



Exploring Self-organisation in Crowd Teams

Ioanna Lykourantzou^{1(✉)}, Antonios Liapis², Costas Papastathis³,
Konstantinos Papangelis⁴, and Costas Vassilakis³

¹ Utrecht University, Utrecht, The Netherlands
i.lykourantzou@uu.nl

² Institute of Digital Games, University of Malta, Msida, Malta
antonios.liapis@um.edu.mt

³ University of Peloponnese, Tripoli, Greece
{cst12079,costas}@uop.gr

⁴ Department of Computer Science and Software Engineering,
Xi'an Jiaotong-Liverpool University, Suzhou, China
K.papangelis@xjtlu.edu.cn

Abstract. Online crowds have the potential to do more complex work in teams, rather than as individuals. Team formation algorithms typically maximize some notion of global utility of team output by allocating people to teams or tasks. However, decisions made by these algorithms do not consider the decisions or preferences of the people themselves. This paper explores a complementary strategy, which relies on the crowd itself to self-organize into effective teams. Our preliminary results show that users perceive the ability to choose their teammate extremely useful in a crowdsourcing setting. We also find that self-organisation makes users feel more productive, creative and responsible for their work product.

Keywords: Crowd teams · Self-organization · Computer-supported collaboration · Creative writing

1 Introduction

As the nature of work is becoming more and more distributed and flexible, creating effective remote teams is becoming an increasingly pertinent problem. Recently, team collaboration has been the subject of increasing research in the crowd work domain. The reason is simple: as the problems that task providers want to delegate to the crowd increase in complexity, individual contributions are not enough and it becomes evident that people from the crowd need to collaborate. Complex problems where crowd collaboration has proven valuable range from mass scale scientific research and article authoring [16], to designing software prototypes [12], and from writing stories [7] to collaborative idea generation [14].

Creating an effective team is well-known to be a problem involving multiple challenges. In a typical workplace setting these include balancing the need

for skill diversity, personality compatibility and schedules, among other parameters [3]. Crowd team formation poses further challenges since crowd team members (i) have usually never worked together before the crowdsourcing task begins, (ii) must perform effectively in relatively little time, and (iii) cannot be assumed to share common values or loyalty to a specific firm (which may be the case when bringing together a remote team from within the same organisation). These challenges mean that the typical methods of team formation, which usually involve pre-profiling team members (in regards to skill, personality etc.) and then placing them to work together, may not suffice in a crowd setting. Instead, new methods are needed to form crowd teams fast and efficiently. In our previous work we proposed “team dating”: a method for crowd team formation in cold-start conditions, i.e. assuming nothing about the profile of the individual workers or the way they will collaborate [9]. Team dating allows people to “try out” different candidate teammates for a number of rounds, before placing them into teams with an appropriate teammate of their choice, with the help of an algorithm that takes into account their in-between evaluations. This ad-hoc manner proved to help create efficient ad-hoc teams on a creative task.

Similarly to most existing algorithm-based methods for crowd management (see for example [11, 13]), our previous method suffered from one important disadvantage: not actively involving the workers in the process, but rather assigning them directly to a task or to a team. However, as latest research in management sciences [8] and also crowdsourcing [18] indicates, too close a monitoring can stifle worker creativity and initiative-taking: two features that are absolutely necessary in creative, complex teamwork.

In this work we explore a new concept: *self-organisation*, which aims to empower crowd workers with the opportunity to choose their teammates, and “guide” the algorithmic process of team formation. Self-organisation is a well-known concept in domains such as Open Source Software Development, or online content co-creation communities like Wikipedia. To the best of our knowledge, however, self-organization has never been explored in a crowd setting. In this first study, we aim to explore how people behave in such a setting, what choices they make, and how.

The rest of this paper is organised as follows. First, we present related literature, focusing on team formation algorithms for crowdsourcing, and findings from management literature on the nature of self organisation. Then we present our methodology, including the description of the interface and algorithm we used to explore crowd team self-organisation, as well as our data collection method. Next we present our preliminary experimental results, focusing on the way people behave when given the choice to select their teammate. We conclude this paper with limitations and future work.

2 Related Work

This paper builds on algorithmic methods for forming teams in a crowd setting but also no self-organization methods for team management. The most critical related work on these two topics is covered below.

