RESEARCH STATEMENT MICHAEL BERNSTEIN

Can computation amplify our collective achievements? From massive crowds to small teams, groups of people convene to work on the world's hardest problems. Each of these groups faces a multiplicity of possible designs for how its members will interact, including how they will coordinate, socialize, and make decisions. Choosing the right design is critical: poor decisions lead to fraught interactions at best and outright dissolution at worst. Unfortunately, we are ill-equipped to choose between these designs: just like human cognition can give rise to biased decision-making at an individual level, the social sciences demonstrate that even experts often exhibit poor decision-making about designing interactions at a group level. The core challenge is emergence: what design will produce individual behaviors that aggregate into the desired group behavior?

My research in human-computer interaction applies a computational lens to help groups achieve their goals. Groups now increasingly interact online—through collaboration tools, crowdsourcing markets, or social network sites—providing an opportunity to inject computational systems that mediate their interactions. Developing these systems sheds light on fundamental design and engineering questions about group behavior: What systems help or hinder coordination? How do we manage anti-social behavior in social environments? Could an adaptive, intelligent system support more effective collaboration? In answering these questions, my goal is to empower ambitious crowdsourcing efforts, lay the foundation for resilient and enjoyable social computing environments, and enable anyone with access to the internet to contribute to challenging and fulfilling collective goals.

CROWDSOURCING SYSTEMS: FROM MICROTASKS TO ORGANIZATIONS. Crowdsourcing is the strategy of inviting contributions from a large group of people online via an open call. This strategy has proven effective when labeling images or classifying types of galaxies, where the goal can be decomposed into a set of small tasks known as microtasks and computation can then weave those microtasks together into a cohesive whole. However, microtasking has hit a complexity limit: it requires goals to be so simple and modular that the path to achieve them can be entirely pre-defined. As a result, crowdsourcing systems have struggled to achieve open-ended and complex goals whose paths cannot be predefined, including design and engineering.

My work enables coordination toward these open-ended and complex goals by designing crowdsourcing systems that draw inspiration from modern modes of coordination in organizational science. Our systems convene on-demand domain experts rather than non-expert workers [Retelny et al. UIST 2014] and coordinate those experts using computational organizational structures rather than microtask workflows [Valentine et al. CHI 2017]. The systems enable the resulting groups, which we call flash teams and flash organizations, to engage in adaptive coordination and achieve goals previously out of reach for crowdsourcing, including product design, software engineering, and game production. This work won best paper awards at UIST 2014 and CHI 2017, has been covered by the New York Times, and has been applied at large organizations (e.g., Accenture) for mobile application production and at startups (e.g., B12.io) for web design.

As colleagues and I argued [Kittur et al. CSCW 2013], it is important to ensure a pro-social future for crowdsourcing's participants. Our work demonstrated the first socio-technical system to support collective action for crowd work, called We Are Dynamo [Salehi et al. CHI 2015]. We then developed Crowd Research [Vaish et al. UIST 2017] to support upward educational and career mobility worldwide through access to online research experiences with faculty at R1 universities. To enable Crowd Research to recognize top contributors at scale, we developed a graph-based peer credit system. Our further work demonstrated systems for crowd worker mentorship and guild accreditation. This research has practical impact: crowd workers on We Are Dynamo succeeded in gaining media attention to their movements, and Crowd Research has enabled over 1,500 people from 62 countries to collaborate on research appearing at top-tier Computer Science venues.

I expect that the lines between crowds and traditional organizations will continue to blur. Viewed from the right angle, an organization is a crowd: a large group of retained experts who can flexibly assemble and reassemble themselves. Leveraging this point of view, I see opportunities for computational systems to enable new forms of collective intelligence that could outstrip modern organizations in their prevalence, impact and achievement. To begin, we are developing data-driven techniques that can help people learn from a massive database of past decisions about how to design collaborations (e.g., How many people? How should we organize them?), as well as rethinking the reputation systems that enable career growth for workers.

DESIGNING SOCIAL SYSTEMS. Crowdsourcing systems focus on coordination, but group activity is fundamentally social. Just like broken windows and negative attributions can ruin a neighborhood, too often our social environments lead to friction, harassment, and frustration. My goal is to create collaborative and social computing systems to support more positive social interactions between participants. Unfortunately, achieving this is not straightforward: designs that succeed with one group will fail with another. So, I study the conditions leading to success and failure, and I apply that knowledge to design new socio-technical systems.

At the team scale, I focus on systems that influence team structures. For example, is it best for a team to be flat or hierarchical? Encouraging or critical? Enforcing equal turn-taking, or not? Poor choices will doom a team to dysfunction. Our work on DreamTeam drew on structural contingency theory to inspire a system where teams self-experiment to rapidly identify the right set of team structures, and introduced a model of temporally-constrained multi-armed bandits to avoid overwhelming teams with too much change [Zhou, Valentine and Bernstein CHI 2018]. Membership is another important determiner, so we introduced network rotation, an algorithm that iteratively weaves a collaboration network by rotating team membership to bring team members into contact with diverse perspectives [Salehi and Bernstein CSCW 2018].

At the web scale, recent events have reminded us that social computing systems such as Facebook, Instagram, and Wikipedia are brittle, manipulated by trolling into anti-social behavior. In contrast to prior work and common wisdom that anti-social behavior is due to a small class of trolls—inveterate sociopaths—our work combined online experiments, data mining of 16 million posts on CNN.com, and predictive modeling to demonstrate that a substantial proportion of trolling behavior is due to ordinary people: that anyone can become a troll [Cheng et al. CSCW 2017].

This work has been awarded best paper at CSCW 2017, used by Mozilla to generate new accessibility concepts, and covered by news outlets including PBS, the Wall Street Journal and Technology Review. In future work, I will deepen our work on anti-social computing: social computing systems amplifying negative social behavior. As a computer security researcher might put it, groups suffer from vulnerabilities: malicious interventions can influence teams to dissolve or community members to lose trust in each other. How might we design social computing systems that make us more resilient to these kinds of attacks? Could such systems help us better manage harassment and misinformation, and forestall a collapse of trust?

SOCIAL COMPUTING, DATA AND A.I. Co-designing social computing and crowdsourcing systems with domain scientists leads to advancements in both fields. We developed crowdsourcing methods that increase training label generation for machine learning by an order of magnitude [Krishna et al. CHI 2006], as well as crowdsourcing methods used to gather data for the computer vision ImageNet Challenge [Russakovsky et al. IJCV 2015] and for a popular dataset of visual scene graphs [Krishna et al. IJCV 2017]. We also work to make data-driven tools more accessible: examples include enabling computational social scientists to combine the flexibility of modern word embeddings with the intelligibility of lexicons [Fast, Chen, and Bernstein CHI 2016], and using fiction to add breadth to knowledge bases of human activity [Fast et al. CHI 2016]. This work won best paper at CHI 2016 and has been adopted by Facebook's Data Science team.

Our future work will aim to amplify collective intelligence with advances in machine learning and data science. For example, we are beginning work in healthcare, where group decision making is challenging and decentralized. How might such systems combine multiple healthcare professionals' points of view?

CONCLUSION. Weaving this research together is a thread connecting technical innovation to group behavior. Whether through crowds, social networks, teams, or data, my research is animated by the goal of designing systems that catalyze the positive interactions that we want to have with each other as groups and as a society.