Human–Computer Interaction Series

Vassillis-Javed Khan
Konstantinos Papangelis
Ioanna Lykourentzou
Panos Markopoulos *Editors*

Macrotask Crowdsourcing

Engaging the Crowds to Address Complex Problems

Springer

Human–Computer Interaction Series

The Human–Computer Interaction Series, launched in 2004, publishes books that advance the science and technology of developing systems which are effective and satisfying for people in a wide variety of contexts. Titles focus on theoretical perspectives (such as formal approaches drawn from a variety of behavioural sciences), practical approaches (such as techniques for effectively integrating user needs in system development), and social issues (such as the determinants of utility, usability and acceptability).

HCI is a multidisciplinary field and focuses on the human aspects in the development of computer technology. As technology becomes increasingly more pervasive the need to take a human-centred approach in the design and development of computer-based systems becomes ever more important.

Titles published within the Human–Computer Interaction Series are included in Thomson Reuters' Book Citation Index, The DBLP Computer Science Bibliography and The HCI Bibliography.

More information about this series at http://www.springer.com/series/6033

Vassillis-Javed Khan · Konstantinos Papangelis ·
Ioanna Lykourentzou · Panos Markopoulos
Editors

Macrotask Crowdsourcing

Engaging the Crowds to Address Complex Problems

Springer

Editors
Vassillis-Javed Khan
Eindhoven University of Technology
Eindhoven, The Netherlands

Konstantinos Papangelis
Xi'an Jiaotong-Liverpool University
Suzhou, China

Ioanna Lykourentzou
Department of Information
and Computing Sciences
Utrecht University
Utrecht, The Netherlands

Panos Markopoulos
Department of Industrial Design
Eindhoven University of Technology
Eindhoven, The Netherlands

ISSN 1571-5035 ISSN 2524-4477 (electronic)
Human–Computer Interaction Series
ISBN 978-3-030-12333-8 ISBN 978-3-030-12334-5 (eBook)
https://doi.org/10.1007/978-3-030-12334-5

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface: Macrotask Crowdsourcing for Advancing the Crowd's Potential

Amazon's launch of Mechanical Turk (MTurk.com) in 2005 kickstarted a new socio-technical phenomenon and a new labor model—that of crowdsourcing. Nowadays, MTurk is just one of several crowdsourcing platforms (à Campo et al. 2018). Such platforms bring two groups of people together: people who request a certain task but lack the skill or the time or the human capital to complete it, aka the *task requesters*; and people who wish to work on such tasks typically for a monetary reward, aka the *crowdworkers*.

The introduction of Application Programming Interfaces (APIs) in crowdsourcing platforms made the process of requesting and completing tasks much easier. These APIs have enabled the emergence of new scientific fields which integrate human effort with computing systems. *Human computation* is one of these fields, which channels human intelligence through the use of computing systems to solve tasks that no known efficient algorithm can yet solve (Von Ahn 2008). *Collective Intelligence* is another neighboring field, which couples human and machine intelligence to solve complex problems which neither humans nor machines can solve on their own (Malone et al. 2010).

The adoption of such platforms from large numbers of both requesters and workers, and the introduction of APIs established crowdsourcing as a fertile ground for researchers. However, due to the fact that widely adopted platforms like MTurk only supported short and easy tasks, known as microtasks (from the Greek word *μικρός*, which means small) research studies so far and most industrial applications have primarily focused on microtask crowdsourcing. To a certain extent, this focus is rightful; microtask crowdsourcing has produced some very impressive results. Examples include labeling images for improving image search and web accessibility (Von Ahn & Dabbish 2004); editing documents for shortening and proofreading (Bernstein et al. 2010); captioning audio in real time for accessibility (Lasecki et al. 2012); getting feedback on articles (Kittur et al. 2008) and designs (Luther et al. 2014). Consequently, most industrial practitioners and researchers today, when thinking of crowdsourcing they automatically think of a large list of small, similar, homogeneous and relatively straightforward to complete tasks—i.e., microtasks.