2.1 Team Formation Algorithms

Team formation in crowdsourcing is often managed by algorithms. Whereas in a traditional work setting a human manager is often enough to decide which person should work with whom in the context of a task, the scale of crowdsourcing often necessitates automation. A number of algorithms managing the crowd team formation process have emerged recently, which can be mapped to two broad categories: (i) *crowd team building algorithms*, which select which worker should collaborate with whom before the task begins, and (ii) *crowd team coordination algorithms*, which control the team processes after the task has begun.

Crowd team building algorithms, such as the ones proposed in [11], view team formation as a mathematical optimisation problem. They tap on the scale of crowdsourcing, which makes it impossible for traditional methods (e.g. a human manager) to put together an effective team. Assuming a large pool of workers with known profiles (e.g. skill level) and a varied pool of tasks, the objective of crowd team building algorithms is to match each task with a group of workers so as to accomplish the task optimally within given constraints (deadline, upper budget threshold etc.).

Crowd team coordination algorithms come to play after the task has begun. For example, the algorithm proposed by Salehi *et al.* [13] rotates workers across teams based on their viewpoint diversity, in an effort to achieve idea cross-fertilisation and thus increase the innovation capacity of the participating collective. Workers however are not asked whether they would like to switch teams or not, and the rotation decision is only taken by the algorithm. Other works in the area are those by Valentine *et al.* [17] and Kim *et al.* [7], who use computational methods to assign crowd workers to specific parts of the work, either in teams or individually, according to a top-down decision manner and pre-defined roles.

The problem with team management algorithms like the above is that they largely micro-manage the workers by assigning them directly to a team. This approach is indeed appropriate for microtask crowdsourcing, where the crowdsourcing task can be clearly delimited to discrete parts and given to specific workers with specific roles. However, when it comes to more complex work, which is usually the type of work that crowd teams are called to address, such algorithms can stifle creativity and initiative-taking, as indicated by recent research in management sciences [8] and crowdsourcing [6, 18].

2.2 Self Organization for Team Building

Self-organisation is a management term often used to describe the functionality of software development teams, either within a company or in Open Source Software Development communities. As part of the Agile Manifesto [1], it is defined as a process followed by teams that manage their own workload, shift tasks based on needs and best fit, and participate in the group decision making [4]. It has been found to improve the performance of participating teams as it “brings the

decision making authority to the operational level, thus increasing speed and accuracy of problem solving” [10].

Self-organizing teams have certain characteristics [15]. First, they are driven by “zero information”, where prior knowledge does not apply. This enables the team members to challenge existing knowledge status quo and have the potential to create something truly novel. Second, they exhibit autonomy, as they do not have a top-down appointed leader; leadership is a property that emerges as the team members divide their roles [5]. Third, the team pursues ambitious goals (self-transcendence), and fourth, team members have a variety of backgrounds, viewpoints and knowledge (cross-fertilization).

In crowd teams, not a lot of works currently exist giving workers the option to self-organise. An early attempt is the work by Lykourantzou *et al.* [9] who partially delegate the team building process to the crowd workers themselves, by enabling them to try out different candidate teammates, evaluate them, and then make crowd teams based on these indications. Although this work does take worker preference into account, it does so indirectly.

In this paper we start exploring the notion of self-organisation for crowd teams. Given the effectiveness of the approach on other types of teams, such as corporate ones, we experiment with allowing individuals from the crowd to explore the “space” of candidate teammates available to them, discover those with whom they might work best, through a trial-and-error approach, and finally directly indicate their preferences.

3 Methodology

This section describes the task for which team self-organization is tested, along with the user interface and self-organization algorithm that facilitate it.

3.1 Interface Design

Current crowdsourcing platforms do not encourage collaboration, let alone self-organisation and choice of who to work with. Therefore, to explore the properties of self-organisation, and how people behave in this context, we designed a tailor-made framework and interface, outlined in Fig. 1.

