

Turkopticon: Interrupting Worker Invisibility in Amazon Mechanical Turk

Lilly C Irani

UC Irvine, Department of Informatics
Irvine, CA 92697
lirani@ics.uci.edu

M. Six Silberman

Bureau of Economic Interpretation
six@economicinterpretation.org

ABSTRACT

As HCI researchers have explored the possibilities of human computation, they have paid less attention to ethics and values of crowdwork. This paper offers an analysis of Amazon Mechanical Turk, a popular human computation system, as a site of technically mediated worker-employer relations. We argue that human computation currently relies on worker invisibility. We then present Turkopticon, an activist system that allows workers to publicize and evaluate their relationships with employers. As a common infrastructure, Turkopticon also enables workers to engage one another in mutual aid. We conclude by discussing the potentials and challenges of sustaining activist technologies that intervene in large, existing socio-technical systems.

Author Keywords

Activism; infrastructure; human computation; Amazon Mechanical Turk; design; ethics

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Crowdsourcing and human computation are often described as a new frontier for HCI research and creativity, and for technological progress more broadly. CHI researchers have built word processors powered by crowds. Others have shown how usability and visualization evaluations can be taken out of the lab and into the natural environments of crowdworkers.

These frontiers, however, are enabled by the novel organization of digital workers, distributed across the world and organized through task markets, APIs, and network connections. This paper looks behind the walls of abstraction that enable human computation in one specific system, Amazon Mechanical Turk (AMT). We present workers' occupational hazards as human computers, and explain the activist project we developed in response. Turkopticon, the project we present, is a tool in its fourth

year of deployment. The system receives 100,000 page views a month and has become a staple tool for many AMT workers, installed over 7,000 times at time of writing.

Turkopticon allows workers to create and use reviews of employers when choosing employers on AMT. Building and maintaining the system, as well as communicating about the system with workers, has offered us a distinct vantage point into the social processes of designing interventions into large-scale, real world systems. Turkopticon supports a thriving collective of workers engaged in mutual aid, brought together by our simple browser extension and web-based technology.

This paper makes several contributions. First, it offers a case study designing an intervention into a highly distributed microlabor system. Second, it shows an example of systems design incorporating tools feminist analysis and reflexivity. Rather than conducting HCI research to reveal and represent values and positions, and then building systems to resolve those political differences, we built a system to make worker-employer relations visible and to provoke ethical and political debate. Third, this paper contributes lessons learned from intervening in existing, large-scale socio-technical systems (here, AMT) from its margins.

METHOD AND OUR STANCE

This paper draws on four years of participant-observation as design activists within AMT worker and technologist communities. Turkopticon grew out of a tactical media art project intended to raise questions about the ethics of human computation. Tactical media, one tradition within activist art, emphasizes developing urgent, culturally provocative interruptions and resistance through the design of media [13, 19, 21]. In addition to the interviews, observation, and casual conversation that feature in many HCI ethnographies, our encounters with Turk workers began through highly mediated "Human Intelligence Tasks" and feedback around Turkopticon. (We began this research in 2008, prior to the growth of popular online worker forums turkernation.com and mturkforum.com.)

We conducted several informal surveys through Mechanical Turk. 67 respondents answered our open-ended question survey about what they would desire as a "Workers' Bill of Rights." Points of agreement among worker respondents on this survey became the basis for the design of Turkopticon.

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.

CHI 2013, April 27–May 2, 2013, Paris, France.

Copyright © 2013 ACM 978-1-4503-1899-0/13/04...\$15.00.

The first author complements participant-observation, as a system builder of Turker tools, with observation and open-ended interviews with AMT employers. She attended a major crowdsourcing conference as well two smaller crowdsourcing meetups. She also conducted open-ended interviews with four employers on AMT and numerous conversations with other employers. These ethnographic data contextualize the data we generate as we design and maintain Turkopticon.

Over the course of this research, each of our stances developed as a result of our own involvement with the workers through the project, and through our evolving understandings of the broader crowdsourcing community. We began highly critical of the fragmentation of labor into hyper-temporary jobs, seeing them as an intensification of decades-old US trends toward part-time, contingent work for employer flexibility and cost-cutting [3, 35]. AMT, it seemed to us produced temporary employees at “the speed of thought,” to borrow Bill Gates’ promissory turn of phrase, precisely by forgetting about ergonomics, repetitive stress injuries, and minimum wage laws. We were biased – decidedly so.

Our biases were validated by some workers and challenged by others. For each one who reported needing the money to pay for rent or groceries, there was another who did it for fun or to “kill time.”

We highlight our own stances under advisement of Borning and Muller who, with many feminist scholars, call for researchers to shed trappings of objective authority and account for how our own contexts and assumptions shape our research practices [7]. However, it is not only that our biases distort our perception of reality that is out there in the world of AMT work. Certainly, we have much to learn about how workers feel about their work and the problems they encounter, as we have published. But we also intervene in AMT by building a technology used by its workers. By intervening in the system as designers and as observers, we change the reality of the system itself [4, 29, 41]. The ethical challenges and issues faced by workers, and the ethical issues we face as researchers, are produced in the encounters between us, the workers, and Turkopticon. This paper offers a snapshot of the lessons we have learned and their implications for design practice at this point in the evolving socio-technical system.

First, we explain AMT, focusing on the kinds of worker-employer relationships enabled by the system. We then describe our motivations for building Turkopticon, the design of the system, and learnings relevant to the design of political and activist technologies.

BACKGROUND: AMAZON MECHANICAL TURK (AMT)

Amazon Mechanical Turk is a website and service operated by Amazon as a meeting place for requesters with large volumes of microtasks and workers who want to do those tasks, usually for money [24]. These tasks often add up to a few dollars an hour for those experienced with the platform.

