

The Future of Crowd Work

Aniket Kittur¹, Jeffrey V. Nickerson², Michael S. Bernstein³,
Elizabeth M. Gerber⁴, Aaron Shaw⁴, John Zimmerman¹, Matthew Lease⁵, and John J. Horton⁶

¹Carnegie Mellon University, ²Stevens Institute of Technology, ³Stanford University,
⁴Northwestern University, ⁵University of Texas at Austin, ⁶oDesk
{nkittur, johnz}@cs.cmu.edu, jnickerson@stevens.edu, msb@cs.stanford.edu, {egerber,
aaronshaw}@northwestern.edu, ml@ischool.utexas.edu, john_horton@odesk.com

ABSTRACT

Paid crowd work offers remarkable opportunities for improving productivity, social mobility, and the global economy by engaging a geographically distributed workforce to complete complex tasks on demand and at scale. But it is also possible that crowd work will fail to achieve its potential, focusing on assembly-line piecework. Can we foresee a future crowd workplace in which we would want our children to participate? This paper frames the major challenges that stand in the way of this goal. Drawing on theory from organizational behavior and distributed computing, as well as direct feedback from workers, we outline a framework that will enable crowd work that is complex, collaborative, and sustainable. The framework lays out research challenges in twelve major areas: workflow, task assignment, hierarchy, real-time response, synchronous collaboration, quality control, crowds guiding AIs, AIs guiding crowds, platforms, job design, reputation, and motivation.

Author Keywords

Crowdsourcing; crowd work; organization design; research vision

ACM Classification Keywords

H.5; K.4.3

INTRODUCTION

Crowdsourcing rapidly mobilizes large numbers of people to accomplish tasks on a global scale. For example, volunteer-based collective projects such as Wikipedia owe their success and longevity to the ongoing efforts of thousands of individual contributors around the world.

Complementing volunteer-based crowdsourcing, a paying *crowd work* industry is now quickly growing in scope and ambition. Crowd work today spans a wide range of skill and pay levels, with commercial vendors providing access

to a range of workers and focused support for various task. For example, anyone with access to the Internet can perform *micro*-tasks on the order of seconds using platforms such as Amazon’s Mechanical Turk, while more skilled workers can complete multi-hour tasks on professional online marketplaces such as oDesk or work for months to solve R&D challenges on open innovation platforms (e.g. Innocentive). Incentives and work structures also vary tremendously, ranging from crowdsourcing contests awarding prizes to winners (e.g. programming tasks on Topcoder) to micro-labor platforms that pay workers per task.

While not all jobs are amenable to being sent down a wire, there are portions of almost any job that can be performed by the crowd [63]. We foresee a world in which crowd work continues to expand, unlocking an incredible number of opportunities for careers and skilled work in online marketplaces. However, we also foresee a serious risk that crowd work will fall into an intellectual framing focused on low-cost results and exploitative labor. With diminished visibility and communication channels vis-a-vis traditional workplaces, workers may be treated as exchangeable and untrustworthy, having low or static skill sets and strong motivations to shirk. Workers may become equally cynical, having fewer bonds, enforceable contracts, and power than with traditional workplaces [130]. Such concerns may grow ever sharper unless this trajectory is somehow altered.

This work originally emerged from the question: “*Can we foresee a future crowd workplace in which we would want our children to participate?*” We suggest that this question has a number of attractive properties as a banner around which to rally research, as well as serving as an anchor to ground speculation. It is simple enough to convey concisely, involving an evaluative component that everyone, with or without children, can make. We intentionally keep the “we” ambiguous so that readers with different values and cultural contexts may “try on” the question, providing a conceptual lens easily refocused, and encourage discussion regarding the challenges this question poses to deciding what is valuable and pride-worthy work. What new services, systems or features are needed for a future of crowd work that the reader would be proud to see

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW '13, February 23–27, 2013, San Antonio, Texas, USA.
Copyright 2013 ACM 978-1-4503-1331-5/13/02...\$15.00.

his or her children take on as their livelihood? Is this a desirable path for the next generation?

Looking toward that positive future, in this paper we contribute an analytic framework for research in crowd work. We lay out a vision for a possible future for crowd work that entails:

- **Worker considerations**, such as motivation, feedback, and pay. These may be addressed by mechanisms to maintain reputation, provide better interaction with requesters, and increase motivation.
- **Requester considerations**, such as coordination, task decomposition, and quality control. These may be addressed through workflow mechanisms including electronically mediated collaboration.

Our analytic framework is organized around a multidisciplinary survey of the literature that speaks to these challenges and helps to envision a positive future. We also include specific comments from crowd workers we surveyed in order to elicit their thoughts and suggestions. We translate our findings into a set of pragmatic design considerations that we believe are crucial in guiding design and motivating research in this field. We are following in the tradition of a set of research contributions which delineated design principles as part of a call for action [69,98,102,114,123].

Crowd Work

A variety of terminology is currently used in regard to crowds, e.g. crowdsourcing, collective intelligence, human computation, serious games, peer production, and citizen science [2,12,90,105]. We focus this paper on *paid, online crowd work*, which we define here as the performance of tasks online by distributed *crowd workers* who are financially compensated by *requesters* (individuals, groups, or organizations). In this sense, crowd work is a *socio-technical work system* constituted through a set of relationships that connect organizations, individuals, technologies and work activities [142]. Online crowd work takes place in marketplaces that allow requesters to seek workers and support workers in finding work. For this paper, we surveyed a number of contemporary, popular crowd work platforms. These platforms include general-purpose marketplaces (e.g., Mechanical Turk, oDesk, Freelancer, Crowdflower, MobileWorks, ManPower) as well as markets for specific expertise (e.g., TopCoder, uTest, 99Designs). While these platforms are intended for legitimate tasks, these and other platforms are sometimes appropriated for illegal or nefarious purposes (e.g., gold farming, CAPTCHA solving, and crowdturfing) [35].

Through our definition, we necessarily omit a wide range of voluntary crowd work, such as wikis [22], games with a purpose [2], captchas [3], and citizen science [31,106,122]. Much has already been written about these systems (e.g., [12,108]). However, not only is paid work the cornerstone of our existing economy and labor markets, but even

volunteers typically engage in some form of paying work in order to sustain themselves. Moreover, we believe there will always be some forms of work needed by society which are not amenable to gamification and volunteering, and for which demand will outstrip supply via unpaid channels. As such, we are interested in developing a future of paid crowd work that extends paid work into the online environment. In addition, we exclude offline crowd work, such as day labor, as it does not possess the same opportunities for distribution and global scalability as online work. For the remainder of the paper, we will use the term *crowd work* to refer to the performance of online tasks by crowd workers who are financially compensated by requesters.

We intentionally focus coverage on areas that may be of greater interest to the CSCW community, especially issues related to computer science, psychology, and organization science. We also draw on other important areas where appropriate (e.g., labor economics, ethics, law) and acknowledge these as critical to the future economy. Many aspects of these issues lie beyond the traditional purview of scientists or designers (e.g., labor regulations); however, we recognize that addressing them will be necessary for a positive future of crowd work.

Pros and Cons

Crowd work has the potential to support a flexible workforce and mitigate challenges such as shortages of experts in specific areas (e.g., IT work) or geographical locations. For individuals, crowd work also creates new opportunities for income and social mobility in regions of the world where local economies may be stagnant and local governmental structures discourage investment.

However, crowd work can be a double-edged sword, equally capable of enhancing or diminishing the quality of workers' lives. We may see echoes of past labor abuses in globally distributed crowd work: extremely low pay for labor, with marketplaces such as Amazon's Mechanical Turk reported to effectively pay an average of \$2/hour [65][126] with no benefits or worker protections. The per-task payment structure used in most crowd work markets is akin to piecework compensation in manufacturing [118], and can offer an invitation for gaming behavior which can negatively influence quality [78]. Moreover, crowds can be deployed in the service of questionable goals: to break captchas, to mine gold in games, and even, potentially, to locate dissidents [158]. The recent film "In Time" (2011) provided a pop culture depiction of how such a society might function where continual performance of menial tasks was literally required for worker survival. Many writers have painted similarly bleak pictures [40,136,137].

Crowd work may also displace current workers and has the potential to replace some forms of skilled labor with unskilled labor as tasks are decomposed into smaller and smaller pieces. Tasks such as speech transcription and copyediting are increasingly being accomplished with

crowd labor, and researchers have found that even some complex and expert tasks such as writing, product design, or translation may be amenable to novice crowd workers with appropriate process design and technological support [75,77,152,153].

This displacement is coupled to a new form of Taylorism [88,141], in which organizations optimize cognitive efficiency [157] at the expense of education and skill development. Taylorism yielded to more enlightened job design after several decades (and protracted struggles by workers), but given the short time commitment between crowd worker and requester, it is easy to imagine heightened exploitation and dehumanization.

As scientists, engineers, and designers, we can propose and evaluate new structures for crowd work and help imagine and bring about more positive futures. We can do so both through the intentional creation of desirable work environments as well as the cultivation of increased demand for work and workers. In particular, we suggest a role for researchers in conceptualizing and prototyping new forms of crowd work that go beyond the simple, independent, and deskilled tasks that are common today, with the goal of blazing a trail for organizations and platforms that will form the foundation of future crowd work.

ENVISIONING FUTURE CROWD WORK

How can we move towards a future of crowd work that is more attractive for both requesters and workers than existing systems? Even more ambitiously, can we design a future of crowd work that is more attractive and more effective than traditional labor systems?

Current crowd work typically consists of small, independent, and homogenous tasks, as shown in Figure 1. Workers are paired with an instance of each task to produce an output. Such simple, small-scale work has engendered low-pay, piece rate reward structures, in part due to the perception that workers are homogenous and unskilled. The current model is also insufficient to support the complexity, creativity, and skills that are needed for many kinds of professional work that take place today. Nor can it drive factors that will lead to increased worker satisfaction, such as improved pay, skill development, and complex work structures.

