A Survey of Ethics in Resource Allocation and Crowdsourcing

Zichao Deng, Jiehuang Zhang, and Han Yu

Joint NTU-UBC Research Centre of Excellence in Active Living for the Elderly,
Nanyang Technological University, Singapore

han.yu@ntu.edu.sg

Abstract

As the forth industrial Revolution is underway, crowdsourcing is increasingly being used to solve real-world problems. Crowdsourcing, in which human intelligence and productivity are dynamically mobilized to tackle tasks too complex for automation, has grown to be an important research topic under artificial intelligence (AI). Over the years, the procedure of crowdsourcing has evolved from simple match-making between workers and tasks into sophisticated data-driven algorithmic management approaches powered by AI. However, there are many ethical considerations that are relatively new to researchers in this field of crowdsourcing. By reviewing full research papers published in major AI conferences including AAAI, AAMAS, HCOMP and IJCAI, we divide this survey paper into 2 major areas with their respective subsections: 1) Crowdsourcing and 2) Ethics. We highlight not only the research challenges in each of these areas, but also the promising future research directions.

Keywords: Ethics, Resource Allocation, Crowdsourcing

I. Introduction

Crowdsourcing

Crowdsourcing is an interdisciplinary technique offering efficient ways to mobilize the effort, time and expertise of human participants to perform tasks which are relatively easy for humans, but difficult for automation [1]. In a typical crowdsourcing system, crowdsourcers propose tasks for workers to complete and offer rewards in exchange for their effort. Over the years, it has been applied in scientific research [2], environment sensing [3] and government service provision [4], etc.

The term ‘algorithmic crowdsourcing’ aims to combine the intelligence of humans and the computing power of machines to solve problems that are difficult to be solved by either humans or machines alone. It makes use of AI to proactively motivate, organize and manage the crowd so as to achieve more efficient utilization of human resources.

Ethics

The issue of ethics has been a challenge for humanity since a long time ago and it has been increasingly so as automation technology becomes ubiquitous. As crowdsourcing is being used to solve problems in fields such as ride sharing algorithms [5][6] and multi-agent systems [2][7][8], they are increasingly widespread experienced by the public. Frequently appearing in popular culture such as movies and television shows, public fascination with AI and crowdsourcing tend to lead the discussions into ethical considerations involving automation of decision making by machines. A topic that has captured public attention is the Artificial General Intelligence (AGI) [9] research, whose objective is to develop AI with capabilities matching and ultimately surpassing human capabilities. However, these fears are unfounded and merely speculation, as AGI is still a distant dream that will take decades or more to materialize. On the other hand, these fears spur us towards taking steps to discuss and define the standards that machine decision making should uphold.

Ethics is a normative practical philosophical discipline of how one should act towards others according to [10], and can be separated into 3 distinct dimensions:

1) Consequentialist ethics: an agent is ethical if and only if it weighs the consequences of each choice and chooses the option which has the most moral outcomes. It is also known as utilitarian ethics as the resulting decisions often aim to produce the best aggregate consequences.

2) Deontological ethics: an agent is ethical if and only if it respects obligations, duties and rights related to given situations. Agents with deontological ethics (also known as duty ethics or obligation ethics) act in accordance to established social norms.
3) Virtue ethics: an agent is ethical if and only if it acts and thinks according to some moral values (e.g. bravery, justice, etc.). Agents with virtue ethics should exhibit an inner drive to be perceived favourably by others.

The research community is perceived to be moving towards an ethically aligned direction, from a standpoint of maximum results previously. For example, in ride-sharing algorithms efforts have been introduced to reduce driver fatigue, and safety considerations. Ethical dilemmas indicate situations where despite any available option or course of action leads to a violation of an accepted ethical principle or rule, the decision still must be made. In the case of decision making by machines and autonomous systems, they must respect the rights of humans and only allow decisions that are acceptable according to the above mentioned ethical standards. Currently the topic of ethics is relatively new to the research community, and technical solutions to implement ethics in machine decision making has not been explored much. In fact, the most recent survey on technical approaches for ethical AI decision making was conducted in 2006 [11].