Amazon legally defines the workers as contractors subject to laws designed for freelancers and consultants; this framing attempts to strip workers of minimum wage requirements in their countries. United States workers are a significant minority, numbering at 46.8% in recent surveys [24]. This framing, however, has not been tested in courts, and courts have deemed similar framings of distributed, non-computer data work illegal [18].

AMT can be described many ways. Explaining it as a microlabor marketplace draws attention to pricing mechanisms, how workers choose tasks, and how transactions are managed. Explaining it as a crowdsourcing platform draws attention to the dynamics of mass collaboration among workers, the aggregation of inputs, and the evaluation of the crowdsourced outputs. Explaining AMT as a source of human computation resources, however, is consistent with how both the computing research community and Amazon’s own marketing frames the system [e.g. 28].

Dividing data work into small components is not itself new. A 1985 case, *Donovan vs DialAmerica*, tells of an earlier version of AMT-style labor. An employer sent cards with names to home workers hired as independent contractors. These contractors had to ascertain the correct phone number for each name; they were paid per task. Courts decided that these workers were in fact employees entitled to minimum wage under the Fair Labor Standards Act (FLSA) [18, p.136]. Since the late 1990s, American companies have hired Business Process Outsourcing firms in English-speaking countries with lower costs of living to perform large volumes of data processing tasks that cannot be fully automated.¹

Humans-as-a-service

AMT jumps beyond these older forms of information work by setting workers up as resources that can be directly integrated into existing computer systems as “human computation.” When Jeff Bezos launched AMT to an MIT audience in 2006, he announced: “You’ve heard of software-as-a-service. Now this is human-as-a-service” [5]. Since launch, AMT has been marketed as one of Amazon’s Web Services, alongside silicon computational cycles and data storage in the cloud. Bloggers and technologists have followed suit, both in published sources and conferences and meetups we attended, calling AMT a “Remote Person Call” (playing off of “Remote Procedure Call”) and “the Human API.” Crowdsourcing company CrowdFlower even coined the neologism “Labor-as-a-service (LaaS)” to market the value of crowdsourced workforces to companies.

This combination of abstraction and service orientation in both the metaphors and infrastructural forms suggest a particular kind of social relationship. “As-a-service” draws

¹ In 2012, a US Mechanical Turk worker filed a class action suit against Crowdflower for violations of FLSA. The outcome of the suit was unknown at time of press.

meaning from commonplace resonances of service. To serve is to make labor and attention available for those served; to promise service is to be bound, by duty or by wage, to the will of the served. Among computer scientists, “as-a-service” builds off of this common sense meaning and more specifically suggests a division of technical labor by which programmers can access computational processing functions housed and maintained on the Internet and by someone else. As long as the service keeps running, programmers need not concern themselves with where the code is running, what kind of machine it runs on, or who keeps the code running, but only the proper protocol for issuing the call through a computer and receiving the response. As-a-service suggests an arrangement of computers, networks, system administrators, and real estate that allows programmers to access a range of computer services remotely and instantly.

Framing workers on AMT as computational services is more than just rhetorical flourish. Through AMT, employers can literally access workers through APIs. Though a web form-based interface is available, the API allows AMT employers can put tasks into the workforce and integrate worker output directly into their algorithms. Techniques for integrating workers into computational technologies in this way have been pioneered in HCI, in databases research, and in industry (see [42] for a summary). Twitter, for example, has recently open sourced a visual toolkit for running human judgment experiments on AMT [12]. These experiments are a key component of developing, evaluating, and training search and ranking algorithms. Twitter’s toolkit offers an interface for building these experiments, providing monitoring tools and visualizations interfacing with AMT’s 24/7, massively distributed workforce through APIs. CrowdFlower also builds atop AMT’s APIs, offering crowdsourced data processing tools tailored to needs common to different industries.

We see here, then, that AMT brings together crowds of workers as a form of infrastructure, rendering employees into reliable sources of computation. As established organizations develop and publicly release tools for the system, they embed computational microwork firmly in existing technological practices and systems. AMT is becoming infrastructure in the sense that Star & Ruhleder have analyzed it: AMT is shared, AMT is incorporated into existing shared practices, and ideally, AMT is ready-to-hand and worked *through* not *on*. Working technological infrastructures, in Star & Ruhleder’s analyses, are used with such fluency that they become taken-for-granted, humming quietly and usefully in the background. The infrastructures kept humming dutifully in the background in AMT are the socio-technical system of workers interacting with employers through APIs, spreadsheets, and minimal web-based task forms.

Ruhleder and Star famously called for going beyond a consideration of *what* is infrastructure to a consideration of *when* is infrastructure [34]. Asking *when* is infrastructure

recognizes that systems that might hum along beyond notice in one moment might break down and require maintenance and repair in another. And a system that might hum along beyond notice for an end-user might be very much the focus of attention for those in charge of maintaining it. The question “when is infrastructure?” then, suggests also asking, “for whom is infrastructure?” When it is working as infrastructure, AMT platform clearly hums along supporting the work of employers — the programmers, managers, and start up hackers who integrate human computation into their technologies. In this light, that the design features and development of AMT has prioritized the needs of employers over workers is not surprising.

Further, by hiding workers behind web forms and APIs, AMT helps employers see themselves as builders of innovative technologies, rather than employers unconcerned with working conditions. Suchman argues that there are “agencies at the interface” that reconfigure the relations among humans and machines, making both what they are [40]. AMT’s power lies in part in how it reconfigures social relations, rendering Turk workers invisible [37], redirecting focus to the innovation of human computation as a field of technological achievement.

