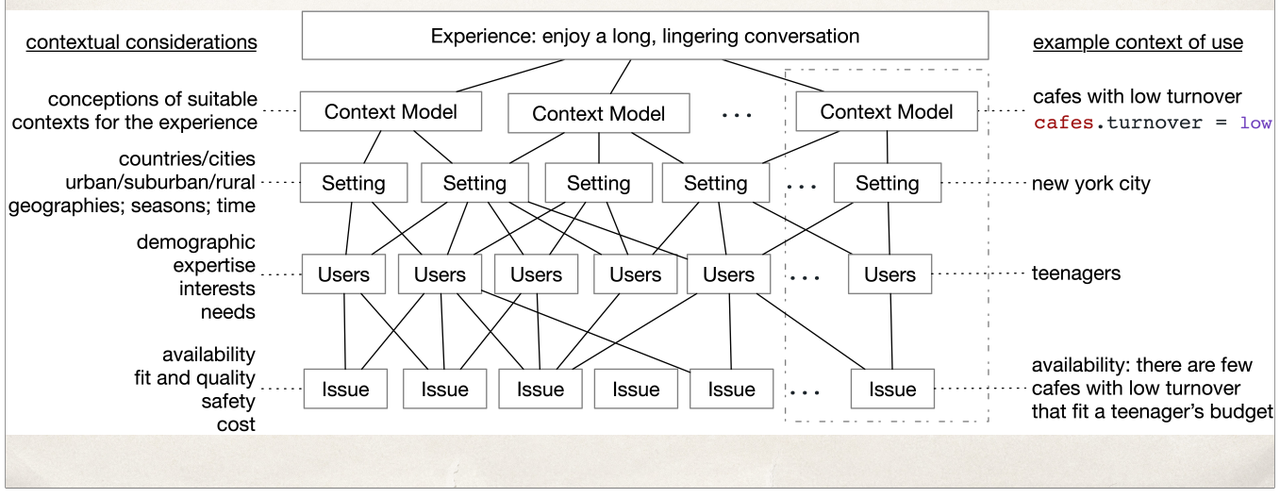

A core issue is that in actuality, there is tremendous variation and complexity in how human experiences are realized across contexts, that is rarely captured nor highlighted in existing systems.

Instead of ignoring this, we created a computational platform that can reason about it explicitly, called Differ.

Differ: An Experiential Computing Platform

[He, Zhang, Liu, Chopra, Lu, Gergle, Z., under review]

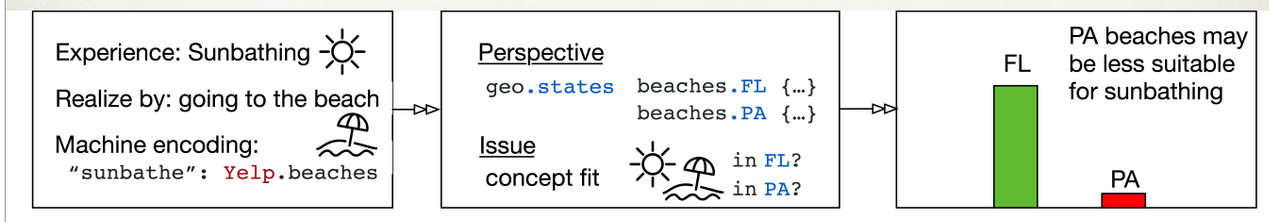

1. Define human experience and its computational encoding

2. Define accountable perspectives and issues of concern

3. Automatically surface differences to designers via visualizations

Given a human experience, Differ can automatically analyze how the experience may play out in different settings, and for different people.