But not all types of work can be accomplished by breaking them down to microtask level (Schmitz & Lykourentzou 2018). Such tasks are complex and would yield meaningless results if decomposed, because of the many interdependencies among knowledge domains that they entail, and the need to maintain the global context while working on them. One can think of writing a story (Kim et al. 2016), a news article, or defining a research methodology (Schmitz & Lykourentzou 2018). In juxtaposition to *microtasks*, this type of tasks are known as *macrotasks* (from the Greek word μακρός (makros) which means 'long, large'), and crowdsourcing research has just started to look into them (Haas et al. 2015, Cheng et al. 2015). Although ground-breaking, the aforecited research in macrotask crowdsourcing has primarily used the term to contrast microtask crowdsourcing and in regards to the size of the task at hand, not its complexity and properties; like decomposition.

Macrotask crowdsourcing can make a more significant impact and to generate more value compared to microtask crowdsourcing, because it directly contributes to solving more challenging problems of both social and economic nature. Furthermore, it also requires salient, lifelong learning skills of the future such as creativity and critical thinking. By primarily focusing on microtasks, we are unnecessarily limiting and underestimating the crowd's potential.

Given the increasing interest of the research community and the industry on what can the crowds achieve, this book is a first effort to underpin this new type of crowd labor model that macrotask crowdsourcing represents and to collect works, of both theory and practice, around this subject that have started to emerge. In addition to researchers and practitioners interested in the evolution of crowdsourcing, it is our hope that this book will also prove useful for researchers and practitioners who are skeptical in regards to what they currently think what crowdsourcing is and what it can accomplish.

We initiate the book with a chapter that aims to properly define the terms *macrotask* and *macrotask crowdsourcing*. The chapter takes into account prior work and relevant theory, and looks deeper into the nature of the task, of worker skills and of crowd labor management, to provide a concrete basis upon which future researchers and practitioners can build upon. The rest of the book is divided into three parts, which together cover a wide range of macrotask crowdsourcing topics: *Coordination and Cooperation*, *The Role of AI and Experts*, and *Macrotasking for Social Good*.

Part I: Coordination and Cooperation

In this first part, the book examines the role of coordination and cooperation in the context of macrotasking. Coordination, in the context of complex work, is not an evident feat. Beyond issues of different time zones, languages, and cultures (issues that might anyway arise in microtasking) the multiple knowledge interdependencies and interactions required among the different workers create novel coordination challenges for macrotasks.

The Chap. 2 aims to advance our understanding on exactly this topic. More specifically, this chapter reviews several popular theories of coordination, examines the current approaches to crowd coordination in the HCI and CSCW literature, and identifies literature shortcomings. Based on these findings, the authors then proceed in proposing a research agenda and design propositions for each of the recommended theories of coordination, thus advancing our understanding of which crowd coordination mechanisms to select when complex macrotask work in involved.

A topic close to crowd coordination is crowd control. Crowdsourcing controls are mechanisms to align crowd workers' actions with predefined standards to achieve a set of goals and objectives set by the task requester. In ordinary microtasking, it is usually enough to address issues of control indirectly through financial incentives. In macrotasking, however, where the task is often performed within groups, more fine-grained behavior influencing control mechanisms are necessary to ensure a successful completion of the macrotask. In Chap. 3, the authors aim to develop a better comprehension of the controls appropriate for macrotask crowdsourcing. To accomplish this, they present and discuss the literature on control theory, identify a series of gaps, and put forth a research agenda to address these shortcomings. The proposed research agenda focuses on understanding how to design controls that are more suitable for macrotasking and the implications that such controls have for future crowdsourcing organizations.

This part of the book ends with an exploration of cooperation among crowd workers. Cooperation is an issue of less importance for microtasking, where workers usually perform tasks individually, but of increasing importance in macrotasking, where workers interact more often. In Chap. 4, the authors aim to leverage cooperation possibilities to improve the data quality of deployed macrotasks. The authors analyze three use cases from the domain of situated crowdsourcing, and use the results of this analysis to propose the design of a novel situated crowdsourcing platform that can effectively support cooperation without alienating solo workers.

