Crowds, Computation, and the Future of Work

Michael Bernstein
Stanford University

ACHIEVING OUR COLLECTIVE GOALS

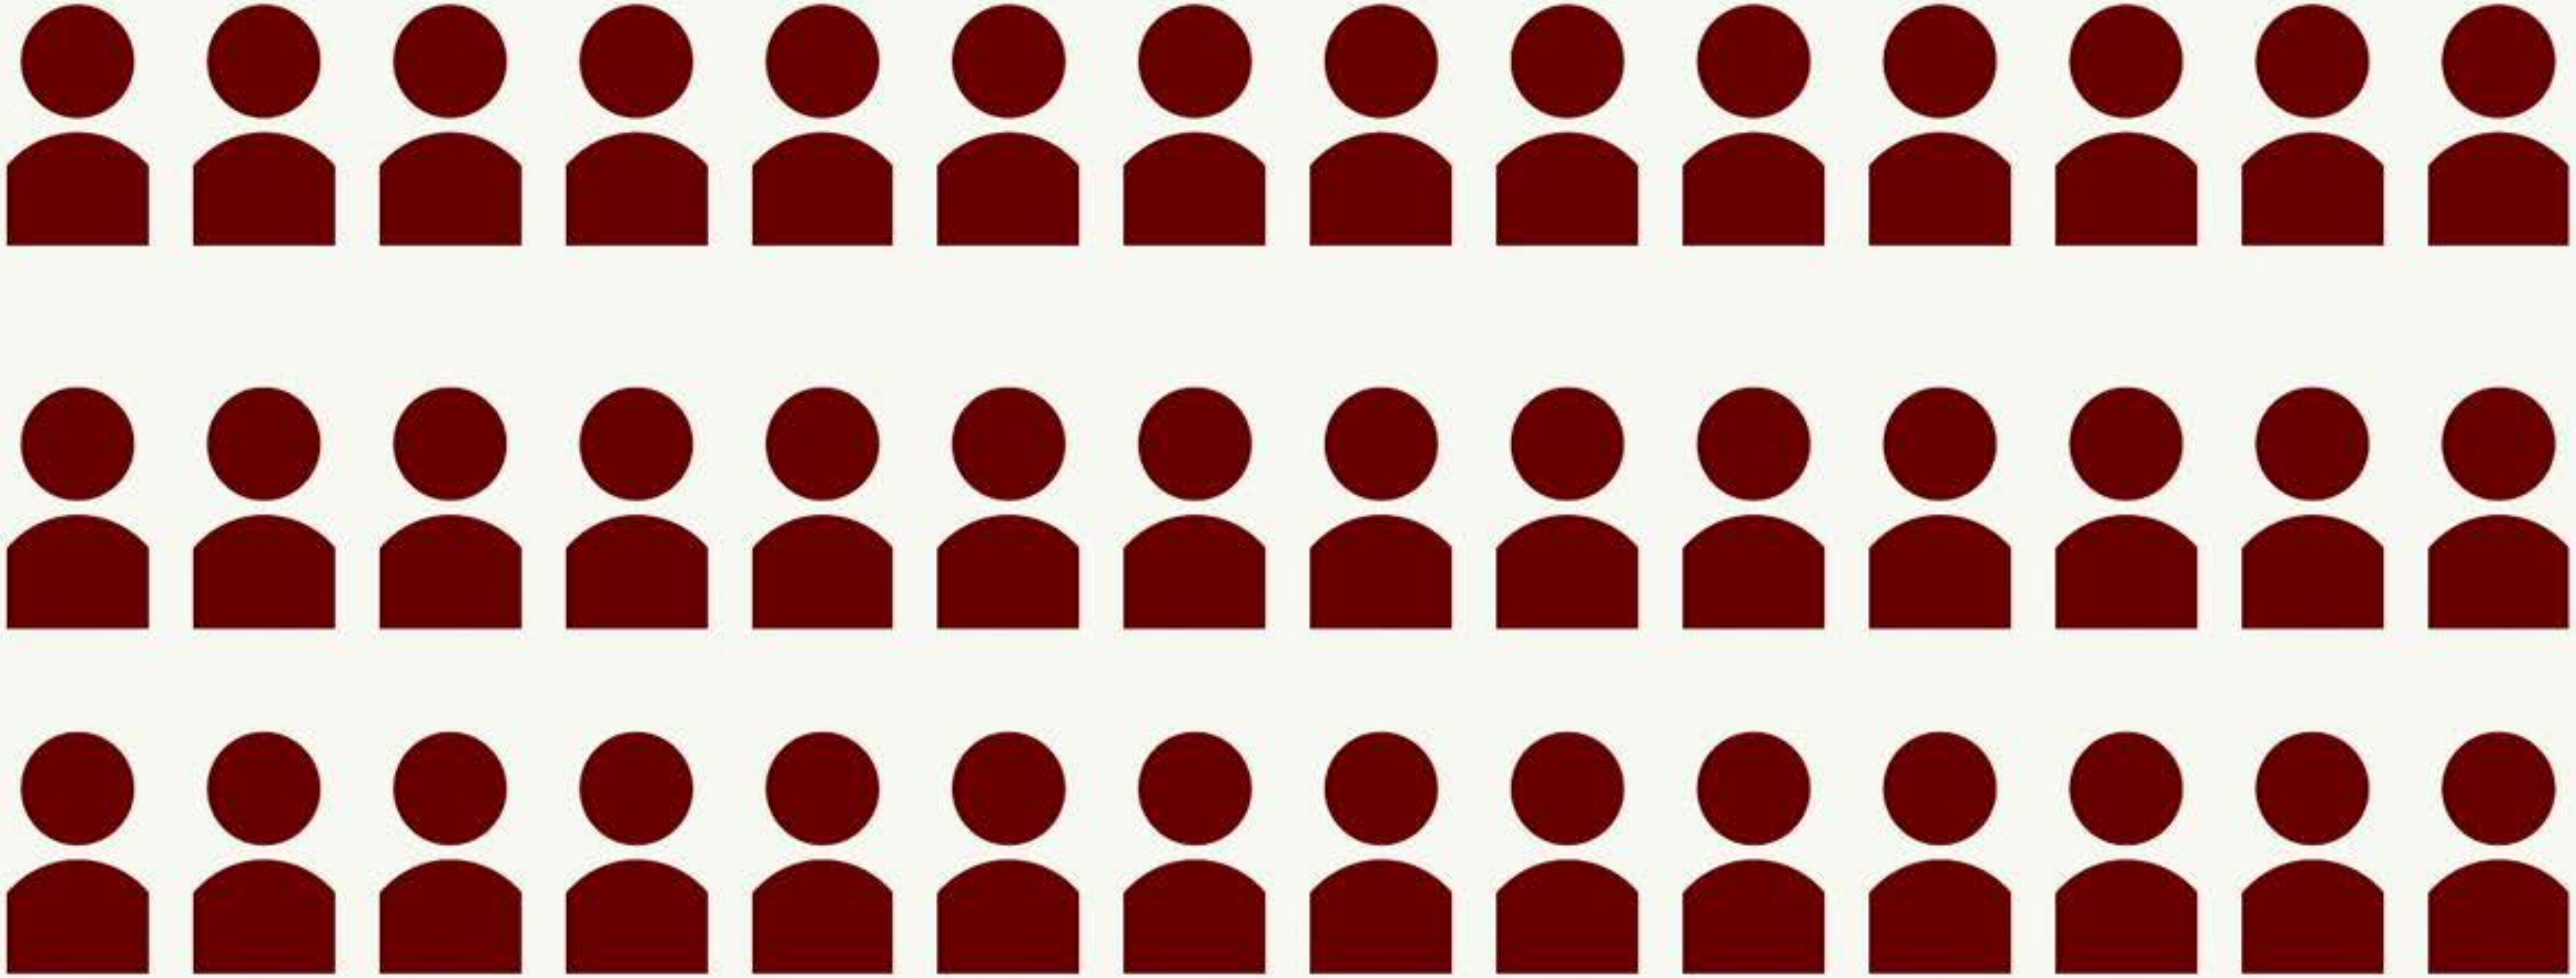

ACHIEVING OUR COLLECTIVE GOALS

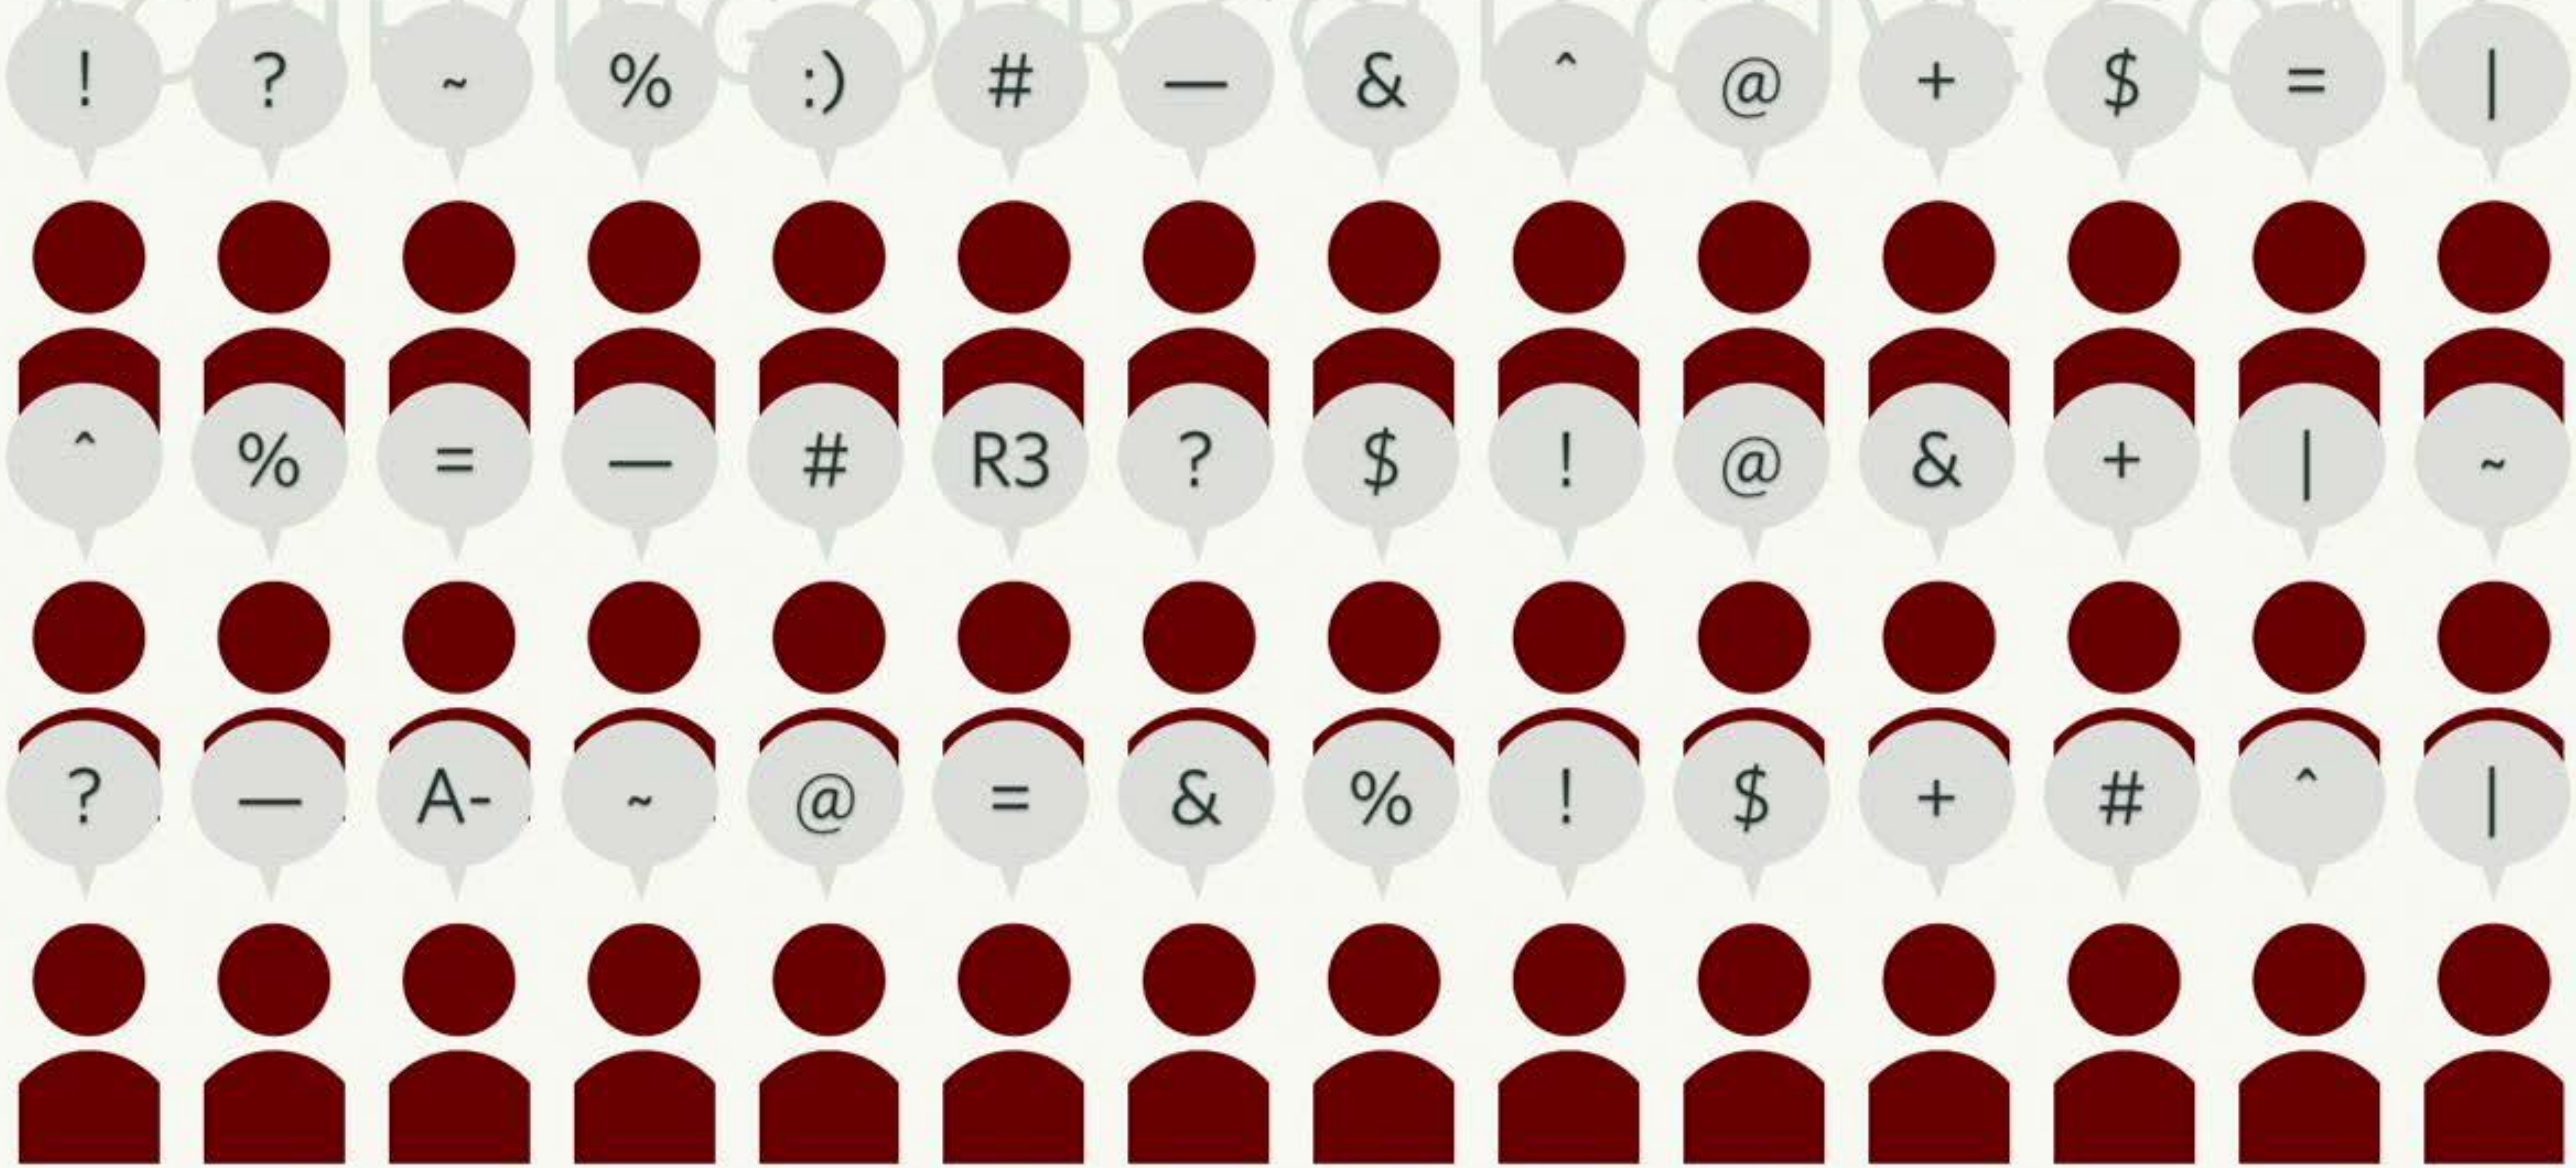

COORDINATION NEGLECT: HOW LAY THEORIES OF ORGANIZING COMPLICATE COORDINATION IN ORGANIZATIONS

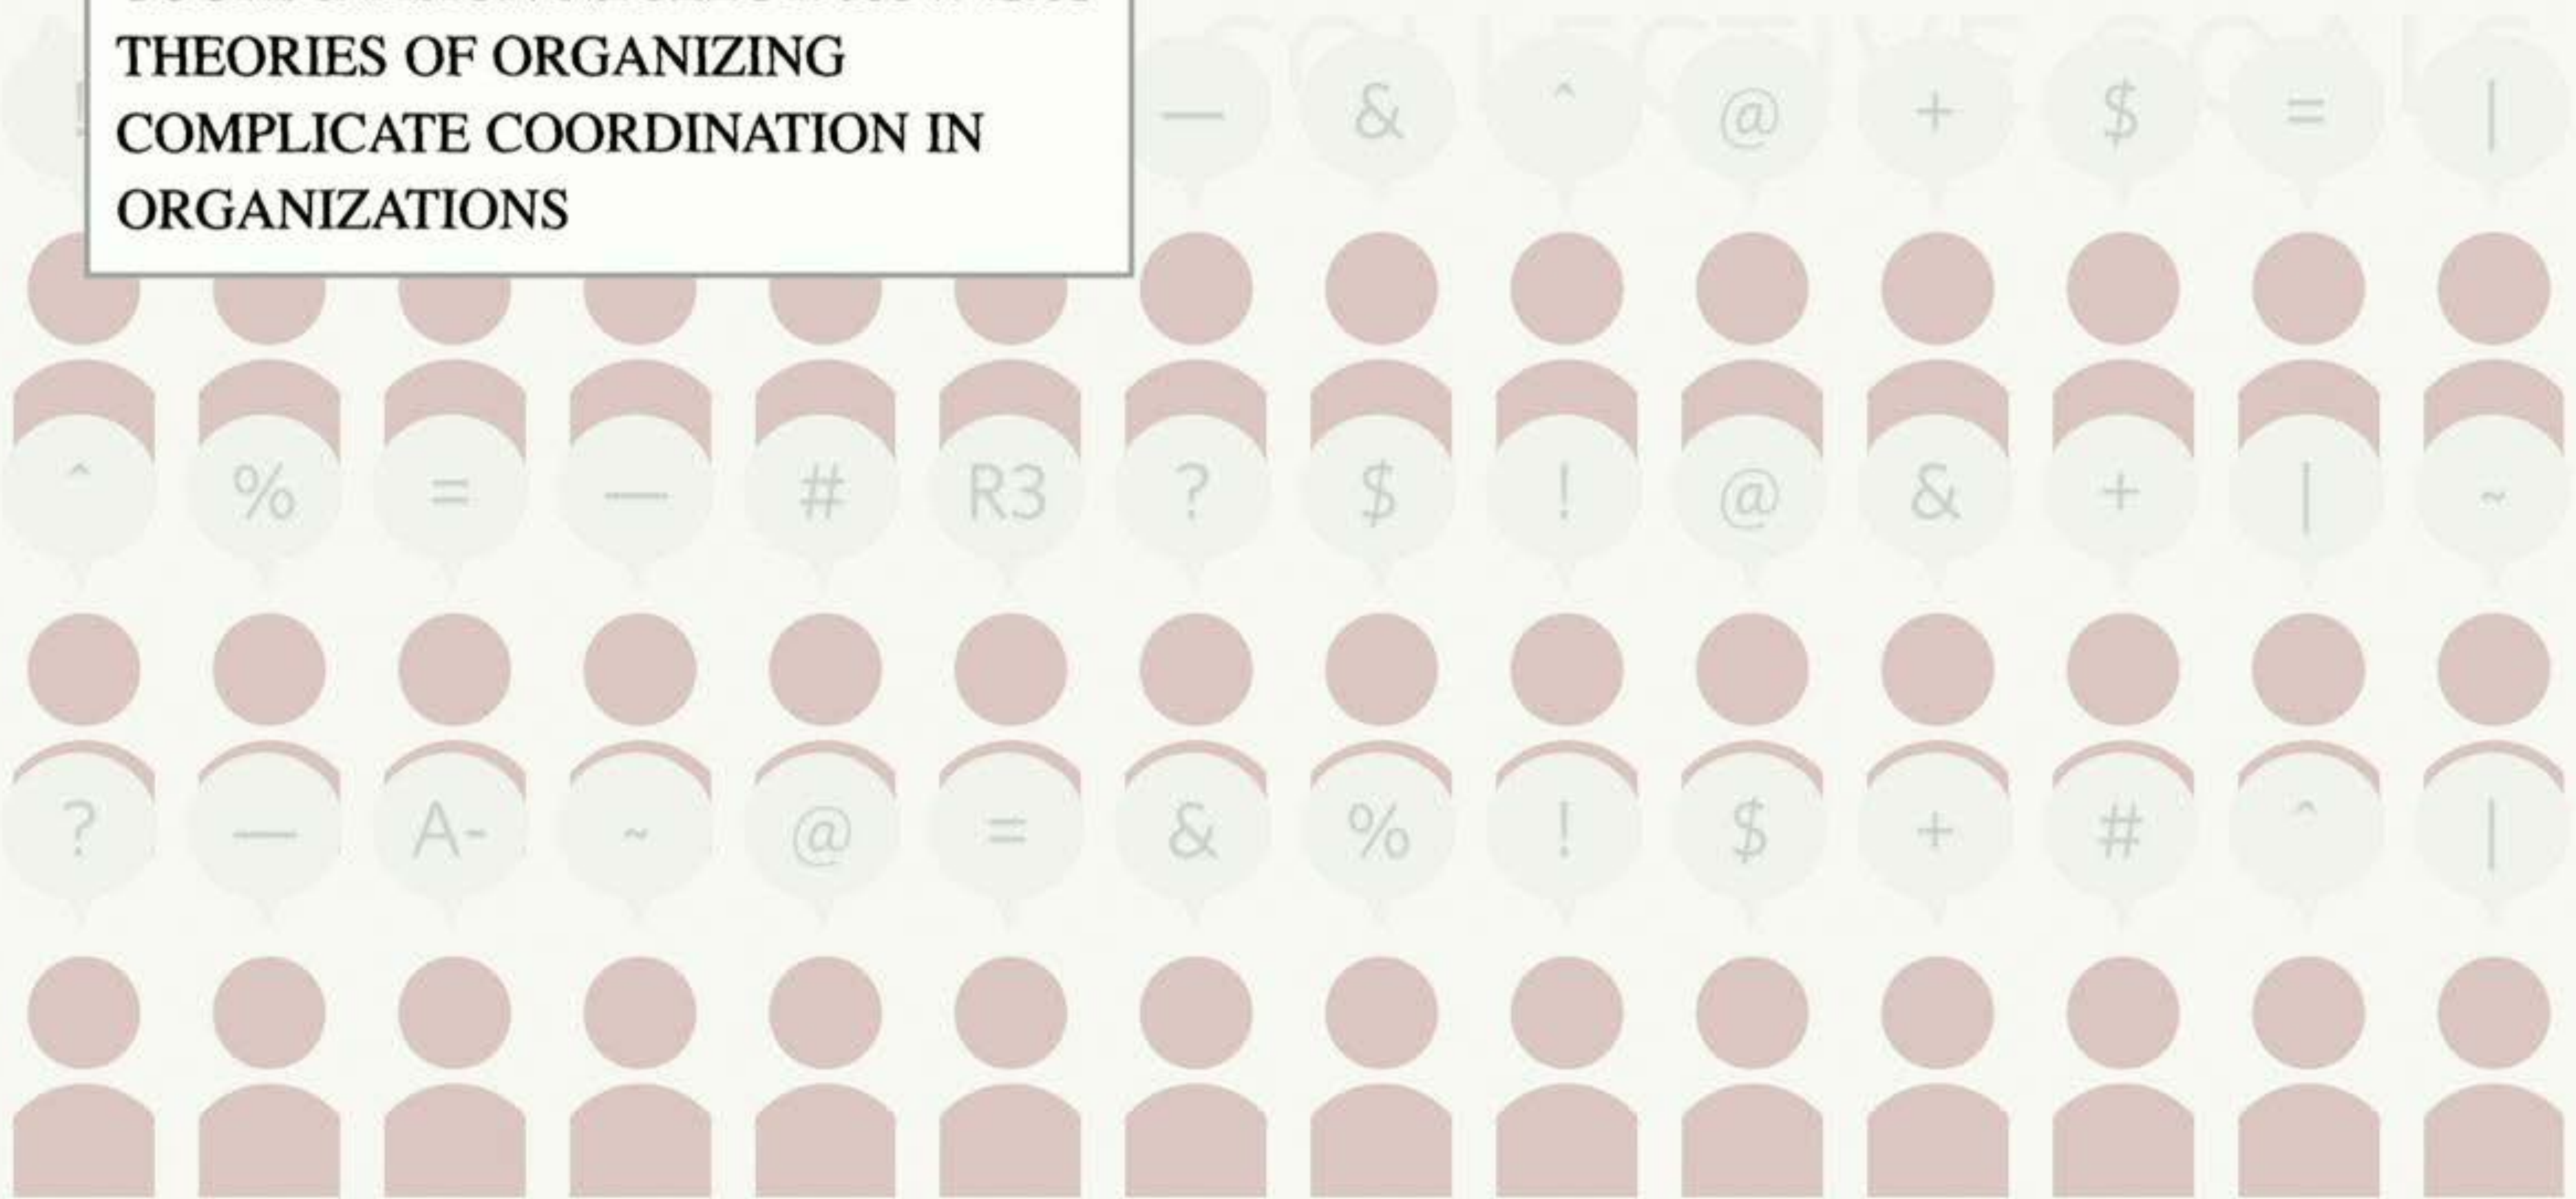

Out of Sight, Out of Sync: Understanding Conflict in Distributed Teams

COORDINATION NEGLECT: HOW LAY THEORIES OF ORGANIZING COMPLICATE COORDINATION IN ORGANIZATIONS

The Mutual Knowledge Problem and Its Consequences for Dispersed Collaboration

The team scaling fallacy: Underestimating the declining efficiency of larger teams

Who's in Charge Here? How Team Authority Structure Shapes Team Leadership

Team Familiarity, Role Experience, and Performance: Evidence from Indian Software Services

The Influence of Shared Mental Models on Team Process and Performance

Some unintended consequences of job design

Structure and Learning in Self-Managed Teams: Why "Bureaucratic" Teams Can Be Better Learners

HOW MIGHT COMPUTING AUGMENT US IN ACHIEVING OUR COLLECTIVE GOALS?

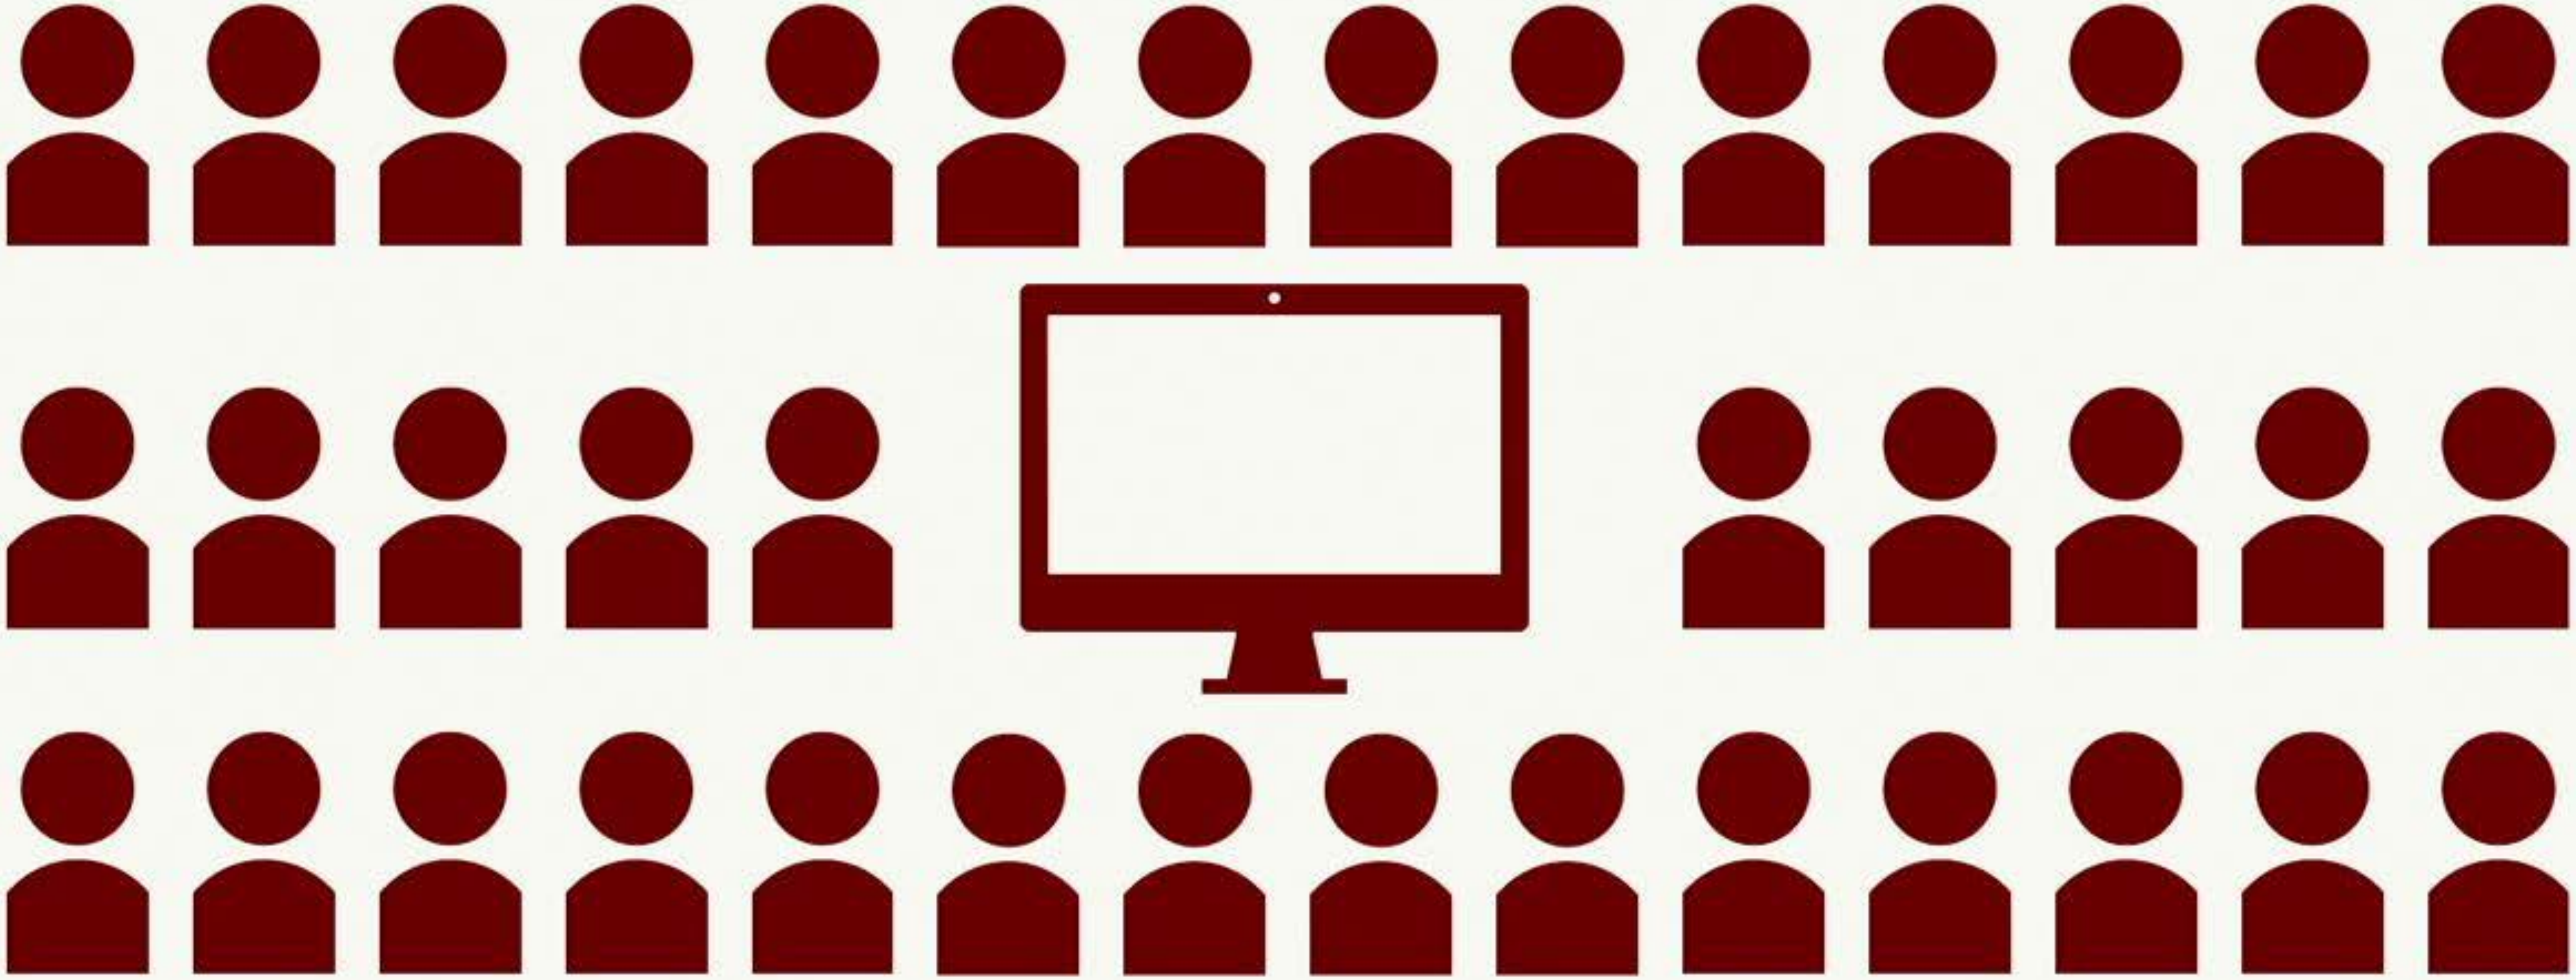

WORKER COLLECTIVE ACTION

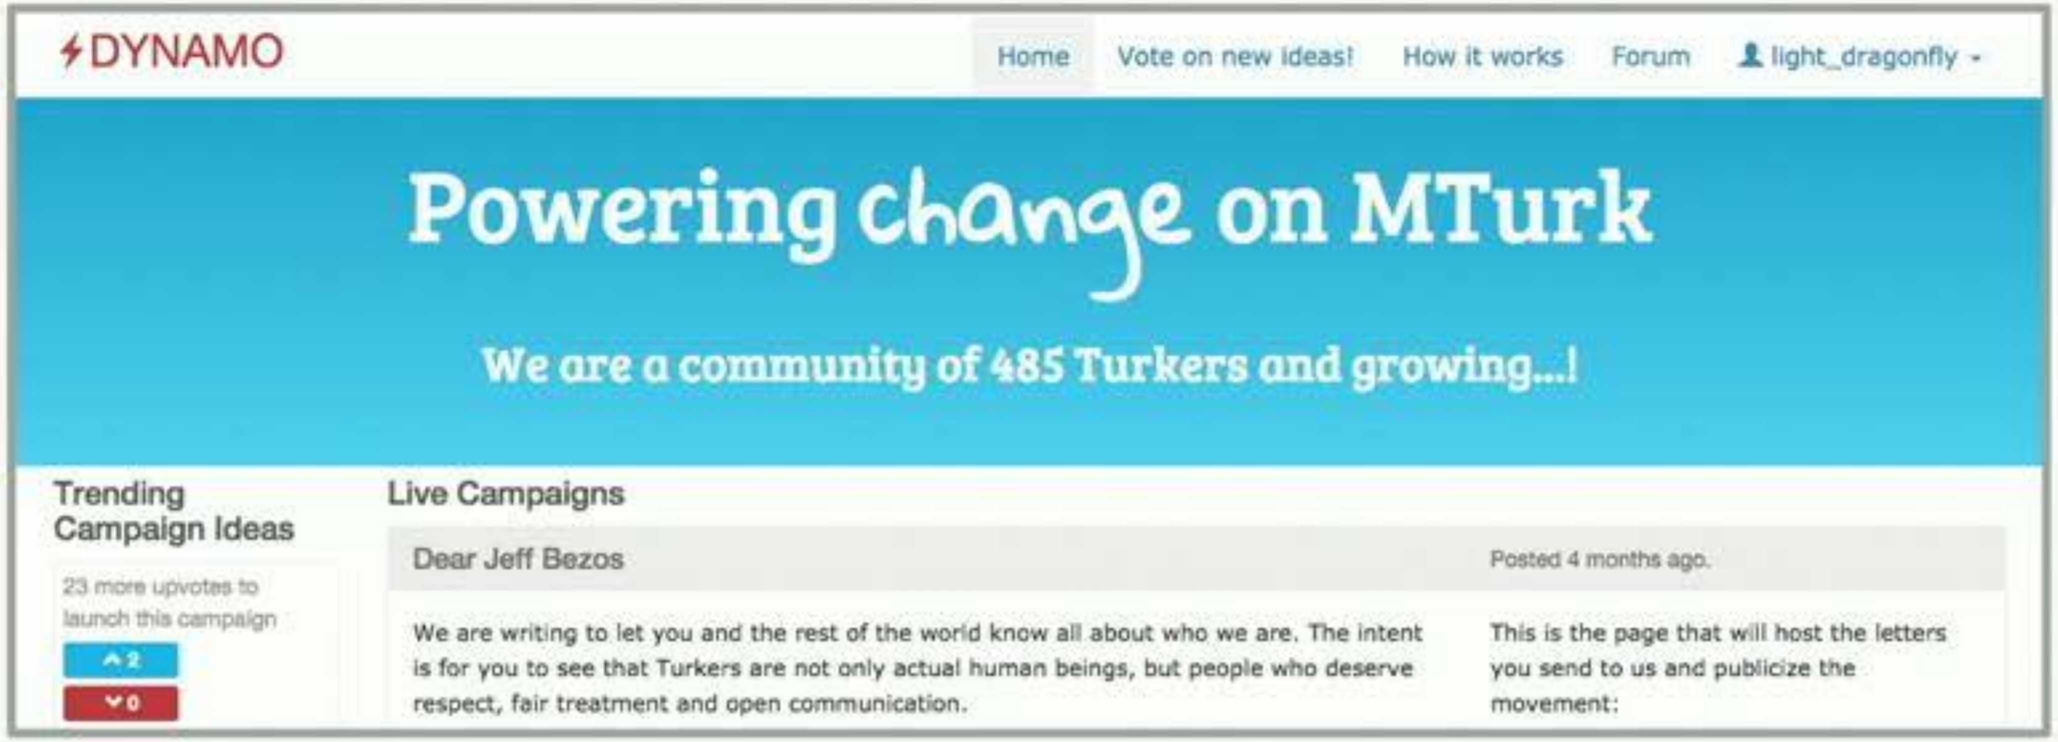


theguardian
Winner of the Pulitzer prize

Amazon.com

Amazon's Mechanical Turk workers protest: 'I am a human being, not an algorithm'

A Christmas email campaign is asking Amazon's CEO Jeff Bezos to improve terms for workers providing cheap digital labour

Socio-technical infrastructure for collective action amongst crowd workers

[Salehi et al. 2015]

FUTURE OF CROWD WORK

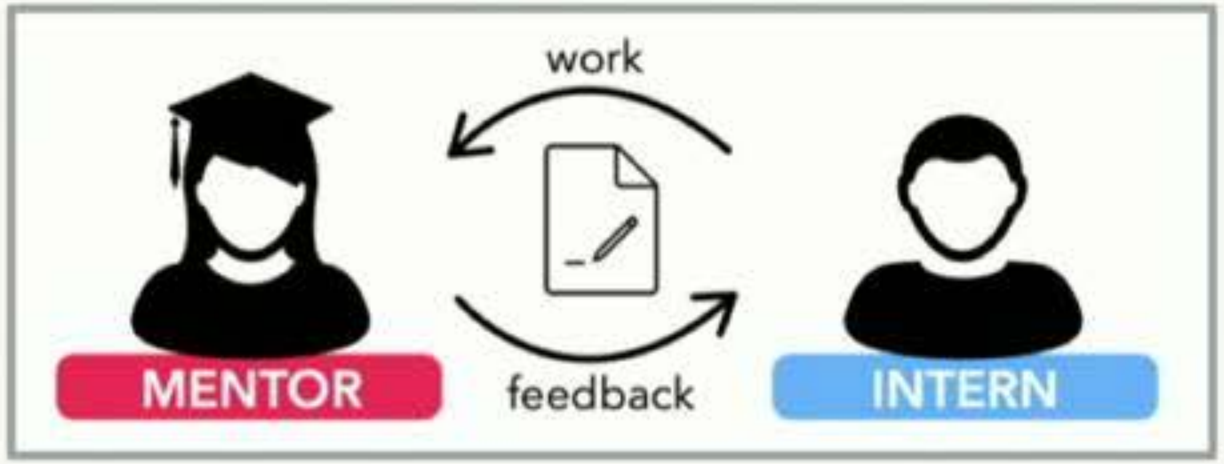

Micro-internships

[Suzuki et al. 2016]

Guild-style collective accreditation

[Whiting et al. 2017]

DOMINANT ARCHITECTURE: ALGORITHMS

Modularize and pre-define all possible behaviors into workflows

Computation decides which behaviors are taken, when, and by whom; optimizes, error-checks, and combines submissions

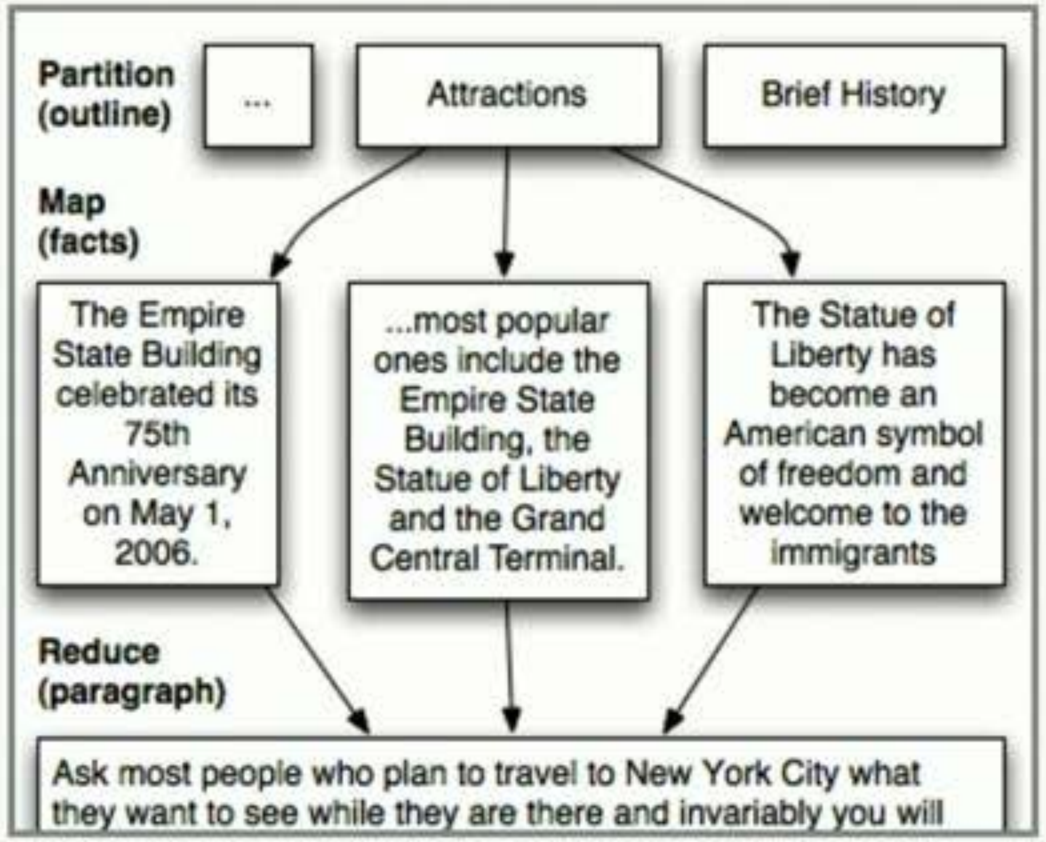

[Kittur 2011]

[Little 2010; Dai, Weld 2010]

LIMITS OF ALGORITHMIC COORDINATION

So far, goals such as invention, production, and engineering have remained largely out of reach [Kittur et al. 2013]

LIMITS OF ALGORITHMIC COORDINATION

So far, goals such as invention, production, and engineering have remained largely out of reach [Kittur et al. 2013]

Reason: open-ended, complex goals are **fundamentally incompatible** with a requirement to modularize and pre-define every behavior [Van de Ven, Delbecq, and Koenig 1976; Rittel and Weber 1973; Schön 1984]

LIMITS OF OPEN SOURCE AND OPEN INNOVATION

"**Peer production is limited** not by the total cost or complexity of a project, but **by its modularity**." [Benkler 2002]

"**With the Linux kernel** [...] **we want to have a system which is as modular as possible**. The open-source development model really requires this, because otherwise you can't easily have people working in parallel." [Torvalds 1999]

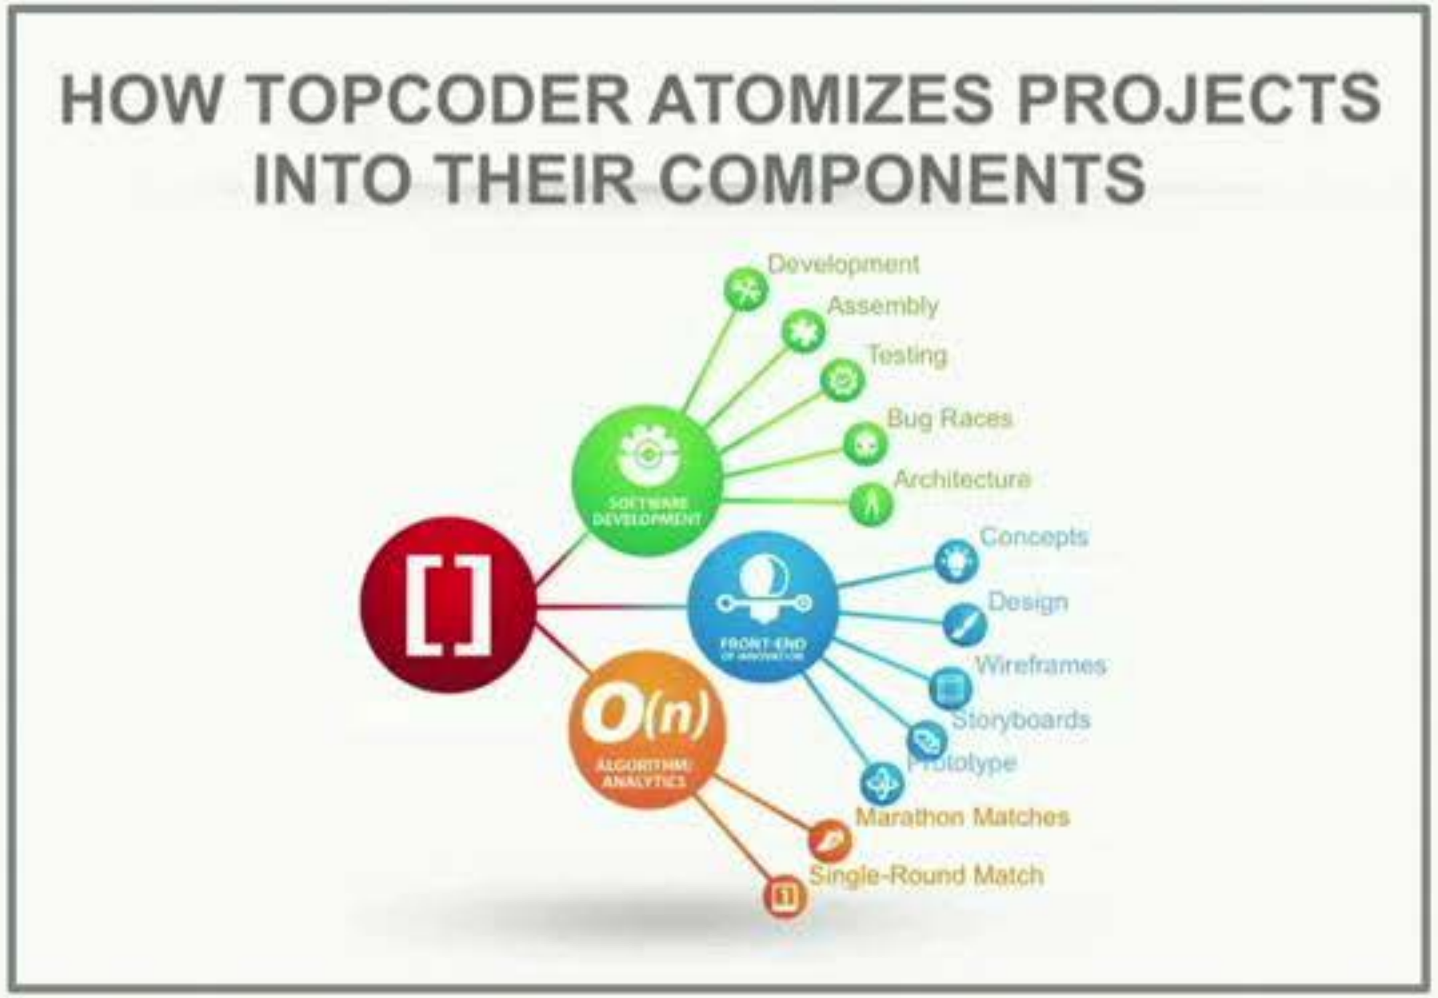

[Boudreau, Lacetera, and Lakhani 2011]

This architecture confines collaborations to goals so predictable that they can be entirely modularized and pre-defined

An alternative architecture: collaborations structured not as algorithms, but as computationally augmented organizations

1) Flash Organizations

Create on-demand organizations capable of complex work

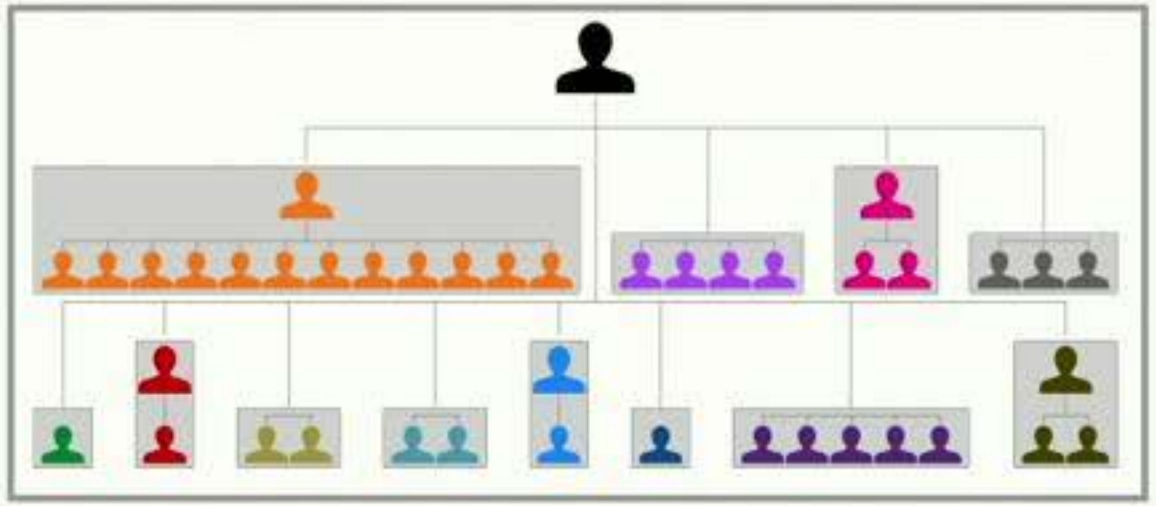

2) Dream Team

Find effective team structures

3) Crowd Research

Provide access to research opportunities to thousands around the world

Flash organizations

Valentine, Retelny, To, Rahmati, Doshi, Bernstein. CHI 2017.

Flash organizations: rapidly assembled and reconfigurable organizations composed of online collaborators

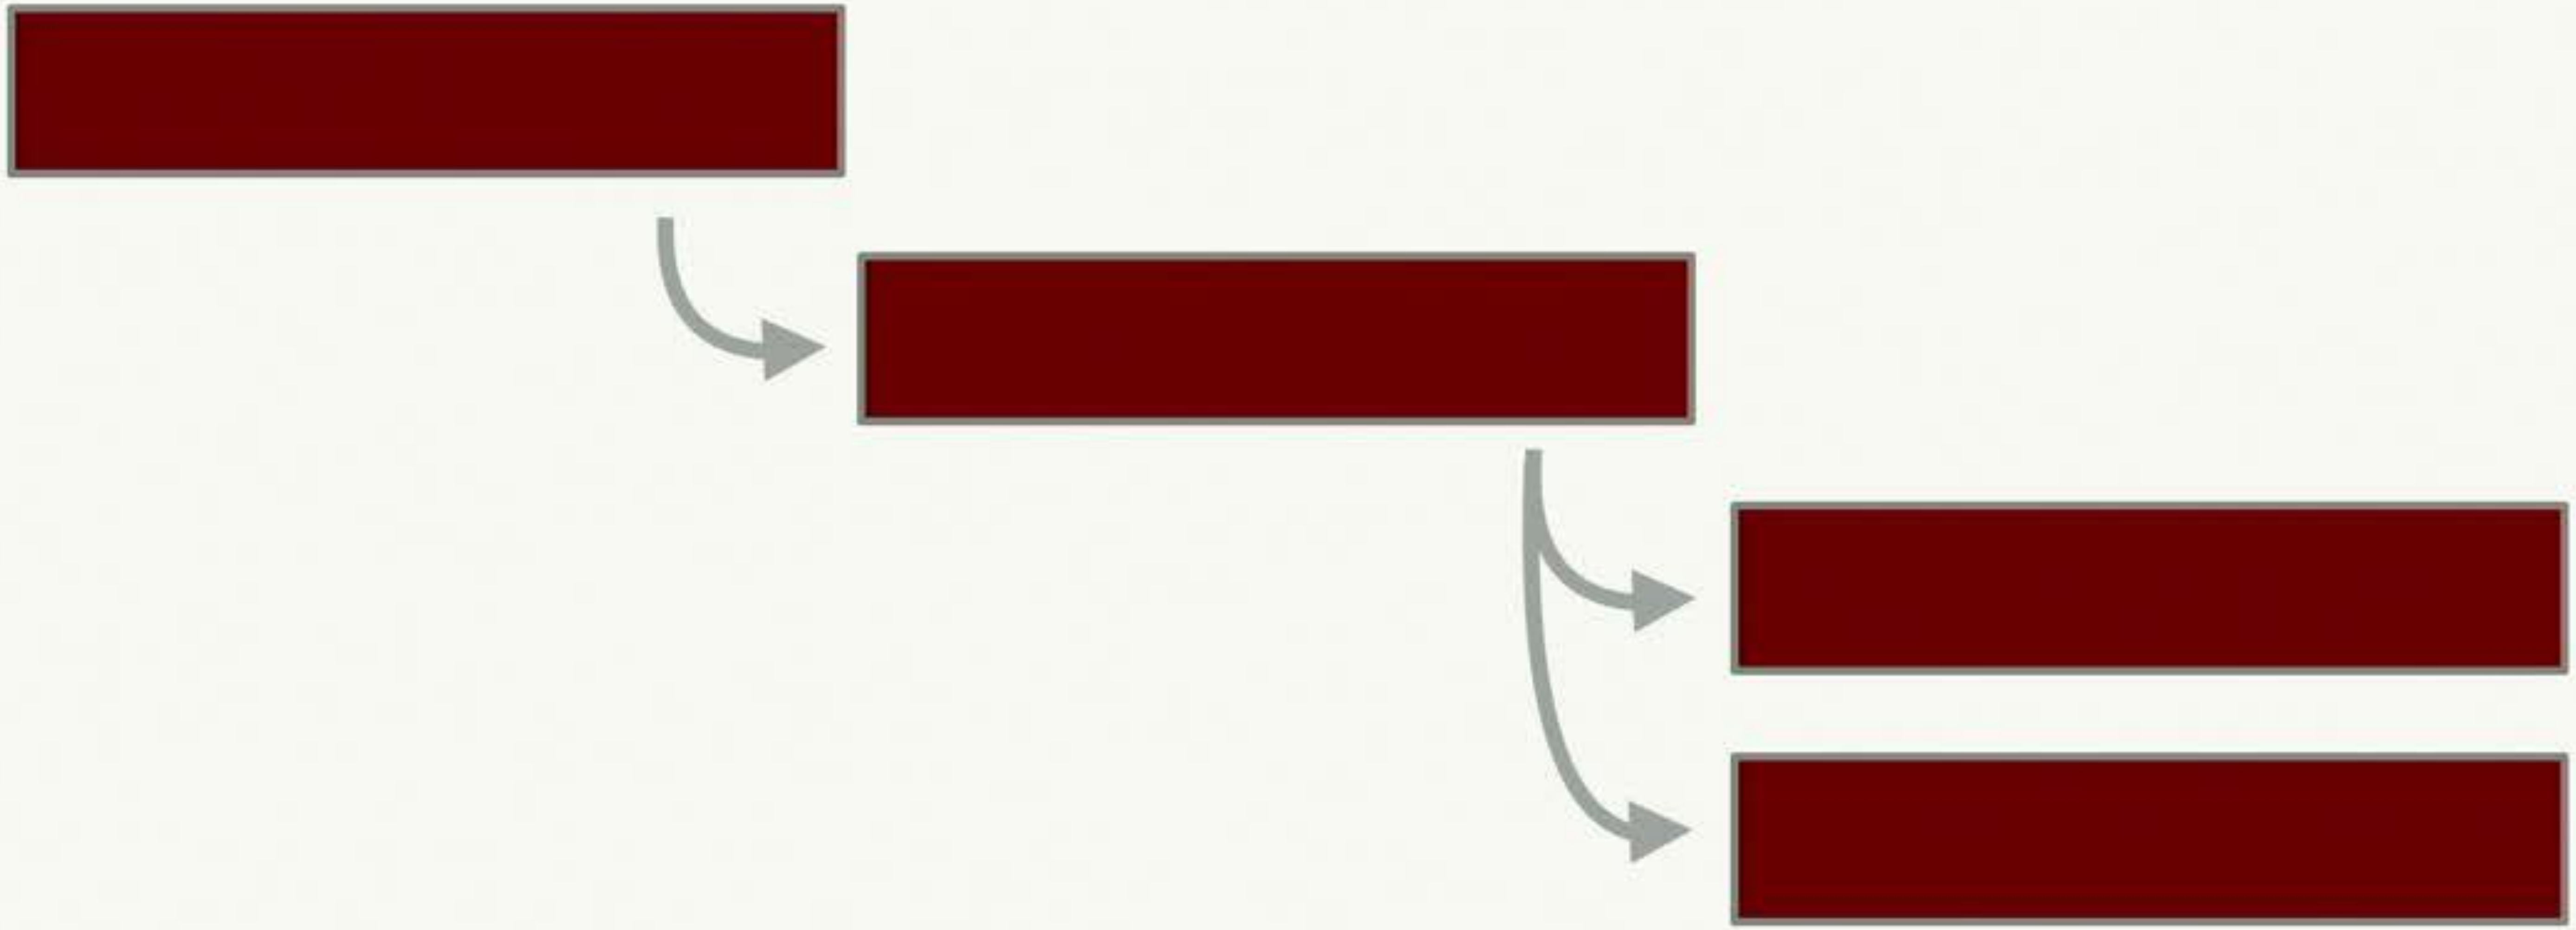

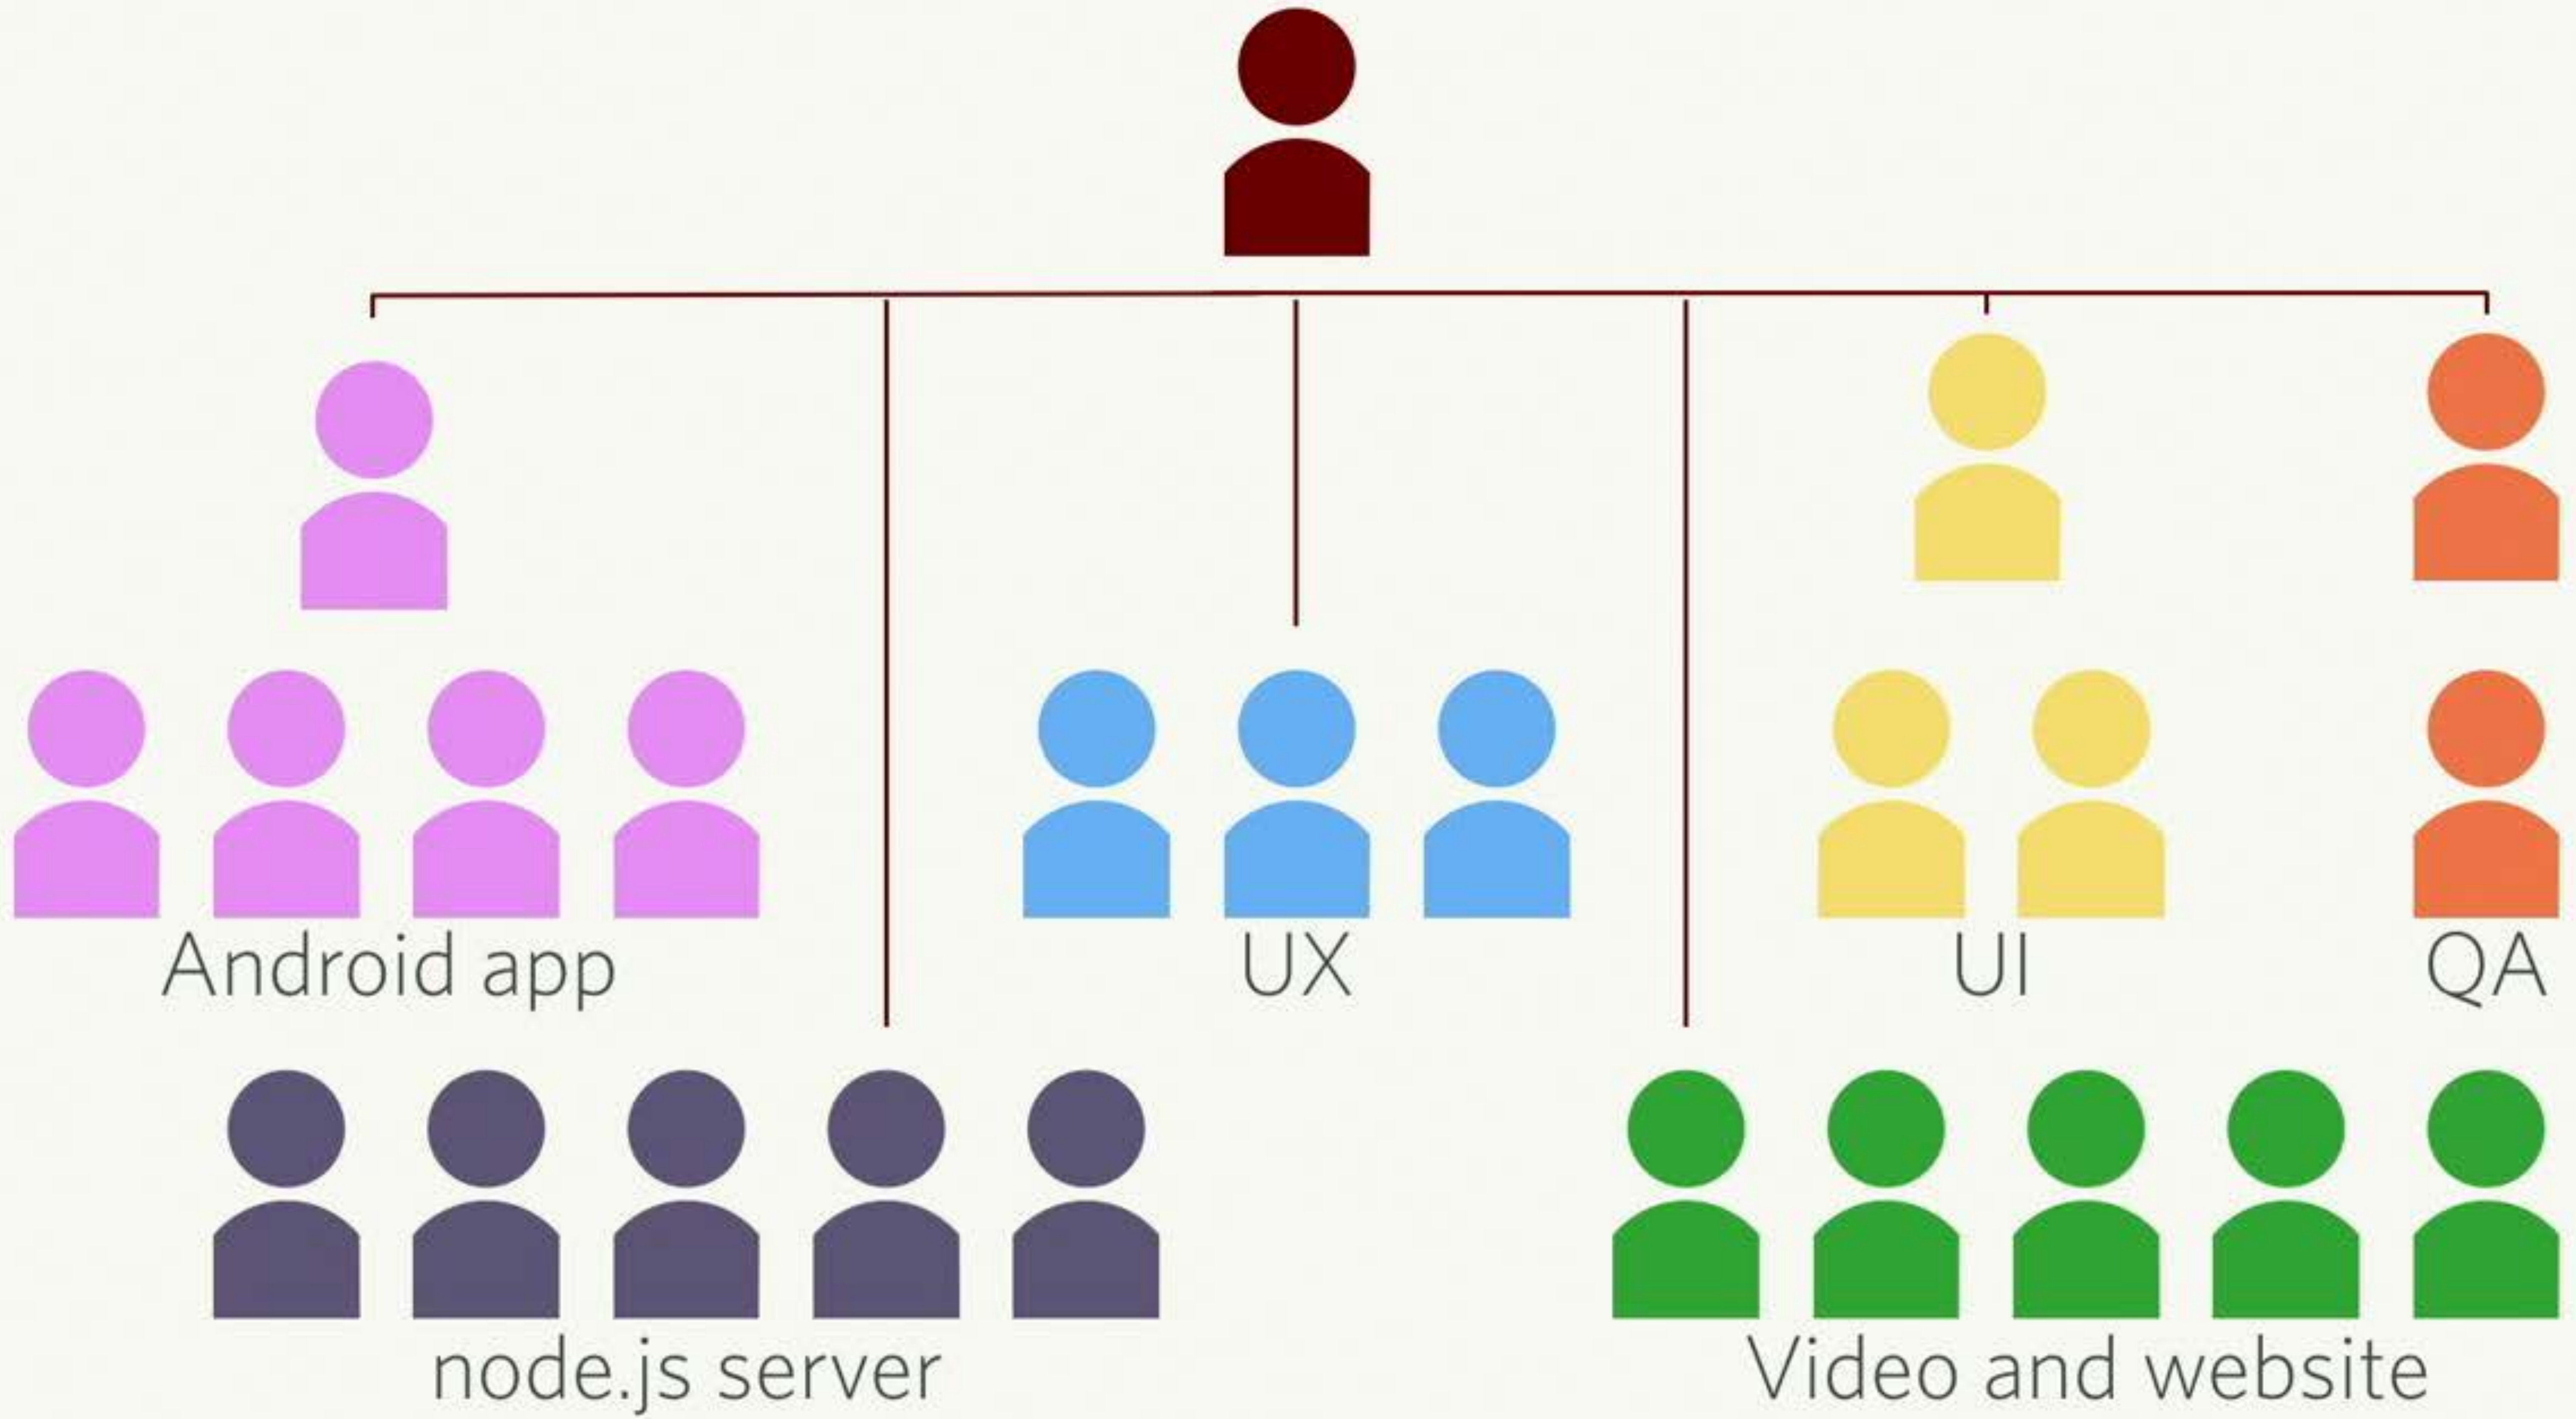
Android app
UX
UI
QA
node.js server
Video and website