Theories of Organizational Behavior and Distributed Computing

Much professional work consists of complex sets of interdependent tasks that need to be coordinated across individuals with different expertise and capabilities [89]. For example, producing a book, an academic paper, or a new car all may involve many individuals working in structured teams, each with different skills and roles, collaborating on a shared output. To address these more complex goals we draw on concepts from both the organizational behavior [89,97,143,148] and the distributed computing literatures [9,132]. We propose that




Figure 1: Current crowd work processes.

crowdsourced labor markets can be viewed as large distributed systems in which each person, such as a worker on Mechanical Turk, is analogous to a processor that can solve a task requiring human intelligence. In this way a crowdsourcing market could be seen as a loosely coupled distributed computing system [9]. Fleshing out this analogy, we develop here the beginnings of a framework for the future of crowd work that integrates the human aspects of organizational behavior with the automation and scalability of the distributed computing literature.

Both distributed organizations and computing systems face many common fundamental challenges in accomplishing complex work. Key challenges in distributed computing include partitioning computations into tasks that can be done in parallel, mapping tasks to processors, and distributing data to and between processors [9,25,96,132]. Many of these challenges map to coordination dependencies identified by Malone & Crowston [89] that also apply to human organizations. Below we discuss two categories of overlap between coordination dependencies discussed in organizational science, their analogs in distributed computing, and their implications for the beginnings of a framework for the future of crowd work.

Managing shared resources

Whenever a limited resource needs to be shared, coordinating how that resource is allocated becomes important. Allocating a fixed pool of workers to multiple tasks that must be completed under a deadline is a classic example of managing shared resources. Malone & Crowston [89] suggest a number of examples of task allocation mechanisms, ranging from first come/first serve, to markets, to managerial decisions. In distributed computing systems, managing shared resources is of similarly vital importance. Tasks must be mapped to processors, requiring functions to govern task partitioning. Reorganization of this mapping must be possible as well, for example if a processor fails or takes a long time to return results (e.g., MapReduce [34]).

Managing producer/consumer and task/subtask relationships

In many situations, one activity produces something required as input for another activity. For example, the structure of an article needs to be decided on before the

sections can be written. These same requirements exist in distributed computing, in which tasks need to be scheduled so that they can be completed in the correct sequence and in a timely manner, with data being transferred between computing elements appropriately. Deciding how to divide a task into subtasks and managing those subtasks is also a challenging problem, especially for complex and interdependent tasks [61,89]. This is true whether a manager in an organization is trying to plan a large project or a programmer is trying to parallelize a complex task. Furthermore, top-down approaches in which a single person (e.g., the task creator) specifies all subtasks a priori may not be possible, or subtasks may change as the task evolves.

Crowd-Specific Factors

Unlike traditional organizations in which workers possess job security and managers can closely supervise and appropriately reward or sanction workers, or distributed computing systems in which processors are usually highly reliable, crowd work poses unique challenges for both workers and requesters ranging from job satisfaction to direction-setting, coordination, and quality control. For example, organizations can maintain high quality work through management, worker incentives, and sanctions. While some of these methods are available in crowd work (e.g., how much to reward workers, whether to reject their work, or impose a reputation penalty) their power is attenuated due to factors such as lack of direct supervision and visibility into their work behavior, lack of nuanced and individualized rewards, and the difficulty of imposing stringent and lasting sanctions (since workers can leave

with fewer repercussions than in traditional organizations, such as to reference letters or work histories). The worker's power is also limited: requesters do not make a long-term commitment to the worker, and endure few penalties if they renege on their agreement to pay for quality work. In distributed computing systems, by contrast, requesters (programmers) have fewer problems with motivating and directing their workers (computers). However, machines cannot match the complexity, creativity, and flexibility that human intelligence manifests. Combining ideas from human and computer organization theories may thus provide complementary benefits and address complementary weaknesses over using either alone.

Framework

Figure 2 presents a framework that integrates the challenges posed by managing shared resources (such as assigning workers to appropriate tasks), managing producer-consumer relationships (such as decomposing tasks and assembling them into a workflow), and crowd-specific factors (such as motivation, rewards, and quality assurance). Many of its elements combine insights from organizational behavior and distributed computing: for example the task decomposition and task assignment functions use both human and computational processes.

The goal of this framework is to envision a future of crowd work that can support more complex, creative, and highly valued work. At the highest level, a **platform** is needed for managing pools of tasks and workers. Complex tasks must be **decomposed** into smaller subtasks, each **designed** with particular needs and characteristics which must be **assigned** to appropriate groups of workers who themselves must be properly **motivated**, selected (e.g., through **reputation**), and organized (e.g., through **hierarchy**). Tasks may be structured through multi-stage **workflows** in which workers may **collaborate** either synchronously or asynchronously. As part of this, **AI** may guide (and be guided by) crowd workers. Finally, **quality assurance** is needed to ensure each worker's output is of high quality and fits together.

Because we are concerned with issues of design – the technical and organizational mechanisms surrounding crowd work – we highlight in the process model twelve specific research foci (Figure 2) that we suggest are necessary for realizing such a future of crowd work. These foci are grouped into three key dimensions: foci relevant to the work process; the computation guiding, guided by, and underlying the work; and the workers themselves. Our 12 foci overlap each other in places. However, in total they provide a wide-ranging multidisciplinary view that covers current and prospective crowd work processes. For example, workflow techniques may be useful for handling the flow of documents through a set of tasks [111], but the effectiveness of these techniques can be amplified through clever job design that divides tasks and allocates incentives in a way that benefits both workers and requesters (cf. [62]).

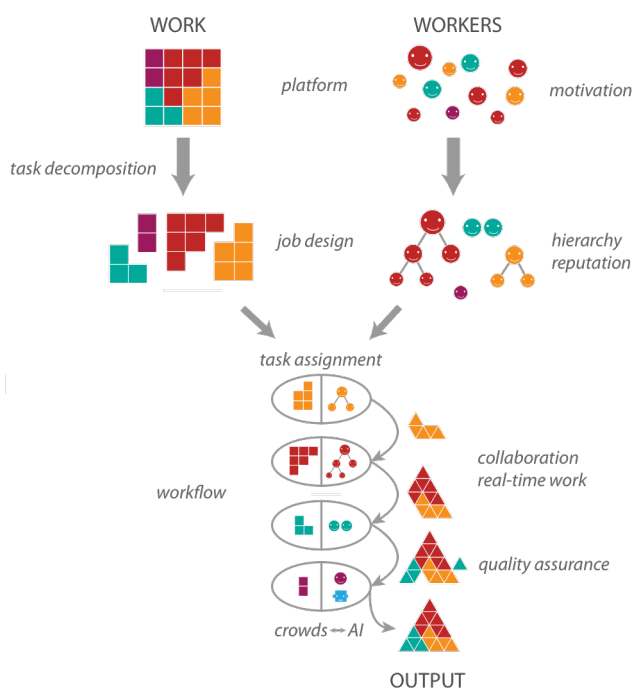


Figure 2: Proposed framework for future crowd work processes to support complex and interdependent work.

Our model is based on empirical as well as theoretical input. In forming this model, we gathered feedback from both task requesters and workers. The authors of this paper have been requesters, have designed crowd workflows, and have worked for platform companies, and so requester and platform company issues are represented. We wished to also represent the voices of workers. We chose one popular crowdsourcing platform, Amazon Mechanical Turk, and asked questions of workers in the two countries with the largest number of crowd workers – the United States and India – who had completed and had approved more than 500 tasks (as enforced by the platform). Workers were paid USD \$2 to comment on and add design suggestions for the 12 foci discussed below. Fifty-two workers responded from the US and 52 from India. Four of the responses were removed from the sample because of incomplete or inconsistent answers. Workers in India had a mean age of 27 ($\sigma=6$). Workers in the US had a mean age of 33 ($\sigma=11$). In India, 27% of workers were female, and in the US, 58% were female. In both countries responders had considerable experience: the mean total lifetime tasks were 6562 ($\sigma=15292$) and 9019 ($\sigma=22460$) for Indian and U.S. workers, respectively. The purpose of the survey was to provide workers a vehicle through which they could contribute their own insights. In general, the workers' responses were thorough and insightful, and we have integrated their ideas into this paper, quoting their answers where appropriate. Even though our survey is informal and relatively small in scale, we believe that it enriches this paper by providing a variety of workers' perspectives on the 12 topics we discuss next.

RESEARCH FOCI

In the sections below, we survey and analyze the 12 research foci that comprise our model. First, we consider the future of the work processes and how the work is organized and accomplished. Second, we consider the integration of crowd work and computation, including the symbiosis between human cognition, artificial intelligence (AI), and computationally-enabled crowd platforms. Finally, we consider crowd workers and how we can develop jobs, reputation systems, motivations, and incentives that will benefit them. In each subsection we state our motivation and goals, briefly review related work, and propose research issues and themes.

The Future of Crowd Work Processes

Increasing the value and meaning of crowd work will require that it move beyond simple deskilled tasks to complex professional work. In this section, we focus on the key challenges that must be met in order to enable complex crowd work processes: designing workflows, assigning tasks, supporting hierarchical structure, enabling real-time crowd work, supporting synchronous collaboration, and controlling quality.

Workflow

Motivation/Goals. Complex crowd work cannot be accomplished using the simple parallel approaches that are

common today, such as aggregating multiple independent judgments through voting or majority rule. Complex tasks have dependencies, changing requirements, and require multiple types of expertise. Instead, *workflows* are needed that facilitate decomposing tasks into subtasks, managing the dependencies between subtasks, and assembling the results. While initial research has shown that enabling more complex workflows can result in large differences in output quality even with small differences in instructions, rewards, and task order [72,127], we have barely begun to understand the broader design space of crowd workflows.