Employing Humans-as-a-Service

In this section, we explain basic features of AMT and show how the design prioritizes the needs of employers.

AMT employers define HITs on AMT by creating web-based forms that specify an information task and allow workers to input a response. Tasks include structuring unstructured data (e.g. entering a given webpage into an employer’s structured form fields), transcribing snippets of audio, and labeling an image (e.g. as pornography or violating given Terms of Service). Employers define the structure of the data workers must input, create instructions, specify the pool of information that must be processed, and set a price. (Ipeirotis [22] offers an excellent background.)

The employer then defines criteria that candidate workers must meet to work on the task. These criteria include the worker’s “approval rating” (the percentage of tasks the worker has performed that employers have approved and, by consequence, paid for), the worker’s self-reported country, and whether the worker has completed certain skill-specific qualification exams offered on the platform. This filter approach to choosing workers, as compared to more individualized evaluation and selection, allows employers to request work from thousands of temporary workers in a matter of hours.

Once a worker submits work, the employer can choose whether to pay for it. This discretion allows employers to reject work that does not meet their needs, but also enables wage theft. Because AMT’s participation agreement grants employers full intellectual property rights over submissions regardless of rejection, workers have no legal recourse against employers who reject work and then go on to use it.

Help Idealist.org Find New Contact Information for Outdated Nonprofit Organizations		View a HIT in this group
Requester: Idealist.org	HIT Expiration Date: Dec 2, 2009 (14 weeks)	Reward: \$0.03
	Time Allotted: 60 minutes	HITs Available: 6071
Description: Search the internet for new contact people for organizations on file on Idealist.org		
Keywords: Idealist , Idealist , org , nonprofit , NGO , Action , without , Borders		
Qualifications Required: HIT approval rate (%) is not less than 95		

Find articles on: Travel		View a HIT in this group
Requester: David Smith	HIT Expiration Date: Aug 28, 2009 (2 days 17 hours)	Reward: \$0.13
	Time Allotted: 30 minutes	HITs Available: 2911

Fig. 1: What a worker sees: the Human intelligence Tasks (HITs) list on AMT.

Employers vet worker outputs through automated approaches such as qualifying workers through test tasks to which the correct answer is known or requesting responses to a single input from several workers and algorithmically eliminating any answers that do not agree with the majority.

Within this large scale, fast moving, and highly mediated workforce, dispute resolution between workers and employers becomes intractable. Workers dissatisfied with a requester’s work rejection can contact the requester through AMT’s web interface. Amazon does not require requesters to respond and many do not; several requesters have noted that a thousand to one worker-to-requester ratio makes responding cost prohibitive. In the logic of massive crowd collaborations, dispute resolution does not scale. Dahn Tamir, a large-scale requester, explained a logic the first author heard from several Turk employers:

“You cannot spend time exchanging email. The time you spent looking at the email costs more than what you paid them. This has to function on autopilot as an algorithmic system...and integrated with your business processes.”

Instead of eliciting a response, workers’ dispute messages become signals to the employer. Rick, a CEO of a crowdsourcing startup, explained to me that messages from workers signal the algorithm’s performance in managing workers and tasks. If a particular way of determining “correctness” for a task results in a large number of disputing messages, Rick’s team will look into revising the algorithm but rarely will retroactively revise decisions. Algorithmic management, here, precludes individually accountable relations.

Workers have limited options for dissent within AMT itself. Resistance through incorrect answers can simply be filtered out through employer’s algorithmic tests of correctness. Dissatisfied workers’ within AMT had little option other than to leave the system altogether. Because AMT treats workers interchangeably and because workers are so numerous (tens of thousands by the most conservative estimates), AMT can sustain the loss of workers who do not accept the system’s terms.

Turkopticon has been designed to offer workers a way to dissent, holding requesters accountable and offering one another mutual aid.

MOTIVATING TURKOPTICON

Turkopticon developed as an ethically-motivated response to workers’ invisibility in the design of AMT. We were troubled by a number of issues in our first encounters with AMT, not only worker invisibility. Workers, even in the US, are paid below minimum wage in many cases. Technologist and research discourse seemed unconcerned with the human costs of human computation. Individuated workers had little opportunity to build solidarity, offering them little chance of creating sufficiently coordinated actions to exert pressure on employers and Amazon.

Rather than working from our own intuitions, however, we took seriously the possibility that this new form of work also might offer workers benefits and pleasures that we did not understand, or cause troubles we could not anticipate. Survey research on Turk worker motivations, for example, reports that though a significant minority of workers rely on their income from the platform to pay for household expenses. At the same time, other workers report working for fun or to pass the time while bored (sometimes even at another job) [23, 22, 33].

Workers’ “Bill of Rights”

To provoke workers’ imaginations about the infrastructural possibilities, we placed a task onto AMT asking workers to articulate a “Worker’s Bill of Rights” from their perspective. We chose this approach over a more neutral battery of questions because of the highly mediated nature of our interactions with workers through the medium of the HIT. Workers paid per task — of which our question was one — provided short answers to open-ended questions based on our past experiences questioning workers in the platform. Asking a provocative question drew stronger, more detailed responses oriented towards concerns of crowdsourcing ethics. We also sought permission from workers to publish their responses on the web in hopes of generating interaction between workers and broader publics concerned with crowdsourcing.