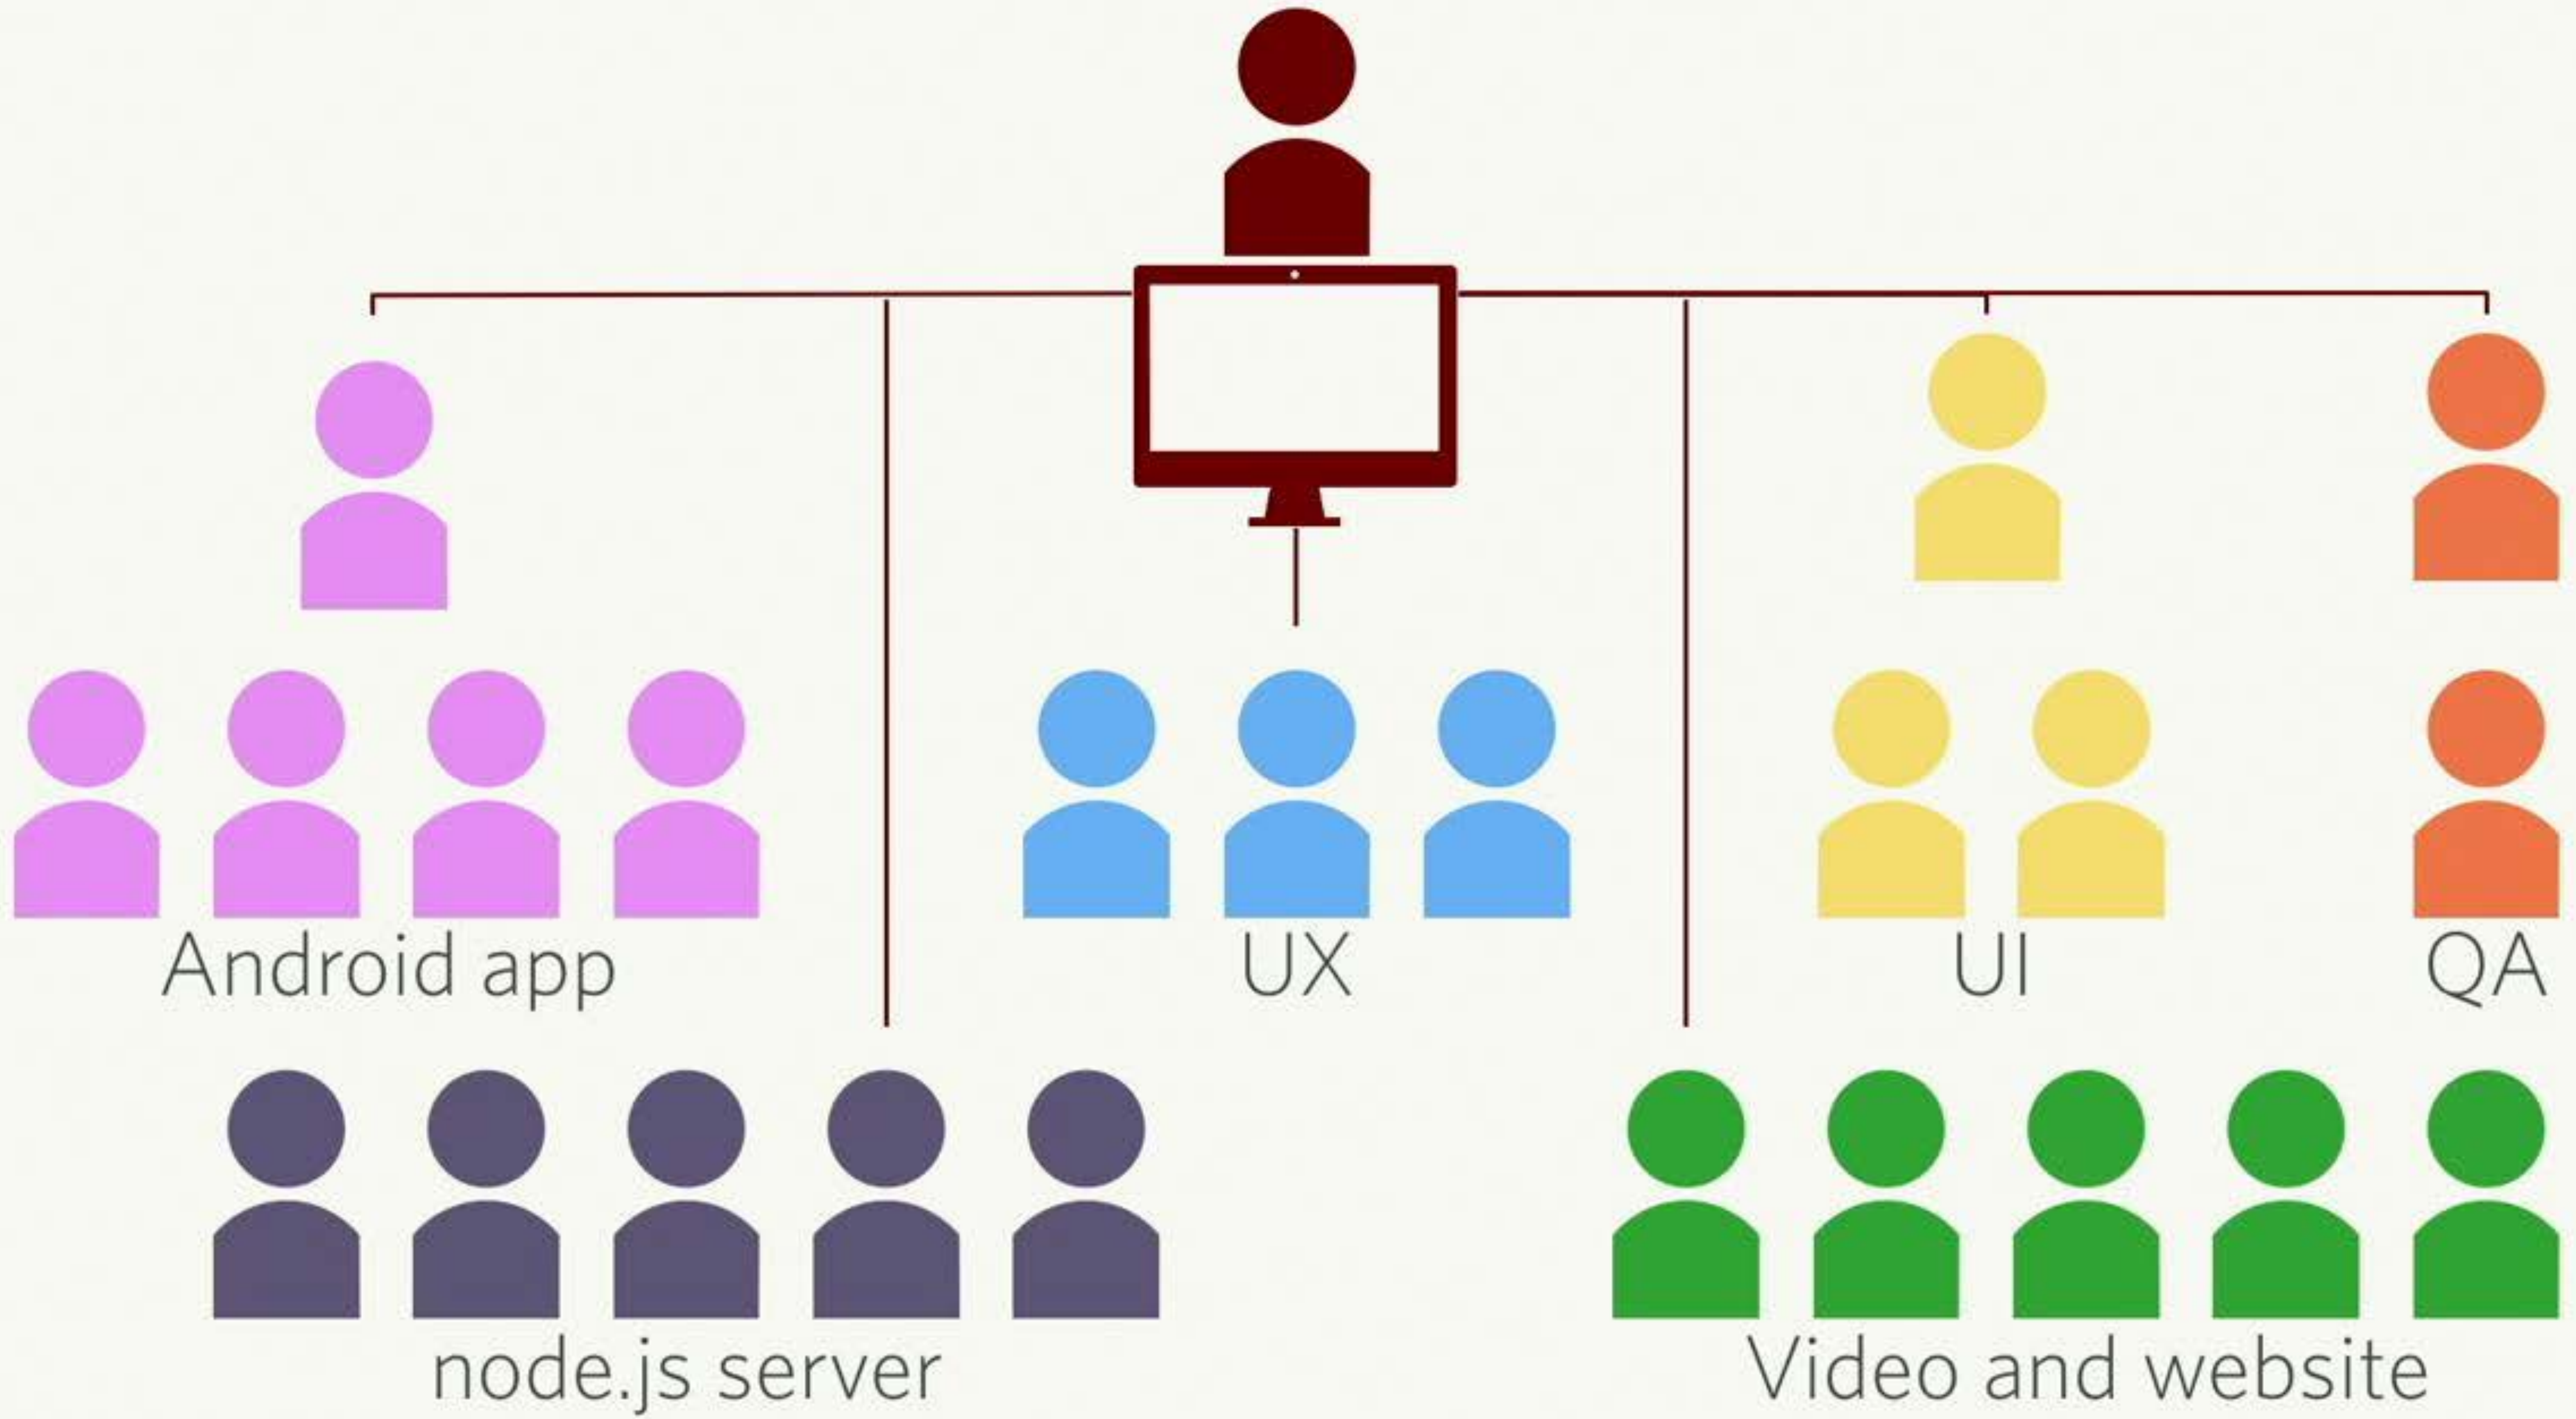
Android app
UX
UI
QA
node.js server
Video and website

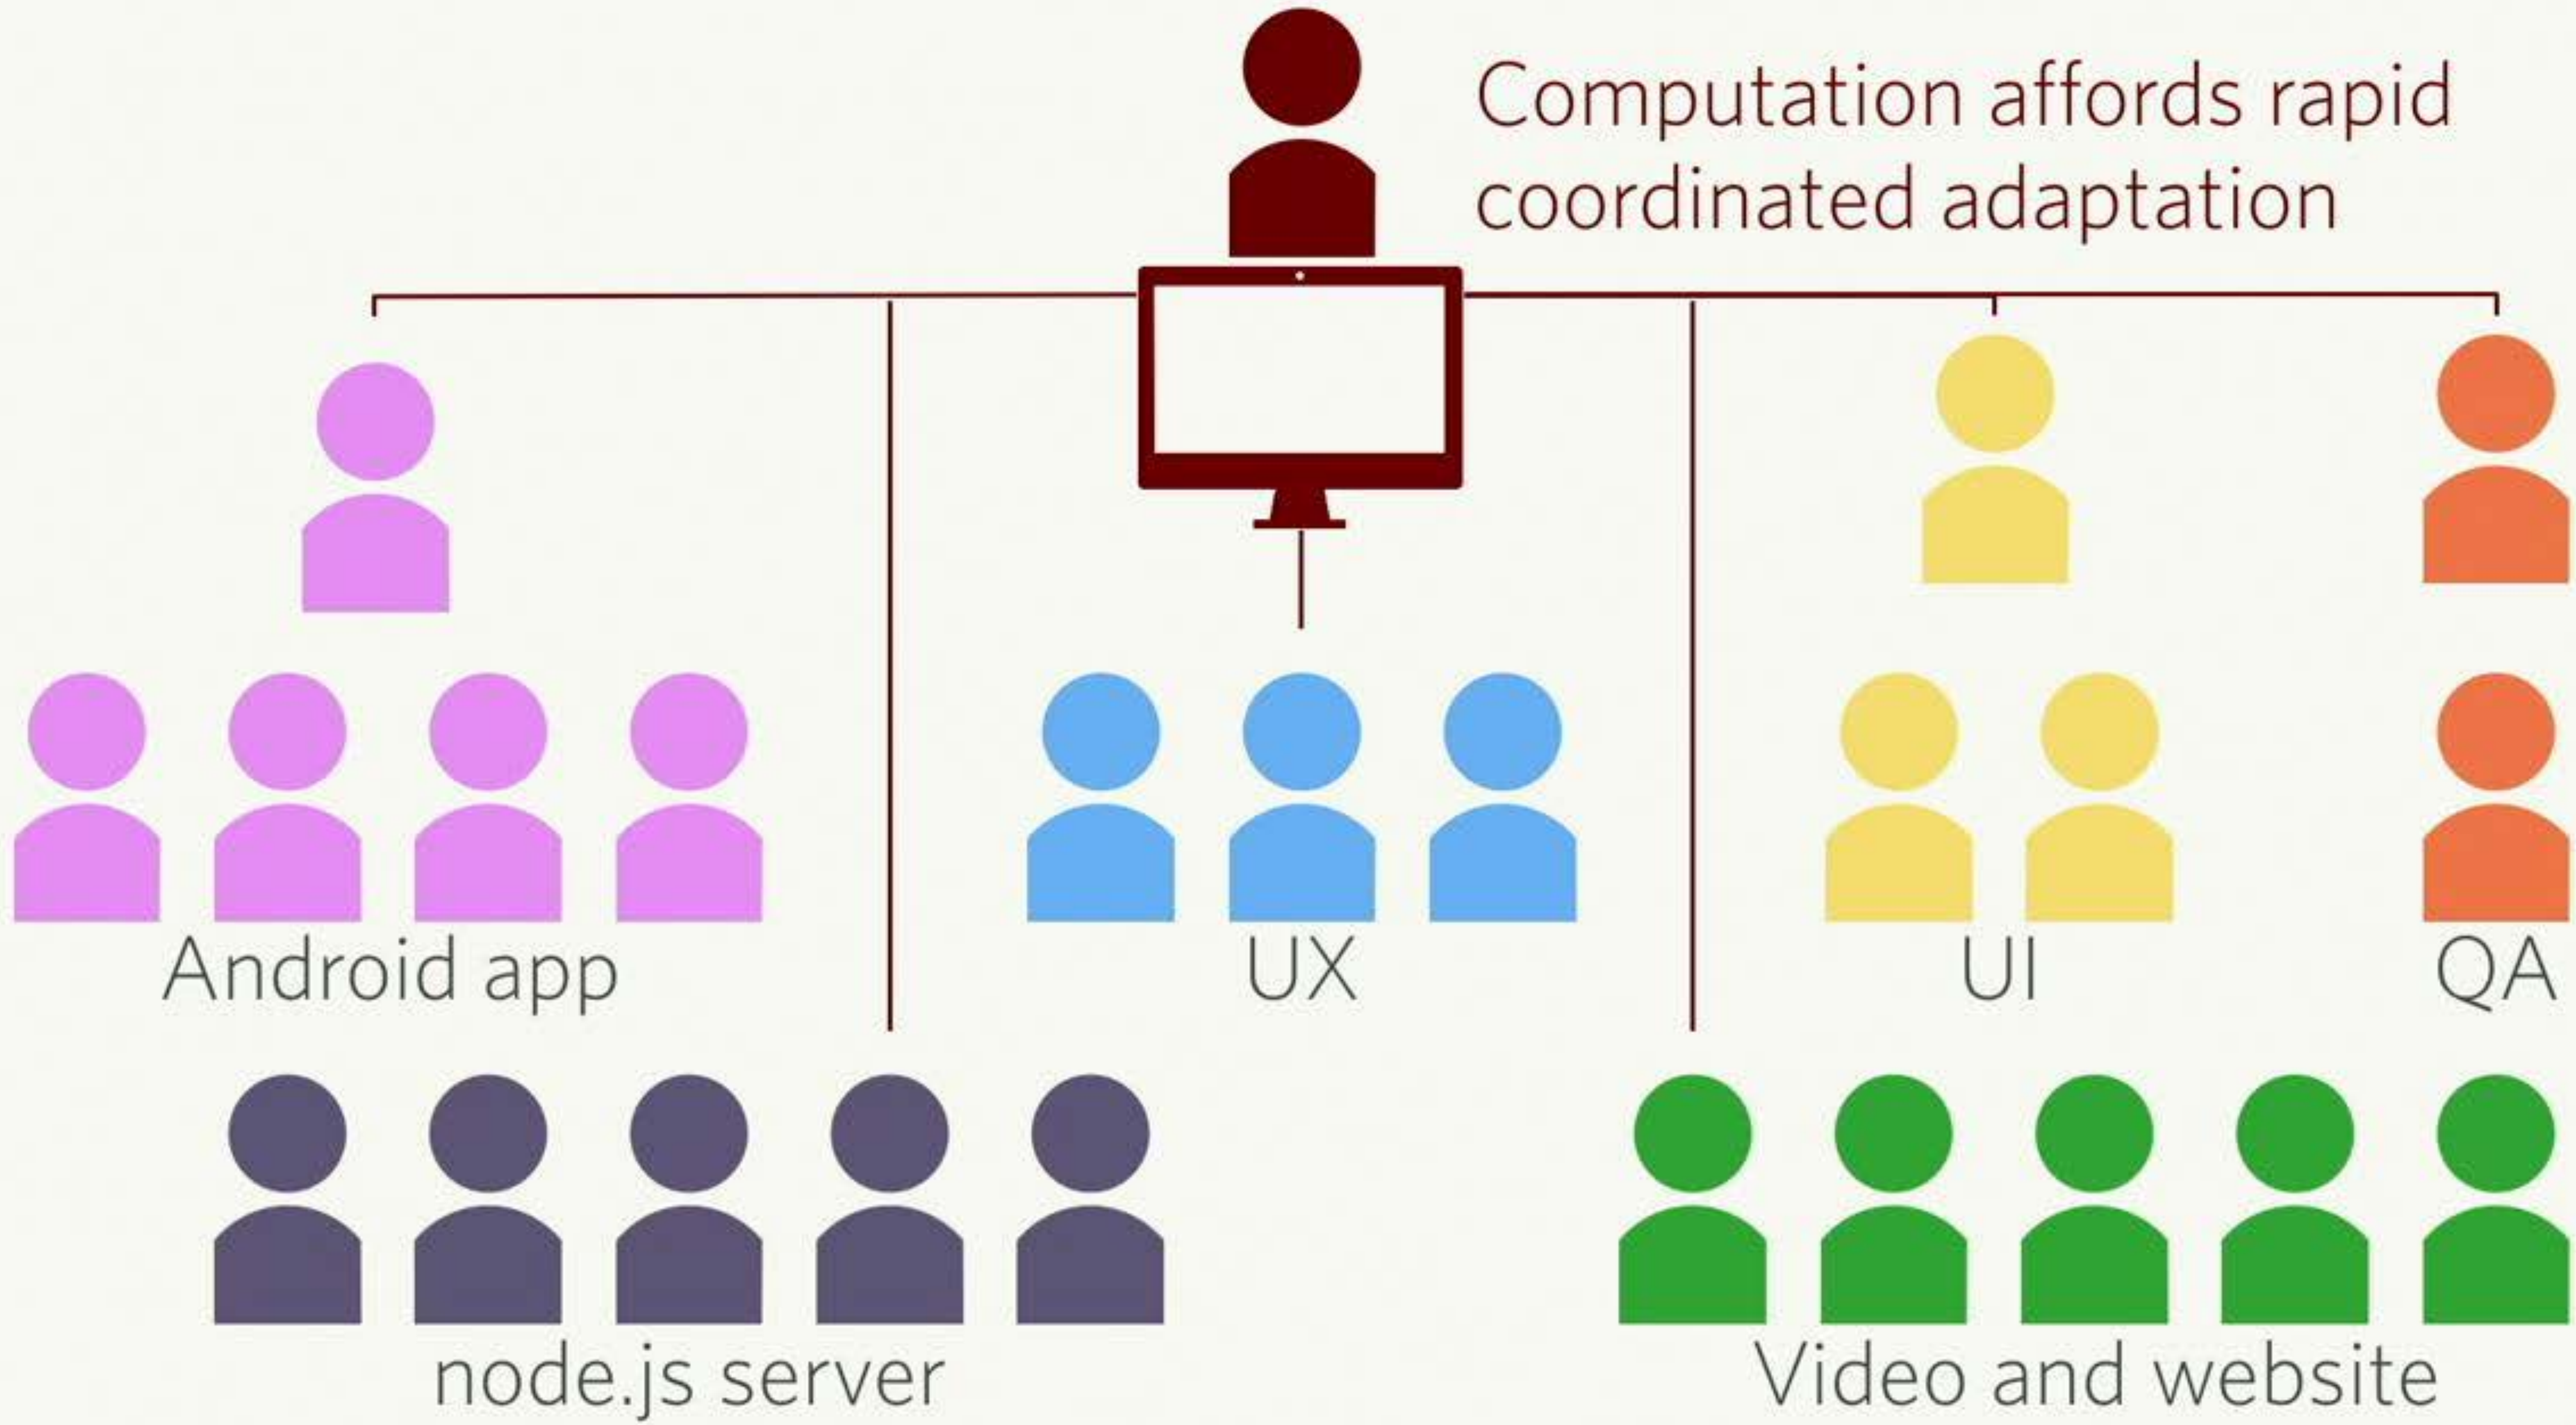
Computation affords rapid coordinated adaptation
Android app
UX
UI
QA
node.js server
Video and website

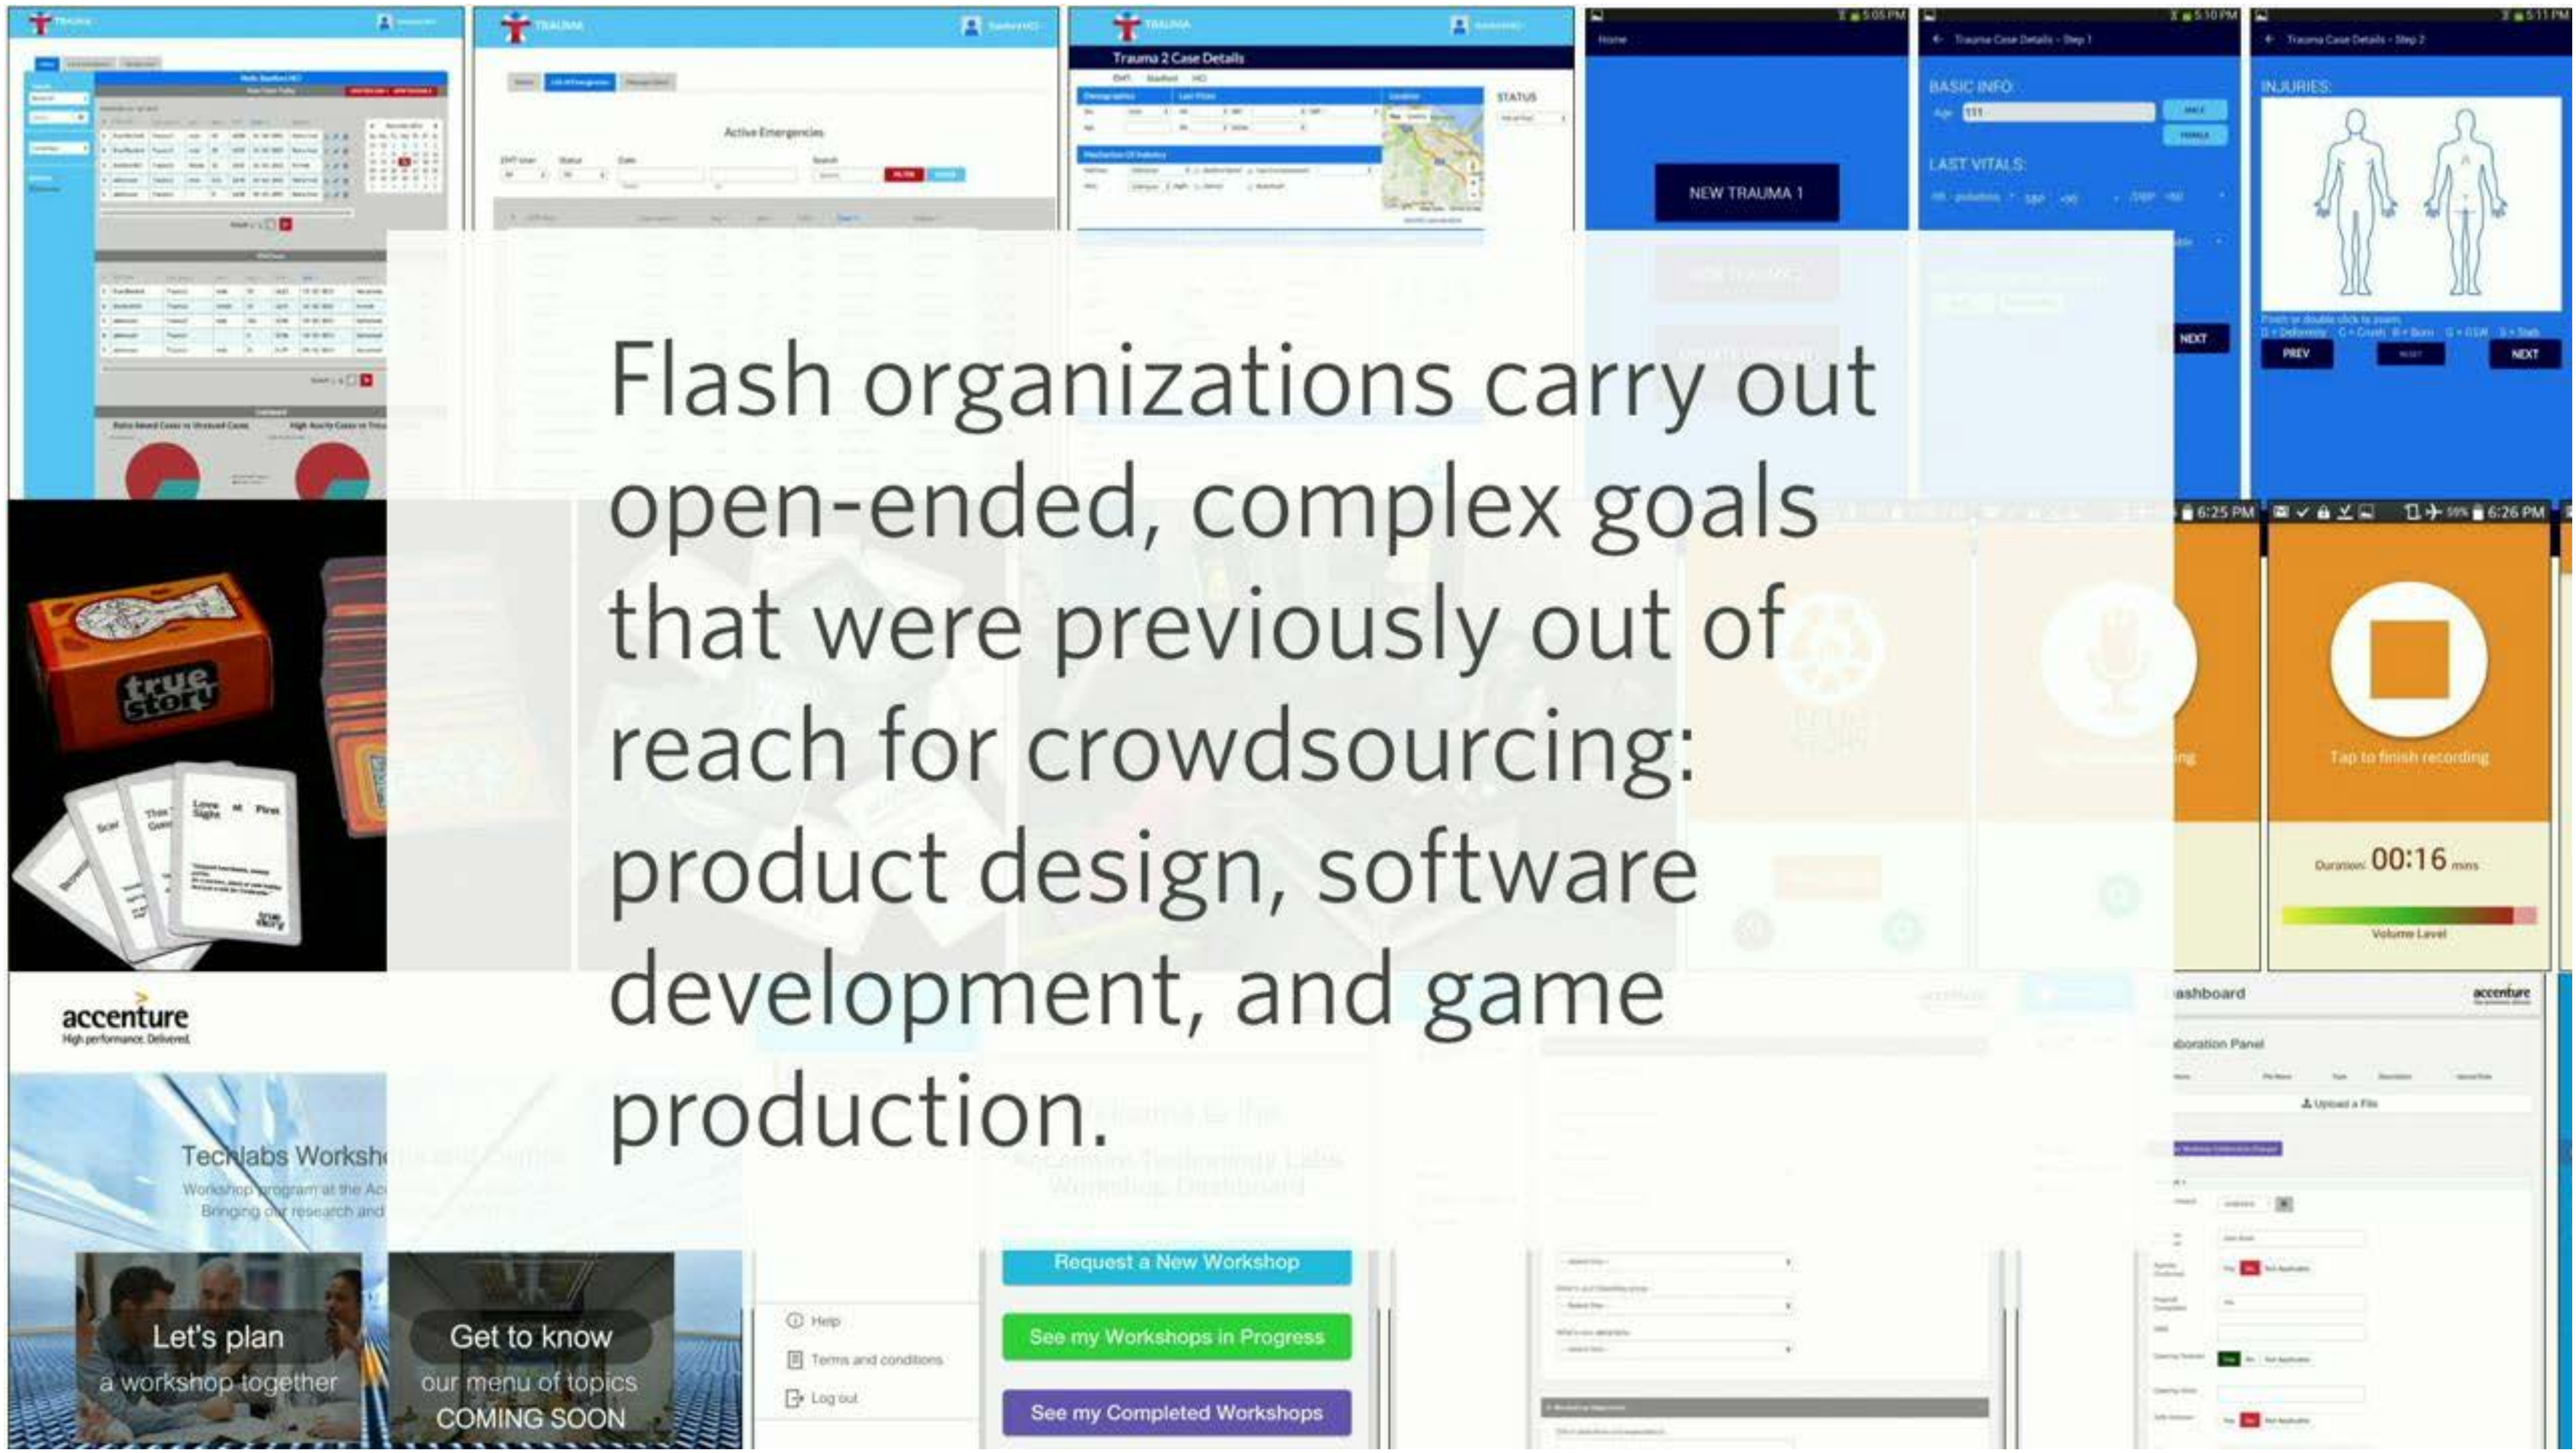

Flash organizations carry out open-ended, complex goals that were previously out of reach for crowdsourcing: product design, software development, and game production.

FOUNDRY

Web platform that supports authoring, reconfiguring, and running flash organizations

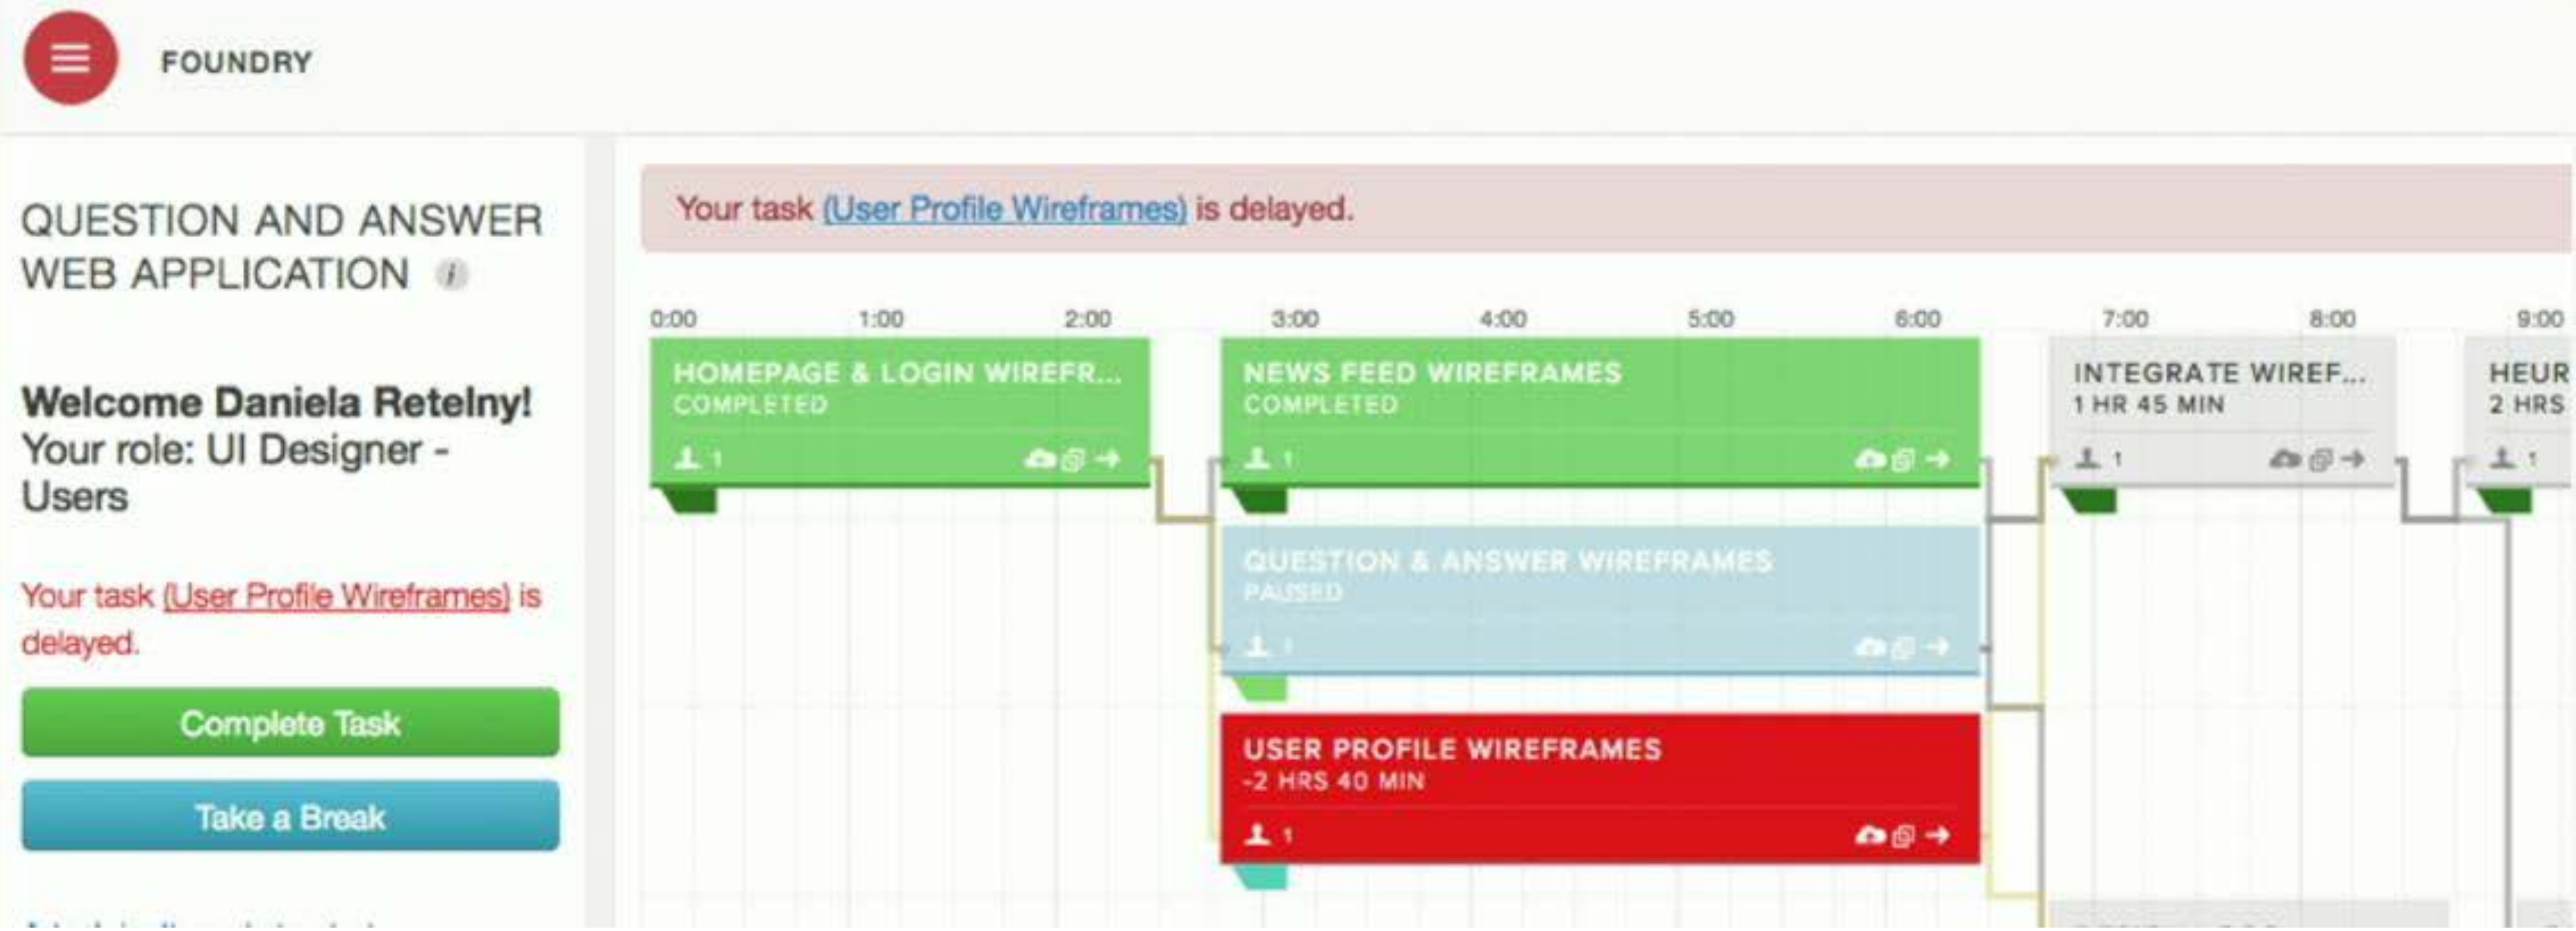

COORDINATION SANS ALGORITHMS

Inspiration: film crews and disaster response teams
[Bigley 2001; Bechky 2006; Klein et. al 2006; Valentine & Edmondson 2015]

Role structures enable interaction based on knowledge of roles rather than asset-specific knowledge of each other

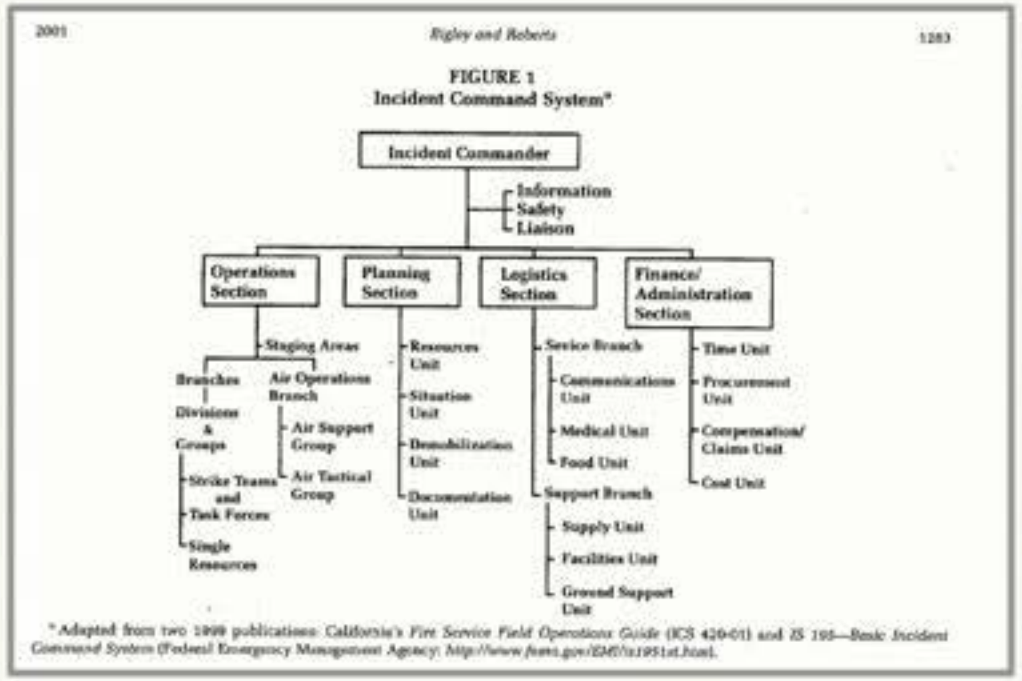

COMPUTATIONAL ORGANIZATIONAL STRUCTURES

Roles: parametrize required expertise

Teams: groups of workers with shared goal

Hierarchy: nested roles that determine decision rights

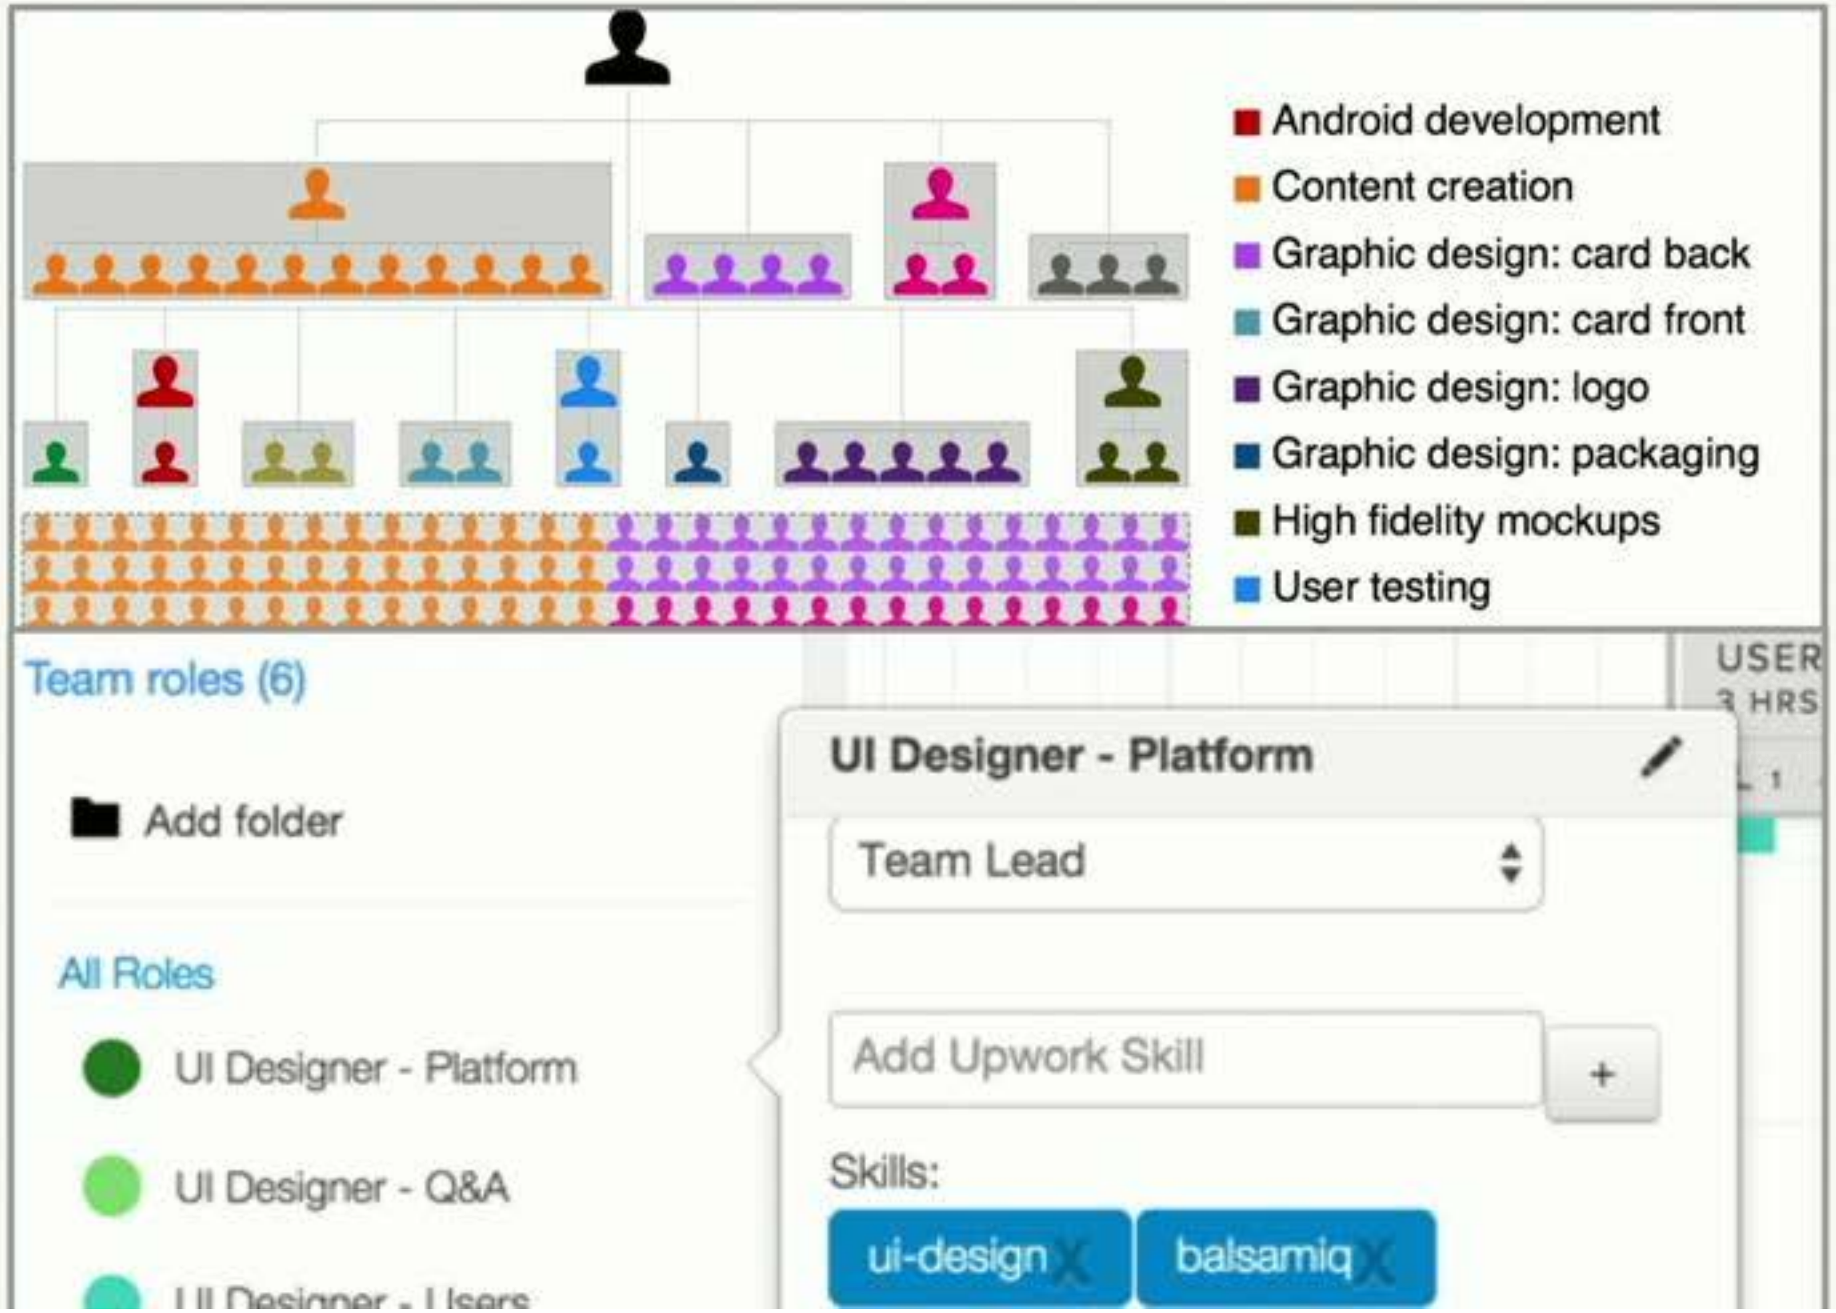

ON-DEMAND HIRING FROM UPWORK

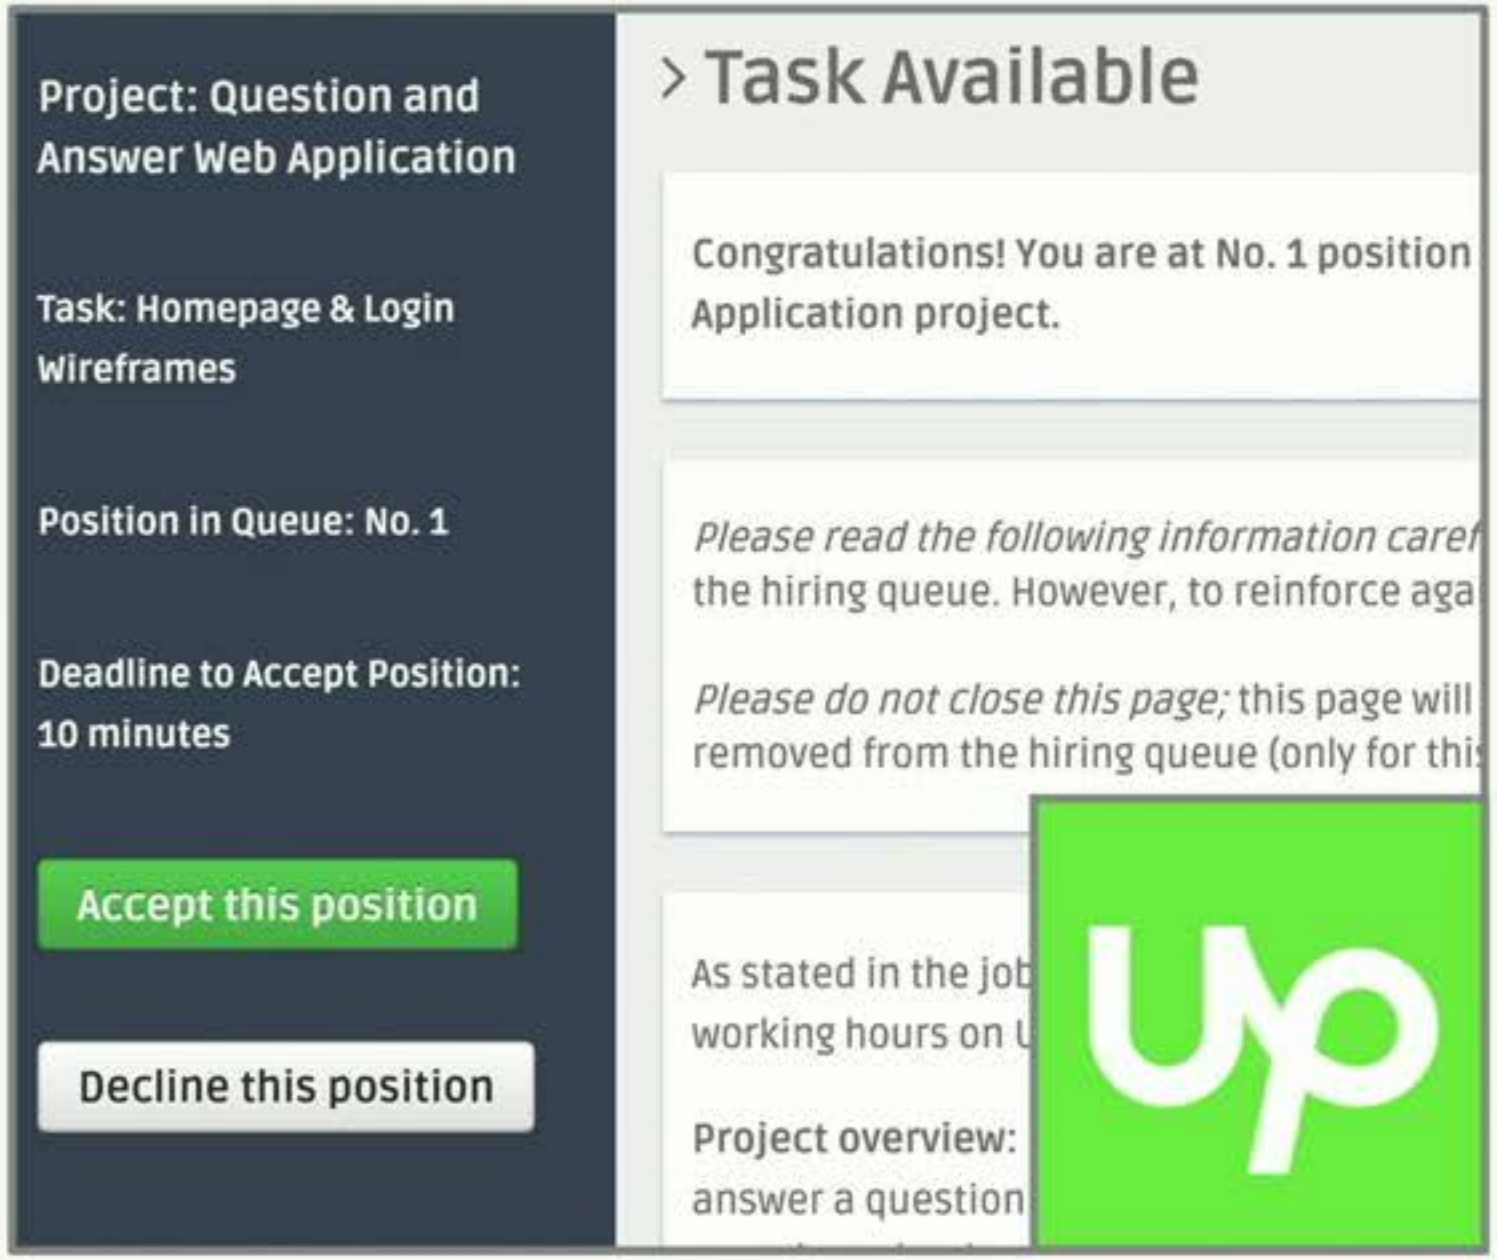

Foundry hiring queue

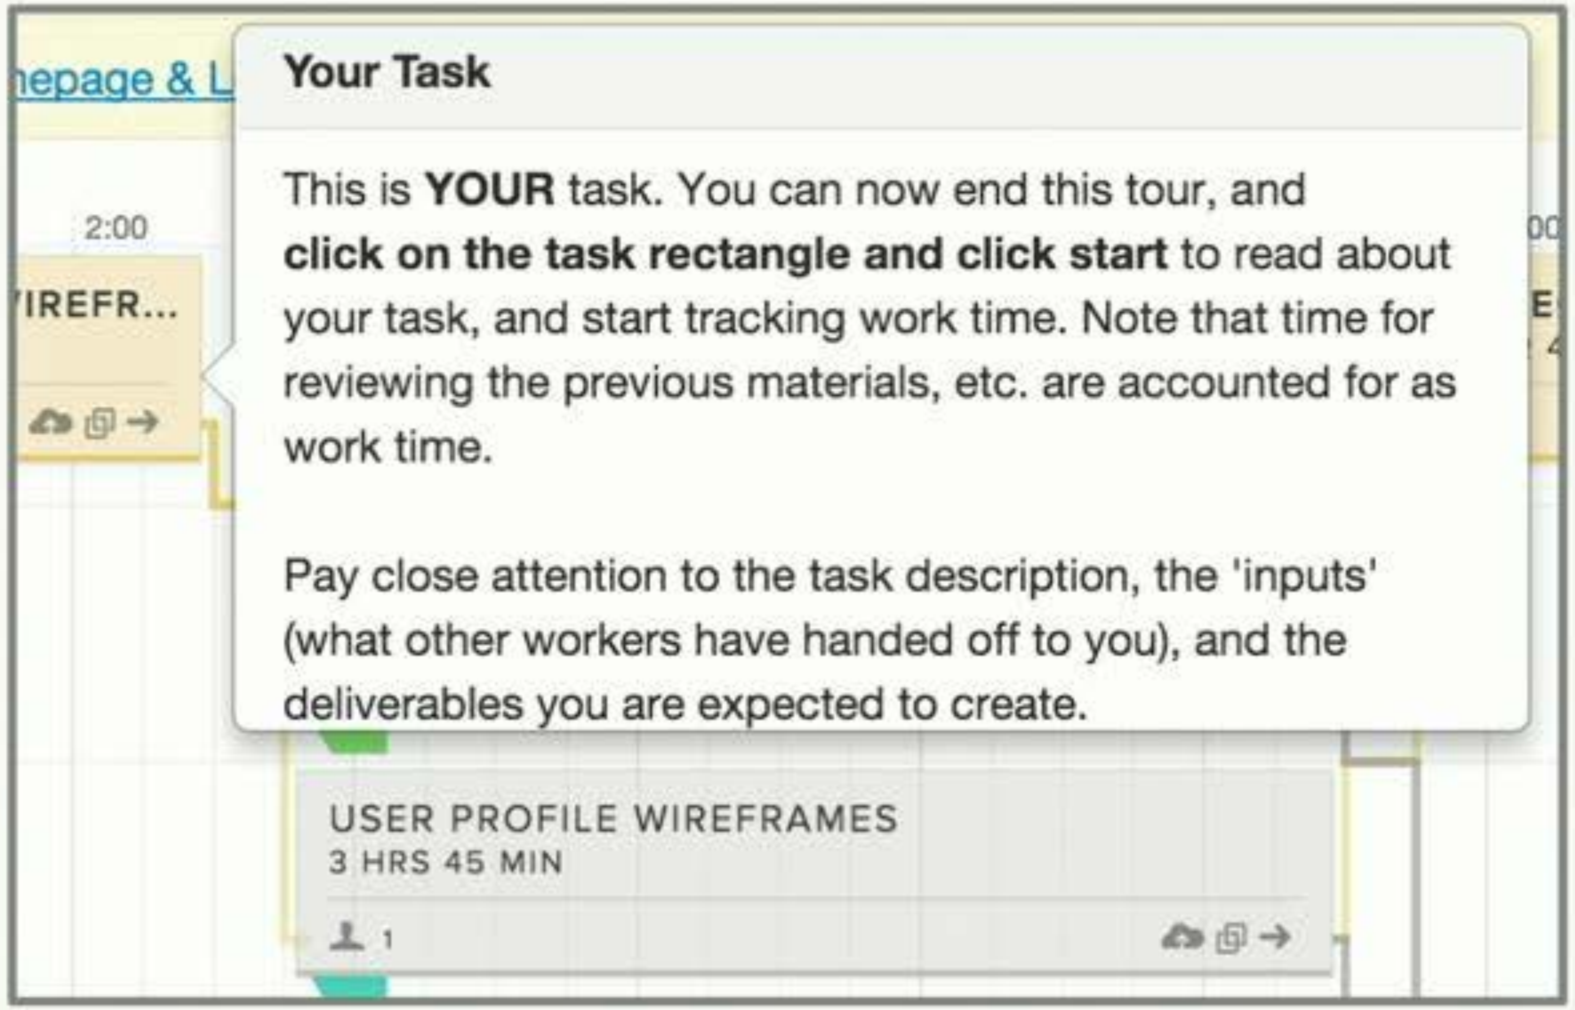

Automated, role-specific onboarding

CHALLENGE: RECONFIGURATION

Organizational structures require constant reconfiguration so that the organization can adapt as it proceeds

How can a computational system keep a distributed crowd in sync as the plan evolves?

