

Computational Ecosystems

Advancing Human Values Through
Integrative Computing and Changing Practice



Haoqi Zhang

Design, Technology, and Research (DTR) | Northwestern University

slides+readings: haoqizhang.com

Hi everyone, thank you for the wonderful introduction, and thank you all for coming.

- My name is Haoqi, and in this talk, I want to share with you my decade long 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 are inspired to realize yours.

And with that, I like to share some thoughts on computational ecosystems, and what I think they can do for advancing human values.

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

But when it comes to my research, I seem to keep hitting upon this question, which is, ___

As someone who designs technologies, this is a really unfortunate question for me to be asking, especially when HCI is totally winning.

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



As a small example of this, I want to briefly show you a picture of my students in the Design, Technology, and Research program at Northwestern, where I work with students on independent research. I have had lots of successful students who have gone on to do amazing things, 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 in the first place, 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 rethinking of the entire sociotechnical configuration.

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

PLANNING LARGE EVENTS

CONNECTING FRIENDS AT DISTANCE

LEARNING TO BUILD SOFTWARE

INNOVATING THROUGH RESEARCH

To illustrate the value of thinking this way, I want to show you four 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



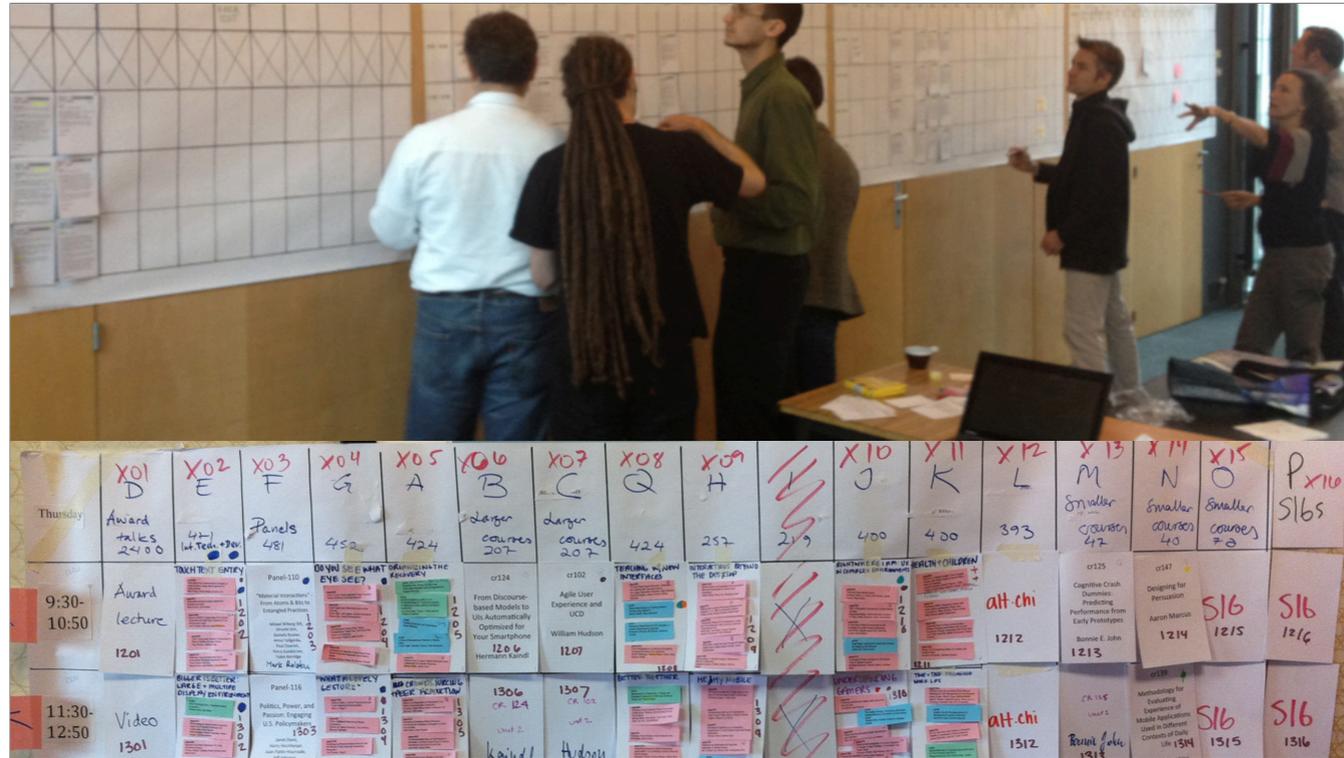
Rob Miller
MIT



Haoqi Zhang
Northwestern

Starting with how we plan 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 to come up with a preliminary schedule for CHI, where they try to construct hundreds of sessions over two days, on paper.

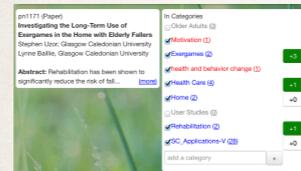
Immediately, you can see the limits of this approach.

- The first is that organizers lack information about the diverse preferences, constraints and knowledge held by community members
- The second, is that organizers 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, you can't easily run an optimizer to find the best schedule, no matter how good your software optimizer is.

So instead of just building better scheduling software, we reimagined the process of planning itself,

New approach: community-informed planning

committeesourcing

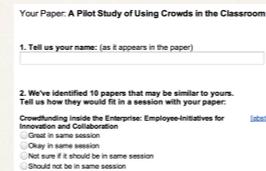


make sessions

[Chilton et al., CHI '14]

expert
categories

authorsourcing

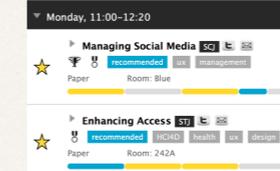


collect affinities

[Andre et al., HCOMP '13]

affinities
as seeds

attendeesourcing



collect preferences

[Bhardwaj et al., HCOMP '14]

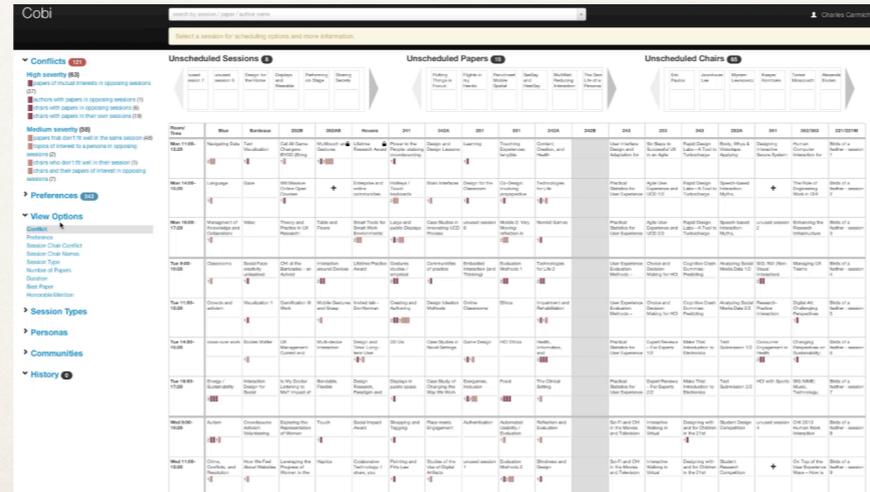
And came up with a new approach that we call, community informed planning.

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

As we did this, we also 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]



Once we gathered all this community data, our second core idea is a community-informed, mixed-initiative interface, that put organizers in the driver's seat and provide a holistic view of all constraints and preferences.

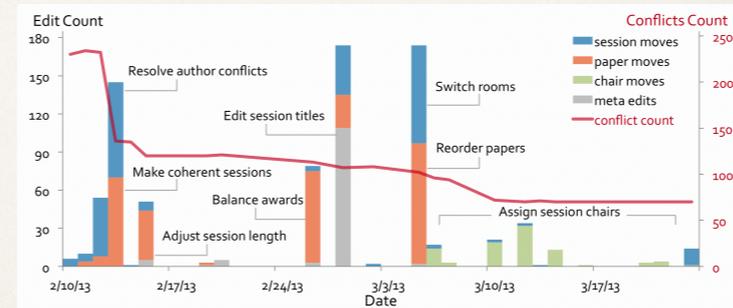
[demo]

This tool gives the organizers a multidimensional view of the schedule, and allows them to make changes 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 useful guidance of what may be good options and what not to do, while retaining control over 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 what you see in the red line on the right...which shows the number of conflicts going down over time as they are resolved, but what's interesting is that organizers improve the schedule even when the number of conflicts aren't going down.

In a way, our process allows organizers to not only engage in the consequential task of resolving conflicts, but to actively exhibit their "care" for creating a "good" schedule and for shaping a "good" conference. This is a really important idea, that I will come back to later in my talk.

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 move on to look at the challenge of connecting with friends and loved ones who are at distance.

This is work led by my former PhD student Ryan Louie, who is currently at Stanford and on the job market, 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 are largely 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

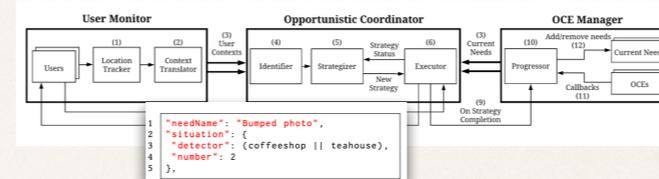
OCEs are computer programs that find opportunities for friends and loved ones to engage in shared activities and to connect at a distance.

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]



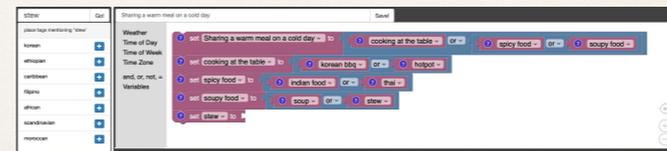
novel interaction model



OCE programming model & coordination engine

OCE Interaction Structures and Examples	Relevant Mechanisms	Design Guidelines
Shared Experiences: Connect People Who Are Doing Similar Activities in Similar Situations in Distributed Contexts. 	Shared attention to stimuli, or awareness that others are attending to the same stimuli, promotes co-experiencing and thinking about an experience. Embedded memory and anchoring, and the general coordination of movements, speech, and activity patterns, enhance the smoothness of interactions and foster liking.	Identify situations that provide a shared context for interacting and for joint actions, e.g., that contain common object affordances or afford similar actions across situations. Support people being aware of shared aspects of their individual situations and contexts to make co-referencing easy. Increase psychological proximity by using visual reminders to increase the salience of an experience.
Interdependent Activities: Construct a Digital Artifact Together or Achieve Collective Goals by Making Contributions through Respective Local Situations. 	Cooperation, or working toward activities with shared achievements, supports bonding and a sense of being part of a group. Positive interdependence in shared activities promotes meaningful interactions and increases social support and social closeness.	Create experiences in which the diversity of situations and contexts across people allows individuals to make unique contributions to achieving the interdependent activity. Create links across individual activities to increase positive interdependence and surface these links so people see how individual activities contribute to shared outcomes and artifacts.