Our work treated crowdsourcing ethics as an open question about a new technology, still under negotiation. In structurationist terms, practices and meanings of the technology had not yet stabilized [32]. Our ethical questions, then, were not trying to get at some underlying, stable truth, but rather at ongoing ethical and political negotiations among participants in crowdsourcing systems. Like Bruckman and Hudson, we gathered empirical data on workers' ethics — here framed as rights — to explore the ethical dimensions of crowdsourcing [9]. Rather than draw firm conclusions here, however, we continue to keep the debate open. We grapple with the problem of advocacy as explained by Bardzell [2], in which Feminist HCI practitioners both seek to bring about social progress, but also question their own images of what such social progress looks like. By publishing responses to our questions and building Turkopticon, as we will discuss, we sought to provoke debate about progress in crowdsourcing and make questions of work conditions visible among technologists, policy makers, and the media.

Workers' responses to the question of a "Bill of Rights" revealed a range of concerns, some broadly expressed among workers and others that polarized.

Of our 67 responses [42], workers recurrently raised a number of issues:

- 35 workers felt that their work was regularly rejected unfairly or arbitrarily
- 26 workers demanded faster payment (Amazon allows employers 30-days to evaluate and pay for work)
- 7 explicitly mentioned a "minimum wage" or "minimum payment" per HIT
- 14 mentioned "fair" compensation generally
- 8 expressed dissatisfaction with employers' and Amazon's lack of response to their concerns

The consequences of these occupational hazards for workers included lost or delayed income, accidental download of malware that damaged their computers, and reduced worker "approval ratings." Approval ratings are one of the few ways employers can filter workers. When an employer rejects an employer's work, whether because it did not meet their needs or simply so they employer did not have to pay, the worker's approval rating goes down. If the rating goes too far down, AMT will hide tasks requiring high ratings from the worker. Lost approval ratings, then, are lost opportunities for work which make it even more difficult to accumulate experiences to raise the rating again.

A Sense of Fairness

Beyond the inconveniences and dangers of doing Turk work, several workers articulated a more general frustration we characterize as a sense of fairness. This sense came through in numerous responses that requesters ought to respond to questions from workers, that requesters ought to

justify their rejections, and that workers have the right to confront employers about those rejections.

A number of workers directed their frustrations towards Amazon itself. One worker was so frustrated that he or she thanked the first author by name for posting the HIT and offering an opportunity to express his anger: "I don't care about the penny I didn't earn for knowing the difference between an apple and a giraffe, but I'm angry that MT will take requester's money but not manage, oversee, or mediate the problems and injustices on their site."

Another worker noted the imbalance in Amazon's priorities as they developed the AMT platform:

"I would also like workers to have more of a say around here, so that they can not easily be taken advantage of, and are treated fairly, as they should be. Amazon seems to pay more credence to the requesters, simply ignoring the fact that without workers, nothing would be done!"

We confirmed this priority with prominent requesters as well as a source close to Amazon who wished to remain anonymous. Because Amazon collects money for task volume, Amazon has little reason to prioritize worker needs in a market with a labor surplus.

Mutual Aid for Accountability

Our exploratory interactions with workers left us with no unified image of what workers are like and what intervention might be "appropriate." Those workers who suggested action offered diverse ways forward. Some were interested in a forum in which Turkers could air concerns publicly without censorship or condescension, and worker visibility and dignity more generally. Others were interested in a way to build long-term work relationships with prolific requesters, and worker-requester relations generally. Several respondents asked for unionization, while several others volunteered their aversion to unions.

There were few shared values and priorities that could guide the development of an infrastructure of mutual aid. There were, however, possibilities for creating partial alliances — points of common cause across diverse workers. Donna Haraway, a feminist STS scholar, argues for partial connections — alliances built on common cause rather than common experience or identity — as a way to sustain political and ethical action across people with irreducible differences [20].² We took inspiration from this approach.

² Haraway's argument responded to criticisms that socialist feminism, a Marxist analysis of gender, claimed white women's experiences of gender marginalization as common cause for *all* women. Crenshaw, for example, countered that women exist at the intersection of race, class, and gender categories; each intersection created specific kinds of vulnerabilities. What Haraway proposed was a way to make progressive interventions without making universalizing claims about the issues of all women. She did this by proposing that women, as irreducibly different "cyborgs," build alliances based on common cause and partial connections [20].

Motivated by responses to the “Bill of Rights,” we designed and built Turkopticon. Turkopticon responded in part to the occupational hazards of Turking listed above. We also built Turkopticon to offer workers ways of supporting one another in context of their existing practices. The system allows workers to make their relationships with employers visible and call those employers to account. As workers build up the record of relationships with employers, they also build up a commons together with other contributors. By explicitly designing for scales beyond the individual or the dyadic relationship, we sought to build up a group of people who see their interests as aligned with others [16, 31]. Dourish called this the design of politics; he calls for moving beyond the user-technology dyad that often defines design interventions to the creation of larger scale collectives and movements building on social software.

The crowd we wanted to mobilize into a collective, however, was constituted by an infrastructure we had no control over – the AMT platform itself. In contrast to the collectives Dourish seeks to mobilize through Facebook, or the Internet hackers Chris Kelty describes as building the infrastructure that make their association possible [26, p.3], our task was to create a means of association people whose common cause was their work on AMT but who lack the technical skills to build infrastructures of assembly. Rather than design a system anew, our work was to graft a new infrastructure onto an existing one.

TURKOPTICON: THE SYSTEM

Turkopticon is a browser extension for Firefox and Chrome that augments workers’ view of their AMT HIT lists with information other workers have provided about employers (“Requesters” in AMT parlance). Workers enter reviews of employers that they have worked with, entering ratings of four qualities of employers as well as an open-ended comment explaining their rating. These reviews are available on the Turkopticon website; workers can view both recent reviews, as well as all reviews for a particular requester, identified by a unique Amazon requester ID.