The task given to the crowd workers is a creative writing challenge inspired by the exquisite corpse method [2], where participants co-create a story by gradually building on each others’ contributions. To give a sense of competition and inspire motivation, the framework has been designed as a game, played in three rounds. Before the start of the game, a pre-authored story is presented to all users. Then during each round, users work in teams of two and collaboratively write a possible continuation to this story. At the end of the round, individuals vote for the best story (they cannot vote for their own team’s story), and decide whether they would like to continue with the same teammate or not. In case they want to change, they can indicate their preference for another person, choosing from the list of all possible candidate teammate profiles. The most voted story is then

appended to the main story, the team that wrote it receives an award (a score bonus), and a new round begins. In the next round, teams are formed with the help of a self-organisation algorithm, which attempts to best satisfy each user’s desired (or undesired) teammate. Each team will have to continue the main story, as it was formed in the previous round(s). This cycle repeats until, at the end of the third round, the final main story is presented, and the user with the highest score is the winner of the challenge. More details on the specifics of the interface, process, and solicited user feedback are presented below.



(a) Initial player profiles

(b) Team collaboration interface

(c) Assessing a player’s team-mate

(d) Updated player profiles

Fig. 1. Screens from the user interface

Start of the Game: Once the system is synchronized to all players, the instructions of the game are presented. Afterwards, players are asked about their demographic information, measuring: age, gender, ethnicity, education level, employment status, whether they have won or participated in a similar writing task before and their self-perceived creativity level. Then, each player is presented with the start of the pre-authored story and is asked to individually write a continuation for it. We use this input as a “writing sample”, and add it to the profile of the individual, for other users to see.

Teammate Selection: Once players have submitted their individual story writing sample, they choose their first teammate from the list of all players' profiles. A player's profile at this stage consists of the player's username, demographic data, as well as their writing samples (Fig. 1a). In the following stages, when players have already collaborated with and have rated one another, profiles will be enriched to include each player's ratings, both from the person looking at the profile and on average. Players can choose up to two possible teammates with whom they would like to work with. The self-organisation algorithm uses these choices to form the pairs that will be invited to work together in the next round.

Team Collaboration: As soon as the algorithm has placed users in pairs, based on their profile choices, the teams must collaborate to continue the story so far ("main story"). In the first round, the main story consists of the initial pre-authored story. Players' individual writing samples are not taken into account, since all teams must start from the same basis. However, each user is free to propose their sample to their teammate, or write something new together. The interface at this stage, shown in Fig. 1b, consists of four main parts: (1) a clock that indicates how long the team has for their collaboration, (2) a chat area, serving as their communication channel, (3) the story so far ("main story"), and (4) a collaborative text writing area, where users can see in real time what their teammate writes, and work together to continue the main story.

Peer Evaluation: Once they complete their collaborative writing, players evaluate their teammate on a Likert scale of 1 to 5 on: Skillfulness ("How skillful was [teammate's username] in continuing the story?"), Collaboration ability ("How good is [teammate's username] as a collaborator?"), Helpfulness ("[teammate's username] comments were helpful"). Players can also rate their own helpfulness level from 1 to 5. Finally, players can tick one or more options regarding how they are similar with their teammate (the options are "commitment to working hard on this task", "how we think the work should be done", "general abilities to do a task like this", "personal values"). These ratings are used to enrich each player's profile, as explained above.

Voting for Best Story: After providing feedback on their teammate, each individual player votes for their preferred story continuation (direct voting method). Given a number of P participants in the experiment, there are $P/2$ candidate story continuations to choose from. Users can vote for any story, except their team's (they cannot see it as an option). Once all players vote, the story with the most votes is presented to them along with the winning team.

Teammate Selection (Enriched). If there are more rounds remaining, users start a new round by first selecting a teammate, as described above. The only difference is that now players are also asked whether they would like to continue with the same teammate or not. This choice is critical for self-organisation: it

renders players responsible for calculating the relative gain they will have from continuing with the same teammate (e.g. lower communication overhead) versus the risk of not working with another teammate that could potentially get them access to a better story in the next round.

Self-organization Algorithm: The aim of the algorithm is to assist the team self-coordination process, by matching users with those teammates that they mostly prefer working with. The algorithm does so by ranking the possible candidate teams based on their pairwise average profile ratings (affinity), and then gradually selecting those teams with the highest affinity values, until all users have been placed into teams. In case of ties, e.g. two candidate teams having the same pairwise affinity score, the algorithm selects randomly among tied options.