Part II: The Role of AI and Experts

The second part of the book examines the role that Artificial Intelligence and Experts play in accomplishing macrotasks. As tasks become more complex, and in order to maintain their quality and scalability, advanced AI is becoming a necessity to efficiently distribute work among expert and nonexpert workers, as well as computational systems. Chapter 5 sheds light on exactly this topic. Using as an example, the macrotask of supporting scientific research at scale, the authors review the state-of-the-art in the intersection of crowdsourcing and AI, and outline how crowd computing research can inform the development of intelligent crowd-powered systems that can efficiently support macrotasking processes.

Selecting suitable workers has always been an important issue, ever since microtask crowdsourcing emerged. This selection is even more important in macrotasking, where the macrotask may require different types and granularities of expertise. In Chap. 6, the authors aim to ensure that the most appropriate workers will participate in the available tasks of a macrotask crowdsourcing marketplace. The authors base their work presenting two novel preselection mechanisms that have been shown to be effective in microtask crowdsourcing, and then proceed to discuss how these mechanisms can be used within macrotasks.

In the dawning age of macrotask crowdsourcing, should experts feel threatened? In the final chapter of this part of the book, the authors of the Chap. 7 present a highly reflective work of how digital technology could allow wider participation whilst preserving the core values of academia. Crucially, they address the question: Is academic resistance to crowdsourcing an elitist fear of the unwashed, or justifiable wariness of incipient poor scholarship?

Part III: Macrotasking for Social Good

As with every technology, macrotask crowdsourcing should eventually bring a positive development to future generations. In this part of the book, we present three chapters that showcase the potential broad benefits that macrotask crowdsourcing could bring to societal challenges.

Changing behaviors is a well-known challenge both widely acknowledged in HCI as well as other scientific fields. In Chap. 8, the authors aim to address this challenge by studying the effects of the content, mode, and style of motivational messages in the context of behavior change. To accomplish this, they use crowdsourcing for collecting a large amount of data to form an accessible database of motivational messages. The authors then report findings on unsupervised explorations of the emotional expressiveness and sound quality (signal-to-noise ratio, SNR) of the crowdsourced motivational speech.

Providing appropriate feedback is a crucial part of the learning process in educational setting. In Chap. 9, the authors aim to investigate how to compliment academic feedback with crowdsourced feedback. To accomplish this, they (1) investigate complimenting academic feedback with "real world feedback" during a course on mobile development, using HCI methods and (2) report the costs and benefits that both staff and students should be aware of, when planning to apply such methods.

Recent disasters due to climate change have been, rightfully so, prominently presented in popular media channels. In Chap. 10, the authors compare and contrast how different online communities employ crowdsourcing to aid disaster response efforts. To accomplish this, they first interview members from Humanitarian OpenStreetMap (HOT) and Public Lab mapping communities. Based on these interviews, they employ OpenStreetMap Analytics and Social Network Analysis, and analyze community strategies and interface logistics involved in the work of both communities.

Eindhoven, The Netherlands — Vassillis-Javed Khan
Suzhou, China — Konstantinos Papangelis
Utrecht, The Netherlands — Ioanna Lykourentzou
Eindhoven, The Netherlands — Panos Markopoulos

References

à Campo, S., Khan, V. J., Papangelis, K., & Markopoulos, P. (2018). Community heuristics for user interface evaluation of crowdsourcing platforms. Future Generation Computer Systems; Vol. 95, pp. 775–789. https://doi.org/10.1016/j.future.2018.02.028.

Bernstein, M. S., Little, G., Miller, R. C., Hartmann, B., Ackerman, M. S., Karger, D. R., Crowell, D. & Panovich, K. (2010). Soylent: A word processor with a crowd inside. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology (UIST '10)* (pp. 313–322). ACM, New York, NY, USA.

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.

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.

Kim, J., Sterman, S., Cohen, A. A. B., & Bernstein, M. S. (2017). Mechanical novel: Crowdsourcing complex work through reflection and revision. In *Proceedings of the 2017 ACM Conference on Computer-supported Cooperative Work and Social Computing* (pp. 233–245). New York, NY: ACM.

Kittur, A., Chi, E. H., & Suh, B. (2008). Crowdsourcing user studies with mechanical turk. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)* (pp. 453–456). ACM, New York, NY, USA

Lasecki, W., Miller, C., Sadilek, A., Abumoussa, A., Borrello, D., Kushalnagar, R., & Bigham, J. (2012). Real-time captioning by groups of non-experts. In *Proceedings of the 25th Annual ACM Symposium on User interface Software and Rechnology* (pp. 23–34). ACM.