Example: Going on a first date

Experience: Going on a First Date over Food or Drinks
Definition: Places that are not very loud, that are romantic, and that serve food or drinks.
Machine Definition:

```
{
    this.business_type matches ["Restaurants", "Bars", "Cafes", "Food"] and
    this.noise_level matches ["Quiet", "Average"] and
    this.ambiance matches "Romantic"
}
```

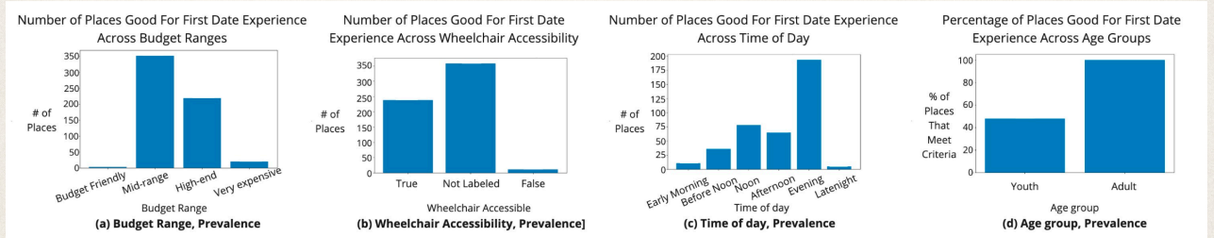

As a quick example, consider the idea of going on a first date, by going somewhere for food or drinks that is considered romantic.

While this description seems very reasonable, what Differ shows us is that it actually leaves us with very few budget friendly places;
few places that are known to be wheelchair accessible;
Or that work for people who are only available early in the day (got a night job);
or that work for younger people.

So even in this one example, what we can already see is a vision for

Vision: Experiential Computing





How can computers help us to experience the world in ways that better reflect its richness and diversity?



Vision for experiential computing, where computers can help us navigate the world in ways that better reflect its richness and diversity.

The reality is that the world is not monotone but colorful.

To appreciate its palette, we need to do more than tweak our algorithms. We need to build an unflattened Internet — one that reflects the richness of our experiences— and that inspires us to connect to the world and ourselves differently.

(Limited) role of technology in advancing human values



1. Technological advances are insufficient
2. We need computational ecosystems

As I approach the end of my talk, I want to emphasize 2 key points about the role of technology in advancing human values that have important implications on how we move forward in computing.

The first is that technological advances are insufficient

And the second, is that we need computational ecosystems

Computational ecosystems rework the entire sociotechnical configuration to enable better practices

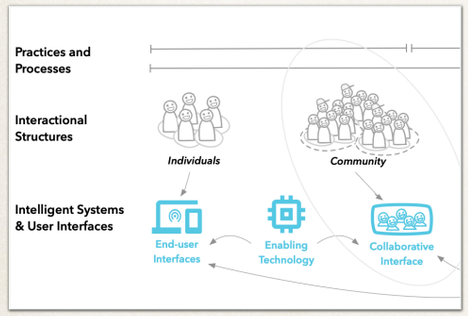

PLAN LARGE EVENTS
COORDINATE ON LOCAL PROBLEMS
CONNECT FRIENDS AT DISTANCE
LEARN TO BUILD SOFTWARE
INNOVATE THROUGH RESEARCH
DEVELOP A SELF/PRACTICE
SUPPORT HUMAN EXPERIENCES

For my first point, as you have seen throughout this talk, computational ecosystems give us a way to fundamentally change how we do things.

When we're willing to rework the entire sociotechnical configuration, we can enable practices that are not only more effective, but that better sustain our values.

Understand the limitations of computers as consequentialist machines



Computers reliably produce desired consequential outcomes

My second point, is that we can do more with computational ecosystems than we can with computers alone.

At their core, computers are input–output machines that are designed to reliably produce desired outcomes.

But what about valuable ways of being human that aren’t defined by the outcomes we achieve?

This isn’t to dismiss what computing has accomplished, but to remind everyone that outcomes alone have never captured all that’s valuable in being human.

Computers can never be the be-all and end-all to promoting human values rooted in *intrinsically valuable human activities*

[Z., CHI 2024]

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.