Related Work. Traditional organizations have developed expertise in workflow design and management. The division of labor is a core tenet of task coordination; Adam Smith in his classic *The Nature and Causes of the Wealth of Nations* [133] described the associated efficacy benefits. Via division of labor, a greater pool of agents can work in parallel, specialize for the tasks they perform, and complete an assignment with less time lost to switching tasks [10]. Coordination is difficult among distributed workers, but organizational coordination techniques can be profitably applied to crowd work (e.g., [72,74,90,97,143]). Systems and formal languages support workflows in traditional firms [39], ranging from purely computational [151] to hybrid approaches where tasks are self-selected and then automatically routed onward [138].

In the context of crowds, workflow could involve a much larger scale of operation, and a much more heterogeneous set of actors. The design space ranges from massively redundant, independent tasks (e.g., contests that choose one entry [8,19,20]) to highly serial processes with work passed from one worker to the next (e.g., passing a task from worker to worker for improvements [87]). Recent systems and toolkits pursued a “flare and focus” approach for complex work by exploring a space of options and then drilling down to flesh out those options [15,75,87,150]. Crowd workers can guide workflows as well [1,75,80,154].

Research Proposal. Crowd workflows are still quite brittle, and are most successful with highly targeted tasks. To improve existing workflows, we must experiment and iterate on a large space of parameters, instructions, incentives and decompositions. Costs of doing so may be reduced through models of worker behavior [120] or by encapsulating and reusing proven design patterns [1][73]. Then, we must push crowd workflows toward more general tasks and wicked problems that have no clearly defined solution [112]. Rather than edit text, for example, crowd workflows should be able to support complex goals such as creativity and brainstorming, essay writing, music composition, or civic planning. Crowd workers reminded us in our survey responses that they also need help managing their own workflows as they juggle tasks from different requesters.

Task Assignment

Motivation/Goals. Sharing limited resources requires

coordination — allocating a fixed pool of workers to multiple tasks with deadlines is a classic example [117]. Ideally, requesters will see their tasks completed quickly, while workers are continuously employed with tasks that match their interests. In the worst case, workers are matched to tasks that are uninteresting or too difficult, and won't make the income they want or deserve.

Related Work. Management scientists have developed techniques such as first come/first serve task queues, markets, and managerial decisions [117]. Computer science research adds useful abstractions drawn from data partitioning, OS scheduling, and failover (e.g., [34]). Currently, workers are usually forced to sort queues by task volume and recency [29]. But some algorithms form teams dynamically based on expertise [5].

Research Proposal. Task assignment has typically involved either a first-come/first-served model (e.g., the ESP game, Galaxy Zoo [2,106]) or a market model (e.g., oDesk, Mechanical Turk). In either case, task designers must guess at the right combination of incentives and iterate until success. This process is both time-consuming and expensive. Better theoretical models, markets or automatic computational matching processes (e.g., [14]) could drastically reduce development costs and address search friction [4], an important issue in labor economics. The assignment of tasks in relation to individuals' abilities has also been studied as part of business workflow research [110,128]. Several workers in our survey complained that they spent too much energy finding appropriate tasks; one suggested that platforms provide automatic recommendations of possible next tasks based on workers' previous task choices. These comments suggest a research question: are the workers or the platforms better-suited to manage task assignment? That is, should tasks be pulled by workers or pushed by platforms?

Hierarchy

Motivation/Goal. Hierarchy has become the primary management strategy in traditional organizations. It benefits coordination, decision making, quality control, and assigns incentives and sanctions [30,92]. Hierarchies decompose large and complex tasks such as developing and manufacturing an automobile by clarifying legitimate authority and workflow across organizations. Hierarchies in crowd work could enable groups of workers to tackle new classes of complex work, increase efficiency, and support consistency and integration. Hierarchies may also allow workers to act more like teams, for example developing accountability standards, decision-making and conflict resolution processes, and review policies.

Related Work. Volunteer crowdsourcing platforms have evolved their own hierarchies and decision-making processes [104,156], appropriating techniques from other online communities where appropriate [101]. Most paid approaches have workers make hierarchical decisions collectively: for example, task decomposition and

integration [75,80], quality oversight of each others' contributions [78,100], and leader elections to represent collective opinions [83]. oDesk and MobileWorks identify and empower workers to serve in leadership roles [79]. No comparative analyses of the effectiveness of these approaches exist yet.

Research Proposal. The fluid nature of crowd work opens the door to new kinds of hierarchy where workers transition between roles continuously. Crowd management could entail a layered tree made up of worker-leaders, requesters, machine learning systems, and algorithms. In such a setup, participants might be leaf nodes (workers) in one job but managers in another. Realizing this vision will require improvements to the platforms (e.g., oDesk teams) and the creation of systems that take advantage of it. Alternatively, organized groups of workers might begin applying to jobs as a single entity. But there may be resistance to hierarchy: one worker wrote, "I like the way it is. There does not seem to be a hierarchy. In fact, this is one of the most satisfying aspects of Mturk. Anyone can be their own boss." This comment suggests the value of further empirical study of the willingness of crowd members to manage or be managed by each other. Perhaps current forms of organizational structure will yield to new ones, in which the processes of managing and being managed will be more intertwined and are conceptually different than what we currently experience.

Realtime Crowd Work

Motivation/Goal. For work with tight completion-time constraints, we will need to create *flash crowds*: groups of individuals who arrive moments after a request and can work synchronously. Any application that wants to embed on-demand crowdsourcing (e.g., [15,17]) is limited by the problem of crowd latency, but current crowd tasks can take hours [15] or days [72].

Related Work. Fast recruitment has been the major research thrust in realtime crowdsourcing so far. Early attempts were motivated by time-limited tasks such as searching for a missing person [59] and timed "competitions" [99,140]. In paid crowdsourcing, researchers began e-mailing a set of workers the night before the study and announced a time for the experiment. Keeping workers busy with old tasks brought wait times down to a half minute to one minute [17], and paying workers a small wage to stay on call is enough to draw a crowd together roughly two to three seconds later [13,14]. This technique can be modeled using queuing theory and adapted to bring together crowds in as little as 500 milliseconds.

Research Proposal. The two core challenges for realtime crowdsourcing will be 1) scaling up to increased demand for realtime workers, and 2) making workers efficient enough to collectively generate results ahead of time deadlines. What happens as more tasks need such responses, and as the size of the crowd increases? Is it

possible to support a large number of realtime crowdsourcing tasks competing for workers' attention [14]?

Workers who arrive quickly can still be slow at completing work [13]. It is possible to design algorithms and workflows to guide workers to complete synchronous tasks quickly. So far, these techniques are restricted to particular domains [13], but general approaches may be possible. What would it take to begin with a sketch on a napkin, find a designer to mock up interface alternatives, a usability analyst to test those prototypes, and a front-end engineer to implement the best one, all in a single afternoon? Workers seem interested in such tasks: one suggested providing "More communications options with the requester, something better than just emailing them, like some sort of immediate chat." Another wrote: "I'd like to see some sort of worker alert system on the dashboard for such events."

Synchronous Collaboration

Motivation/Goal. Many tasks worth completing require cooperation – yet crowdsourcing has largely focused on independent work. Distributed teams have always faced challenges in cultural differences and coordination [60], but crowd collaboration now must create rapport over much shorter timescales (e.g., one hour) and possibly wider cultural or socioeconomic gaps.

Related Work. Many attempts at collaborative crowd work highly structure the communication between participants. For example, crowds can solve distributed problems by observing the behavior of neighbors [91], letting the system choose which suggestions to focus on [13], continuously elect new leaders [83] and pass on knowledge to new members [84]. These techniques reduce the damage that a single poor or malicious crowd worker can do, while also limiting the types of collaboration possible. They also set up opportunities for feedback and learning [42].

Unstructured collaboration also shows promise, for example by giving workers a task and placing them in collaborative text authoring software [77]. These techniques draw on synchronous collaboration research (e.g., [53,66]).

Research Proposal. To shift from independent workers to teams of on-demand collaborators, we must revisit and extend traditional CSCW work on distributed teamwork. Short periods of intense crowd collaboration call for fast teambuilding and may require the automatic assignment of group members to maximize collective intelligence [5,149]. Finally, it will be a major research undertaking to invent and describe the tasks and techniques that succeed with synchronous collaboration.

Quality Control

Motivation/Goal. Quality problems are a serious challenge to the mass adoption of crowd work. The most appealing aspects of crowd work — such as high throughput, low transaction costs, and complex/subjective tasks — also make it susceptible to quality control issues. Workers *satisfice*, minimizing the amount of effort they expend, and

in the extreme cheat or game the system. For example, 30% or more of submissions on Mechanical Turk may be low quality [15,72]. One worker warned us: "People collude to agree on wrong answers to trick the system." That is, quality filters based on consensus may be fooled by workers who agree to coordinate answers. Workers who have low expertise and requesters who provide unclear instructions also contribute to subpar responses. Problems arise even for workers who are highly motivated: these "eager beavers" often make well-intentioned but counterproductive contributions [15]. In our survey, workers saw quality control as a major issue that affected their compensation, and they expressed a dislike for their peers who lowered quality standards through misbehavior. However, many complained about requesters: one said: "Too often the job itself is badly designed or is messed up and there is a degree of misunderstanding between the worker and the job engineer."

Related Work. Of the research foci, quality control has arguably received the most attention so far. Approaches for quality control largely fall into two camps: up-front task design and post-hoc result analysis. Task design aims to design tasks that are resistant to low-quality work. For example, requesters can split work into fault-tolerant subtasks [15,75,103], apply peer-review or agreement filters [2,15,17,42,63,75], optimize instructions [43,72,127], and manipulate incentives [26,115,127].

Worker output approaches filter out poor quality work after results have been submitted. Workers' results can be compared to gold standard data on a pre-labeled set of the tasks [24,43,85]. Including gold standards can prevent workers' inherent biases from dominating the results [37]. However, authoring gold data can be burdensome, and gold standards may not be possible for subjective or generative tasks (e.g., writing an essay). Other common methods scale the influence of a submission according to how well that worker agrees with others [24,36,64,134] or according to the workers' votes. However, recruiting multiple workers costs more, agreement may not be possible for subjective or generative tasks, and the approach is susceptible to collusion [41]. False identities are increasingly being created as an attack on quality assurance methods.