Luther, K., Pavel, A., Wu, W., Tolentino, J. L., Agrawala, M., Hartmann, B., & Dow, S. P. (2014). CrowdCrit: Crowdsourcing and aggregating visual design critique. In *Proceedings of the Companion Publication of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing* (pp. 21–24). ACM.

Malone, T. W., Laubacher, R., & Dellarocas, C. (2010). The collective intelligence genome. *MIT Sloan Management Review, 51*(3), 21–31.

Schmitz, H., & Lykourentzou, I. (2018). Online sequencing of non-decomposable macrotasks in expert crowdsourcing. *ACM Transactions on Social Computing, 1*(1), 1.

Von Ahn, L., & Dabbish, L. (2004). Labeling images with a computer game. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 319–326). ACM.

Von Ahn, L. (2008). Human computation. In *Proceedings of the 2008 IEEE 24th International Conference on Data Engineering* (pp. 1–2). IEEE Computer Society.

Contents

Chapter 1
Macrotask Crowdsourcing: An Integrated Definition

Ioanna Lykourentzou, Vassillis-Javed Khan, Konstantinos Papangelis and Panos Markopoulos

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 has 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.

1.1 Introduction

The distinction between microtasks and macrotasks was made relatively early on in the crowdsourcing literature. Grier (2013) emphasized 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

I. Lykourentzou (✉)
Department of Information and Computing Sciences, Utrecht University, Princetonplein 5, De Uithof, 3584 CC Utrecht, The Netherlands
e-mail: i.lykourentzou@uu.nl

V.-J. Khan · P. Markopoulos
Eindhoven University of Technology, Eindhoven, The Netherlands

K. Papangelis
Xi'an Jiaotong-Liverpool University, Suzhou, China
e-mail: K.papangelis@xjtlu.edu.cn

V.-J. Khan et al. (eds.), *Macrotask Crowdsourcing*,
Human–Computer Interaction Series, https://doi.org/10.1007/978-3-030-12334-5_1

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, like other authors after him, introduced macrotasks in juxtaposition to microtasks in terms of the magnitude of the task. These works go as far as to propose a checklist for defining a macrotask as follows: a macrotask is a task that can be carried out independently without support by the requester, which is simple to describe with clear criteria of completion, which has a clear and concrete deadline, and which requires special skills that 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 why it 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. 2013). 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. (2013) compared annotating video for a single object per crowdworker which they considered as a microtask to annotating a video segment for a whole set of objects which they considered to be a macrotask. They noted 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 labels of higher quality. Furthermore, errors in coding specific objects are distributed over different segments and handled by different coworkers, 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 (2013) 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 microtasks. 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 microtasks 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 it 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 that 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. when 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 microcontributions, there can be substantial benefits in recruiting 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. (2015) 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 or 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 out these, 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 prespecify 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 prespecifying 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: (1) 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 incentivizing workers. For example, personal preferences or biases may color assessments of solution quality. Xie and Lui (2018) propose an optimization approach for incentivizing workers to provide high-quality contributions and empirically evaluate the effectiveness and efficiency of their approach.

1.2 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. 1.1).

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 a strong degree of domain interaction. **Decomposability** measures whether the problem can be divided into subproblems, and the granularity that this division can reach. Decomposable problems can be broken down to separate subproblems, each drawing from distinct knowledge sets, which can be solved independently and with little communication or collaboration among problem solvers. Non-decomposable problems on