In this paper, we review published full research papers from major AI conferences including AAAI, AAMAS, HCOMP and IJCAI on the topic of algorithmic crowdsourcing and ethics. Firstly, we propose a taxonomy which divides algorithmic crowdsourcing into 3 major areas: 1) Motivating workers, 2) Task allocation, and 3) Quality control, and highlight the intuitions and insights behind representative works. Then, we explore 4 main areas in ethics [12]: 1) Exploring ethical dilemmas, 2) Individual ethical decision frameworks, 3) Collective ethical decision frameworks, and 4) Ethics in human-AI interactions, and discuss how they are being used to incorporate ethics into crowdsourcing problems.

II. Motivating Workers

There are two broad categories of approaches proposed to motivate crowdsourcing workers: 1) Volunteer-based approaches and 2) Payment-based approaches.

Volunteer-based Approaches

Volunteer-based crowdsourcing mainly depends on engaging workers and maintaining such engagement over time. In [13], the authors proposed a methodology to enhance workers' engagement via a combination of machine learning and intervention design. The methodology consists of a platform for using real-time predictions about future disengagement to guide interventions. Different messages based on the proximity to the predicted time of disengagement were delivered to workers to study their impact. The result indicates that messages highlighting workers' helpfulness, when delivered on the predicted time of disengagement, significantly increased their contributions.

Some approaches leverage the nature of crowdsourcing tasks to design intrinsic incentives to motivate workers' participation. In [14], the authors designed a probabilistic maximum coverage based optimization approach to improve the possibility for workers engaged in birding watching to encounter various bird species based on their given locations. This approach enhances workers' sense of achievement and retains their interest in the task.

Psychological principles have also been applied to enhance workers' engagement. In [15], the authors designed three methods to improve workers' engagement on tasks for which ground truth information is not available based on the psychological theory of commitment. Workers commit to tasks by 1) signing a contract, 2) listening to a recording, and 3) recording a personal commitment. It was found that method 2 and 3 are more effective in improving workers' engagement.

Payment-based Approaches

In order to design incentive mechanisms that can motivate and retain workers, it is important to understand the dynamics of workers' responses to various incentives in practice. In [16], the authors found that more financial incentives can improve the quality of task results by attracting a larger pool of workers, but do not always translate into more effort output by workers.

In [17], the authors studied four incentive structures on social network referral behaviors: 1) No extra incentive, 2) A small fixed number of points for a referral, 3) A large fixed number of points for a referral, and 4) Multi-level rewards for direct and indirect referrals. The results indicate only mechanism 3 significantly increased referral behaviors and sign-ups in crowdfunding. In [18], the authors studied the problem of motivating workers in incentive networks. In such networks, a worker’s reward depends not only on his own contributions to the tasks, but also in part on the contributions made by his friends. Therefore, the efficiency of reward structures is significantly effected by workers’ degree of direct altruism in incentive networks. Incentive network-based rewards are more advantageous than contributions-based rewards.

In addition to fixed rewards, some studies suggest that bonuses can be used to reward better than expected performance. In [19], the authors studied various incentive designs on retaining workers including fixed bonuses, training bonuses, increasing bonuses, milestone bonuses, using qualifications as an exclusion factor, and random bonuses. The results indicate that prompt bonuses given to workers who reached a predefined object improve workers’ retention. Overpaying workers at the beginning of tasks which require training also improves workers’ retention and reduces task latency. In [20], the authors proposed a Hidden Markov Model-based approach to determine whether to reward workers with bonuses in a working session to maximize the crowdsourcer’s expected utility. The timing of distributing bonuses is the key parameter here.

Incentives can also encourage workers to focus on high valued tasks. In [21], the authors proposed an incentive mechanism for the generalized setting of an unknown set of workers with non-deterministic availabilities and stochastically rational reporting behavior. This mechanism can produce optimal stochastic solutions on payment distribution and direct workers’ effort towards more important tasks.

Some approaches not only encourage good quality workers to participate, but also discourage workers with poor performance from participating. In [22], the authors proposed a boosting scheme that makes a distinction between incentives for high and low quality work to prevent random answers.

III. Task Allocation

The task allocation in algorithmic crowdsourcing generally balance the considerations among three entities: 1) Crowdsourcers, whose aim is to maximize the quality of tasks results and minimize the latency time subject to a limited budget; 2) Workers who are heterogeneous in terms of availability, motivation to work, skill level, and productivity, whose aim is to maximize the rewards they receive after completing tasks. Workers may be strategic in terms of how to price their effort and how to spend their effort; and 3) Tasks which are associated with deadlines before which results must be received from workers. Complex tasks also consider skill sets from workers.