A promising approach that addresses some worker output issues examines the way that workers do their work rather than the output itself, using machine learning and/or visualization to predict the quality of a worker's output from their behavior [119,120]. Similar but simpler approaches provide requesters more visibility into worker behavior, such as oDesk's Worker Diary which periodically takes snapshots of workers' computer screens. While powerful, such techniques must address privacy and autonomy concerns if widely deployed.

Research Proposal. While quality control is improving for tasks with a closed set of possible answers, we still have few techniques for open-ended work and highly skilled

tasks. Is it possible to robustly gauge workers' skills at tasks such as audio engineering, art critique, or poetry? Should we rely on peer evaluation, or data mine low-level activity to predict output quality? Do quality metrics map well across different marketplaces?

In the long term, we must move from reducing poor quality work to scaffolding truly excellent work. To do so, we can optimize workflows for creativity, innovation, and discovery. This vision will involve recruiting experts; metrics that enable us to evaluate such factors; and reward structures that value high-quality artifacts.

The Future of Crowd Computation

Crowd labor is already mediated by computation. However, computation has the opportunity to step into a much more active role in helping recruit and manage workers as well as contribute directly to work processes. Hybrid human-computer systems could tap into the best of both human and machine intelligence. We structure our discussion around the potential for reciprocal benefits between crowd workers and computational systems. *Crowds guiding AIs* considers how crowd intelligence can help train, supervise, and supplement automation. *AIs guiding crowds* considers how machine intelligence can help make the crowd more efficient, skilled, and accurate. Finally, we also consider the design and evaluation of crowdsourcing platforms.

Crowds Guiding AIs

Motivation/Goal. In human computation, people act as computational components and perform the work that AI systems lack the skills to complete [2]. By tapping into crowd intelligence, computational systems can support a much broader set of tasks. It should be understood that this area, out of all areas, least excites crowd workers, and perhaps with good reason: crowd workers may end up training machines to replace them.

Related Work. Paid crowds have gathered large amounts of data to train algorithms. For example, crowds can: 1) match expert annotations on natural language processing tasks such as word sense disambiguation [23,134], 2) generate speech corpora for spoken language research [23,95], 3) annotate objects and people in images [135], and 4) help with graphics tasks such as identifying depth layers [51]. Crowds can also solve algorithmic problems such as graph coloring [68,91].

Research Proposal. While algorithms will continue to benefit from crowd-generated training data, there are opportunities to integrate crowds more deeply into algorithms. Rather than treating crowd data as ground truth labels, it may be profitable to understand and model the biases and intuitions that human cognition brings [147]. It may be possible to design machine learning algorithms that more deeply understand the human nature of these labels. Algorithms may also more directly model the tradeoff between cost and performance: for example, by using a

combination of active learning and semi-supervised learning to collect the most informational labels [155].

AIs Guiding Crowds

Motivation/Goal. While large groups are increasingly adept at completing straightforward parallel tasks [54,139], they can struggle with complex work. Participants have varied skill levels, even well-intentioned contributions can introduce errors, and errors are amplified as they propagate through the crowd. It is possible to integrate crowds directly inside of software [15,17,46], and use the software to help guide crowd work; for example, a machine learning model can determine which work products may still be improved, and then assign workers most likely to make such improvements [32,33]. These systems may also be able to predict their expertise needs in advance, then train and adapt workers in an online fashion via automated tutoring [47,113] or peer learning [18].

Related Work. Computational approaches for designing and integrating workflow, incentive, and instruction patterns have shown promise [32,125], as well as techniques that trade off the strengths of crowds and artificial intelligence [67,146]. AIs could also serve as a reflective aids, encouraging the crowd to learn by pointing out what others have done in similar contexts (e.g., [109]).

Research Proposal. The research community should examine whether algorithmic management is an improvement over traditional organization management techniques. In an algorithmic organization crowds will need to be able to raise exceptions, as well as halt and restart processes. And AIs will need to know when then can proceed, when they need human help, and when they need to help the workers. In addition, workers should be able to modify their supporting AIs as needed. In the end, both workers and AIs should improve. To accomplish this, we will need to move from a setting where simple AIs completely determine a workflow to a richer, mixed-initiative setting where crowds and AIs jointly teach each other, and jointly control the work process.

Crowdsourcing Platforms

Motivation/Goal. Crowdsourcing platforms provide the central nexus between requesters and workers. As a result, platform design offers the opportunity to change our perceptions and understanding of crowd work in general, as well as shape the relationships and practical interactions between workers and requesters in practice. While several platforms already exist, limiting ourselves to existing platforms greatly restricts the scope and nature of change we can enact. By contrast, novel platforms can help drive the diffusion of new designs and techniques for crowd work.

Related Work. Platform research tends to optimize existing processes or reach out to new populations. For example, CrowdFlower experiments on its own gold standard metrics to minimize the number of times a question is asked [129].

Platforms may also manage a large number of real time requests or route tasks to ensure that all tasks have steady streams of incoming workers [14]. MobileWorks promotes its own workers to management positions based on their performance [79]. To find new populations, MobileWorks and mClerk engage with the developing world through specifically-tailored mobile phone interfaces [56,100]. Incentives can be used to recruit local experts as well: for example, a vending machine served up exam grading tasks and dispersed candy as a reward [58].

Research Proposal. While designing and building a new platform may seem daunting, examples like Brin and Page's Google [21] show that two graduate students can disrupt the commercial landscape and dramatically change the way people work. We challenge the community to similarly revolutionize our conception of what a crowdsourcing platform is and can achieve. Innovation should articulate a vision for crowd work that is effective, efficient, and fair. Beyond mere technology, negotiating the balance of power between interested parties is central to platforms and markets [11,130]. As crowd work already faces challenges related to power inequalities similar to those encountered in offline labor markets, future platforms may shape or be constrained by future regulatory intervention, on, for example, the use of independent contracting for regular, recurring work [45]. Beyond money, platforms might also support labor exchanged for virtual goods or performed in virtual environments, which raises additional policy issues [44]. Privacy is another issue: how can platforms disclose enough information to be trusted as a source of worker quality while also maintaining privacy? Experience with markets such as eBay and Amazon suggests that greater transparency may be helpful, but such mechanisms must be carefully managed to avoid abuse [38].

Security concerns will also continue to grow in importance. For example, recognition of trusted workers may lead to identity theft and fraudulent use of compromised accounts [41]. Security research must consider both new attacks on platforms and use of platforms for launching attacks [144].

The Future of Crowd Workers

Crowd work involves a partnership between requesters and workers. Thus, when designing the future of crowd work, it is important to develop tools to support not only the work itself but also those performing the work. Below we identify and discuss three important research challenges for supporting the crowd workers of the future: job design, reputation and credentials, and motivation and rewards.

Job Design

Motivation/Goals. Some tasks that need to be done are just dull. Motivating workers to accomplish such tasks can be challenging, and may lead to reduced engagement with the system: "It would be better if some of the task assignments weren't so monotonous...I don't see the long-term payoff and it discourages me." While dressing up such tasks as

games may reduce boredom, entertainment represents a fairly superficial form of work satisfaction. We believe the future of crowd work depends on creating jobs that achieve both organizational performance and worker satisfaction.

Related Work. In traditional firms, when managers design jobs that provide skill variety, task identity, and task significance, workers find their jobs more meaningful [57]. When workers have autonomy in their jobs and receive feedback, they experience increased responsibility for and understanding of the impact of their work. Combined, these lead to increased performance for the firm and reduced employee problems such as absenteeism and turnover [57].

Unfortunately, many paid crowd work platforms do not provide as much skill variety, task identity, and task significance as volunteer platforms such as Wikipedia. There are direct payoffs when requesters convey the identity and significance of tasks to crowd workers [26,115]. Timely and task-specific feedback from peers and requesters, as well as opportunities for self-assessment, help workers to learn, persevere, and produce better results [42].

Research Proposal. An ideal crowd work system would allow workers to complete a whole and identifiable piece of work in a way that is satisfying and measurable (cf. [62]). The system would explain the significance of the job, offer peer-to-peer and expert feedback, and encourage self-assessment. These systems could offer a variety of ways to complete the task, and thereby not only provide autonomy for the worker, but also reduce errors [55]. Achieving this vision will require communication with workers about the scope and impact of their work, as well as verification that workers understand their impact.

Providing more context has tradeoffs, however. On the one hand, more context enables workers to better judge how the fruits of their effort will be used so they can make informed decisions about whether or not to perform work. On the other hand, reducing context may streamline work, leading to greater efficiencies for both requesters and workers. Moreover, requesters may not always want to fully disclose the context of work due to privacy, security, or intellectual property concerns. This suggests a need to balance distinct and competing concerns: how much information a worker needs before consenting to provide labor, how much information should be shared to motivate and retain workers, how to share context without introducing inefficiency and how much information may be legitimately withheld to protect interests of requesters.

Reputation and Credentials

Motivation/Goals. In traditional firms, reputation and credentials (e.g., letters of reference, certifications, work history) are critical to recruiting, creating lasting rewards and sanctions, and managing work quality. For example, institutions with brand names such as Google and Apple are likely to attract software engineers to apply for jobs, and, in turn, such names on the resumes of software engineers

attract other prospective employers. However, many existing crowd work systems have only coarse reputation and credentialing systems. For example, a worker's history on Mechanical Turk primarily measures the percentage of work that has been approved. In contrast, traditional employers judge a prospective employee's education and work history through a variety of instruments, including interviews, transcripts, and references. Likewise, applicants can investigate employers' reputation (and crowd workers in our survey wanted better reputation rankings for employers to be established within the platform [130]).

Related Work. Robust mechanisms for supporting trust, assurance, and accountability can support volunteer-based crowds and collaboration online [28,76]. Likewise, in the context of monetary rewards, reputation can have serious financial consequences [111], motivating people to manipulate systems for their own benefit. For example, workers may build networks of Sybil identities — that is, create pseudonyms [41] — to enhance their reputations and to foil quality control methods. Workers may also boost their ratings by agreeing implicitly or explicitly to recommend each other reciprocally [38]. The design and evaluation of future crowd work reputation systems must address these issues.

