

An Endorsement-based Reputation System for Trustworthy Crowdsourcing

Chunchun Wu*, Tie Luo†, Fan Wu*, Guihai Chen*

*Department of Computer Science and Engineering, Shanghai Jiao Tong University, China

†Institute for Infocomm Research, A*STAR, Singapore

Email: bubble_chun@sjtu.edu.cn, luot@i2r.a-star.edu.sg, {fwu, gchen}@cs.sjtu.edu.cn

Abstract—Crowdsourcing is a new distributed computing paradigm that leverages the wisdom of crowd and the voluntary human effort to solve problems or collect data. In this paradigm of soliciting user contributions, the trustworthiness of contributions becomes a matter of crucial importance to the viability of crowdsourcing. Prior mechanisms either do not consider the trustworthiness of contributions or assess the quality of contributions only after the event, resulting in irreversible effort exertion and distorted player utilities. In this paper, we propose a reputation system to not only assess but also predict the trustworthiness of user contributions. In particular, we explore an inter-worker relationship called endorsement to improve trustworthiness prediction using machine learning methods, while taking into account the heterogeneity of both workers and tasks.

I. INTRODUCTION

Crowdsourcing is a new problem-solving model that solicits solutions to various tasks from the crowd, particularly online labor markets. The diversity and pervasiveness of crowd have brought substantial benefits but created some serious challenges as well. Two crucial issues are incentive and trustworthiness. The former refers to motivating users to continuously participate in tasks and the latter refers to ensuring solution or data quality since the workers are diverse in abilities and may even be malicious. There is a sizeable body of prior work dedicated to addressing these two issues. However, the heterogeneity of both tasks and workers are largely overlooked; furthermore, in prior work a task requester has to assess the quality of contributions [1] after workers submit them, which incurs irreversible expenses of effort for workers, and distorted, sometimes even negative, utilities for both requester and workers.

In this paper, we propose a reputation system to assess and predict the trustworthiness of contributions without ex-ante expenditure. We use reputation to motivate participation because it is easier to deploy in practice than direct monetary rewards. The novelty of our work lies in the following. First, we identify and exploit an inter-worker relationship called endorsement [2] to improve trustworthiness prediction as well as assessment. Endorsement casts a social-network perspective on workers in crowdsourcing rather than treating them individually. To the best of our knowledge, we are the first to make use of connections among workers to design a trustworthy crowdsourcing system. Second, to avoid wasting irreversible worker effort and distorted or negative player utilities, our mechanism sets out to predict the trustworthiness of contributions rather than requiring workers to pay effort upfront like in all-pay auctions [3]. The predicting method is based on machine learning approaches and takes into account the heterogeneity of both workers and tasks. Finally, we allow the weight of endorsements to vary and model the variation by resembling it to the evolution of human relationship in general.

II. SYSTEM OVERVIEW

An illustrative crowdsourcing platform is depicted in Fig. 1. A diverse population of workers participate in various types of tasks which may require different domain-specific knowledge or expertise. Workers are connected to one another via endorsement links which represent a trust or support relationship. Every endorsement link has a weight, indicating how much confidence the endorser has in the endorsee. The worker population can thus be represented by a directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{L}, \mathcal{D})$, where \mathcal{N} is the set of workers, \mathcal{L} is the set of endorsement links, and \mathcal{D} is the set of weights. A weight $d_{ij} \in \mathcal{D}$ represents the degree of endorsement for any $i, j \in \mathcal{N}$.