Turkopticon is named for panopticon, a prison surveillance design most famously analyzed by Foucault. The prison is round with a guard tower in the center. The tower does not reveal whether the guard is present, so prisoners must assume they could be monitored at any moment. The *possibility* of surveillance, the theory goes, induces prisoners to discipline themselves. Turkopticon’s name cheekily references the panopticon, pointing to our hope that the site could not only hold employers accountable, but induce better behavior.

Going beyond simply a review site, we designed Turkopticon to fit into workers existing Turking workflow. The browser extension – a Javascript userscript packaged for both Firefox and Chrome – works by searching the document object model (DOM) of AMT pages as the worker browses. We locate links that contain requester IDs in their target URLs, choose the link that is a requester

name, and insert a small CSS button next to the name that, on mouse-over, launches more details on that requester. The extension issues an XMLHttpRequest for details we have on the requester that then load in the background as the rest of the “Available HIT” page renders.

The embedded review overlay contains both averaged ratings of the requester, and a link to view all reviews and open-ended comments on the requester on our website. (See figure 2.) From this overlay, workers can also review requesters. When the worker clicks the requester review link, we take them to Turkopticon’s requester review form with the requester ID we strip from page’s underlying HTML pre-populating the review’s form field. The embedded overlay is available anywhere in the AMT interface where a worker might see a requester: both at points where they are selecting HITs and where they are checking approval and payment status for submitted HITs.

Standardizing Requester Reputations

We now turn to how the kind of data we decided to collect on requesters. Because the AMT model often has workers doing HITs from a large number of employers in a session, we needed to offer workers a quick way to assess employers. We also saw in the Bill of Rights that workers were not unified in what they valued in an employer. Some wanted a short response time while others did not care, for example. By taking ratings on various qualities rather than taking an aggregating rating in the style of product review sites, we offered workers discretion in evaluating the ratings.

Turkopticon collects quantitative ratings from reviewers on four qualities that we hypothesized would be relevant based on the Workers’ Bill of Rights survey.

- **Communicativity:** How responsive has this requester been to communications or concerns you have raised?
- **Generosity:** How well has this requester paid for the amount of time their HITs take?
- **Fairness:** How fair has this requester been in approving or rejecting your work?
- **Promptness:** How promptly has this requester approved your work and paid?

A score of “0” means we have no data for that attribute

We also require workers to enter a free-form text comment to contextualize their scores. We provide the free-form box so that workers can share more nuanced, fine-grained stories of their experiences. We require workers to fill it, however, because the substance of testimonials is one of the ways other workers can evaluate other workers credibility.

Bootstrapping a Collective System

Turkopticon is nothing without users and their reviews; like many CSCW systems, it requires a “critical mass” to serve users at all [1]. How to launch a brand new system when the system has no content? Generating community around the project was difficult because workers were so invisible,

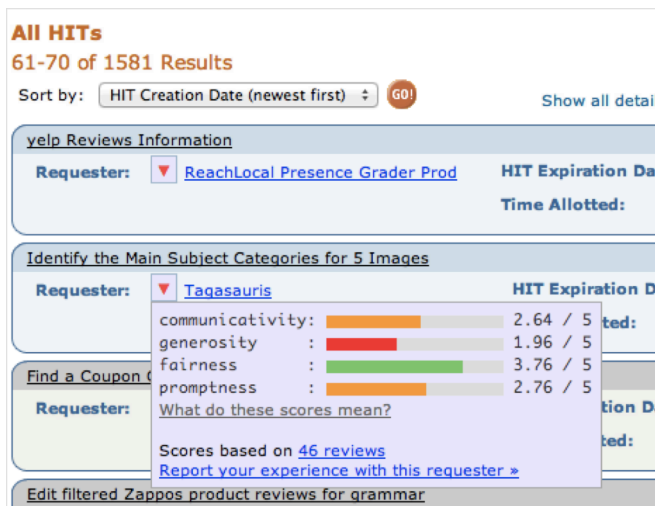


Fig. 2: The Turkopticon browser add-on adds information about requesters provided by other workers.

largely separated from one another, and today’s prominent worker forums (e.g. TurkerNation or mturkforum) were much smaller.

We overcame this problem by enlisting the support of DoloresLabs, a crowdsourcing company that builds custom toolkits for employers wishing to employ Mechanical Turk labor. DoloresLabs created a task for our team with a list of prominent requesters and solicited 300 initial reviews for which it compensated workers. The initial reviews seeded our database so new users installing Turkopticon could immediately integrate the tool into their workflow. Rather than requiring initial users to produce reviews, our bootstrapping allowed for users to consume the reviews we hoped they would eventually produce and improve upon.

Making alliance with a prominent employer in the Mechanical Turk system was a double-edged sword. DoloresLabs supported us because they believed that crowd labor industries would benefit from a fairer labor market; Turkopticon promised to remedy the information asymmetry between workers and employers, repairing Mechanical Turk into a more “transparent” marketplace [6]. Our team, by contrast, built Turkopticon in part to draw attention to commodification and exploitation in large-scale crowdsourcing markets. Just as the Turkopticon tool was a way of building partial connections across workers, the Turkopticon design process made partial connections despite different visions for the future of crowdsourcing.

Reputation without Retribution

Turkopticon attempts to prevent employers from retaliating against workers writing reviews by obfuscating workers’ email addresses. As we designed Turkopticon, we anticipated that workers would fear retribution for writing critical reviews. Our discussions with workers on forums have confirmed this at least for some workers. At tension with the need for anonymity, however, is the need for reputation among users of the system. There are high

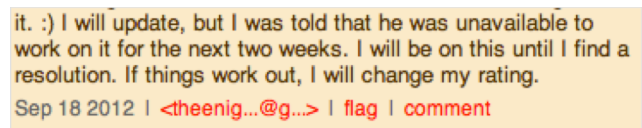


Fig. 3: Workers are protected from retribution by the obfuscation of their email addresses.

incentives for requesters to write positive reviews for themselves and, in practice, requesters attempt to do this.