Research Proposal. The core challenge with reputation is trading off the benefits of pseudonymous, low transaction cost hiring with the richer but higher transaction cost hiring decisions most firms make today. For example, employers know little about workers on Mechanical Turk, but such workers can be hired and engaged in work nearly instantaneously. Meanwhile, platforms such as oDesk with richer reputation systems incur correspondingly higher transaction costs in hiring workers, involving negotiation and handshakes between workers and employers.

To address this challenge reputation systems will need to be robust to cheating and gaming while preserving the benefits of pseudonymity and supporting low-transaction cost hiring. One possible solution would be to create a web of trust in which requesters and workers validate each other as trustworthy [79]. However, malicious workers and requesters can infiltrate the community and spread their own web of trust. Interfaces may support detection of bad actors: for example, by highlighting topology [121,144], statistical patterns [93] and behavior [120].

The creation of technical tools for sharing information about workers should be coupled with more robust systems for monitoring and reporting requester abuses [131]. Lastly, these initiatives towards enhanced reputation systems should be balanced with the need to preserve privacy, as well as the potential payoffs of anonymous collaboration for both workers and requesters [16].

Motivation and Rewards

Motivation/Goals. Requesters often envision crowd workers as either anonymous individuals who are motivated to

complete piecework by small monetary rewards or highly-skilled professionals who work on large, higher-paid tasks without management oversight. However, crowd workers are a diverse and multifaceted population with a range of motives and experiences. Many workers in micro-task environments are ambitious individuals, desiring to be CEOs or top-rate school teachers [48]. Yet few researchers have grappled with the diversity and richness of the motives of the individuals comprising the crowd. One worker reminded us that requesters, as well as workers, need to be motivated:

We could definitely use more motivation, we perform task[s] for mere pennies. Mturk should encourage and reward requesters that provide clear instructions, quick payment, and higher pay. Rewarding them would create more worthwhile tasks that we would take more seriously and work hard on. Good, credible HITs are few and far between.

Related Work. Research in psychology, sociology, management, and marketing provide insights into human motivation that are applicable to crowd work. Management research illustrates the challenge of clearly understanding, communicating, and rewarding desired behavior [70]. Workers seek to understand which activities are rewarded and then tend to do those activities to the virtual exclusion of others. Other studies find mixed results on the effect of financial incentives on the quality of workers' outputs, and underscore the performance and satisfaction benefits of harnessing intrinsic motivations in task design such as non-financial awards and recognition, meaningfulness of tasks, and the feeling of contributing towards the greater good [26,72,86,94,115,127].

Past research suggests requesters should (1) clearly understand and communicate desired behaviors, (2) understand and align worker motivations and incentives with these desired behaviors, and (3) design the requests and incentive structures in order to achieve both effective task completion and worker satisfaction. This also requires requesters to understand variations in crowd workers' motivations (e.g., [7,11,71,116]), for example in competence, enjoyment, connectedness, prosocial orientation, and autonomy [49].

Research Proposal. The future of crowd work requires that requesters and platform designers consider a broad set of motivations, not just financial incentives. We must create frameworks that acknowledge the dynamic nature of motivation and its dependence on context. For example, it is not clear that payment alone will be the optimal motivator for expert crowdsourcing markets. Such frameworks should enable us to move from analysis to design of new motivational schemes. Research should overcome the dichotomous emphases on dehumanizing piece-work and frictionless virtual collaboration in order to provide a more holistic framework within which to

understand and build systems that support workers' diverse motivations (e.g., [22,50,52,81,82,107,124]).

NEXT STEPS

Many of the really difficult problems that workers and requesters face will require advances on multiple foci at the same time. Below we describe three design goals demonstrating how the integration of multiple foci can lead to concrete next steps and calls to actions that will create a better environment for crowd workers, and a better set of human resources for requesters.

Create Career Ladders - [*motivation, job design, reputation, hierarchy*]

Crowd work today is largely a dead-end job, offering few opportunities for career advancement and economic mobility. As workers demonstrate aptitude and diligence, it is to the advantage of both requesters and workers to recognize workers' potential to take on new tasks requiring greater knowledge, effectiveness, and responsibility. As all organizations benefit from making the best use of available talent and workers' diverse abilities, more skilled workers should be paid for their expertise, and encouraged to train less skilled workers: for example, MobileWorks promotes workers to management jobs based on performance [79].

As workers demonstrate proficiency, they might be invited to create gold labels for verifying the quality of work from other, less established workers, which could be used to build webs of trusted workers. Proficient, trusted workers could also manage other workers, respond to reported issues, and provide first-pass triage of new task designs (to catch problems or make suggestions before new tasks go live). Crowd workers could eventually become employees themselves, or develop the skills needed to launch their own business using crowd work. For example, a career trajectory might proceed as: 1) entry-level, untrusted worker; 2) trusted worker; 3) hourly contractor; 4) employee. First steps toward building such a ladder include studying worker *motivations* in order to develop better *job designs*; creating lasting and transferable mechanisms for *reputation and credentialing* for workers; and building greater support for *hierarchy* in the form of structured teams that provide training for novices by skilled workers.

Improve Task Design through Better Communication - [*quality assurance, job design, task assignment, realtime crowd work, synchronous collaboration, platform*]

There is a popular myth that the poor quality of some crowd work stems largely from workers being lazy, stupid, or deceitful. In practice, both we and our surveyed workers have observed many cases where poor quality work instead arises from poorly designed crowdsourcing tasks. For example, a requester might assume a task obvious to them should be equally obvious to everyone else. However, even highly-educated workers may have difficulty understanding exactly what the requester actually wants. Task instructions are often incomplete or ambiguous, do not address boundary cases, and do not provide clarifying input-output

examples of what is expected. Task interfaces may be poorly designed or even have bugs that make it impossible to complete a task. Gold labels used as trap questions may be far more subjective than requesters realize, leading to mistaken rejection of work.

What might we do to address these problems? Designers can make it easier and faster for requesters to create effective tasks. For example, platforms might provide task templates showing examples of proven task designs [27] to encourage shared mental models and improve *quality assurance*. Platforms could also help educate requesters about the impact of *job design* and *task assignment* on resulting quality, best practices, and common errors to be avoided. Platforms might even offer a "first pass" service in which a set of trusted workers (known to the platform) test out the task and report any issues encountered.

Platforms might also provide a wider array of communication channels between requesters and workers supporting *synchronous collaboration* and *real-time crowd work*. Workers we surveyed were adamant that the perception of poor crowd work quality was due, at least in part, to unclear instructions and insufficient feedback, and that they need more guidance to better understand what is expected. They suggested instant chat with requesters to clarify jobs, though we note this would require the continual presence of requesters while work is being performed. They also requested feedback during or just after a task. Both would be consistent with the practices of good managers and workers in other labor contexts, and so more experimenting with channels of communication could potentially have a large effect on both worker satisfaction and job quality. As a first step, requesters might provide ways for workers to clarify tasks in real time – with the requester, or with a more experienced worker. And, when work is in progress, they might provide informal feedback through these same channels.

Facilitate Learning - [*quality assurance, AIs guiding crowds, crowds guiding AIs, task assignment, reputation and credentials, platform*]

Crowd work naturally involves learning and assessment. Workers may need to acquire new skills to perform unfamiliar tasks, before or in the midst of performing the actual work. Workers may also polish and refine existing skills while completing more familiar tasks. Requesters must continually engage in *quality assurance*. Such a training-assessment cycle of work offers potentially exciting synergies with online, education by-doing. For example, DuoLingo (duolingo.com) explores this direction for foreign language learning. This idea can be generalized much further; for example, content generation tasks could be designed to better assess and enhance writing skills. A self-sustaining cycle might involve *AIs guiding crowds* on which tasks to complete (*task assignment*) depending on the worker's and requester's skill development and quality assurance goals, and then using the crowd-generated data to automate some of the simpler tasks (*crowds guiding AIs*).

The potential of crowd work-based education is enormous and multi-faceted, benefiting all parties by producing more skilled and employable workers. Online tutoring systems, perhaps augmented with human tutoring, could provide a path toward delivering more scalable education to the public at large [6,146]. Moreover, tracking and mining of work history could support personalized instruction and feedback, as well as recommending new tasks and learning modules. As workers master new skills and are assessed, badges or *credentials* could document this proficiency so that others can recognize and utilize this enhanced skill set. *Platforms* themselves can also be an important element of learning: data from crowd work can reveal what kinds of requests attract talented workers, patterns of learning and skill building among workers over time, the valuation and interaction between extrinsic and intrinsic rewards, and which types of tasks are most appropriate for which types of workers.

CONCLUSION

Crowd work may take place in the scale of minutes, but the impact of crowd work may be felt for generations. We have asked: what will it take for us, the stakeholders in crowd work – including requesters, workers, researchers – to feel proud of our own children when they enter such a work force? Answering this question has led to a discussion of crowd work from a longer-term perspective.

Specifically, we have synthesized a framework for the future of crowd work based on theory from organizational behavior and distributed computing, and informed by the concerns of crowd work stakeholders. Our hope is that this framework and the corresponding research challenges will spur discussion, experiment, and further insight. In sum, we envision a future where many cognitively complex, large-scale tasks can be decomposed into workflows and executed by crowds consisting of novices, experts and algorithms, and that the crowd work environment can be designed in such a way that satisfies the needs of both workers and requesters. After identifying the twelve major research foci that constitute the framework, we probed their current state, identified gaps, and created a call for action demonstrating ways they can be profitably integrated.