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 what facilitates social connections, and thought about how the theory of social connections and how they happen in person, can be translated to this remote setting, when people for instance, can't always 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.

LEARNING TO BUILD SOFTWARE



Josh Hibschan



Sarah Lim



Gobi Dasu

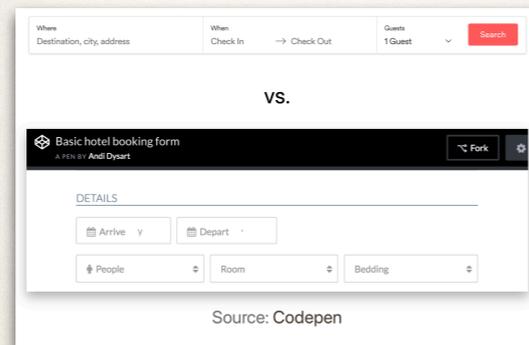
Dissertations: Josh Hibschan. *Readily Available Learning Experiences in Production Code*, 2018.

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

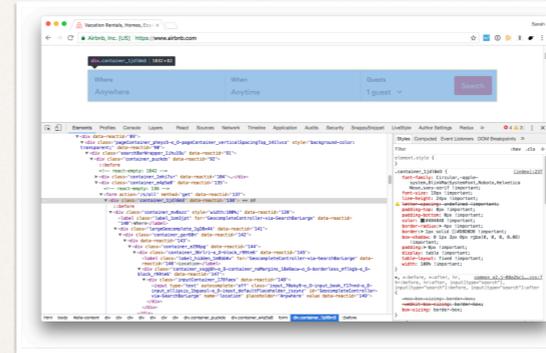
Funding: National Science Foundation, Cyberlearning EXP: *Readily Available Learning Experiences: Turning the Entire Web into Progressive Examples to Bridge Conceptual Knowledge Gaps for Novice Web Developers*.

Moving on, let's talk about learning to build software. For about a decade now, my students Josh Hibschan, Sarah Lim (who is now at Berkeley), and Gobi Dasu have been reimagining how novice developers learn the conceptual knowledge and ways of thinking that professionals have—and what that transition could look like.

Tutorials aren't quite right.



Any webpage can be inspected...



But tools are overwhelming for novices

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

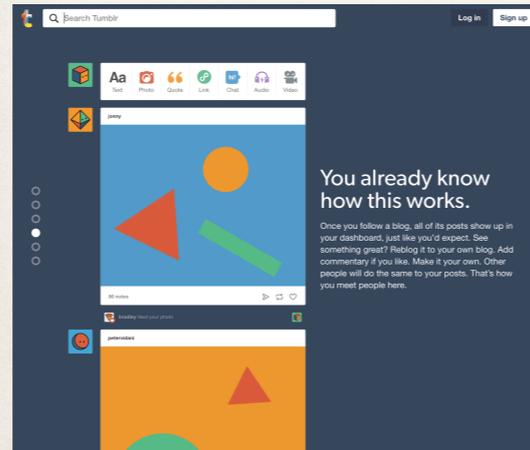
On one hand, we have tutorials, but they aren't quite right. So on the top left you are looking at is Airbnb's search bar, and on the bottom is a tutorial for making something similar. But the Airbnb one is just nicer, and knowing how to build the tutorial one on the bottom, doesn't mean you can build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

[Hirschman, 2018]

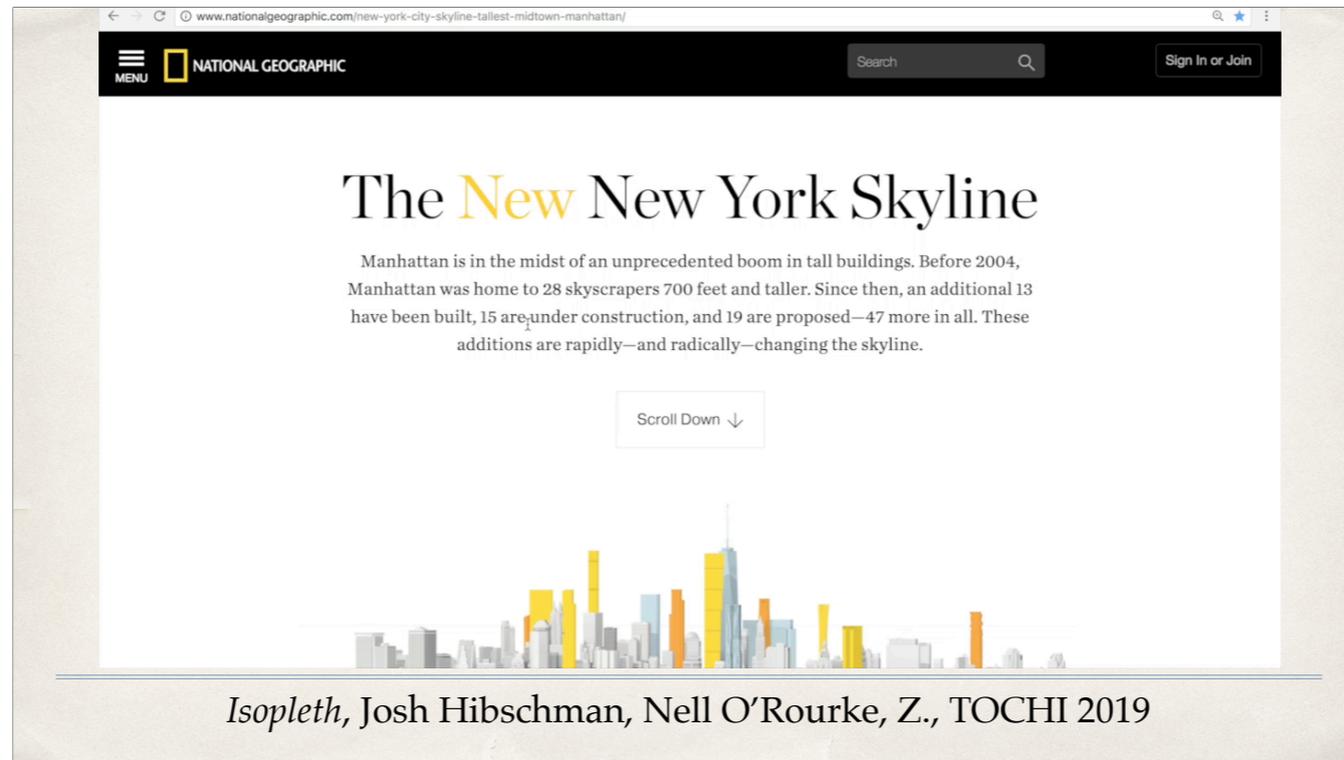


Transform *all* professional websites into a learning resource.

Support novices making sense of complex code.

So in our work, 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. Instead of relegating novices to simple tutorials, we are going to support them to make sense of complex code directly.



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

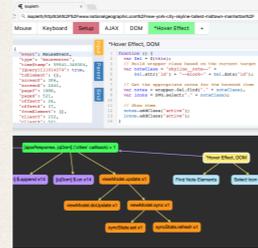
What you are seeing here is an interactive website of the New York City Skyline from National Geographic.

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

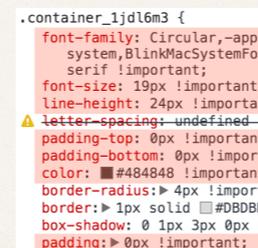
Isopleth [Hibschman et al., TOCHI 2019]



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, what we find is that they are really effective.

We found significant improvements in both program comprehension and replication tasks 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 actually went further and found that novices struggled not only with having good tools, but with having good strategies for sensemaking, and for learning.

For example, when faced with a set of professional examples of web layouts, we found that novices lacked good strategies for looking beyond surface differences, building their understanding over time, and transferring understanding across multiple examples.

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

The diagram illustrates a process for learning from professional source code, structured into six main steps:

- Step 1:** List layout features. This involves identifying examples (Ex. A through Ex. F) and making initial groupings based on prior knowledge and intuition about visual features (VFs). A sub-step, Step 1a, involves finding examples that have those visual features and categorizing them in.
- Step 2:** Compare and contrast similar examples to identify granular differences (Schwartz). This involves listing differences (diff 1, diff 2, diff 3).
- Step 3-4:** Given the intuitive visual start (Quintana), now deep dive into identifying the code behind visual similarities and differences. This involves identifying groups (Ex. 1, Ex. 2, Ex. 3) and differences (diff 1, diff 2, diff 3). A specific example of CSS Techniques/VFs Left is shown: `grid-template-columns...`.
- Step 5:** Identify groups of visual features and CSS techniques not yet covered and jump back to step 1 to cover more of them through a new cycle with a new overarching visual similarity.
- Step 6:** Review and find cases where similar visual features are produced by different CSS techniques, and cases where similar CSS techniques produce vastly different visual outcomes.

On the left side of the diagram, several professional source code examples are shown, including:

- Smashing Magazine
- Flat Icons
- Masterclass
- Hero Icons
- CSS Tricks

This led us to focus not only on giving novices better tools, but on creating scaffolds for managing their sense making process.

We then embedded these scaffolds in software, to completely change how novices approach professional examples in the first place.

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

Step 1 - Make initial groupings based on prior knowledge and intuition about visual features (VFs).

Step 1a - Find examples that have those visual features and categorize them in.

Step 2 - Compare and contrast similar examples to identify granular differences (Schwartz).

Step 3-4 - Given the intuitive visual start (Quintana), now deep dive into identifying the code behind visual similarities and differences.

Step 5 - Identify groups of visual features and CSS techniques not yet covered and jump back to step 1 to cover more of them through a new overarching visual similarity.

Step 6 - Review and find cases where similar visual features are produced by different CSS techniques, and cases where similar CSS techniques produce vastly different visual outcomes.

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

Participant	Pre-test Score	Post-test Score
P2	10	30
P4	10	35
P5	15	45
P7	15	45
P9	20	50

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

Participant	Pre-test Score	Post-test Score
P1	10	15
P3	10	15
P6	15	30
P8	15	25
P10	15	25

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.

So what we did here wasn't to build a better tutorial or a better code inspector, but to change the practice of how novices engage with expert work. In 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.



Going back to the question of how we train the next generation of researchers, I want us to think about why this is so hard, and in particular, what's needed for students to drive research themselves.

Students need regulation skills

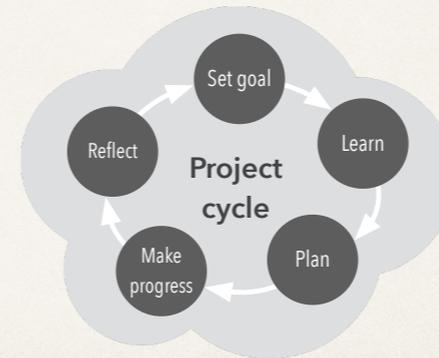
- ❖ **Regulation skills:** cognitive, metacognitive, motivational, and emotional skills for reaching a goal [Jarvela & Hadwin. 2013]
- ❖ Independent research requires regulation skills including **planning** and **seeking help** to overcome challenges.
- ❖ Students lacking these skills are confined to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

To truly become self-directed researchers, students need what are called regulation skills — or cognitive, metacognitive, motivational, and emotional skills needed to plan, and seek help while doing their research. 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 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 while they do research, and there's got to be a better way.