VERSION CONTROL

To enable reconfiguration of the organizational structures:
branching and merging inspired by version control

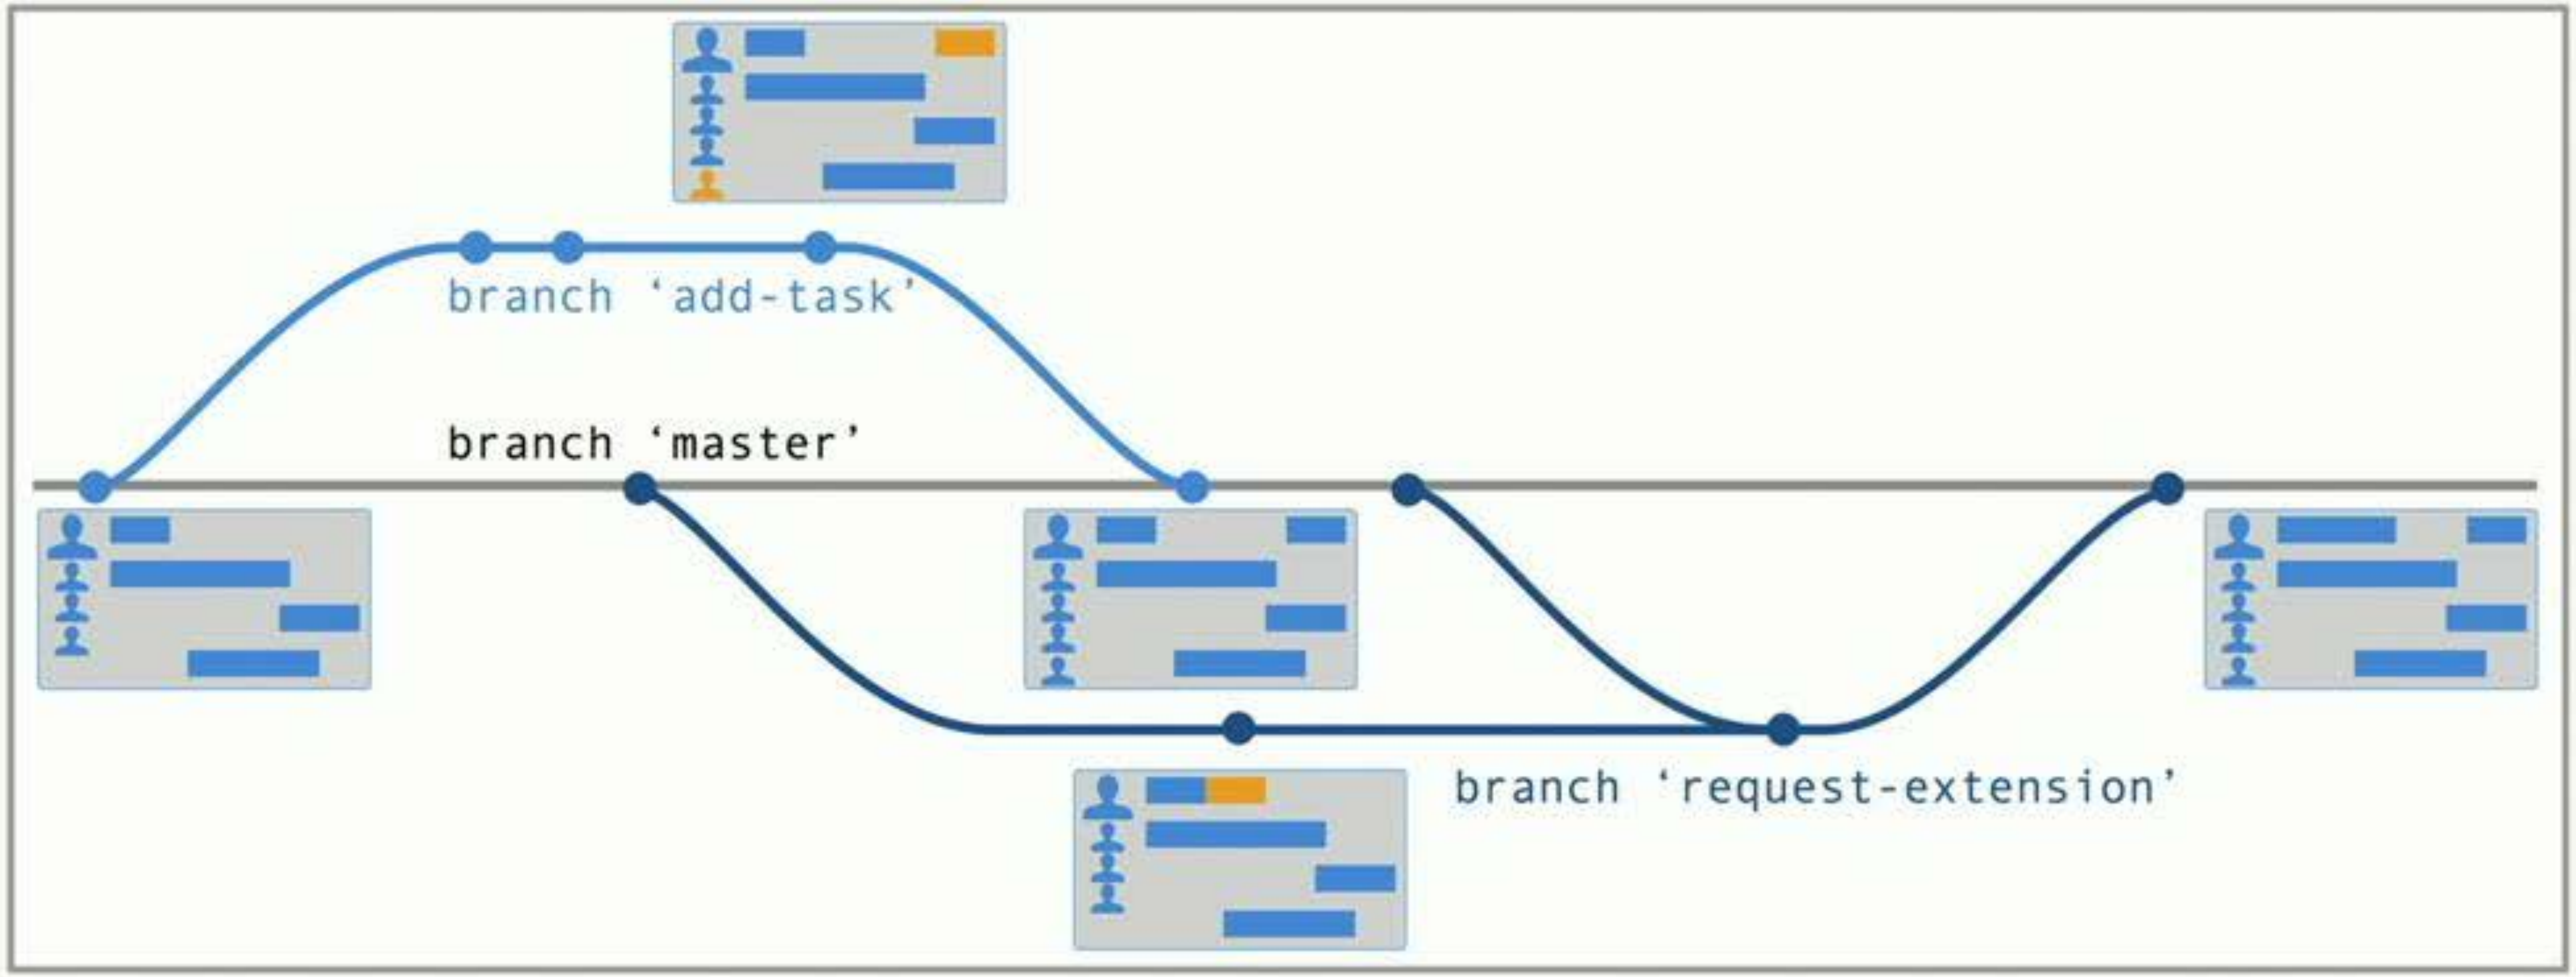

VERSION CONTROL IN FOUNDRY

Any member can **branch, edit, and make pull requests** against any organizational structure: roles, teams, hierarchy, tasks

Diff view for reviewing

Pull requests are reviewed up the hierarchy and merged through a three-way diff

EVALUATION

Field study: System deployment with outside leaders willing to crowdsource their complex open-ended goals

	EMS Report
Leader	Medical resident
Open-ended goal	Develop prototype application for EMTs to transmit patient information en route to hospital

EVALUATION

Field study: System deployment with outside leaders willing to crowdsource their complex open-ended goals

	EMS Report	True Story
Leader	Medical resident	Storytelling podcast kickstarter team
Open-ended goal	Develop prototype application for EMTs to transmit patient information en route to hospital	Design and manufacture a storytelling card game with accompanying mobile application

EVALUATION

Field study: System deployment with outside leaders willing to crowdsource their complex open-ended goals

	EMS Report	True Story	Workshop Portal
Leader	Medical resident	Storytelling podcast kickstarter team	Tech lab employee of a large company
Open-ended goal	Develop prototype application for EMTs to transmit patient information en route to hospital	Design and manufacture a storytelling card game with accompanying mobile application	Develop a workshop planning portal consistent with enterprise standards and branding

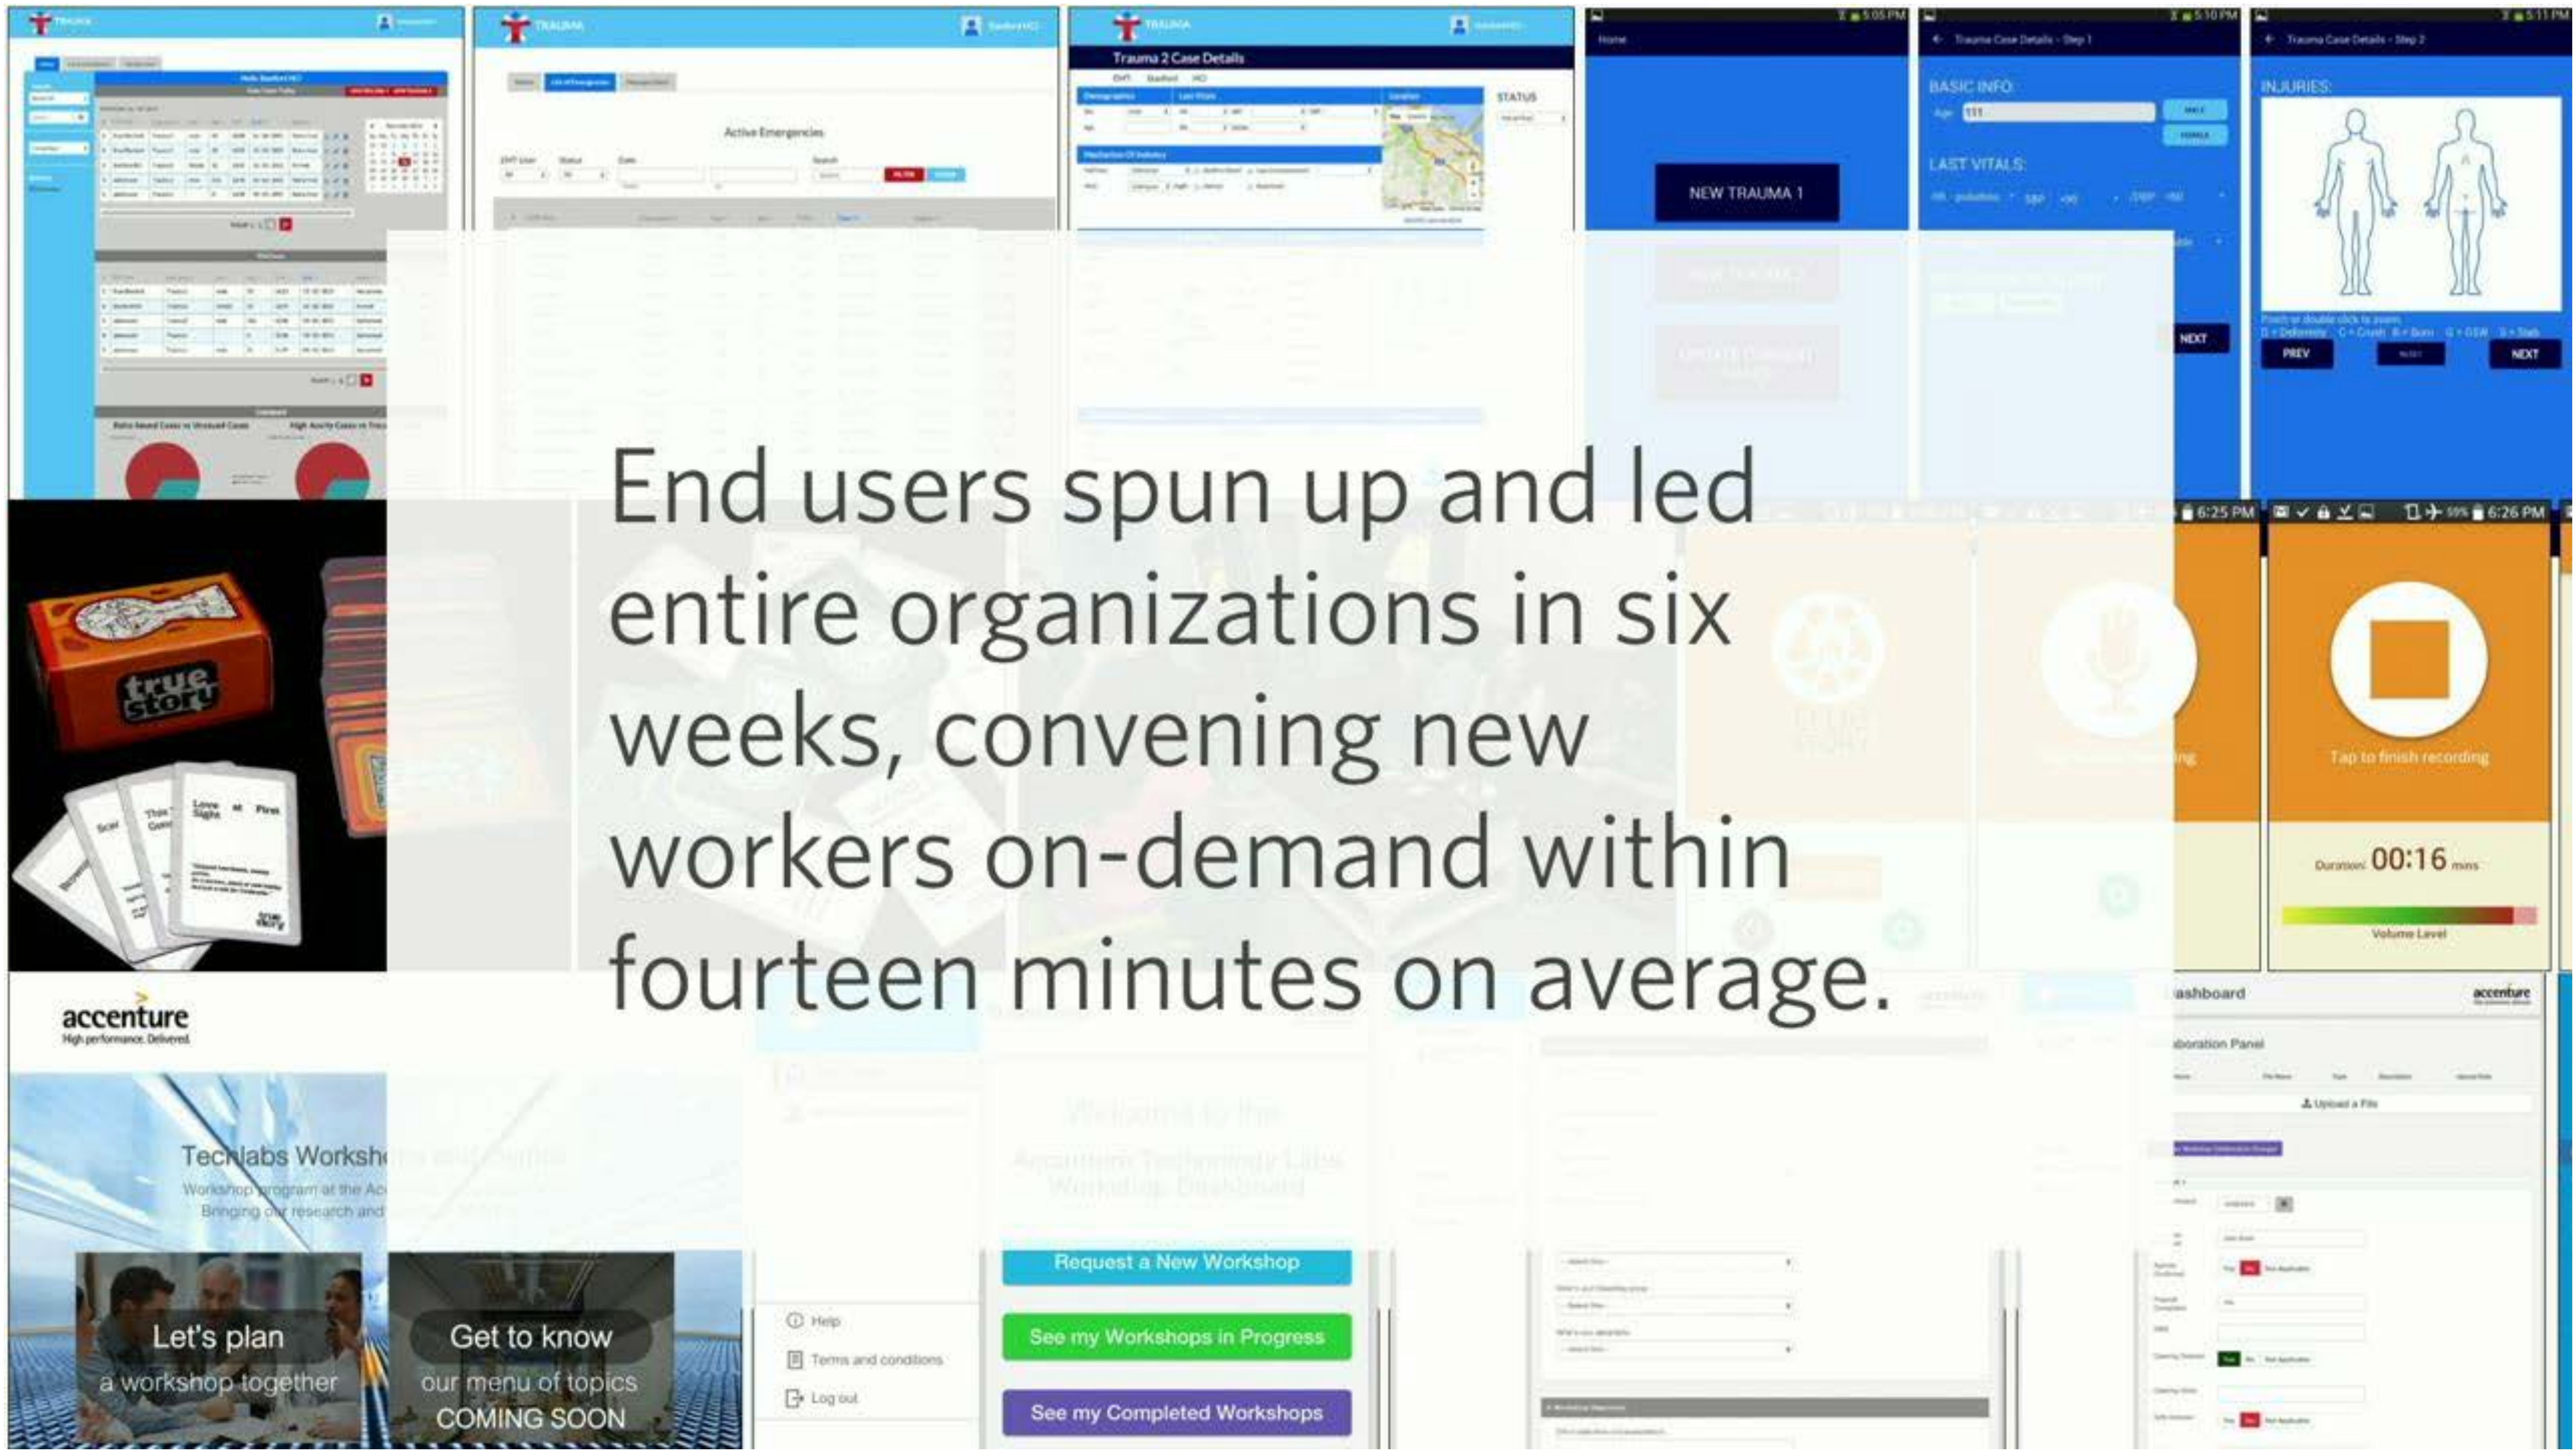

End users spun up and led entire organizations in six weeks, convening new workers on-demand within fourteen minutes on average.

2 mobile applications, 3 full-stack web applications consisting of 52,000 lines of code, 2 illustrated card decks

639 tasks across 566 pull requests, 3261 person-hours of work across 35–46 days including engineers, designers, testers, and poets

Passed quality review by neutral experts

TRUE STORY GAME

TRUE STORY GAME

CRUSHING

Subtle looks, pounding pulse
However long the hover lasts
Between friend zone and fun zone

TRUE STORY GAME

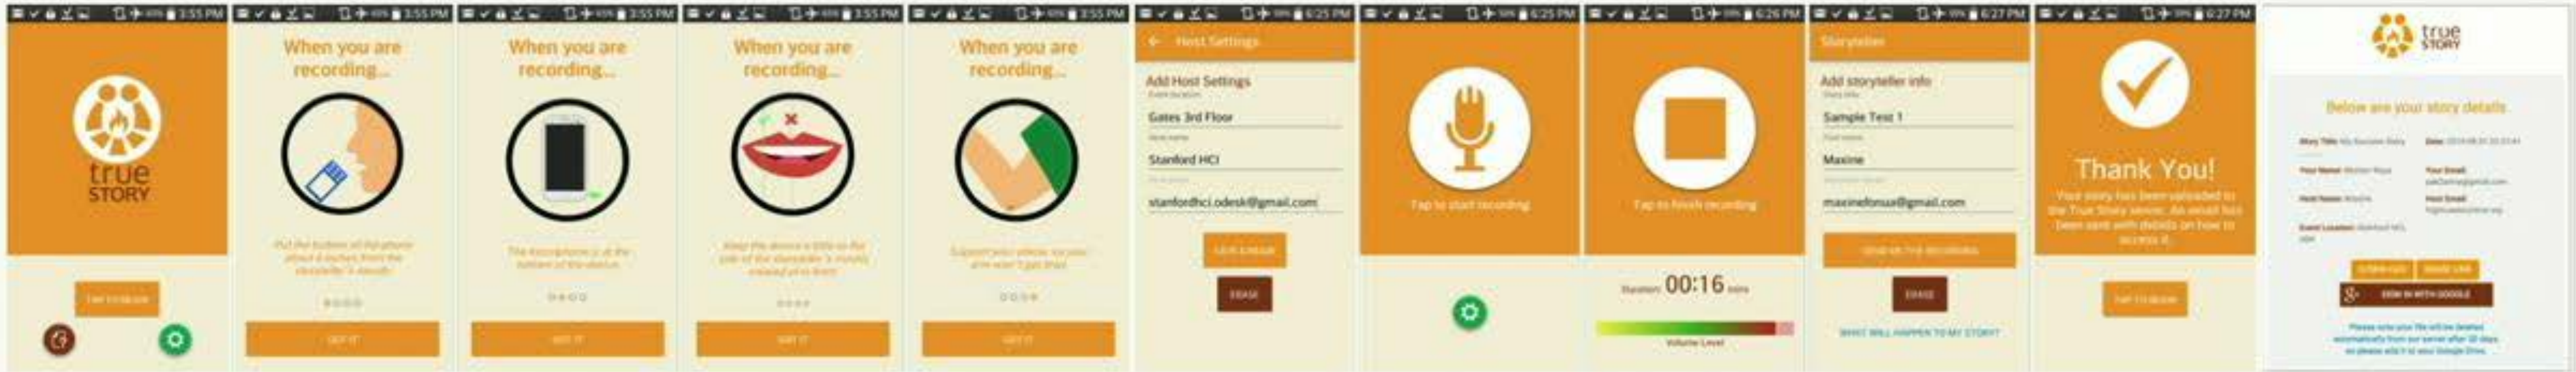

Android companion app spun up in the final week

ENTERPRISE WORKSHOP PORTAL

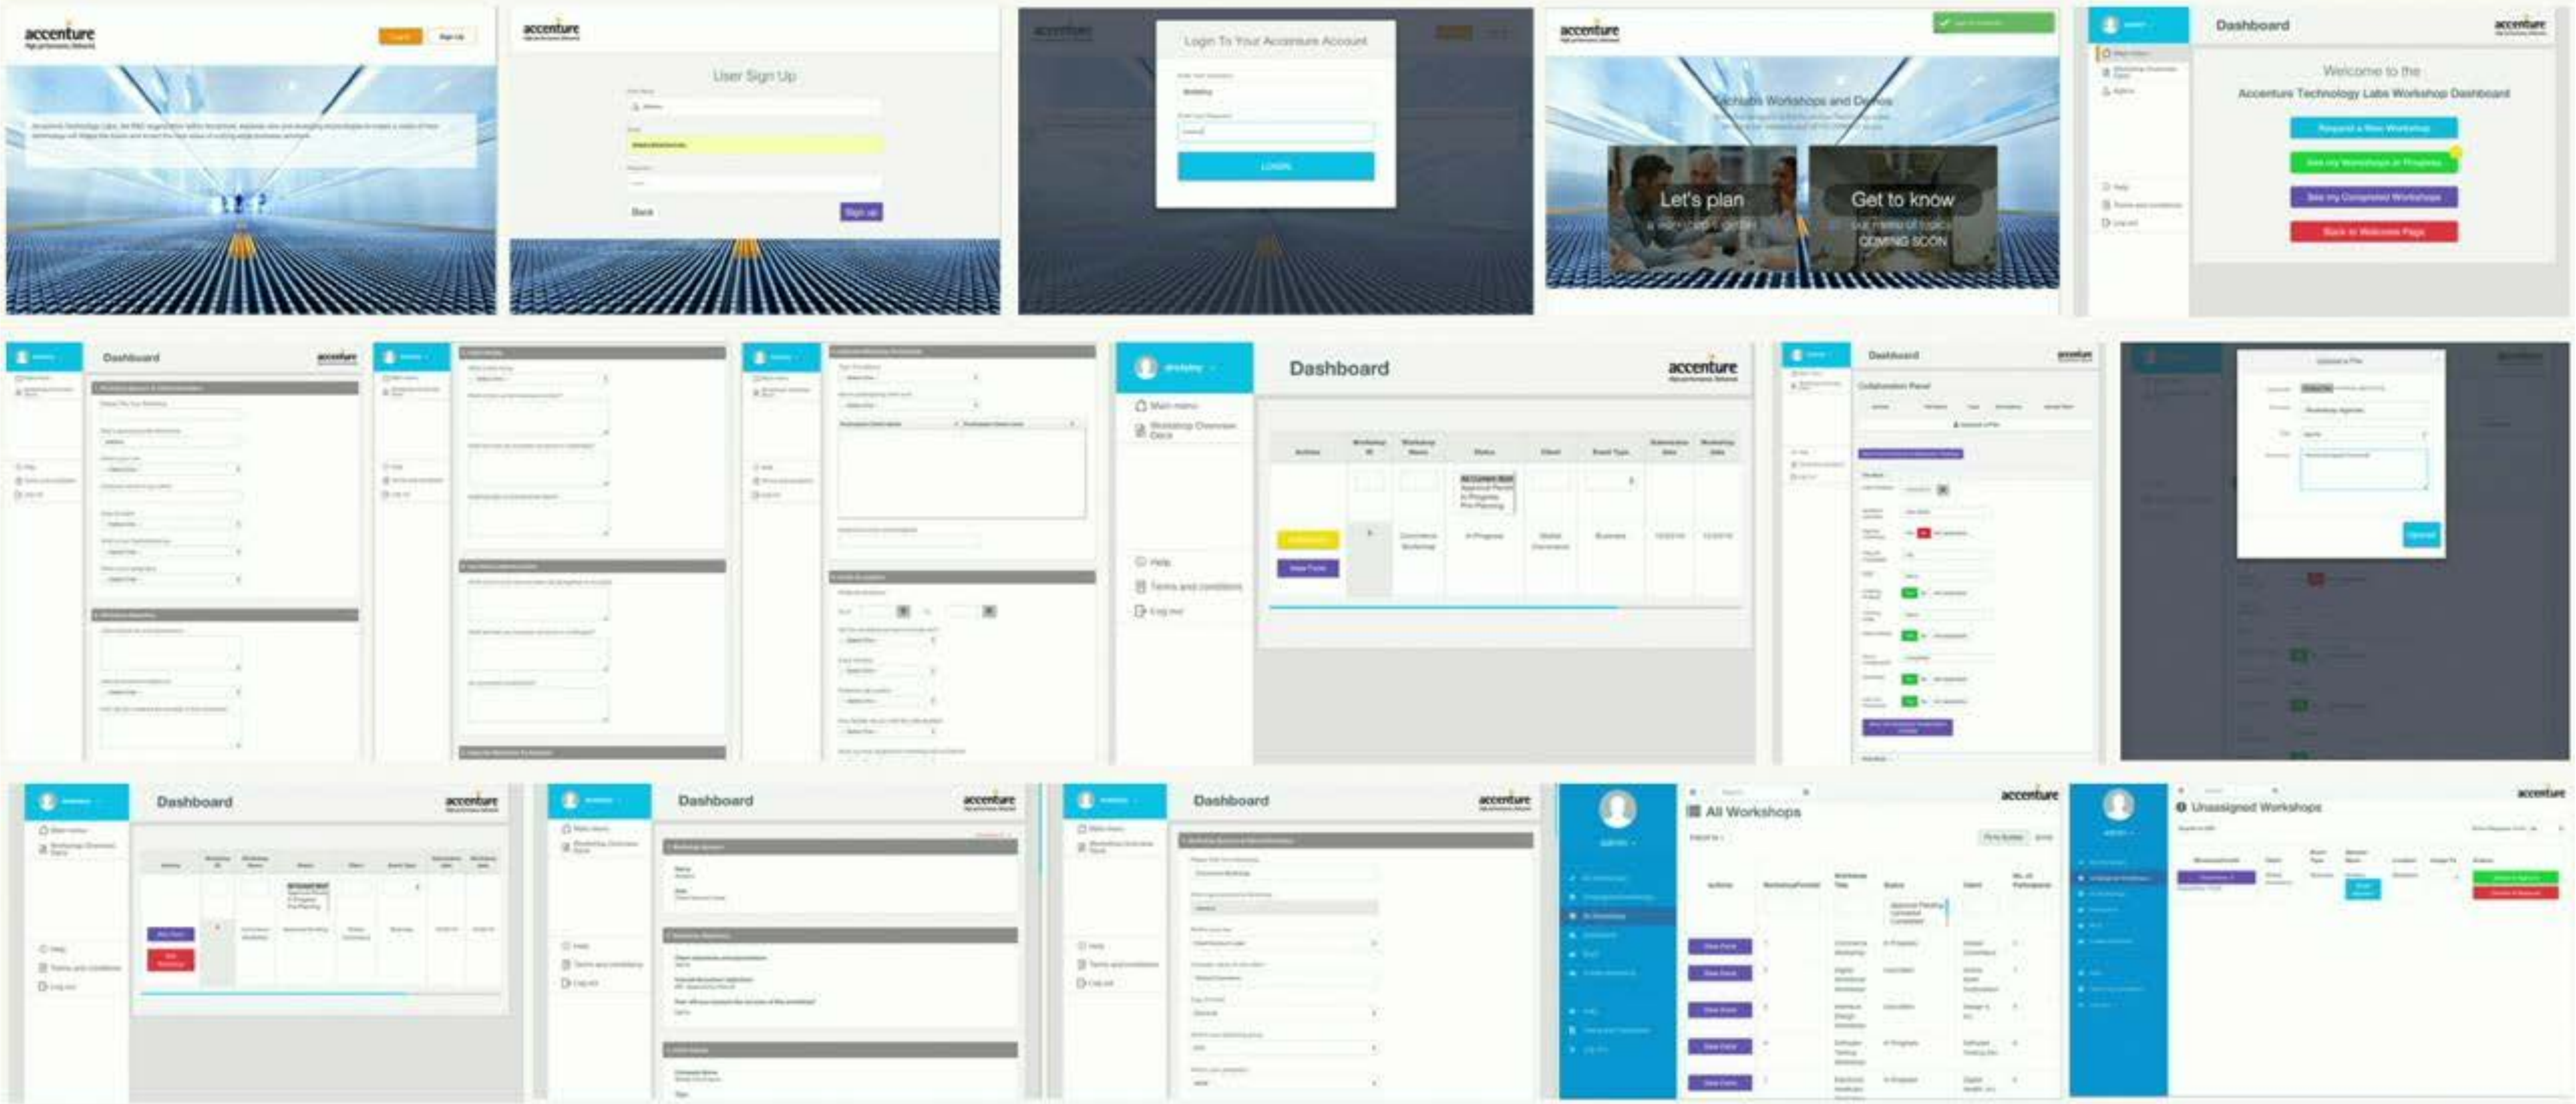

FLASH ORGANIZATIONS: REFLECTION

Flash organizations suggest a future where organizations are **fluidly assembled and adapted** within globally networked collectives.

Open questions:

- Might data help suggest effective organizational structures?
- How can researchers support industry norms, labor organization, and legislation to encourage a prosocial future of work?
- Do flash organizations change the transaction costs core to the Theory of the Firm?

EMS TRAUMA REPORT

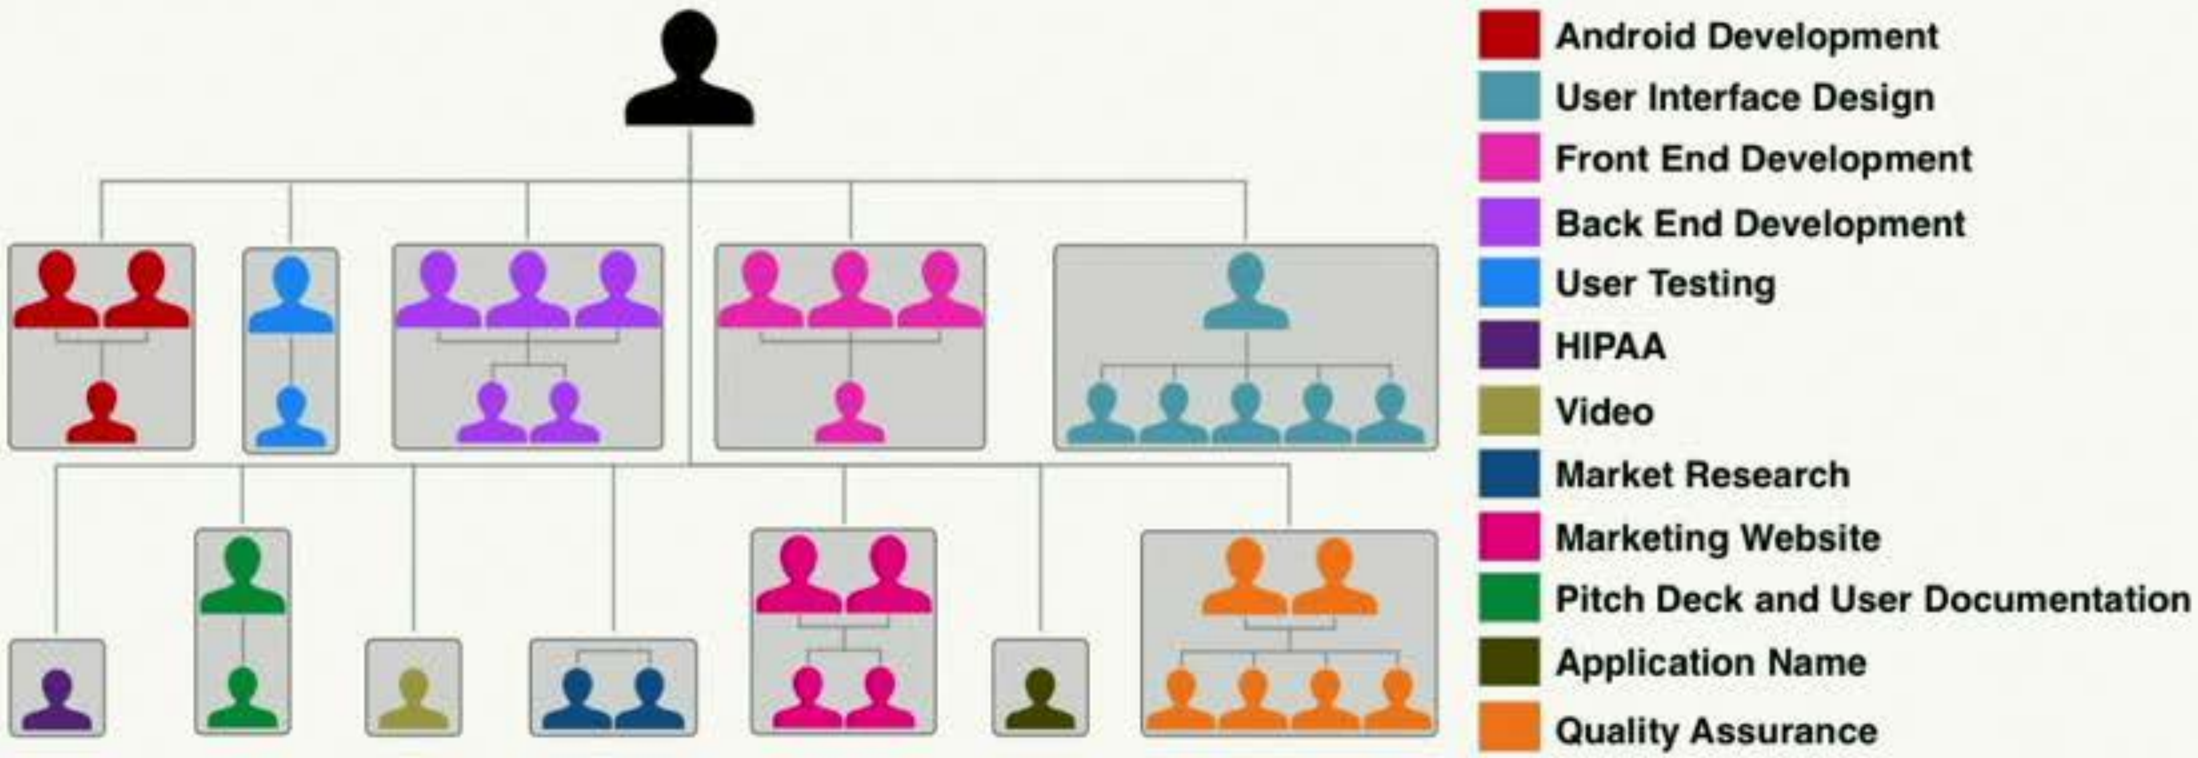

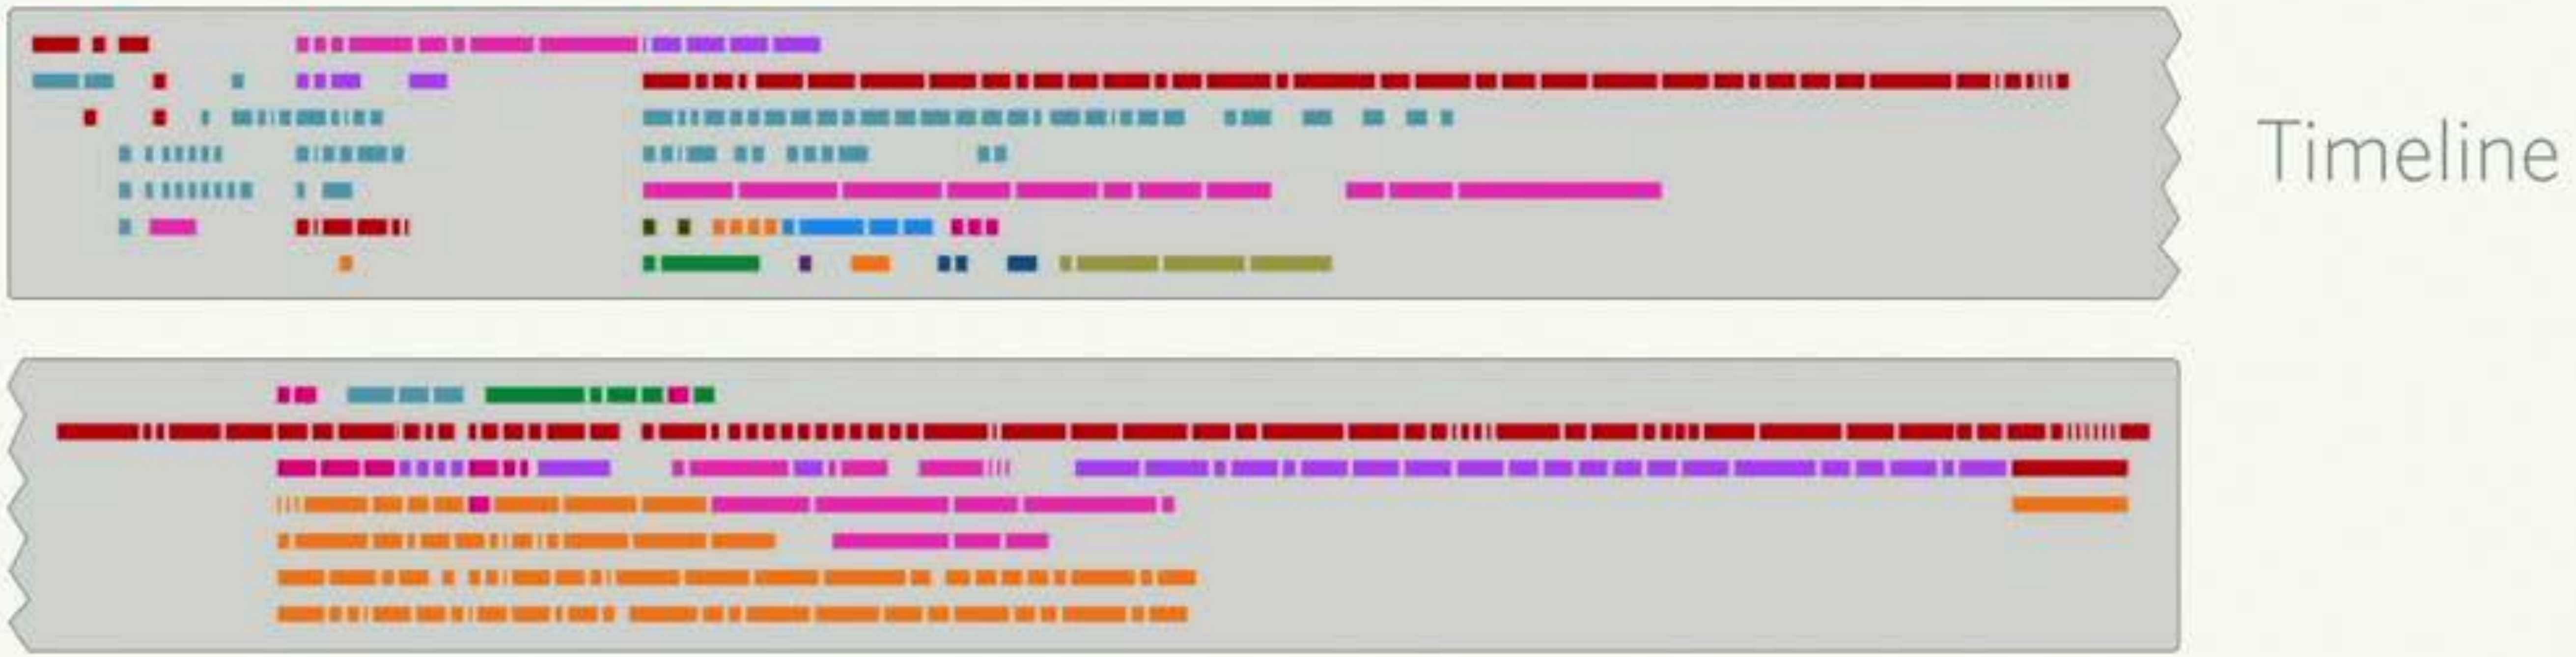

Dream Team

Zhou, Valentine, Bernstein. CHI 2018.

WHAT IS THE BEST WAY FOR TEAMS TO ORGANIZE THEMSELVES?

Organizations rely on teaming and ad-hoc collaboration
[Edmondson 2012]

WHAT IS THE BEST WAY FOR TEAMS TO ORGANIZE THEMSELVES?

Organizations rely on teaming and ad-hoc collaboration
[Edmondson 2012]

But: Should teams be flat or hierarchical? Encouraging or critical? Enforcing equal turn-taking?

WHAT IS THE BEST WAY FOR TEAMS TO ORGANIZE THEMSELVES?

Organizations rely on teaming and ad-hoc collaboration
[Edmondson 2012]

But: Should teams be flat or hierarchical? Encouraging or critical? Enforcing equal turn-taking?

- These roles, norms, and interaction patterns define **team structures**
  [Ilgen et al. 2005]

WHAT IS THE BEST WAY FOR TEAMS TO ORGANIZE THEMSELVES?

Organizations rely on teaming and ad-hoc collaboration [Edmondson 2012]

But: Should teams be flat or hierarchical? Encouraging or critical? Enforcing equal turn-taking?

- These roles, norms, and interaction patterns define **team structures** [Ilgen et al. 2005]
- Researchers theorize ideal structures, then build systems nudging teams toward those structures [Olsen & Olsen 2000; Ackerman 2000; Dourish & Bellotti 1992; Erickson & Kellogg 2000; Winograd 1986; Lykourentzou et al. 2017]

ORG. BEHAVIOR: THERE ARE NO UNIVERSALLY IDEAL STRUCTURES

Structural contingency theory: the best team structures depend on the task and the team members [Donaldson 1999]

ORG. BEHAVIOR: THERE ARE NO UNIVERSALLY IDEAL STRUCTURES

Structural contingency theory: the best team structures depend on the task and the team members [Donaldson 1999]

The wrong structures will doom teams to dysfunction
[Ilgen et al. 2005; Schippers, Edmondson, & West 2014; Ancona, Okhuysen & Perlow 2001]

ORG. BEHAVIOR: THERE ARE NO UNIVERSALLY IDEAL STRUCTURES

Structural contingency theory: the best team structures depend on the task and the team members [Donaldson 1999]

The wrong structures will doom teams to dysfunction [Ilgen et al. 2005; Schippers, Edmondson, & West 2014; Ancona, Okhuysen & Perlow 2001]

Managers — who are trained and paid for choosing effective team structures — are not effective at the task [de Brujin, Ten Heuvelhof, & In 't Veld 2002]

DREAMTEAM

Rapid self-experimentation with different team structures to converge on one that works well for the team and task

DREAMTEAM

Rapid self-experimentation with different team structures to converge on one that works well for the team and task

DREAMTEAM

Rapid self-experimentation with different team structures to converge on one that works well for the team and task

DREAMTEAM

Rapid self-experimentation with different team structures to converge on one that works well for the team and task

DreamTeam

DREAMTEAM

Rapid self-experimentation with different team structures to converge on one that works well for the team and task

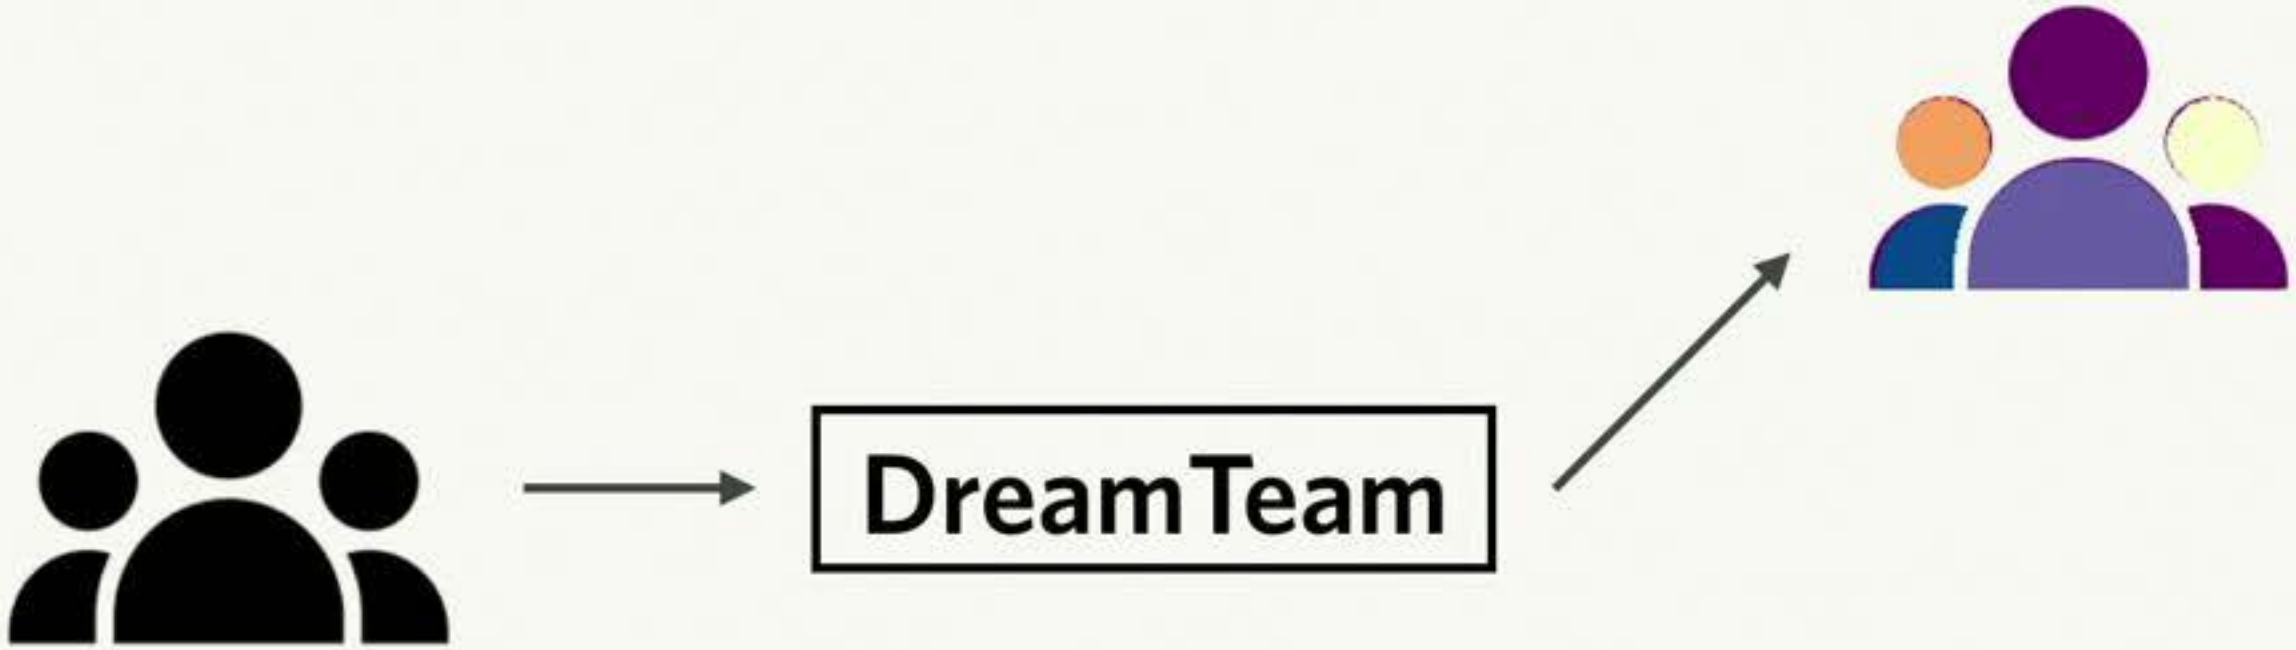

DREAMTEAM

Rapid self-experimentation with different team structures to converge on one that works well for the team and task

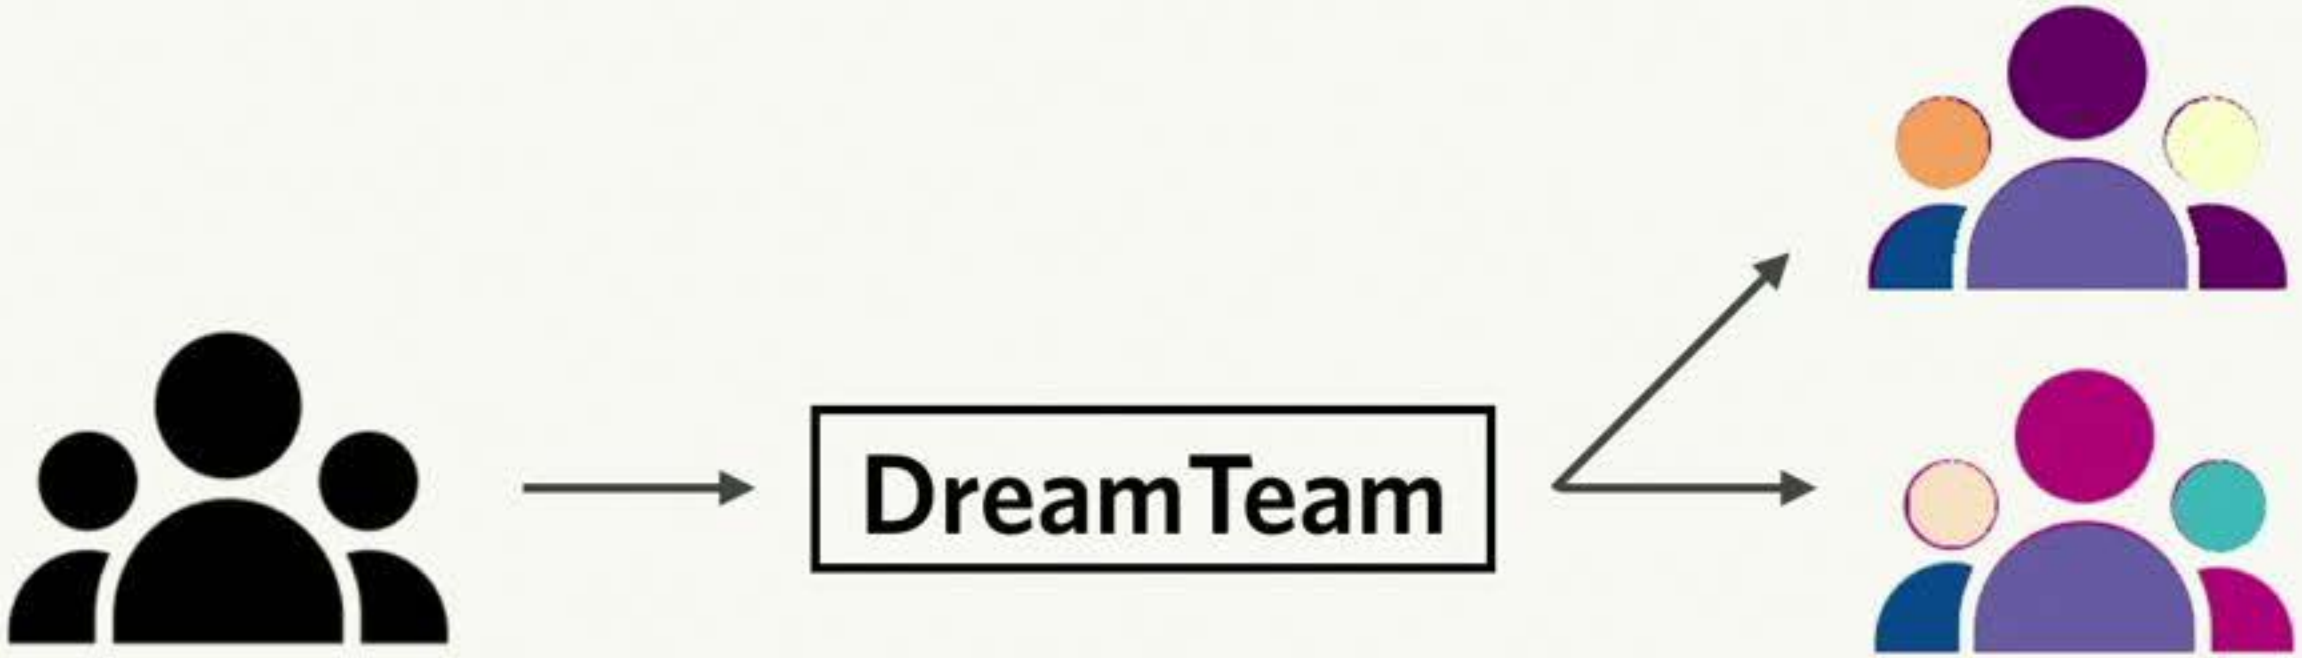

DREAMTEAM

Rapid self-experimentation with different team structures to converge on one that works well for the team and task