Calls for action must be made responsibly. Crowd work marketplaces are complex socio-technical systems, composed of many people and a changing technical infrastructure, with emergent organizational forms, new incentives being offered, and shifting labor pools. As a result of this complexity, there is the possibility of unpredictable side effects. For example, innovations that effectively train workers through micro-tasks may have ramifications for the world's educational institutions, and thereby for society as a whole. Hybrid combinations of workers and artificial intelligence that seek to build collective intelligence may instead lead to mechanized workers or human-imitating machines.

Our call for action, then, calls for both exciting innovation

and also close observation of its effects. In crowd work, we have two important affordances: an ability to constitute new forms of organization in short amounts of time, and an ability to situate these organizations in an experimental context. While organization science has been built slowly based on observation, the proliferation of crowd work makes large-scale organizational experiments comparing distinct management strategies and task designs possible. These comparisons will help us understand how to improve crowd platforms, workers' skills, and requesters' assignments. Perhaps instead of a hadron collider, the field of crowd work needs a "social collider" in which different forms of organization can be tested. The goal should be better systems, better requests, better work, and better experience. We hope the community's observational, experimental, design and technical skills will play a vital role in shaping the future of crowd work and the next generation of workers.

ACKNOWLEDGEMENTS

We thank all the participants in the CHI 2012 CrowdCamp Workshop for their suggestions, and in particular Paul Andre who spearheaded the effort. We are pleased to acknowledge the crowd workers of the present: for their opinions and creative suggestions in response to our queries, and, more generally, for their pioneering spirit. Finally, we acknowledge Ashima, the nascent inspiration for this paper.

This work was supported under National Science Foundation awards IIS-0968561, IIS-0855995, OCI-0943148, IIS-0968484, IIS-1111124, IIS-1149797, IIS-1217096, and IIS-1217559. This work was also supported by a DARPA Young Faculty Award N66001-12-1-4256, a Temple Fellowship, Northwestern University, and the Center for the Future of Work, Heinz College, Carnegie Mellon University.

Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the funding agencies.

REFERENCES

1. Ahmad, S., Battle, A., Malkani, Z., and Kamvar, S. The jabberwocky programming environment for structured social computing. *Proc. UIST '11*, (2011).
2. Von Ahn, L. and Dabbish, L. Labeling images with a computer game. *Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM (2004), 319–326.
3. Ahn, L.V., Blum, M., Hopper, N.J., and Langford, J. CAPTCHA: Using hard AI problems for security. *Proceedings of the 22nd international conference on Theory and applications of cryptographic techniques*, (2003), 294–311.
4. Albrecht, J. The 2011 Nobel Memorial Prize in Search Theory. *Department of Economics*,

- Georgetown University. [http://www9.georgetown.edu/faculty/albrecht/SJE_Survey.pdf], (2011).
5. Anagnostopoulos, A., Becchetti, L., Castillo, C., Gionis, A., and Leonardi, S. Online team formation in social networks. *Proceedings of the 21st international conference on World Wide Web*, ACM (2012), 839–848.
 6. Anderson, M. Crowdsourcing Higher Education: A Design Proposal for Distributed Learning. *MERLOT Journal of Online Learning and Teaching* 7, 4 (2011), 576–590.
 7. Antin, J. and Shaw, A. Social desirability bias and self-reports of motivation: a study of amazon mechanical turk in the US and India. *Proc. CHI '12*, (2012).
 8. Archak, N. and Sundararajan, A. Optimal Design of Crowdsourcing Contests. *ICIS 2009 Proceedings*, (2009), 200.
 9. Bal, H.E., Steiner, J.G., and Tanenbaum, A.S. Programming languages for distributed computing systems. *ACM Computing Surveys (CSUR)* 21, 3 (1989), 261–322.
 10. Becker, G.S. and Murphy, K.M. The division of labor, coordination costs, and knowledge. *The Quarterly Journal of Economics* 107, 4 (1992), 1137–1160.
 11. Bederson, B.B. and Quinn, A.J. Web workers unite! addressing challenges of online laborers. *Extended Abstracts CHI '11*, (2011).
 12. Benkler, Y. *The wealth of networks: How social production transforms markets and freedom*. Yale Univ Pr, 2006.
 13. Bernstein, M.S., Brandt, J., Miller, R.C., and Karger, D.R. Crowds in two seconds: Enabling realtime crowd-powered interfaces. *Proc. UIST '11*, (2011).
 14. Bernstein, M.S., Karger, D.R., Miller, R.C., and Brandt, J. Analytic Methods for Optimizing Realtime Crowdsourcing. *Proc. Collective Intelligence 2012*, (2012).
 15. Bernstein, M.S., Little, G., Miller, R.C., et al. Soylent: A Word Processor with a Crowd Inside. *Proc. UIST '10*, (2010).
 16. Bernstein, M.S., Monroy-Hernández, A., Harry, D., André, P., Panovich, K., and Vargas, G. 4chan and /b/: An Analysis of Anonymity and Ephemerality in a Large Online Community. *Fifth International AAI Conference on Weblogs and Social Media*, AAAI Publications (2011).
 17. Bigham, J.P., Jayant, C., Ji, H., et al. VizWiz: Nearly Real-time Answers to Visual Questions. *Proc. UIST '10*, (2010).
 18. Boud, D., Cohen, R., and Sampson, J. Peer learning and assessment. *Assessment & Evaluation in Higher Education* 24, 4 (1999), 413–426.
 19. Boudreau, K.J., Lacetera, N., and Lakhani, K.R. Incentives and problem uncertainty in innovation contests: An empirical analysis. *Management Science* 57, 5 (2011), 843.
 20. Brabham, D.C. Moving the crowd at iStockphoto: The composition of the crowd and motivations for participation in a crowdsourcing application. *First Monday* 13, 6 (2008), 1–22.
 21. Brin, S. and Page, L. The anatomy of a large-scale hypertextual Web search engine. *Computer networks and ISDN systems* 30, 1-7 (1998), 107–117.
 22. Bryant, S.L., Forte, A., and Bruckman, A. Becoming Wikipedian: transformation of participation in a collaborative online encyclopedia. *GROUP '05*, ACM Press (2005), 1–10.
 23. Callison-Burch, C. and Dredze, M. Creating speech and language data with Amazon's Mechanical Turk. *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*, Association for Computational Linguistics (2010), 1–12.
 24. Callison-Burch, C. Fast, cheap, and creative: evaluating translation quality using Amazon's Mechanical Turk. *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1-Volume 1*, (2009), 286–295.
 25. Casavant, T.L., Braun, T.A., Kaliannan, S., Scheetz, T.E., Munn, K.J., and Birkett, C.L. A parallel/distributed architecture for hierarchically heterogeneous web-based cooperative applications. *Future Generation Computer Systems* 17, 6 (2001), 783–793.
 26. Chandler, D. and Kapelner, A. Breaking monotony with meaning: Motivation in crowdsourcing markets. *University of Chicago mimeo*, (2010).
 27. Chen, J.J., Menezes, N.J., and Bradley, A.D. Opportunities for Crowdsourcing Research on Amazon Mechanical Turk. *Interfaces* 5, (2011), 3.
 28. Cheshire, C. Online Trust, Trustworthiness, or Assurance? *Daedalus* 140, 4 (2011), 49–58.
 29. Chilton, L., Horton, J., Miller, R.C., and Azenkot, S. Task search in a human computation market. *Proc. HCOMP '10*, (2010).
 30. Coase, R.H. The Nature of the Firm. *Economica* 4, 16 (1937), 386–405.
 31. Cooper, S., Khatib, F., Treuille, A., et al. Predicting protein structures with a multiplayer online game. *Nature* 466, 7307 (2010), 756–760.
 32. Dai, P., Mausam, and Weld, D. Decision-theoretic control of crowd-sourced workflows. *Proc. AAAI '10*, (2010).
 33. Dai, P., Mausam, and Weld, D.S. Artificial intelligence for artificial intelligence. *Twenty-Fifth AAAI Conference on Artificial Intelligence*, (2011).
 34. Dean, J. and Ghemawat, S. MapReduce: Simplified Data Processing on Large Clusters. *To appear in OSDI*, (2004), 1.
 35. Debeauvais, T., Nardi, B.A., Lopes, C.V., Yee, N.,