Agile Research Studio (ARS) [Z. et al., 2017]

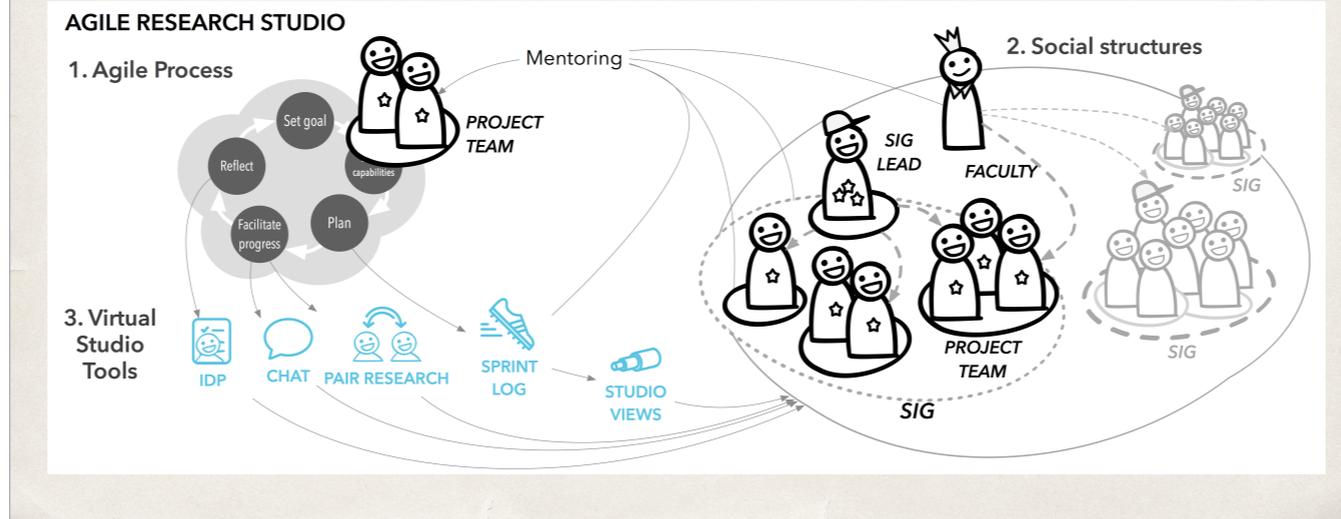
- ❖ Model for research training in a learning community
- ❖ All students, regardless of seniority, conduct independent research and receive authentic research practice.



To address this challenge, we created a sociotechnical system that we call Agile Research Studios, that provides a model for research training in a learning community.

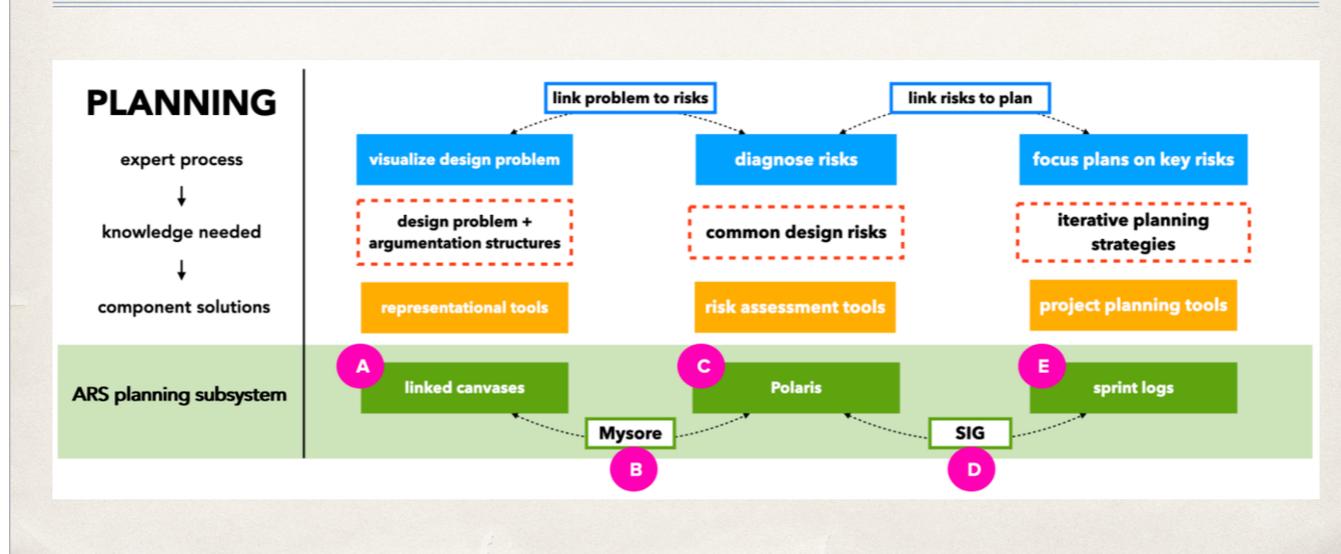
All students, regardless of how senior they are, receive authentic research practice, by which I mean that they drive the self-directed learning cycle, where they set goals, learn, plan, make progress, and reflect on what they have done and learned before setting goals again.

ARS is a sociotechnical model for developing regulation skills



At its core, ARS is a socio-technical model with processes, social structures, and tools that work together to help students develop regulation skills. To do this, we fundamentally reconfigured the roles and relationships of student researchers and mentors, and created technologies that facilitate this reconfiguration.

ARS: planning [Maliakal, 2023]



Let me first describe how the ARS model supports students learning to plan research work. 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, diagnosing risks, and then focusing plans on those key risks.

We created entire subsystems, or sociotechnical configurations, for 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 deepen their practice of how they drive their research themselves.

ARS: help & collaboration

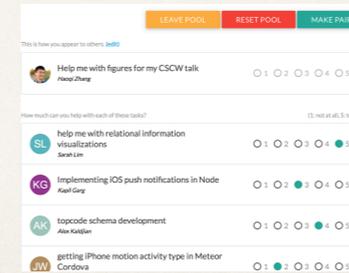
Process:
Distributed help



Social structure:
Studio meeting



Studio tool:
Pair research

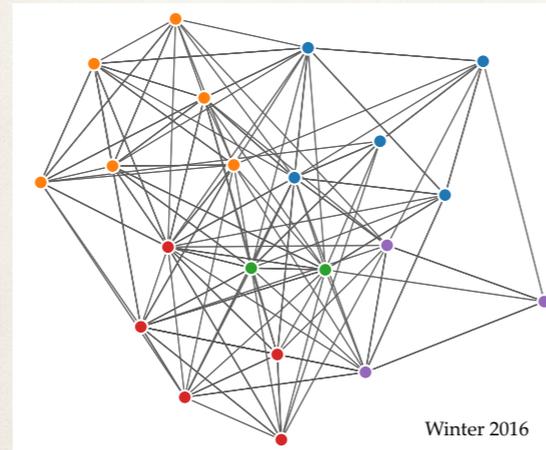


[Miller, Z., Gilbert, Gerber, CSCW '14]

On the help side, instead of relying solely on faculty mentors, ARS distributes help across the entire community. For instance, 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 those 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]

Outcomes (11 yrs)

- ❖ 184 students (160 UG, 11 MA, 15 PhDs) who led 70+ research projects.
- ❖ 50% women!
- ❖ 75 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.



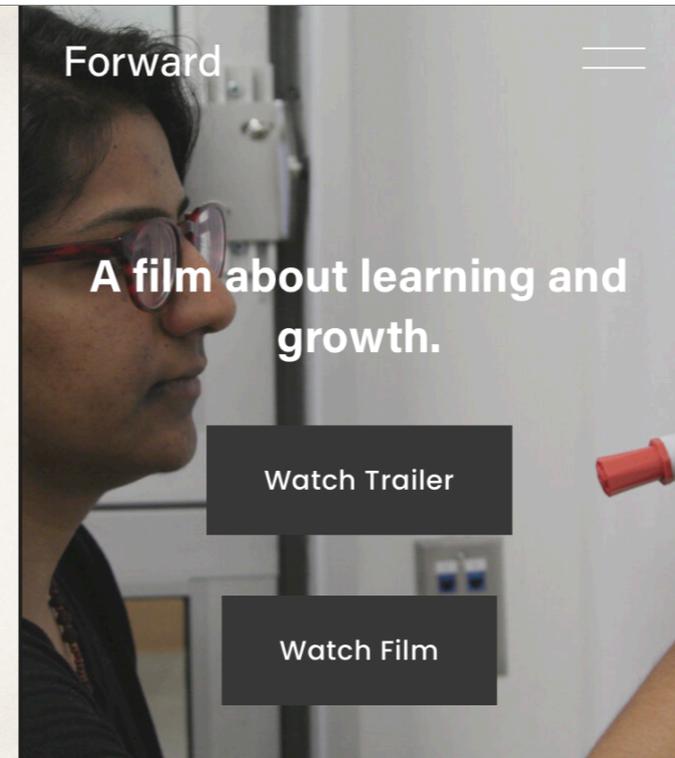
In 11 years, in my own studio at Northwestern, I have directly mentored 184 students,

[Slide]

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

Success beyond Northwestern

- ❖ Founded Agile Research University (ARU) to support 70+ faculty at universities across the world using the Agile Research Studios model (even in the humanities!)
- ❖ Produced the DTR documentary, *Forward*
See: <http://forward.movie>
- ❖ DTR annual letters to start deeper conversations on mentoring and learning
See: <http://dtr.northwestern.edu/letters>



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

In 2022 we released the DTR documentary, *Forward*, to further share our culture of learning and growing.

I also began to write annual letters to initiate dialogue on learning and growing, and the challenges we all face, as students and educators.