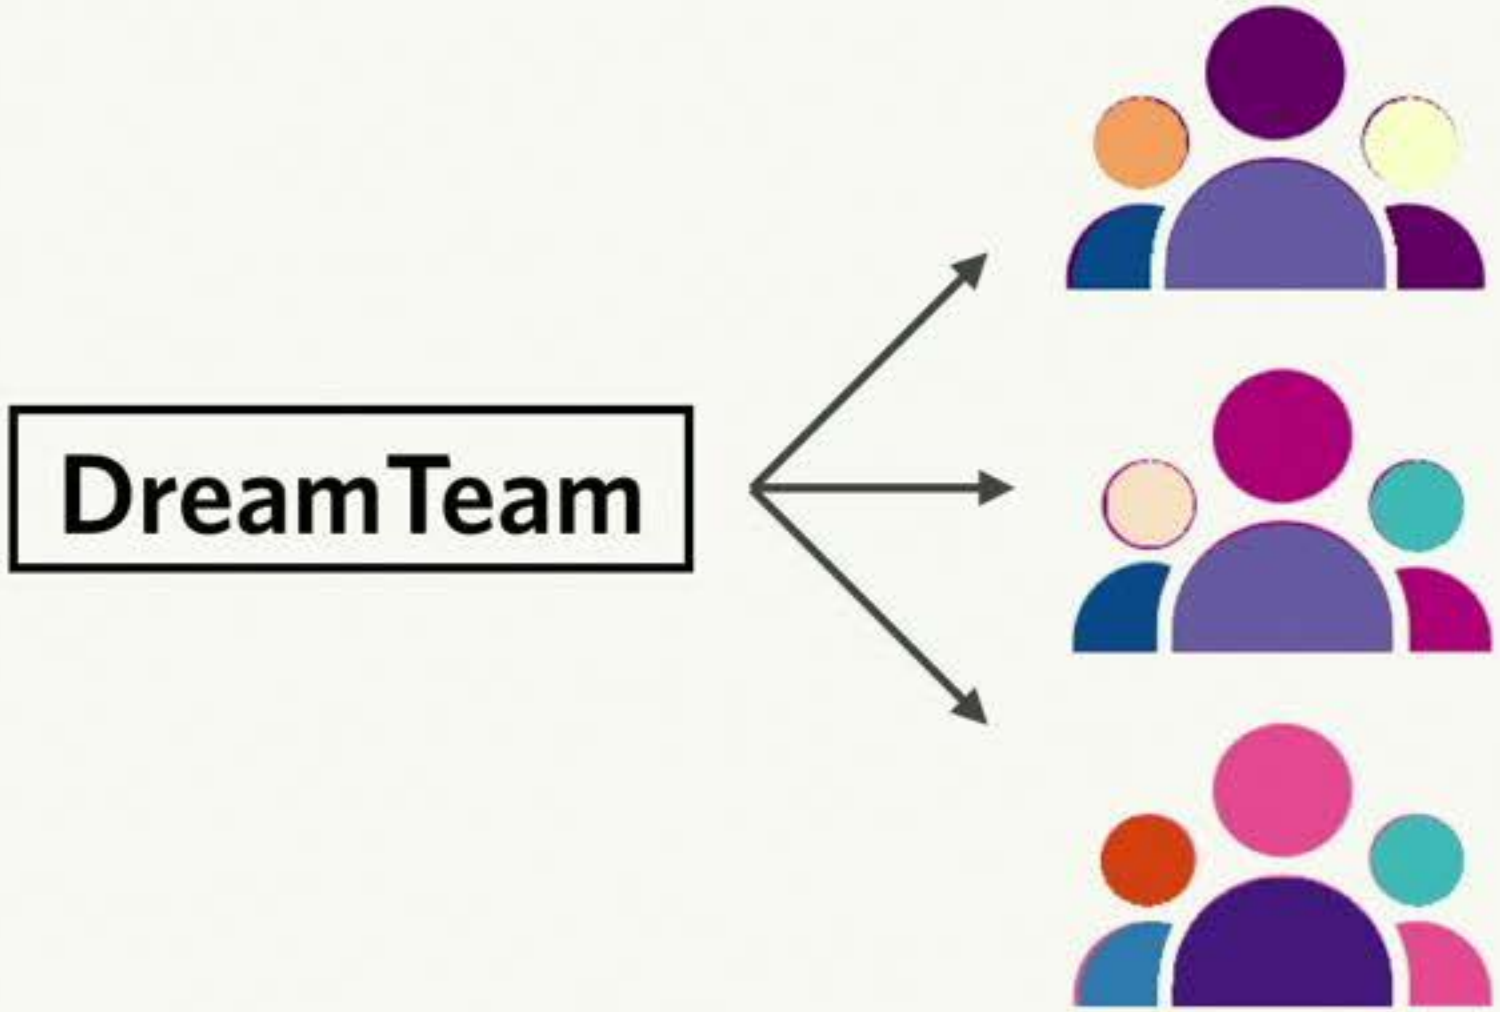

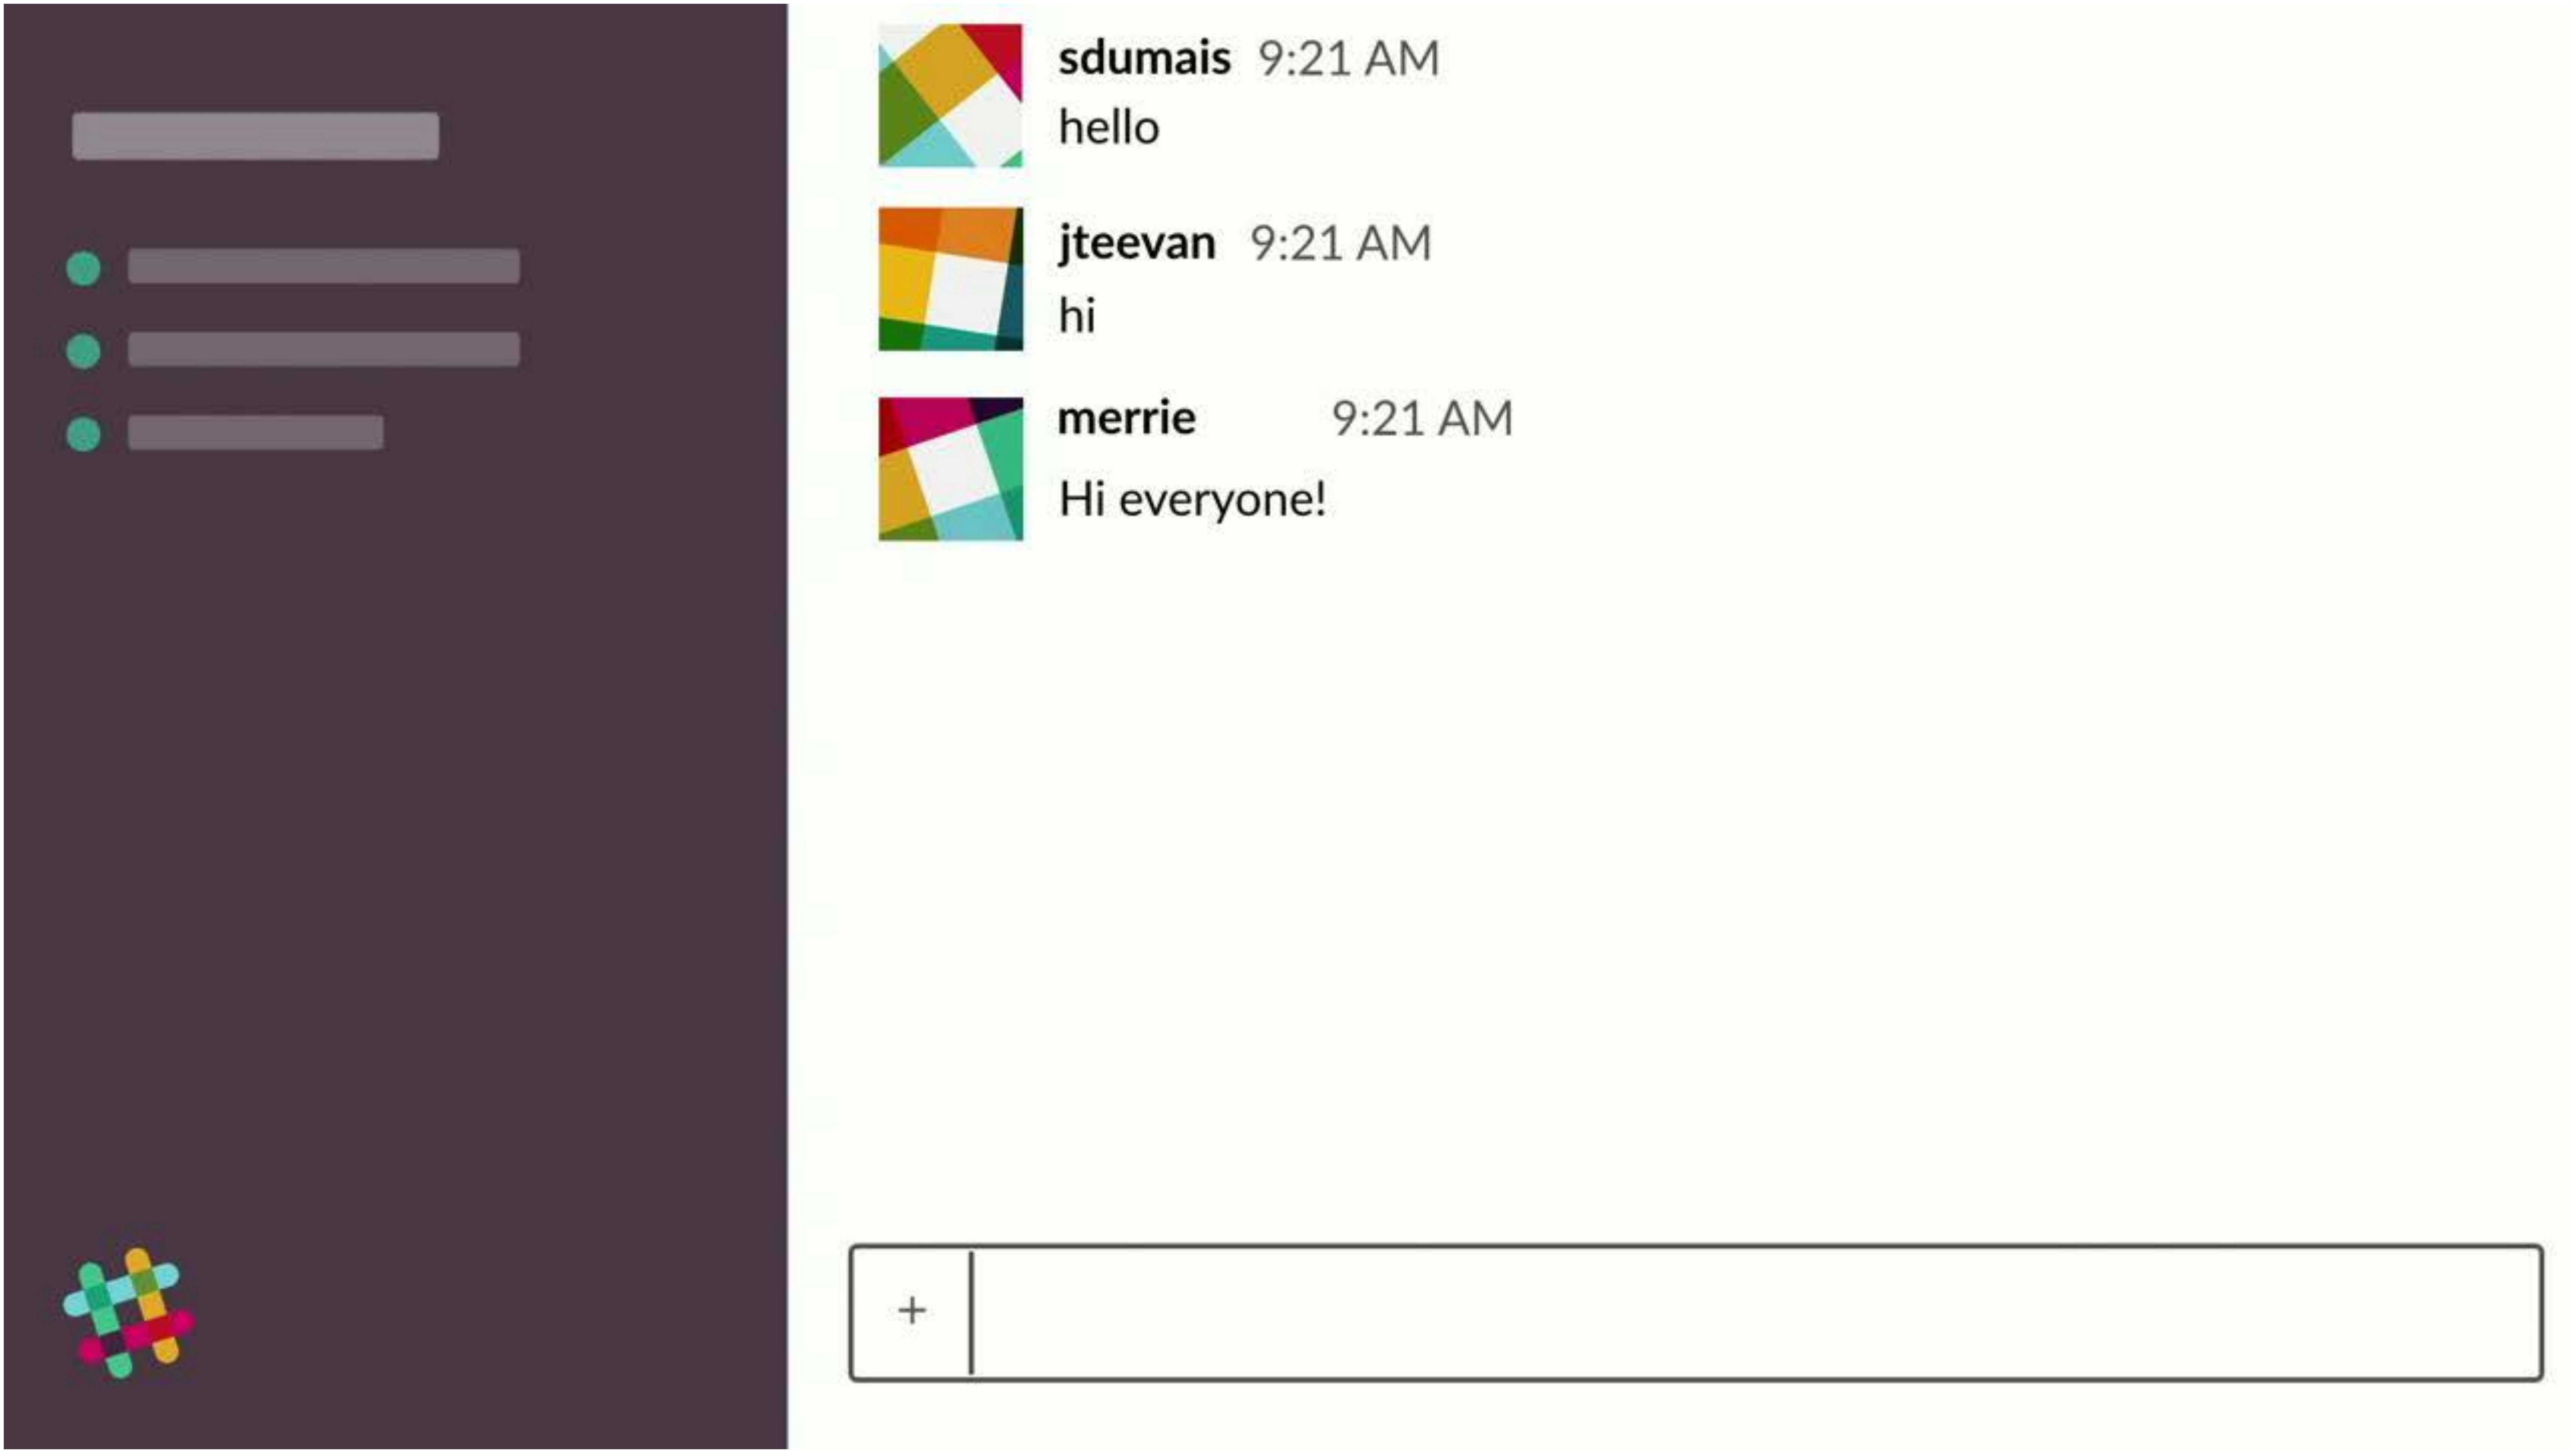
sdumais 9:21 AM
hello
jteevan 9:21 AM
hi
merrie 9:21 AM
Hi everyone!
+

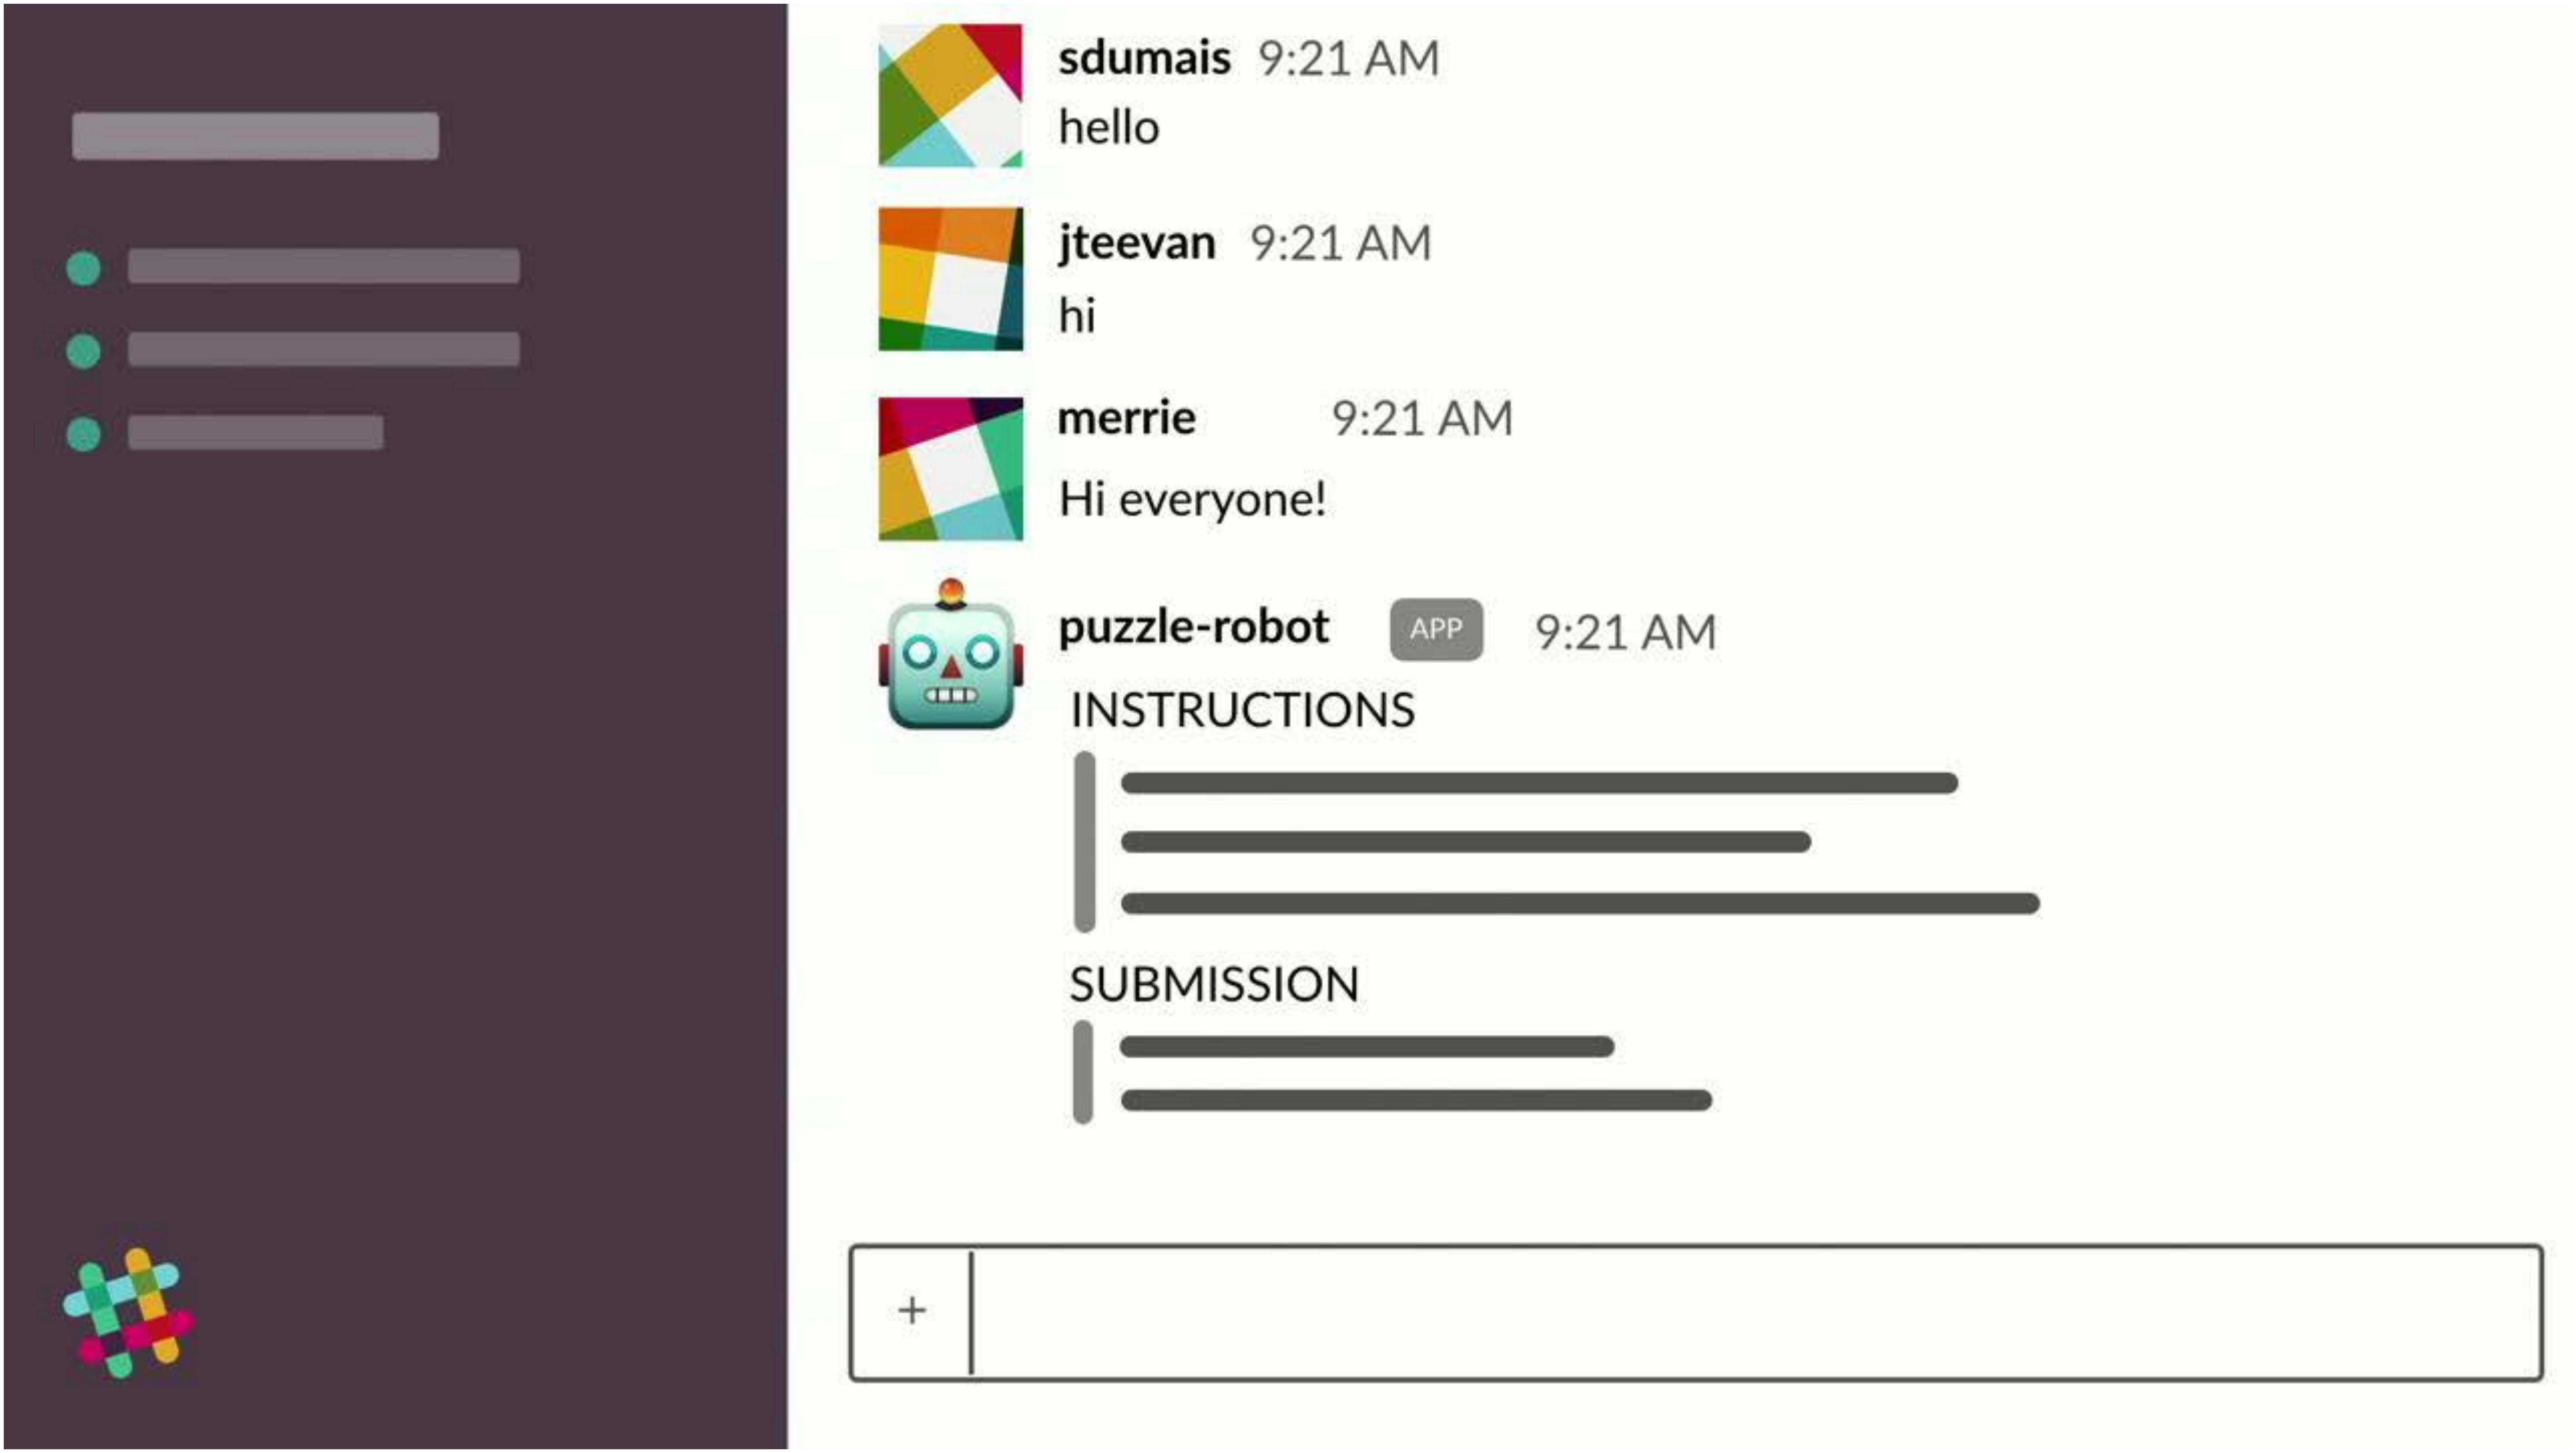
sdumais 9:21 AM
hello
jteevan 9:21 AM
hi
merrie 9:21 AM
Hi everyone!
puzzle-robot APP 9:21 AM
INSTRUCTIONS
SUBMISSION
+

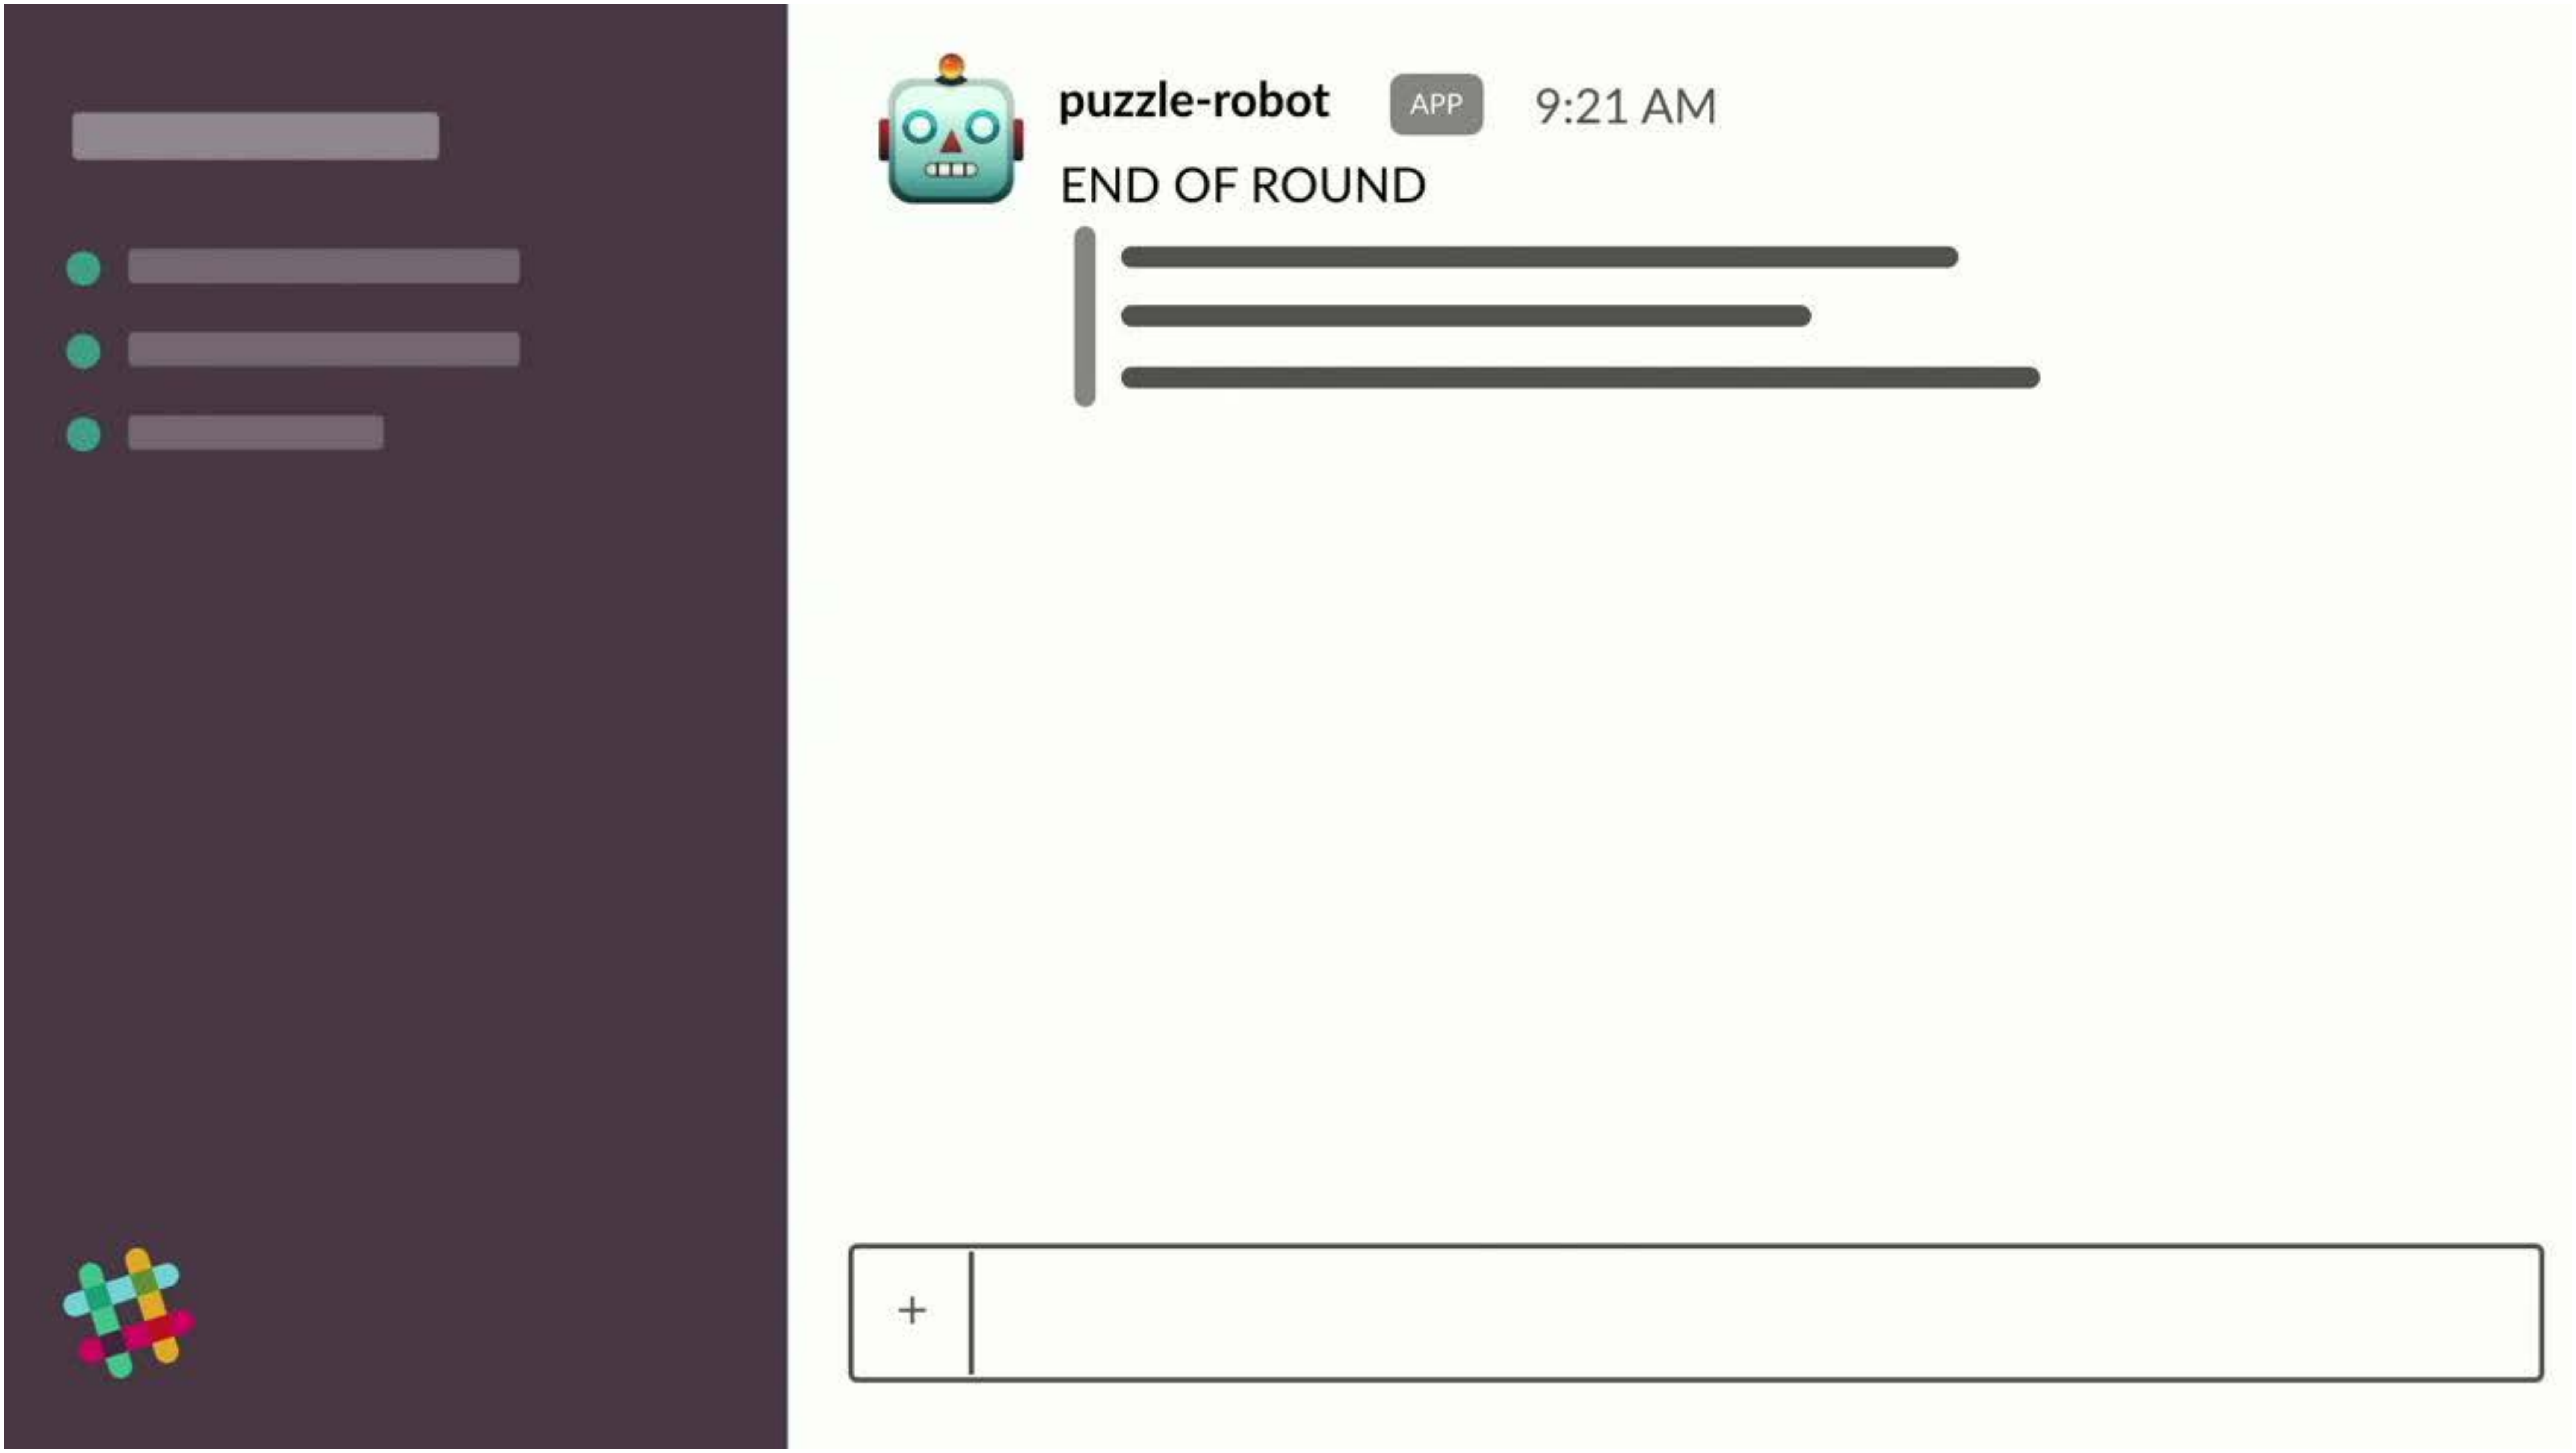
puzzle-robot
APP
9:21 AM
END OF ROUND
+

puzzle-robot APP 9:21 AM

END OF ROUND

<feedback to DreamTeam system>

puzzle-robot APP 9:21 AM

END OF ROUND

<feedback to DreamTeam system>

dreamteam-robot 9:21 AM

This round change the following...

> Be super cheery! Make sure to write encouraging comments to all your teammates, despite any losses!

+

Hierarchy
None, Centralized, Decentralized

Interaction Patterns
Emergent, Round-robin, Equally distributed

Norms of Engagement
None, Professional, Informal

Decision-Making Norms
None, Divergent, Convergent, Informed, Rapid

Feedback Norms
None, Encouraging, Critical

Time

Hierarchy
None, Centralized, Decentralized

Interaction Patterns
Emergent, Round-robin, Equally distributed

Norms of Engagement
None, Professional, Informal

Decision-Making Norms
None, Divergent, Convergent, Informed, Rapid

Feedback Norms
None, Encouraging, Critical

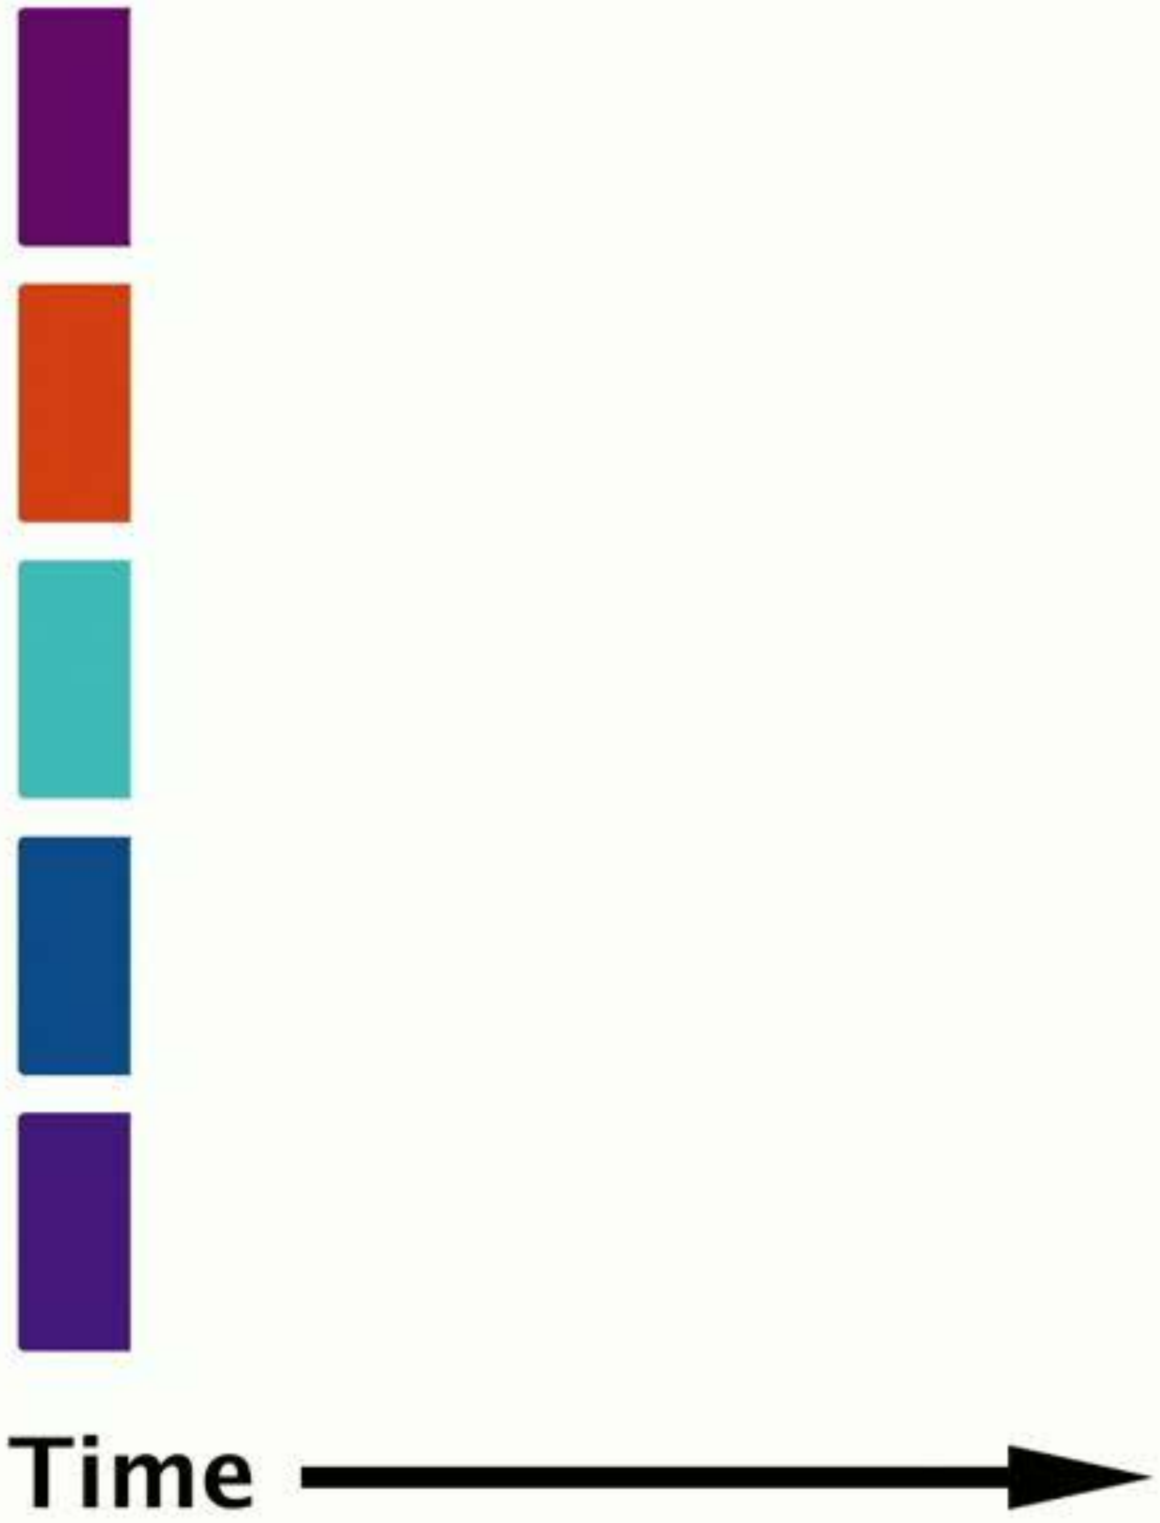

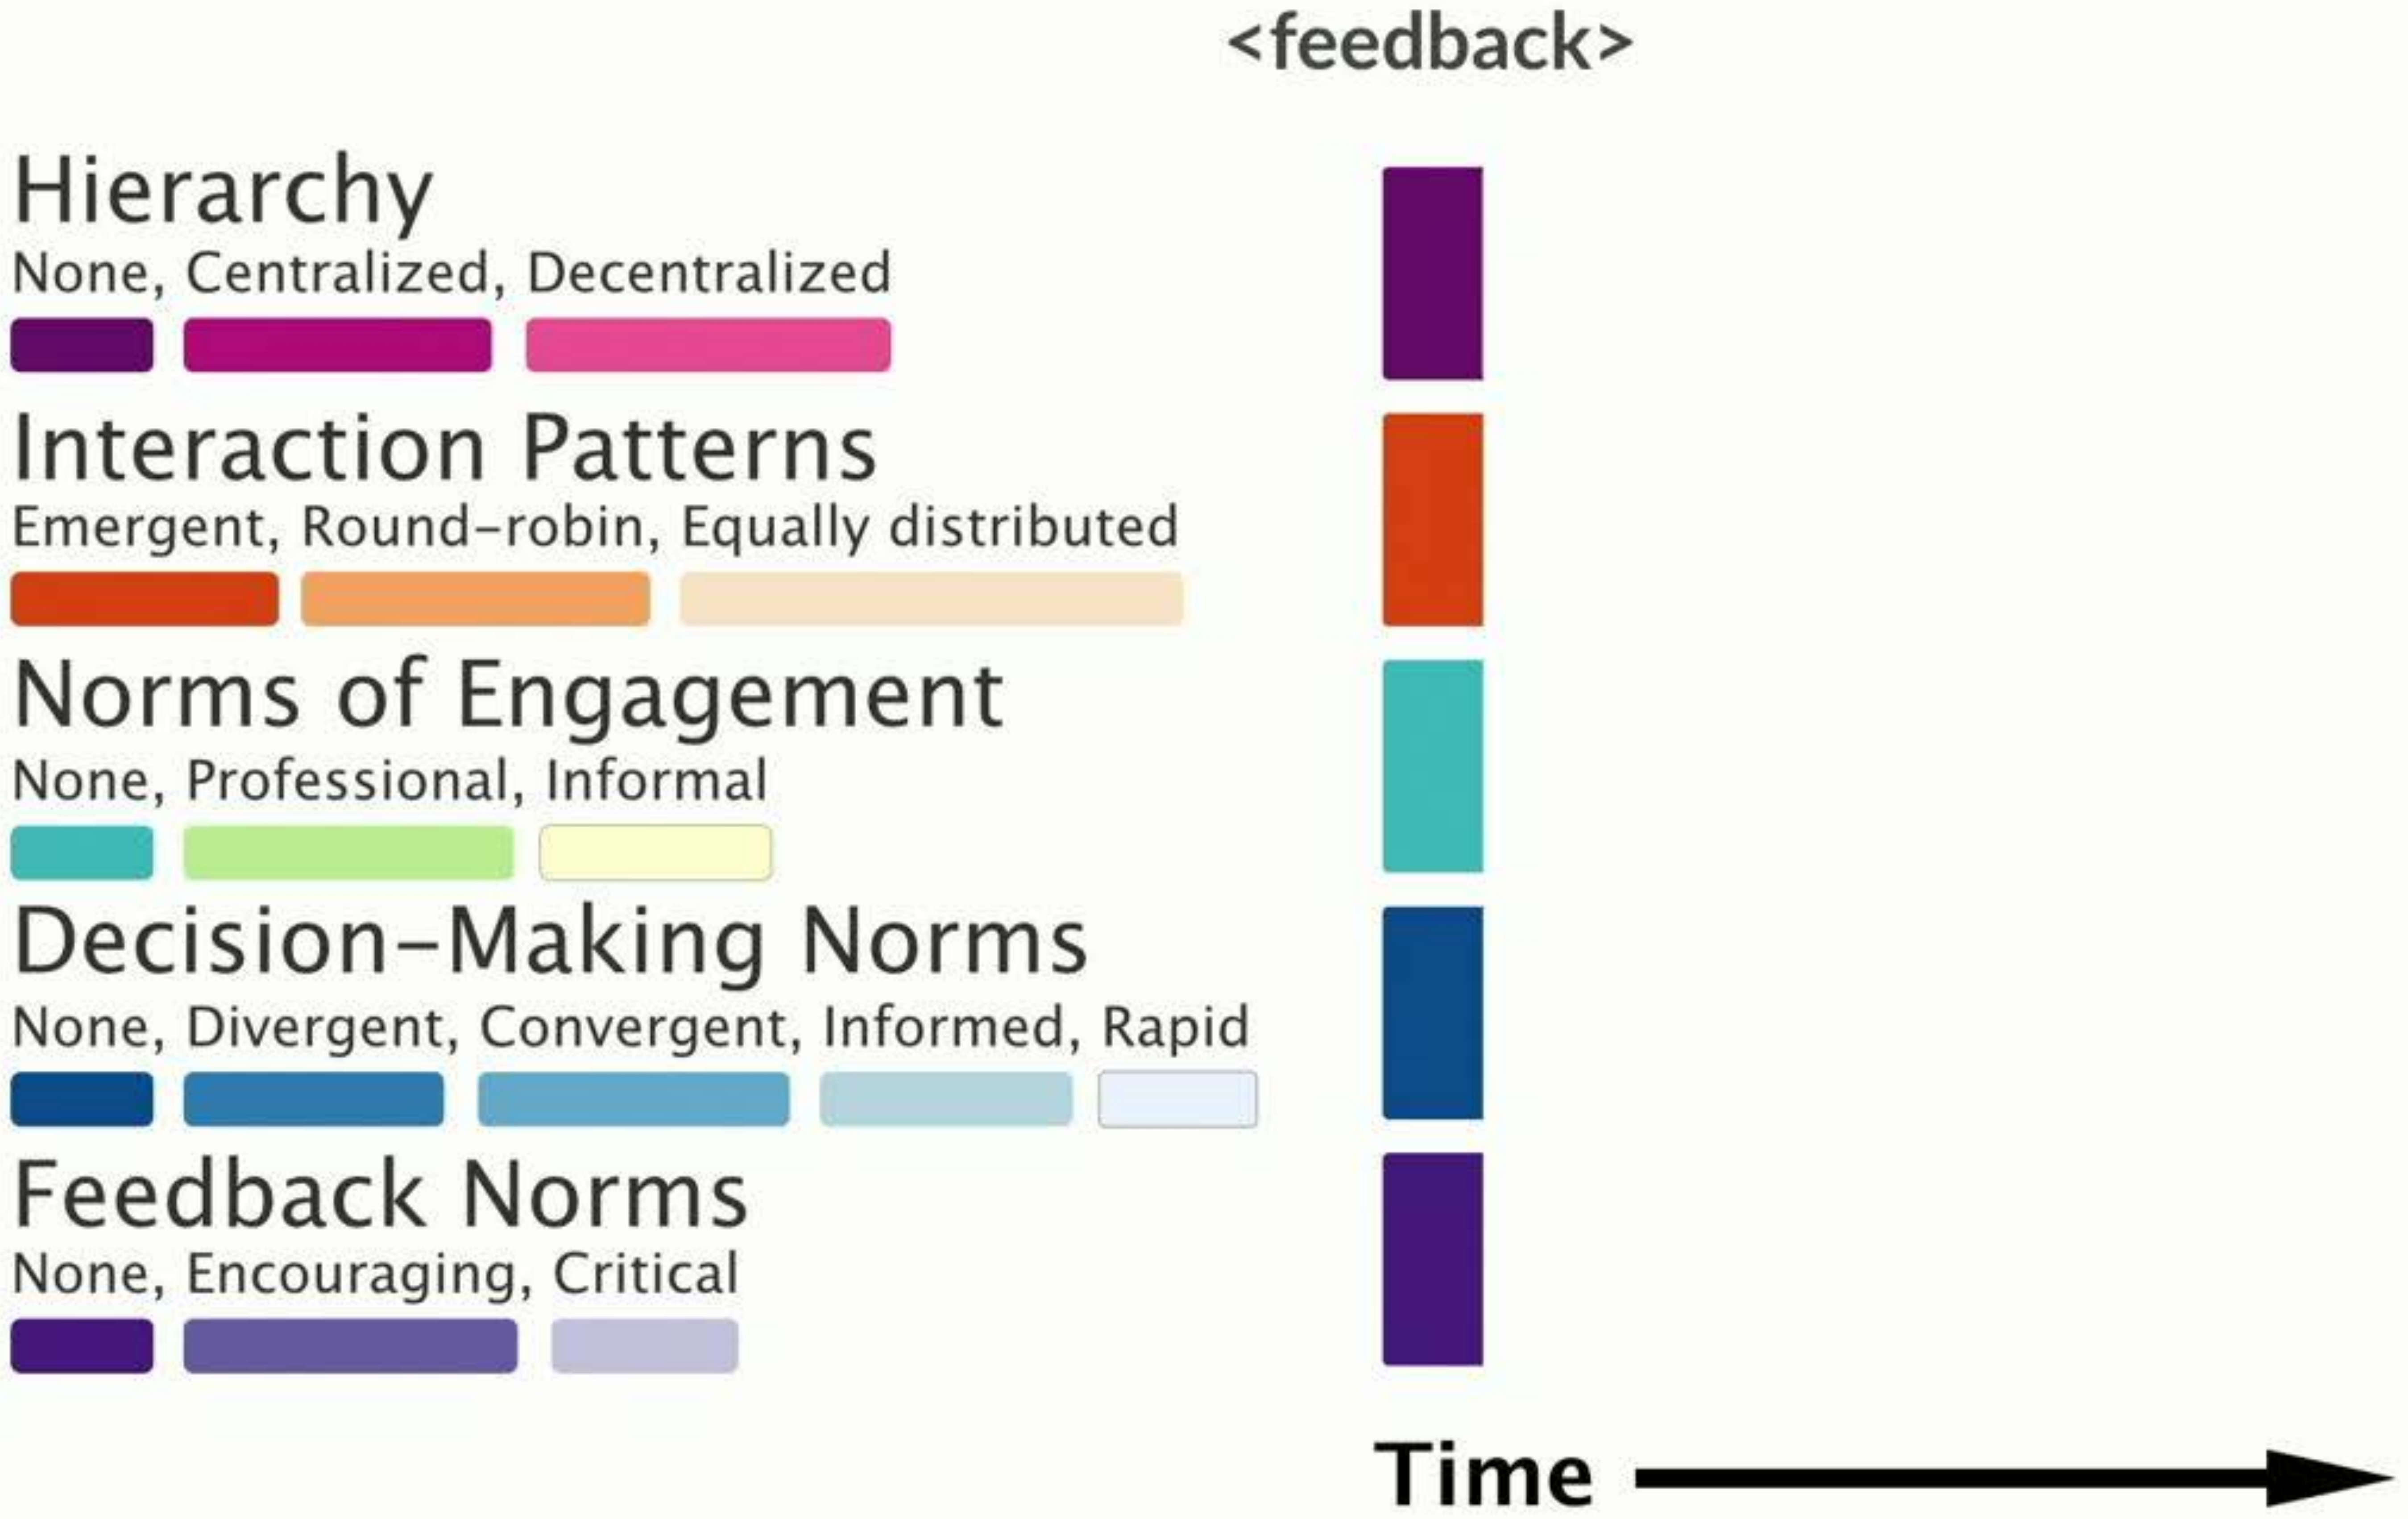
<feedback>
Hierarchy
None, Centralized, Decentralized
Interaction Patterns
Emergent, Round-robin, Equally distributed
Norms of Engagement
None, Professional, Informal
Decision-Making Norms
None, Divergent, Convergent, Informed, Rapid
Feedback Norms
None, Encouraging, Critical
Time

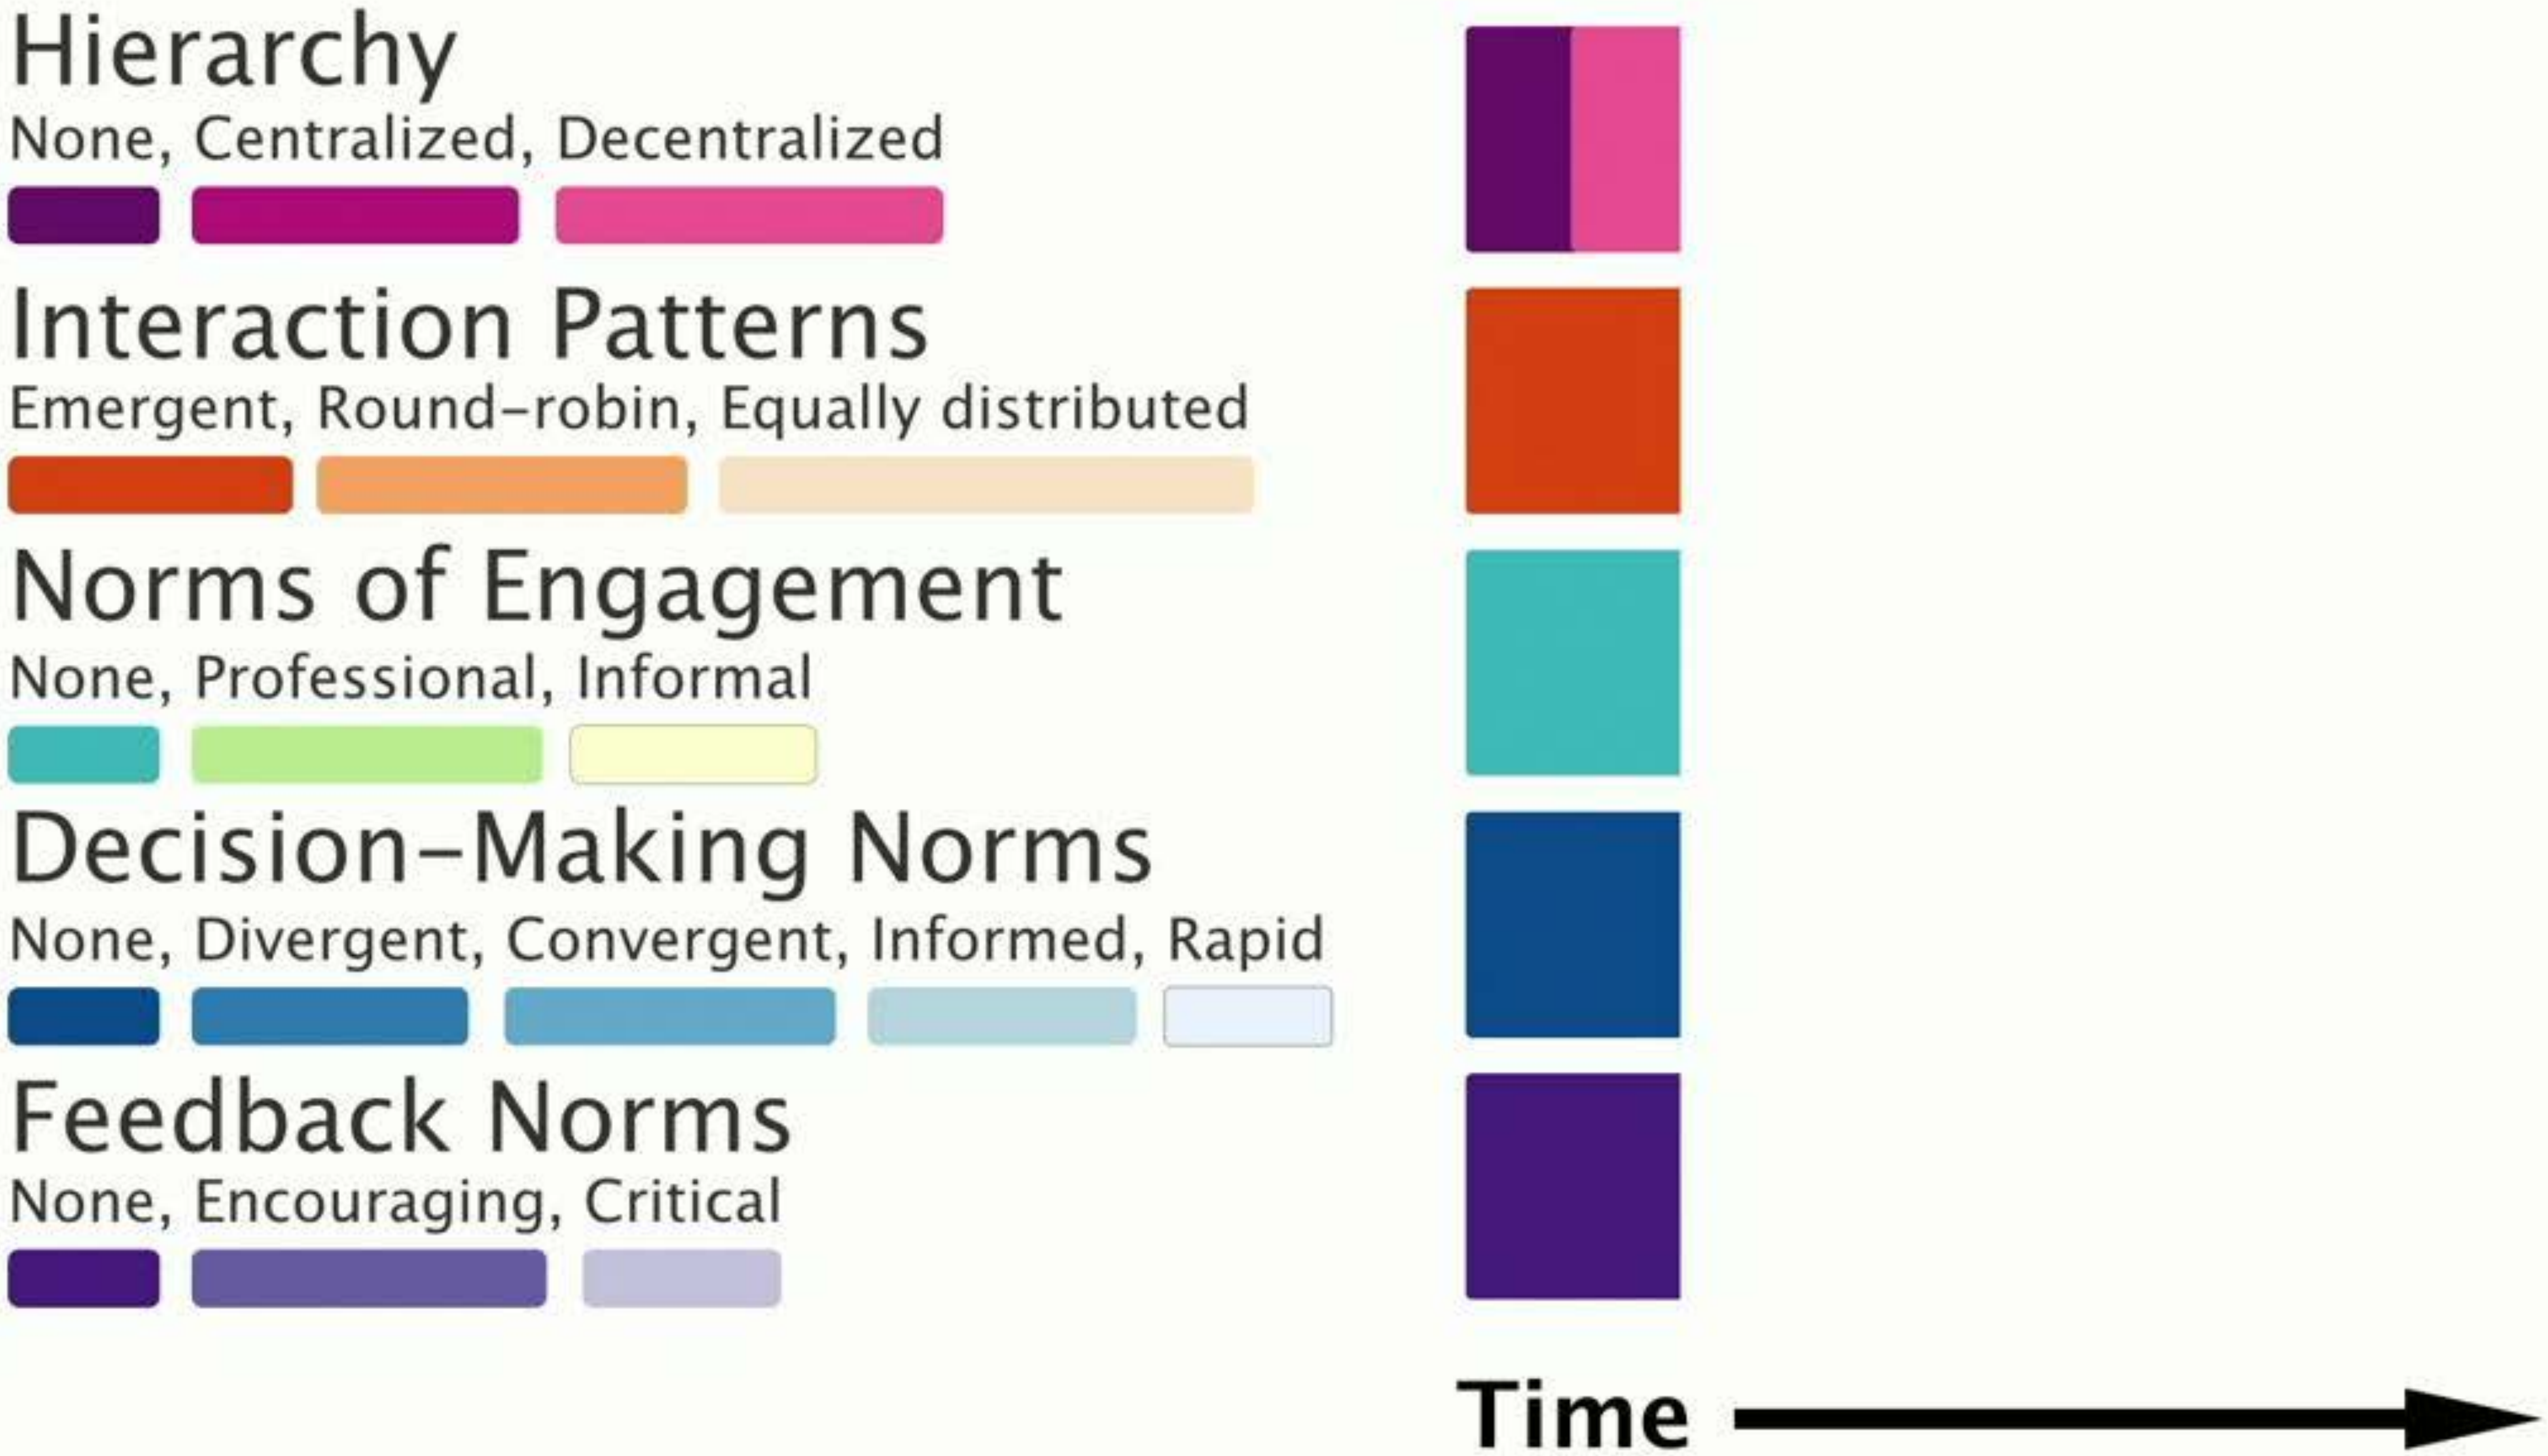
Hierarchy
None, Centralized, Decentralized
Interaction Patterns
Emergent, Round-robin, Equally distributed
Norms of Engagement
None, Professional, Informal
Decision-Making Norms
None, Divergent, Convergent, Informed, Rapid
Feedback Norms
None, Encouraging, Critical
Time