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

In a philosophy paper I published at CHI, I argued that computers can never be the be-all and end-all to promoting human values rooted in intrinsically valuable human activities.

I also lay out arguments for how HCI too, has been a largely consequentialist enterprise, and as a result, have largely neglected to design for intrinsically valuable human activities.

If you are interested in learning more, I also brought hard copies of the paper.

The real value of computational ecosystems is not “merely”consequential *[Z., CHI 2024]*

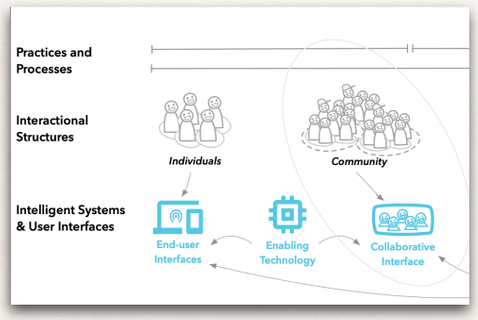

- Computational Ecosystems produce desired goods and services; they solve consequential problems
- Computational Ecosystems also promote engagement in intrinsically valuable human activities

And this leads me back to computational ecosystems. What’s important to remember about them is that they’re valuable not just for their productive capacities, but for inspiring us to engage in activities that are meaningful in and of themselves.

In the two examples I showed you in the second half of my talk: when my students and I learn to see ourselves more clearly and engage more deeply in our research — and when we seek more meaningful ways to experience the world around us— the real contribution of our designs isn’t the outcomes they achieved, but the kinds of lives that they invite us to live.