PLAN LARGE EVENTS

CONNECT FRIENDS AT DISTANCE

LEARN TO BUILD SOFTWARE

INNOVATE THROUGH RESEARCH

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

[Before I summarize what we did across these four systems, any questions about any particular system? PAUSE QUESTIONS]

PLAN LARGE EVENTS

CONNECT FRIENDS AT DISTANCE

LEARN TO BUILD SOFTWARE

INNOVATE THROUGH RESEARCH

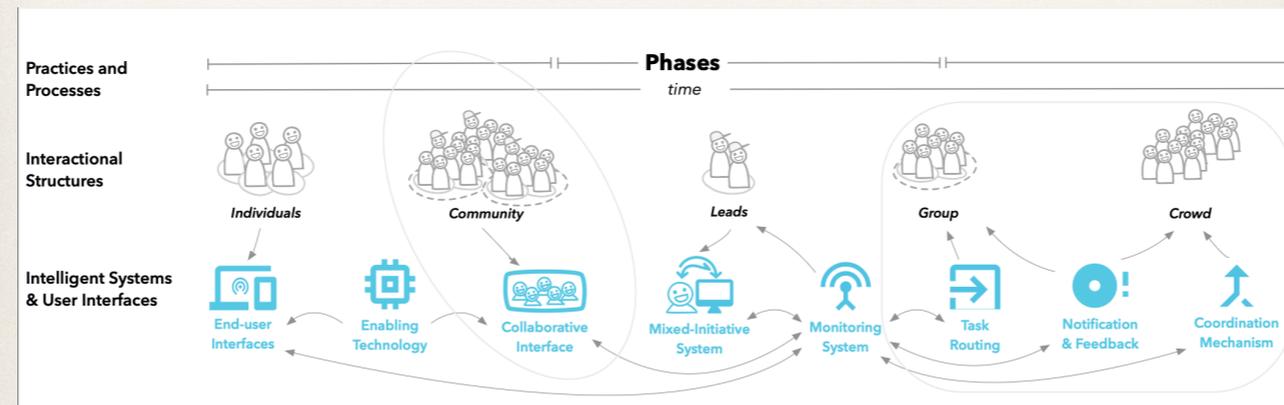
- ❖ There were fundamental limitations with existing approaches.
- ❖ These limitations could not be overcome with technology alone.
- ❖ We took a different approach.
- ❖ We created **entirely new sociotechnical configurations** to support it.

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

So what we did is to take a new approach.

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

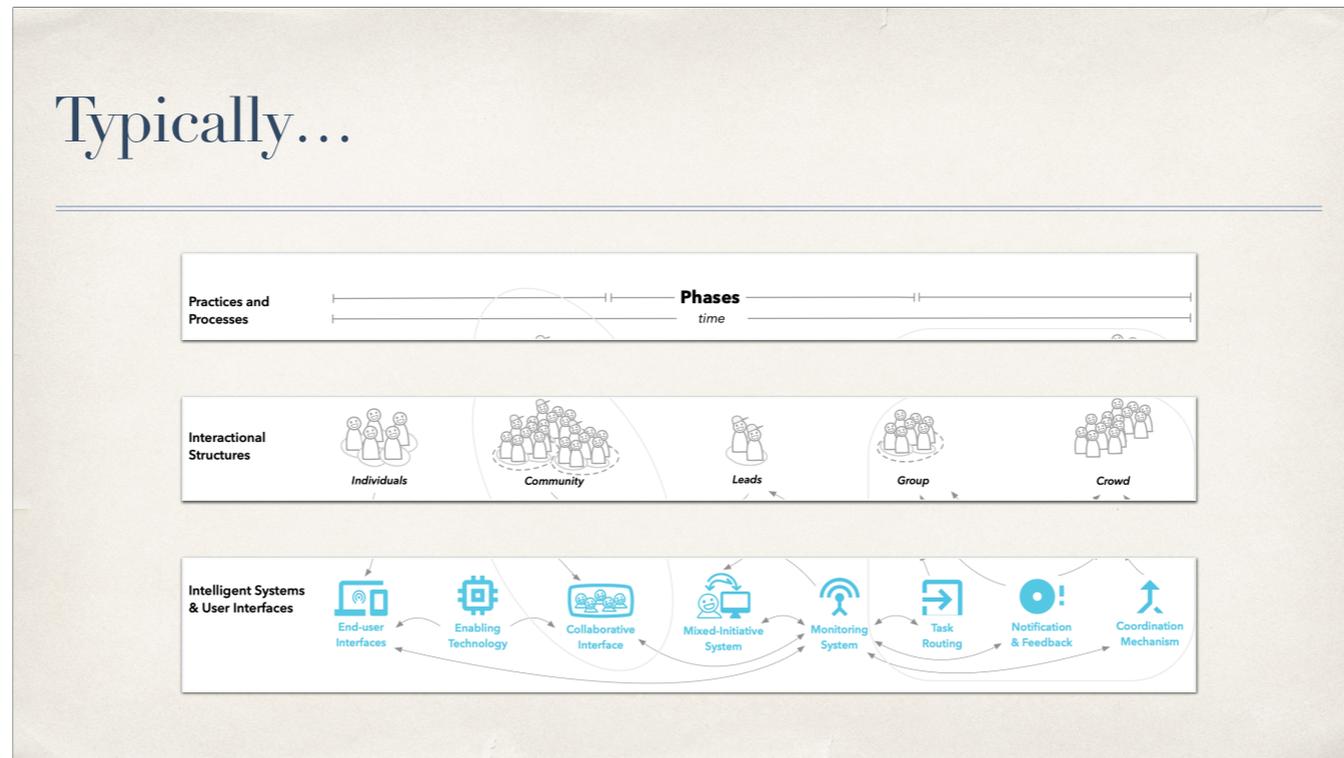
Computational Ecosystems



Computational ecosystems.

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

Typically...

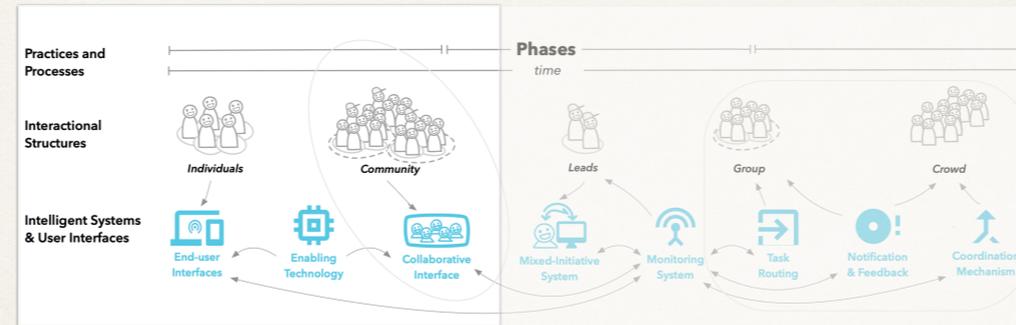


Typically, we design sociotechnical systems layer by layer:

- we define a process,
- then shape the social structures around it,
- and finally build tools to support both.

And that's 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 four 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.

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 is the great task of our generation of physicians and scientists. I would go further to say that 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 HCI, 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.

PLAN LARGE EVENTS

CONNECT FRIENDS AT DISTANCE

LEARN TO BUILD SOFTWARE

INNOVATE THROUGH RESEARCH

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

Any questions?

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 broader 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.

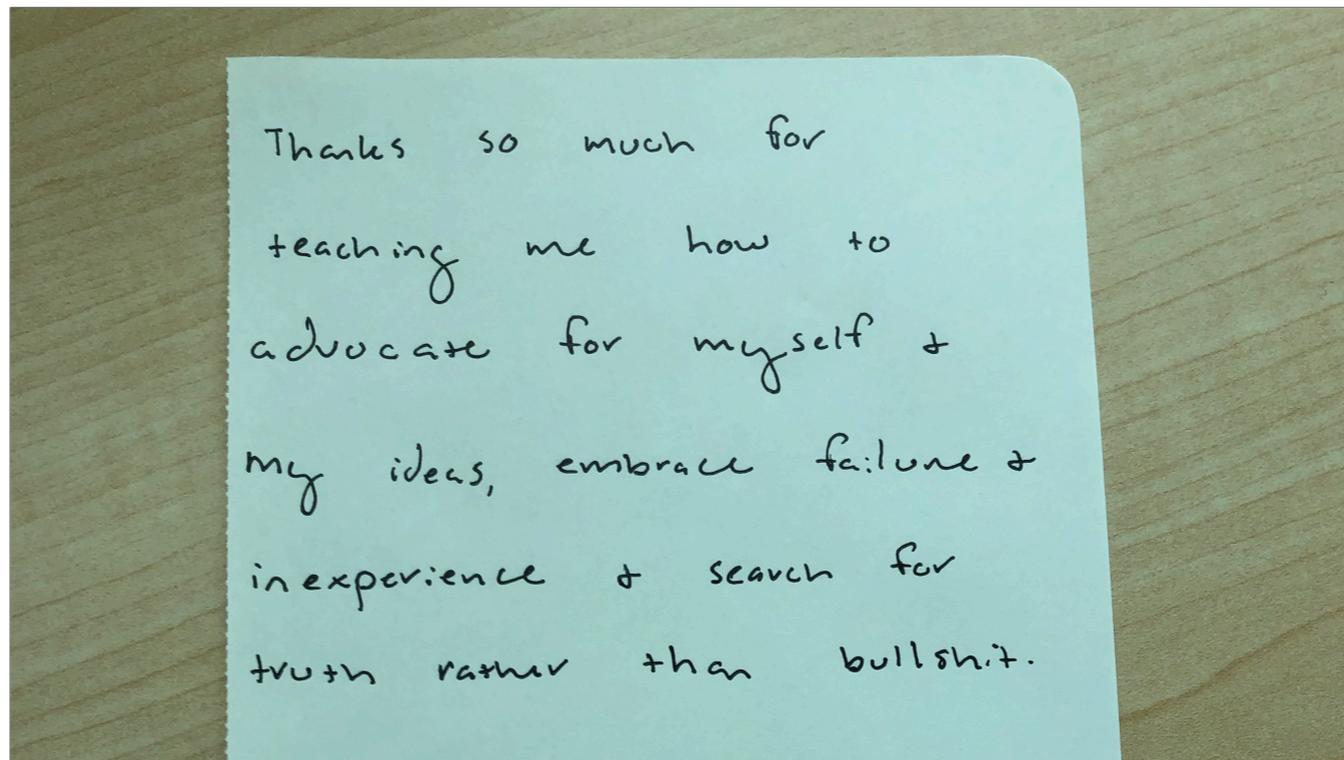
The first is that, computational ecosystems can create spaces for humans to grow, by supporting their developing a self and a practice.

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



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.

Instead, we use independent 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) practice

Cognitive skills

- * representing problem and solution spaces
- * assessing risks
- * critical thinking and argumentation
- * core design, research, and STEM methods

Metacognitive skills and dispositions

- * planning: forming feasible plans and planning effective iterations
- * help-seeking: leveraging resources; seeking help; communication skills
- * reflection: awareness of one's own skills, abilities, and metacognitive blockers

Emotional regulation and disposition toward self and learning

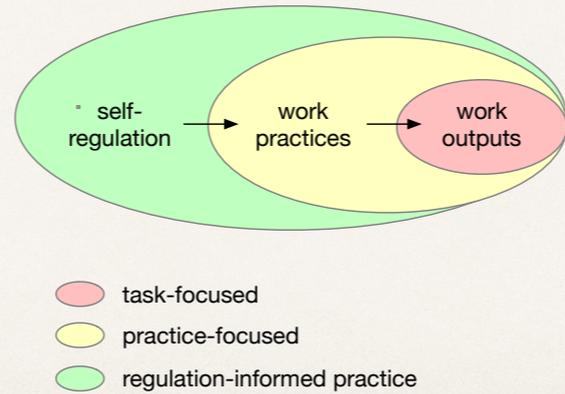
- * emotional regulation: understanding one's fears and anxieties
- * disposition: dealing with failure, embracing challenges, embracing self-direction

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.