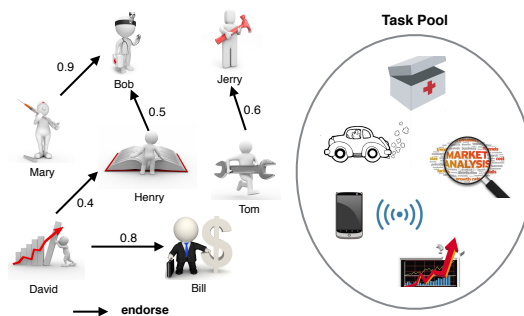


Fig. 1. An illustrative crowdsourcing platform with endorsement links

The workflow of a crowdsourcing system that uses our endorsement-based reputation system is outlined in Algorithm 1. When a requester posts a task with a description, workers who are interested express their willingness to join in by signing up. Lines 6-10 predict the trustworthiness of contributions (ToC) for all the applicants, where two determining factors are each applicant's own reputation and the endorsement impact from his endorsers (Line 10). In calculating the endorsement impact, the heterogeneity of tasks and domain-specific expertise of workers, which are overlooked by the vast literature, play an important role.

Line 8 measures a worker's expertise. Although an intuitive method of assessing one's expertise is to observe his past performance in this particular task, such historical information may not be available. Therefore, we take a machine learning approach called *matrix factorization* [4], which is so far the most superior *collaborative filtering* [5] methodology. To do this, the system maintains a rating matrix with one dimension representing workers and the other dimension representing tasks. An entry r_{ik} in the matrix indicates the performance rating of worker i in task k given by the requester after i performs task k . Naturally, not all r_{ik} have values.

Matrix factorization maps both workers and tasks to a joint latent factor space, where each worker i is characterized by a vector \vec{p}_i and each task k is characterized by a vector

Algorithm 1: Workflow of the Endorsement-based Reputation System

Input: The directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{L}, \mathcal{D})$.

- 1 Initialize reputation scores RS_i for every worker;
- 2 **while** *The platform is in service* **do**
- 3 A requester posts a task k ;
- 4 $m \leftarrow$ number of workers needed;
- 5 a set of workers W_k sign up;
- 6 **foreach** $i \in W_k$ **do**
- 7 **foreach** j who endorses i **do**
- 8 Predict e_{jk} as j 's expertise in task k , according to (2);
- 9 **end**
- 10 $ToC_{ik} = RS_i + \sum_{j \in N_i^n} RS_j \cdot d_{ji} \cdot e_{jk}$;
- 11 **end**
- 12 Sort workers in non-descending order of ToC_{ik} :
- 13 $ToC_{1k} \geq ToC_{2k} \geq \dots \geq ToC_{|W_k|k}$;
- 14 $n \leftarrow \min(m, |W_k|)$;
- 15 The requester selects first n workers in W_k ;
- 16 **foreach** i who just participated in task k **do**
- 17 Requester gives rating r_{ik} ;
- 18 $RS_i = \frac{\sum_{k \in CoT_i} r_{ik}}{|CoT_i|}$;
- 19 **end**
- 20 Update \mathcal{D} ;
- 21 **end**

\vec{q}_k . Both vectors correspond to the same set of features that describe a task on the crowdsourcing platform. For example, a question-and-answer forum may describe a task (i.e., a question) according to its relevance to fashion, healthcare, education, and computers, which constitutes \vec{q}_k . The vector \vec{p}_i then describe worker i in terms of his performance with respect to the same features. These latent features (as contained in \vec{p}_i and \vec{q}_k) are not manually assigned but automatically inferred from workers' past performance, by minimizing the *regularized squared error* on the set of known ratings κ :

$$\operatorname{argmin}_{\vec{q}_*, \vec{p}_*} \sum_{(i,k) \in \kappa} (r_{ik} - \vec{q}_k^T \vec{p}_i)^2 + \lambda (||\vec{q}_k||^2 + ||\vec{p}_i||^2), \quad (1)$$

where the dot product $\vec{q}_k^T \vec{p}_i$ approximates worker i 's performance in task k , and the regularization term $\lambda (||\vec{q}_k||^2 + ||\vec{p}_i||^2)$ is to restrict the magnitudes of \vec{q}_k and \vec{p}_i in order to prevent overfitting.