Task allocation approaches consider the workload balancing among workers as well as their availability, skill level, and productivity to make quality-time-cost trade-offs. They can be divided into three categories: 1) Allocating simple tasks which can be completed by one worker; 2) Allocating complex tasks which require collaboration by workers with complementary skills; and 3) Optimizing workflows comprising many simple tasks.

Allocating Simple Tasks

A series of research works concentrate on the simple setting of allocating homogeneous micro-tasks. In [23], the authors proposed a decision theoretic approach which leverages workers' profile information to optimize task allocation among them to improve the result quality. In [24], the authors proposed an approach which combines the decision theoretic approach and game theory to reach the expected result quality subject to a limited budget in allocating contest-based crowdsourcing tasks to workers. In [25], the authors proposed a decision theoretic approach to effectively organize workers with various capabilities to balance the workload distribution among them.

Queueing systems theory can also be applied in allocating simple tasks. In [26][7], the authors treated workers as first-come-first-served task queues and proposed centralized and distributed Lyapunov optimization-based approaches which direct each worker on how many new tasks to accept at any given time to maximize time-averaged system level expected task result quality and minimize the latency time.

The workers' unknown qualities and costs make task allocation a non-trivial problem in most crowdsourcing applications. In [27], the authors proposed a multi-armed bandit approach to strategically select a subset of workers with unknown qualities in order to achieve the expected accuracy for binary labelling tasks. In [28], the authors further assumed that workers are strategic about their costs, and proposed an interval cover mechanism which leverages the nature linear ordering structure in crowdsourcing applications to compute optimal task allocation plans in polynomial time.

In [29], the authors proposed a dynamic price mechanism and an auction-based mechanism by considering the constraint that unreliable and strategic workers arrive into the system over time. In [30], the authors proposed a heuristic algorithm to determine how to sequence mobile/spatial crowdsourcing tasks by considering workers' current movement trajectories in order to minimize interrupts to their movement patterns.

A few works concentrated on allocating heterogeneous tasks which require a set of unique skills. In [31], the authors constructed a bipartite graph to solve the problem of matching workers with different skills and interests with heterogeneous tasks. They proposed a task allocation mechanism to achieve budget feasibility, incentive-compatibility, and near-optimal utility. In [32], the authors proposed an online task allocation algorithm from the perspective of crowdsourcers who have heterogeneous tasks to allocate. Workers are assumed to arrive one-by-one and each declares tasks he can solve and desired payments. This algorithm was proven to have an upper bound over an arbitrary sequence of workers who declare small payments relative to the crowdsourcer's total budget.

Allocating Complex Tasks

The problem of allocating complex tasks which require workers with different skills to collaborate remains challenging. In [33], this problem was proven to be NP-hard. They proposed a set of approximation algorithms which attempt to allocate tasks to the most appropriate workers.

In [34], the authors proposed a Lyapunov optimization-based approach for dynamically forming teams of workers with appropriate combinations of skills to tackle complex tasks. This approach makes efficient quality-time-cost trade-offs to maximize the expected results and minimize the latency time with a limited budget. This approach can achieve close to optimal collective profit if workers adhere to it.

Optimizing Workflows

Hierarchical crowdsourcing in which tasks pass through different workers in a network of trust naturally forms dynamic workflows. In [35], the authors proposed a game theoretic approach to determine how to distribute workers' effort between concentrating on tasks and finding other workers to pass tasks in order to improve overall efficiency. In [36], the authors focused on hierarchical consensus tasks which look for correct answers to a hierarchy of subtasks, where branching depends on answers at preceding levels of the hierarchy. The authors proposed a set of hierarchical classification models which combine machine and human effort on different subtasks, and used Monte Carlo planning to exploit task structures to constrain the policy space for improved tractability. In [37][2], the authors proposed a Lyapunov optimization approach to determine how many new tasks a worker should accept, how many existing tasks a worker should sub-delegate to others, and how much to price a worker's effort based on situational contexts in order to maximize the task results quality and minimize the latency time.

In [38], the authors proposed a graphic framework to monitor crowdsourcing processes and represent the workflows of crowdsourcing applications. In [39], the authors proposed an efficient budget allocation algorithm for multiple complex workflows which consist of multiple inter-dependent tasks.