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**.

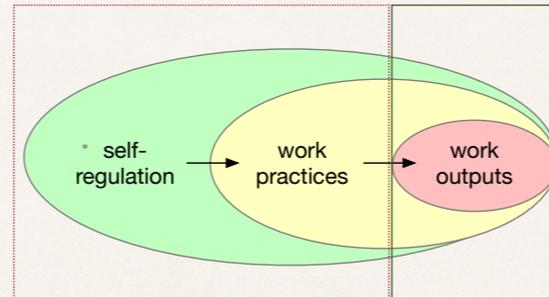


A core challenge is that understanding and addressing the gaps in students' practice and regulation is really complicated.

We understand that students' work outputs aren't everything -- and that if we want students to grow, we need them 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**.

not enough attention here mostly here

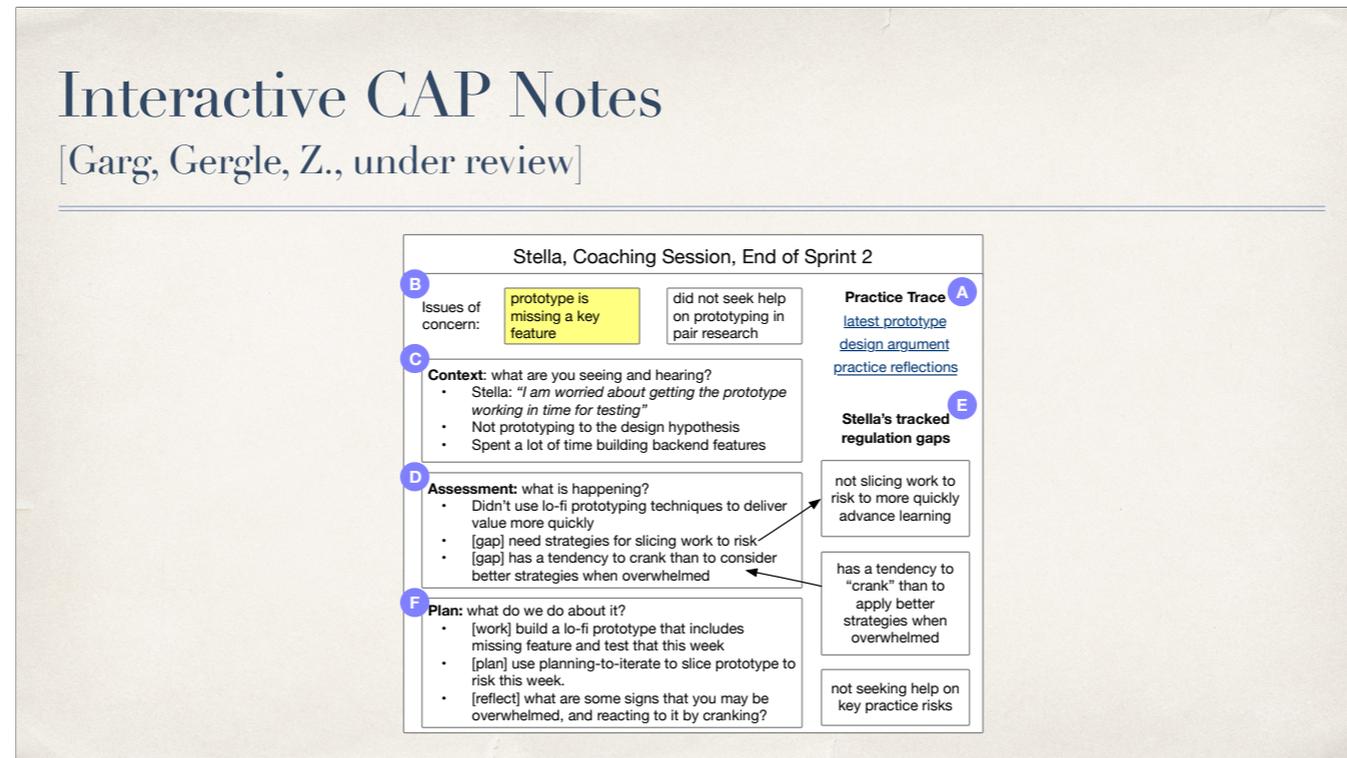


- task-focused
- practice-focused
- regulation-informed practice

However, this is often not what we focus on in project-based learning environments. A lot of effort is spent on understanding 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., under review]

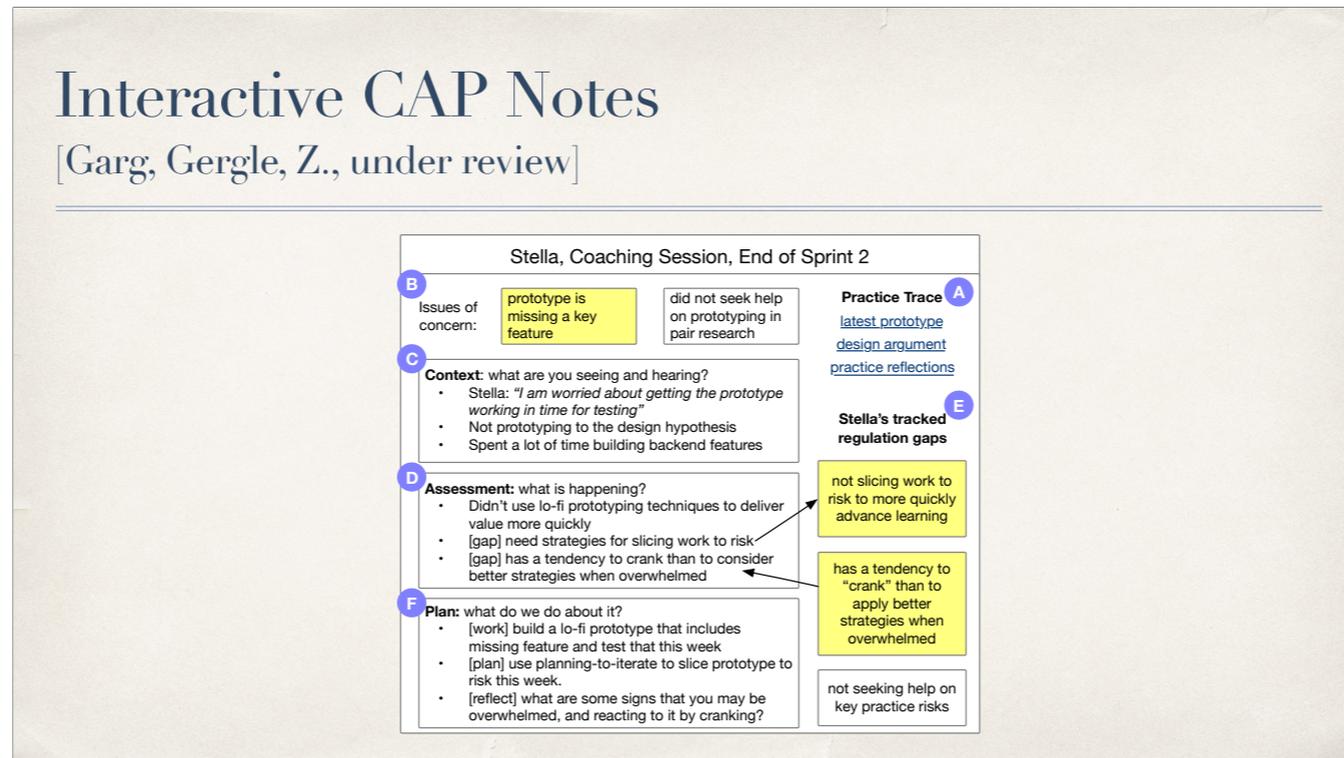


To address this, we developed a mentoring practice and tool called CAP notes, where coaching focuses on addressing issues in students practice and regulation than just their work output.

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

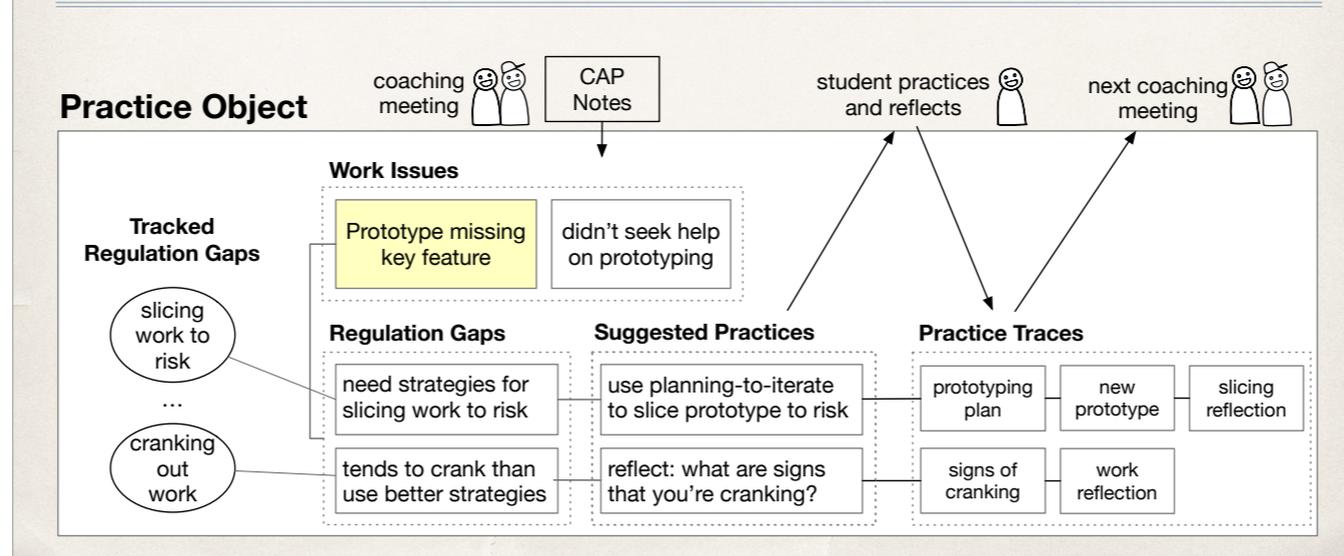
[Garg, Gergle, Z., under review]



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

So with this process and tool, a coach focuses on building a rich model of students' regulation behaviors and practices, and on addressing the regulation gaps that they discover, and not only on advancing the work output.

Practice Objects for tracking practice and regulation [Garg, Gergle, Z., under review]



To make this possible, we created practice objects, which are computational representations that capture how a student's regulation and practice evolves over time.

[Example]

This gives us a way to work with students on their underlying patterns, rather than just addressing one work issue after another across weeks.