- and Ducheneaut, N. 10,000 Gold for 20 Dollars: An exploratory study of World of Warcraft gold buyers. *Proceedings of the International Conference on the Foundations of Digital Games*, (2012), 105–112.
36. Dekel, O. and Shamir, O. Vox populi: Collecting high-quality labels from a crowd. *Proc. 22nd Annual Conference on Learning Theory*, (2009).
 37. Della Penn, N. and Reid, M.D. Crowd & Prejudice: An Impossibility Theorem for Crowd Labelling without a Gold Standard. *Collective Intelligence*, (2012).
 38. Dellarocas, C. Analyzing the economic efficiency of eBay-like online reputation reporting mechanisms. *Proceedings of the 3rd ACM Conference on Electronic Commerce*, (2001), 171–179.
 39. Van Der Aalst, W.M.P., Ter Hofstede, A.H.M., Kiepuszewski, B., and Barros, A.P. Workflow patterns. *Distributed and parallel databases 14*, 1 (2003), 5–51.
 40. Doctorow, C. *For the Win*. Voyager, 2010.
 41. Douceur, J. The sybil attack. *Peer-to-peer Systems*, (2002), 251–260.
 42. Dow, S., Kulkarni, A., Klemmer, S., and Hartmann, B. Shepherding the crowd yields better work. *Proc. CSCW '12*, (2012).
 43. Downs, J.S., Holbrook, M.B., Sheng, S., and Cranor, L.F. Are your participants gaming the system?: screening mechanical turk workers. *Proceedings of the 28th international conference on Human factors in computing systems*, (2010), 2399–2402.
 44. Felstiner, A. Sweatshop or Paper Route?: Child Labor Laws and In-Game Work. *Proceedings of CrowdConf*, (2010).
 45. Felstiner, A. Working the Crowd: Employment and Labor Law in the Crowdsourcing Industry. *Berkeley J. Emp. & Lab. L.* 32, (2011), 143–143.
 46. Franklin, M., Kossmann, D., Kraska, T., Ramesh, S., and Xin, R. CrowdDB: answering queries with crowdsourcing. *Proc. SIGMOD '11*, (2011).
 47. Gal, Y., Yamangil, E., Shieber, S., Rubin, A., and Grosz, B. Towards collaborative intelligent tutors: Automated recognition of users' strategies. *Intelligent Tutoring Systems*, (2008), 162–172.
 48. Gerber, E. and Dontcheva, M. *Career Aspirations for Crowdworkers*. In preparation.
 49. Gerber, E., Hui, J., and Kuo, P. *Crowdfunding: Why creators and supporters participate*. Segal Design Institute, Evanston, IL, 2012.
 50. Ghosh, R. and Glott, R. *Free/Libre and Open Source Software: Survey and Study*. European Commission, 2002.
 51. Gingold, Y., Shamir, A., and Cohen-Or, D. Micro Perceptual Human Computation. *To appear in ACM Transactions on Graphics (TOG)*, (2012).
 52. Glott, R., Ghosh, R., and Schmidt, P. *Wikipedia Survey*. UNU-MERIT, Maastricht, Netherlands, 2010.
 53. Greenberg, S. and Bohnet, R. GroupSketch: A multi-user sketchpad for geographically-distributed small groups. *Proc. Graphics Interface '91*, (1991).
 54. Grier, D.A. *When Computers Were Human*. Princeton University Press, 2005.
 55. Grier, D.A. Error Identification and Correction in Human Computation: Lessons from the WPA. *Proc. HCOMP '11*, (2011).
 56. Gupta, A., Thies, W., Cutrell, E., and Balakrishnan, R. mClerk: Enabling Mobile Crowdsourcing in Developing Regions. *Proc. CHI '12*, (2012).
 57. Hackman, J.R. and Oldham, G.R. Motivation through the design of work: Test of a theory. *Organizational behavior and human performance 16*, 2 (1976), 250–279.
 58. Heimerl, K., Gawalt, B., Chen, K., Parikh, T.S., and Hartmann, B. Communitysourcing: Engaging Local Crowds to Perform Expert Work Via Physical Kiosks. *Proc. CHI '12*, (2012).
 59. Hellerstein, J.M. and Tennenhouse, D.L. Searching for Jim Gray: a technical overview. *Communications of the ACM 54*, 7 (2011), 77–87.
 60. Hinds, P. *Distributed work*. The MIT Press, 2002.
 61. Von Hippel, E. Task partitioning: An innovation process variable. *Research policy 19*, 5 (1990), 407–418.
 62. Holmstrom, B. and Milgrom, P. Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *JL Econ. & Org.* 7, (1991), 24.
 63. Horton, J.J. and Chilton, L.B. The labor economics of paid crowdsourcing. *Proceedings of the 11th ACM conference on Electronic commerce*, (2010), 209–218.
 64. Ipeirotis, P.G., Provost, F., and Wang, J. Quality management on amazon mechanical turk. *Proc. HCOMP '10*, (2010).
 65. Ipeirotis, P.G. *Mechanical Turk, Low Wages, and the Market for Lemons*. <http://www.behind-the-enemy-lines.com/2010/07/mechanical-turk-low-wages-and-market.html>, 2010.
 66. Ishii, H. and Kobayashi, M. ClearBoard: a seamless medium for shared drawing and conversation with eye contact. *Proc. CHI '92*, (1992), 525–532.
 67. Kamar, E., Hacker, S., and Horvitz, E. Combining Human and Machine Intelligence in Large-scale Crowdsourcing. *Proc. AAMAS '12*, (2012).
 68. Kearns, M., Suri, S., and Montfort, N. An experimental study of the coloring problem on human subject networks. *Science 313*, 5788 (2006), 824–827.
 69. Kensing, F. and Blomberg, J. Participatory Design: Issues and Concerns. *Computer Supported Cooperative Work (CSCW) 7*, 3 (1998), 167–185.
 70. Kerr, S. On the folly of rewarding A, while hoping for B. *Academy of Management Journal*, (1975), 769–783.
 71. Khanna, S., Ratan, A., Davis, J., and Thies, W.

- Evaluating and improving the usability of Mechanical Turk for low-income workers in India. *Proc. ACM Symposium on Computing for Development '10*, (2010).
72. Kittur, A., Chi, E.H., and Suh, B. Crowdsourcing user studies with Mechanical Turk. *Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, (2008), 453–456.
 73. Kittur, A., Khamkar, S., André, P., and Kraut, R.E. CrowdWeaver: visually managing complex crowd work. *Proc. CSCW '12*, (2012).
 74. Kittur, A., Lee, B., and Kraut, R.E. Coordination in collective intelligence: the role of team structure and task interdependence. *Proceedings of the 27th international conference on Human factors in computing systems*, (2009), 1495–1504.
 75. Kittur, A., Smus, B., Khamkar, S., and Kraut, R.E. Crowdforge: Crowdsourcing complex work. *Proceedings of the 24th annual ACM symposium on User interface software and technology*, (2011), 43–52.
 76. Kittur, A., Suh, B., and Chi, E.H. Can you ever trust a wiki?: impacting perceived trustworthiness in wikipedia. *Proceedings of the 2008 ACM conference on Computer supported cooperative work*, (2008), 477–480.
 77. Kittur, A. Crowdsourcing, collaboration and creativity. *XRDS 17*, 2 (2010), 22–26.
 78. Kochhar, S., Mazzocchi, S., and Paritosh, P. The anatomy of a large-scale human computation engine. *Proceedings of the ACM SIGKDD Workshop on Human Computation*, (2010), 10–17.
 79. Kulkarni, A., Gutheim, P., Narula, P., Rolnitzky, D., Parikh, T.S., and Hartmann, B. MobileWorks: Designing for Quality in a Managed Crowdsourcing Architecture. *IEEE Internet Computing To appear*, 2012.
 80. Kulkarni, A.P., Can, M., and Hartmann, B. Turkomatic: automatic recursive task and workflow design for mechanical turk. *Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems*, (2011), 2053–2058.
 81. Kuznetsov, S. Motivations of contributors to Wikipedia. *SIGCAS Comput. Soc.* 36, 2 (2006), 1.
 82. Lakhani, K. and Wolf, B. Why hackers do what they do: Understanding motivation and effort in free/open source software projects. In J. Feller, B. Fitzgerald, S.A. Hissam and K.R. Lakhani, eds., *Perspectives on Free and Open Source Software*. MIT Press, 2005, 3–22.
 83. Lasecki, W.S., Murray, K.I., White, S., Miller, R.C., and Bigham, J.P. Real-time crowd control of existing interfaces. *Proc. UIST '11*, ACM Press (2011).
 84. Lasecki, W.S., White, S.C., Murray, K.I., and Bigham, J.P. Crowd memory: Learning in the collective. *Proc. Collective Intelligence 2012*, (2012).
 85. Le, J., Edmonds, A., Hester, V., and Biewald, L. Ensuring quality in crowdsourced search relevance evaluation: The effects of training question distribution. *Proc. SIGIR 2010 Workshop on Crowdsourcing for Search Evaluation*, (2010), 21–26.
 86. Lewis, S., Dontcheva, M., and Gerber, E. Affective computational priming and creativity. *Proceedings of the 2011 annual conference on Human factors in computing systems*, (2011), 735–744.
 87. Little, G., Chilton, L.B., Goldman, M., and Miller, R.C. Turkit: human computation algorithms on mechanical turk. *Proceedings of the 23rd annual ACM symposium on User interface software and technology*, (2010), 57–66.
 88. Littler, C.R. Understanding Taylorism. *British Journal of Sociology*, (1978), 185–202.
 89. Malone, T.W. and Crowston, K. The interdisciplinary study of coordination. *ACM Computing Surveys (CSUR)* 26, 1 (1994), 87–119.
 90. Malone, T.W., Yates, J., and Benjamin, R.I. Electronic markets and electronic hierarchies. *Communications of the ACM* 30, 6 (1987), 484–497.
 91. Mao, A., Parkes, D.C., Procaccia, A.D., and Zhang, H. Human Computation and Multiagent Systems: An Algorithmic Perspective. *Proc. AAAI '11*, (2011).
 92. March, J.G. and Simon, H.A. Organizations. (1958).
 93. Marcus, A., Karger, D.R., Madden, S., Miller, R.C., and Oh, S. Counting with the Crowd. *In Submission to VLDB*, (2012).
 94. Mason, W. and Watts, D.J. Financial Incentives and the “Performance of Crowds”. *Proc. HCOMP '09*, ACM Press (2009).
 95. McGraw, I., Lee, C., Hetherington, L., Seneff, S., and Glass, J.R. Collecting voices from the cloud. *Proc. LREC*, (2010).
 96. Milner, R. *Communicating and mobile systems: the [symbol for pi]-calculus*. Cambridge Univ Pr, 1999.
 97. Mintzberg, H. An emerging strategy of “direct” research. *Administrative science quarterly* 24, 4 (1979), 582–589.
 98. Moran, T.P. and Anderson, R.J. The workaday world as a paradigm for CSCW design. *Proceedings of the 1990 ACM Conference on Computer-supported Cooperative Work*, (1990), 381–393.
 99. Nalimov, E.V., Wirth, C., Haworth, G.M.C., and Others. KQQKQQ and the Kasparov-World Game. *ICGA Journal* 22, 4 (1999), 195–212.
 100. Narula, P., Gutheim, P., Rolnitzky, D., Kulkarni, A., and Hartmann, B. MobileWorks: A Mobile Crowdsourcing Platform for Workers at the Bottom of the Pyramid. *Proc. HCOMP '11*, (2011).
 101. Nickerson, J.V. and Monroy-Hernandez, A. Appropriation and creativity: User-Initiated contests in scratch. *System Sciences (HICSS), 2011 44th Hawaii International Conference on*, (2011), 1–10.
 102. Norman, D.A. and Draper, S.W. *User centered*