We balanced the need for anonymity with reputation by displaying users’ reviews signed with a partially obfuscated email address. This email address is then linked to a page that shows all reviews written by that user. Readers of reviews can make judgments about the credibility of workers by evaluating other contributions by the user and making their own decision about whether to engage the employer.

Comment Moderation

After two years of running the tool unmoderated, we developed a set of user interface designs to allow selected users to moderate comments on the site. The mechanism is technically simple, leaning on existing social practices and community reputation. Any Turkopticon user can flag a review. A moderator has to add a second flag to hide the review from the site.

We selected our first cohort of moderators by calculating the most prolific reviewers on the site, emailing them invitations to moderate Turkopticon, and posting the list of those who accepted invitations to a widely read worker forum. We left nominations up for a week and received no objections, so we proceeded.

In selecting moderators, we also attempted to align Turkopticon with other worker forums in two ways. First, we selected moderators from the worker community who were engaged in debates and movements in worker forums that we, as non-workers, had little visibility into. By letting moderators in, we also gave them visibility and input into our design processes; based on this inside view, these moderators have been able to vouch for us during critical junctures where a bug or misunderstood feature triggers suspicions among users.

Along with moderation, we also introduced an option for workers to take on screen names – self-chosen identifiers – in place of their obfuscated email addresses. This simple measure has made it possible for reviewers to choose to harmonize their Turkopticon identity with their identity in other forums. We do not, however, force any harmonization.

We rely on primarily social moderation, by a small number of moderators, for several reasons. First, automated approaches are difficult to implement in practice because they cannot account for community-specific and emergent norms [38]. Within the space of social moderation, broad-based community moderation (e.g. Slashdot [27]) is susceptible to vandalism because our users are potentially

from two competing classes with opposed incentives. Requesters could easily make an account and begin flagging negative reviews they have received, or even pay Turk workers to down vote their reviews. Moderators draw on knowledge from their involvement in other worker forums to judge the credibility of reviews in question.

DISCUSSION

Strengthening Ties Through Maintenance and Repair

Though HCI has conventionally been concerned with the design, deployment, and evaluation of technological artifacts, the social and technical life of Turkopticon, like any technology, depends on ongoing maintenance and repair [25]. Certainly, we do ongoing technical maintenance. For example, we have to rebuild the extension when Firefox and Chrome release versions with new requirements of add-ons; server load that grew with use demanded that we rewrite code to make more efficient use of our servers resources.

Less remarked on, however, is the work of keeping up with changing design requirements as worker and requester practices change. Comment moderation to cull increasing requester reviews and profanity was one such change, already discussed. We also recently augmented the requester review form with a toggle indicating whether a requester violates Amazon's Terms and Conditions. These design changes reflect changing norms as the kinds of tasks and practices on AMT shift.

As important as the specific design features that we add and upkeep are the community relationships we build and strengthen through this ongoing maintenance of Turkopticon. We learn of concerns and confusions through our user support forum, through our email, and through our moderators who face emerging review practices on the frontline of the Turkopticon reviews page. We, as systems designers and maintainers, gain from highly engaged workers who help us understand what it means to see like a Turk worker and keep up with changes to their evolving practices. We enlist moderators in discussions of web site policy and interaction design, and alter and repair the technology in response to their requests and observations. Moderators here are not objects to be observed by us, but experts in their own right who participate in the collective activism of keeping Turkopticon thriving. (Bardzell and Bardzell have also argued for the incorporation of experts into activist design.)

This work of maintenance and upgrading, undertaken with the participation of workers, does more than offer insight into needs and requirements. This work strengthens ties and builds solidarity among workers collaborating on the practical, shared, and political circumstances they face as crowdworkers. Dourish has argued that HCI research often takes market framings for granted, individuating users as decision-makers to be persuaded or empowered [16]. Framings of social computing that emphasize networks and interaction can similarly frame collectivity as an aggregation of individuals. We call on HCI researchers to instead see

technology design's potential for sustaining new polities that can become powerful foundations for social change.

Tactical Quantification

Although quantification has myriad problems as a description of lived practice, Turkopticon employs *tactical quantification* to enable worker interaction and employer accountability while integrating into the rhythms of AMT. Tactical quantification is a use of numbers not because they are more accurate, rational, or optimizable [10], but because they are partial, fast, and cheap – a way of making do in highly constrained circumstances.

We were skeptical of quantifying workers' rich experiences and diverse frustrations, conditioned by their diverse social positions and needs. HCI researchers have raised a number of critiques of quantification in computational systems. Quantification has been associated with failed, injurious modernist attempts to model, rationalize, and optimize messy real world systems. These models necessarily universalize and simplify [10]. In the hands of powerful actors, quantifying, approximate models can drive policies that attempt to form the world in models' images [17].

The use of Turkopticon in the wild has, unsurprisingly, borne out some of these concerns. The "generosity" category, for example, has strained under the weight of representing such a subjective assessment. Workers in India accustomed to much lower salaries and cost of living than Americans may feel that a job averaging \$2 an hour is generous, while an American might balk at such a rate.

Standardizing ratings into quantified buckets was instead a compromise we made to our own values as designers in negotiating the power relations of the AMT ecosystem. AMT emphasizes speed and scale [11, 36]. To attract and retain users, we had to begin with the norms of the infrastructure in which we intervened, lest we push too far and become incompatible.