In [40], the authors proposed a workflow model which considers workers' abilities, the difficulty to improve task results, and the preference of the crowdsourcers to bridge the gap that the characteristics of workflow in relation to tasks and the environment are not fully understood.

IV. Quality Control

Many current algorithmic crowdsourcing research concentrates on how to improve the expected quality of task results.

Task Preparation

Appropriate design of task instructions is the first step to improve the task results quality. Experienced crowdsourcers have written best practices guidelines on how to design task instructions. In [41], the authors found that most tasks did not state the acceptance criteria clearly, but workers appear to view these tasks favorably. In addition, more task instructions affect worker uptake and increase the latency time as they need to understand more information on tasks.

In [42], the authors proposed a second level of crowdsourcing approach which engages workers to monitor the task instructions to prevent crowdsourcers from posting illegal or objectionable tasks. In [43], the authors proposed to expose the worker population to a small set of sample tasks to build a model of how difficult tasks are to the workers. The learnt model can be further applied to future tasks to allocate hard and easy tasks to workers with different capability levels.

Expert Review

A common quality control method is assigning a task to multiple workers and aggregating results to get a higher quality result, but it is redundant to review all results. In [44], the authors proposed a semi-supervised learning algorithm to choose the most informative result instances in labelling tasks, and invite experts to review these results. In [45], the authors proposed a creator-evaluator approach to dynamically invite workers to review complex task results.

Worker Reputation

Workers' reputation, or performance track records, can be used to improve task results quality. In [46], the authors proposed Sembler approach which combines majority voting, reputation modelling for workers, and linguistic contexts to achieve the design objective. In [47], the authors used workers' reputation to compute weighted average answers for multiple choice tasks.

However, it is costly for crowdsourcers to provide feedbacks on the quality of all results obtained because crowdsourcing tasks are small and allocated to multiple workers redundantly. Therefore, reputation-based approaches are required to adapt to the lack of direct feedbacks. In [48][49], the authors incorporated a trust model into a fusion method which combines task results based on the trust parameters. They proposed an inference algorithm which computes the aggregated results and workers' individual trustworthiness based on the maximum likelihood framework.

In addition, workers' self-reported confidence can also help to improve the task results quality. In [50], the authors proposed a method for workers to declare their confidence indirectly by choosing from a set of reward plans to avoid workers being over-confident or untruthful.

In winner-take-all contest-based crowdsourcing, workers whose entries are not selected receive no reward, which decreases the social welfare due to wasted effort. In [51], the authors proposed a discrete choice model to obtain qualities of workers' output first and filter out low-expertise workers before they produce a solution to join the contest.

V. Exploring Ethical Dilemmas

GenEth

In resource allocation or other algorithmic problems, we need to design systems that model their decision making closely to human ethical principles. In the recent years, software has been developed to explore different scenarios where dilemmas are present. According to Anderson [9], the GenEth ethical dilemma analyzer can be used to address ethical issues pertaining to intelligent systems. They are likely to exceed the ability of the original system designers and is used as a platform to include ethicists into the loop to codify principles into the systems. GenEth also uses a graphic user interface to discuss ethical dilemmas in a specific scenario and applies inductive logic programming to infer principles of ethical actions.

Moral Machine Project

The Massachusetts Institute of Technology (MIT) initiated the Moral Machine Program (MMP), [52] which leverages on crowdsourcing to find resolutions for ethical dilemmas. It focuses on the perception of autonomous vehicles (AV), which are potentially to cause harm if they malfunction. The MMP allows the crowds to judge various ethical dilemma and select which outcomes they resonate with. The outcomes are then analyzed according to 8 different preferences. Also, the MMP allows a graphic user interface for users to design their own ethical dilemmas and seek opinion from other participants.

VI. Individual Ethical Decision Frameworks

The research community largely concurs that generalized frameworks are preferred over ad-hoc ones. This is because flexible incorporation of ethical standards enables morally right usage and prevent unethical use since ethical bounds can be contextual and changes depending on time and context. MoralDM is proposed to deal with ethical dilemmas by leveraging on 2 mechanisms: 1) First-principles reasoning and 2) Analogical reasoning. However, as the number of past cases increases exponentially, the approach used by MoralDM becomes computationally impossible.

Therefore, in [54], MoralDM is extended with structure mapping which trims the search space.