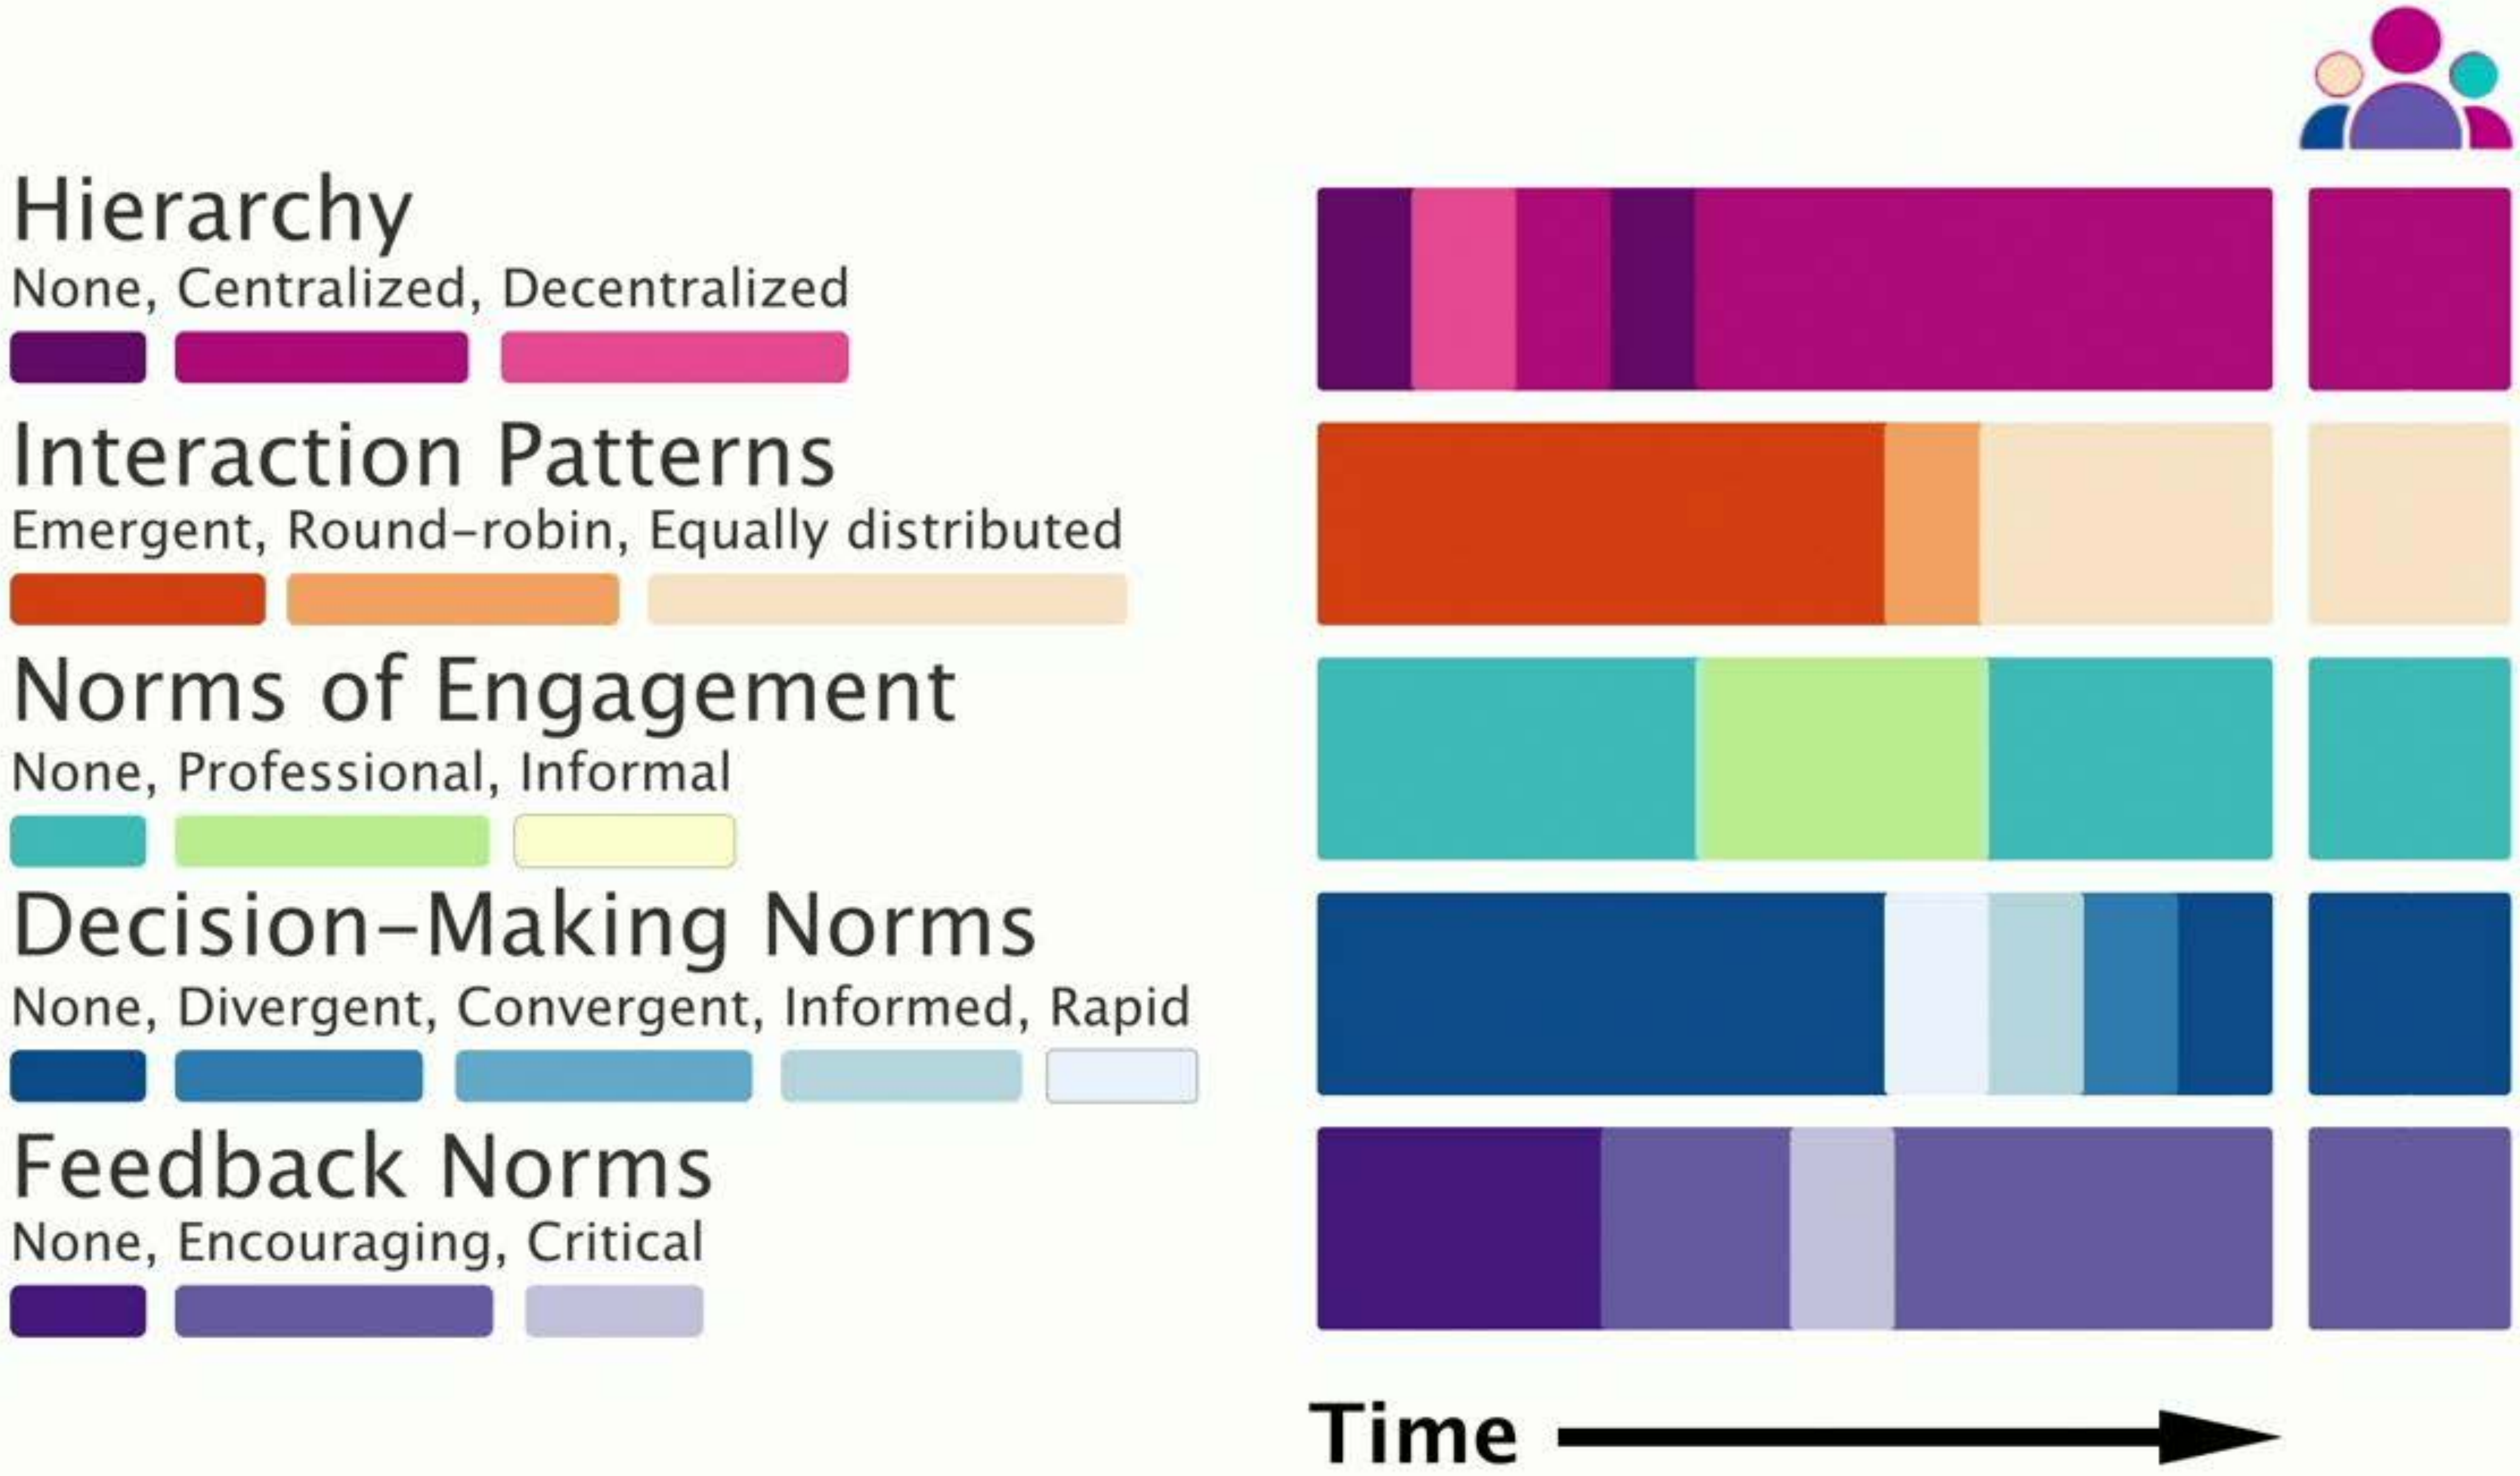
Hierarchy
None, Centralized, Decentralized
Interaction Patterns
Emergent, Round-robin, Equally distributed
Norms of Engagement
None, Professional, Informal
Decision-Making Norms
None, Divergent, Convergent, Informed, Rapid
Feedback Norms
None, Encouraging, Critical
Time

NETWORK OF MULTI-ARMED BANDITS

Multi-armed bandits efficiently explore multiple options over time.

However, this results in so much simultaneous change that **teams become quickly overwhelmed**

NETWORK OF MULTI-ARMED BANDITS

Multi-armed bandits efficiently explore multiple options over time.

However, this results in so much simultaneous change that **teams become quickly overwhelmed**

TEMPORALLY CONSTRAINED BANDITS

Model when teams are open to change, and how much change they are open to simultaneously

e.g., teams are most open to change at the midpoint of their progress [Okhuysen and Waller 2002]

e.g., teams are resilient to exploring hierarchical structures early on, but less resilient to changing them later [Marks, Mathieu, & Zaccaro 2001]

TEMPORALLY CONSTRAINED BANDITS

Model when teams are open to change, and how much change they are open to simultaneously

e.g., teams are most open to change at the midpoint of their progress [Okhuysen and Waller 2002]

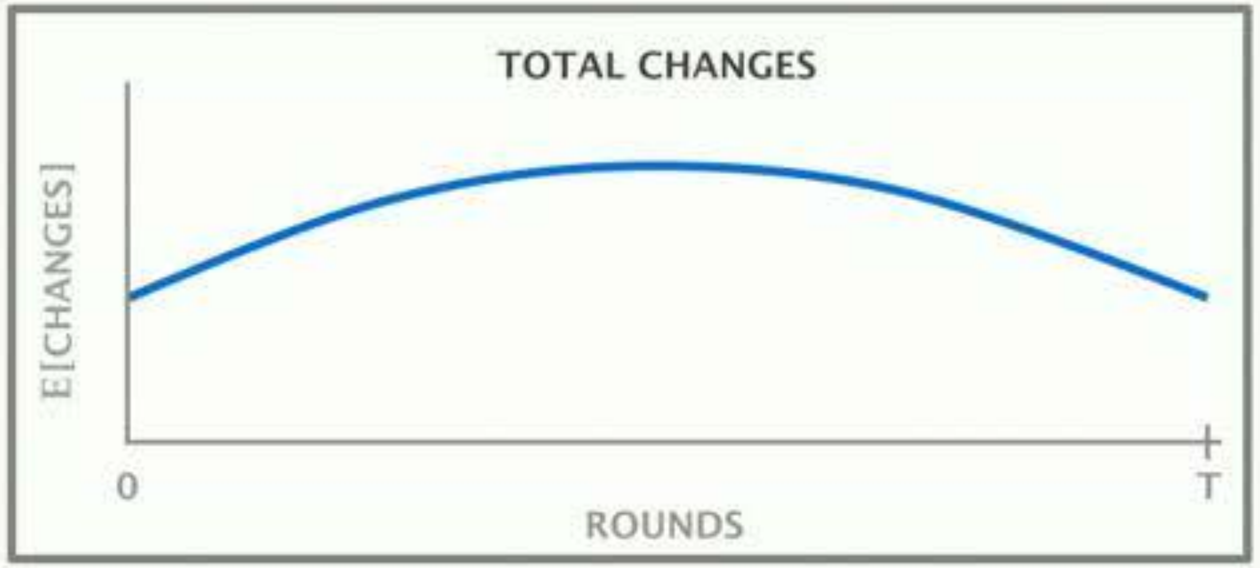

e.g., teams are resilient to exploring hierarchical structures early on, but less resilient to changing them later [Marks, Mathieu, & Zaccaro 2001]

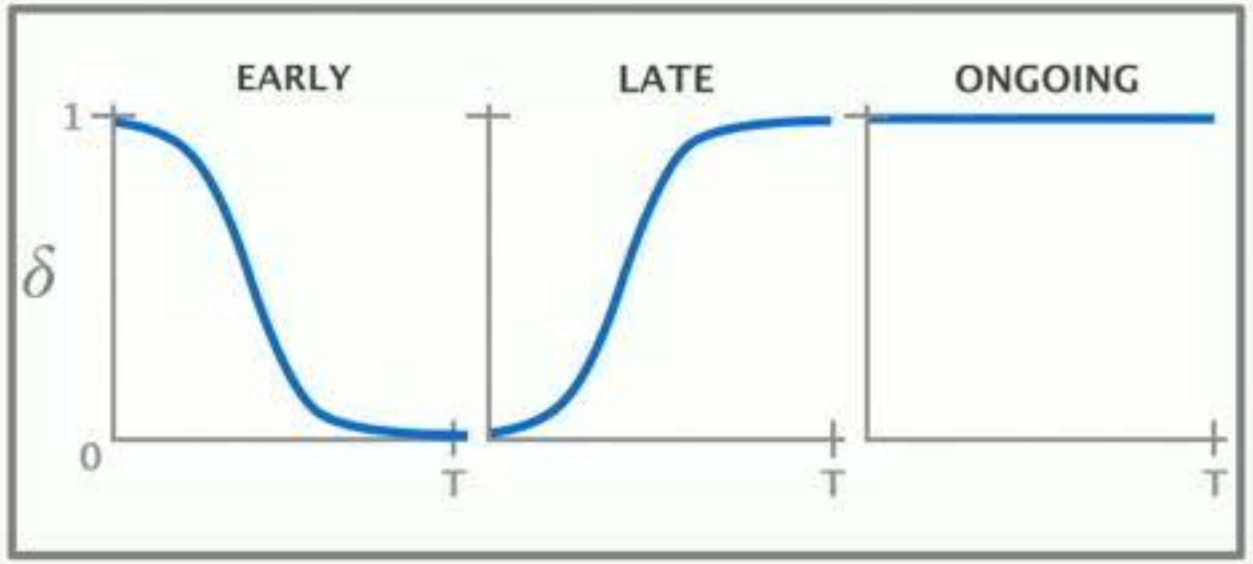

TEMPORALLY CONSTRAINED BANDITS

Redistribute the probability of arm selection via Thompson sampling to respect desired expected value of changes

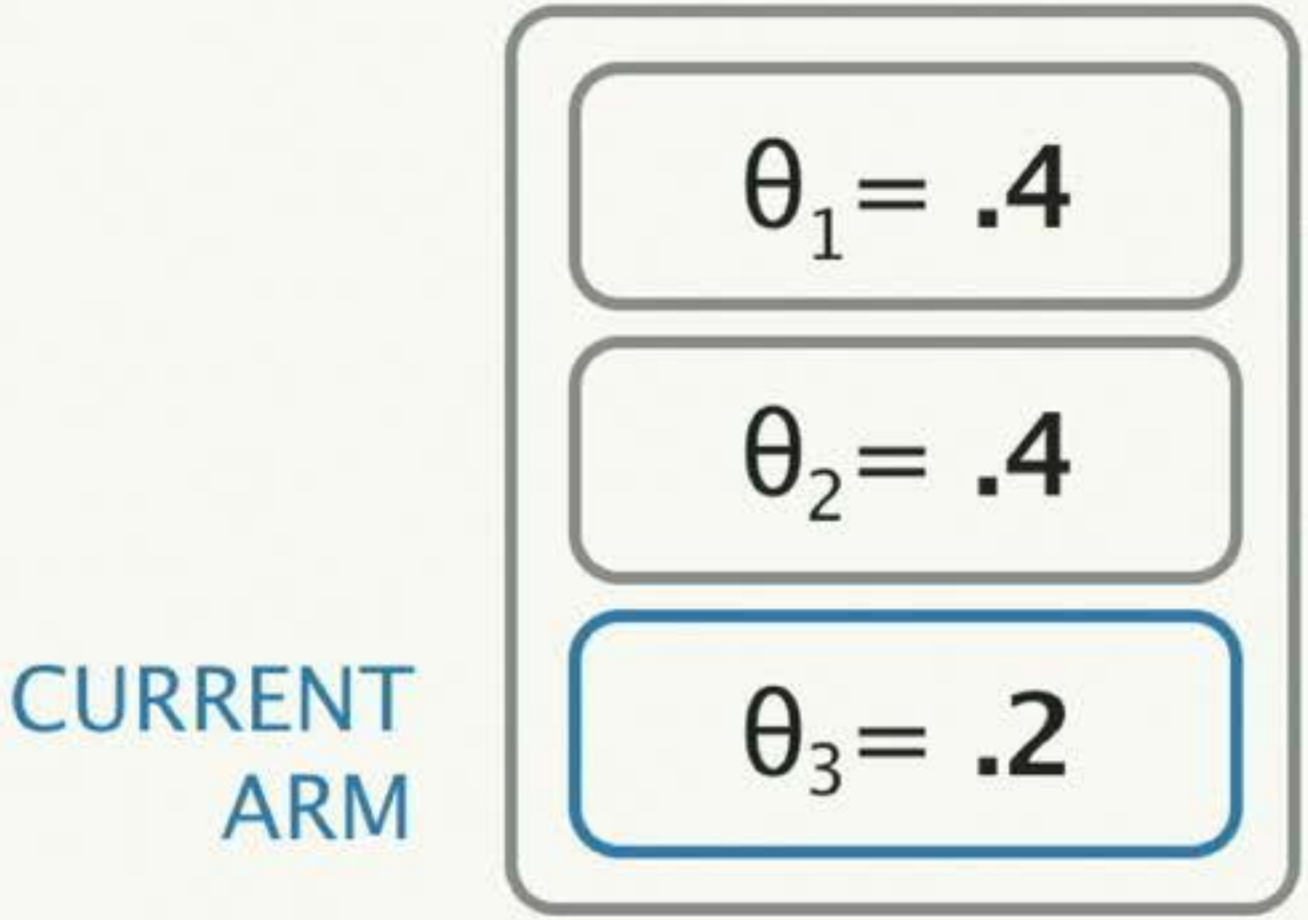

TEMPORALLY CONSTRAINED BANDITS

Redistribute the probability of arm selection via Thompson sampling to respect desired expected value of changes

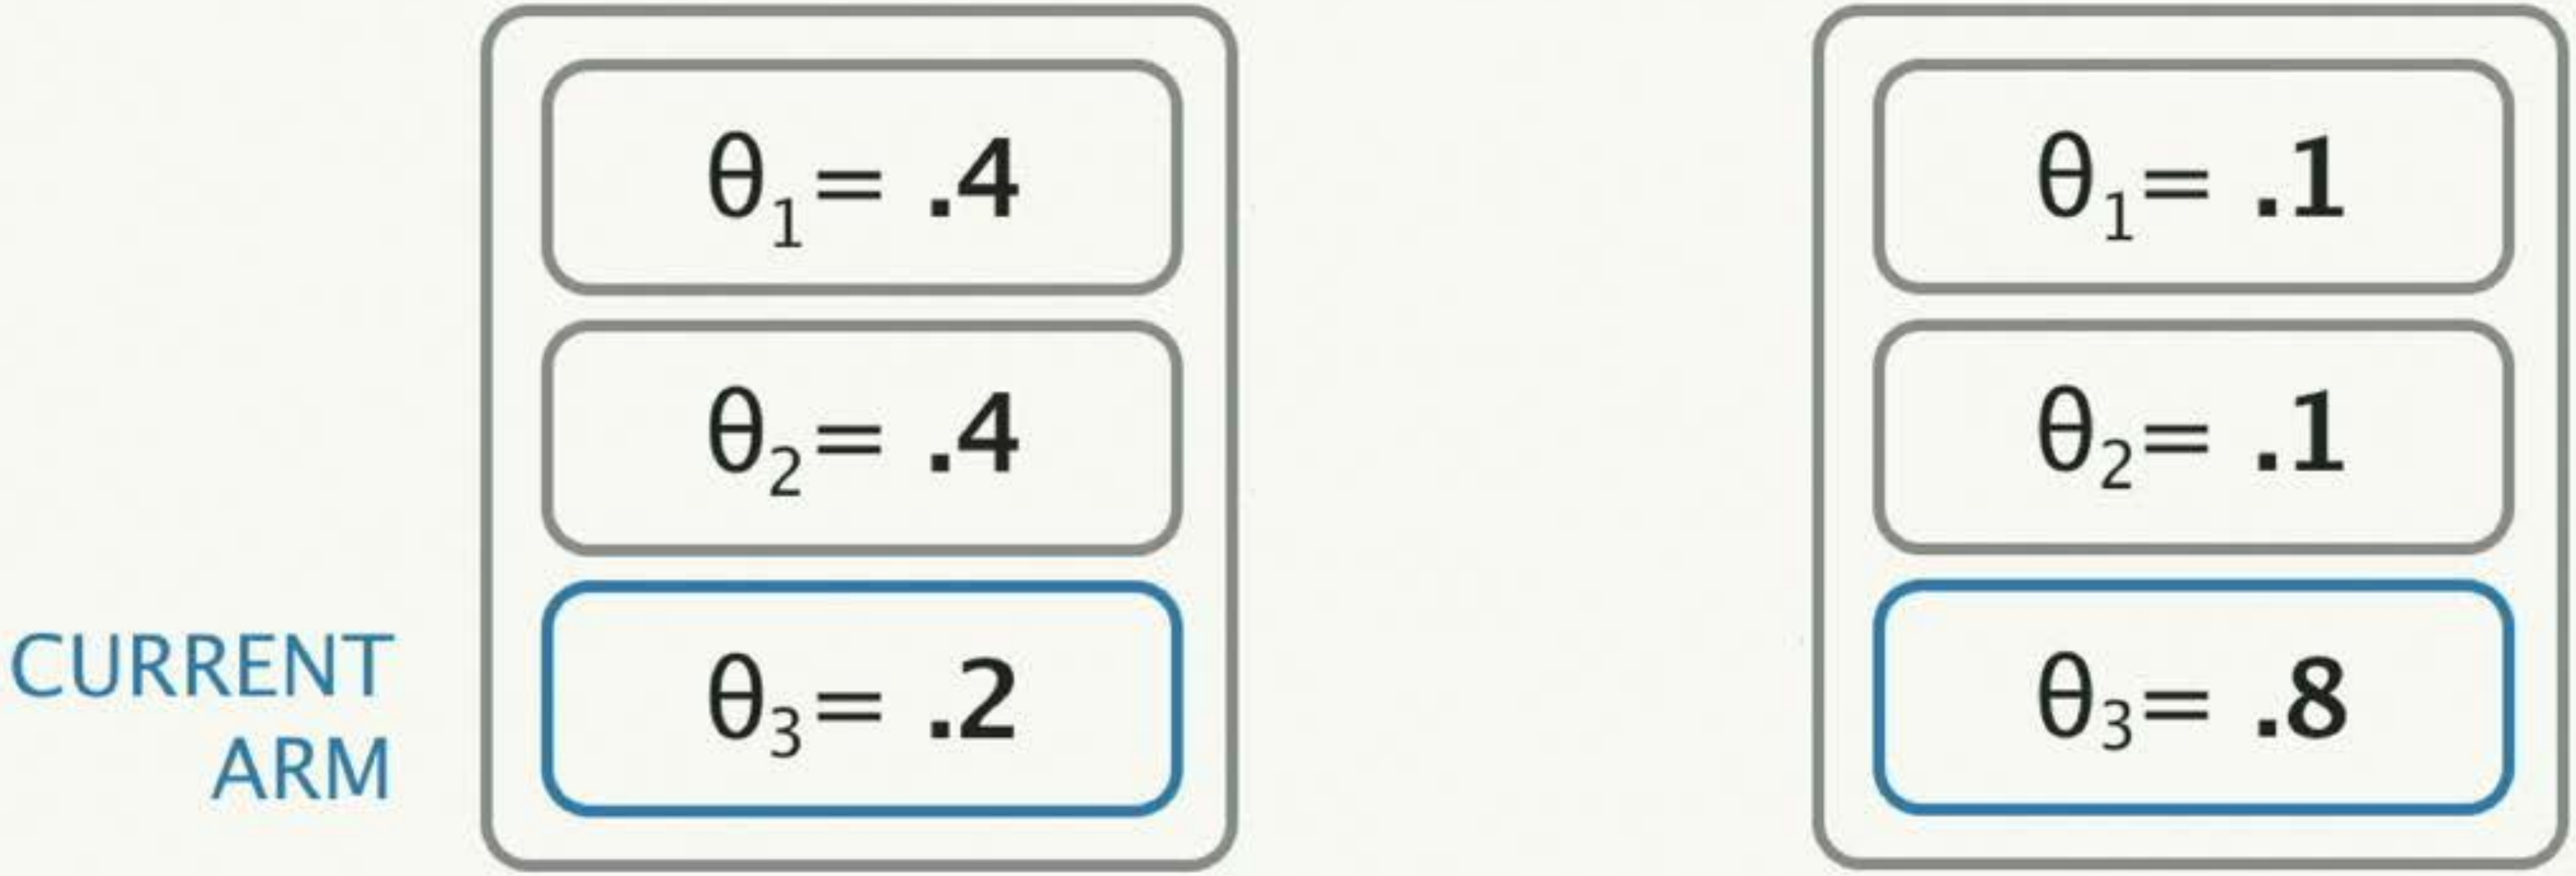

TEMPORALLY CONSTRAINED BANDITS

Redistribute the probability of arm selection via Thompson sampling to respect desired expected value of changes

CURRENT ARM

$\theta_1 = .4$

$\theta_2 = .4$

$\theta_3 = .2$

$\delta = .25$

$\theta_1 = .1$

$\theta_2 = .1$

$\theta_3 = .8$

TEMPORALLY CONSTRAINED BANDITS

Redistribute the probability of arm selection via Thompson sampling to respect desired expected value of changes

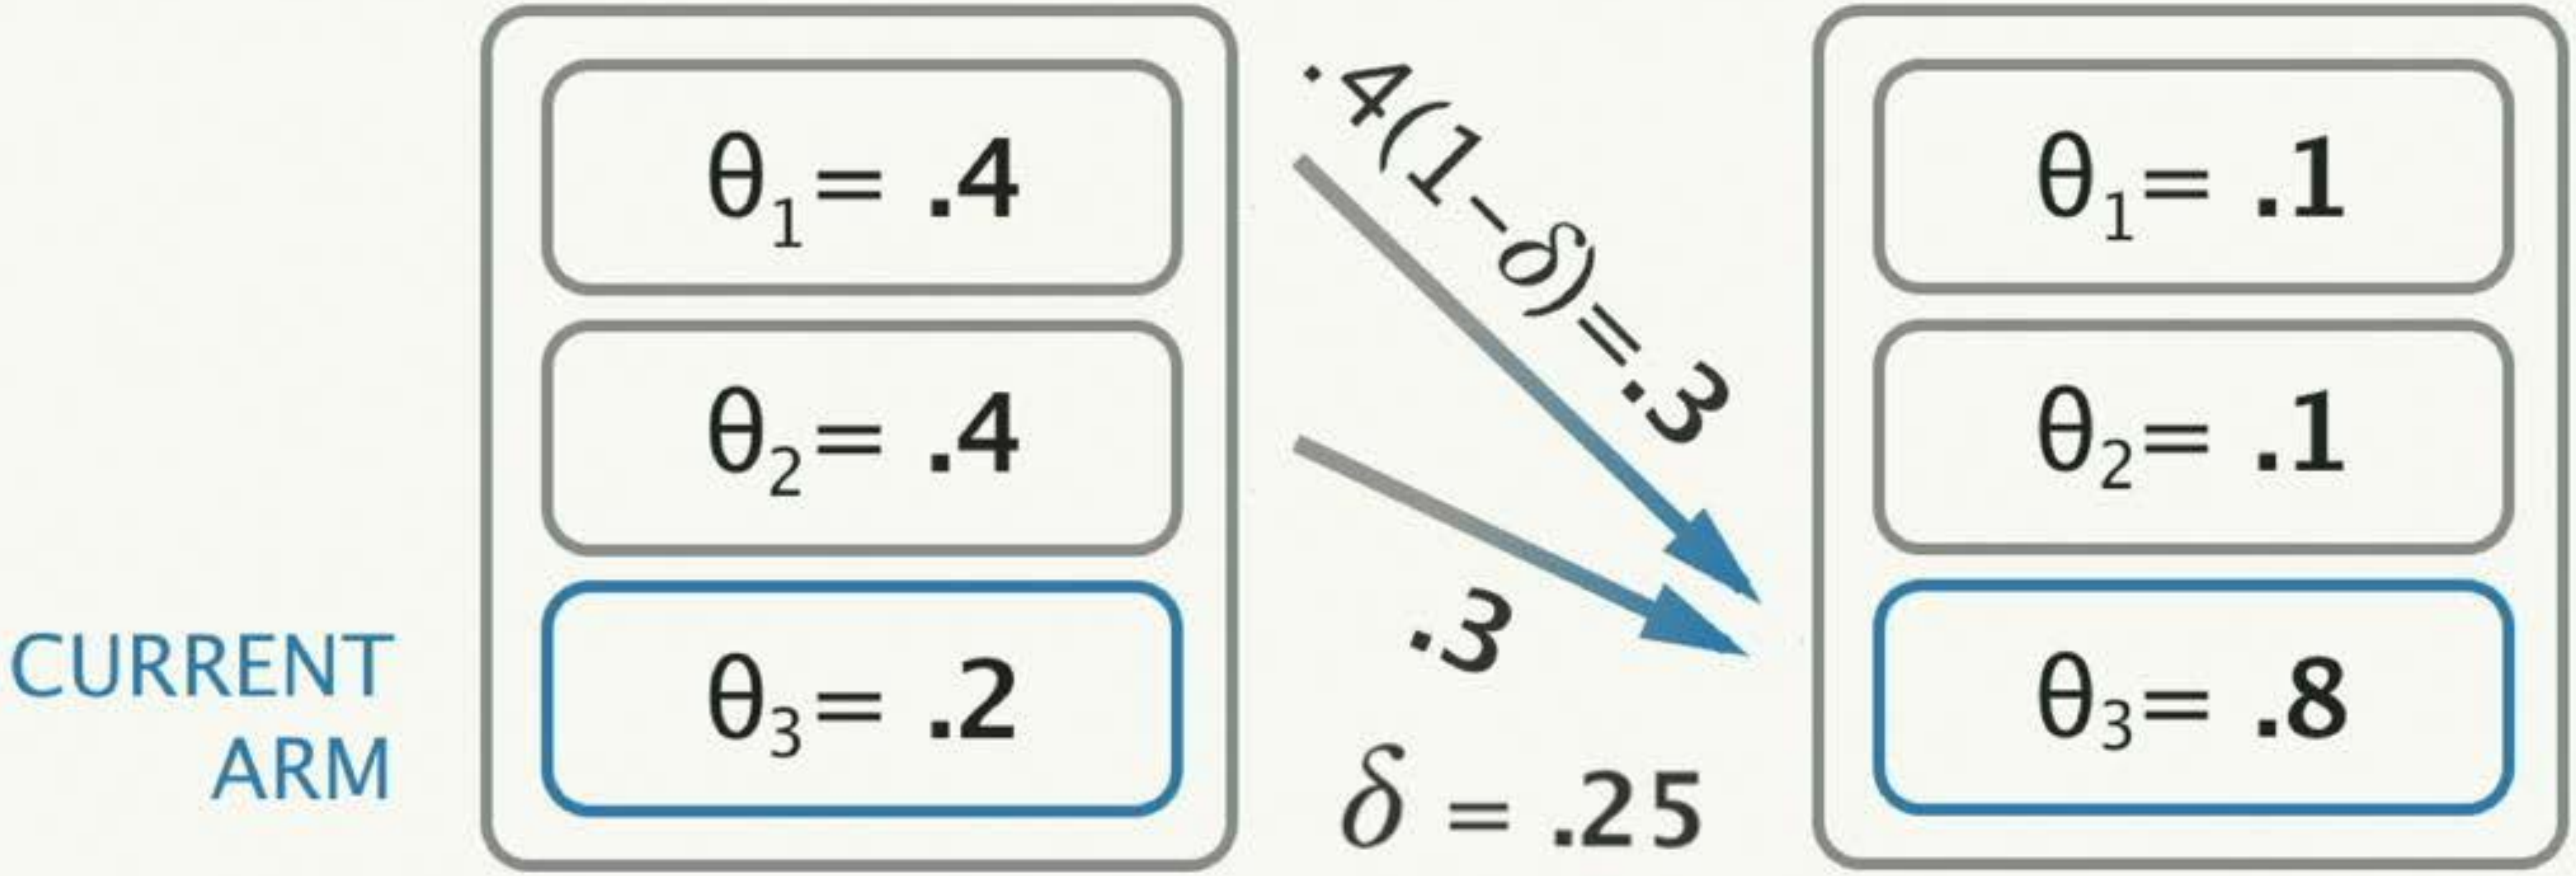

TEMPORALLY CONSTRAINED BANDITS

$$\theta_i' = \begin{cases} \theta_i \delta, \text{ if } i \neq c \\ \theta_i + \sum_{j \neq c} \theta_j (1-\delta), \text{ if } i = c \end{cases}$$

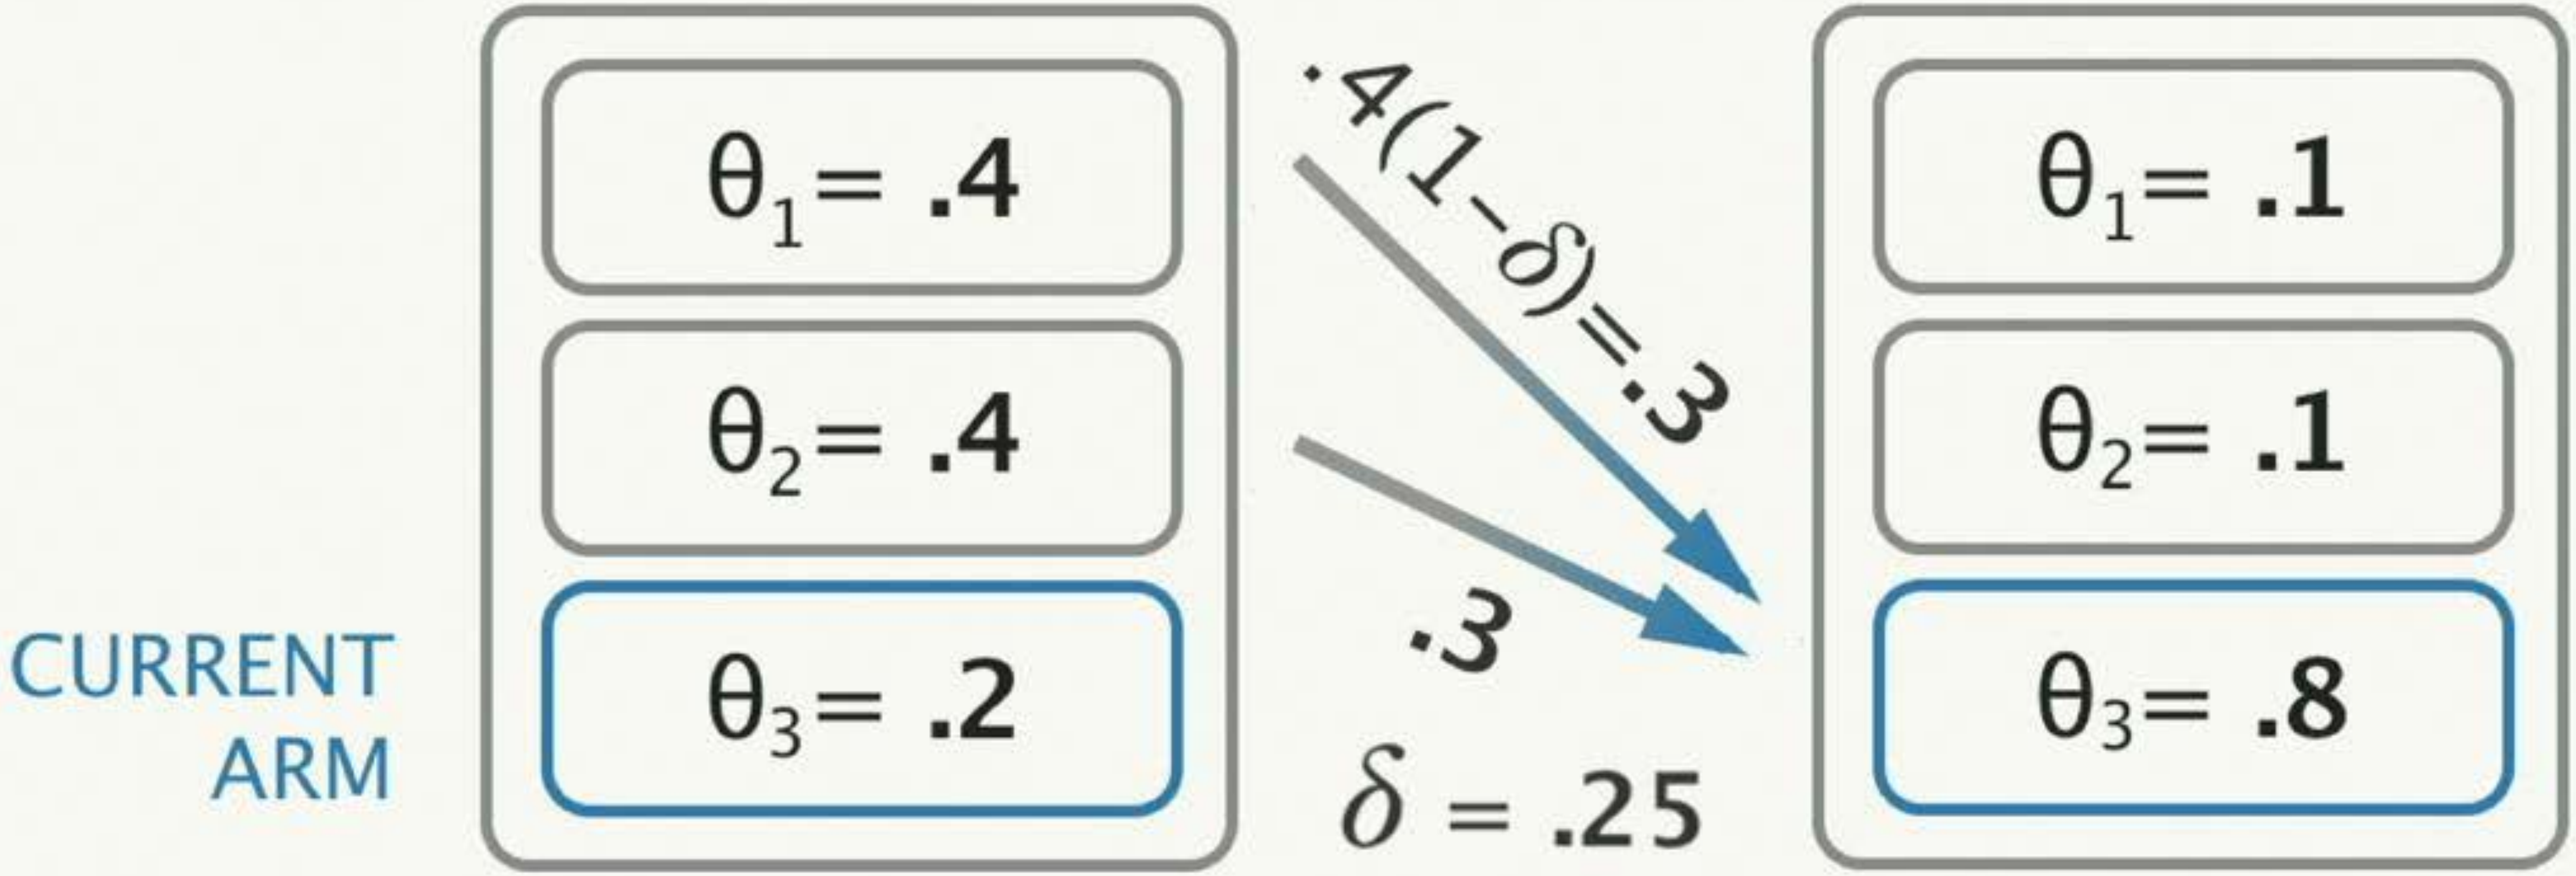

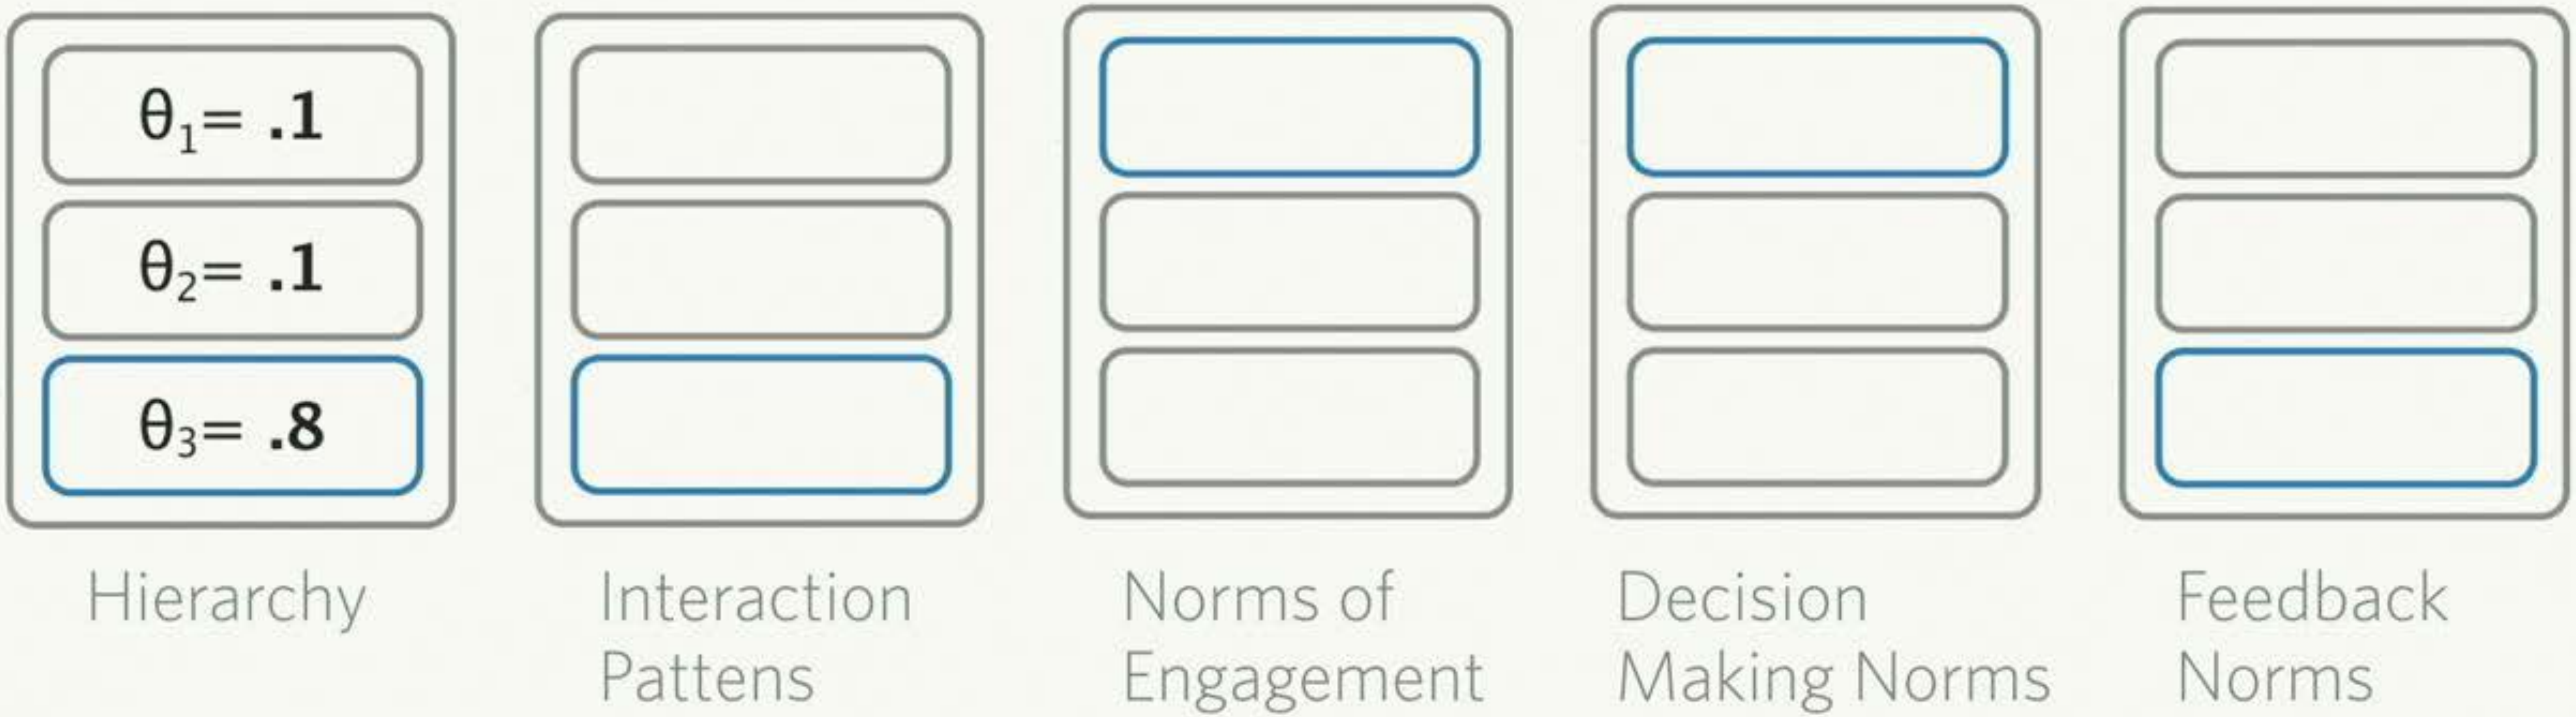

$\theta_1 = .1$
$\theta_2 = .1$
$\theta_3 = .8$
Hierarchy
Interaction Pattens
Norms of Engagement
Decision Making Norms
Feedback Norms

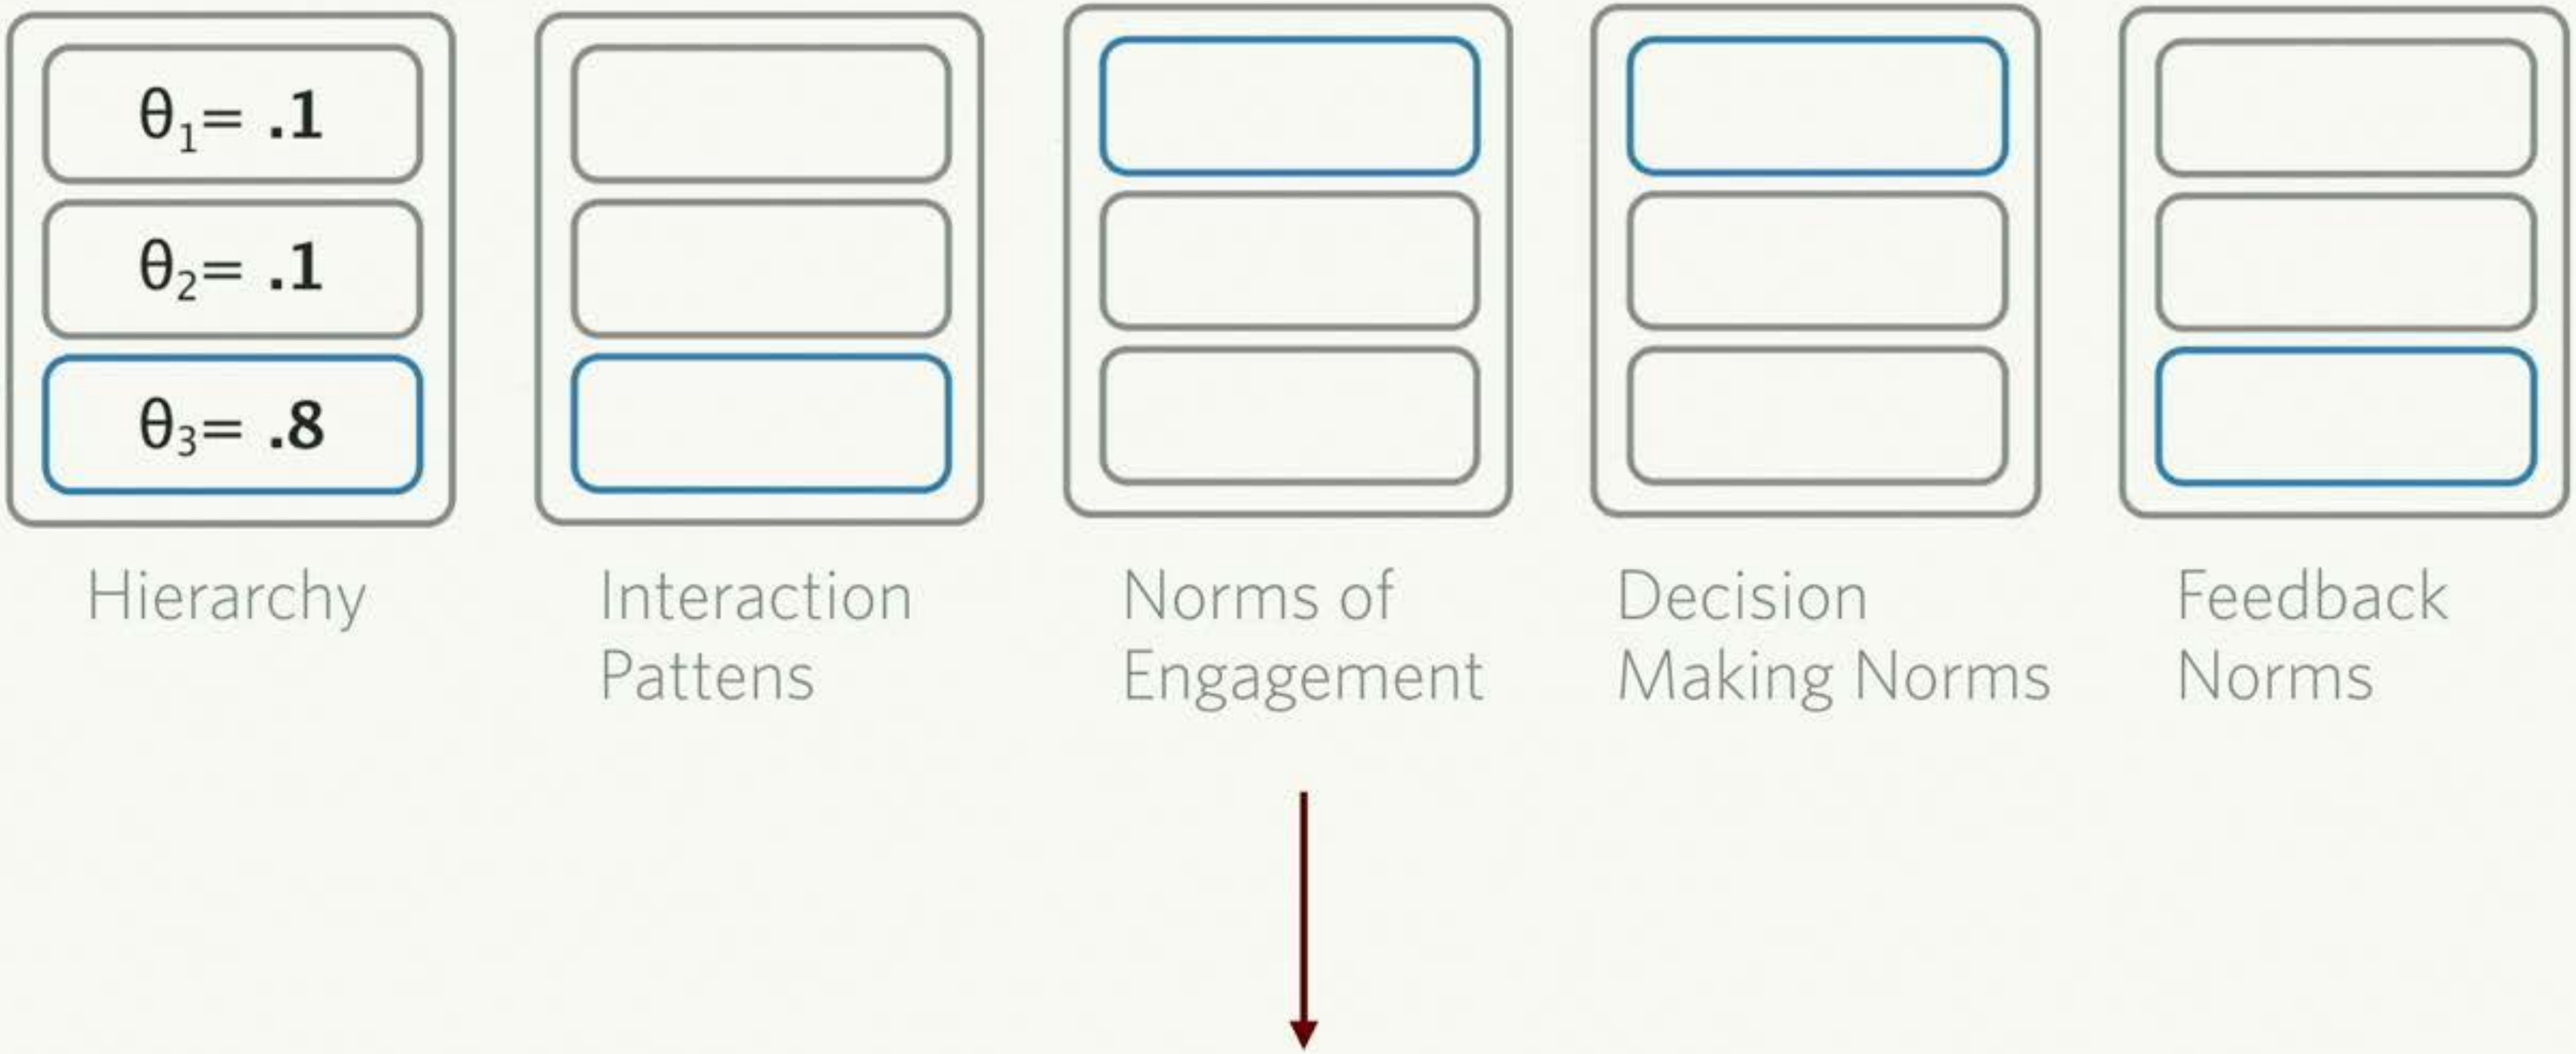

Global constraint on expected number of changes

Prioritize which bandits can change and when

EVALUATION METHOD

135 workers on Mechanical Turk randomized into 45 teams

Measure & bandit reward: team performance on a collaborative intellective task — score on Codewords puzzle

1 training round, 1 baseline round, 10 performance rounds

EVALUATION METHOD

135 workers on Mechanical Turk randomized into 45 teams

Measure & bandit reward: team performance on a collaborative intellective task — score on Codewords puzzle

1 training round, 1 baseline round, 10 performance rounds

Conditions:

1. Manager-chosen
2. Collectively-chosen
3. Control

EVALUATION METHOD

135 workers on Mechanical Turk randomized into 45 teams

Measure & bandit reward: team performance on a collaborative intellective task — score on Codewords puzzle

1 training round, 1 baseline round, 10 performance rounds

Conditions:

1. Manager-chosen
2. Collectively-chosen
3. Control
4. Unconstrained bandit-chosen
5. DreamTeam-chosen

DREAMTEAM TEAMS OUTPERFORM OTHER CONDITIONS BY ~40%

DreamTeam outperformed:

Manager-chosen	by 46%
Collectively-chosen	by 45%
Unconstrained bandit-chosen	by 41%
Control	by 38%

Repeated measures ANCOVA $p<0.05$, all post-hoc Tukey pairwise comparisons to Dreamteam *p<0.05*. N=45. Non-intervention training round used as a covariate to adjust for teams' initial performance.

DREAMTEAM: REFLECTION

The heuristics we use to decide on our team structures can be risk-averse, avoiding fruitful exploration and adaptation

But, raw algorithms overcompensate and overwhelm, leading people to ignore them. Design can help.

Open questions:

- How might we combine voices equitably in the reward feedback?
- Can we adapt when when membership changes? When tasks change? Over the long term? In traditional organizations?

DREAMTEAM TEAMS OUTPERFORM OTHER CONDITIONS BY ~40%

DreamTeam outperformed:

Manager-chosen	by 46%
Collectively-chosen	by 45%
Unconstrained bandit-chosen	by 41%
Control	by 38%

Repeated measures ANCOVA $p<0.05$, all post-hoc Tukey pairwise comparisons to Dreamteam *$p<0.05$*. N=45. Non-intervention training round used as a covariate to adjust for teams' initial performance.

DREAMTEAM: REFLECTION

The heuristics we use to decide on our team structures can be risk-averse, avoiding fruitful exploration and adaptation

But, raw algorithms overcompensate and overwhelm, leading people to ignore them. Design can help.

Open questions:

- How might we combine voices equitably in the reward feedback?
- Can we adapt when when membership changes? When tasks change? Over the long term? In traditional organizations?

Crowd research

Vaish, Gaikwad, Veit, Krishna, Ibarra, Simoiu, Wilber, Belongie, Davis, Goel, Bernstein. UIST 2017.

RESEARCH: THE DOMAIN OF THE PRIVILEGED FEW

Those able to attend prestigious universities can access research experiences that support open-ended inquiry and launch careers [Russell et al. 2007]

...but the vast majority of people cannot [Bowen and Bok 2016; Bianchini 2011]

Top 50 global universities, US News 2017

RESEARCH: THE DOMAIN OF THE PRIVILEGED FEW

Those able to attend prestigious universities can access research experiences that support open-ended inquiry and launch careers [Russell et al. 2007]

...but the vast majority of people cannot [Bowen and Bok 2016; Bianchini 2011]

Top 50 global universities, US News 2017

RESEARCH: THE DOMAIN OF THE PRIVILEGED FEW

Those able to attend prestigious universities can access research experiences that support open-ended inquiry and launch careers [Russell et al. 2007]

...but the vast majority of people cannot [Bowen and Bok 2016; Bianchini 2011]

A research ecosystem that **under-represents minorities and developing regions**, and a literature that overlooks their perspectives

Top 50 global universities, US News 2017

CROWD RESEARCH

A **crowdsourcing technique** enabling a global crowd to work together on an **open-ended research project**

Participants collaborate as one large team to brainstorm, execute and publish the project under the leadership of a PI

GOALS

Enable access to training and research experiences in support of upward career and educational mobility

Convene hundreds or thousands of people on a single ambitious project

WE ARE NOT EQUIPPED FOR LARGE SCALE, OPEN ENDED RESEARCH

Research is not a linear path from idea to result: it is an iterative process of exploration
[Gowers 2000]

In contrast, citizen science efforts today focus on **pre-defined goals** in order to structure the crowd's contributions

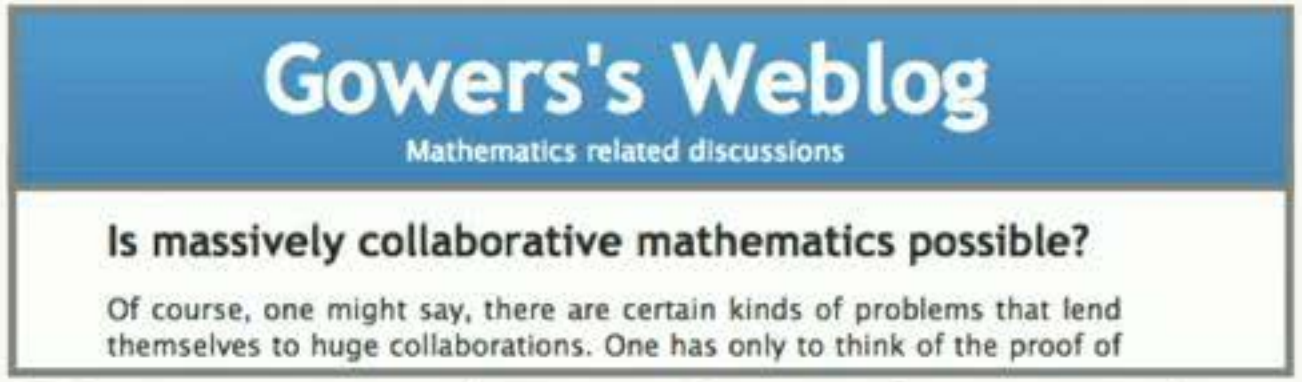

[Gowers and Tao; Cranshaw and Kittur 2011]

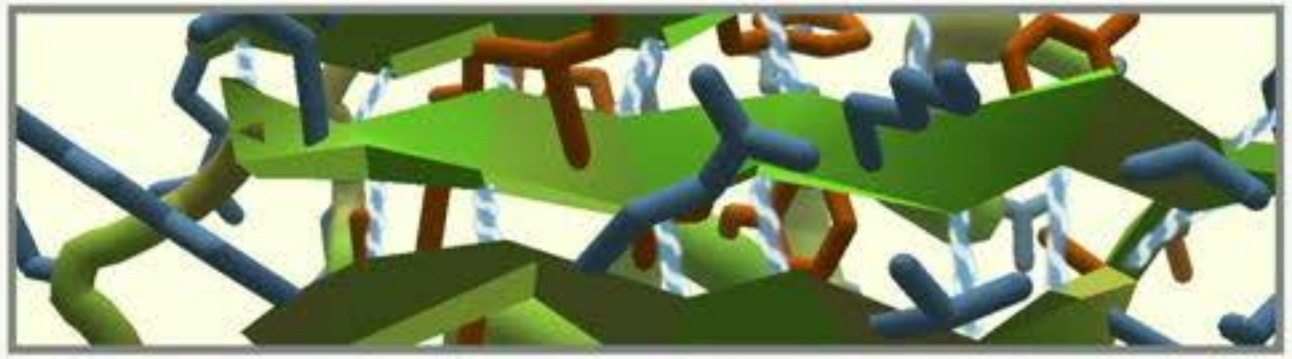

[Cooper et al. 2010]

PROBLEMS

Coordination:
How do we prevent the project from moving in 1,000 directions at once, across easily 6,000 messages per week?

Credit:
How can we provide proof that participants made substantial contributions to the project, when no one central authority can assert this?

