Chapter 1. Macrotask Crowdsourcing: An Integrated Definition

Ioanna Lykourentzou, Utrecht University
Vassilis-Javed Khan, Eindhoven University of Technology
Konstantinos Papangelis, Xi'an Jiaotong-Liverpool University
Panos Markopoulos, Eindhoven University of Technology

Abstract

The conceptual distinction between microtasks and macrotasks has been made relatively early on in the crowdsourcing literature. However, only recently a handful of research works have explored it explicitly. These works, for the most part, have focused on simply discussing macrotasks within the confines of their own work (e.g. in terms of creativity), without taking into account the multiple facets that working with such tasks involves. This has resulted in the term "macrotask" to be severely convoluted and largely meaning different things to different individuals. More importantly, it has resulted in disregarding macrotask crowdsourcing as a new labor model of its own right. To address this scholarly gap, in this paper we discuss macrotask crowdsourcing from a multitude of dimensions, namely the nature of the problem it can solve, the crowdworker skills it involves, and the work management structures it necessitates. In view of our analysis, we provide a first integrated definition of macrotask crowdsourcing.

Introduction

The distinction between microtasks and macrotasks was made relatively early on in the crowdsourcing literature. Grier (2013) emphasized on the skills and expertise of workers when discussing macrotasks, which he considers as "the professional form of crowdsourcing" and "freelancing on a global scale", which happens in an open, public market contrary to microtasks, which are brief tasks that do not require advanced skills. Crowdsourcing platforms help manage the relationship between the requester who owns the problem and the worker who will execute it, they handle payments, and support practical challenges such as verifying the time worked. Grier as other authors after him, introduced macrotasks in juxtaposition to microtasks in terms of their magnitude of the task and goes as far as to propose a checklist for defining a macrotask: the macrotasker can carry out the work independently without support from the requester, it is simple to describe with clear criteria of completion, it has a clear and concrete deadline, and it requires special skills the requester's organization does not possess. This practical and down to earth guidance helps get one on the way with macrotasking but does not shed much light into how macrotasking differs and needs to be addressed differently than microtasking.

One of the early investigations of task decomposition in crowdsourcing was presented in the case of video annotation (Vondrick et al. 2014). Video annotation is a canonical example of a crowdsourcing task where valuable results are obtained by combining small contributions by many crowdworkers. To assess the value of task decomposition Vondrick et al. (2014) compared annotating video for a single object per crowdworker which they consider as a microtask to annotating a video segment for a whole set of objects which they consider to be a macrotask. They note how video annotation of a segment for all objects may cost more time but it allows the crowdworker to develop ownership of the result and deliver higher quality labels. Furthermore, errors in coding specific objects are distributed over different segments and handled by different co-workers, while the effort a crowdworker invests to visually decode a scene is committed only once for all objects that need to be identified. Beyond video annotation, Machado et al. (2014) discuss crowdsourcing in the context of software development, where in line with Grier et al discussed above they consider macrotasks as larger than microtasks and requiring specific knowledge from the crowdworker. They propose software testing as an example of a macrotask and discuss macrotasking practices by the Brazilian company Crowdtest or the American Utest.

Cheng et al. (2015) is the first (and to this point the only) empirical study that focuses explicitly on the trade-offs involved in decomposing macrotasks to micro-tasks. They examined task performance for three types of tasks, which included simple arithmetic, sorting text and audio transcription. Their results suggest that decomposing macrotasks to smaller parts, may make the total task completion time longer but it enhances the task quality and makes work easier. The experiment and their whole discussion considers macro- and micro-tasks as relative descriptions, the latter being a decomposition of the former. The macrotasks in their experiment are very simple, namely adding 10 numbers, sorting 7 lines of text or transcribing 30 seconds of audio. This helps test the decomposition decision very directly in the experiment, but does not help transposing the conclusions of this experiment to situations where leadership, creativity, initiative, coordination might be manifested, as is often the case in what one might consider a more complex task in real life. Cheng et al. (2015), also considered how interruptions may affect the task completion time arguing that macrotasks are less resilient to interruptions. However, this result may indeed be very specific to the nature of the experimental tasks they used where task decomposition translates directly to lower demands on short term memory - which is challenged during interruptions. Arguably decomposing macrotasks of much larger scale such as creating a logo, which might take minutes or hours rather than seconds, is not likely to produce similar gains.