In this sense, Turkopticon is not an expression of our own values, or even the values of the users we interviewed, but a compromise between those values and the weight of the existing infrastructural norms that torqued our design decisions as we intervened in this powerful, working real world system. In their analyses of the consequences of infrastructural classifications, Bowker and Star use the concept of torque to describe the way people's lives can be twisted and shaped as they are forced to fit classification systems and infrastructures, such as racial classifications on government documents or disease categorizations. People live messy, fluid lives that can fall out of sync with the rhythms, categories, and temporality of the infrastructure [8,p.190]. Bowker and Star note that more powerful actors do not experience torque as they determine the categories of the infrastructure and often experience those categories as natural. We were situated at the margins of a large, working sociotechnical system, trying to insert ourselves in. The design of Turkopticon, then, had to be as much an expression of the standards and rhythms set by a large,

corporate infrastructure as it was of designer and user values, desires, and politics. By intervening in a working, real world collaborative technological system, we did not enjoy the luxuries of ethics- and values-oriented design projects that design technologies anew.

Publics and their Means of Assembly

A number of researchers have argued that design activities can generate publics – groups that coalesce around identification with a common problem and a shared effort to resolve the problem [14, 30]. Activities such as exploratory prototyping or future-envisioning engage diverse stakeholders in identifying causes of common concern. Design engagement offers one way of collectively inquiring into assumptions, dependencies, and paths forward.

Our early work on Turkopticon – especially the Workers’ Bill of Rights – shared this spirit of engaging workers in imagining alternative ways of doing microlabor. Workers’ responses revealed vastly disparate visions and self-understandings when it came to issues of minimum wage, relations with requesters, and desire for additional forms of support. Moreover, workers distributed across the world faced vastly different circumstances. Indian Turkers, for example, tend to be highly educated and face lower costs of living than Americans. Bringing these workers together as a public to engage in shared inquiry and democratic interchange would require speaking across cultures, ideologies, and vastly different life circumstances. Turkopticon performs an intermediate step in the formation of publics by bringing people together around practical, broadly shared concerns. By creating infrastructures for mutual aid, we bolster the social interchange and interdependency that can become a foundation for a more issue-oriented public. There have been calls in HCI for *representing* interdependence as a way of working towards more ethical and sustainable practices [31]. AMT’s labor market, however, individuates by design; workers are independent by default. Turkopticon provides an infrastructure through which workers can engage in *practices of interdependence*, here as mutual aid.

THE AMBIVALENCE OF SUCCESS IN ACTIVIST TECHNOLOGIES

Turkopticon has succeeded in attracting a growing base of users that sustain it as a platform for an information-sharing community. In part because of its practical embeddedness, it has drawn sustained attention to ethical questions in crowdsourcing over the course of its operation. This attention comes not only in the crowdsourcing community, but also in broader public fora. We have been invited to speak at industry meetups and on Commonwealth Club panels on crowdsourcing. We have also attracted attention from journalists writing pieces on crowdsourcing in venues such as *O’Reilly Radar*, *The Sacramento Bee*, *AlterNet*, and *The San Jose Mercury News*. As a media piece, Turkopticon’s sustained dissent over the last four years has qualities of adversarial design [15]; the system stands as a visible reminder of the microlabors that sustain crowd

platforms. This agonistic reminder disrupts the optimism that surrounds crowdpowered systems.

However, Turkopticon’s existence sustains and legitimizes AMT by helping safeguard its workers. AMT relies on an ecosystem of third party developers to provide functional enhancements to AMT (e.g. CrowdFlower, SamaSource, Twitter). Turkopticon is a squeaky but reliable part of this ecosystem. Ideally, however, we hoped that Amazon would change its systems design to include worker safeguards. This has not happened. Instead, Turkopticon has become a piece of software that workers rely on funded through subsidies from academic research – an unsustainable foundation for such a critical tool.

To stay vital, our team plans on developing new media interventions to give the Turkopticon community greater visibility to the press, to policy makers, and to organizers. Through the design of layered infrastructures, we can support complex and overlapping publics that open up questions about possible futures once again.

This paper has offered an account of an activist systems development intervention into the crowdsourcing system AMT. We argued that AMT is predicated on infrastructuring and hiding human labor, rendering it a reliable computational resource for technologists. Based on a “Workers’ Bill of Rights” meant to evoke workers’ imaginations, we identify hazards of crowdwork and our response as designers to those hazards – Turkopticon. The challenges of developing Turkopticon shows the challenges of developing real-world technologies that intervene in existing, large-scale sociotechnical systems. Such activism takes design out of the studio and into the wild, not only testing the seeds of possible technological futures, but attempting to steer and shift the existing practices and infrastructures of our technological present.

ACKNOWLEDGEMENTS

We dedicate this paper to the memory of Beatriz da Costa, the tactical media artist and professor who pushed us to take the plunge from imagining to building and maintaining. We thank Chris Countryman, Paul Dourish, Gillian Hayes, Lynn Dombrowski, Karen Cheng, Khai Truong, and anonymous reviewers for feedback. This work was supported by NSF Graduate Research Fellowship and NSF award 1025761.

REFERENCES

1. Ackerman, M. 2000. The intellectual challenge of CSCW: The gap between social requirements and technical feasibility. *Human-computer interaction*, 15(2), 179–203.
2. Bardzell, S. 2010. Feminist HCI: Taking Stock and Outlining an Agenda for Design, *Proc. CHI*, 1301–1310.
3. Barley, S.R. and Kunda, G. 2004. *Gurus, hired guns, and warm bodies: itinerant experts in a knowledge economy*. Princeton University Press.