CROWD RESEARCH

Iterative crowdsourcing technique:
Weekly cycle of open contribution, synchronous collaboration, and peer assessment

Decentralized credit:
Participants allocate finite credits to each other, enabling a graph centrality algorithm to determine credit and author order

THREE PARALLEL PROJECTS

HCI

Michael Bernstein, Stanford

Building a new crowd marketplace

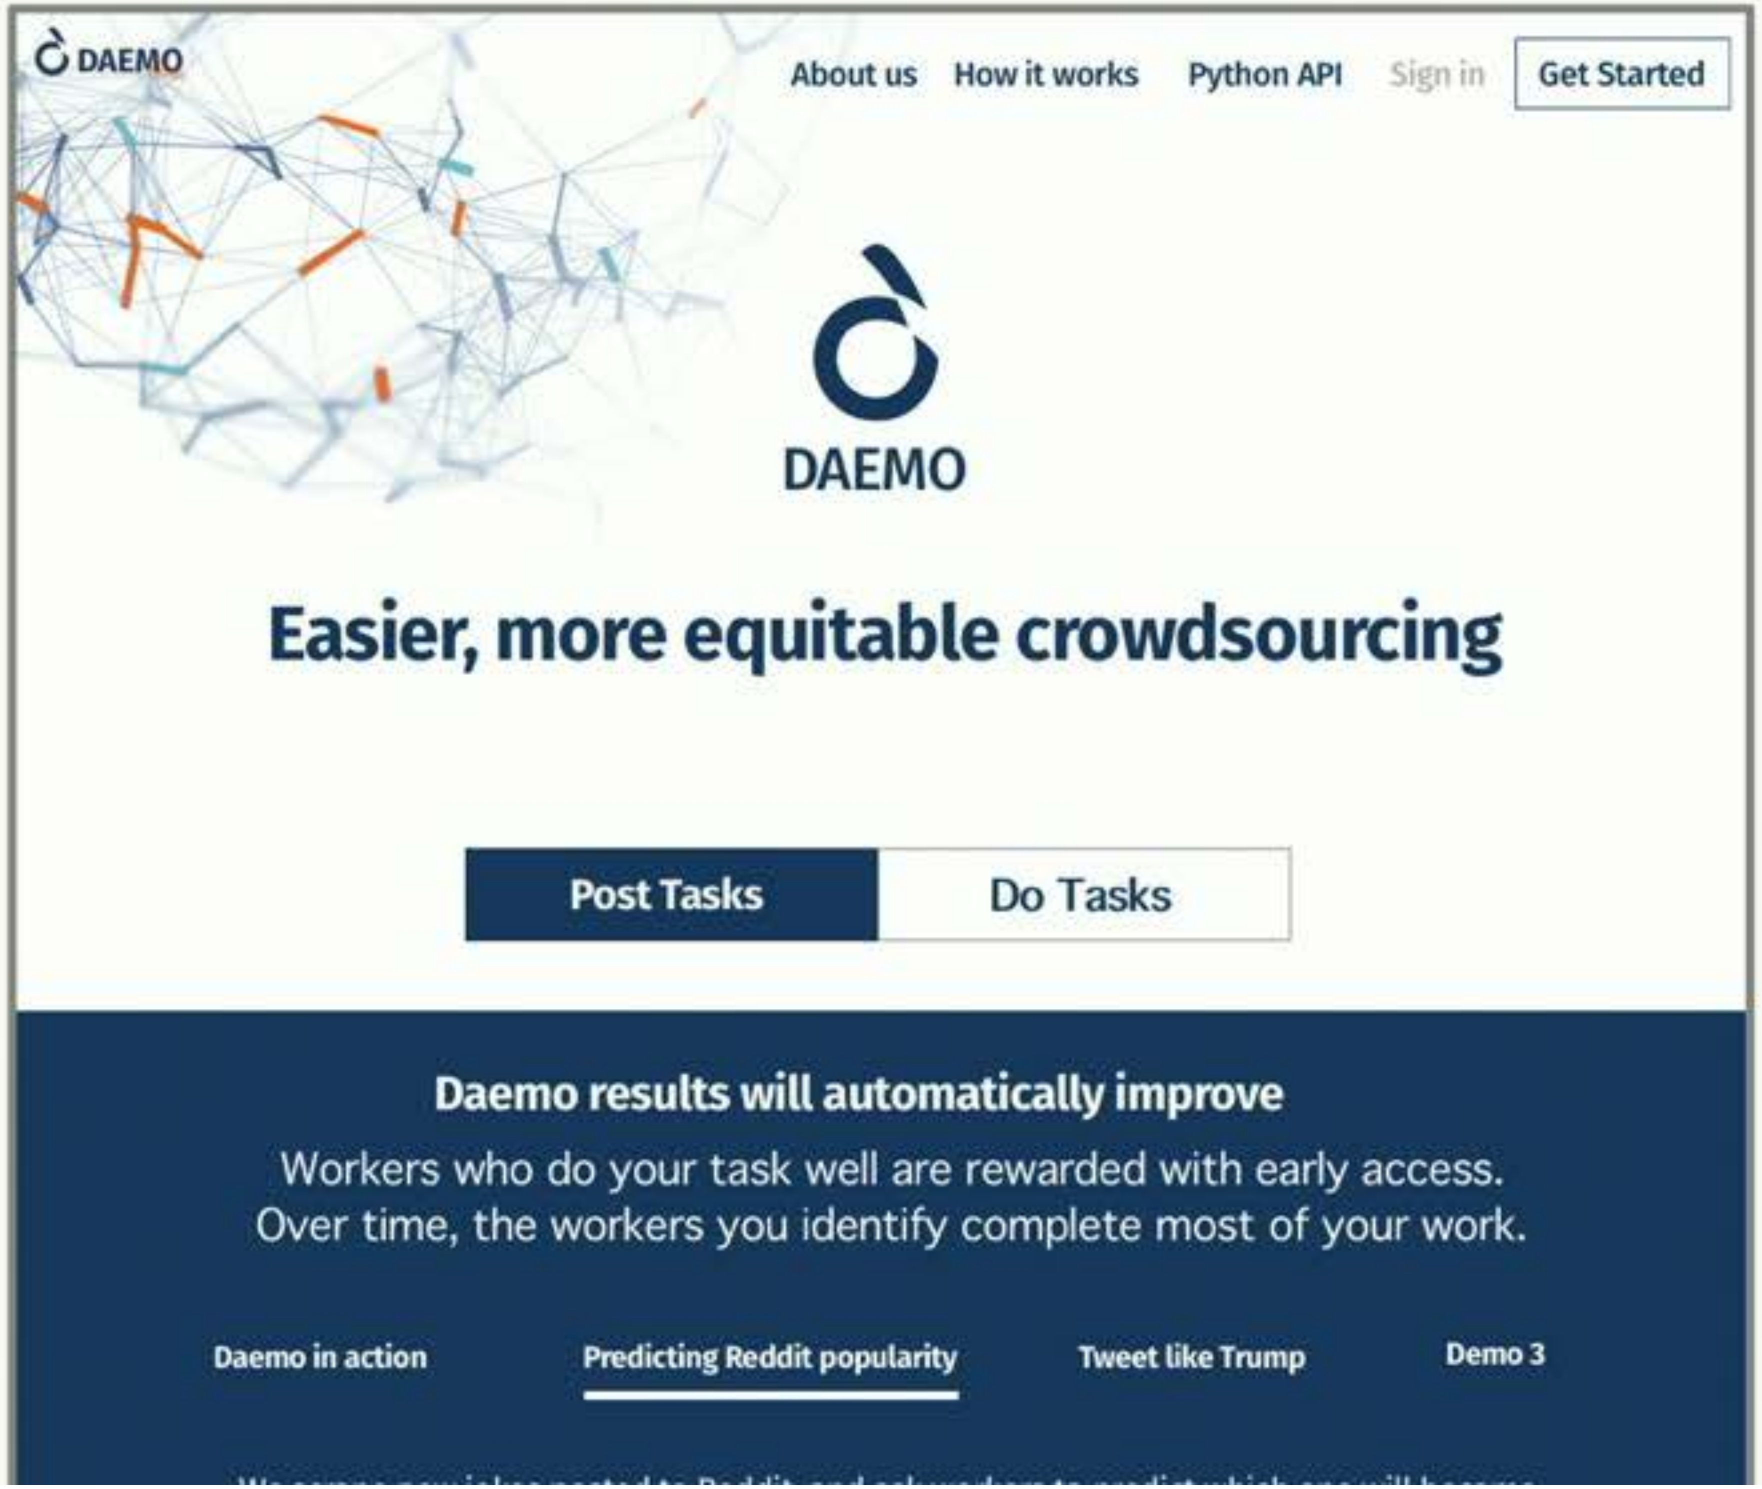

THREE PARALLEL PROJECTS

Computer vision

James Davis, UCSC
Serge Belongie, Cornell Tech

Hybrid crowd-computer vision algorithms

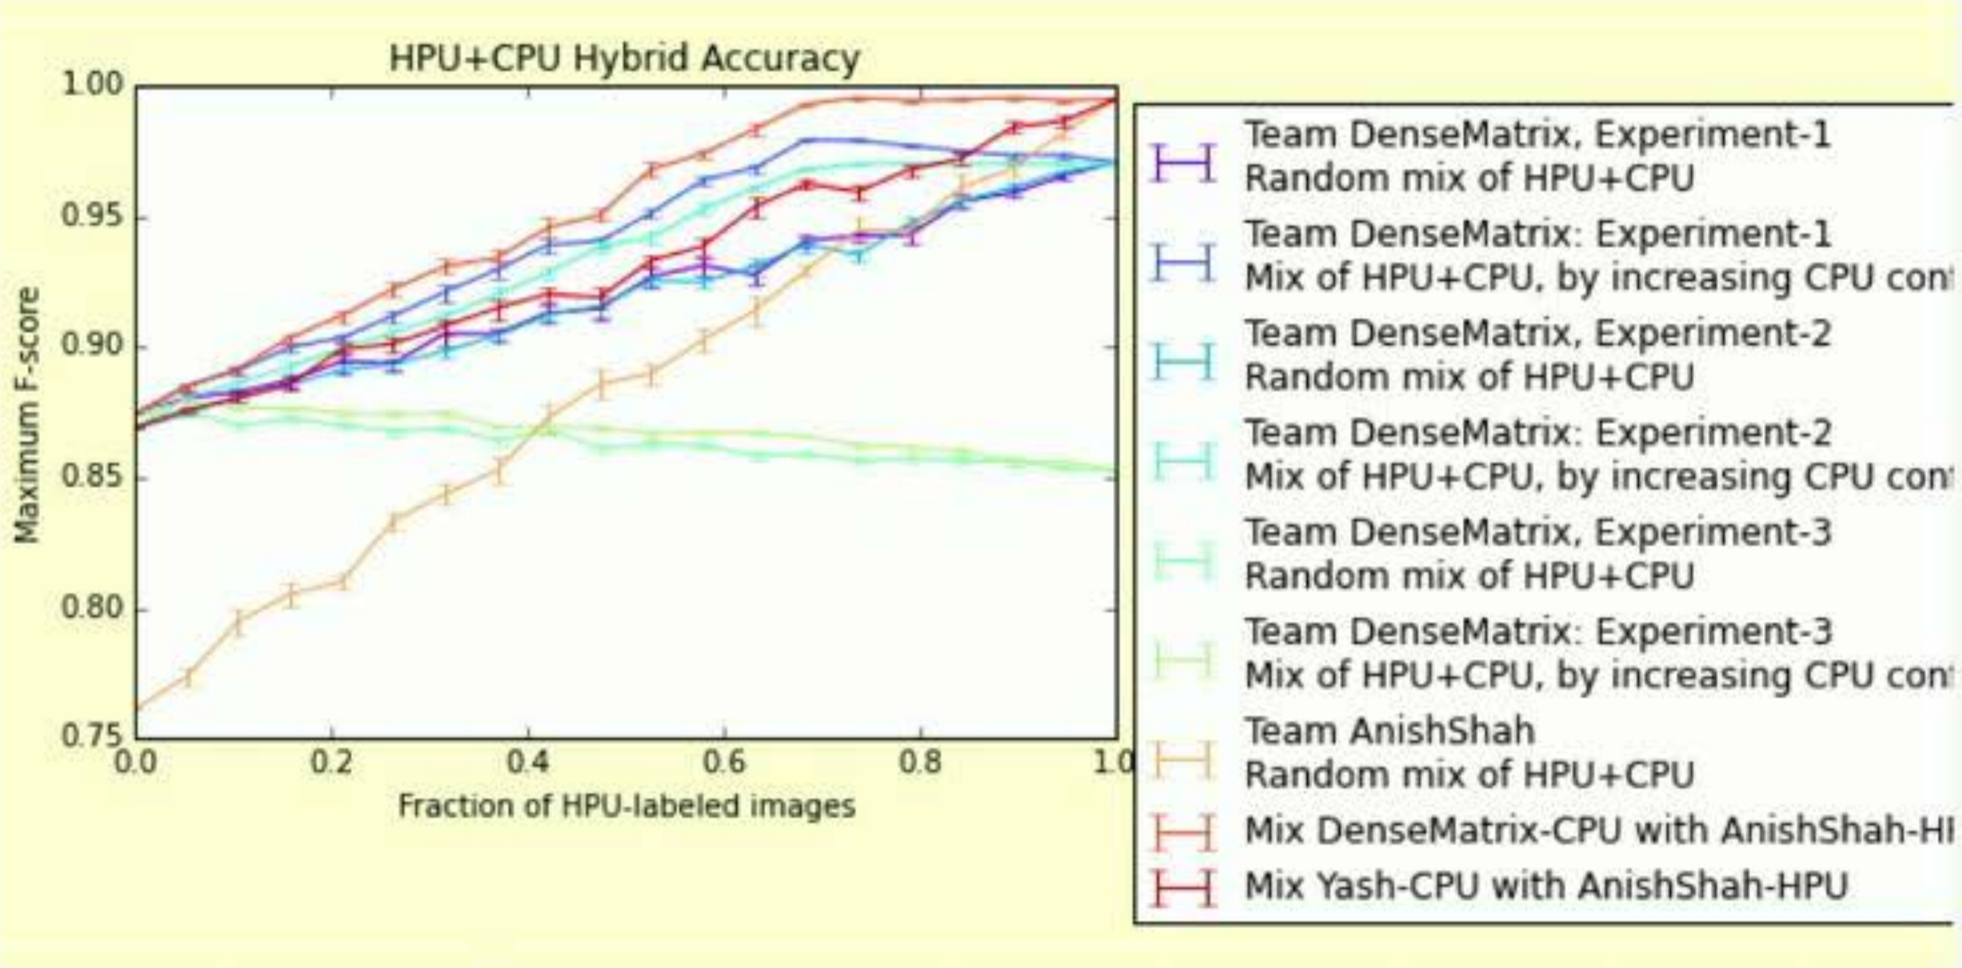

THREE PARALLEL PROJECTS

Data science

Sharad Goel, Stanford

Hundreds of experiments testing the wisdom of the crowd

CROWDSOURCING PROCESS

open call

CROWDSOURCING PROCESS

open call → group meeting

CROWDSOURCING PROCESS

open call → group meeting ↔ milestone deadline

CROWDSOURCING PROCESS

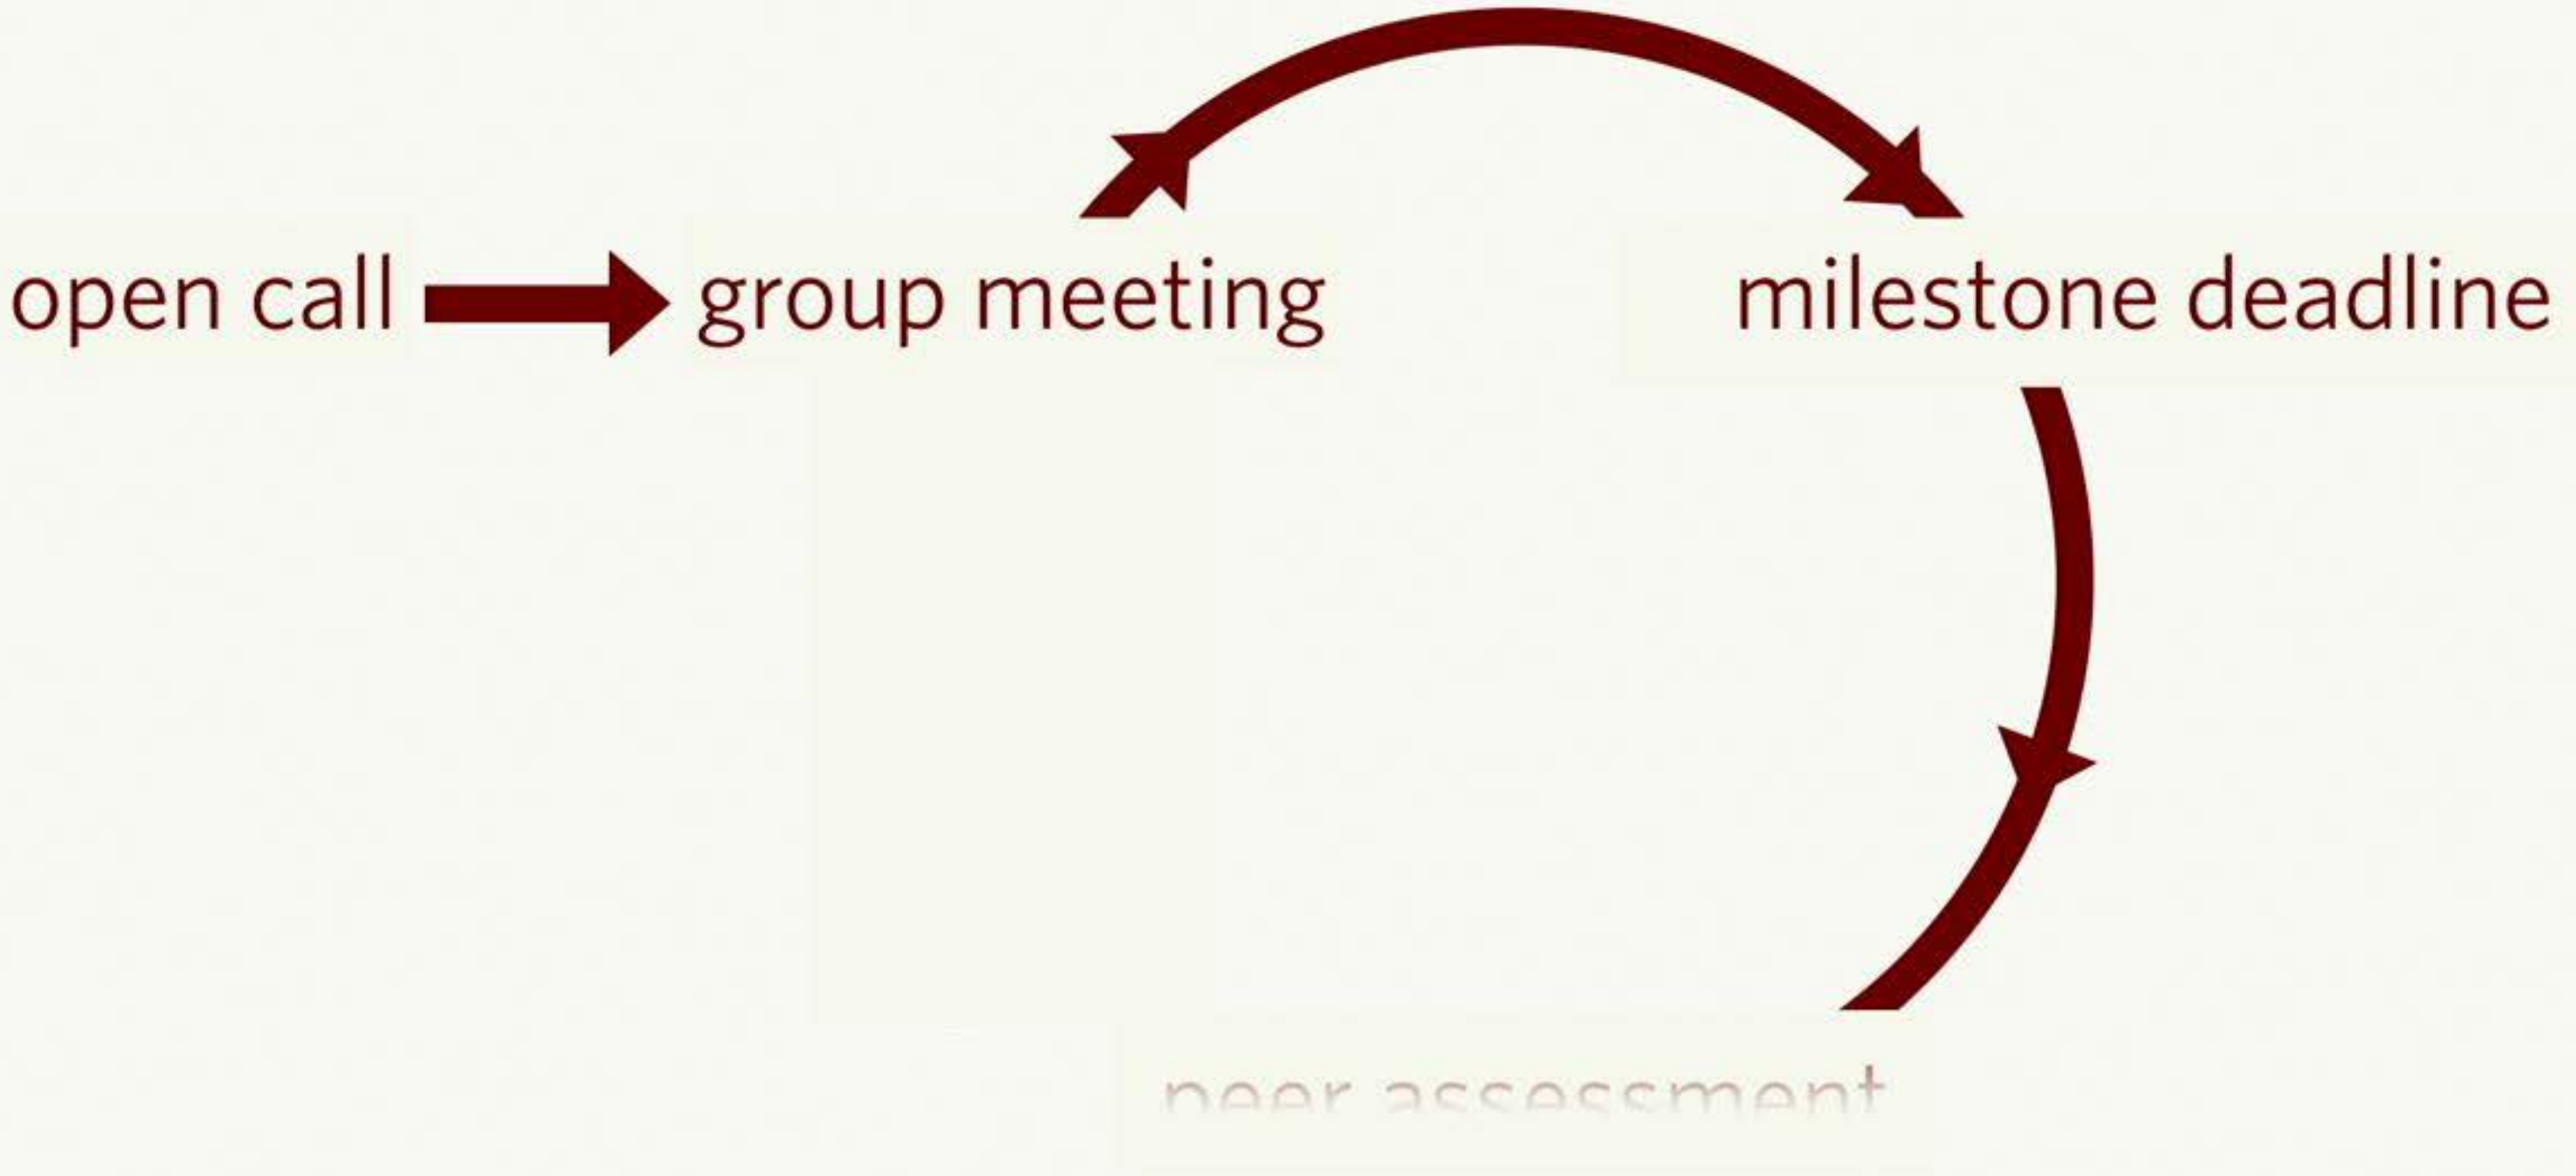

CROWDSOURCING PROCESS

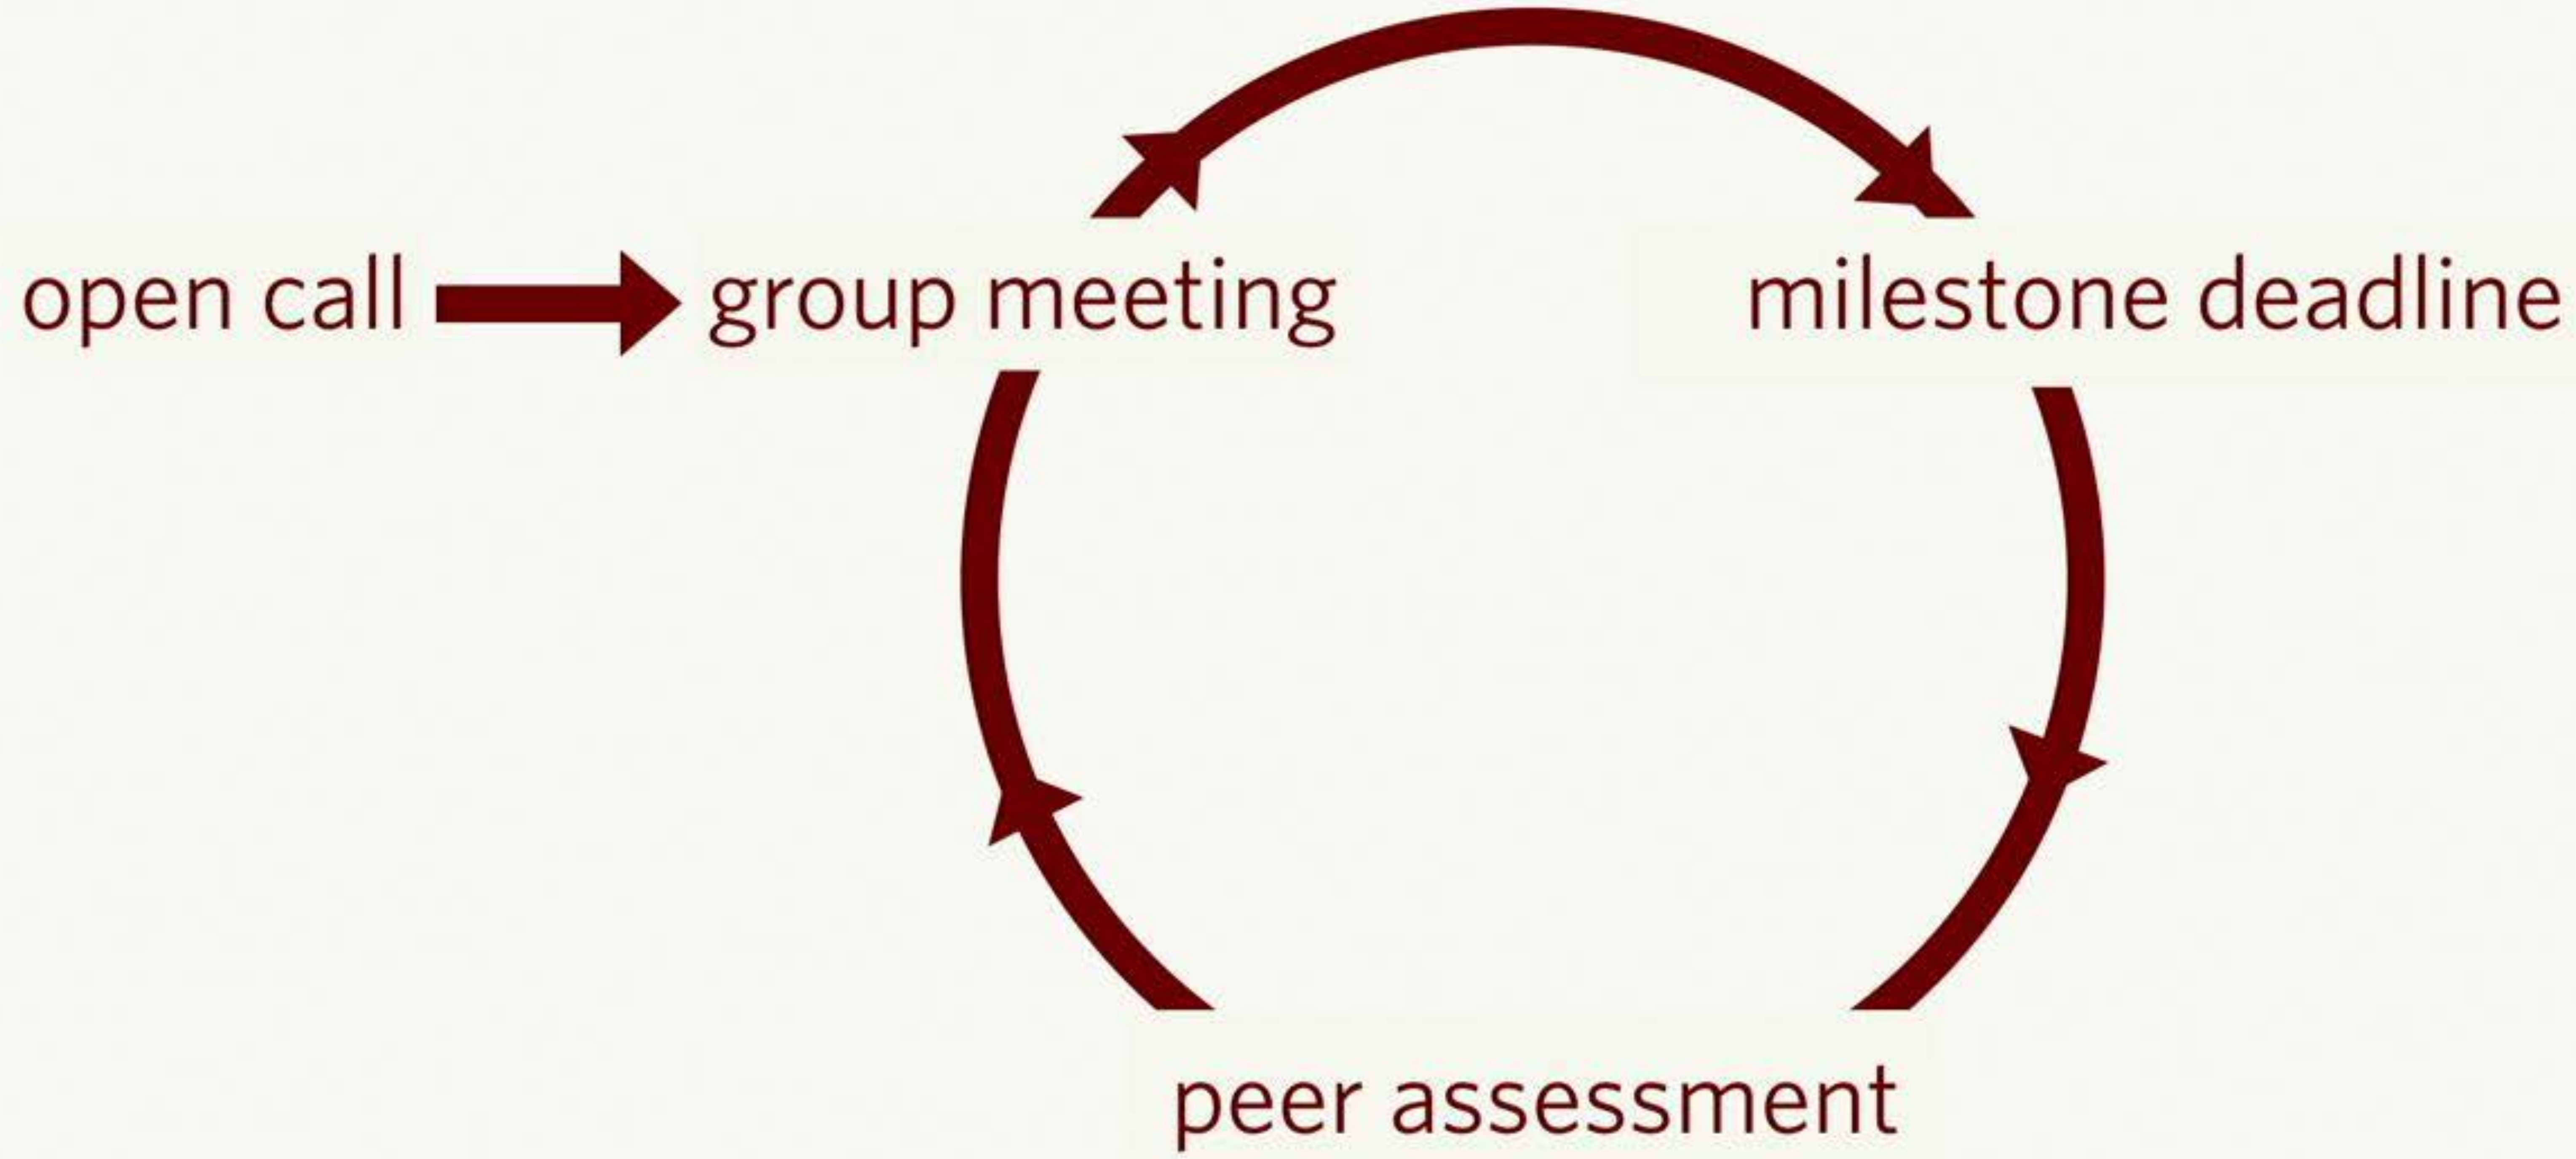

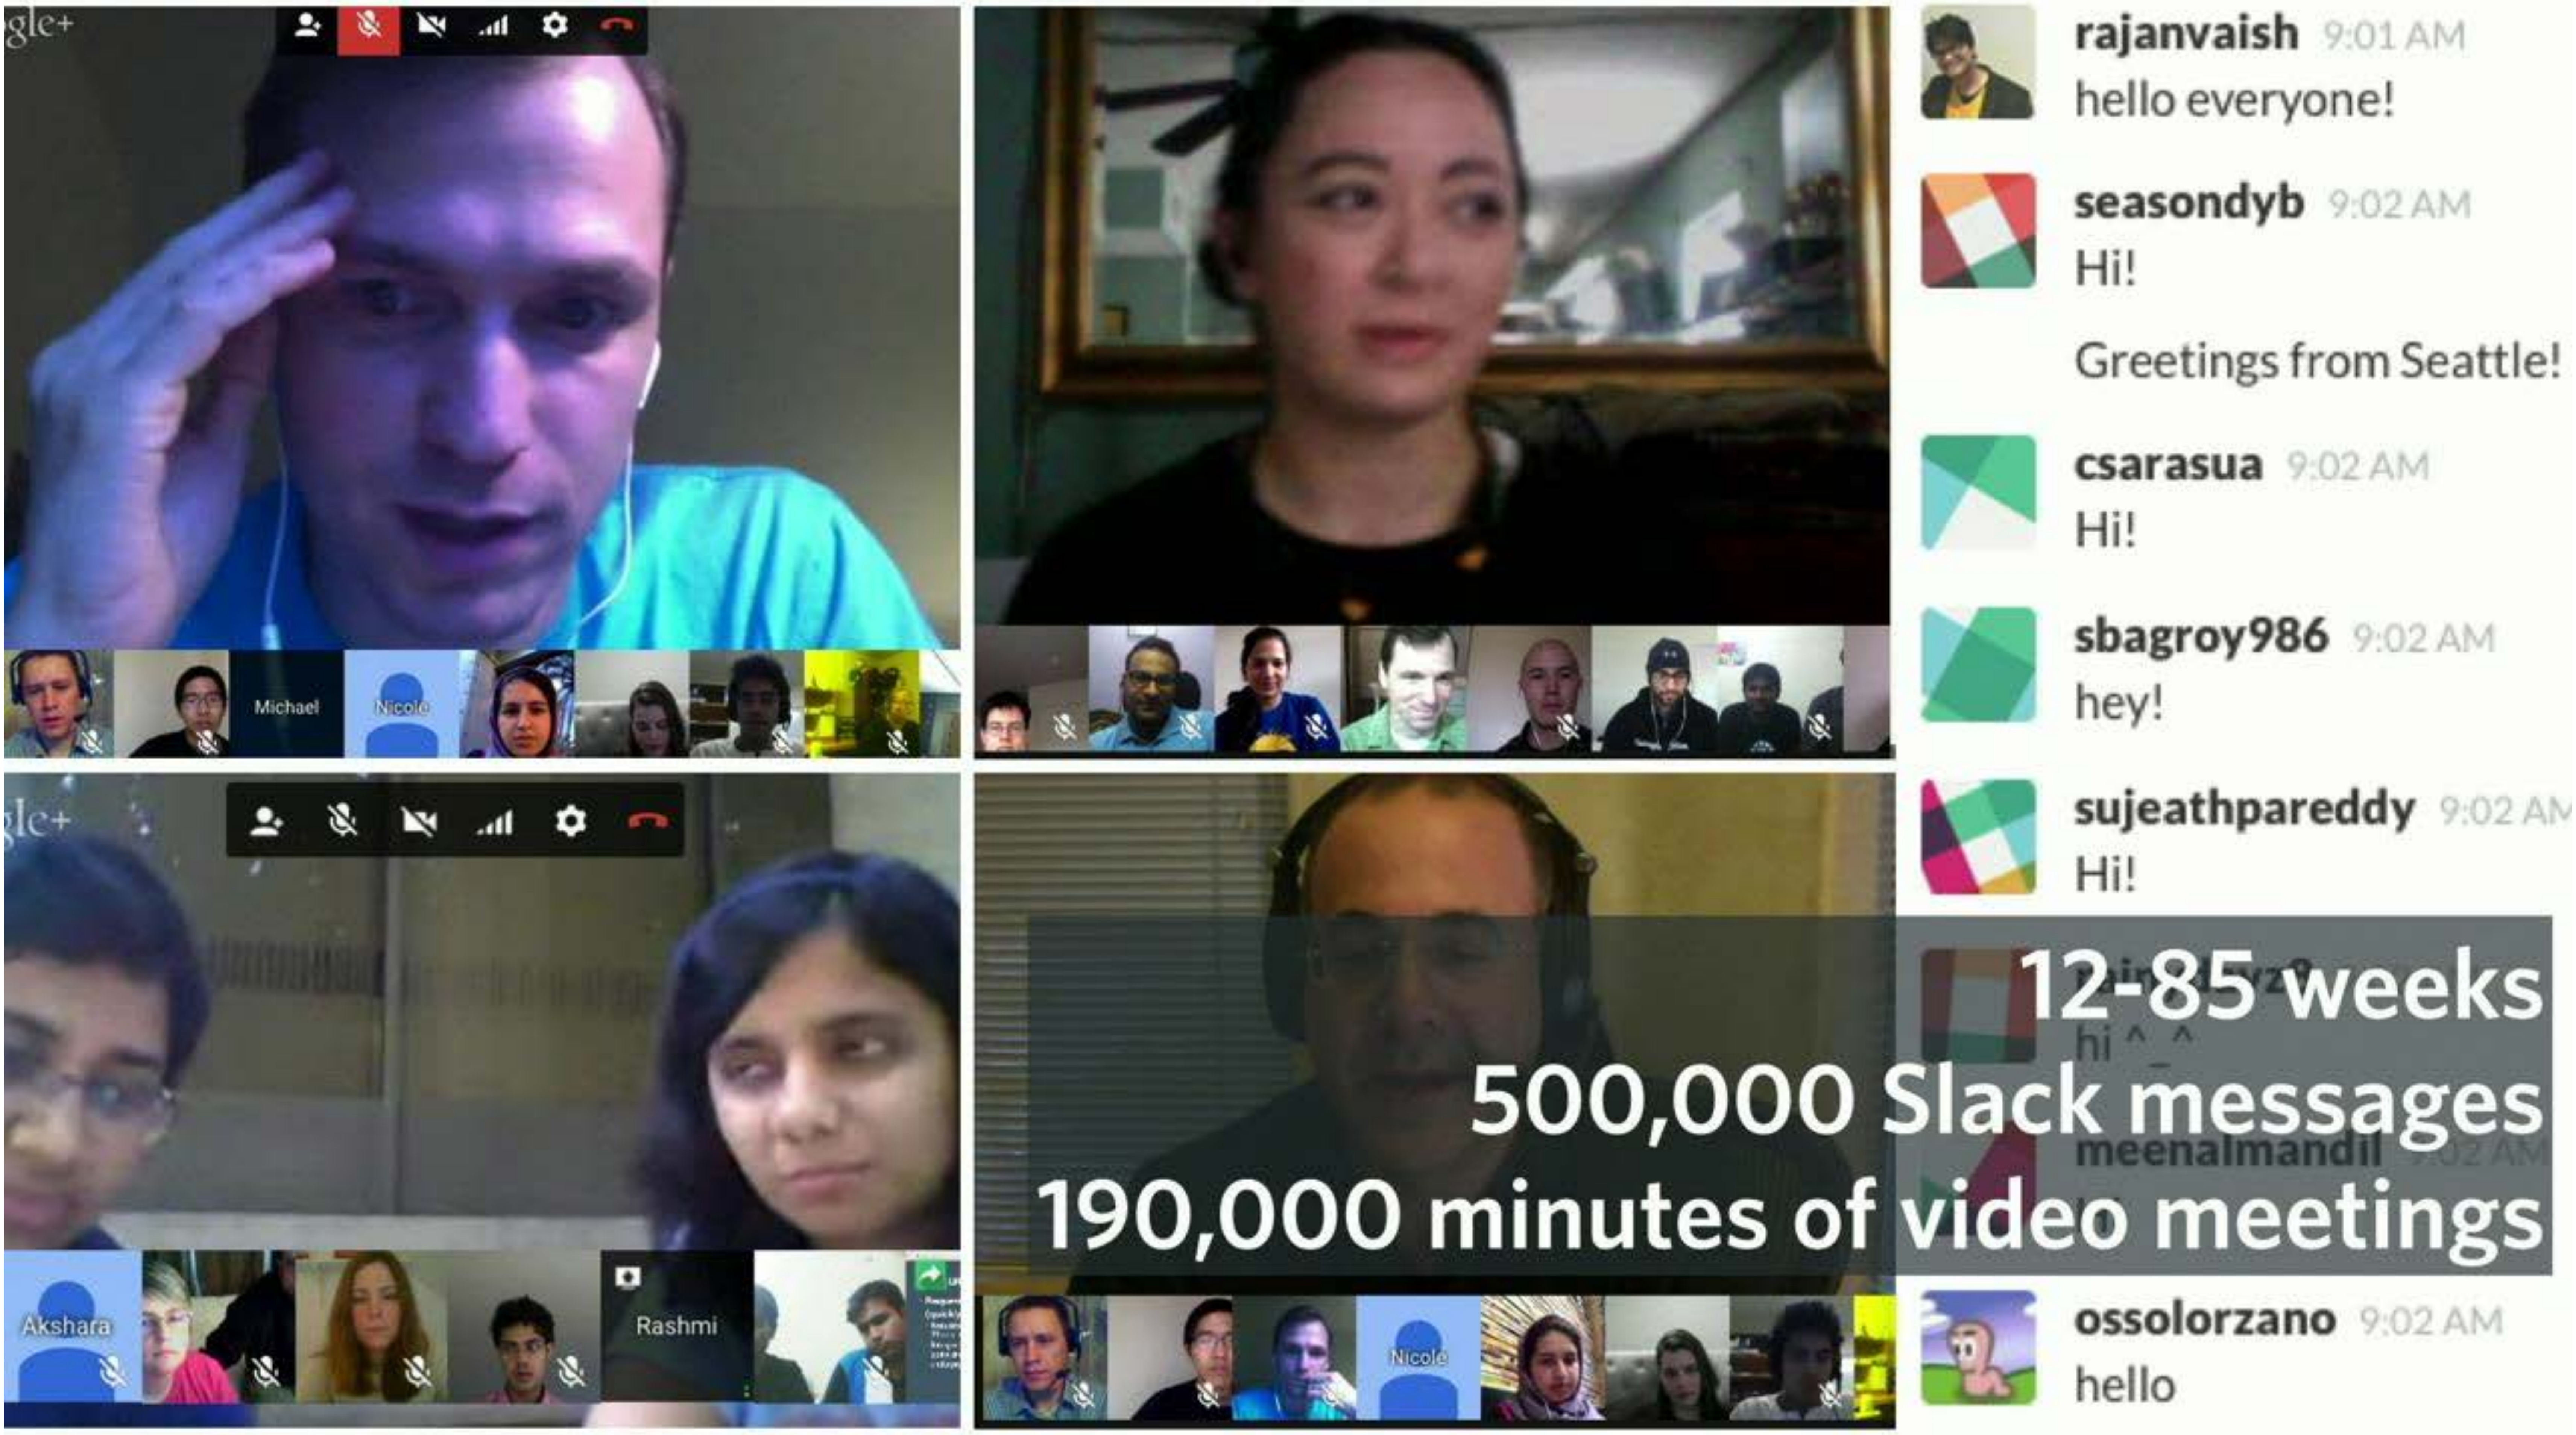
rajanvaish 9:01 AM
hello everyone!
seasondyb 9:02 AM
Hi!
Greetings from Seattle!
csarasua 9:02 AM
Hi!
sbagroy986 9:02 AM
hey!
sujeathpareddy
Hi!
ossolorzano 9:02 AM
hello
Michael
Nicole
Akshara
Rashmi
Nicole
12-85 weeks
500,000 Slack messages
190,000 minutes of video meetings

TASK PLANNING

MILESTONES

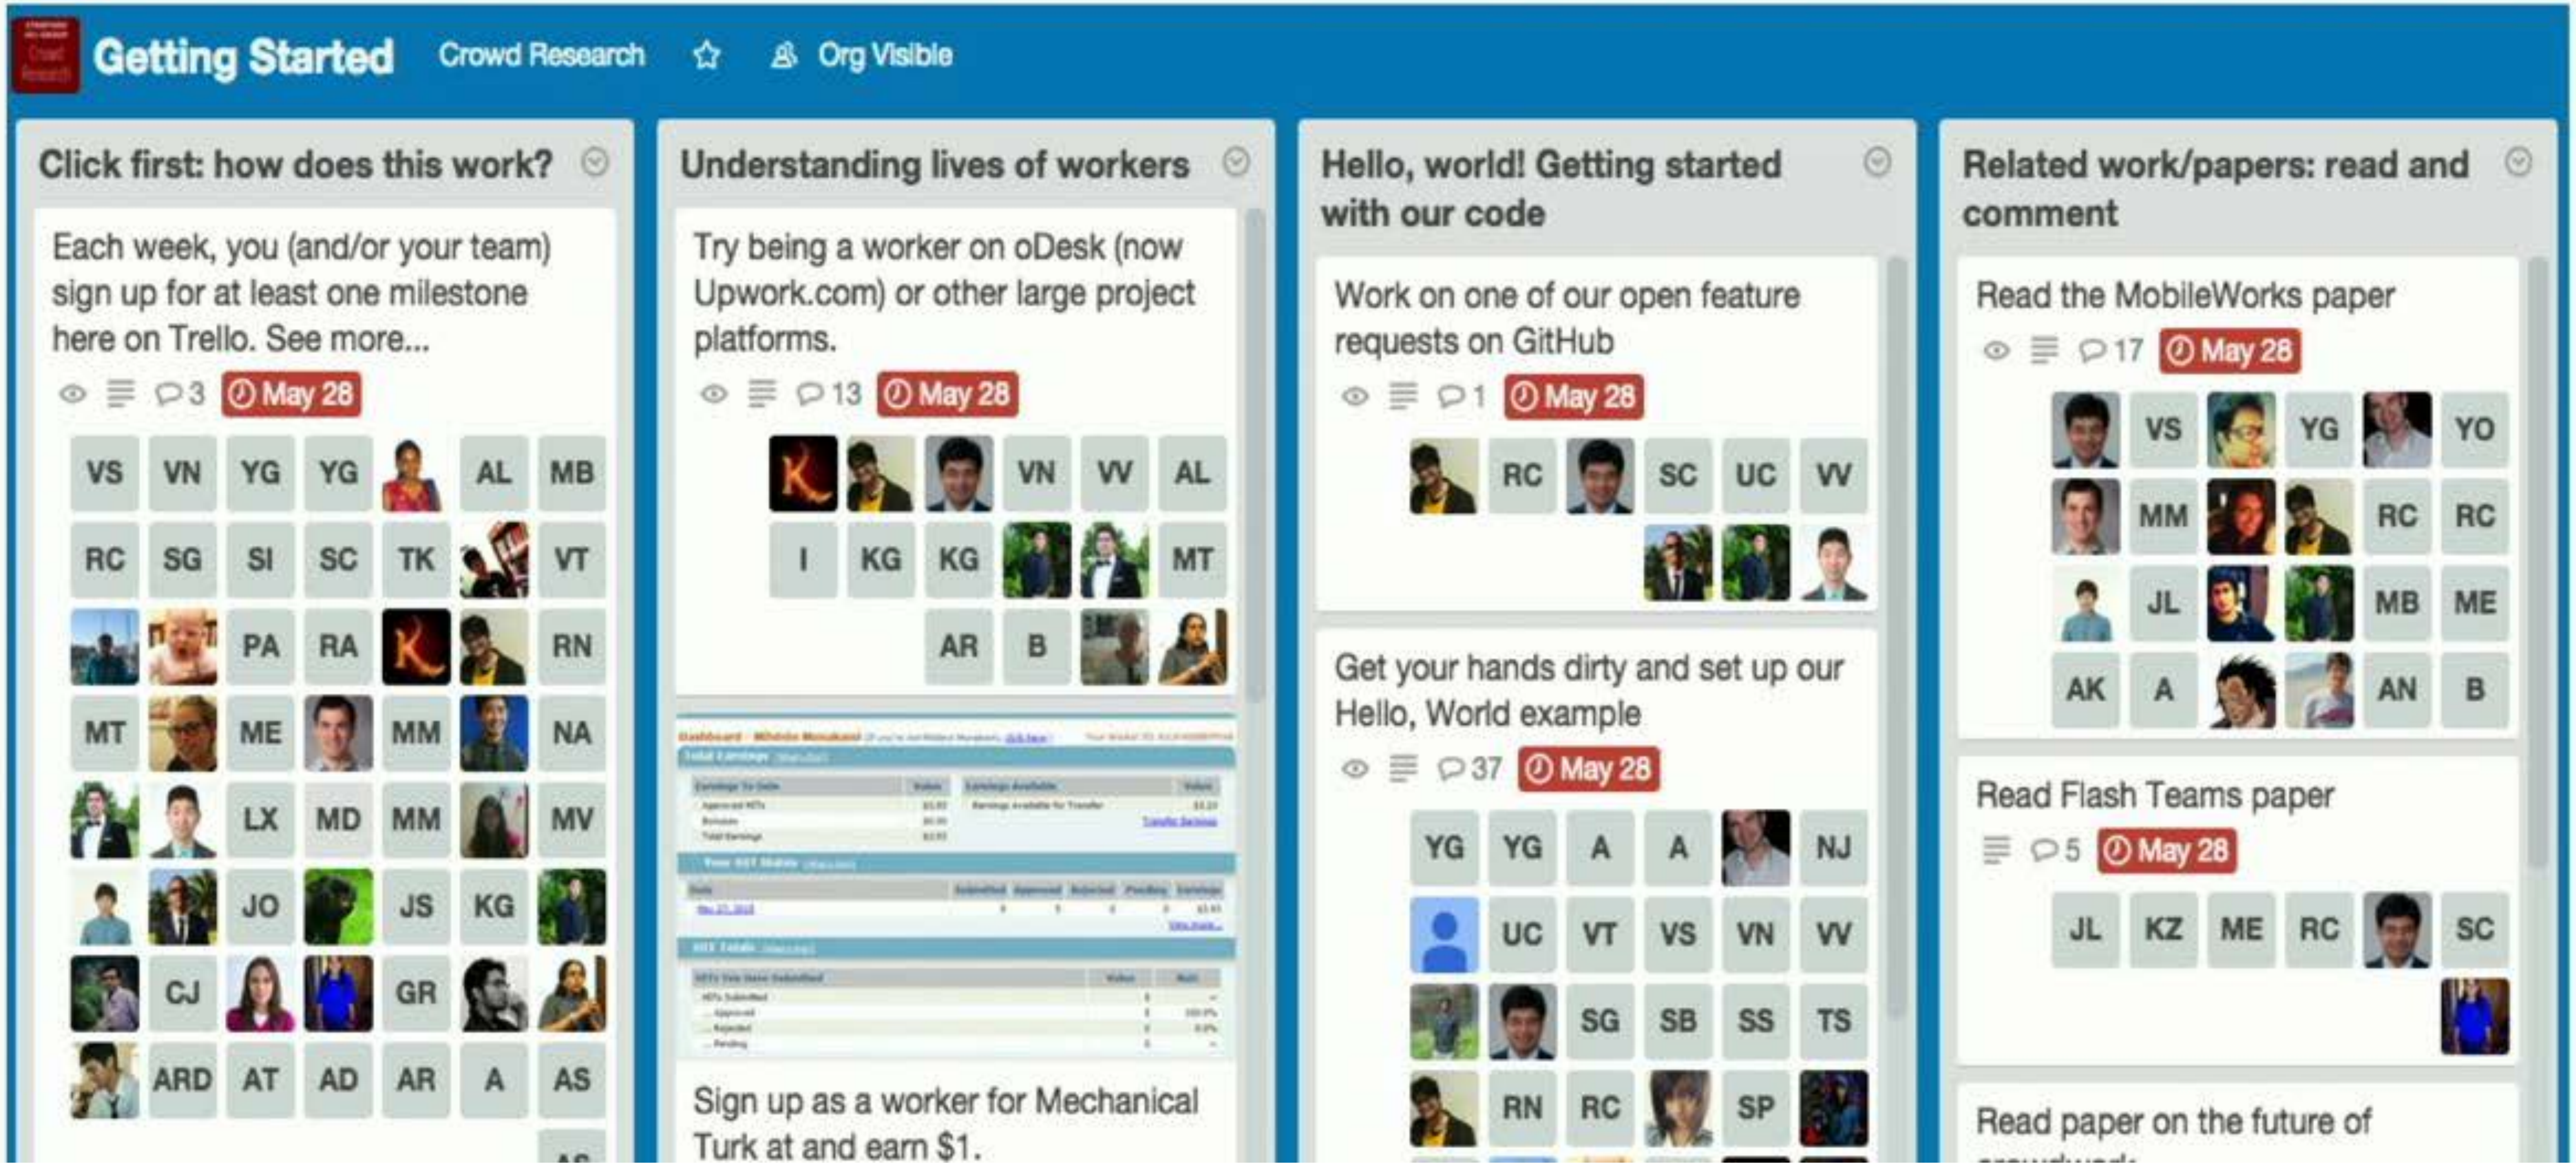

ENGINEERING

MILESTONES

PROTOTYPING

MILESTONES

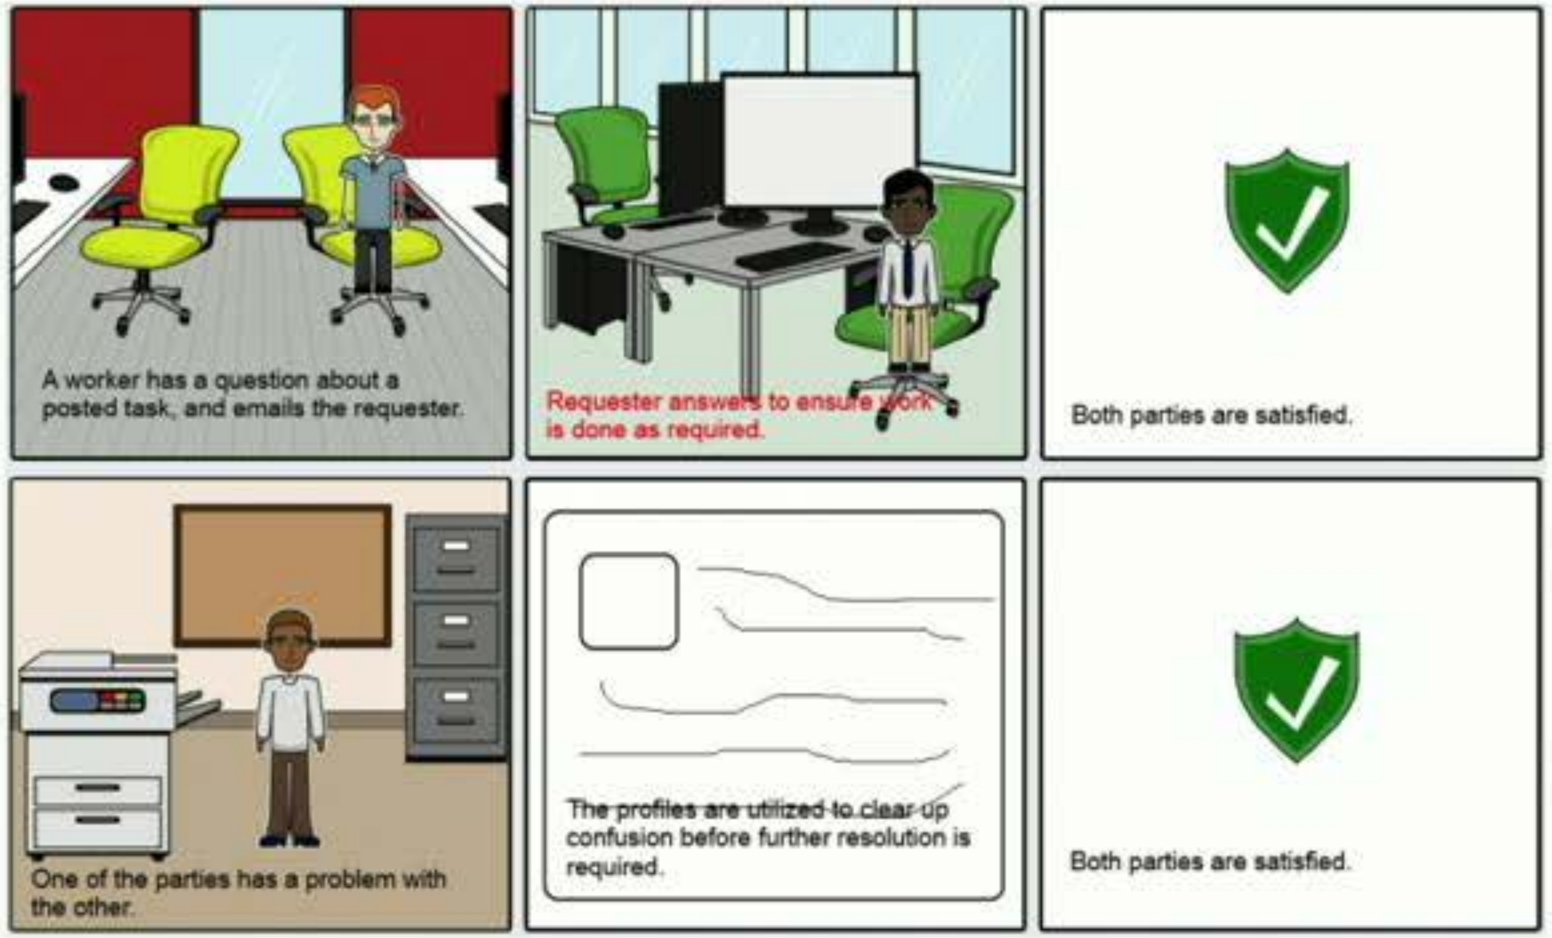

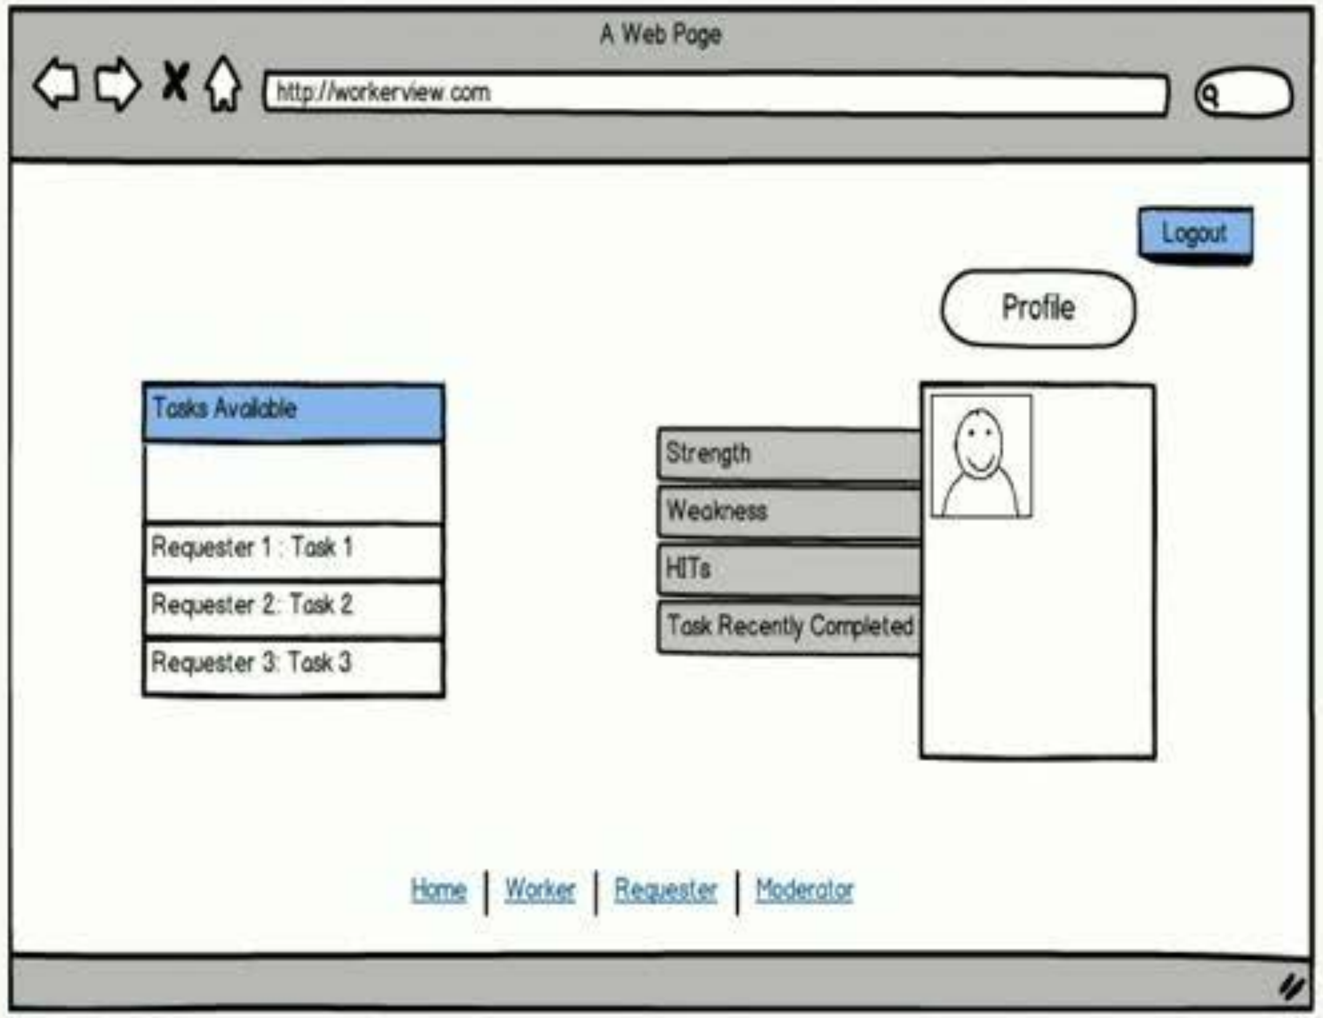

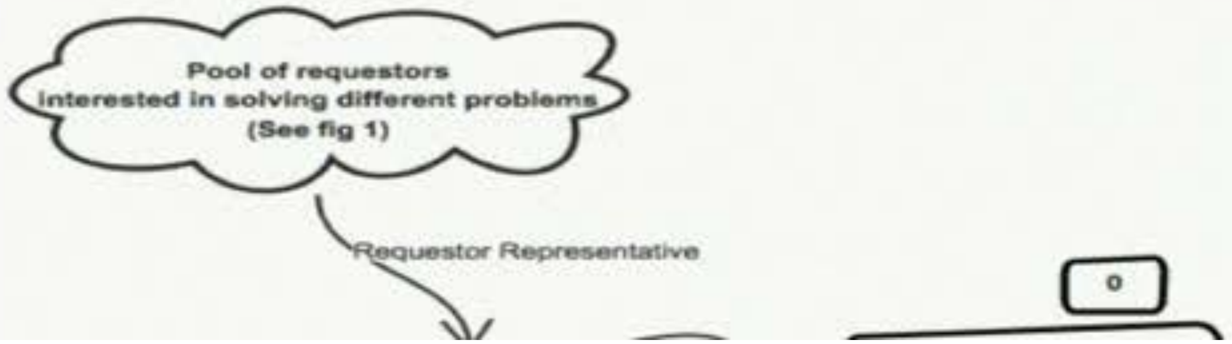

DATA ANALYSIS

MILESTONES

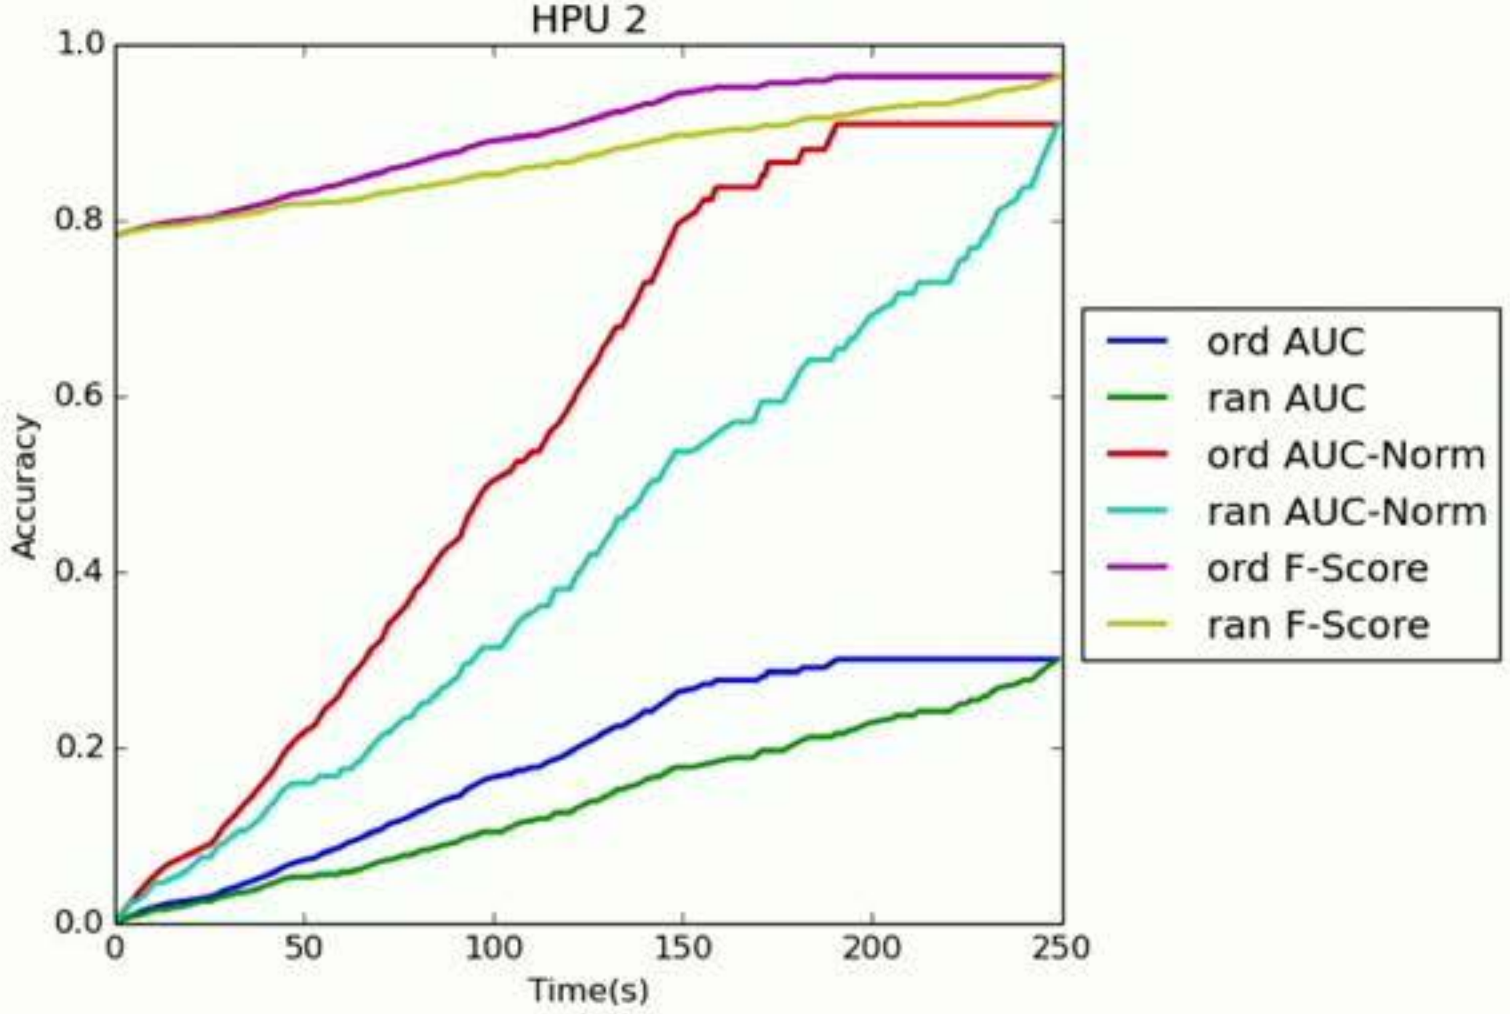

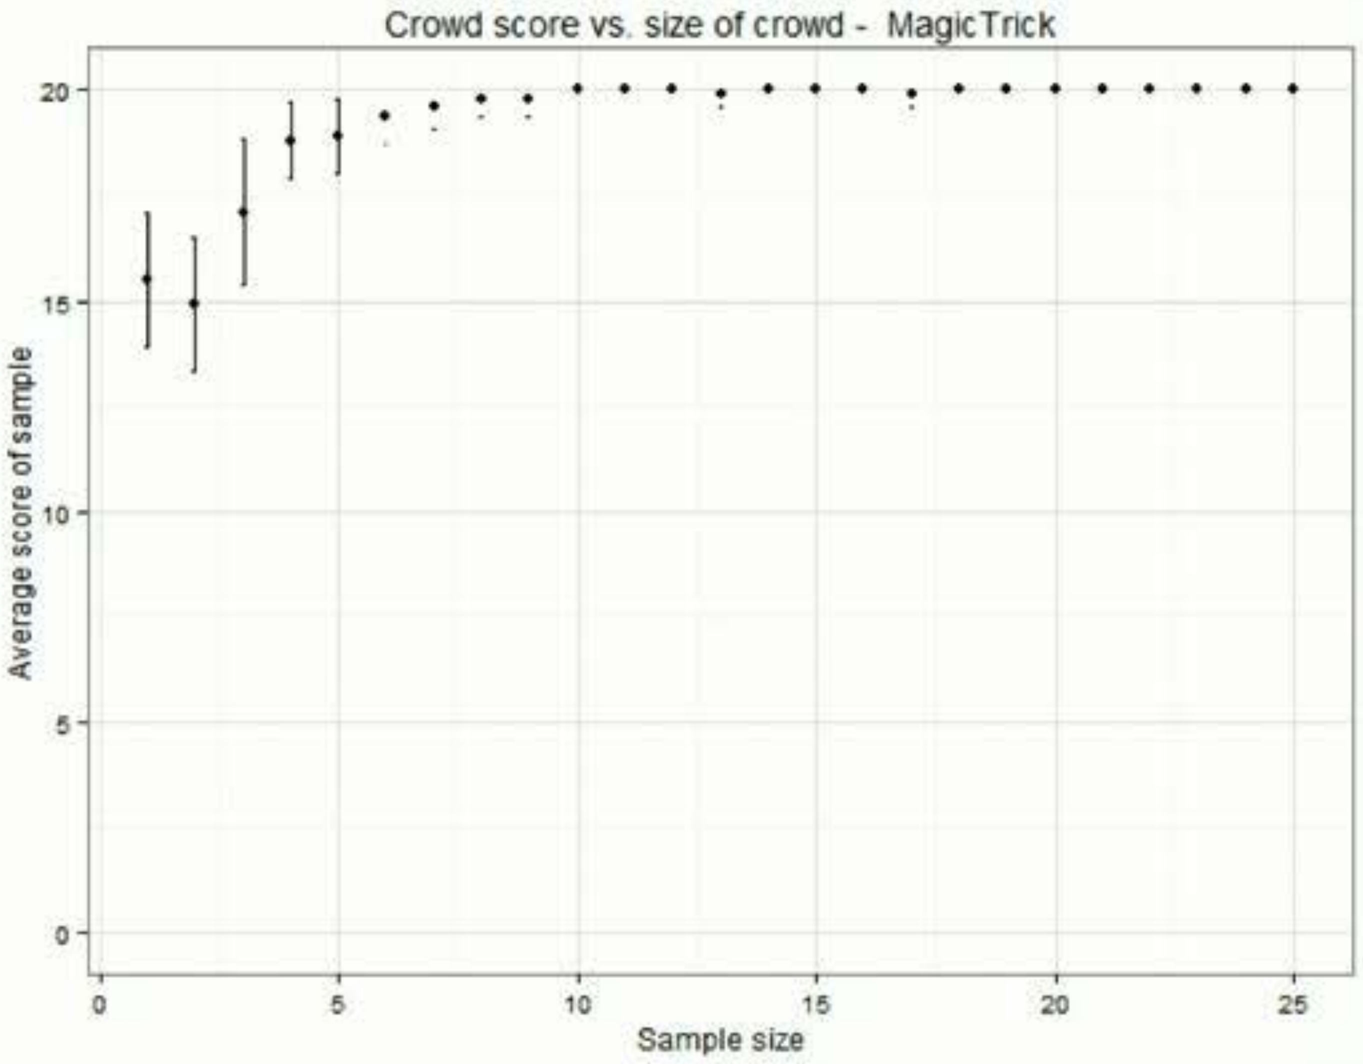

WRITING

MILESTONES

1. Anyone can pitch an idea. If it gets enough support, it goes to the next election and needs majority support from both workers+requesters.
    - **(original) Direct democracy**: anyone can pitch a policy idea, and once it gets past a threshold of support (e.g., 1000 votes), it goes up on a ballot. Twice a year, ideas go out to a direct vote for everyone on the platform. If it gets majority support from both workers and requesters, it passes.
2. Members get elected as worker or requester representatives (3 each) to a panel. Tiebreaking from a 7th member (jointly elected president).
    - **(original) Representative democracy:** once a year, members of the platform can be elected as either worker or requester representatives for a small panel (e.g, six people). Anybody can pitch a policy idea, and once it gets past a threshold of support (e.g., 1000 votes), the elected representatives must discuss it and vote on it.
3. Wikimocracy: the site's rules and policies are a wiki. Anyone can discuss, and if they edit, policies change directly.
4. Any idea that gets enough support enters a public one-month voting period. It's completely voluntary to vote. (Like a Kickstarter campaign.)
    - **Original: Fast-paced referendums:** similar concept as direct democracy, but instead of per year, you do it as vote thresholds within a month (within time of posting), and it's completely voluntary to vote. Kinda like a campaign on kickstarter. Fast pace and flexible deadlines will help the ideas continually flowing in.
5. For low-level changes, highlight the interface and suggest changes directly. Upvote/downvote directly on the interface.

majority of workers and requestors and not only one side. this could help balancing the platform.