Haas et al. 2015 identify quality control as one of the major challenges in setting up workflows involving macrotasking. They consider macrotasks as ones that cannot be easily decomposed, or where larger context (e.g. domain knowledge) or a significant initial investment of time is needed before workers can engage in task execution in order to develop a global context, e.g,. authoring a paper or a presentation. They point out that while crowdsourcing researchers have sought efficiency and quality gains in the algorithmic decomposition of tasks and synthesis of individual crowdworker micro-contributions, there can be substantial benefits to recruit task workers to perform macrotasks that last longer and which apply more flexible compensation schemes, combining some of the benefits of microtasks and traditional freelance work. Haas et al introduce Argonaut a framework for managing macrotask based workflows that addresses a major challenge for automating macrotask work, which is to ensure the quality of the work. The Argonaut framework profiles workers in terms of the work quality they deliver and their speed, and uses these profiles to sustain a hierarchy of roles (workers, reviewers and top-tier reviewers). Workers are assigned suitable roles within the macrotask workflow and are promoted, demoted dynamically depending on task availability.

Li et al. (2016) consider macrotasks as those lasting several hours. They argue that workers are not easily motivated to carry these out, and that they are challenging to define/decompose. For this, they suggest that macrotasking is an important topic for future research.

Valentine et al. (2017) report on an approach for handling a specific class of macrotasks that are complex and open-ended, and which are difficult to crowdsource using microtasking because it is difficult to articulate, modularize, and pre-specify the actions needed to achieve them. To do so, they propose ways to structure the crowd in 'flash organizations' that involve defining formal structures such as roles, teams, and hierarchies that delineate responsibilities, interdependencies and information flow without pre-specifying all actions. Their approach is characterized by a) a de-individualized role hierarchy (as can be found in organizations like movie crews, disaster response teams, or the army) where collaboration is based on workers' knowledge of the roles rather than their knowledge of each other. b) a continuous reconfiguration of the organization e.g., by changing roles or adding teams. Valentine et al. (2017) demonstrate the feasibility of their approach through three case studies concerning respectively creating an application for emergency medical technicians (EMTs) to report trauma injuries from an ambulance en route to the hospital designing, manufacturing, and playtesting a storytelling card game and an accompanying mobile application, and creating an enterprise web portal to administer client workshops.

Implementing such organizational structures in crowdsourcing in order to support macrotasks brings about challenges related to incentivising workers. For example, personal preferences or biases may colour assessments of solution quality. Xie et al. (2018) propose an optimization approach for incentivizing workers to provide high-quality contributions and empirically evaluate the effectiveness and efficiency of their approach.

On the Nature of the Problem

To understand the reasons that may necessitate a shift from microtasking to macrotasking, one must first understand the problems that each crowdsourcing model can and cannot solve. Drawing from organizational management literature, below we classify crowdsourcing models according to the problem attributes that each can solve (Fig. 1).

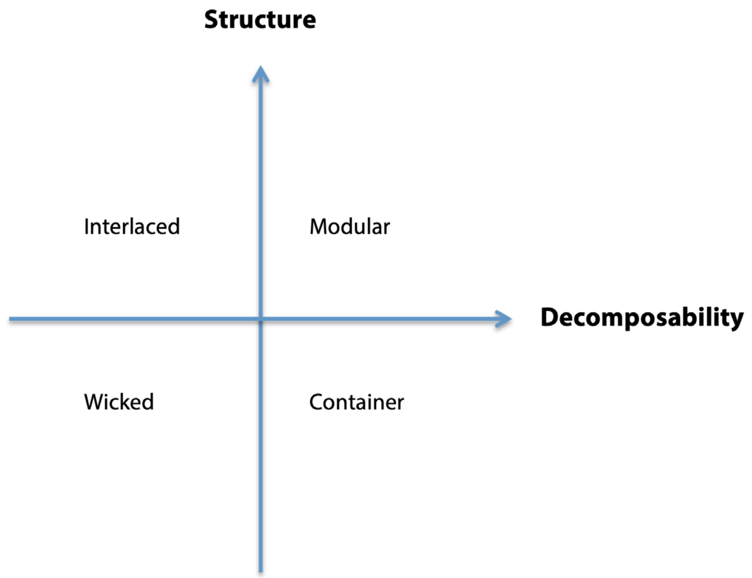

Figure 1: The macrotask dimension space. To draw this diagram we assume all macrotasks are complex. Then we have a cartesian space of them with dimensions structure & decomposability. This space characterizes four types of macrotasks: interlaced, modular, wicked and container.