End of Game: After three rounds, all users are prompted to the “Final Survey” screen, which shows the final story with all winning stories appended, the ranking of the users (based on a point system rewarding each round’s winning teams), and a questionnaire that users must fill regarding their overall experience.

3.2 Data Collection

As a first exploratory study, we conducted five experiments with teams of crowdworkers, using only the self-organisation team formation algorithm described above. In total, 34 crowdworkers participated. While more crowdworkers joined each experiment, they swiftly dropped out and were not part of any phase of the team-based tasks; their data is omitted without affecting the validity of the findings. Three of the sessions had 6 participants working in 3 teams, while two sessions had 8 participants working in 4 teams. The majority of participants self-identified as aged 18–24 (94%), Asian (91%), male (62%) and with a Bachelors (53%) or High school degree or equivalent (29%). While only a few participants (18%) reported some prior experience with creative writing, most of them reported being highly creative, with an average score of 3.9 out of 5.

4 Results

This section compiles the findings of the five experiments detailed in Sect. 3.2.

4.1 Research Questions

We wish to analyse the participants’ opinions on self-organization as supported in this experiment and implications in their selection of team-mate, relating to the following questions:

1. Do people like being able to choose their teammate?
2. How does having the choice of teammate affect participants?

3. What matters the most when having the option to select a teammate?
4. How did players decide which team's story to vote for?

The questions are analysed at the level of the individual, using the questionnaire data that players filled in at the end of the experiment. Since we are dealing with a single population, we employ the one sample *t*-test and Chi-Square Goodness-of-Fit Test methods.

Table 1. Reasons for choosing a teammate.

What mattered the most when choosing teammate?	
Choosing the person that would make me win	27%
I chose randomly	13%
I chose the people whose initial story I liked the most	27%
I chose the people whose profile information I liked the most	27%
Other	6%

4.2 Player Behavior in a Self-organisation Context

RQ1: Users Find It Useful to Choose Their Teammates. We conduct a one-sample *t*-test on their answers to the question “How useful was being able to choose your teammate?” of the final questionnaire, against the hypothetical mean of 3.0 (denoting a neutral opinion). Results show that users show a significantly higher than neutral preference to being able to choose their teammate, with $t(29) = 2.134, p < 0.005$. A mediation analysis revealed that this preference is not affected by being in a winning team or not.

Table 2. Effects of choice.

Having the choice to select who I work with, made me...	
More productive	27%
Feel in control of my own work	13%
Be more responsible for what I write	27%
More creative	27%
More tired	3%
Other	3%

RQ2: Having the Choice to Select One’s Teammates Makes Users Feel Productive, Responsible and Creative. A Chi-Square Goodness-of-Fit Test revealed that there are significant differences among the different ways that this choice could affect users, with $x^2(5) = 12.00, p < 0.005$. Specifically, having the choice to select one’s teammates made users feel more productive, creative, responsible for what they write and more in control of their work (Table 2). Post-hoc binomial pairwise tests revealed statistical significant differences across all user-reported answers with at least $p < 0.05$, except between answers “more tired” and “other” where no statistically significant difference was found. No statistically significant differences were found on whether the user was in a winning team or not.

RQ3: Users Selected Teammates Based on a Variety of Reasons. A Chi-Square Goodness-of-Fit Test revealed that there are no statistically significant differences among the different ways that users select teammates, with $x^2(4) = 5.33, p = 0.26$. Player choices from our sample seem to be distributed primarily among selecting a teammate based on winning potential, writing sample and profile, and to a lesser extent randomly or for other reasons (Table 1).

Table 3. Story voting criteria.

How did you decide which story to vote?	
I always vote for the story I like the most	83%
I sometimes voted for a not good story, to have more chances to win	3%
I voted randomly	14%
Other	0%

RQ4: Users Voted for the Story They Liked the Most, Not Strategically. A Chi-Square Goodness-of-Fit Test revealed that there are very significant differences among the different ways that users select the winning story, with $x^2(3) = 133.625, p < 0.001$. The dominant reason for voting for a story is how much users liked it (Table 3). Users did not vote strategically (i.e. for a worse story in hope that their would win) or randomly. Post-hoc binomial pairwise tests revealed statistical significant differences only across the dominant reason mentioned above (“I voted for the story I liked the most”) and the rest, but not among the other reasons.