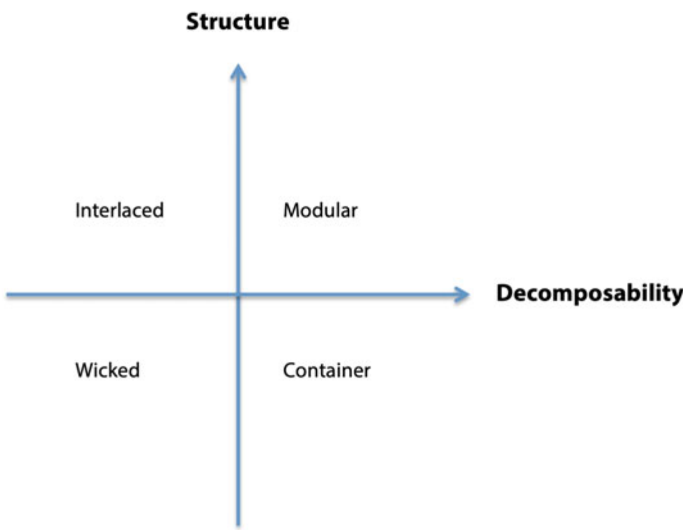

Fig. 1.1 The macrotask dimension space. To draw this diagram we assume all macrotasks are complex. Then we have a cartesian space of them with the dimensions of structure and decomposability. This space characterizes four types of macrotasks: interlaced, modular, wicked and container

the other hand, are impractical or even impossible to subdivide into separate subproblems, 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 a 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 optimal; 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.

1.2.1 Macrotask Type 1 (Modular): Well-Structured, High-Decomposability Problems

The first type of macrotasks 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 microtask level. Then, the distinct microtasks are assigned in parallel to multiple workers, and finally they are recomposed to a final output by combining the separate smaller subtasks.

The difference with what we might call "vanilla" microtasking is that, because of the problem complexity, the way of breaking down the problem to microtasks 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 suffers from non-generalization. Because the workflows are usually tailored to the very specific problem, they cannot be generalized easily to handle other problem instances.

The resulting microtasks 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).

1.2.2 Macrotask Type 2 (Interlaced): Well-Structured, Low-Decomposability Problems

The second type of macrotasks aims to tackle problems that are well-structured but 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 subtasks 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 would be: defining a research methodology or formulating an R&D approach.

1.2.3 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 handled through 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 illustrates 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.

An example of a type 3 macrotask is end-to-end innovation production.

1.2.4 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 organizational 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.

1.3 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 and 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 specialized 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 the platform). Microtasking platforms do not usually store other worker skills (Ho and Vaughan 2012). In case requesters need workers to have a specialized 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 specialized worker skills. This practice however is costly, and not generalizable.

In addition, microtasking usually relies on *skill homogeneousness*: the problem is decomposed to microtasks that all require the same type of nonexpert skill. Consequently, currently not a lot of works in existing crowdsourcing literature analyze 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 twenty-first-century skills, as well as valid skill identification and evaluation mechanisms. Examples of higher order cognitive and twenty-first-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 and Gabrilovich 2014), drawing from a wide range of approaches (from computerized 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 and 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-fertilization.

1.4 On the Nature of Management

When referring to crowdsourcing, scalability is the key. Unlike traditional management settings, where the human manager needs to organize 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 optimization 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 organizational perspective, viewing crowds as organizations, 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 optimization. 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 and Worley 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, multilevel ways of organization, 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 multilevel 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 and 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; examples might include: providing work feedback, and scaffolding workers' career growth (Edmondson et al. 2001). To ensure that this research will have practical impact, crowdsourcing platforms need to raise awareness and educate requesters about the importance of offering such incentives and support them in the process of doing so.

1.5 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.

1.6 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 worker skills required for macrotask crowdsourcing, emphasizing the need for skill diversity, fine-grained skill types, expert and twenty-first-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 in this expanding field.

References

Altshuller, G. (2005). *The innovation algorithm: TRIZ, systematic innovation and technical creativity*. Worchesrter, MA: Technical Innovation Center.

Ambati, V., Vogel, S., & Carbonell, J. (2012). Collaborative workflow for crowdsourcing translation. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (CSCW '12)* (pp. 1191–1194). ACM, New York, NY, USA.

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.

Bernstein, M. S., Little, G., Miller, R. C., Hartmann, B., Ackerman, M. S., Karger, D. R., Crowell, D. & Panovich, K. (2010). Soylent: A word processor with a crowd inside. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology (UIST '10)* (pp. 313–322). ACM, New York, NY, USA.