Knowledge problems can be categorized based on three attributes: complexity, decomposability and structure (Nickerson and Zenger, 2004; Huang and Holden, 2016). **Complexity** refers to the number of knowledge domains that are relevant to the problem, and the strength of their interactions. Simple problems tend to involve few knowledge domains, with a low degree of domain interdependency. More complex problems involve a large number of knowledge domains, which share an strong degree of domain interaction. **Decomposability** measures whether the problem can be divided into sub-problems, and the granularity that this division can reach. Decomposable problems can be broken down to separate sub-problems, each drawing from distinct knowledge sets, which can be solved independently and with little communication or collaboration among problem solvers. Non-decomposable problems on the other hand, are impractical or even impossible to sub-divide into separate sub-problems, because the interdependencies among their knowledge domains are too extensive. For such problems, if a solution is to be found, this needs to be an overall solution, which enables problem solvers to maintain the global problem context. **Structure** is the degree to which one can determine all the knowledge domains relevant to the problem, the expertise needed to solve it, and the interrelations between the identified domains. Well-structured problems consist of an clear set of relevant knowledge domains. The boundaries and interactions among these domains can be easily understood, and there are explicit and widely accepted approaches to solve the problem. On the other hand, ill-structured problems are those where the relevant knowledge domains, necessary to solve the problem, are not evident, the boundaries among these domains are ambiguous and their in-between interactions are very poorly understood. Conversely, consensus approaches may not be the optimal approach; rather these problems often benefit from "spontaneous" disruptive innovations, which often challenge scientific and industrial status quos and offer new ways of interpreting the problem and its solution.

This classification enables us to position existing and future crowdsourcing models with respect to the problems that they can solve, and the problems for which they are not suitable.

Tasks related to data such as: categorization, curation, or enrichment (Kittur et al. 2008; Musthag and Ganesan 2013) tackle problems that are simple, well-structured and decomposable. The bulk of tasks in most commercial crowdsourcing platforms are of that sort.

Macrotask Type 1 (Modular): Well-structured, High-decomposability problems

The first macrotasking type is meant to solve problems that are*, decomposable,* and *well-structured.* These form the majority of complex problems that current crowdsourcing literature and applications focus on, and understandably so, since these problems can be addressed using a "*divide and conquer approach*". The problem is first broken down to smaller, distinct work units, i.e. at micro-task level. Then, the distinct micro-tasks are assigned in parallel to multiple workers, and finally they are re-composed to a final output by combining the separate smaller subtasks.

The difference with what we might call 'vanilla' micro-tasking is that, because of the problem complexity, the way of breaking down the problem to micro-tasks is not evident and may require the *involvement of experts, who design tailor-made workflows for the crowd to follow*. These experts in collaboration with the task requester, often determine how the macrotask should be decomposed into smaller chunks, and how to recompose these once completed. Because of the involvement of experts the decomposition of microtask level can be costly (Kim et al. 2014; Chan et al. 2016). Nevertheless, once the workflow has been designed, it can be very effective (Teevan et al. 2016). That being said, this approach suffer from non-generalization. Because the workflows are usually tailored to the very specific problem, they cannot be generalized easily to handle other problem instances but it is non-generalizable as it is, for the most part, tailed to a very specific problem.

The resulting macrotasks may not be homogeneous in terms of size, or skill requirement. Examples of macrotask type 1 include: taxonomy creation (Chilton et al. 2013), itinerary planning (Zhang et al. 2012), editing and correcting a document (Bernstein et al. 2010), or aggregating multiple word or sentence-level translations to form a larger corpus (Ambati et al. 2012; Zaidan and Callison-Burch 2011).