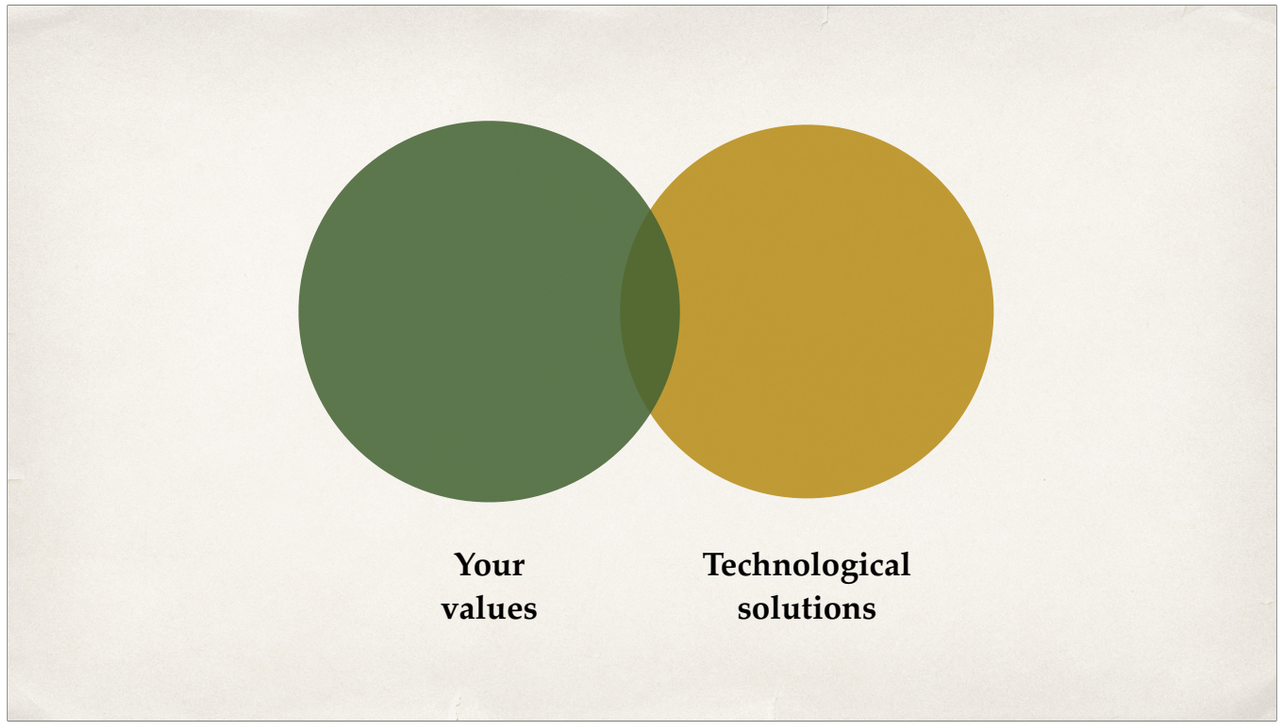

As I end my talk, I want to leave you with this question: Are your values aligned with the solutions our technologies provide?

If they are, that’s wonderful, and perhaps you’ve found a way to live a life that feels truly yours, within the sociotechnical configuration in which we all find ourselves.

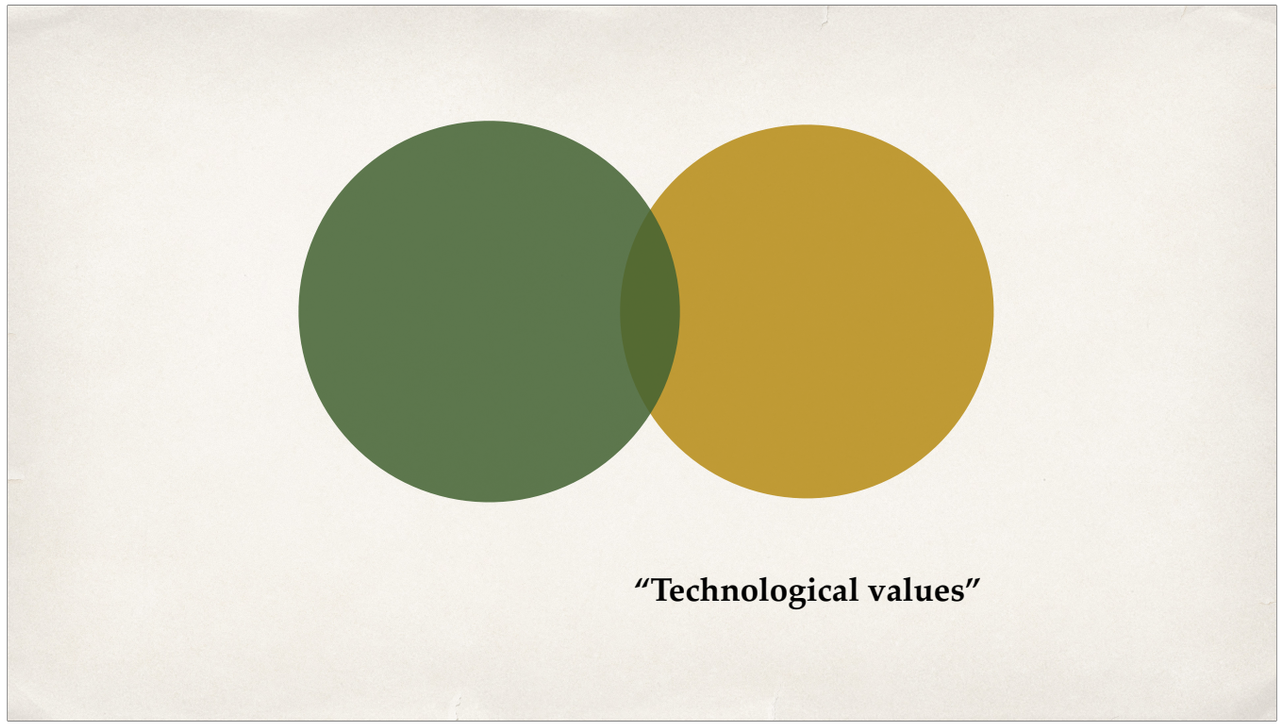

But if it’s not, it may be tempting to shift your values to align with what technology provides, and to live up to what I call “technological values”

But I think we’d all lose something if we settle for that.