Chan, J., Dang, S., & Dow, S. P. (2016). Improving crowd innovation with expert facilitation. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)* (pp. 1223–1235). ACM, New York, NY, USA.

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.

Chilton, L. B., Little, G., Edge, D., Weld, D. S., & Landay, J. A. (2013). Cascade: Crowdsourcing taxonomy creation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)* (pp. 1999–2008). ACM, New York, NY, USA.

Creative and Cultural Skills. (2017). Building a creative nation: Current and future skills need. 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.

Goel, G., Nikzad, A., & Singla, A. (2014). Allocating tasks to workers with matching constraints: Truthful mechanisms for crowdsourcing markets. In *Proceedings of the Companion Publication of the 23rd International Conference on World Wide Web Companion (WWW Companion '14)* (pp. 279–280). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland.

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*. Wiley.

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*.

Huang, S., & Holden, D. (2016). The R&D boundaries of the firm: A problem solving perspective. *International Journal of the Economics of Business, 23*(3), 287–317.

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.

Jansson, D. G., & Smith, S. M. (1991). Design fixation. *Design Studies, 12*(1), 3–11.

Kim, J., Cheng, J., Bernstein, & M.S. (2014) 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. https://doi.org/10.1145/2531602.2531638, http://doi.acm.org/10.1145/2531602.2531638.
Kittur, A., Chi, E. H., & Suh, B. (2008). Crowdsourcing user studies with mechanical turk. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)* (pp. 453–456). ACM, New York, NY, USA
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.
Little, G. Chilton, L. B., Goldman, M., & Miller, R. C. (2010). Exploring Iterative and parallel human computation processes. In *Proceedings of the ACM SIGKDD Workshop on Human Computation (HCOMP '10)*. ACM, New York, NY, USA, 68–76.
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.
Majchrzak, A., & Malhotra, A. (2013). Towards an information systems perspective and research agenda on crowdsourcing for innovation. *The Journal of Strategic Information Systems, 22*(4), 257–268.
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.
Musthag, M., & Ganesan, D. (2013). Labor dynamics in a mobile micro-task market. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)* (pp. 641–650). ACM, New York, NY, USA.
Nickerson, J. A., & Zenger, T. R. (2004). A knowledge-based theory of the firm–the problem-solving perspective. *Organization Science, 15*(6), 617–632.
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. In *Proceedings of the ACM on Human-Computer Interaction*, *1*(CSCW), 89.
Salehi, N., & Bernstein, M. S. (2018). Hive: Collective design through network rotation. In *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.
Sieg, J. H, Wallin, M. W., & Von Krogh, G. (2010). Managerial challenges in open innovation: a study of innovation intermediation in the chemical industry. *R&D Management, 40*(3), 281–291.
Teevan, J., Iqbal, S. T., Cai, C. J., Bigham, J. P., Bernstein, M. S., & Gerber, E. M. (2016). Productivity decomposed: Getting big things done with microtasks. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '16)* (pp. 3500–3507). ACM, New York, NY, USA.

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, P., & Von dem Bussche, A. (2017). The EU General Data Protection Regulation (GDPR). In *A practical guide*, 1st edn. Cham: Springer.

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.

Zhang, H., Law, E., Miller, R., Gajos, K., Parkes, D., & Horvitz, E. (2012). Human computation tasks with global constraints. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)* (pp. 217–226). ACM, New York, NY, USA.

Zaidan, O. F., & Callison-Burch, C. (2011). Crowdsourcing translation: Professional quality from non-professionals. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1 (HLT '11)* (pp. 1220–1229). Association for Computational Linguistics, Stroudsburg, PA, USA.

Part I
Coordination and Cooperation

Chapter 2
Crowdsourcing Coordination: A Review and Research Agenda for Crowdsourcing Coordination Used for Macro-tasks

Sangmi Kim and Lionel P. Robert Jr.