Macrotask Type 2 (Interlaced): Well-structured, Low-decomposability problems

The second type of macrotasks aims to tackle problems that well-structured but they are non-decomposable. In general, these are problems often found at the beginning of creative projects (e.g. when the broad objectives and solution criteria need to be set) and are, for the most part, only processed manually, even if the rest of the project can be broken down into sub-tasks and potentially crowdsourced (Sieg et al. 2010).These problems can be solved through a "continuity of useful action" (Altshuller 2005) where each consecutive contributor maintains the global context and full semantic overview of the problem while iteratively refining it until an acceptable solution is found .

Examples of type 2 macrotasks include: defining a research methodology or formulating an R&D approach.

Macrotask type 3 (Wicked): ill-structured, Low-decomposability problems

The third type of macrotask problems are the so-called "wicked problems" or "holy grail" problems. These are ill-structured tasks, for which the interactions among the relevant knowledge domains (or even the exact required knowledge domains themselves), are not well understood, and the requirements are incomplete, contradictory, and in some cases ever-changing. Wicked problems, in a crowdsourcing context, tend to be innovation idea contests (Majchrzak and Malhotra 2013), where the purpose is to collect as many ideas as possible in search for the few breakthrough ideas, rather than an iterative idea development. There has been limited research on how to process and tackle wicked problems through crowdsourcing. Evidence illustrate that using a sequential process could lead to problems such as fixation with one solution (Jansson and Smith 1991) or solution confounding (Little et al. 2010). However, further research is necessary to shed light on the issue.

Examples of type 3 macrotasks include: open innovation contests such as Lego Mindstorms, Open IDEO etc.

Macrotask Type 4 (Container): ill-structured, High-decomposability problems

The final macrotask type aims to tackle problems that are ill-structured and highly decomposable. Although such problems are not directly addressed in the literature, one could conceptually identify them based on the structure/decomposability matrix that organisational research suggests. Such problems could be those for which the required expertise cannot be determined automatically a priori, but it can be determined with the help of an expert or team of experts. For example, in a crowdsourcing context, such a problem is the coordination of a team of crowd workers. Very recent literature (Wood et al. 2019) has indeed touched upon this phenomenon, reporting that high-reputation crowd workers delegate complex work to other crowd workers or other workers from their social circles. They also often explain the tasks and train (in the form of instructions) their delegates on how to accomplish the (part of) complex work. This method of understanding the ill-structured problem, and then decomposing and delegating it based on experience, could be a precursor of more complex workflows that are needed to handle this type of tasks. Future work is required to research such problems in more detail, and understand which crowdsourcing workflows can be designed to address them.

On the Nature of Skills

Few works in existing microtask crowdsourcing literature focus on workers skills. Although very recent works in the area do try to understand better the needs of the crowdworkers, for example by examining their working conditions or the context they find themselves into (Gray et al. 2016; Irani & Silberman 2013; Martin et al. 2014), these works do not examine which skills a worker has or needs to have. This research gap may be partially attributed to the fact that, apart from language (e.g. English) skills and general perception skills, workers in microtask crowdsourcing are usually not required to have very specialised skills to perform their work. Consequently, microtasking platforms also usually store only worker demographics and the percentage of tasks the worker has successfully completed (number of HITs, Levels, or other name depending on platform). Microtasking platforms do not usually store other worker skills (Ho & Vaughan 2012). In case requesters need workers to have a specialised skill, they mention it in an open field, which workers fill in based on self-assessment. Self-assessment may be biased and its validity as a metric of skill quality is low since not all workers have the same perception of their skills. Less often, requesters may develop a tailor-made test, prior to the actual microtask, to test specialised worker skills. This practice however is costly, and not generalisable.

In addition, microtasking usually relies on *skill homogeneousness*: the problem is decomposed to microtasks that all require the same type of non-expert skill. Consequently, currently not a lot of works in existing crowdsourcing literature analyse the spectrum of worker skills across a variety of possible problems that they could solve. The only works that usually assume a variety of different skills are based on simulations, either across different domains of the same level (Basu Roy et al. 2015), or even across hierarchical skills levels (Mavridis et al. 2016).