Creating space for learning about one's self

- ❖ See students and build strong relationships
- ❖ Hold space for reflection, sharing, and acceptance
- ❖ Recognize that patterns + beliefs recur, and can be quite sticky
- ❖ Value developing regulation skills over production
- ❖ Problematize how students approach problems and look at themselves (see "Interruption, Discord, and Drama" in 2025 DTR letter)

Beyond providing computational supports, what really shifted in us 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 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
- ❖ ...

see: “The Good” and
“Beyond Production” in
the 2024 & 2025
DTR annual letters

These goods are largely not about producing research,
but about how we engage with it, deeply.

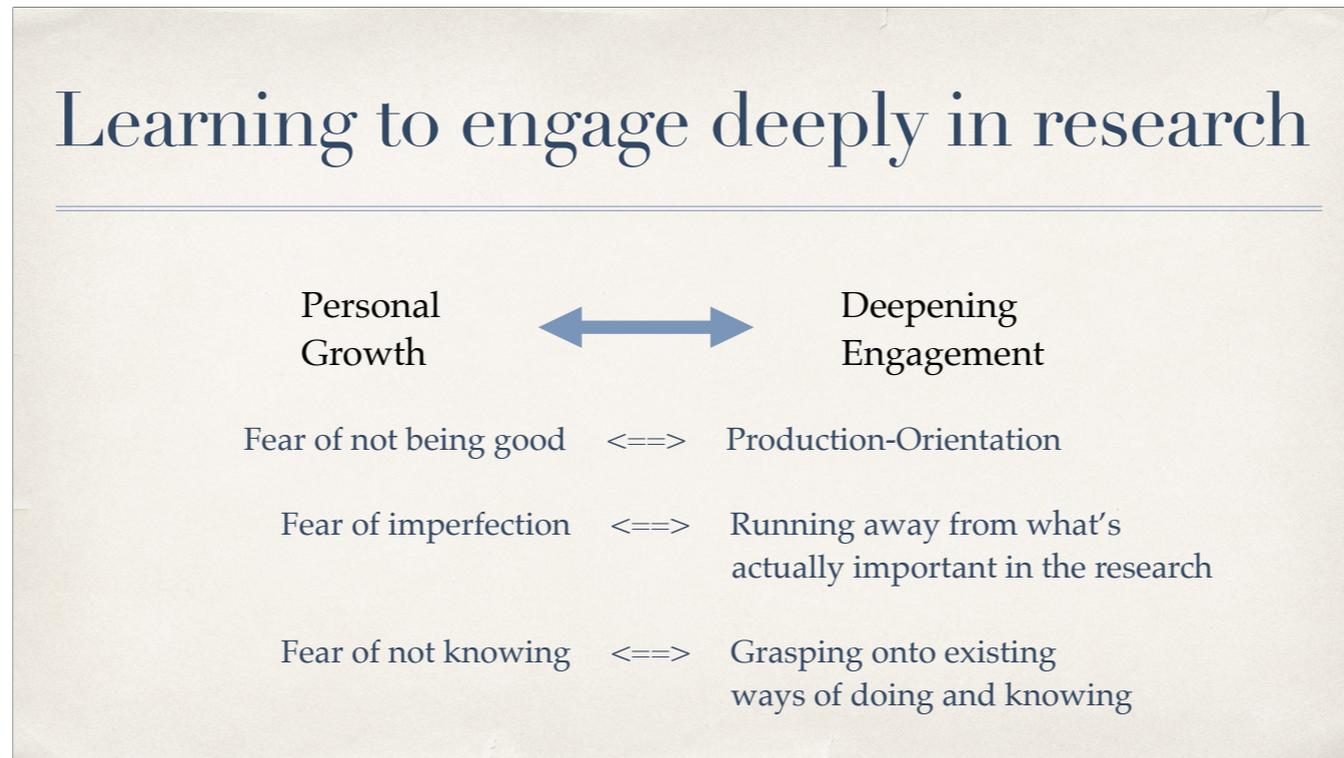
In addition to learning regulation skills, we focused on the challenge of helping students to engage more deeply in research.

What I value about research is that it can give us the 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.

If you are interested in learning more, please refer to these essays in the 2024 and 2025 DTR annual letter.

Learning to engage deeply in research



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 their 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 ~~me~~ 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 with you this note I recently received from a student, who said that until DTR,... [quote]

This note really strikes me, because for her and for all of us: How can we 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 makes these 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 also, our growing into who we are becoming.

A focus on becoming



Kalina Silverman



Meg Grasse



Katherine Lin



And this focus on BECOMING 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 growing into themselves and developing the courage and skills needed to pursuing a life that is fully their own.

As examples, Kalina started Big Talk, a social movement to foster more meaningful conversations

Meg went from leading an engineering 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 just spent the last year at the whale and dolphin conservation in Massachusetts.

I am excited for these students personally, but broadly, for shifting the dialogue on what the value of research experiences can be for our students, and the meaning it can have in their lives.

SUPPORTING HUMAN EXPERIENCES



Jackie He



Shirley Zhang

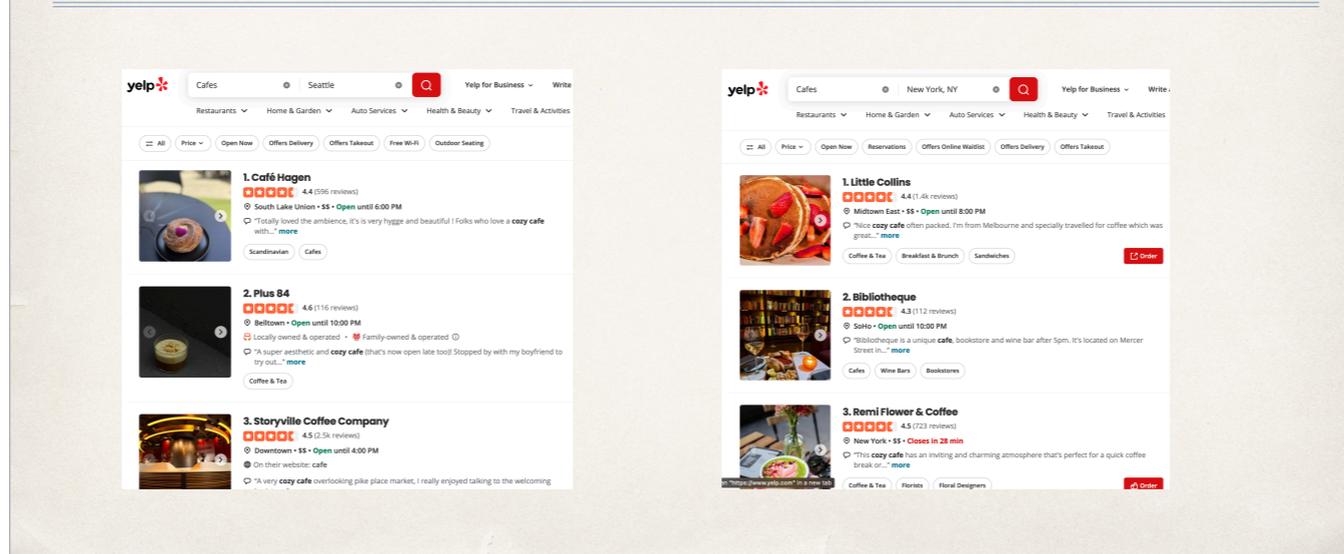
Dissertation: [yes please!]

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 navigate the outer world — by building computers that can understand and support human experience.

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

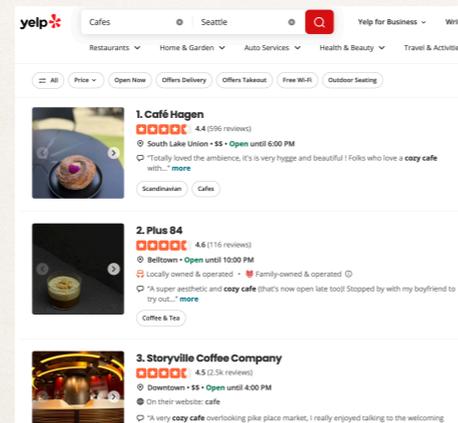


When we look at the world through the internet, it's just very flat..

You can search for Cafes in Seattle, you can search for them in New York, And you will get the best cafes there, 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?

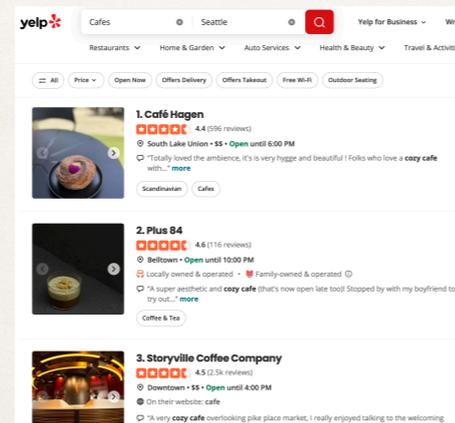
How does that differ across user
populations, geographies,
cultures?

And this is the cost of ubiquity, right? We built systems that work beautifully everywhere — but in doing so, we've replaced questions like

What makes this place special,
what experiences does it afford,
and how does that differ across user populations, geographies, and cultures?

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 be guided and consumed by searching for what's best, and valuing what's best over anything else.

Computational Understanding of Human Experiences

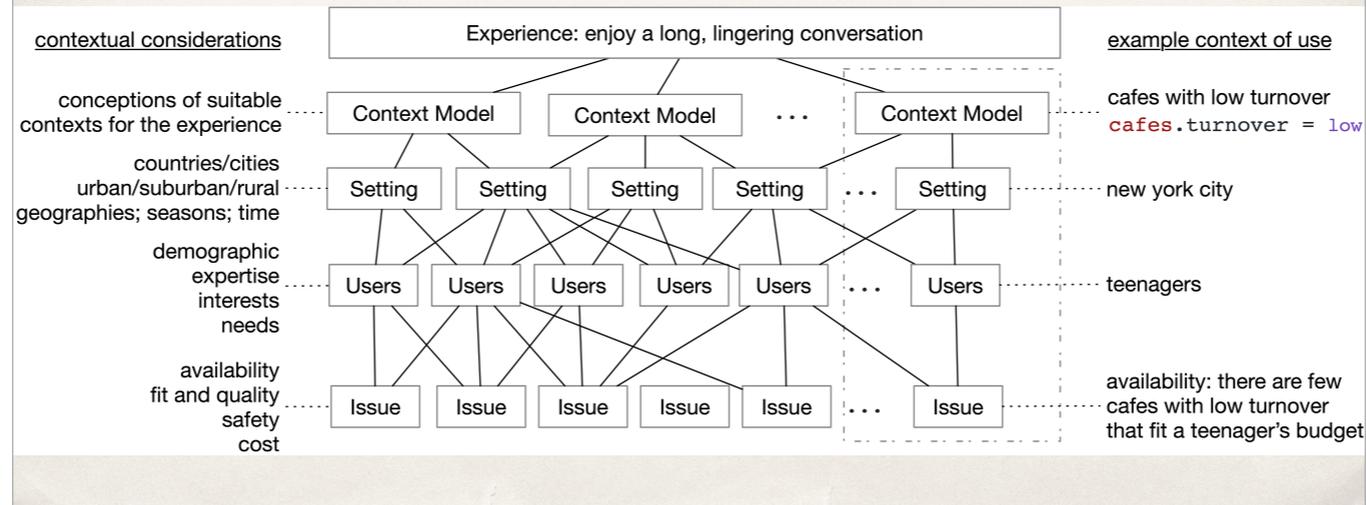
*“an increasing understanding on the part of system developers that human activities are enmeshed in a variety of [cultural and social] practices and relations that make them **meaningful** by setting a context within which they can be understood...”*