Abstract Crowdsourcing has become a widely accepted approach to leveraging the skills and expertise of others to accomplish work. Despite the potential of crowdsourcing to tackle complex problems, it has often been used to address simple micro-tasks. To tackle more complex macro-tasks, more attention is needed to better comprehend crowd coordination. Crowd coordination is defined as the synchronization of crowd workers in an attempt to direct and align their efforts in pursuit of a shared goal. The goal of this chapter is to advance our understanding of crowd coordination to tackle complex macro-tasks. To accomplish this, we have three objectives. First, we review popular theories of coordination. Second, we examine the current approaches to crowd coordination in the HCI and CSCW literature. Finally, the chapter identifies shortcomings in the literature and proposes a research agenda directed at advancing our understanding of crowd coordination needed to address complex macro-tasks.

2.1 Introduction

Crowdsourcing has become a widely accepted approach to leveraging the skills and expertise of others to accomplish work (Robert and Romero 2015, 2017). Crowdsourcing has many definitions but was first defined by Jeff Howe as the outsourcing of work to a crowd (Howe 2006). Typical modern definitions of crowdsourcing involve two attributes: (1) a crowd, or group of people, and (2) online work. Crowdsourcing platforms such as Mechanical Turk (http://www.mturk.com) and CrowdFlower (http://www.crowdflower.com) attract large groups of people who can work online via these digital platforms. These platforms and the people who work on them (i.e., crowd workers) provide access to a wealth of knowledge and expertise that can be leveraged to tackle complex problems.

S. Kim (✉) · L. P. Robert Jr.
University of Michigan School of Information, Ann Arbor, USA
e-mail: sangmik@umich.edu

L. P. Robert Jr.
e-mail: lprobert@umich.edu

V.-J. Khan et al. (eds.), *Macrotask Crowdsourcing*,
Human–Computer Interaction Series, https://doi.org/10.1007/978-3-030-12334-5_2

Despite the potential of crowdsourcing to tackle complex problems, it has often been used to address rather simple micro-tasks. Micro-tasks are standalone simple tasks that do not require the coordination of work among individuals (Schmitz and Lykourentzou 2018). To tackle more complex problems, crowdsourcing must address macro-tasking. Macro-tasking can be described as complex crowd work that is sometimes but not always decomposable to micro-tasks (Schmitz and Lykourentzou 2018). Crowdsourcing macro-tasks is more challenging than crowdsourcing micro-tasks. Macro-tasking requires work processes needed to tackle complex problem-solving involving activities such as the generation and integration of diverse ideas along with group decision-making. Macro-tasking also requires crowd workers to coordinate in order to both divide their labor and aggregate the outputs of their labor.

In the human–computer interaction/computer-supported cooperative work (HCI/CSCW) fields, crowd coordination is typically handled by the requestor and results in micro-tasking. Requestors divide and assign work prior to any crowd involvement and in many cases the work is never aggregated. Unfortunately, this approach to crowd coordination limits the potential of crowds to solve complex problems and reach their full potential.

Consider the following scenario: An organization wants to use crowdsourcing to identify its next new product. The organization puts forth a call to the public for new ideas and gives a specific deadline. The organization receives many great ideas and asks the crowd to vote on the best idea for a new product. The votes are tallied and the winner is announced. This approach to crowdsourcing is oriented toward micro-tasking. The work process is reasonably well formulated and easy to understand by all crowd workers. Although the outcome might not be predictable, the work process is very predictable. The crowdsourcing tasks require little interaction or dependence among crowd workers, so coordination is of little importance.

Now consider a different scenario: An organization wants to crowdsource the development of the marketing plan for this new product. Because there are many ways to accomplish this task, the work is not easily nor reasonably well formulated. Both the work process and the outcome are not as predictable as in the last scenario. Because the crowd is expected to produce one marketing plan, the crowd workers must decide how the work is to be divided and how or whether the work needs to be aggregated. To accomplish this task, crowd workers need to work together. This approach to crowdsourcing is oriented toward macro-tasking and requires interaction and greater dependence among crowd workers; therefore, coordination is of the utmost importance. Clearly, to fully leverage crowdsourcing, more work is needed on coordinating the crowdsourcing of macro-tasks.

There are many definitions of coordination (Robert 2016). For the sake of clarity, this chapter defines coordination generally as:

> The synchronization of individuals in an attempt to direct and align their efforts in pursuit of a shared goal.

And crowd coordination specifically as:

> The synchronization of crowd workers in an attempt to direct and align their efforts in pursuit of a shared goal.