

Fairness and Transparency in Crowdsourcing

Ria Mae Borromeo, Thomas Laurent,
Motomichi Toyama
Keio University, Yokohama, Japan
{riamae,thomas.laurent,toyama}@keio.jp

Sihem Amer-Yahia
Univ. Grenoble Alpes, CNRS, LIG,
F-38000 Grenoble, France
sihem.amer-yahia@imag.fr

ABSTRACT

Despite the success of crowdsourcing, the question of ethics has not yet been addressed in its entirety. Existing efforts have studied *fairness* in worker compensation and in helping requesters detect malevolent workers. In this paper, we propose fairness axioms that generalize existing work and pave the way to studying fairness for task assignment, task completion, and worker compensation. *Transparency* on the other hand, has been addressed with the development of plug-ins and forums to track workers' performance and rate requesters. Similarly to fairness, we define transparency axioms and advocate the need to address it in a holistic manner by providing declarative specifications. We also discuss how fairness and transparency could be enforced and evaluated in a crowdsourcing platform.

Keywords

Crowdsourcing, Fairness, Declarative Transparency

1. INTRODUCTION

The success of crowdsourcing is undeniable. Many tasks ranging from image recognition to sentiment analysis, are routinely deployed and completed by a pool of workers ready to be solicited. It is therefore timely to start addressing *fairness and transparency* in crowdsourcing, two key questions that are of interest today in ethics.¹ Existing work on fairness has primarily focused on studying worker compensation or on helping requesters identify malevolent workers [2, 17, 21, 20, 19]. For transparency, tools and plug-ins have been developed to disclose computed information such as workers' performance and requesters' ratings [3, 6, 9, 15]. In this paper, we argue that a holistic approach to both fairness and transparency is necessary because of the dependencies between crowdsourcing processes. We define fairness and transparency axioms that serve as a basis for our framework and discuss implementation and evaluation.

¹<http://www.fatml.org/>

Our first endeavor is to understand fair crowdsourcing.

Discrimination against individuals is generally defined according to the attributes of those individuals [5]. For example, Google's advertising displays ads for high-income jobs to men much more often than it does to women; and ads for arrest records are most often associated to search queries for common African-American names [18]. In crowdsourcing, even if workers are assigned tasks fairly, the attributes used in task assignment may not have been inferred fairly. It is therefore crucial to characterize fairness in a holistic fashion. We define a set of fairness axioms that capture and generalize existing approaches. For example, we state that in task assignment, two workers with the same qualifications should have access to the same tasks. Similarly, comparable tasks offered by two different requesters should be equally visible to workers. In task completion, fairness to workers means letting them complete tasks without interruption.

Our second question is about transparent crowdsourcing.

Intuitively, a crowdsourcing platform that provides better transparency would generate less frustration among workers and see better worker retention. This realization is not new, and several proposals have addressed transparency in crowdsourcing from requester and platform perspectives. Requester transparency reveals details such as recruitment criteria, the conditions under which work may be rejected, and the time before workers' contributions are approved. Platform transparency, e.g., providing feedback to workers on their performance, has also been addressed [12]. Several tools and forums have been developed to disclose information to workers. For example, Turkopticon [9] provides a plug-in to AMT that helps workers determine which HITs do not pay fairly and which requesters have been reviewed by other workers. Turker Nation² is an online forum that lets workers exchange information about the latest available HITs and their opinion on requesters. CrowdFlower³ displays a panel with the worker's estimated accuracy so far. In this paper, we advocate that a single framework is needed to express and enforce transparency. We believe it is essential to provide declarative languages to help requesters and platform developers express what they want to make transparent. Such a solution would also facilitate sharing and comparing transparency choices across platforms.

The question of *how to validate fairness and transparency* in crowdsourcing also merits attention. A common approach is to design appropriate user studies that gather the experience of workers and requesters with specific implementa-

²<http://www.turkernation.com/>

³<https://www.crowdflower.com/>

tions of fairness and transparency. Such an approach was used in [12] to validate that feedback contributes to increasing workers' motivation. In this paper, we wish to define a validation protocol based on objective measures and propose to quantify measures such as contributions quality for fairness and worker retention for transparency.