Macrotasking on the other hand is innately linked with *skill diversity*, and more fine-grained skill types, including expert and 21st century skills , as well as valid skill identification and evaluation mechanisms. Examples of higher-order cognitive and 21st century skills that macrotask workers might need include: creativity, curiosity and imagination, critical thinking and problem-solving (Creative and Cultural Skills 2017), effective oral and written communication skills, information analysis ability, agility, adaptability and the capacity to learn new knowledge fast, collaboration ability, communication skills, taking initiative, leadership and people management skills (Wagner 2014). Expert skills can be obtained by direct training and “learning by doing”, and naturally include the whole spectrum of today’s and tomorrow’s expertise, with some prominent examples being coding, graphic design skills, research methodology skills, business marketing and communications etc.

Although microtask crowdsourcing practice tends to consider workers as an endless, homogeneous and replaceable mass, the truth is that complex skills and crowd workers who possess them are inevitably expected to be less frequent. Therefore, for macrotask crowdsourcing, it is important to ensure the following:

- **Skill structure and assessment.** Develop mechanisms to assess macrotasking skills with validity, and in a scalable manner (Ipeirotis & Gabrilovich 2014), drawing from a wide range of approaches (from computerised to peer assessment), as well as the skill assessment scientific domain.
- **Develop training opportunities.** Workers who are not at the right skill level should not be excluded at face value. Rather, macrotasking platforms should support worker skill development, by offering training opportunities and scaffolded learning.
- **Access to skill data and skill data sharing.** Provide workers with expert skills with an access to and ownership of their skill data, and the opportunity to share them across platforms. This approach is not only in line with latest data management ethics (see the recent EU GDPR rules, see Voigt & Bussche 2017), but it is also expected to give workers a sense of control, the ability to indicate their skill pertinency, and promote workers mobility and platform cross-fertilisation.

On the Nature of Management

When referring to crowdsourcing, scalability is the key. Unlike traditional management settings, where the human manager needs to organise the work of a few people (up to the level of dozens), the scale of crowdsourcing necessitates automation. For this reason, recent works have focused on algorithm-based human resource allocation in crowdsourcing settings, from two perspectives. From the mathematical optimisation perspective, such algorithms assume a large pool of worker profiles (skills, availability etc.) and a large pool of tasks with certain characteristics (e.g. knowledge domain), and constraints (deadline, budget etc.). In this setting, the objective of the algorithms is to match each task with one or more workers, to accomplish the task optimally (e.g. in terms of quality) with the given constraints (e.g. Basu Roy et al. 2015; Goel et al. 2014; Schmitz and Lykourentzou, 2018). From an organisational perspective, viewing crowds as organisations, algorithms coordinate the automated hiring of workers for different roles, and computationally structure their activities around complex workflows (Retelny et al. 2014, Kim et al. 2014; Valentine et al. 2017). Other types of algorithms, focusing more on teamwork, computationally rotate workers in different team combinations, to mix their viewpoints and ideas (Salehi and Bernstein, 2018).

The problem with existing crowd management algorithms, is that they tend to **micro-manage the workers**, by assigning them directly on a specific task or team. Existing algorithms also tend to focus on computational efficiency and optimisation. This approach is appropriate for microtasking, but it has drawbacks when it comes to macrotasks, as it can stifle creativity and initiative-taking, as indicated by recent research in management sciences (Lawler 2006) and crowdsourcing (Retelny et al. 2017). Future research is therefore needed to explore flexible algorithms that avoid micromanaging the workers, and explore ways to empower them.

Furthermore on crowd management, current crowdsourcing platforms have usually two management levels, i.e. the requester and the worker. Very recent works, indicate that new, multi-level ways of organisation, such as re-outsourcing (Wood et al. 2019) and subcontracting (Morris et al. 2017), and Upwork’s agency structures are emerging. Although the above works are applied on microtasking and freelance work, the multi-level management approach that they propose could be especially beneficial for the needs of macrotasking (see macrotask types 2, 3, and 4 above). Future research could explore this dimension.