Another framework that empowers agents to make judgments on the morality of its and other agents' actions was introduced in [8]. The theories of good and right was represented and was based on awareness and evaluation. Based on the beliefs and goals of the agent, the evaluation process generates the possible actions and desirable actions. The goodness process then computes the ethical actions based on the agents' beliefs, desires, actions and moral value rules.

According to [11], the authors proposed two methods towards developing a general ethical decision-making framework for AI based on game theory and machine learning. For the game theory-based framework, the authors proposed the extensive form as a foundation to represent dilemmas. On the other hand, for machine learning based ethical decision making, it classifies whether a given action under a given scenario is morally right or wrong.

VII. Collective Ethical Decision Frameworks

In this section, we discuss the decision-making frameworks which help a group of autonomous entities to choose ethical actions collectively. There is a question in the research community of whether it is enough to create a society of well-coordinated and collaborative agents acting with human wellbeing as their primary concern, by allowing individual agents to behave ethically and judge the ethics of other agents' actions? According to [12], it is not enough, and requires a set of primary rules governing social norms and allowing the creation, modification and suppression of the rules as different scenarios arise. According to [52], the authors proposed a framework that uses social norms to govern autonomous entities' behaviors. While in [53], the authors suggested to imbue the individual agents with ethical decision making mechanisms, a population of agents can use different roles when evaluating the choices of actions with moral considerations in a given scenario. While [55] suggested a voting-based system for autonomous entities to make collective ethical decisions.

VIII. Ethics in Human-AI Interactions

In this section, we discuss the AI applications which attempt to persuade people's behaviors. These principles for influencing human actions have principles established by the Belmont Report [55]. The three key requirements for the principles include: 1) People's personal autonomy should not be violated; 2) Benefits brought about by the technology should outweigh the risks; and 3) Benefits and risk should be distributed fairly among the users.

References

[1] Anhai Doan, Raghu Ramakrishnan, and Alon Y. Halevy. Crowdsourcing systems on the world wide web. Communications of the ACM, 54(4):86–96, 2011.

[2] Han Yu, Chunyan Miao, Cyril Leung, Yiqiang Chen, Simon Fauvel, Victor R. Lesser, and Qiang Yang. Mitigating herding in hierarchical crowdsourcing networks. Scientific Reports, 6(4), 2016.

[3] Zhengxiang Pan, Han Yu, Chunyan Miao, and Cyril Leung. Crowdsensing air quality with camera-enabled mobile devices. In IAAI, pages 4728–4733, 2017.

[4] Yongqing Zheng, Han Yu, Lizhen Cui, Chunyan Miao, Cyril Leung, and Qiang Yang. SmartHS: An AI platform for improving government service provision. In IAAI, 2018.

[5] Altshuler, T., Katoshevski, R., & Shiftan, Y. (2017). Ride Sharing and Dynamic Networks Analysis. CoRR, abs/1706.00581.

[6] Bei, X., & Zhang, S. (2018). Algorithms for Trip-Vehicle Assignment in Ride-Sharing. AAAI.

[7] Han Yu, Chunyan Miao, Yiqiang Chen, Simon Fauvel, Xiaoming Li, and Victor R. Lesser. Algorithmic management for improving collective productivity in crowdsourcing. Scientific Reports, 7(12541), 2017.

[8] Nicolas Cointe, Gr´egory Bonnet, and Olivier Boissier. Ethical judgment of agents' behaviors in multi-agent systems. In AAMAS, pages 1106–1114, 2016.

[9] Ben Goertzel and Cassio Pennachin. Artificial General Intelligence. Springer, Berlin, Heidelberg, 2007.

[10] Bruce M. Mclaren. Computational models of ethical reasoning: Challenges, initial steps, and future directions. IEEE Intell. Syst., 21(4):29–37, 2006.

[11] Vincent Conitzer, Walter Sinnott-Armstrong, Jana Schaich Borg, Yuan Deng, and Max Kramer. Moral decision making frameworks for artificial intelligence. In AAAI, pages 4831–4835, 2017.

[12] Ugo Pagallo. Even angels need the rules: AI, roboethics, and the law. In ECAI, pages 209–215, 2016.

[13] Avi Segal, Ya'akov Gal, Ece Kamar, Eric Horvitz, Alex Bowyer, and Grant Miller. Intervention strategies for increasing engagement in crowdsourcing: Platform, predictions, and experiments. In IJCAI, pages 3861–3867, 2016.