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



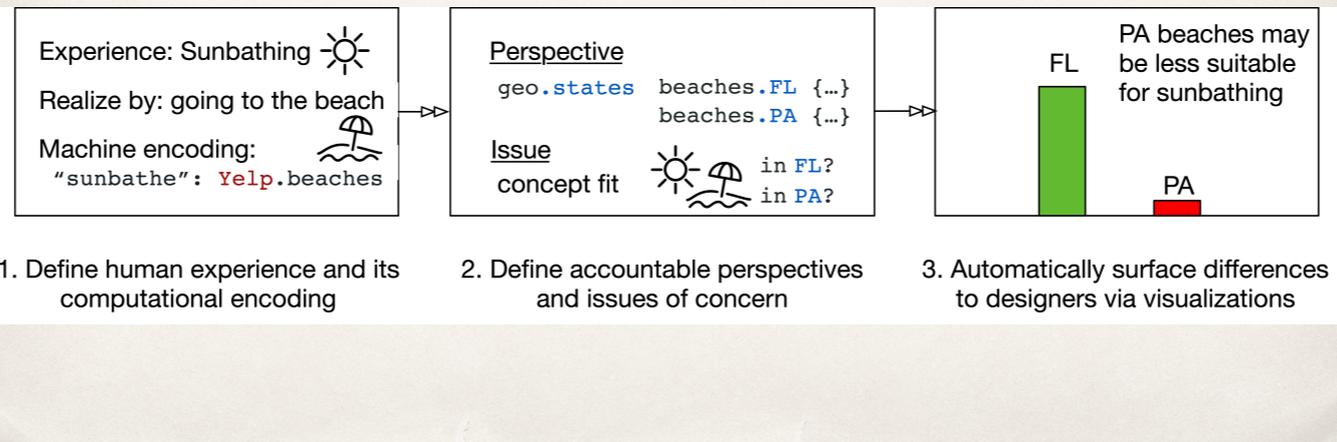
It turns out, that conceptualizing a human experience across many contexts is just really hard.

Unlike cafes, even a simple experience like enjoying a long, lingering conversation isn't really the same everywhere, and is conceptualized differently across different settings, and for different user populations.

This task is too hard for humans to think about across many contexts, and yet we don't really have machine solutions for this, which leaves us with a pretty nasty socio-technical gap.

Differ: An Experiential Computing Platform

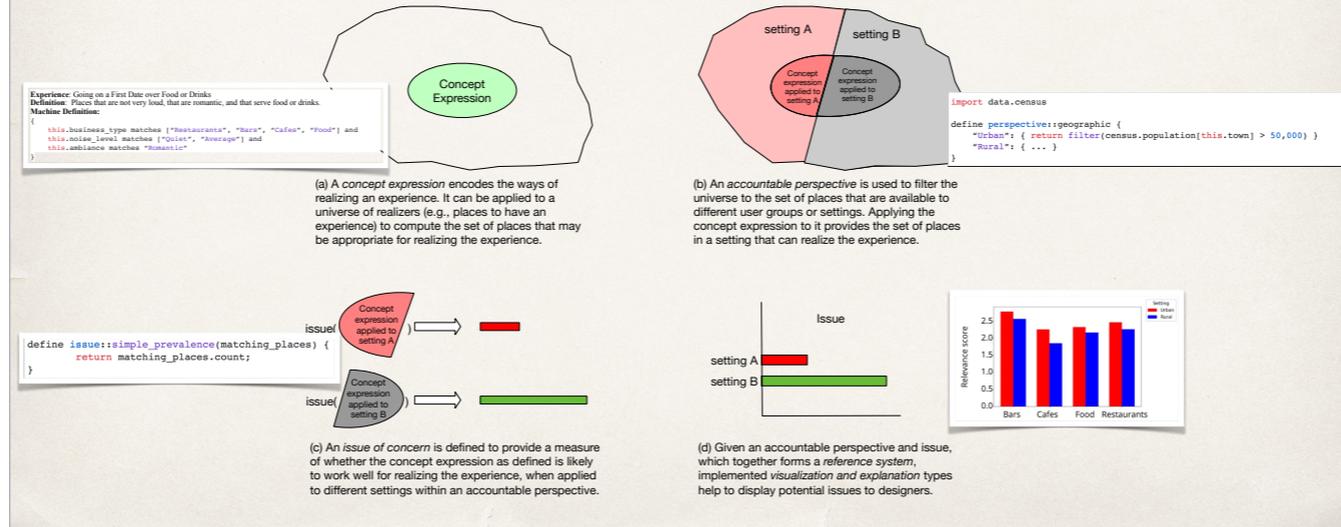
[He, Zhang, Gergle, Z., in preparation]



What I want to show you briefly is a preview of a computational platform called Differ.

What Differ does, is that given a human experience, it can automatically analyze how that experience may play out in different settings, and for different people.

Computational Abstractions for Reasoning about Human Experiences [He, Zhang, Gergle, Z., in preparation]



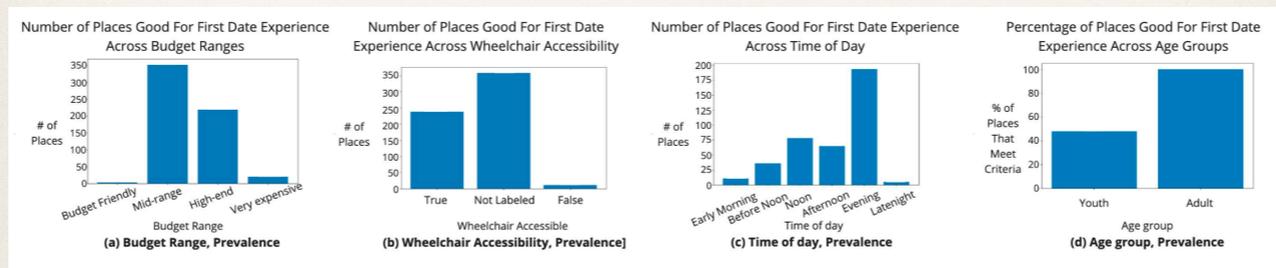
To support this, Differ contributes a set of computational abstractions for reasoning about human experiences.

I won't go into details right now, but the key idea here is that instead of relying on labeled datasets of experiences (which are rarely available across contexts), we created a computational architecture for integrating many available data sources and models to inform different aspects of experiential understanding, so as to form a composite picture of how any experience may play out for different people and in different places.