There are several approaches to solving the optimization problem (1), among which two most commonly adopted ones are *stochastic gradient descent* and *alternating least squares*. After solving for \vec{p}_i and \vec{q}_k , we can determine each endorser j 's expertise in task k , for all j, k :

$$e_{jk} = \begin{cases} r_{jk} & \text{If } (j, k) \in \kappa, \\ \vec{q}_k^T \vec{p}_j & \text{Otherwise.} \end{cases} \quad (2)$$

Then, Line 10 combines a worker's own reputation and his endorsers influence, which demonstrates how the endorsement impact is incorporated when predicting the trustworthiness of contribution for a worker on a task, before he actually does it.

At Lines 12–14, the requester selects a subset of workers who are the top n workers in terms of predicted ToC, to perform the task. After the task is completed, the requester

shall give each worker i a rating r_{ik} , which will then cause the reputation score of each worker to be updated as per Line 17, where CoT_i is the collection of tasks performed by i so far.

In the meantime, the endorsers of each worker will also have their degree of endorsement adjusted, as at Line 19. To do this, we draw an analogy to human relationship which generally evolves in the following manner: it develops slowly between strangers, and then slightly accelerates and makes steady progress between acquaintances, and after a sufficiently long period of time, converges to intimate friendship. Thus, we use the *generalized logistic function* to model this behavior. This function corresponds to a flexible S-shape curve, which captures the progression from slow start to acceleration and eventually approaching a climax:

$$D_{ji}(R_i) = A_{ji} + \frac{B_{ji} - A_{ji}}{(1 + Qe^{-K(R_i - M)})^{1/v}} \quad (3)$$

In the above, $D_{ji} \in [0, 1]$ represents the degree of endorsement from worker j to i , R_i is worker i 's overall performance by aggregating historical information:

$$R_i = \sum_{k \in CoT_i} \alpha^{t_0 - t_k} \cdot (r_{ik} - \frac{\sum_{j \in CoW_k} r_{jk}}{|CoW_k|}) \quad (4)$$

where $\alpha^{t_0 - t_k}$ ($\alpha < 1$, t_0 being current time) is a discounting factor that prioritizes recent information over the past, CoW_k is the collection of workers who performed task k . In (3), A_{ji} and B_{ji} are the lower and upper asymptotes of D_{ji} , respectively, which are determined by

$$A_{ji} = \max\{0, \bar{d}_{ji} - c_{ji}\} \quad (5)$$

$$B_{ji} = \min\{1, \bar{d}_{ji} + c_{ji}\} \quad (6)$$

where \bar{d}_{ji} is the initial degree of endorsement (trust or support) assigned when the endorsement link l_{ji} was established, c_{ji} is the tolerance level indicating to what (maximum) extent the endorser would allow the initial trust to change.

III. CONCLUSION

In this paper, we leverage endorsement relationships between workers to design a reputation system for trustworthy crowdsourcing, taking into account the heterogeneity of workers and tasks. We predict the trustworthiness of contributions via collaborative filtering methods, which avoids workers' irreversible efforts and distorted or negative player utilities. Our mechanism is also a decent solution to the cold-start problem for reputation-based communities, so that a promising new worker will not be buried in the crowd although he hasn't built his own reputation yet.

REFERENCES

- [1] C.-K. Tham and T. Luo, "Quality of contributed service and market equilibrium for participatory sensing," in *IEEE DCOSS*, 2013, pp. 133–140.
- [2] T. Luo, S. S. Kanhere, and H.-P. Tan, "SEW-ing a simple endorsement web to incentivize trustworthy participatory sensing," in *IEEE SECON*, 2014, pp. 636–644.
- [3] T. Luo, H.-P. Tan, and L. Xia, "Profit-maximizing incentive for participatory sensing," in *IEEE INFOCOM*, 2014, pp. 127–135.
- [4] Y. Koren, R. Bell, and C. Volinsky, "Matrix factorization techniques for recommender systems," *Computer*, no. 8, pp. 30–37, 2009.
- [5] X. Su and T. M. Khoshgoftaar, "A survey of collaborative filtering techniques," *Advances in artificial intelligence*, vol. 2009, p. 4, 2009.