[14] Yexiang Xue, Bistra Dilkina, Theodoros Damoulas, Daniel Fink, Carla Gomes, and Steve Kelling. Improving your chances: Boosting citizen science discovery. In HCOMP, pages 198–206, 2013.

[15] Avshalom Elmalech and Barbara J. Grosz. "But you promised": Methods to improve crowd engagement in non-ground truth tasks. In HCOMP, pages 21–30, 2017.

[16] Ricardo Matsumura Araujo. 99designs: An analysis of creative competition in crowdsourced design. In HCOMP, pages 17–24, 2013.

[17] Victor Naroditskiy, Sebastian Stein, Mirco Tonin, Long Tran-Thanh, Michael Vlassopoulos and Nicholas R. Jennings. Referral incentives in crowdfunding. In HCOMP, pages 171–183, 2014.

[18] Yuezhou Lv and Thomas Moscibroda. Incentive networks. In AAAI, pages 1270–1276, 2015.

[19] Djellel Eddine Difallah, Michele Catasta, Gianluca Demartini, and Philippe Cudr´e-Mauroux. Scaling-up the crowd: Micro-task pricing schemes for worker retention and latency improvement. In HCOMP, 2014.

[20] Ming Yin and Yiling Chen. Bonus or not? learn to reward in crowdsourcing. In IJCAI, pages 201–207, 2015.

[21] Arpita Biswas, Deepthi Chander, Koustuv Dasgupta, Koyel Mukherjee, Mridula Singh, and Tridib Mukherjee. PISCES: Participatory incentive strategies for effective community engagement in smart cities. In HCOMP, pages 22–31, 2015.

[22] Goran Radanovic and Boi Faltings. Learning to scale payments in crowdsourcing with PropeRBoost. In HCOMP, pages 179–188, 2016.

[23] Dafna Shahaf and Eric Horvitz. Generalized task markets for human and machine computation. In AAAI, pages 986–993, 2010.

[24] Ruggiero Cavallo and Shaili Jain. Efficient crowdsourcing contests. In AAMAS, pages 677–

686, 2012.

[25] Hoda Heidari and Michael Kearns. Depth-workload tradeoffs for workforce organization. In HCOMP, pages 60–68, 2013.

[26] Han Yu, Chunyan Miao, Bo An, Cyril Leung, and Victor R. Lesser. A reputation management approach for resource constrained trustee agents. In IJCAI, pages 418–424, 2013.

[27] Shweta Jain, Sujit Gujar, Onno Zoeter, and Y. Narahari. A quality assuring multi-armed bandit crowdsourcing mechanism with incentive compatible learning. In AAMAS, pages 1609–1610, 2014.

[28] Arpita Biswas, Shweta Jain, Debmalya Mandal, and Y. Narahari. A truthful budget feasible multi-armed bandit mechanism for crowdsourcing time critical tasks. In AAMAS, pages 1101–1109, 2015.

[29] Praphul Chandra, Yadati Narahari, Debmalya Mandal, and Prasenjit Dey. Novel mechanisms for online crowdsourcing with unreliable, strategic agents. In AAAI, pages 1256–1262, 2015.

[30] Cen Chen, Shih-Fen Cheng, Aldy Gunawan, Archan Misra, Koustuv Dasgupta, and Deepthi Chander. TRACCS: Trajectory-aware coordinated urban crowd-sourcing. In HCOMP, pages 30–40, 2014.

[31] Gagan Goel, Afshin Nikzad, and Adish Singla. Mechanism design for crowdsourcing markets with heterogeneous tasks. In HCOMP, pages 77–86, 2014.

[32] Sepehr Assadi, Justin Hsu, and Shahin Jabbari. Online assignment of heterogeneous tasks in crowdsourcing markets. In HCOMP, pages 12–21, 2015.

[33] Jonathan Bragg, Andrey Kolobov, Mausam, and Daniel S. Weld. Parallel task routing for crowdsourcing. In HCOMP, pages 11–21, 2014.

[34] Zhengxiang Pan, Han Yu, Chunyan Miao, and Cyril Leung. Efficient collaborative crowdsourcing. In AAAI, pages 4248–4249, 2016.

[35] Swaprava Nath and Balakrishnan (Murali) Narayanaswamy. Productive output in hierarchical crowdsourcing. In AAMAS, pages 469–476, 2014.