4. Berg, M. 1998. The politics of technology: On bringing social theory into technological design. *ST&HV*, 23(4), 456–490.
5. Bezos, J. 2006. Opening Keynote. *MIT Emerging Technologies Conference*. Accessed: <http://video.mit.edu/watch/opening-keynote-and-keynote-interview-with-jeff-bezos-9197/>
6. Biewald, L. 2009. “Turkopticon.” *Crowdfower Blog*. Accessed: <http://blog.crowdfower.com/2009/02/turkopticon/>
7. Borning, A. and Muller, M. 2012. Next steps for value sensitive design. *Proc. CHI*, 1125–1134.
8. Bowker, G. & Star, S.L. 2000. *Sorting Things Out: Classification and Its Consequences*. MIT Press.
9. Bruckman, A. and Hudson, J.M. 2005. Using Empirical Data to Reason about Internet Research Ethics. *Proc. ECSCW*, 287–306.
10. Brynjarsdottir, H., Håkansson, M., Pierce, J., Baumer, E., DiSalvo, C., and Sengers, P. 2012. Sustainably unpersuaded. *Proc. CHI*, 947–956.
11. Crenshaw, K. 1991. Mapping the Margins: Intersectionality, Identity Politics, and Violence Against Women of Color. *Stanford Law Review*, 43(6), 1241–1299.
12. Crowdsourced data analysis with Clockwork Raven. Accessed at <http://engineering.twitter.com/2012/08/crowdsourced-data-analysis-with.html>
13. da Costa, B. and Philip, K., eds. 2008. *Tactical Biopolitics: Art, Activism, Technoscience*. MIT Press.
14. DiSalvo, C. 2009. Design and the Construction of Publics. *Design Issues*, 25(1), 48–63.
15. DiSalvo, C. 2012. *Adversarial Design*. MIT Press.
16. Dourish, P. 2010. HCI and environmental sustainability: the Design of Politics and Politics of Design. *Proc. DIS 2010*, 1–10.
17. Dourish, P. and Mainwaring, S. 2012. Ubicomp’s Colonial Impulse. *Proc. UbiComp*.
18. Felstiner, A. 2010. Working the Crowd : Employment and Labor Law in the Crowdsourcing Industry. *Berkeley Journal of Employment & Labor Law*, 32 (1), 143–204.
19. Garcia, D. and Lovink, G. 1997. The ABC of tactical media. nettime listserv. Accessed: <http://www.nettime.org/Lists-Archives/nettime-l-9705/msg00096.html>
20. Haraway, D.J. 1990. *Simians, Cyborgs, and Women: The Reinvention of Nature*. Routledge.
21. Hirsch, T. 2009. Learning from activists. *Interactions*, 16(3), 31–33.
22. Ipeirotis, P. 2010. Demographics of Mechanical Turk.. NYU Working Paper No. CEDER-10-01
23. Ipeirotis, P. 2010. Analyzing the Mechanical Turk Marketplace. *XRDS*, 17(2), 16–21.
24. Ipeirotis, P. 2008. Why People Participate in Mechanical Turk? Accessed: <http://www.behind-the-enemy-lines.com/2008/03/why-people-participate-on-mechanical.html>
25. Jackson, S., Pompe, A., and Krieshok, G. 2011. Things fall apart: maintenance, repair, and technology for education initiatives in rural Namibia. *Proc iConference*, 83–90.
26. Keltz, C. 2012. *Two Bits: The Cultural Significance of Free Software*. Duke University Press.
27. Lampe, C. and Resnick, P. 2004. Slash(dot) and burn. *Proc SIGCHI 2004*, 543–550.
28. Law, E. and Von Ahn, L. 2011. *Human Computation*. Morgan Claypool Publishers.
29. Law, J. 2004. *After method: mess in social science research*. Routledge.
30. LeDantec, C. 2012. Participation and Publics: Supporting Community Engagement, *Proc. CHI 2012*, 1351–1360.
31. Light, A. 2011. Digital interdependence and how to design for it. *Interactions*, 18 (2), 34.
32. Orlikowski, W.J. 1992. The Duality of Technology: Rethinking the Concept of Technology in Organizations. *Organization Science*, 3(3), 398–427.
33. Ross, J., Irani, L., Silberman, M.S., et al. 2010. Who are the crowdworkers?, *EA CHI 2010 (alt.chi)*, 2863–2872.
34. Ruhleder, K. and Star, S.L. 2001. Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces. In J. Yates and J. Van Maanen, eds., *Information Technology and Organizational Transformation: History, Rhetoric, and Practice*. Sage, 305–343.
35. Smith, V. 1997. New Forms of Work Organization. *Annual Review of Sociology*, 23, 315–339.
36. Snow, R., O’Connor, B., Jurafsky, D., & Ng, A.Y. 2008. Cheap and fast---but is it good?: Evaluating non-expert annotations for natural language tasks. *Proc. EMNLP*, 254–263.
37. Star, S.L. and Strauss, A. 1999. Layers of Silence, Arenas of Voice: The Ecology of Visible and Invisible Work. *Proc. CSCW*, 8, 9–30.
38. Sood, S., Antin, J., and Churchill, E. 2012. Profanity use in online communities. *Proc. CHI 2012*, 1481–1490.
39. Suchman, L. 1995. Making Work Visible. *CACM*, 38(9).
40. Suchman, L. 2006. *Human Machine Reconfigurations*, Cambridge Univ. Press.
41. Taylor, A. 2011. Out There. *Proc CHI*, 685–694.
42. TurkWork. <http://turkwork.differenceengines.com>