4.3 Sample Session

To shed more light onto the process followed by different crowd workers when collaboratively writing in teams, we analyse an indicative session with 8 crowd

workers. The session is one of the five sessions analyzed in the quantitative results section above. The winning story of the first round is the one by users *shuyuan-liu16* and *yihao15*. Interestingly, this story is very unlike to the initial writing sample that any of the two users contributed. The harmonious collaboration of this team was rewarded by other players, and the winning team chose to stay together, with both users rating each other highly (pairwise rating of 4.5 out of 5). From the other competing teams, at least one member per team declared that they wanted to change teammate. It is also interesting to note that both members of the winning team declared to have 3 out of 4 possible collaboration style elements in common (“task commitment”, “way of thinking”, “general abilities”, differing only in the “personal values” element). In contrast, the members of the losing teams either declared no collaboration style element in common (2 out of the 3 losing teams), or they had a large imbalance in the collaboration style elements they declared (in the last team, one member declared to share no elements, while the other member declared sharing all four). This observation indicates that the dissimilarity in regards to the collaboration styles may play an important role to the perceptions of the team members and their eventual team’s performance. Further research is however needed to statistically examine this observation.

Moving forward to the next rounds, we observe that the winning story of the next round is that of a newly formed team, that of users *HambinQin* and *grizzHuang*. Similarly to the previous winning team, this team declared to share many collaboration style elements (4outof4) and its members rated each other highly (4.25 out of 5). From the losing teams, two out of three also shared at least one collaboration style element, and these teams declared that they wanted to continue working together, despite having lost. The members of these teams also rated each other highly (4.75 out of 5, and 5 out of 5 respectively). The losing team that expressed a wish to change teammates was the team that shared no common collaboration style element, and teammates rated each other very low (2 out of 5). This qualitative observation also indicates that there seems to be a connection between the number of collaboration style elements a team shares, their inter-team evaluations and their wish to stay together, as well as their performance. Further research is needed to confirm this observation. Given that all but one teams wanted to stay together, all four teams remained unchanged by the algorithm in the third and final round. The winning story was again by a different team, this time by the team of users *nakashimaritsu* and *Hecate*, who had rated each other highly in the previous round. The inter-team ratings and collaboration styles that the teams had declared to share in the previous round remained unchanged.

A final interesting observation with regards to this sample is the fact that the number of total user votes on the winning stories decreases from round to round. After the first round, all 8 users voted for their preferred story. However in rounds 2 and 3 only 6 users voted. This pattern seems to repeat in the other experiments. Although users declared that they did not intentionally vote strategically, this finding indicates that perhaps some of them avoided to vote,

to give more chances to their team by not taking any action “against” it. Further research is nevertheless required as to the reasons why some users choose not to vote, and whether this finding is supported by statistical evidence.

5 Conclusion, Limitations and Future Work

The results of our exploratory study show that users consider being able to choose their teammates, and thus being able to affect the decision of the team formation process, extremely useful in a crowdsourcing setting. Being involved in the process—and not a mere component of it—made users feel significantly more productive, responsible for their choices and creative. These findings are fully in line with what research in organisational sciences has shown, i.e. that involving users (in this case workers) in the decision-making process helps motivate feelings of empowerment, responsibility and ownership of one’s work, with eventual positive effects on team performance.

As an exploratory work, this study has a number of limitations which can eventually form part of future work. The study only explored the behavior of users in a self-organisation context, without a control condition. This is appropriate to gain insights regarding the process, but it also means that further work is needed to examine, for instance, the differences between self-organised team formation and more traditional team building methods. Moreover, the results of this study relied only on the final questionnaire that users answered. Much richer information can be extracted by quantitatively analysing their pairwise evaluations or their profile choices after each round, and this can also be the subject of future work. Such an analysis can help reveal what makes a “winning team” and if these teams share some common characteristics compared to non-winning teams. There is also a need for deeper analysis of the quality of the stories themselves (winning or not) to reveal the extent to which users may vote strategically, even without realising it. Finally, a sentiment and text analysis can be performed on the text produced during team collaboration phases, on the shared collaborative document or on the chat. Such an analysis can reveal more fine-grained elements of team collaboration, such as turn-taking styles, signs of social or cultural affinity, trust and common ground building, equality of communication etc. Exploring these elements, especially between rounds, can help reveal why some teams persist while others dissolve, and therefore help understand to a deeper level the properties of human self-organised collaboration.