[36] Ece Kamar and Eric Horvitz. Planning for crowdsourcing hierarchical tasks. In AAMAS, pages

1191–1199, 2015.

[37] Han Yu, Chunyan Miao, Zhiqi Shen, Cyril Leung, Yiqiang Chen, and Qiang Yang. Efficient task sub-delegation for crowdsourcing. In AAAI, pages 1305–1311, 2015.

[38] David S´anchez-Charles, Victor Munt´es-Mulero, Marc Sol´e, and Jordi Nin. CrowdWON: A modelling language for crowd processes based on workflow nets. In AAAI, pages 1284–1290, 2015.

[39] Long Tran-Thanh, Trung Dong Huynh, Avi Rosenfeld, Sarvapali D. Ramchurn, and Nicholas R. Jennings. Crowdsourcing complex workflows under budget constraints. In AAAI, pages 1298–1304, 2015.

[40] Shinsuke Goto, Toru Ishida, and Donghui Lin. Understanding crowdsourcing workflow: Modeling and optimizing iterative and parallel processes. In HCOMP, pages 52–58, 2016.

[41] Meng-Han Wu and Alexander J. Quinn. Confusing the crowd: Task instruction quality on Amazon Mechanical Turk. In HCOMP, pages 206–215, 2017.

[42] Yukino Baba, Hisashi Kashima, Kei Kinoshita, Goushi Yamaguchi, and Yosuke Akiyoshi. Leveraging crowdsourcing to detect improper tasks in crowdsourcing marketplaces. In AAAI, pages 1487–1492, 2013.

[43] WeiWang, Xiang-Yu Guo, Shao-Yuan Li, Yuan Jiang, and Zhi-Hua Zhou. Obtaining high-quality label by distinguishing between easy and hard items in crowdsourcing. In IJCAI, pages 2964–2970, 2017.

[44] Mengchen Liu, Liu Jiang, Junlin Liu, Xiting Wang, Jun Zhu, and Shixia Liu. Improving learning-from-crowds through expert validation. In IJCAI, pages 2329–2336, 2017.

[45] Takeru Sunahase, Yukino Baba, and Hisashi Kashima. Pairwise HITS: Quality estimation from pairwise comparisons in creator-evaluator crowdsourcing process. In AAAI, pages 977–983, 2017.

[46] Xian Wu, Wei Fan, and Yong Yu. Sembler: Ensembling crowd sequential labeling for improved quality. In AAAI, pages 1713–1719, 2012.

[47] Bahadir Ismail Aydin, Yavuz Selim Yilmaz, Yaliang Li, Qi Li, Jing Gao, and Murat Demirbas. Crowdsourcing for multiple-choice question answering. In AAAI, pages 2946–2953, 2014.

[48] Satoshi Oyama, Yukino Baba, Yuko Sakurai, and Hisashi Kashima. Accurate integration of crowdsourced labels using workers' self-reported confidence scores. In IJCAI, pages 2554–2560, 2013.

[49] Matteo Venanzi, Alex Rogers, and Nicholas R. Jennings. Trust-based fusion of untrustworthy information in crowdsourcing applications. In AAMAS, pages 829–836, 2013.

[50] Yuko Sakurai, Tenda Okimoto, Masaaki Oka, Masato Shinoda, and Maokoto Yokoo. Quality-control mechanism utilizing worker's confidence for crowdsourced tasks. In AAMAS, pages 1347–1348, 2013.

[51] Haifeng Xu and Kate Larson. Improving the efficiency of crowdsourcing contests. In AAMAS, pages 461–468, 2014.

[52] Munindar P. Singh. Norms as a basis for governing sociotechnical systems. ACM Trans. Intell. Syst. Technol., 5(1):21:1–21:23, 2014.

[53] Joshua Greene, Francesca Rossi, John Tasioulas, Kristen Brent Venable, and Brian Williams. Embedding ethical principles in collective decision support systems. In AAAI, page 4147–4151, 2016.

[54] Ritesh Noothigattu, Snehalkumar 'Neil' S. Gaikwad, Edmond Awad, Sohan Dsouza, Iyad Rahwan, Pradeep Ravikumar, and Ariel D. Procaccia. A voting-based system for ethical decision making. In AAAI, 2018.s 4147–4151, 2016.

[55] The Belmont Report. Technical report, 1978.