We believe that our proposal paves the way for checking fairness and transparency in existing crowdsourcing systems and also for enforcing them by design in newly developed systems. Section 2 contains a review of fairness and transparency in crowdsourcing and other related areas. Section 3 illustrates the need for fairness and transparency using key use cases, and formalizes our proposal. Section 4 discusses validation.

2. RELATED WORK

2.1 Fairness

Fairness in crowdsourcing has mainly been considered in providing fair wages and managing malicious workers. In [2] and [17], wage discrimination is viewed as the wrongful rejection of work, unfair compensation amount, or delayed payment. In [21], a quality-based reward scheme provides compensation that depends on the quality of a worker's contribution. Vuurens et al. proposed measures to detect and counter malicious users since they observed that nearly 40% of the answers they received from AMT were from malicious users [20].

Studies that address malicious workers through task assignment and worker reputation focus on the quality, reliability, and total cost of worker contributions. Examples of existing task assignment schemes include offering low-cost, reliable answers [7, 11], and accounting for worker skills to maximize the requester's total gain from the completed work [8]. These schemes are requester-centric and do not guarantee fair task assignment to workers.

Overall, we observed that while some work have developed ways of enforcing fair wages and helping requesters detect malicious workers, no holistic approach has been developed to address fairness as a whole for all crowdsourcing processes.

2.2 Transparency

Bederson et al. claimed that higher transparency in working conditions such as hourly wage, or in requester expectations such as work quality metrics, lead to fairness [2]. They asserted the need for requester and platform transparencies to address discrimination but did not tackle its systematic implementation.

Requester transparency has been shown to have positive effects on worker engagement. Studies show that providing workers with information about the requester leads to higher engagement and more effort in task completion [16]. Moreover, providing workers with information about the crowdsourcing workflow and helping them feel part of a group, result in more contributions and higher accuracy [13].

Different initiatives implement transparency in crowdsourcing platforms through plug-ins. Turkocticon [9] is a plugin for AMT that lets workers review tasks and requesters. Crowd-Workers [3] and Turkbench [6] provide expected hourly wages when workers browse tasks. The MobileWorks platform [15] facilitates worker-to-worker communication and assigns manager roles to some workers, allowing workers to

monitor each other and benefit from each other's experience, which results in higher quality contributions.

Transparency for workers comes from worker initiatives and communities mainly through forums such as Turker Nation and Mturk Forum⁴ where workers share information about tasks, requesters and tools to enhance their experience. These tools are often worker-made scripts that disclose information hidden by the platform such as the time until automatic approval of a submission on AMT.

In summary, workers strive for transparency as it is often fragmented and external to platforms. In this paper, we advocate a systematic way of expressing and enforcing transparency in crowdsourcing platforms through a formalization of fairness and transparency axioms.

3. PROPOSAL

In this section, we first discuss scenarios where discrimination and opacity can hinder workers' and requesters' experience. We then formalize our framework and discuss how we can enforce fairness and transparency in crowdsourcing systems.

3.1 Scenarios

3.1.1 Discrimination

In Task Assignment. Task assignment, the process through which workers find tasks to complete, is central to crowdsourcing. In platforms such as AMT and CrowdFlower, requesters post tasks, and qualified workers choose the ones they like. This simple task assignment mechanism could be characterized as fair because workers have access to the same set of tasks.

Aside from self-appointment, many task assignment algorithms have been designed to optimize a particular objective. However, these algorithms can be discriminatory [14]. For instance, requester-centric task assignment allocates tasks to workers so as to maximize the total gain of the requester. This could be discriminatory to workers. On the other hand, a worker-centric assignment that allocates tasks based on workers' preferences is more likely to be fair to workers, by favoring their expected compensation, but may be unfavorable to requesters.