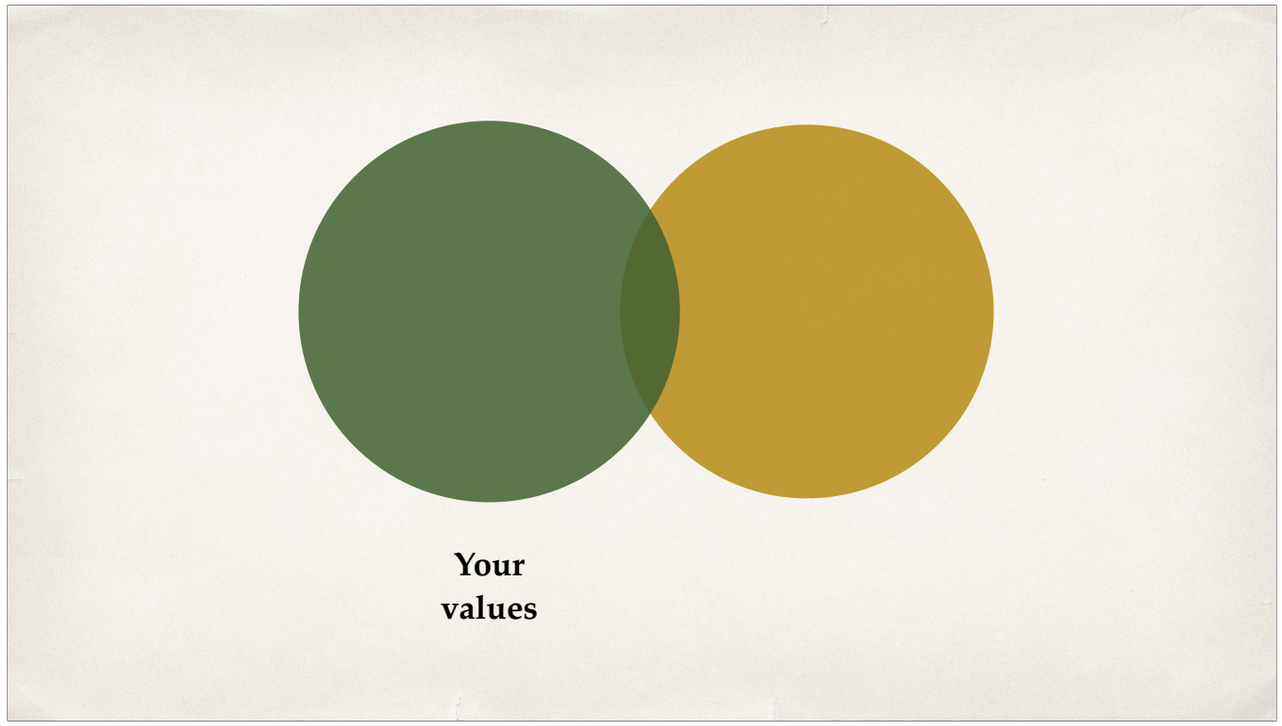

Instead, I hope you allow yourself time to find ways to realize your values.

It will require a ton of patience and new ways of seeing and doing, but it is still the more interesting question to ask, and at least for me, what I found to be the most fulfilling.

Hold onto our values.
Know the limits of technology.
Build entirely new ecosystems.