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"
}
```



Let me show you an example of what Differ can automatically generate. What we see here is that when we define going on a first date as going somewhere for food or drinks that is considered romantic,

What Differ shows us that we are actually left with very few budget friendly places;
much fewer 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 navigate the world in ways that better reflect our shared experiences, and our differences?



Vision for experiential computing, where computers can help us navigate the world in ways that better reflect both our shared experiences, and our differences.

The reality is that the world is colorful: it's beautifully diverse and full of nuance.

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 see that richness when we use it.

(Limited) role of technology in advancing human values



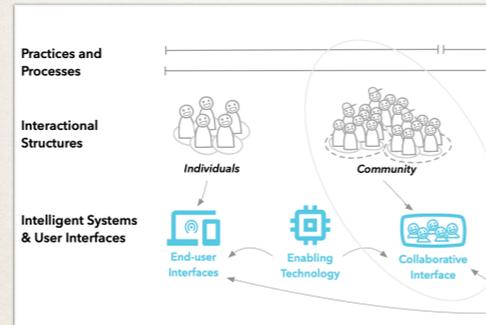
- ❖ Advancing individual technologies is insufficient
- ❖ We need computational ecosystems not just for consequential aims, but for advancing human values

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 HCI.

The first point is that advancing individual technologies is insufficient.

And the second, is that we need computational ecosystems not just for consequential aims, but for advancing human values.

Computational ecosystems rethink the entire sociotechnical configuration to enable better practices



PLAN LARGE EVENTS

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.

In HCI, we are still too focused on components. What I've shown today is that when we're willing to rethink the entire sociotechnical configuration, we can create solutions 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.

Computers, by their very nature, are consequentialist machines. At their core, computers are input–output systems designed to reliably produce desired outcomes.

But what about valuable ways of being human that aren't defined by what we produce, or by the outcomes we achieve?

This isn't to dismiss what computing has accomplished, but to remind us that outcomes alone don't capture 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

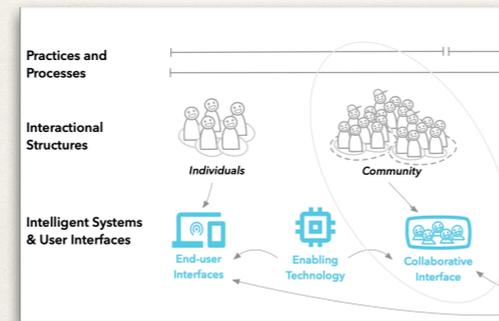
This paper examines the pervasiveness of consequentialist thinking in human-computer interaction (HCI), and forefronts the value of non-consequential, dialectical activities in human life. *Dialectical activities* are human endeavors in which the value of the activity is intrinsic to itself, including 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.

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

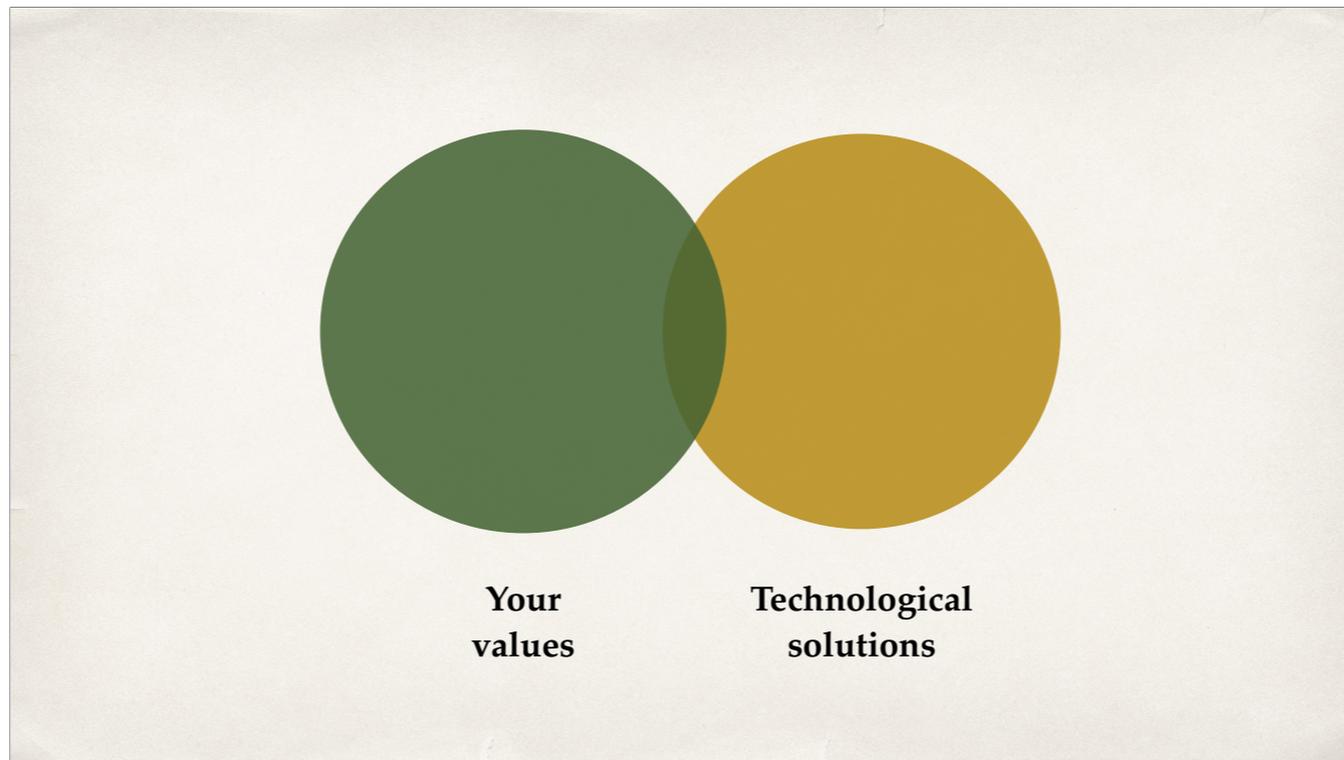


- ❖ Computational Ecosystems produce desired goods and services; they solve consequential problems
- ❖ Computational Ecosystems also promote engagement in intrinsically valuable human activities

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 how they challenge and inspire us to engage in activities that are meaningful in and of themselves.

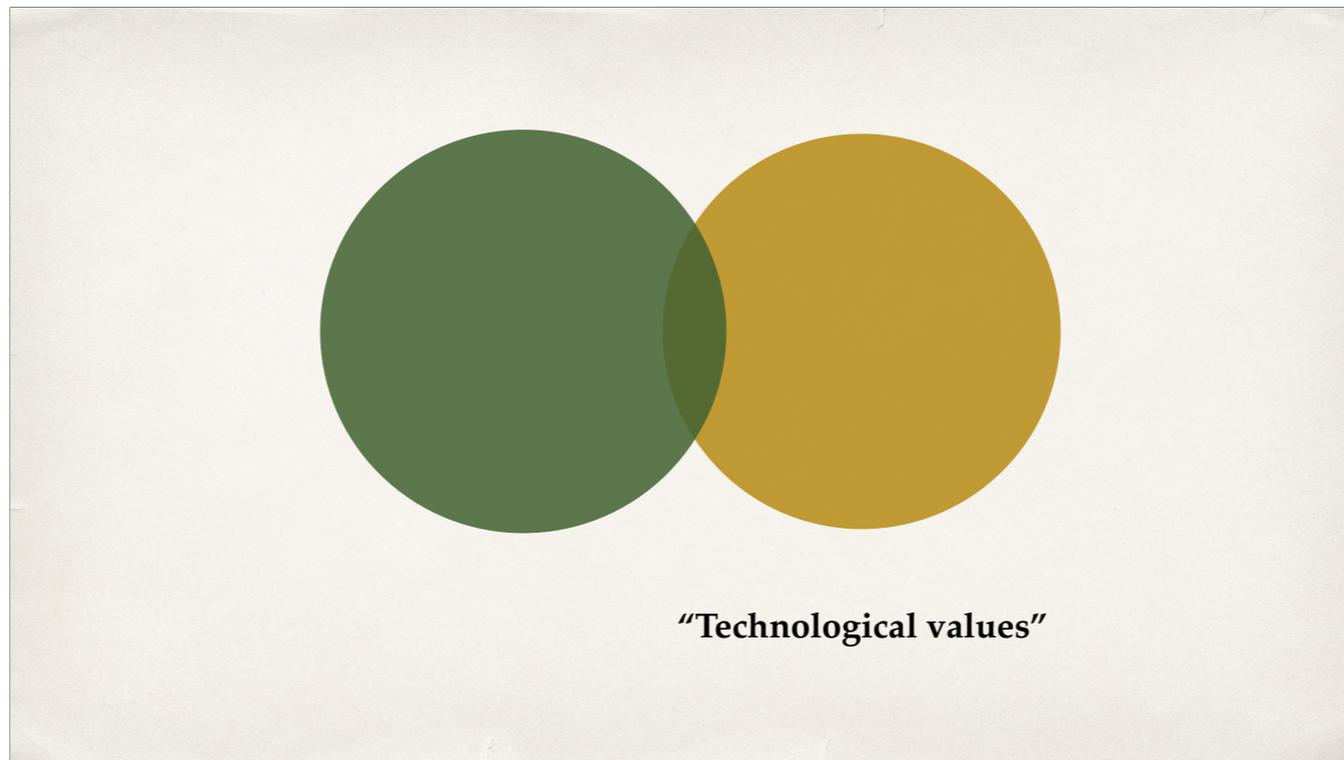
In the two examples I showed you in the second half of my talk: when my students and I learn to see ourselves more clearly and engage more deeply in our research — and when we seek more meaningful ways to experience the world around us— the real contribution of our designs isn’t just the outcomes they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

So the real question for all of us in HCI isn’t just what our sociotechnical systems can produce, but also, what kinds of lives do they invite us to live?



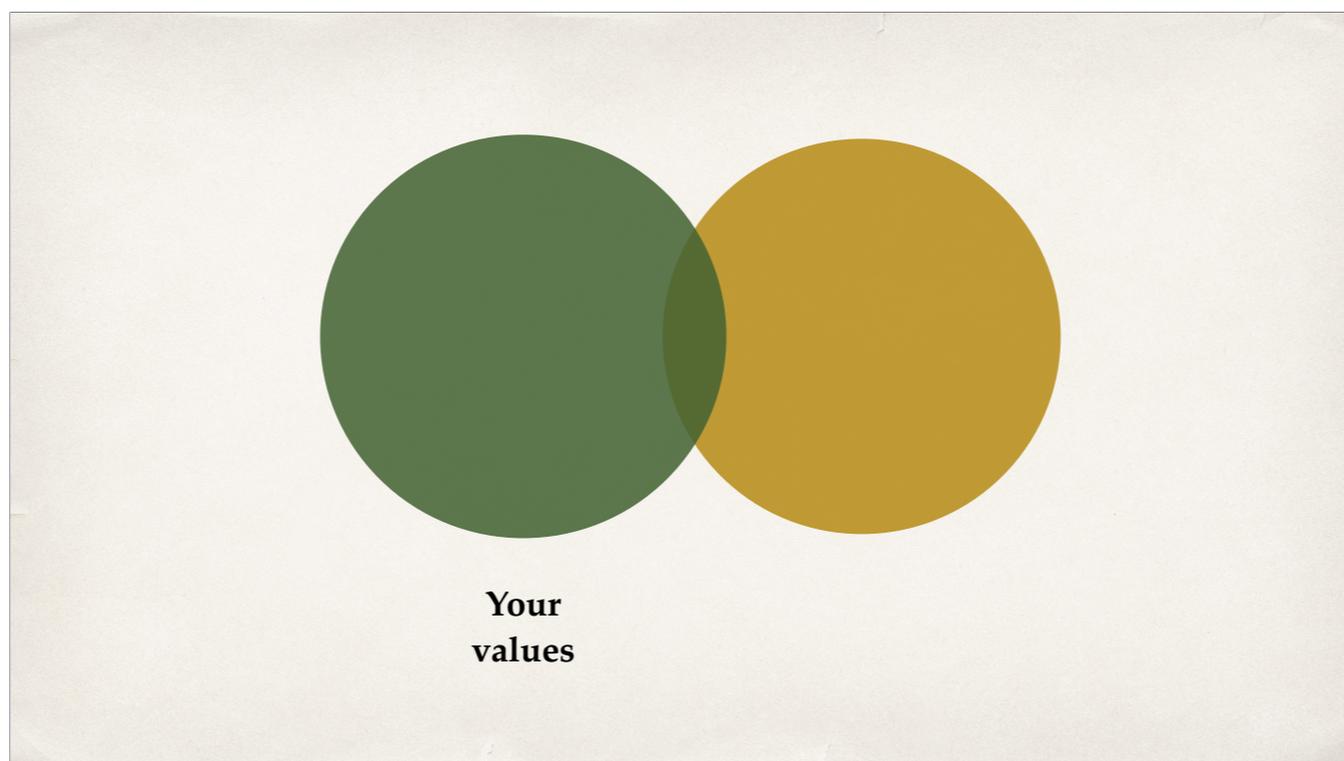
As I end my talk, I want to leave you with this question: Are your values aligned with the solutions our technologies provide?

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 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"

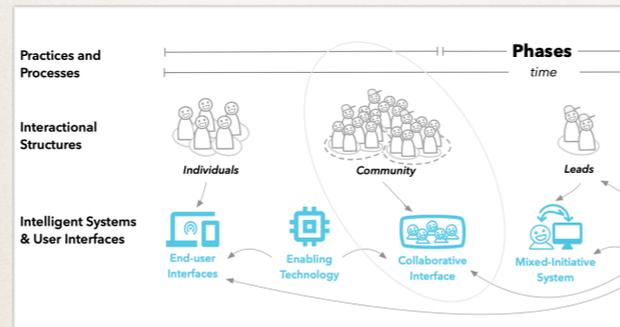
But I think we'd all lose something if we settle for that.



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 all our human and machine abilities to bear, and we will need to think about the larger ecology we are creating, and in which we are embedded.

In doing all this, **my hope is that we will bridge the ways that we have embraced 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

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 Hibschan, 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 Hibschan, Haoqi Zhang, and Eleanor O'Rourke. Ply: A visual web inspector for learning from professional webpages. UIST 2018.
- ✦ Josh Hibschan. 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. Orchestration Scripts: A System for Encoding an Organization's Ways of Working to Support Situated Work. CHI 2023.
- ✦ Kapil Garg, Darren Gergle, and Haoqi Zhang. Understanding the Practices and Challenges of Networked Orchestration in Research Communities of Practice. CSCW 2022.
- ✦ 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

Experiential Computing

- ✦ Jackie He, Shirley Zhang, 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]