In Task Completion. In task completion, workers and requesters have different goals. Requesters aim to get enough good results while workers' objectives range from getting paid to improving their skills, spending their time wisely, or signaling their presence and achievements to others [12]. These goals may be advantageous to one but unfair to the other.

For example, in survey tasks, requesters usually publish more HITs than necessary to get a good number of responses. There are cases when a requester cancels tasks when she gets the target number of acceptable responses. Requesters do so to reduce their waiting time and avoid paying more than needed. However, this would be unfair to a worker who has partially completed a task but is not paid for her efforts. A requester may also experience discrimination during task completion in the case of malevolent workers.

⁴<http://www.mturkforum.com/>

In Worker Compensation. Discriminatory compensation has been identified as one of the major problems for crowd workers [2, 17]. For instance, in AMT, a requester may reject valid work and not pay the worker. In some other cases, a requester promises to provide a bonus when a worker completes a series of tasks but does not do so in the end. In collaborative tasks, a worker may contribute more than another and still receive the same amount of payment.

3.1.2 Opacity

Requester Opacity. In Turk Nation, workers often complain about requesters who reject their contribution without providing feedback. For example, a requester who posts a text summarization task may not publish how a worker’s contribution will be evaluated. This requester opacity does not only negatively affect workers’ experiences but also affects other crowdsourcing processes. If a contribution is rejected, it is reflected in the worker’s history and statistics thus it may limit future task assignment opportunities. If a worker is provided means to post a review of a requester, this may encourage requesters to be more transparent.

Platform Opacity. Since a platform facilitates the entire crowdsourcing process, it must provide valuable information to help requesters and workers achieve their goals [2]. For requesters, it is important to see worker statistics and progress to help them monitor tasks. For workers, it is beneficial to have access to various information that could help them select and complete tasks such as requester reviews and ratings, payment schedules, and estimated worker performance in comparison with other workers. Currently, CrowdFlower shows ratings per task in its task browsing interface and the Turkopticon plug-in shows requester ratings in AMT [9]. Nevertheless, there is currently no systematic way for platform developers and for requesters to specify which information should be made transparent.

3.2 Fairness and Transparency Axioms

In this section, we attempt to define and formalize fairness and transparency axioms. Our proposal does not aim to be exhaustive. Rather, it provides a framework to define and extend a series of axioms that govern checking if a crowdsourcing system abides by fairness and transparency goals, in a principled fashion, or for designing a fair and transparent platform from scratch.

We consider a set of tasks $\mathcal{T} = \{t_1, \dots, t_n\}$, a set of workers $\mathcal{W} = \{w_1, \dots, w_p\}$ and a set of skill keywords $\mathcal{S} = \{s_1, \dots, s_m\}$.

Tasks. A task t is a tuple (id_t, id_r, S_t, d_t) where id_t is a unique task identifier, id_r a unique requester identifier, and S_t is a vector $\langle t(s_1), t(s_2), \dots, t(s_m) \rangle$ where each $t(s_j)$ is a Boolean value that denotes the requirement or not of having skill s_j to qualify for task t . A reward d_t is given to a worker who completes t . To capture a variety of tasks, skill keywords may be interpreted as expected workers’ interests or qualifications.

Workers. A worker w is a tuple (id_w, A_w, C_w, S_w) where id_w is the worker’s unique id, A_w is a set of self-declared worker attributes such as demographics and location, C_w is a set of computed worker attributes such as a worker’s

acceptance ratio, and S_w is a skill vector $\langle w(s_1), \dots, w(s_m) \rangle$ where each $w(s_j)$ is a Boolean value capturing the interest of w in the skill keyword s_j .

3.2.1 Fairness

We define fairness axioms for task assignment, worker compensation and task completion.

AXIOM 1 (WORKER FAIRNESS IN TASK ASSIGNMENT).
Given two different workers w_i and w_j , if A_{w_i} is similar to A_{w_j} and C_{w_i} is similar to C_{w_j} , and S_{w_i} is similar to S_{w_j} , then w_i and w_j should have access to the same tasks.