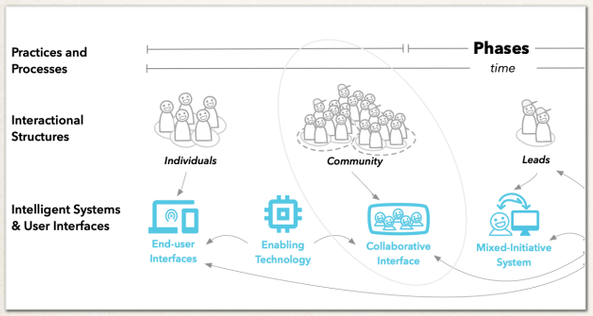

- **Computational thinking**: combining wedges of human and machine intelligence to make effective practices feasible
- **Ecological thinking**: creating processes, interactions, and communities that sustain and deepen practice

To do that, we'll need to hold onto our values, know the limits of what technology alone can do, and build entirely new ecosystems that sustain meaningful practices and ways of living.

We will need computational thinking to bring our human and machine abilities to bear, and to think about the ecology that we are creating, and in which we are embedded.

As we do this, **my hope is that we bridge the ways that we already embrace design and technology's power, with our deepest needs as humans.**

thank you

dtr.northwestern.edu/letters
forward.movie
agileresearch.io
dtr.northwestern.edu



slides+readings: haoqizhang.com

Thank you so much for being here. I am happy to take questions.

References

Critical Technical Practice and System-Level Thinking

- Philip Agre. Computation and Human Experience. Cambridge University Press, 1997.
- Atul Gawande. How do we heal medicine? TED, 2012.
- George Furnas. Future design mindful of the MoRAS. Human-Computer Interaction 15.2 (2000): 205-261.

References

Community-Informed Planning

- Lydia Chilton, Juho Kim, Paul André, Felicia Cordeiro, James Landay, Dan Weld, Steven P. Dow, Robert C. Miller, Haoqi Zhang. Frenzy: Collaborative Data Organization for Creating Conference Sessions. In Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2014), 1255–1264.
- Paul André, Haoqi Zhang, Juho Kim, Lydia B. Chilton, Steven P. Dow, and Robert C. Miller. Community clustering: Leveraging an academic crowd to form coherent conference sessions. HCOMP 2013.
- Anant Bhardwaj, Juho Kim, Steven P. Dow, David Karger, Sam Madden, Robert C. Miller, Haoqi Zhang. Attendee-sourcing: Exploring the Design Space of Community-Informed Conference Scheduling. HCOMP 2014.
- 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 Proceedings of the 26th Annual ACM symposium on User Interface Software and Technology (UIST 2013), 173–182.
- Cobi: Communitysourcing Large Scale Conference Scheduling. http://projectcobi.com.

References

Flexible Coordination

- Yongsung Kim. Designing Flexible Coordination Systems to Advance Individual and Collective Goals in Physical Crowdsourcing. Doctoral Dissertation at Northwestern University, 2020.
- Yongsung Kim, Darren Gergle, and Haoqi Zhang. Hit-or-wait: Coordinating opportunistic low-effort contributions to achieve global outcomes in on-the-go crowdsourcing. CHI 2018.
- Kapil Garg, Yongsung Kim, Darren Gergle, and Haoqi Zhang. 4x: A hybrid approach for scaffolding data collection and interest in low-effort participatory sensing. CSCW 2019.

References

Opportunistic Collective Experiences

- 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. CSCW 2020.
- Ryan Louie, Darren Gergle, and Haoqi Zhang. Affinder: Expressing concepts of situations that afford activities using context-detectors. CHI 2022.
- Ryan Louie. Human-AI Interface Layers: Enhancing Communication of Intent for AI-Assisted Creative Pursuits and Social Experiences. Northwestern Doctoral Dissertation, 2023.

References

Readily Available Learning Experiences

- Josh Hibschman, Darren Gergle, Eleanor O'Rourke, and Haoqi Zhang. Isopleth: Supporting sensemaking of professional web applications to create readily available learning experiences. TOCHI 2019.
- Sarah Lim, Josh Hibschman, Haoqi Zhang, and Eleanor O'Rourke. Ply: A visual web inspector for learning from professional webpages. UIST 2018.
- Josh Hibschman. Readily Available Learning Experiences in Production Code. Northwestern Dissertation, 2017.
- Gobi Dasu, Eleanor O'Rourke, and Haoqi Zhang. Process Management for Learning from Professional Source Code: Cultivating Experts in the Age of AI. ACM TOCE, 2026.