References

1. Beck, K., et al.: The agile manifesto (2001). <https://agilemanifesto.org/>
2. Brotchie, A., Gooding, M.: *A Book of Surrealist Games*. Redstone Press, London (1995)
3. Gilley, J.W., Morris, M.L., Waite, A.M., Coates, T., Veliquette, A.: Integrated theoretical model for building effective teams. *Adv. Dev. Hum. Resour.* **12**, 7–28 (2010). <https://doi.org/10.1177/1523422310365309>

4. Highsmith, J.: *Agile Project Management: Creating Innovative Products*. Addison-Wesley, Boston (2004)
5. Hoda, R., Noble, J., Marshall, S.: Organizing self-organizing teams. *Assoc. Comput. Manuf. J.* (2010). <https://doi.org/10.1109/NSREC.2017.8115448>
6. Jahanbakhsh, F., Fu, W.T., Karahalios, K., Marinov, D., Bailey, B.: You want me to work with who?: Stakeholder perceptions of automated team formation in project-based courses. In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, CHI 2017*, pp. 3201–3212. ACM, New York (2017). <https://doi.org/10.1145/3025453.3026011>
7. Kim, J., Cheng, J., Bernstein, M.S.: Ensemble: exploring complementary strengths of leaders and crowds in creative collaboration. In: *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pp. 745–755 (2014). <https://doi.org/10.1145/2531602.2531638>
8. Lawler III, E.E., Worley, C.G.: Designing organizations that are built to change. *MIT Sloan Manag. Rev.* **48**, 19–23 (2006)
9. Lykourantzou, I., Wang, S., Kraut, R.E., Dow, S.P.: Team dating: a self-organized team formation strategy for collaborative crowdsourcing. In: *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pp. 1243–1249. ACM (2016)
10. Moe, N.B., Dingsøy, T.: Scrum and team effectiveness: theory and practice. In: Abrahamsson, P., Baskerville, R., Conboy, K., Fitzgerald, B., Morgan, L., Wang, X. (eds.) *XP 2008. LNIP*, vol. 9, pp. 11–20. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-68255-4_2
11. Rahman, H., Roy, S.B., Thirumuruganathan, S., Amer-Yahia, S., Das, G.: Optimized group formation for solving collaborative tasks. *VLDB J.* **28**(1), 1–23 (2019). <https://doi.org/10.1007/s00778-018-0516-7>
12. Retelny, D., et al.: Expert crowdsourcing with flash teams. In: *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (2014). <https://doi.org/10.1145/2642918.2647409>
13. Salehi, N., Bernstein, M.S.: Hive: collective design through network rotation. In: *Proceedings of the ACM on Human-Computer Interaction - CSCW*, pp. 151:1–151:26 (2018)
14. Siangliulue, P., Chan, J., Dow, S.P., Gajos, K.Z.: IdeaHound: improving large-scale collaborative ideation with crowd-powered real-time semantic modeling. In: *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, pp. 609–624 (2016). <https://doi.org/10.1145/2984511.2984578>
15. Takeuchi, H., Nonaka, I.: The new product development game. *J. Prod. Innov. Manag.* (1986). [https://doi.org/10.1016/0737-6782\(86\)90053-6](https://doi.org/10.1016/0737-6782(86)90053-6)
16. Vaish, R., et al.: Crowd research: open and scalable university laboratories. In: *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology* (2017). <https://doi.org/10.1145/3126594.3126648>
17. Valentine, M.A., Retelny, D., To, A., Rahmati, N., Doshi, T., Bernstein, M.S.: Flash organizations: crowdsourcing complex work by structuring crowds as organizations. In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 3523–3537 (2017). <https://doi.org/10.1145/3025453.3025811>
18. Xia, H., et al.: No workflow can ever be enough: how crowdsourcing workflows constrain complex work. *Proc. ACM Hum.-Comput. Interact.* (2017). <https://doi.org/10.1145/3134724>