Similarity can be platform-dependent and ranges from perfect equality to threshold-based similarity.

AXIOM 2 (REQUESTER FAIRNESS IN TASK ASSIGNMENT).
Given two tasks t_i and t_j posted by different requesters id_{r_i} and id_{r_j} , if the required skills for the two tasks S_{t_i} and S_{t_j} are similar, and the two tasks offer comparable rewards d_{t_i} and d_{t_j} , then t_i and t_j should be shown to the same set of workers.

Skill similarity can be computed using different measures such as cosine similarity.

AXIOM 3 (FAIRNESS IN WORKER COMPENSATION).
Given two distinct workers w_i and w_j who contributed to the same task t , if their contributions are similar, they should receive the same reward d_t .

Different measures could be used to compute similarity of contributions depending on the nature of those contributions, e.g., for textual contributions, n-grams could be used [4], for ranked lists, using measures such as Discounted Cumulative Gain [10] would be more appropriate.

AXIOM 4 (REQUESTER FAIRNESS IN TASK COMPLETION).
Requesters must be able to detect workers behaving maliciously during task completion.

AXIOM 5 (WORKER FAIRNESS IN TASK COMPLETION).
A worker who started completing a task should not be interrupted.

3.2.2 Transparency

It is believed that limiting worker and requester anonymity may lead to fairer labor practices [2]. Therefore, letting workers and requesters reveal information about themselves, ranging from their true identity, to historical worker performance, for example, may help raise everyone’s trust in the platform. Transparency axioms govern what requesters and platforms should make available to workers in order to ensure their fair treatment.

AXIOM 6 (REQUESTER TRANSPARENCY).
A Requester must make available requester-dependent working conditions such as hourly wage and time between submission of work and payment, and task-dependent working conditions such as recruitment criteria and rejection criteria.

AXIOM 7 (PLATFORM TRANSPARENCY).
The platform must disclose, for each worker w , computed attributes C_w such as performance and acceptance ratio.

3.3 Implementation

We discuss our preliminary thoughts on how fairness and transparency could be implemented and enforced.

3.3.1 Fairness

Our axioms form a framework to check how fair an existing crowdsourcing system is and also develop guidelines for designing fair crowdsourcing processes from scratch. Using the axioms discussed in Section 3.2, we intend to develop *fairness check* benchmarks and algorithms for existing crowdsourcing systems.

Particular attention needs to be given to checking fairness due to the inter-dependencies between crowdsourcing processes. For instance, an algorithm that checks worker fairness in task assignment must check the fairness of derived computed attributes such as worker’s performance.

3.3.2 Transparency

We advocate the use of a declarative high-level language to specify fairness rules. Such rules can be used by requesters to disclose task requirements, recruitment criteria, evaluation scheme, and payment schedule. Platform designers can use these rules to disclose relevant information that they want to show both workers and requesters. Rules can also be translated into human-readable descriptions for workers’ consumption. Last but not least, the declarative nature of those rules will allow easy comparison across platforms. Guiding principles for such a language can be found in works on privacy policy declaration such as in [1].

4. DISCUSSION

4.1 Evaluation

When measuring fairness and transparency, objective measures such as quality of worker contribution and worker retention, can be used in controlled experiments to quantify the level of fairness and transparency of a system as well as its effectiveness.

4.2 Research agenda

Regarding fairness, our immediate agenda is to review existing algorithms for task assignment, strategies for worker compensation, and approaches for task completion, to assess their discriminatory power.

Regarding transparency, we plan to run a user study to validate what kind of transparency choices workers are most sensitive to. Meanwhile, we started designing a declarative language in which transparency rules can be expressed.