Adam Marcus
5:05 AM May 9
Resolve

equal representation of workers and requesters? pro: seems fair, con: might run into the same sorts of paralysis issues the FEC is in now (http://mobile.nytimes.com/2015/05/03/us/politics/fec-cant-curb-2016-election-abuse-commission-chief-says.html)

Saloni Kogta
11:52 AM May 13
Resolve

I am not sure how "fair" these elections would be. Money and power could play a major role here. I may be referring to a case that has extremely small possibility of occurring, but, what if the intentions of the elected members are changed or are influenced by some other party?

Reply...

Anonymous

PEER ASSESSMENT

MILESTONES

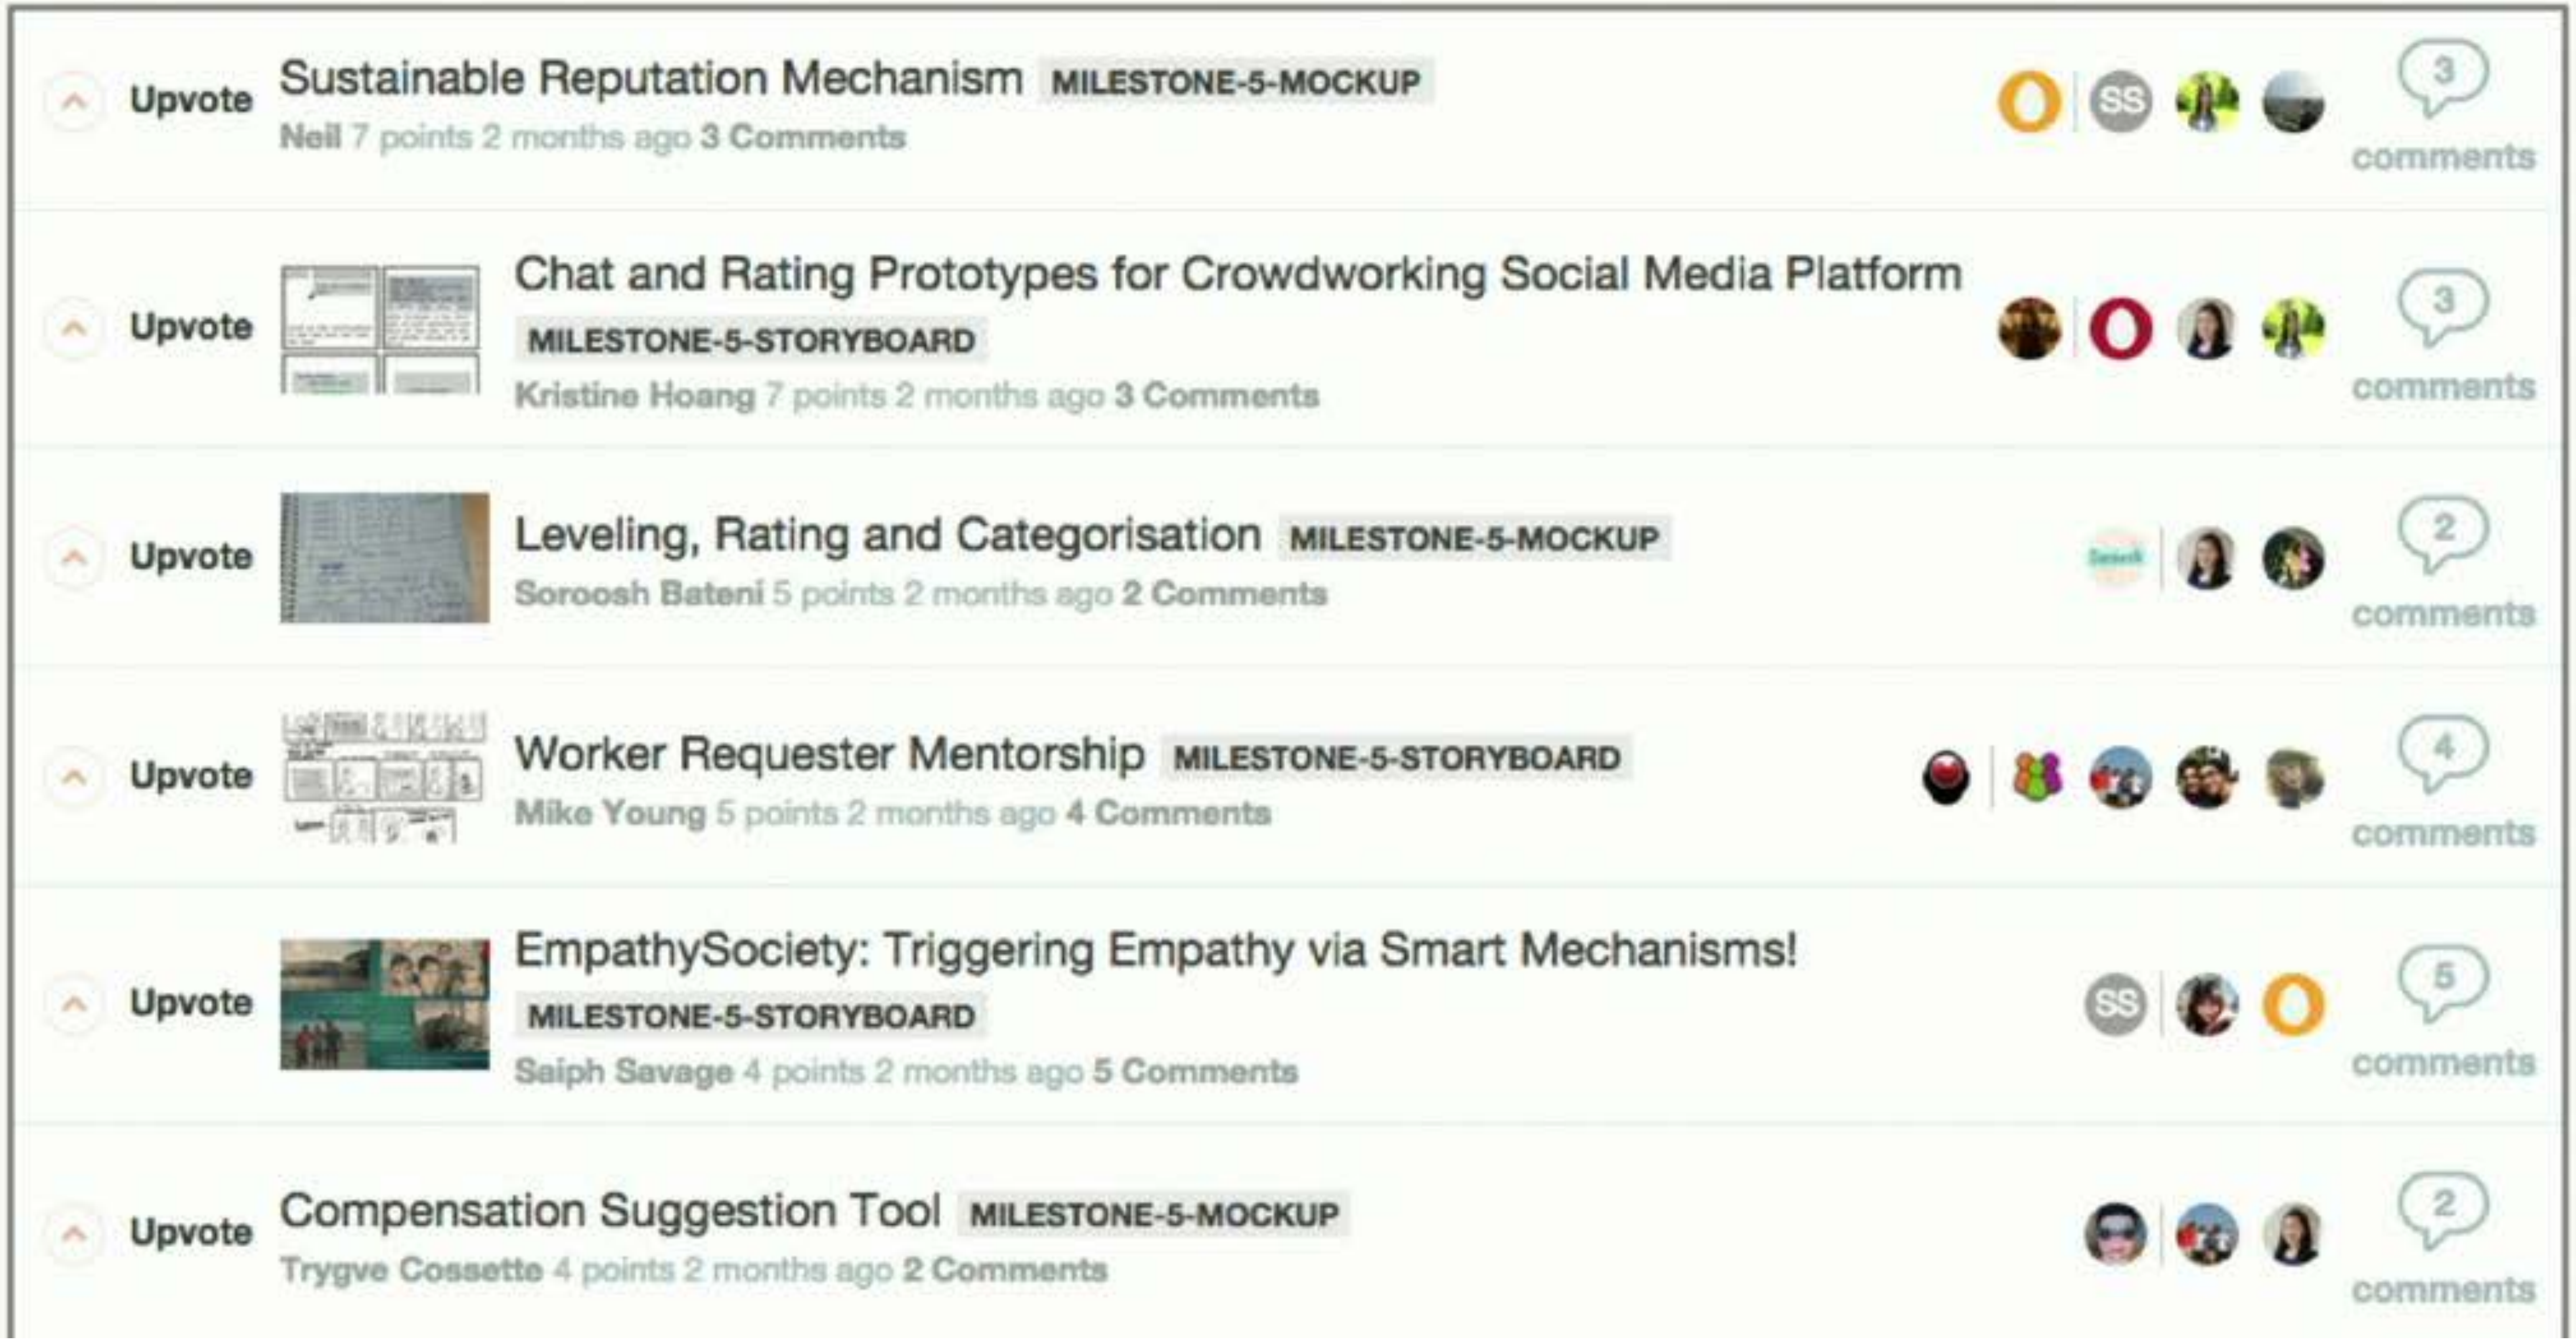

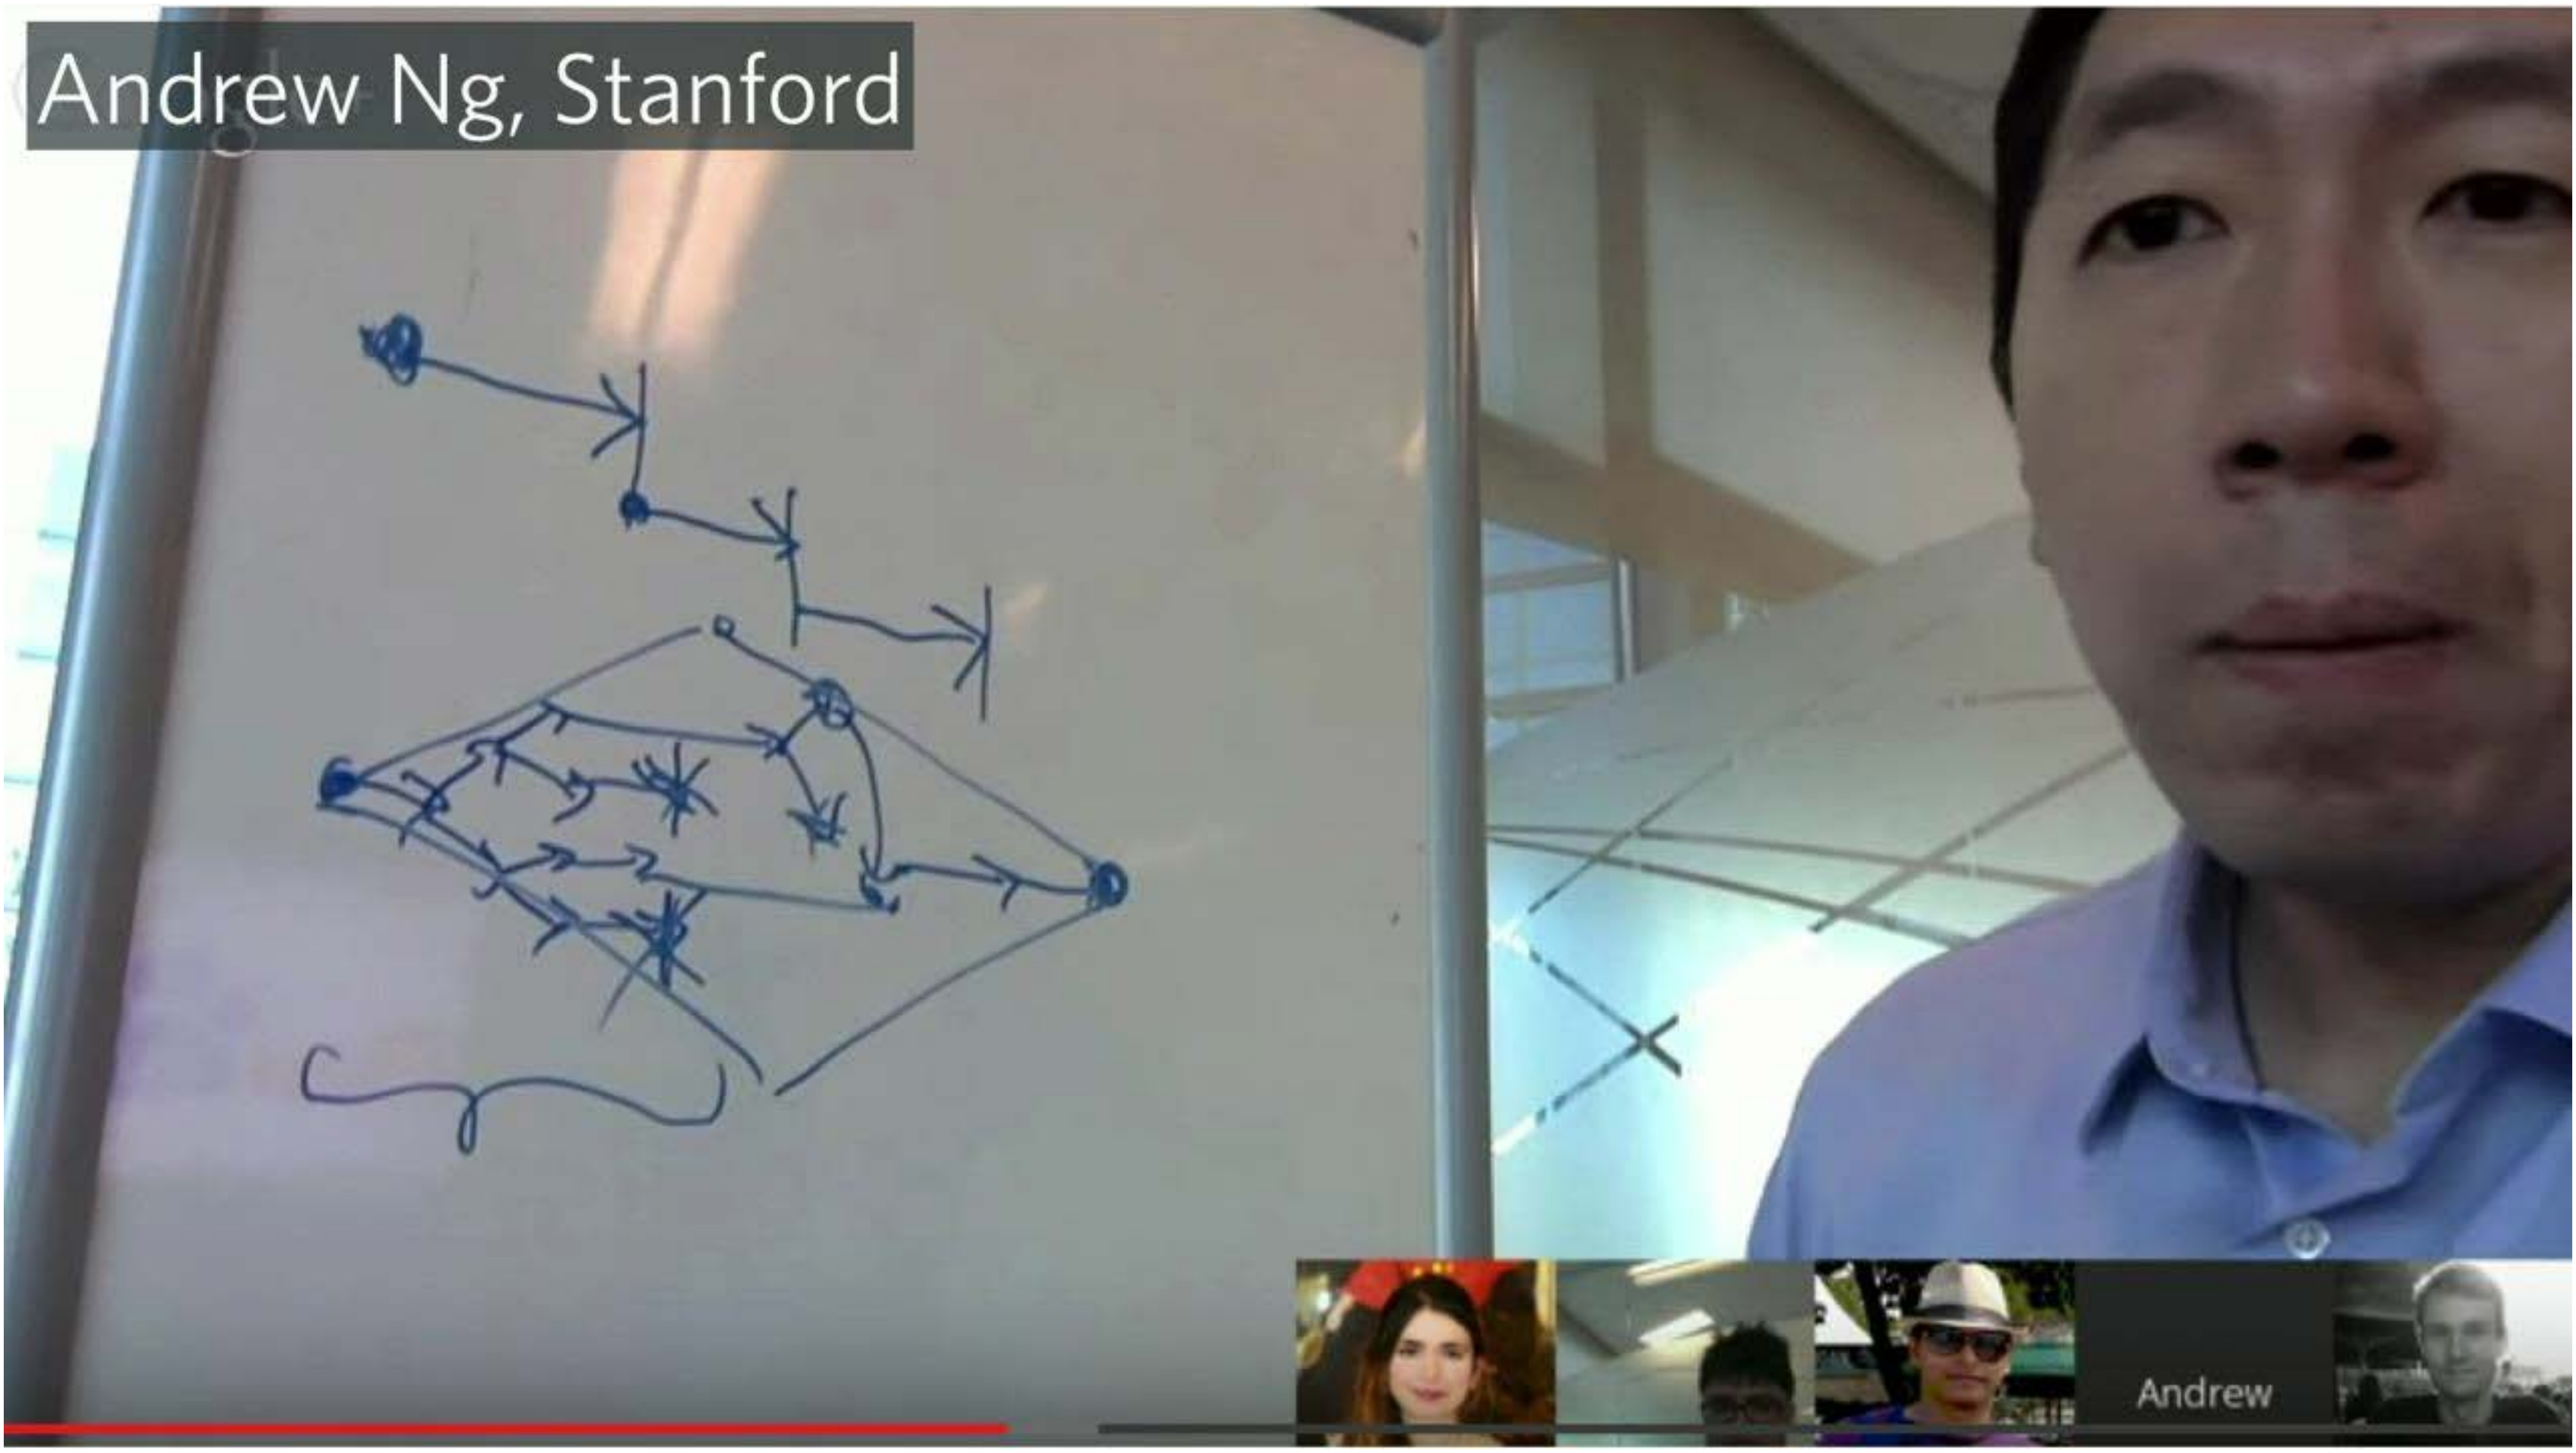
Andrew Ng, Stanford
Andrew

Peter Norvig, Google
Peter


Anant Agarwal, MIT and EdX

1500 participants from six continents

2% high school, 73% undergrad, 22% master's, 3% PhD

RECRUITMENT: PROVIDING ACCESS

Matching affiliations to Times Higher Education Global Rankings: **75% come from universities ranked below 500.**

Considering country of origin: **66% come from countries ranked over 50 in GDP per capita**

Boomerang: Rebounding the Consequences of Reputation Feedback on Crowdsourcing Platforms

S.S. Gaikwad, D. Morina, A. Ginzberg, C. Mullings, S. Goyal, D. Gamage, C. Diemert, M. Burton, S. Zhou, M. Whiting, K. Ziulkoski, A. Ballav, A. Gilbee, S.S. Niranga, V. Sehgal, J. Lin, L. Kristianto, A. Richmond-Fuller, J. Regino, N. Chhibber, D. Majeti, S. Sharma, K. Mananova, D. Dhakal, W. Dai, V. Purynova, S. Sandeep, V. Chandrakanthan, T. Sarma, S. Matin, A. Nassar, R. Nistala, A. Stolzoff, K. Milland, V. Mathur, R. Vaish, M.S. Bernstein

Stanford Crowd Research Collective, Stanford University
daemo@cs.stanford.edu

ABSTRACT

Paid crowdsourcing platforms suffer from low-quality work and unfair rejections, but paradoxically, most workers and requesters have high reputation scores. These inflated scores, which make high-quality work and workers difficult to find, stem from social pressure to avoid giving negative feedback. We introduce Boomerang, a reputation system for crowdsourcing platforms that elicits more accurate feedback by rebounding the consequences of feedback directly back onto the person who gave it. With Boomerang, requesters find that their highly-rated workers gain earliest access to their future tasks, and workers find tasks from their highly-rated requesters at the top of their task feed. Field experiments verify that Boomerang

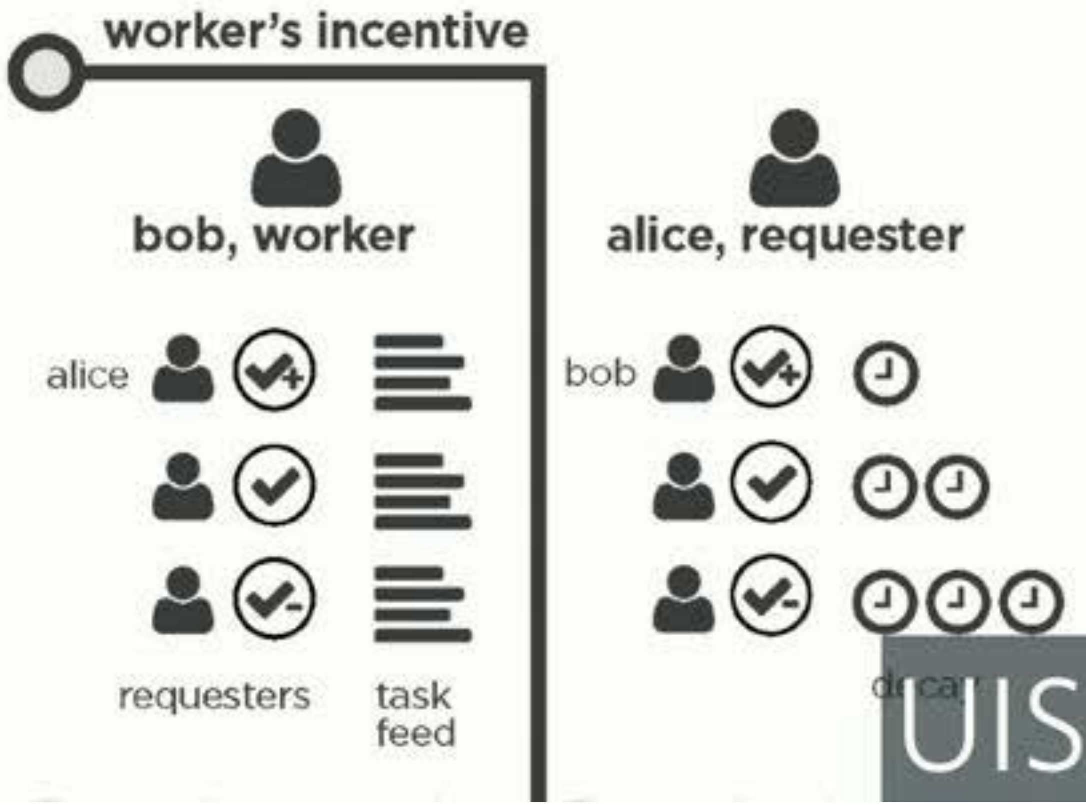

Crowd Guilds: Worker-led Reputation and Feedback on Crowdsourcing Platforms

Mark E. Whiting, Dilrukshi Gamage, Snehalkumar (Neil) S. Gaikwad, Aaron Gilbee, Shirish Goyal, Alipta Ballav, Dinesh Majeti, Nalin Chhibber, Angela Richmond-Fuller, Freddie Vargus, Tejas Seshadri Sarma, Varshine Chandrakanthan, Teogenes Moura, Mohamed Hashim Salih, Gabriel Bayomi Tinoco Kalejaiye, Adam Ginzberg, Catherine A. Mullings, Yoni Dayan, Kristy Milland, Henrique Orefice, Jeff Regino, Sayna Parsi, Kunz Mainali, Vibhor Sehgal, Sekandar Matin, Akshansh Sinha, Rajan Vaish, Michael S. Bernstein

Stanford Crowd Research Collective

daemo@cs.stanford.edu

ABSTRACT

Crowd workers are distributed and decentralized. While decentralization is designed to utilize independent judgment to promote high-quality results, it paradoxically undercuts behaviors and institutions that are critical to high-quality work. Reputation is one central example: crowdsourcing systems depend on reputation scores from decentralized workers and requesters, but these scores are notoriously inflated and uninformative. In this paper, we draw inspiration from historical worker guilds (e.g., in the silk trade) to design and implement

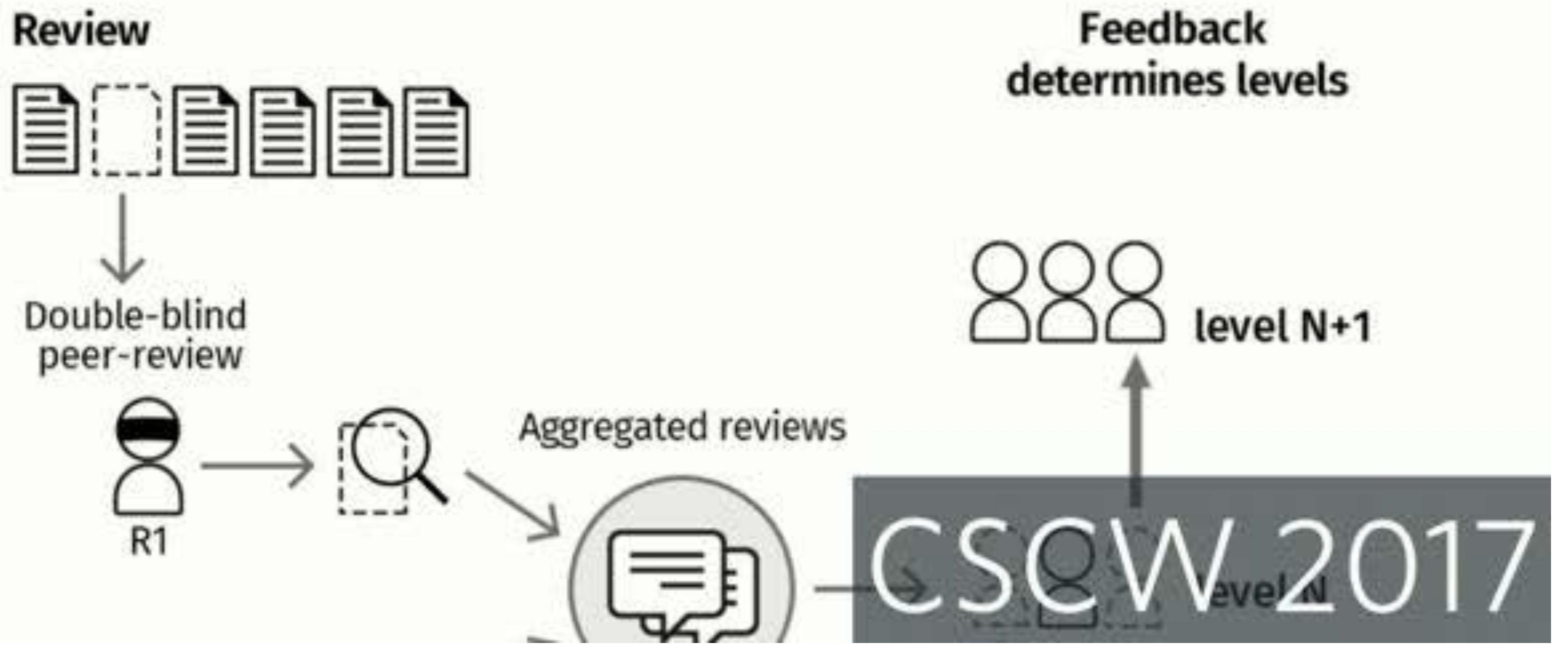

WORKS-IN-PROGRESS

Daemo: a Self-Governed Crowdsourcing Marketplace

Stanford Crowd Research Collective
Stanford HCI Group daemo@cs.stanford.edu

HCI

On Optimizing Human-Machine Task Assignments

Computer Vision

Investigating the "Wisdom of Crowds" at Scale

Alok Shankar Mysore
PES Institute of Technology
alok.shankar@pesit.pes.edu

Vikas S Yaligar
National Institute of Technology Karnataka
vikasyaligar@ieee.org

Imanol Arrieta Ibarra
Stanford University
imanol@stanford.edu

Camelia Simoiu
Stanford University
csimoiu@stanford.edu

Sharad Goel
Stanford University
sgoel@stanford.edu

Additional Authors
Various Institutions

Data Science

INCREASED ACCESS TO RESEARCH

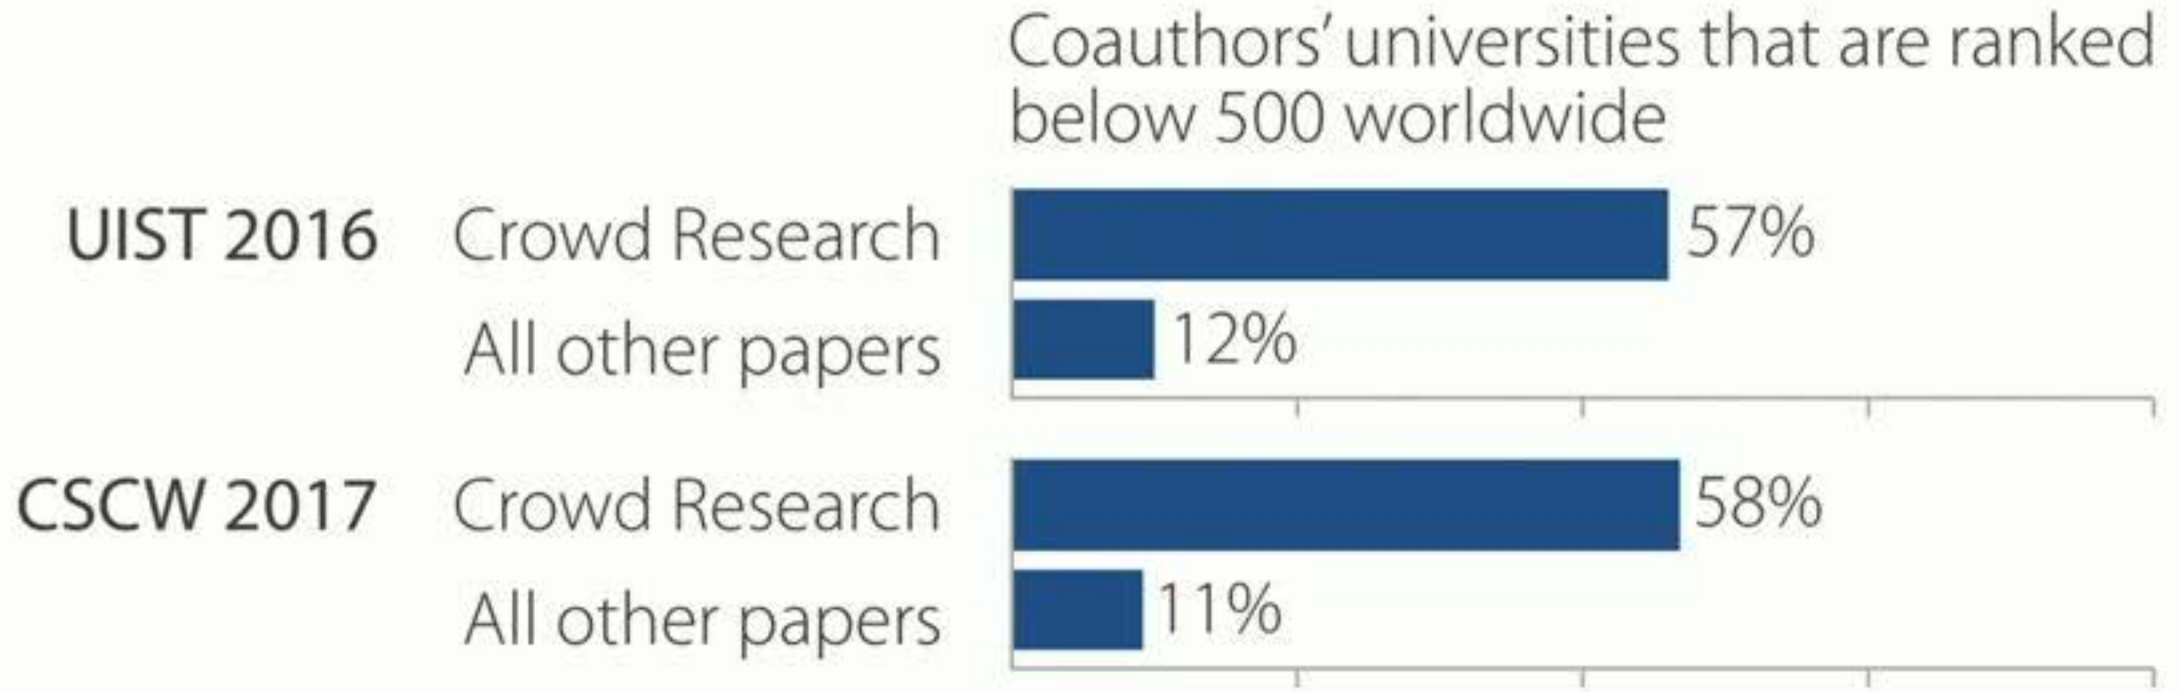

INCREASED ACCESS TO RESEARCH

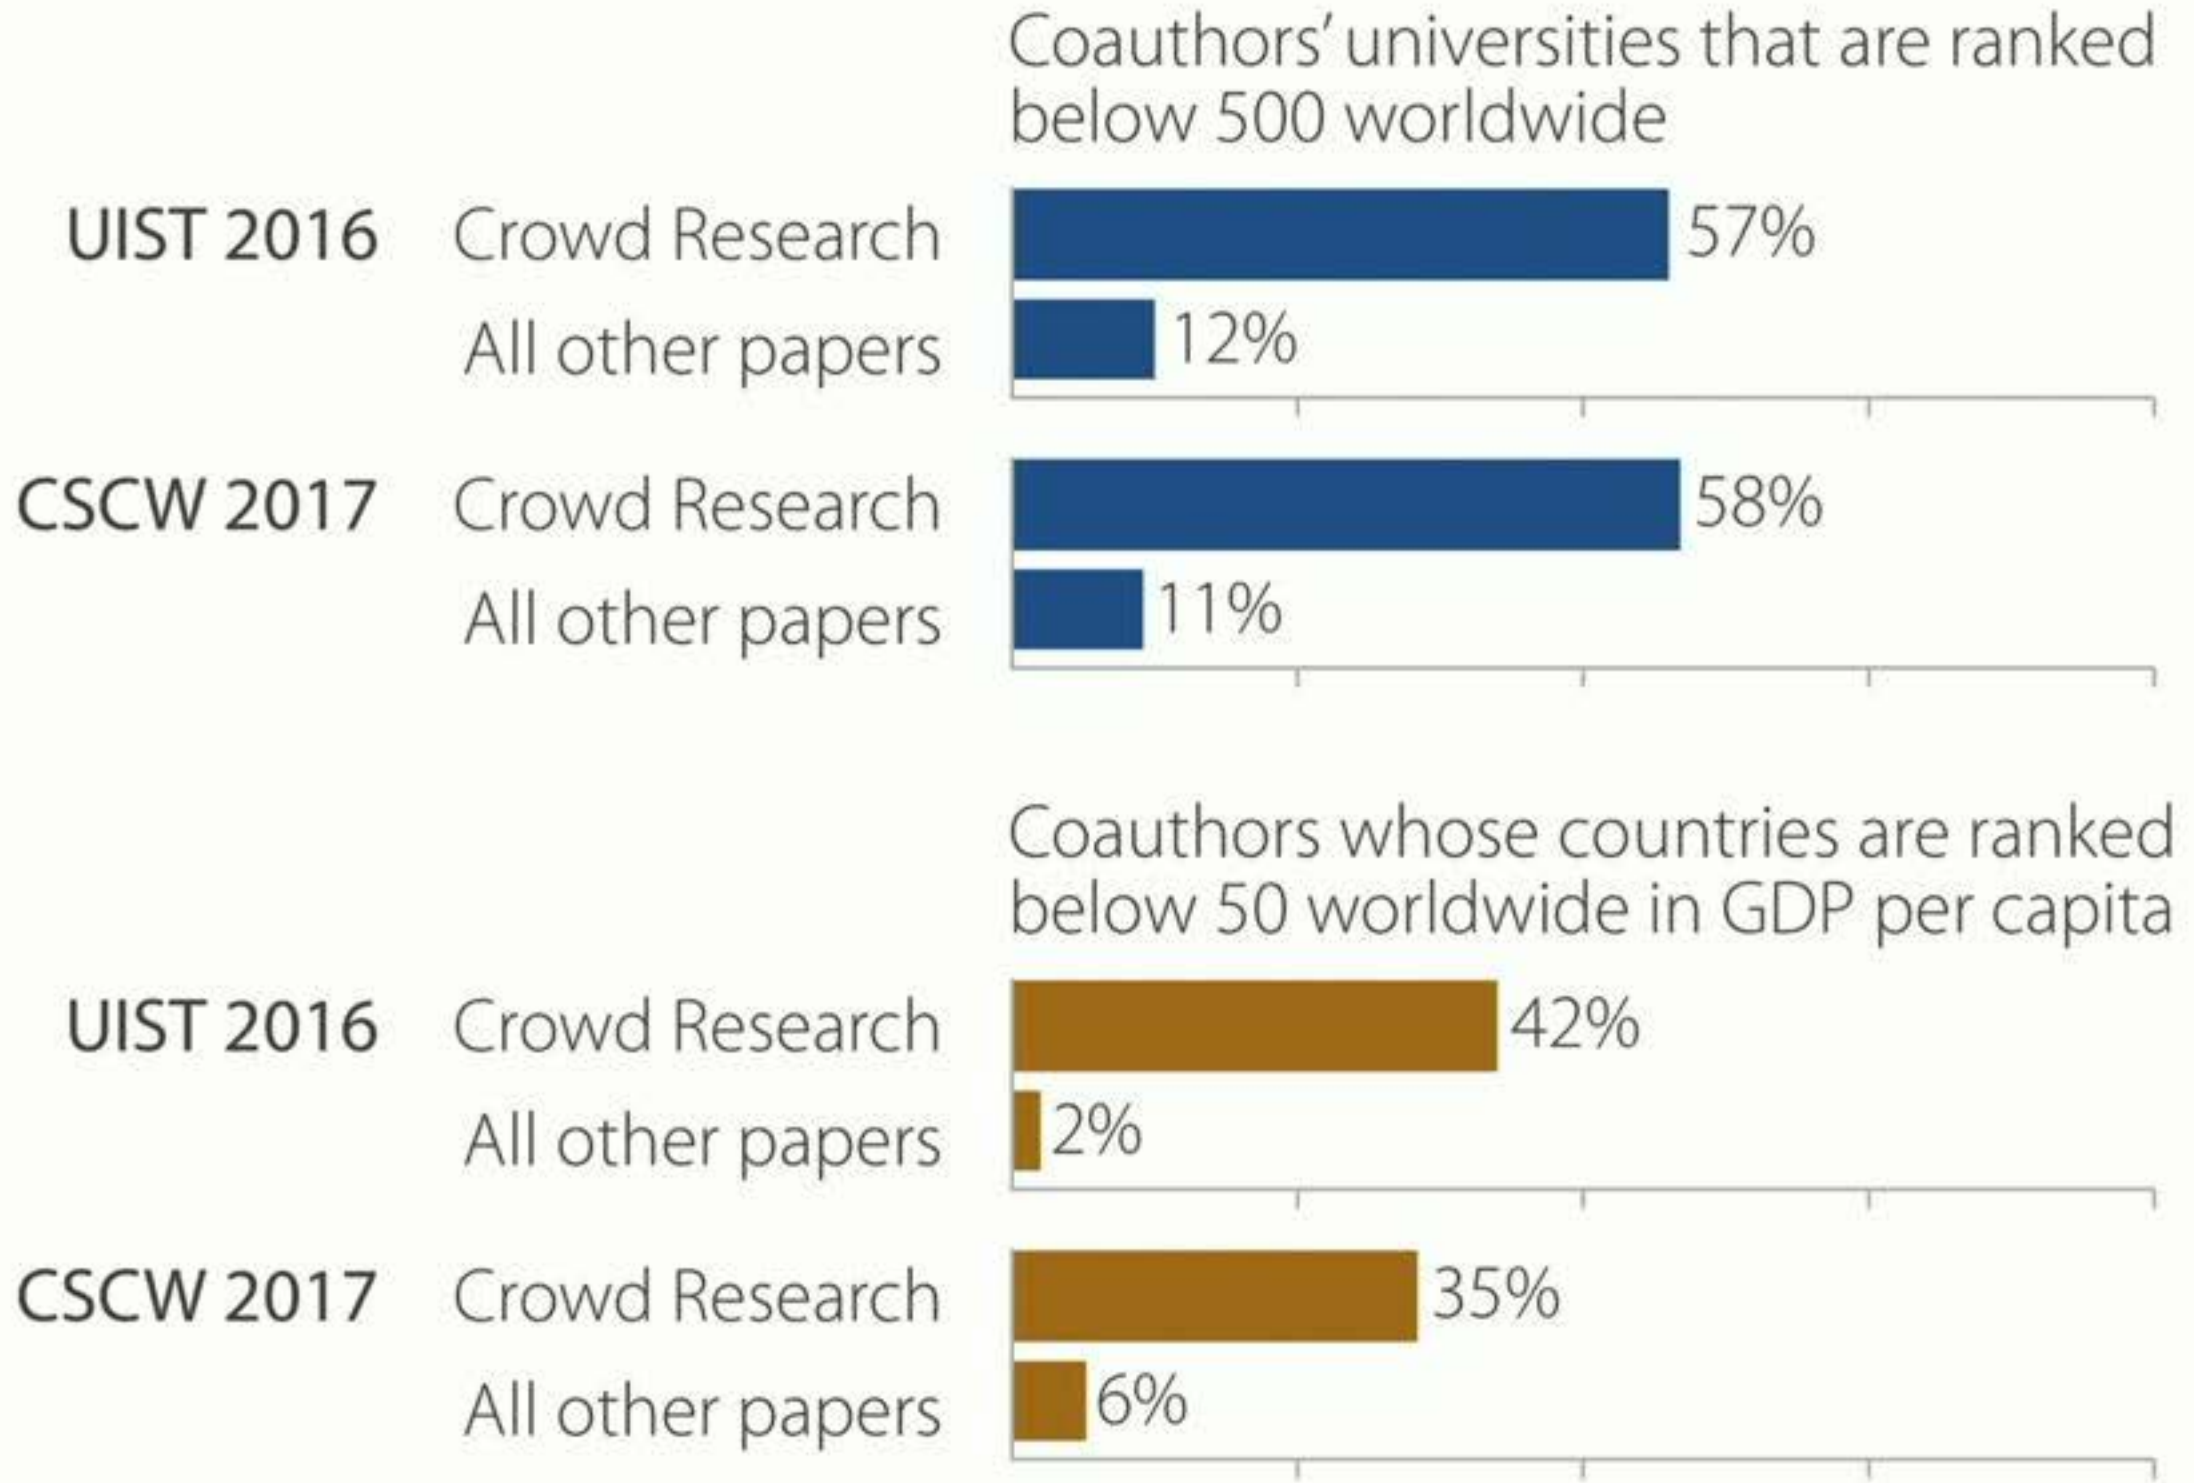

ATTAINMENT: PROVIDING ACCESS

Participants have gone on to programs at UC Berkeley, and Carnegie Mellon University, and MIT.

21 of 33 surveyed were admitted to at least one program, despite a **median of 0 other letter writers** from institutions ranked better than 500 worldwide.

Crowd Guilds: Worker-led Reputation and Feedback on Crowdsourcing Platforms

Mark E. Whiting, Dilrukshi Gamage, Snehalkumar (Neil) S. Gaikwad, Aaron Gilbee, Shirish Goyal, Alipta Ballav, Dinesh Majeti, Nalin Chhibber, Angela Richmond-Fuller, Freddie Vargus, Tejas Seshadri Sarma, Varshine Chandrakanthan, Teogenes Moura, Mohamed Hashim Salih, Gabriel Bayomi Tinoco Kalejaiye, Adam Ginzberg, Catherine A. Mullings, Yoni Dayan, Kristy Milland, Henrique Orefice, Jeff Regino, Sayna Parsi, Kunz Mainali, Vibhor Sehgal, Sekandar Matin, Akshansh Sinha, Rajan Vaish, Michael S. Bernstein
Stanford Crowd Research Collective
daemo@cs.stanford.edu

ABSTRACT

Crowd workers are distributed and decentralized. While decentralization is designed to utilize independent judgment to promote high-quality results, it paradoxically undercuts behaviors and institutions that are critical to high-quality work. Reputation is one central example: crowdsourcing systems depend on reputation scores from decentralized workers and requesters, but these scores are notoriously inflated and uninformative. In this paper, we draw inspiration from historical worker guilds (e.g., in the silk trade) to design and implement

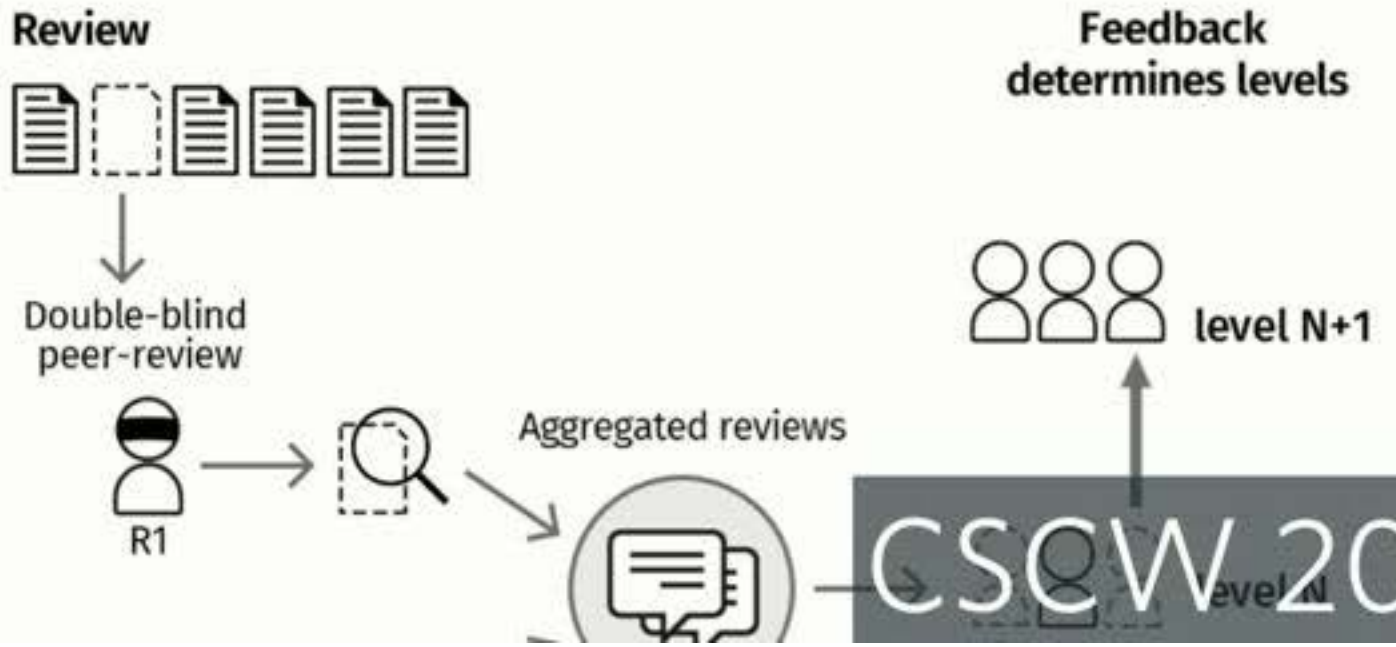

CSCW 2017

Boomerang: Rebounding the Consequences of Reputation Feedback on Crowdsourcing Platforms

S.S. Gaikwad, D. Morina, A. Ginzberg, C. Mullings, S. Goyal, D. Gamage, C. Diemert, M. Burton, S. Zhou, M. Whiting, K. Ziulkoski, A. Ballav, A. Gilbee, S.S. Niranga, V. Sehgal, J. Lin, L. Kristianto, A. Richmond-Fuller, J. Regino, N. Chhibber, D. Majeti, S. Sharma, K. Mananova, D. Dhakal, W. Dai, V. Purynova, S. Sandeep, V. Chandrakanthan, T. Sarma, S. Matin, A. Nassar, R. Nistala, A. Stolzoff, K. Milland, V. Mathur, R. Vaish, M.S. Bernstein

Stanford Crowd Research Collective, Stanford University
daemo@cs.stanford.edu

ABSTRACT

Paid crowdsourcing platforms suffer from low-quality work and unfair rejections, but paradoxically, most workers and requesters have high reputation scores. These inflated scores, which make high-quality work and workers difficult to find, stem from social pressure to avoid giving negative feedback. We introduce Boomerang, a reputation system for crowdsourcing platforms that elicits more accurate feedback by rebounding the consequences of feedback directly back onto the person who gave it. With Boomerang, requesters find that their highly-rated workers gain earliest access to their future tasks, and workers find tasks from their highly-rated requesters at the top of their task feed. Field experiments verify that Boomerang

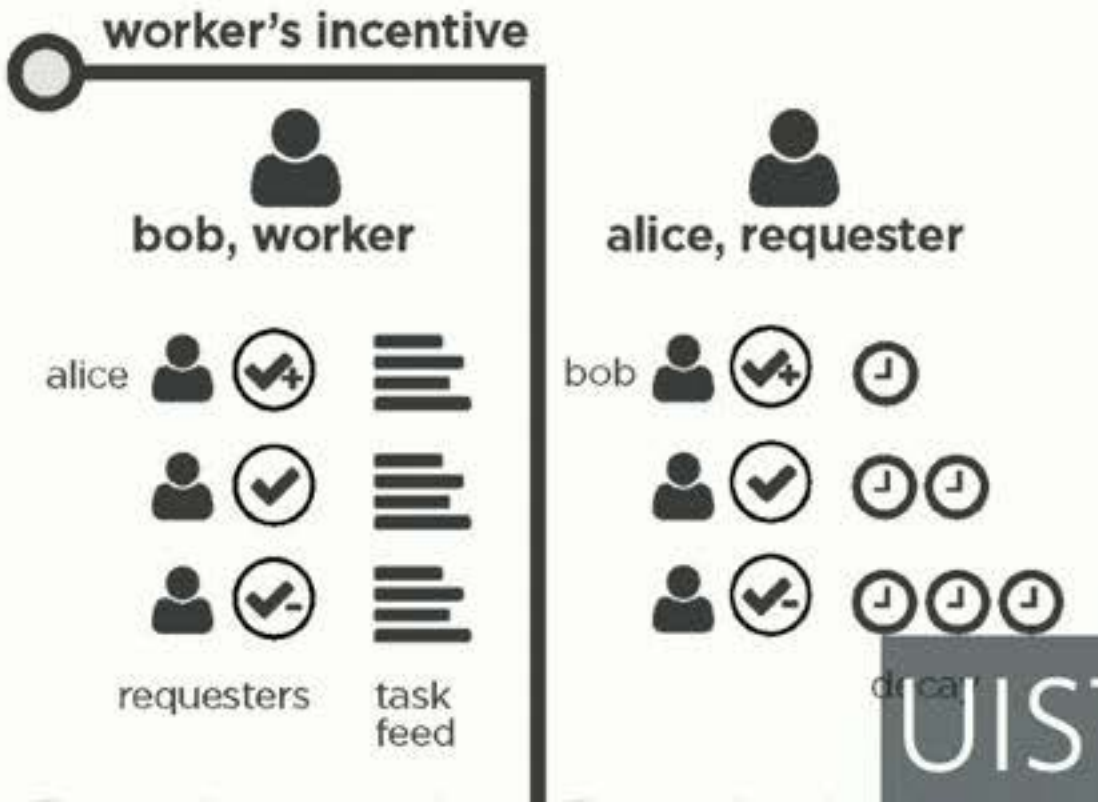

UIST 2016

WORKS-IN-PROGRESS


Daemo: a Self-Governed Crowdsourcing Marketplace

Stanford Crowd Research Collective
Stanford HCI Group daemo@cs.stanford.edu

HCI

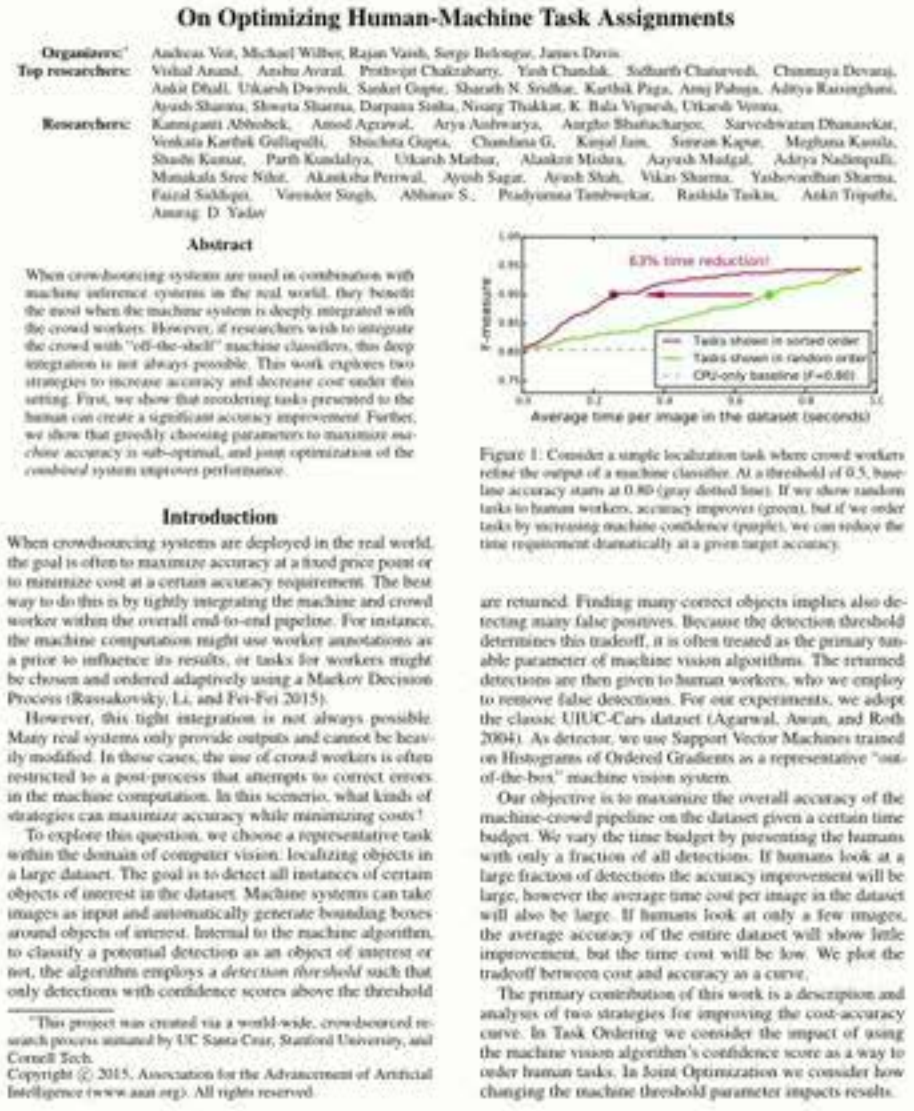
On Optimizing Human-Machine Task Assignments