- system design; new perspectives on human-computer interaction*. L. Erlbaum Associates Inc., 1986.
103. Noronha, J., Hysen, E., Zhang, H., and Gajos, K.Z. Platemate: crowdsourcing nutritional analysis from food photographs. *Proc. UIST '11*, (2011).
 104. Preece, J. and Shneiderman, B. The Reader-to-Leader Framework: Motivating Technology-Mediated Social Participation. *AIS Transactions on Human-Computer Interaction 1*, 1 (2009), 13–32.
 105. Quinn, A.J. and Bederson, B.B. Human computation: a survey and taxonomy of a growing field. *Proc. CHI '11*, (2011).
 106. Raddick, J., Lintott, C., Bamford, S., et al. Galaxy Zoo: Motivations of Citizen Scientists. *Bulletin of the American Astronomical Society*, (2008), 240.
 107. Rafaeli, S. and Ariel, Y. Online Motivational Factors: Incentives for Participation and Contribution in Wikipedia. In A. Barak, ed., *Psychological Aspects of Cyberspace*. Cambridge University Press, New York, NY, 2008.
 108. Raymond, E. The cathedral and the bazaar. *Knowledge, Technology & Policy 12*, 3 (1999), 23–49.
 109. Redmiles, D. and Nakakoji, K. Supporting reflective practitioners. *Software Engineering, 2004. ICSE 2004. Proceedings. 26th International Conference on*, (2004), 688–690.
 110. Reijers, H., Jansen-Vullers, M., Zur Muehlen, M., and Appl, W. Workflow management systems+ swarm intelligence= dynamic task assignment for emergency management applications. *Business Process Management*, (2007), 125–140.
 111. Resnick, P. and Zeckhauser, R. Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system. *Advances in Applied Microeconomics 11*, (2002), 127–157.
 112. Rittel, H.W.J. and Webber, M.M. Dilemmas in a general theory of planning. *Policy sciences 4*, 2 (1973), 155–169.
 113. Ritter, S., Anderson, J.R., Koedinger, K.R., and Corbett, A. Cognitive Tutor: Applied research in mathematics education. *Psychonomic bulletin & review 14*, 2 (2007), 249–255.
 114. Rogers, Y. Exploring obstacles: integrating CSCW in evolving organisations. *Proceedings of the 1994 ACM conference on Computer supported cooperative work*, (1994), 67–77.
 115. Rogstadius, J., Kostakos, V., Kittur, A., Smus, B., Laredo, J., and Vukovic, M. An Assessment of Intrinsic and Extrinsic Motivation on Task Performance in Crowdsourcing Markets. *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media: Barcelona, Spain*, (2011).
 116. Ross, J., Irani, L., Silberman, M.S., Zaldivar, A., and Tomlinson, B. Who Are the Crowdworkers? Shifting Demographics in Amazon Mechanical Turk. *alt.chi '10*, ACM Press (2010).
 117. Roth, A.E. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica 70*, 4 (2002), 1341–1378.
 118. Roy, D.F. Work satisfaction and social reward in quota achievement: An analysis of piecework incentive. *American Sociological Review 18*, 5 (1953), 507–514.
 119. Rzeszotarski, J. and Kittur, A. CrowdScape: interactively visualizing user behavior and output. *Proceedings of the 25th annual ACM symposium on User interface software and technology*, (2012), 55–62.
 120. Rzeszotarski, J.M. and Kittur, A. Instrumenting the crowd: using implicit behavioral measures to predict task performance. *Proc. UIST '11*, (2011).
 121. Sadlon, E., Barrett, S., Sakamoto, Y., and Nickerson, J.V. The Karma of Digg: Reciprocity in Online Social Networks. Paris, December 2008. *Proceedings of the 18th Annual Workshop on Information Technologies and Systems (WITS)*, (2008).
 122. Savage, N. Gaining wisdom from crowds. *Communications of the ACM 55*, 3 (2012), 13–15.
 123. Schmidt, K. and Bannon, L. Taking CSCW seriously. *Computer Supported Cooperative Work (CSCW) 1*, 1 (1992), 7–40.
 124. Schroer, J. and Hertel, G. Voluntary Engagement in an Open Web-Based Encyclopedia: Wikipedians and Why They Do It. *Media Psychology 12*, 1 (2009), 96.
 125. Shahaf, D. and Horvitz, E. Generalized task markets for human and machine computation. (2010).
 126. Shamir, B. and Salomon, I. Work-at-Home and the Quality of Working Life. *The Academy of Management Review 10*, 3 (1985), 455–464.
 127. Shaw, A.D., Horton, J.J., and Chen, D.L. Designing incentives for inexpert human raters. *Proceedings of the ACM 2011 conference on Computer supported cooperative work*, (2011), 275–284.
 128. Shen, M., Tzeng, G.H., and Liu, D.R. Multi-criteria task assignment in workflow management systems. *System Sciences, 2003. Proceedings of the 36th Annual Hawaii International Conference on Social Sciences*, (2003).
 129. Sheng, V.S., Provost, F., and Ipeirotis, P.G. Get another label? improving data quality and data mining using multiple, noisy labelers. *Proc. KDD '08*, ACM (2008), 614–622.
 130. Silberman, M.S., Irani, L., and Ross, J. Ethics and tactics of professional crowdwork. *XRDS 17*, 2 (2010), 39–43.
 131. Silberman, M.S., Ross, J., Irani, L., and Tomlinson, B. Sellers' problems in human computation markets. *Proc. HCOMP '10*, (2010).
 132. Skillicorn, D.B. and Talia, D. Models and languages for parallel computation. *ACM Computing Surveys (CSUR) 30*, 2 (1998), 123–169.

133. Smith, A. *The Wealth of Nations* (1776). *New York: Modern Library*, (1937), 740.
134. Snow, R., O'Connor, B., Jurafsky, D., and Ng, A.Y. Cheap and fast—but is it good?: evaluating non-expert annotations for natural language tasks. *Proc. ACL '08*, (2008).
135. Sorokin, A. and Forsyth, D. Utility data annotation with Amazon Mechanical Turk. *Proc. CVPR '08*, (2008).
136. Sterling, B. *A good old-fashioned future: stories*. Spectra, 1999.
137. Sterling, B. *Shaping things*. 2005.
138. Stohr, E.A. and Zhao, J.L. Workflow automation: Overview and research issues. *Information Systems Frontiers* 3, 3 (2001), 281–296.
139. Surowiecki, J. *The Wisdom of Crowds*. Random House, New York, 2005.
140. Tang, J.C., Cebrian, M., Giacobe, N.A., Kim, H.-W., Kim, T., and Wickert, D. “Beaker”. Reflecting on the DARPA Red Balloon Challenge. *Communications of the ACM* 54, 4 (2011), 78.
141. Taylor, F.W. *The principles of scientific management*. Harper & Brothers, New York, 1911.
142. Trist, E.L. *The evolution of socio-technical systems: A conceptual framework and an action research program*. Ontario Quality of Working Life Center.
143. Van de Ven, A.H., Delbecq, A.L., and Koenig Jr, R. Determinants of coordination modes within organizations. *American sociological review*, (1976), 322–338.
144. Viswanath, B., Mondal, M., Clement, A., et al. Exploring the design space of social network-based Sybil defenses. *Proceedings of the 4th International Conference on Communication Systems and Network (COMSNETS'12)*, (2012).
145. Wang, G., Wilson, C., Zhao, X., et al. Serf and Turf: Crowdturfing for Fun and Profit. *Arxiv preprint arXiv:1111.5654*, (2011).
146. Weld, D.S., Adar, E., Chilton, L., Hoffmann, R., and Horvitz, E. Personalized Online Education—A Crowdsourcing Challenge. *Workshops at the Twenty-Sixth AAAI Conference on Artificial Intelligence*, (2012).
147. Welinder, P., Branson, S., Belongie, S., and Perona, P. The multidimensional wisdom of crowds. *Neural Information Processing Systems* 6, 7 (2010), 1—9.
148. Williamson, O.E. The economics of organization: The transaction cost approach. *American journal of sociology* 87, 3 (1981), 548–577.
149. Woolley, A.W., Chabris, C.F., Pentland, A., Hashmi, N., and Malone, T.W. Evidence for a collective intelligence factor in the performance of human groups. *science* 330, 6004 (2010), 686–688.
150. Yu, H., Estrin, D., and Govindan, R. A hierarchical proxy architecture for Internet-scale event services. *Enabling Technologies: Infrastructure for Collaborative Enterprises, 1999. (WET ICE '99) Proceedings. IEEE 8th International Workshops on*, (1999), 78–83.
151. Yu, J. and Buyya, R. A taxonomy of scientific workflow systems for grid computing. *ACM Sigmod Record* 34, 3 (2005), 44–49.
152. Yu, L. and Nickerson, J.V. Cooks or cobblers?: crowd creativity through combination. *Proceedings of the 2011 annual conference on Human factors in computing systems*, (2011), 1393–1402.
153. Yu, L. and Nickerson, J.V. An Internet-Scale Idea Generation System. *ACM Transactions on Interactive Intelligent Systems* 3, 1 (2013).
154. Zhang, H., Law, E., Miller, R.C., Gajos, K.Z., Parkes, D.C., and Horvitz, E. Human Computation Tasks with Global Constraints. *Proc. CHI '12*, (2012).
155. Zhao, L., Sukthankar, G., and Sukthankar, R. Robust Active Learning Using Crowdsourced Annotations for Activity Recognition. *Workshops at the Twenty-Fifth AAAI Conference on Artificial Intelligence*, (2011).
156. Zhu, H., Kraut, R.E., Wang, Y.-C., and Kittur, A. Identifying shared leadership in Wikipedia. *Proceedings of the 2011 annual conference on Human factors in computing systems*, ACM (2011), 3431–3434.
157. Zittrain, J. Ubiquitous human computing. *Philosophical Transactions of the Royal Society A* 366, (2008), 3813–3821.
158. Zittrain, J. Human Computing's Oppenheimer Question. *Proceedings of Collective Intelligence*, (2012).