References

Agile Research Studios

- Leesha Maliakal Shah. Agile Research Studios: Learning Ecosystems to Scale Effective Research Training. Doctoral Dissertation at Northwestern University, 2023.
- Haoqi Zhang, Matthew W. Easterday, Elizabeth Gerber, Daniel Rees Lewis, and Leesha Maliakal. Agile Research Studios: Orchestrating Communities of Practice to Advance Research Training at Scale. Proceedings of the 20th ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW 2017), 220–232.
- Robert C. Miller, Haoqi Zhang, Eric Gilbert, and Elizabeth Gerber. Pair Research: Matching People for Collaboration, Learning, and Productivity. Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '14), 2014.
- DTR Annual Letters. http://dtr.northwestern.edu/letters
- Forward: A Story About Learning and Growth. http://forward.movie, 2020.
- Agile Research University: Supporting Authentic Research Experiences at Scale. http://agileresearch.io

References

Situated Practice Systems

- Kapil Garg, Darren Gergle, and Haoqi Zhang. Situated Practice Systems: A Computational System for Supporting the Coaching and Practice of Regulation Skills for Innovation Work. CSCW 2026.
- Kapil Garg, Darren Gergle, and Haoqi Zhang. Orchestration Scripts: A System for Encoding an Organization's Ways of Working to Support Situated Work. CHI 2023.
- Kapil Garg. Situated Practice Systems: Developing Worker's Capabilities for Complex Work in Networked Workplaces. Doctoral Dissertation at Northwestern University, 2024.
- DTR Annual Letters. http://dtr.northwestern.edu/letters
- Forward: A Story About Learning and Growth. http://forward.movie, 2020.

References

Situated Practice Systems

- Molly Pribble, Neha Sharma, Haoqi Zhang, Leesha Maliakal Shah. MindYoga: Scaffolding the Metacognitive Reflection Process within Learning Ecosystems. CHI 2022 Late Breaking Work.

- Spencer E. Carlson, Daniel G. Rees Lewis, Leesha V. Maliakal, Elizabeth M. Gerber, Matthew W. Easterday. The design risks framework: Understanding Metacognition for iteration. Design Studies 2020.

- Spencer E. Carlson, Leesha V. Maliakal, Daniel G. Rees Lewis, Jamie Gorson, Elizabeth M. Gerber, Matthew W. Easterday. Defining and Assessing Risk Analysis: The Key to Strategic Iteration in Real-World Problem Solving. ICLS 2018

- Daniel G. Rees Lewis, Jamie Gorson, Leesha V. Maliakal, Spencer E. Carlson, Elizabeth M. Gerber, Christopher K. Riesbeck, Matthew W. Easterday. Planning to Iterate: Supporting Iterative Practices for Real-world Ill- structured Problem-solving
ICLS 2018

References

Experiential Computing

- Jackie He, Shirley Zhang, Elara Liu, Medini Chopra, Nicole Lu, Darren Gergle, and Haoqi Zhang. Differ: A Platform for Experiential Computing. In preparation, 2026.

References

Technology's role in advancing human values

- Haoqi Zhang. Searching for the Non-Consequential: Dialectical Activities in HCI and the Limits of Computers. CHI 2024.
- Talbot Brewer. Retrieval of Ethics. Oxford University Press, 2009.
- [See also: EM Forster. The Machine Stops. feedbooks, 1909]
- [See also: Nicholas Carr. The Shallows: What the Internet is doing to our brains. WW Norton & Company, 2011]
- [See also: Batya Friedman and Peter H. Kahn. Human values, ethics, and design. The human-computer interaction handbook (2003): 1177-1201]