A final note on crowd management is incentives engineering. Current microtasking crowdsourcing primarily relies on monetary rewards. Prior research in this domain has shown that higher payment indeed leads to faster completion time of the microtasks, but not necessarily to higher quality (Mason & Watts 2009). Initial research shows that purely extrinsic motivators, such as money, are not enough (Zheng et al. 2011). Macrotasking, which often involves open-ended and innovation- oriented work, and which for this reason relies on workers’ creativity and expertise, needs to find the right balance between extrinsic and intrinsic incentives. Earlier studies have offered “implications for the design of mobile workforce services, including future services that do not necessarily rely on monetary compensation” (Teodoro et al. 2014). For this reason, further work is needed to explore which intrinsic incentives platforms could offer to motivate quality macrotask work, with some examples including providing work feedback, and scaffolding workers’ career growth (Edmondson et al. 2001). A least, but absolutely more trivial point, to ensure that this future research has a practical impact, is the need for platform to raise awareness and educate requesters as to the need and concrete means of offering such incentives.

Macrotask Crowdsourcing Definition

Taking into account the aforementioned dimensions, on the nature of the task, the skills of the workers, and the management principles, we provide below a first integrated definition of macrotask crowdsourcing:

> *"Macrotask crowdsourcing refers to crowdsourcing that is designed to handle complex work of different degrees of structure and decomposability, assumes varying levels of (expert) knowledge over one or more domains, requires a range of 21st century skills, benefits from worker communication, collaboration, and training, and incorporates flexible work management processes that potentially involve the workers."*

Conclusion

In this chapter we discuss macrotask crowdsourcing in terms of three dimensions: i) the complex *problems* this labor model can solve, ii) the worker *skills* it requires and iii) the *management structures* it benefits from. In regards to the first dimension, we define four types of macrotasks - modular, interlaced, wicked, and container. Each type can solve a different problem, based on two problem axes: decomposability and structure. Regarding the second dimension, we touched upon the workers skills required for macrotask crowdsourcing, emphasizing the need for skill diversity, fine-grained skill types, expert and 21st century skills, as well as for skill development and evaluation mechanisms. Finally, in regards to the third dimension, we discussed the work management structures that are appropriate for this new type of work, highlighting the need to avoid micromanaging the workers but rather providing them with more initiative and actively involving them in the management of their work. We conclude this chapter with a definition, for the first time, of macrotask crowdsourcing. Our aim in providing this definition is to assist future researchers to better position their work, and inspire future developments on this expanding field.

References

Basu Roy, S., Lykourentzou, I., Thirumuruganathan, S., Amer-Yahia, S., & Das, G. (2015). Task assignment optimization in knowledge-intensive crowdsourcing. *The VLDB Journal—The International Journal on Very Large Data Bases*, *24*(4), 467-491.

Cheng, J., Teevan, J., Iqbal, S. T., & Bernstein, M. S. (2015, April). Break it down: A comparison of macro-and microtasks. In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (pp. 4061-4064). ACM.

Creative and Cultural Skills. 2017. Building a Creative Nation: Current and future skills need [online] Available from: https://ccskills.org.uk/downloads/Building_a_Creative_Nation_-_Current_and_Future_Skills_Needs.pdf

Edmondson, A. C., Bohmer, R. M., & Pisano, G. P. (2001). Disrupted routines: Team learning and new technology implementation in hospitals. Administrative Science Quarterly, 46(4), 685-716.

Gray, M. L., Suri, S., Ali, S. S., & Kulkarni, D. (2016, February). The crowd is a collaborative network. In *Proceedings of the 19th ACM conference on computer-supported cooperative work & social computing* (pp. 134-147). ACM.

Grier, D. A. (2013). Crowdsourcing for dummies. John Wiley & Sons.

Haas, D., Ansel, J., Gu, L., & Marcus, A. (2015). Argonaut: macrotask crowdsourcing for complex data processing. *Proceedings of the VLDB Endowment*, *8*(12), 1642-1653.

Ho, C. J., & Vaughan, J. W. (2012, July). Online task assignment in crowdsourcing markets. In Twenty-Sixth AAAI Conference on Artificial Intelligence

Ipeirotis, P. G. (2010). Analyzing the amazon mechanical turk marketplace. *XRDS: Crossroads, The ACM Magazine for Students, Forthcoming*.

Ipeirotis, P. G., & Gabrilovich, E. (2014, April). Quizz: targeted crowdsourcing with a billion (potential) users. In Proceedings of the 23rd international conference on World wide web (pp. 143-154). ACM.

Irani, L. C., & Silberman, M. (2013, April). Turkopticon: Interrupting worker invisibility in amazon mechanical turk. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 611-620). ACM.