Computer Vision


Investigating the "Wisdom of Crowds" at Scale

Alok Shankar Mysore, PES Institute of Technology, alok.shankar@pesit.pes.edu
Vikas S Yaligar, National Institute of Technology Karnataka, vikasyaligar@ieee.org
Imanol Arrieta Ibarra, Stanford University, imanol@stanford.edu
Camelia Simoiu, Stanford University, csimoiu@stanford.edu
Sharad Goel, Stanford University, sgoel@stanford.edu
Additional Authors, Various Institutions

Data Science

INCREASED ACCESS TO RESEARCH

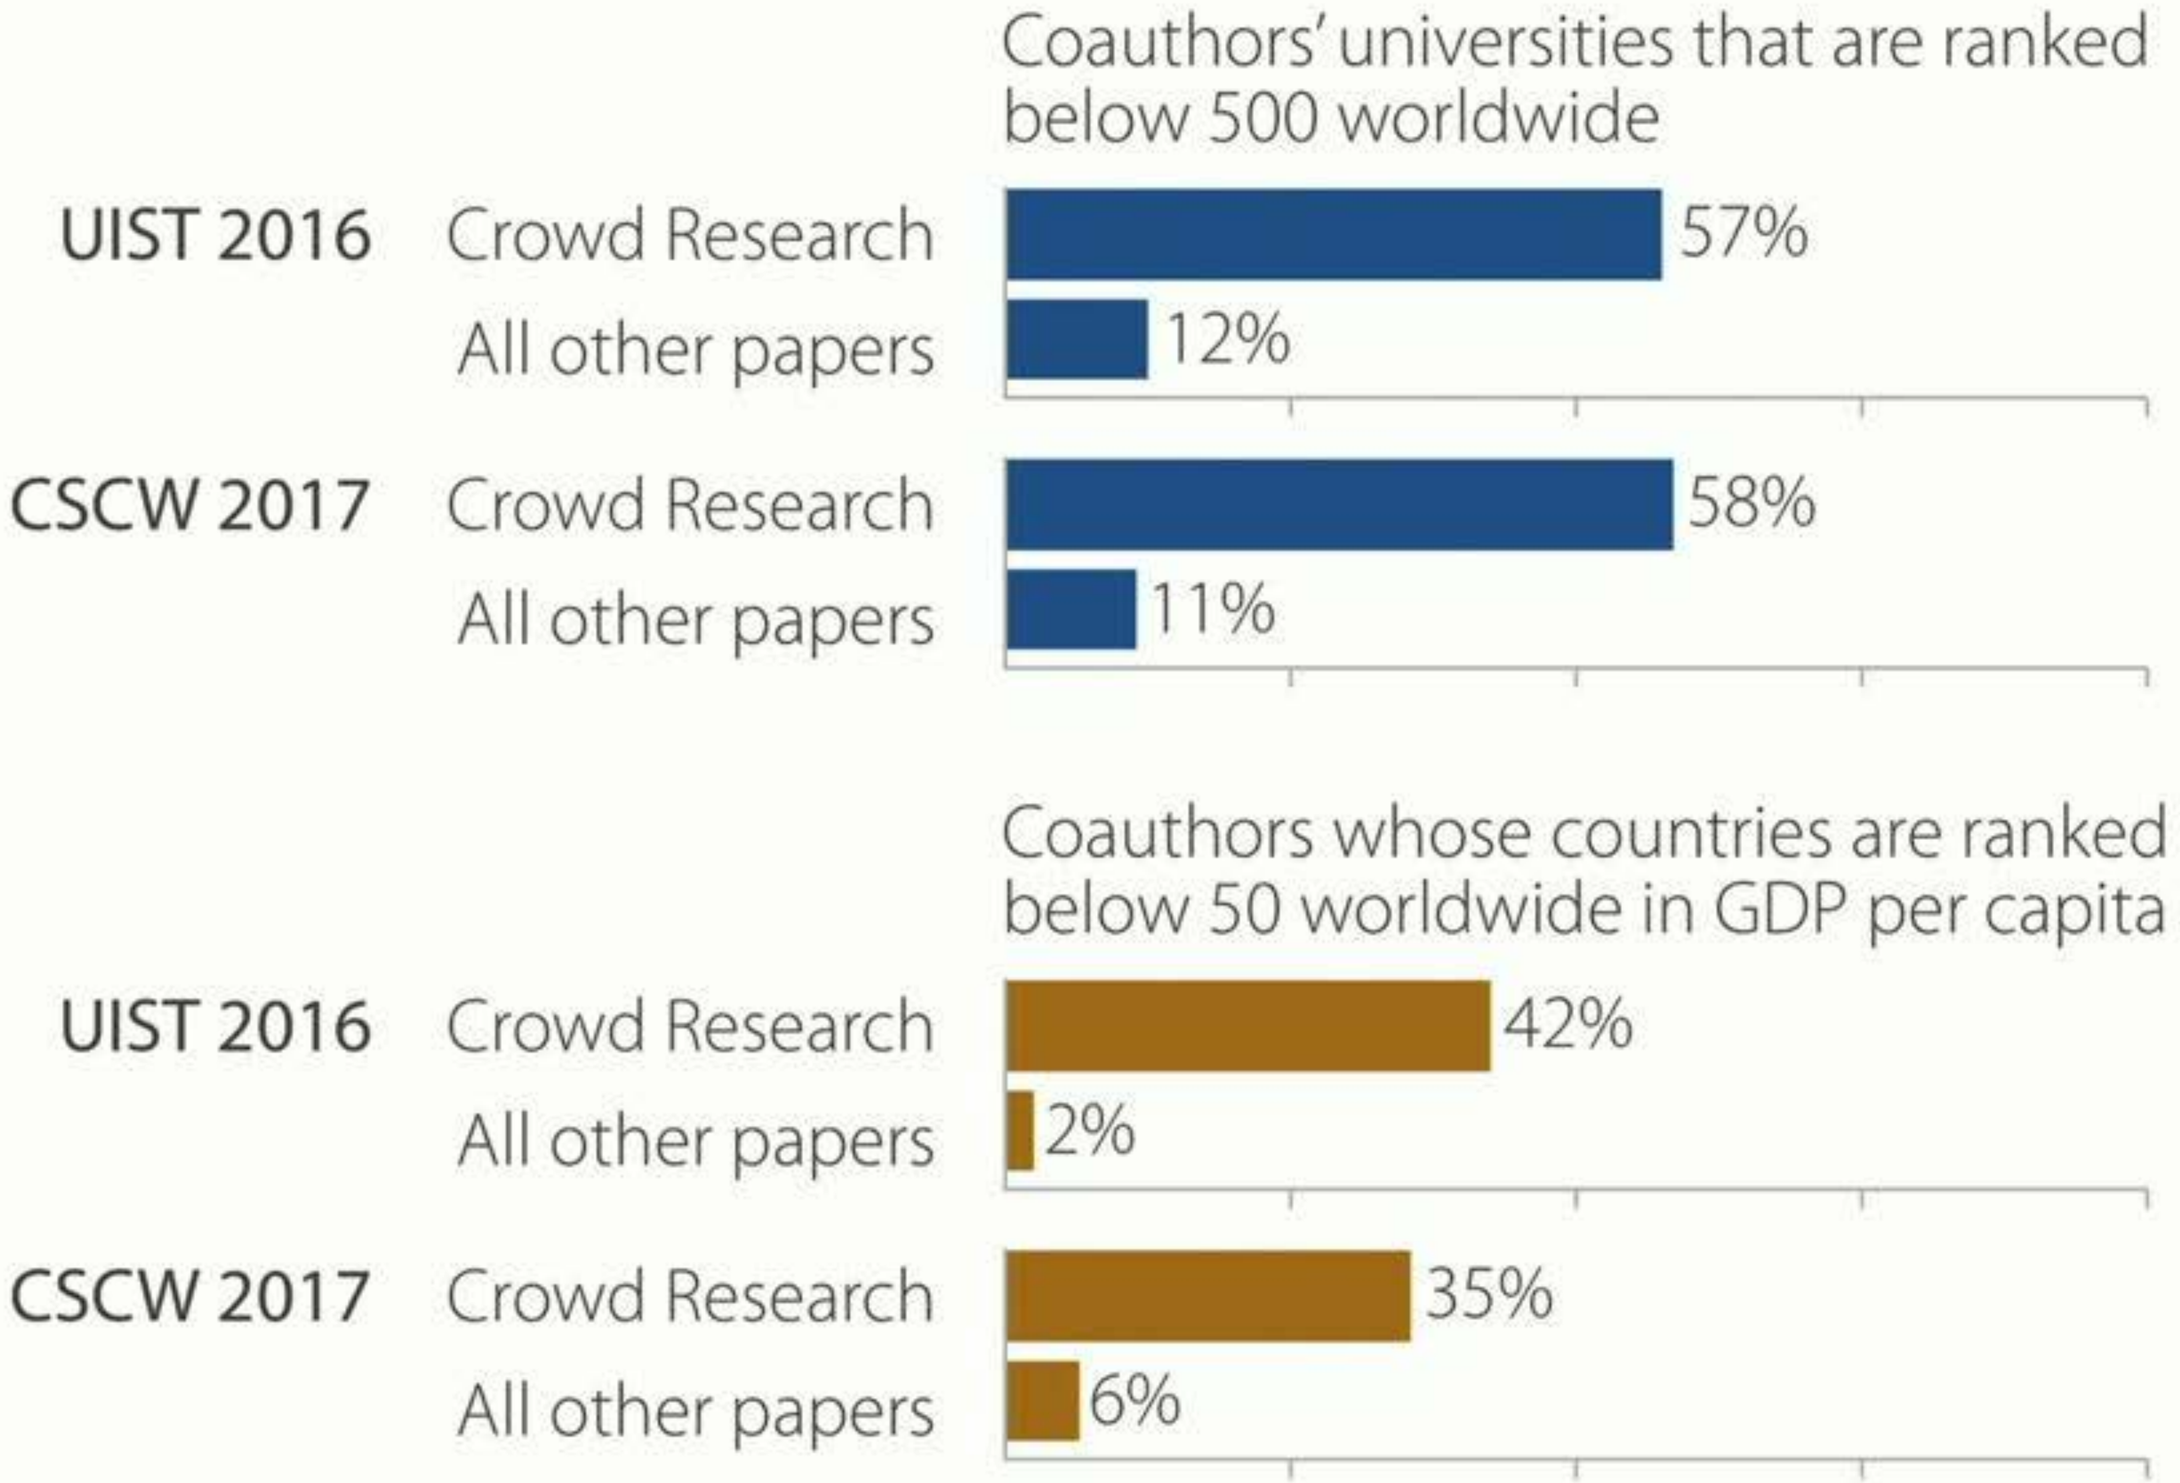

ATTAINMENT: PROVIDING ACCESS

Participants have gone on to programs at UC Berkeley, and Carnegie Mellon University, and MIT.

21 of 33 surveyed were admitted to at least one program, despite a **median of 0 other letter writers** from institutions ranked better than 500 worldwide.

DECENTRALIZED CREDIT: TRANSLATE INTO GRAPH PROBLEM

Each participants allocates **credit points** to other participants based on their assessment of who impacted the project

Resulted: weighted directed graph

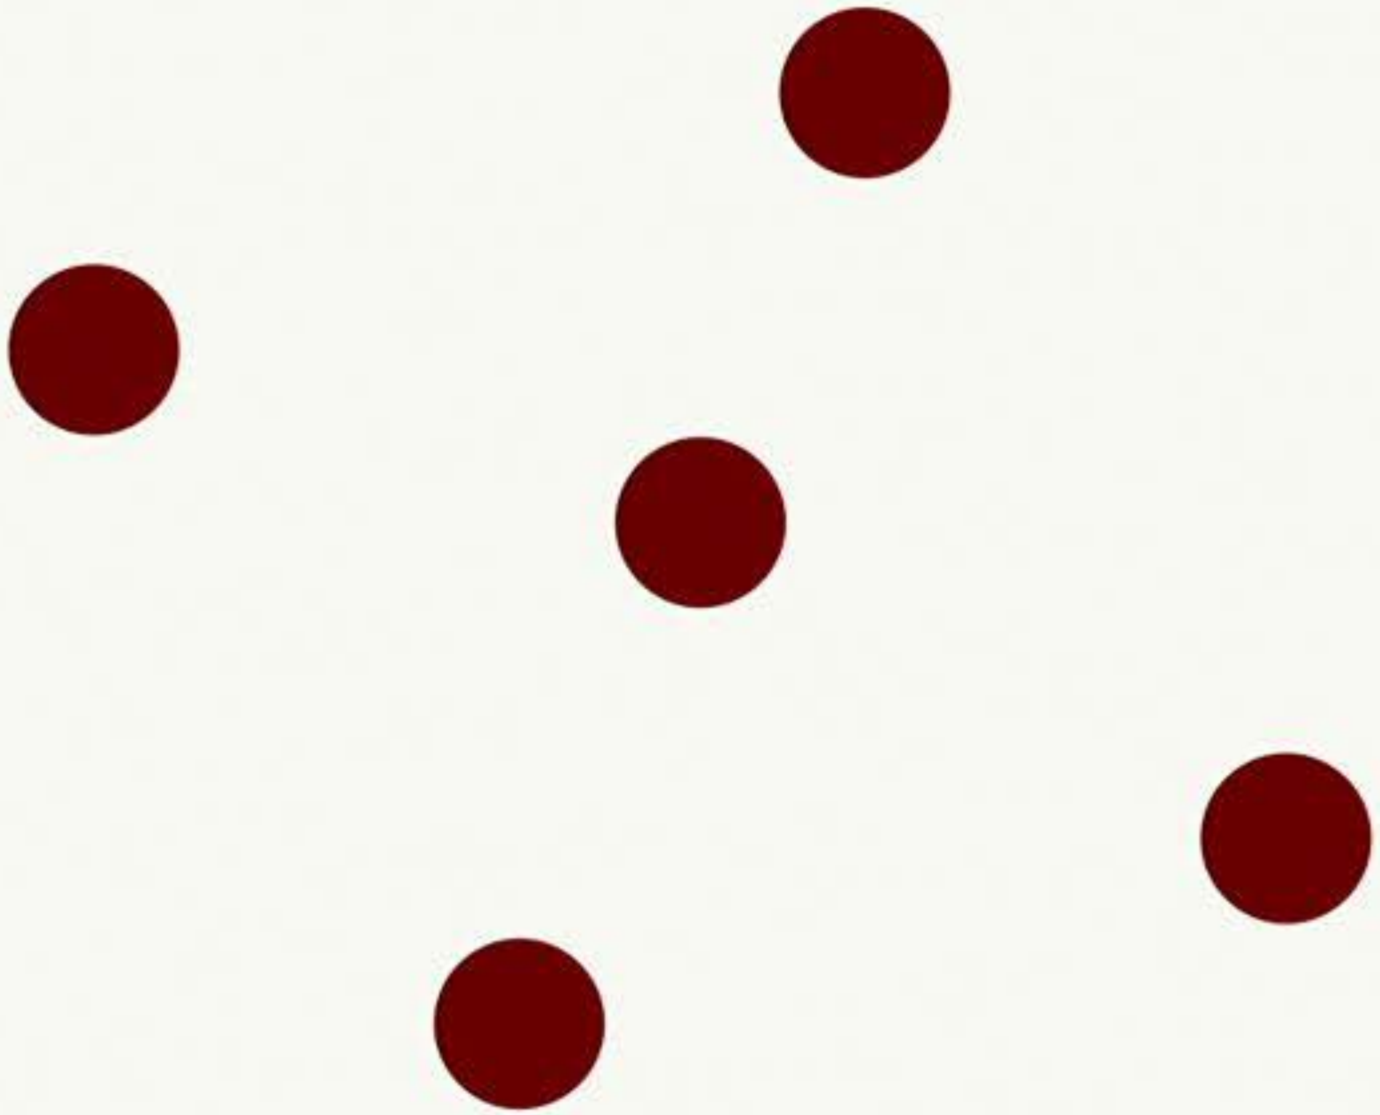

DECENTRALIZED CREDIT: TRANSLATE INTO GRAPH PROBLEM

Each participants allocates **credit points** to other participants based on their assessment of who impacted the project

Resulted: weighted directed graph

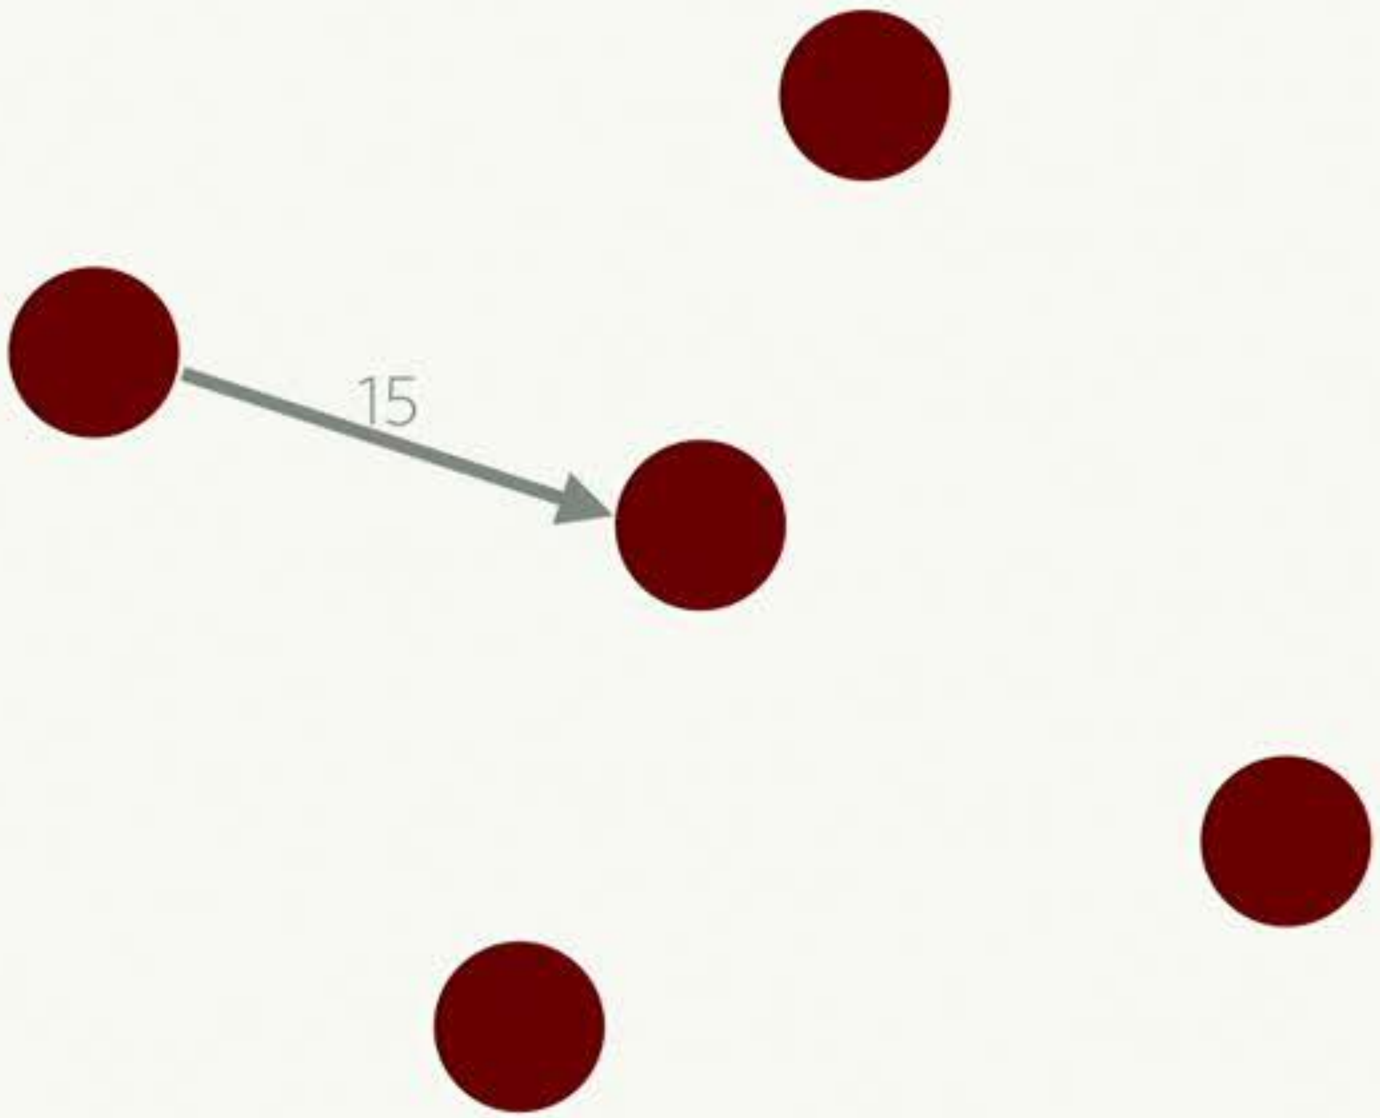

DECENTRALIZED CREDIT: TRANSLATE INTO GRAPH PROBLEM

Each participants allocates **credit points** to other participants based on their assessment of who impacted the project

Resulted: weighted directed graph

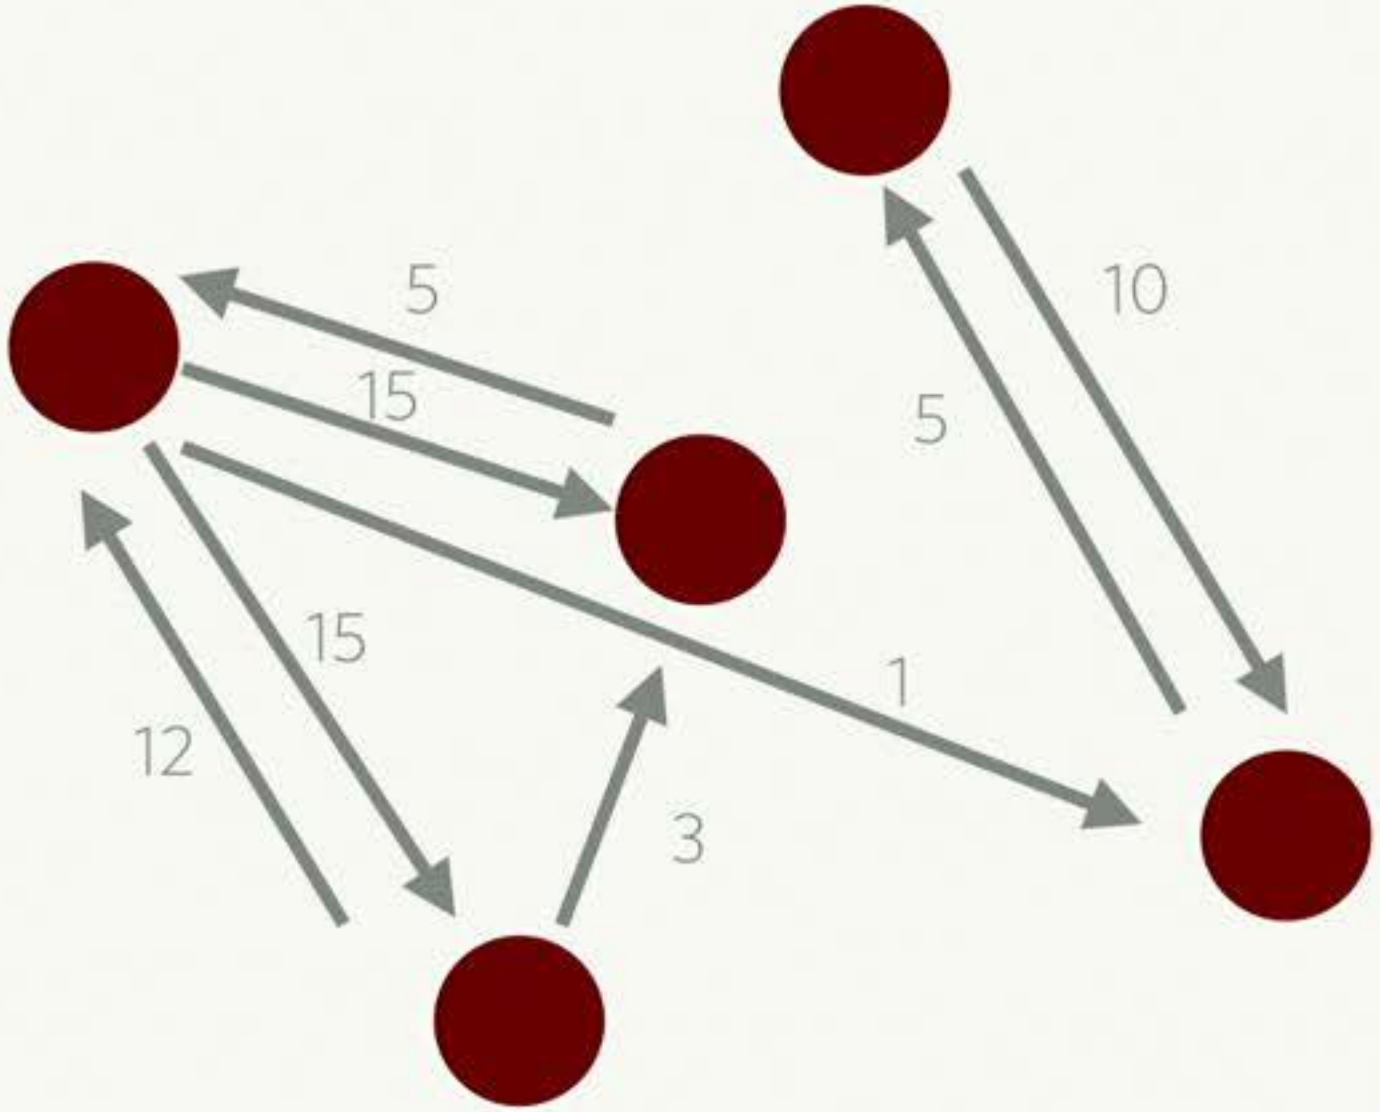

GRAPH CENTRALITY: PAGERANK

Intuition: identify nodes that are receiving large amounts of credit, weigh those nodes' allocations heavily, and iterate until convergence

Propagate each node's score in proportion to its outgoing wedge weights

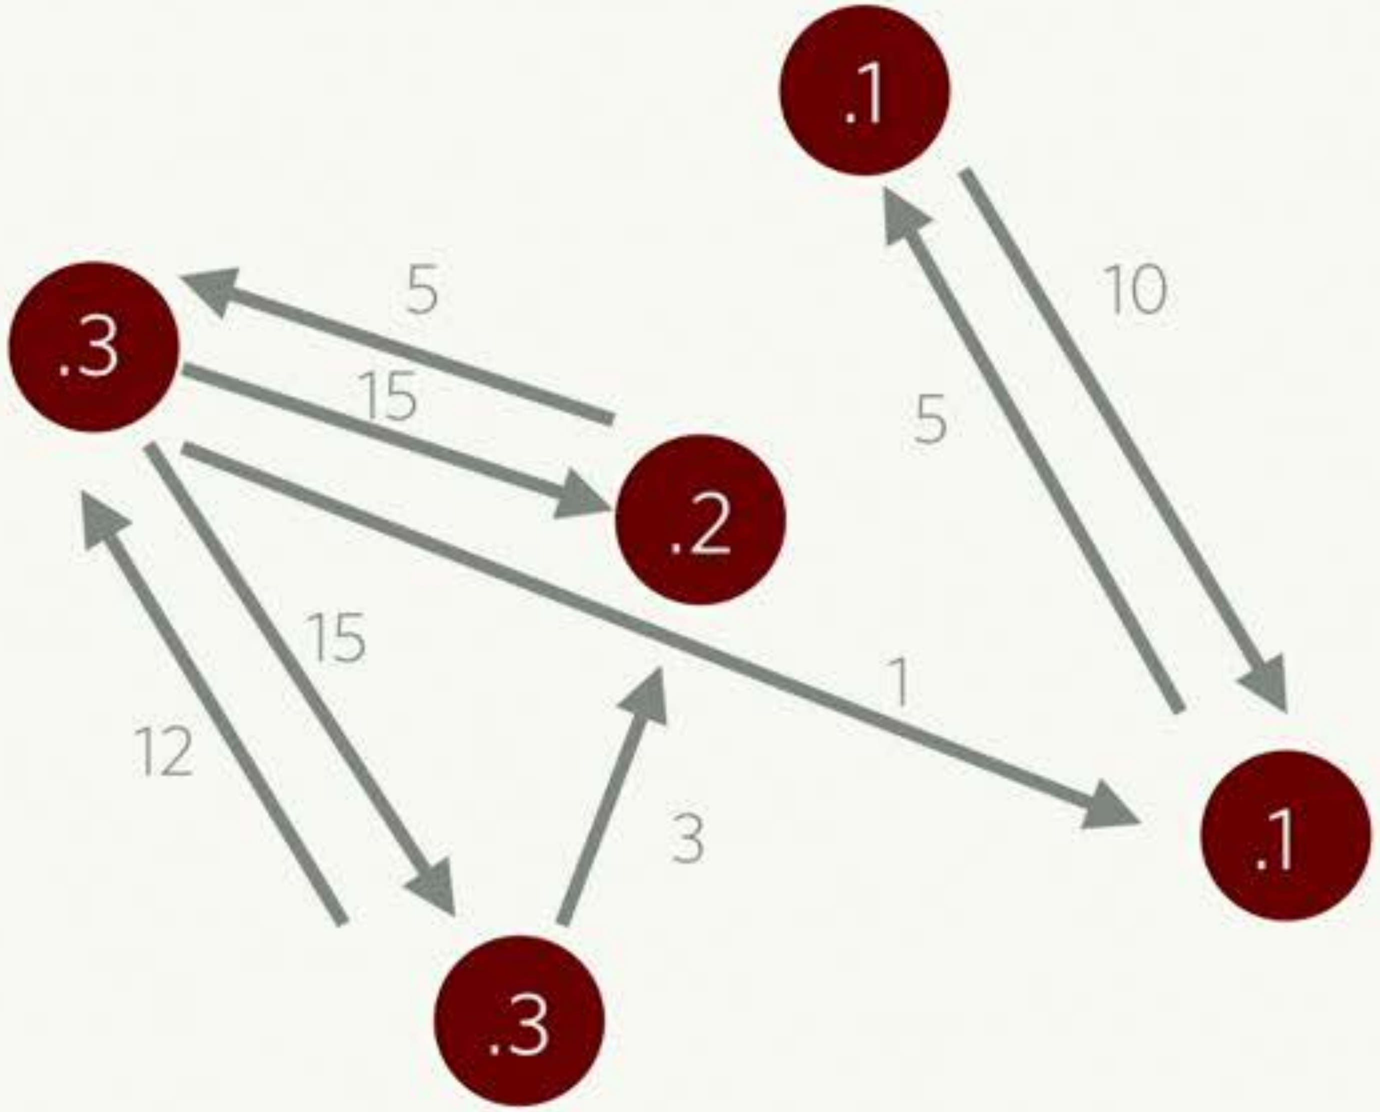

STRATEGIC BEHAVIOR

Speaking different languages or otherwise interacting with only a small part of the crowd: link ring

Strategically directing credit toward those who will return credit to you: such attacks occur in 360-degree reviews

Formulations of centrality algorithms can correct for most of these attacks

creddit.stanford.edu

ANALYZING CREDDIT'S EFFECT

What impact did Creddit have on credit distribution?

Method: normalize raw summed credit scores, and Creddit-adjusted scores, to sum to 1.0

Regress both raw score and Creddit score on observable collaboration behaviors, and compare β estimates across the regressions

LESS TALKING, MORE DOING

Participation Measure	Creddit: β_C	Raw Votes: β_{raw}	β_C-β_{raw}
# Hangouts	0.0694***	0.0438*	0.0256
# Files Uploaded	0.0352**	0.0293*	0.0059
# GitHub commits	0.0171	-0.024*	**0.0411*** **
# Slack messages	0.0351*	0.1122***	**-0.0770*** **
# self-organized meetings	0.0239*	0.0115	0.0123
Milestone leader (binary)	0.0360***	0.0059	**0.0300** **
Weeks active	0.0252*	0.0141	0.011

All variables standardized

EFFECTS ON AUTHOR ORDER

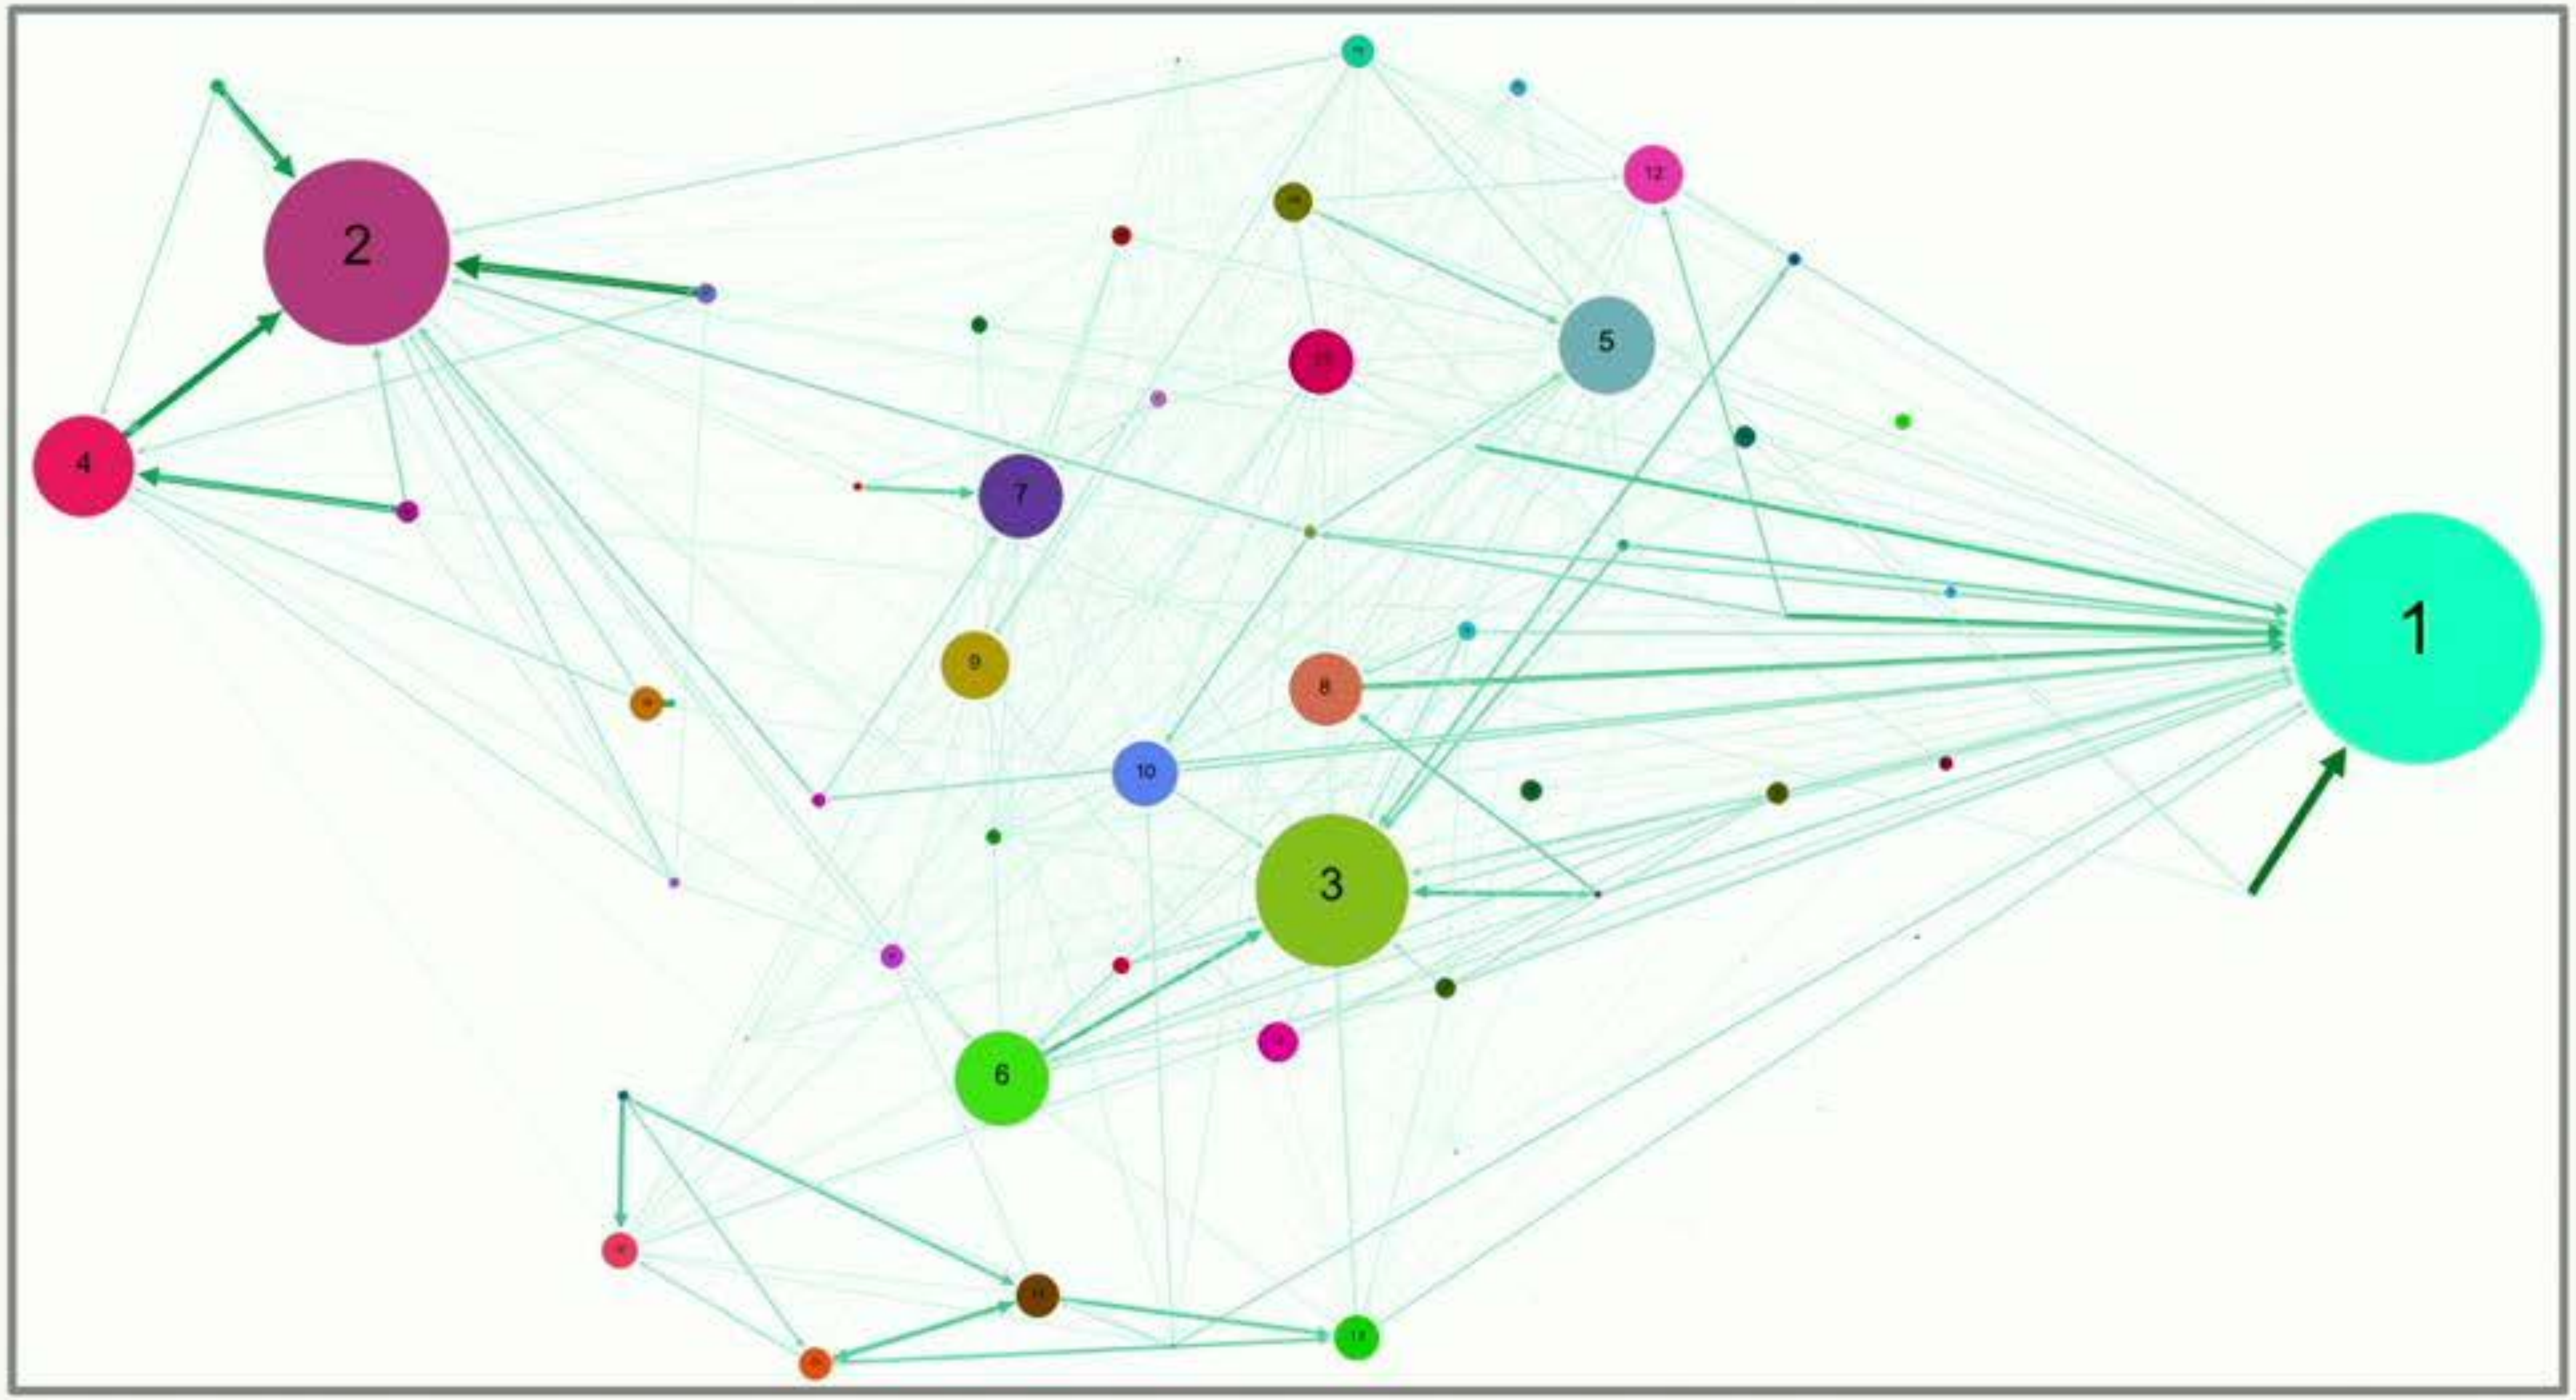

Raw vote ranking

#2 and #4 have a high rank due to link ring

EFFECTS ON AUTHOR ORDER

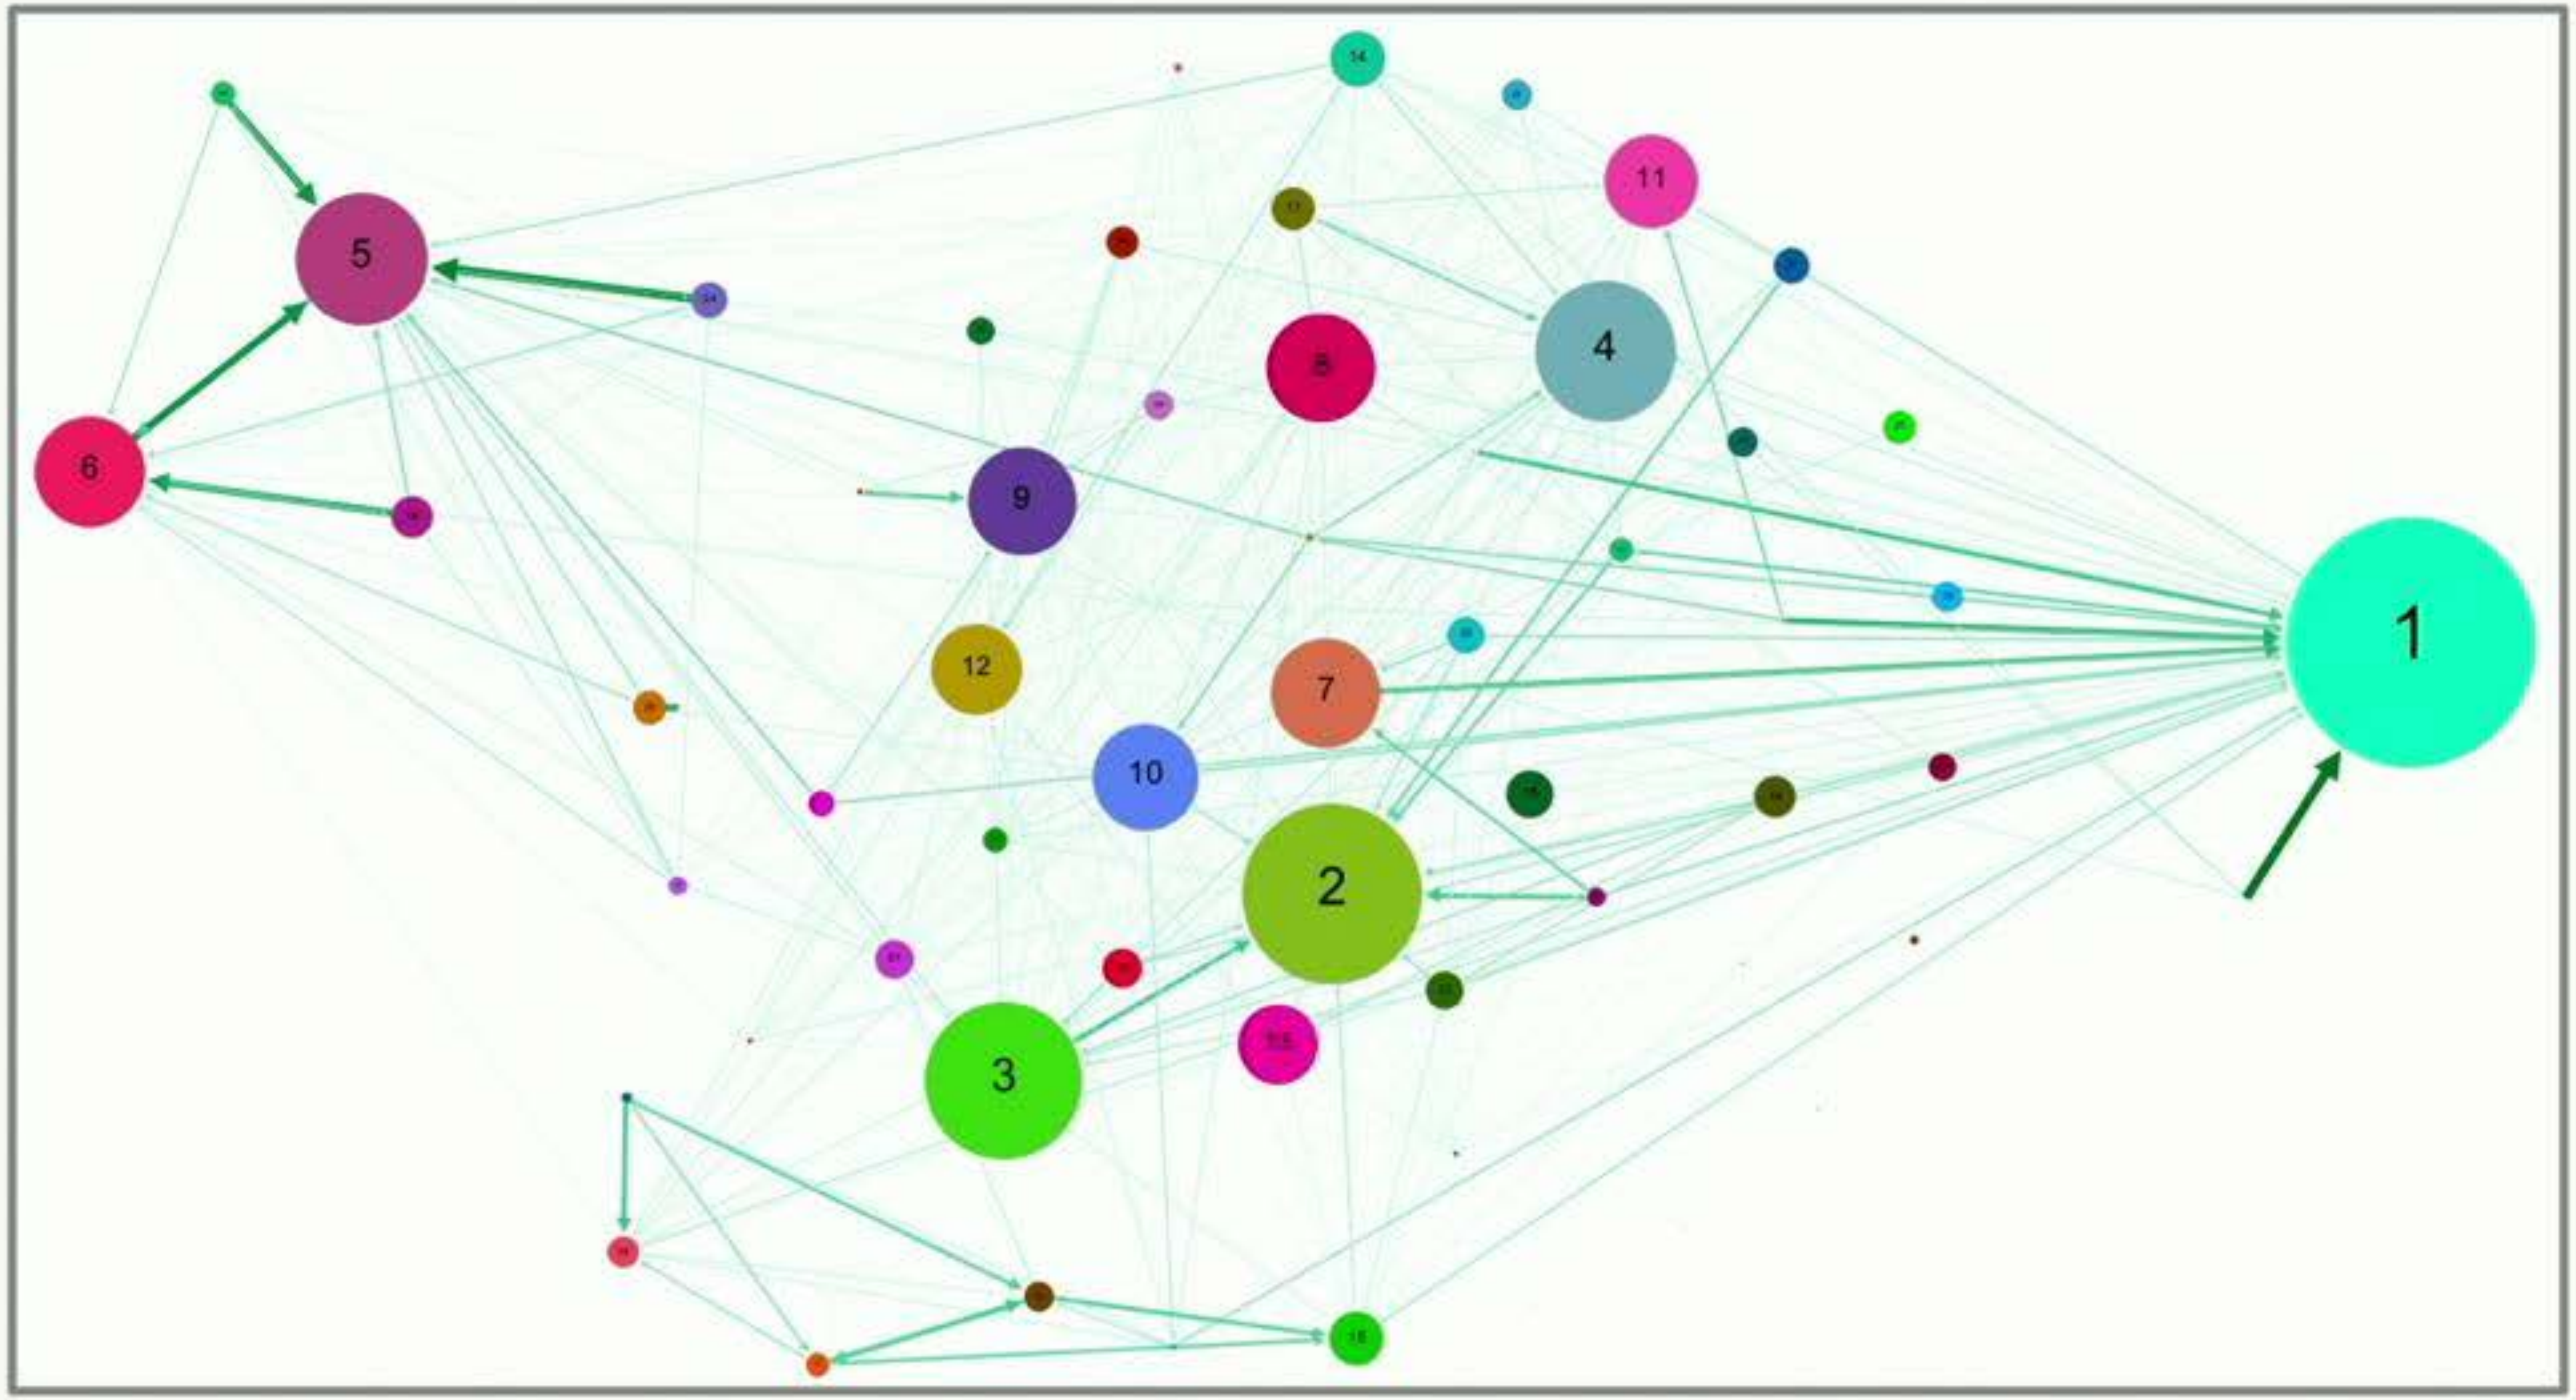

PageRank-corrected author order

Influential coauthors reduced impact of link ring

CROWD RESEARCH: REFLECTIONS

What impact would decentralized credit have on traditional teams or organizations?

What kinds of research projects can operate at a larger technical scale than traditional CS research?

Rather than structuring the future of work as an algorithm, how might we design computationally augmented organizational structures?

Rather than structuring the future of work as an algorithm, how might we design computationally augmented organizational structures?

Organizations were originally designed with inspiration by mechanical systems. What might a computational infrastructure offer them?

Crowds, Computation, and the Future of Work

Thanks to…

Amazing students: Daniela Retelny, Niloufar Salehi, Rajan Vaish, Ranjay Krishna, Sharon Zhou

Amazing colleagues: Melissa Valentine, Fei-Fei Li, James Davis, Sharad Goel

Amazing supporters: NSF, Sloan, Accenture ATL, Microsoft, Toyota Human-Centered AI, Stanford Cyber Initiative, ONR, DARPA, Hasso Plattner Design Thinking Program

Amazing questions

THREE PARALLEL PROJECTS

HCI

Michael Bernstein, Stanford

Building a new crowd marketplace

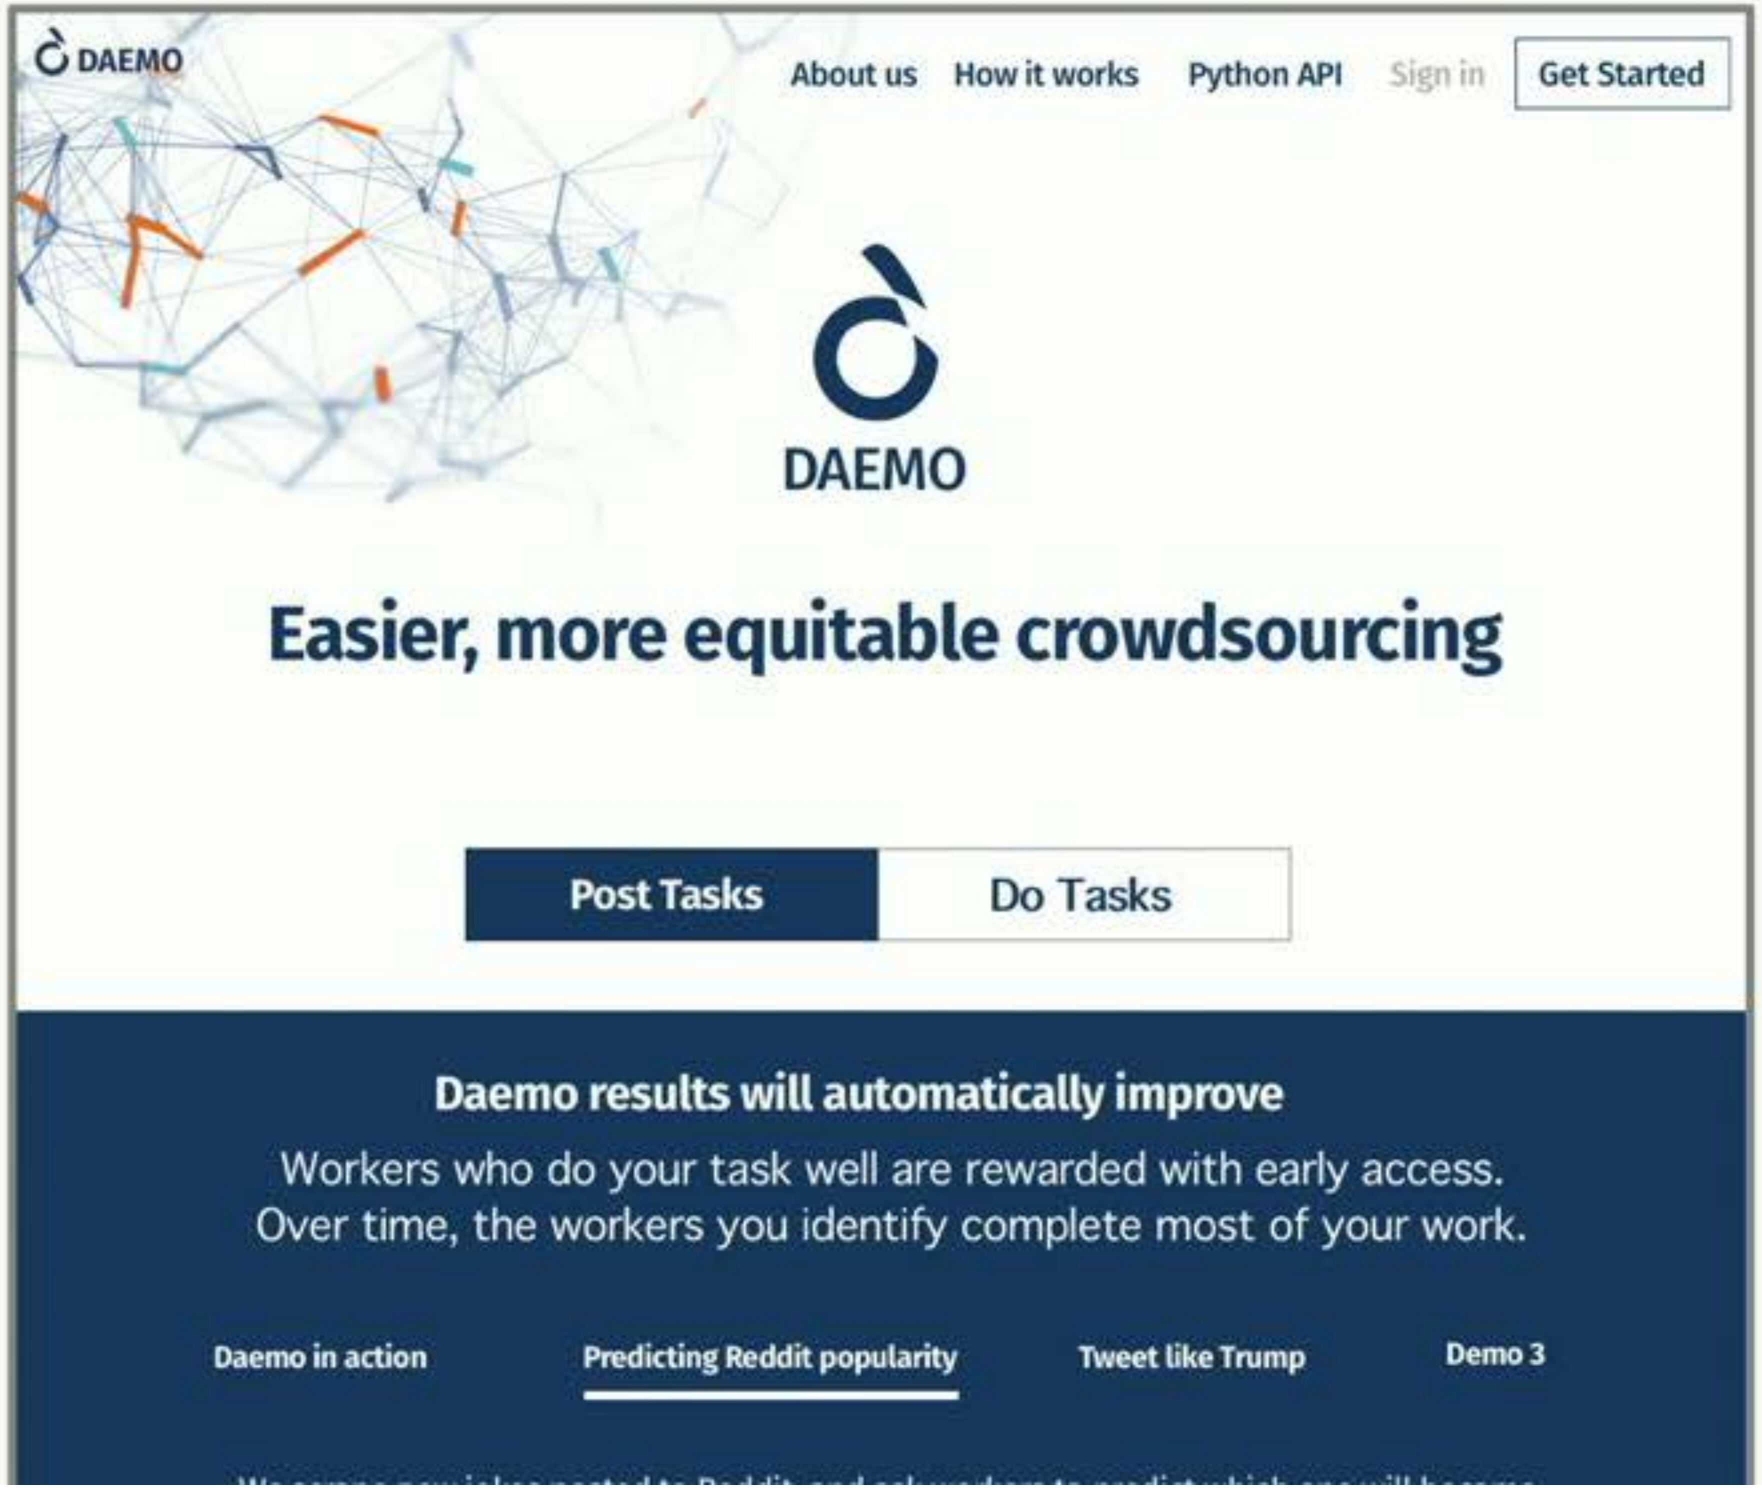

Crowds, Computation, and the Future of Work

Thanks to…

Amazing students: Daniela Retelny, Niloufar Salehi, Rajan Vaish, Ranjay Krishna, Sharon Zhou

Amazing colleagues: Melissa Valentine, Fei-Fei Li, James Davis, Sharad Goel

Amazing supporters: NSF, Sloan, Accenture ATL, Microsoft, Toyota Human-Centered AI, Stanford Cyber Initiative, ONR, DARPA, Hasso Plattner Design Thinking Program

Amazing questions