5. REFERENCES

- [1] M. Y. Becker, A. Malkis, and L. Bussard. A practical generic privacy language. In *ICISS*, pages 125–139. Springer, 2010.
- [2] B. B. Bederson and A. J. Quinn. Web workers unite! addressing challenges of online laborers. In *ACM CHI ’11 Extended Abstracts on Human Factors in Computing Systems*, New York, NY, USA, 2011.
- [3] C. Callison-Burch. Crowd-workers: Aggregating information across turkers to help them find higher paying work. In *Second AAAI*, 2014.
- [4] M. Damashek. Gauging similarity with n-grams: Language-independent categorization of text. *Science*, 267(5199):843, 1995.
- [5] S. Hajian, F. Bonchi, and C. Castillo. Algorithmic bias: From discrimination discovery to fairness-aware data mining. In *Proceedings of ACM SIGKDD*, pages 2125–2126, 2016.
- [6] B. V. Hanrahan, J. K. Willamowski, S. Swaminathan, and D. B. Martin. Turkbench: Rendering the market for turkers. In *Proc. CHI*, pages 1613–1616. ACM, 2015.
- [7] C.-J. Ho, S. Jabbari, and J. W. Vaughan. Adaptive task assignment for crowdsourced classification. In *Proc. ICML 2013*, pages 534–542, 2013.
- [8] C.-J. Ho and J. W. Vaughan. Online task assignment in crowdsourcing markets. In *AAAI*, volume 12, pages 45–51, 2012.
- [9] L. C. Irani and M. Silberman. Turkopticon: interrupting worker invisibility in amazon mechanical turk. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 611–620. ACM, 2013.
- [10] K. Järvelin and J. Kekäläinen. Cumulated gain-based evaluation of ir techniques. *ACM Transactions on Information Systems (TOIS)*, 20(4):422–446, 2002.
- [11] D. R. Karger, S. Oh, and D. Shah. Budget-optimal task allocation for reliable crowdsourcing systems. *Operations Research*, 62(1):1–24, 2014.
- [12] N. Kaufmann, T. Schulze, and D. Veit. More than fun and money. worker motivation in crowdsourcing-a study on mechanical turk. In *AMCIS*, volume 11, pages 1–11, 2011.
- [13] P. Kinnaird, L. Dabbish, and S. Kiesler. Workflow transparency in a microtask marketplace. In *Proceedings of the 17th ACM international conference on Supporting group work*, pages 281–284. ACM, 2012.
- [14] K. Kirkpatrick. Battling algorithmic bias: how do we ensure algorithms treat us fairly? *Communications of the ACM*, 59(10):16–17, 2016.
- [15] A. Kulkarni, P. Gutheim, P. Narula, D. Rolnitzky, T. Parikh, and B. Hartmann. Mobileworks: Designing for quality in a managed crowdsourcing architecture. *IEEE Internet Computing*, 16(5):28–35, 2012.
- [16] J. Marlow and L. A. Dabbish. Who’s the boss?: requester transparency and motivation in a microtask marketplace. In *CHI’14 Extended Abstracts on Human Factors in Computing Systems*, pages 2533–2538. ACM, 2014.
- [17] B. McInnis, D. Cosley, C. Nam, and G. Leshed. Taking a hit: Designing around rejection, mistrust, risk, and workers’ experiences in amazon mechanical turk. In *Proc. CHI*, pages 2271–2282. ACM, 2016.
- [18] L. Sweeney. Discrimination in online ad delivery. *Commun. ACM*, 56(5):44–54, May 2013.
- [19] B. Valeri, S. Elbassuoni, and S. Amer-Yahia. Crowdsourcing reliable ratings for underexposed items. In *Proceedings of WEBIST 2016*, pages 75–86, 2016.
- [20] J. Vuurens, A. P. de Vries, and C. Eickhoff. How much spam can you take? an analysis of crowdsourcing results to increase accuracy. In *Proc. ACM SIGIR CIR’11*, pages 21–26, 2011.
- [21] J. Wang, P. G. Ipeirotis, and F. Provost. Quality-based pricing for crowdsourced workers. *NYU Working Paper*, 2013.