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. CSCW ’14, ACM, New York, NY, USA(2014). https://doi.org/10.1145/2531602.2531638,http://doi.acm.org/10.1145/2531602.2531638

Lawler, E. E., & Worley, C. G. (2006). Designing organizations that are built to change. MIT Sloan Management Review, 48(1), 19–23.

Li, G., Wang, J., Zheng, Y., & Franklin, M. J. (2016). Crowdsourced data management: A survey. *IEEE Transactions on Knowledge and Data Engineering*, *28*(9), 2296-2319.

Machado, L., Pereira, G., Prikladnicki, R., Carmel, E., & de Souza, C. R. (2014, November). Crowdsourcing in the Brazilian IT industry: what we know and what we don't know. In Proceedings of the 1st International Workshop on Crowd-based Software Development Methods and Technologies (pp. 7-12). ACM.

Martin, D., Hanrahan, B. V., O'Neill, J., & Gupta, N. (2014, February). Being a turker. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing* (pp. 224-235). ACM.

Mason, W., & Watts, D. J. (2009, June). Financial incentives and the performance of crowds. In Proceedings of the ACM SIGKDD workshop on human computation (pp. 77-85). ACM.

Mavridis, P., Gross-Amblard, D., & Miklós, Z. (2016, April). Using hierarchical skills for optimized task assignment in knowledge-intensive crowdsourcing. In *Proceedings of the 25th International Conference on World Wide Web* (pp. 843-853). International World Wide Web Conferences Steering Committee.

Morris, M. R., Bigham, J. P., Brewer, R., Bragg, J., Kulkarni, A., Li, J., & Savage, S. (2017, May). Subcontracting microwork. In *Proceedings of the 2017 CHI conference on human factors in computing systems* (pp. 1867-1876). ACM.

Retelny, D., Robaszkiewicz, S., To, A., Lasecki, W. S., Patel, J., Rahmati, N., ... & Bernstein, M. S. (2014, October). Expert crowdsourcing with flash teams. In *Proceedings of the 27th annual ACM symposium on User interface software and technology* (pp. 75-85). ACM.

Retelny, D., Bernstein, M. S., & Valentine, M. A. (2017). No workflow can ever be enough: How crowdsourcing workflows constrain complex work. *Proceedings of the ACM on Human-Computer Interaction*, *1*(CSCW), 89.

Salehi, N., & Bernstein, M. S. (2018). Hive: Collective Design Through Network Rotation. *Proceedings of the ACM on Human-Computer Interaction*, *2*(CSCW), 151.

Schmitz, H., & Lykourentzou, I. (2018). Online Sequencing of Non-Decomposable Macrotasks in Expert Crowdsourcing. ACM Transactions on Social Computing, 1(1), 1.

Teodoro, R., Ozturk, P., Naaman, M., Mason, W., & Lindqvist, J. (2014, February). The motivations and experiences of the on-demand mobile workforce. In Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing (pp. 236-247). ACM.

Valentine, M. A., Retelny, D., To, A., Rahmati, N., Doshi, T., & Bernstein, M. S. (2017, May). 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). ACM.

Voigt, Paul, and Axel Von dem Bussche. "The EU General Data Protection Regulation (GDPR)." *A Practical Guide, 1st Ed., Cham: Springer International Publishing* (2017).

Vondrick, C., Patterson, D., & Ramanan, D. (2013). Efficiently scaling up crowdsourced video annotation. International Journal of Computer Vision, 101(1), 184-204.

Wagner, T. (2014). *The global achievement gap: Why even our best schools don't teach the new survival skills our children need—and what we can do about it*. Basic Books.

Wood, A. J., Graham, M., Lehdonvirta, V., & Hjorth, I. (2019). Networked but Commodified: The (Dis) Embeddedness of Digital Labour in the Gig Economy. *Sociology*, 0038038519828906.

Xie, H., & Lui, J. C. (2018). Incentive mechanism and rating system design for crowdsourcing systems: Analysis, tradeoffs and inference. IEEE Transactions on Services Computing, 11(1), 90-102.

Zheng, H., Li, D., & Hou, W. (2011). Task design, motivation, and participation in crowdsourcing contests. International Journal of Electronic Commerce, 15(4), 57-88.