Development of a Crowd-Powered System Architecture for Knowledge Discovery in Scientific Domains

António Correia[*1,2], Benjamim Fonseca[1], Hugo Paredes[1], Daniel Schneider[3], and Shoaib Jameel[2]
[1] INESC TEC and University of Trás-os-Montes e Alto Douro, Vila Real, Portugal
[2] University of Kent, School of Computing, Medway Campus, UK
[3] Tércio Pacitti Institute of Computer Applications and Research (NCE), Federal University of Rio de Janeiro, Brazil
[*] Corresponding author. Email address: antonio.g.correia@inesctec.pt

***Abstract*—A substantial amount of work is often overlooked due to the exponential rate of growth in global scientific output across all disciplines. Current approaches for addressing this issue are usually limited in scope and often restrict the possibility of obtaining multidisciplinary views in practice. To tackle this problem, researchers can now leverage an ecosystem of citizens, volunteers and crowd workers to perform complex tasks that are either difficult for humans and machines to solve alone. Motivated by the idea that human crowds and computer algorithms have complementary strengths, we present an approach where the machine will learn from crowd behavior in an iterative way. This approach is embodied in the architecture of SciCrowd, a crowd-powered human-machine hybrid system designed to improve the analysis and processing of large amounts of publication records. To validate the proposal's feasibility, a prototype was developed and an initial evaluation was conducted to measure its robustness and reliability. We conclude this paper with a set of implications for design.**

Keywords—collaborative conceptual modeling; crowd-machine hybrid interaction; crowd science; human-centered AI; massively collaborative science; scientific knowledge discovery

I. Introduction

Current scientific work practices are characterized by large amounts of data [1], resulting in a vast set of neglected research domains and an increasing demand for efficient ways to process dynamic information from multiple data sources. The initial trigger for this work relies on the daunting task of keeping up-to-date about the millions of scientific papers that are published every year [21]. Over the last decades, researchers have put a substantial amount of effort into data-intensive research practices that range from the collection of scientific records to the analysis of multidimensional data fields. Such processes cannot be fully automated due to the semantic limitations of technologies for handling the heterogeneity of datasets. Thus, managing dynamic situations in scientific knowledge discovery is a recurrent challenge for software practitioners. Seen from a socio-technical perspective, the role of software as an integral component of the work in science "is not commonly an object of inquiry in studies of scientific infrastructures" [3]. This creates a vast set of opportunities that come with novel challenges that require new problem formulations and methods for analyzing complex knowledge representations.

According to Gregory et al. [15], "open research data are heralded as having the potential to increase effectiveness, productivity and reproducibility in science". In connection with this aspect, crowd science provides a new framework to operate under uncertainty in research settings [4]. The term 'crowd science' was used by Sauermann and Franzoni [10] in order to define "scientific research performed with the involvement of the broader public (the crowd)". In this regard, scientific work can be partly or entirely conducted by volunteer (unpaid) amateur scientists intended to provide meaningful insights comparable to those produced by scientific experts. Nevertheless, there is a need to understand crowd work and how to engineer crowd-enabled systems, creating new paths for improving research by processing large amounts of data without spatial or temporal barriers. This challenged the authors to consider the socio-technical aspects of designing a tool and architecture to mobilize crowds for supporting research tasks in global ways.

The main contribution of this paper lies in the concrete implementation of the SciCrowd system's architecture, including an initial evaluation to demonstrate the performance of the proposed platform. The work presented here contributes to the field of crowd-powered Human-Computer Interaction (HCI) research by addressing situated scientific practices through the contribution of online crowds. In this sense, we introduce a crowd-enabled architecture that uses an open participation model in which researchers and crowd members can contribute for improving the way in which scientific data are filtered and processed for achieving new discoveries.

Unlike other crowdsourcing frameworks, SciCrowd relies on faceted search and mass collaboration mechanisms as central tenets. For example, researchers can add or validate metadata, datasets, experiments, and any other form of research outcomes for further interpretation and reuse by the crowd. We believe that this provides a more diversified interpretation of scientific results through meaningful categories, filters, and observations generated through the combination of long-term volunteer participation and machine intelligence [9]. In this paper, we assume that the integration of crowd-AI hybrids can be an important instrument by enabling us to identify new relations between topics, authors, groups, among other entities.

II. RELATED WORK

Crowdsourcing systems have become popular in the software industry and offer advantages such as the lower cost and shorter execution time of complex tasks [5]. As argued by Weiss [7], crowdsourcing systems differ in terms of task complexity and type of tasks assigned to the crowd, incentives, amount of time spent, and level of collaboration between members. As previously noted, the wide accessible networks of crowd participants have been harnessed as an efficient approach to aid scientific practices [11]. Crowd-powered systems can help answer research questions while supporting various stages of the scientific process [4]. For instance, academic crowds have been leveraged for performing tasks such as organizing conference sessions while aiding the extraction of categories and clusters from high-dimensional data [13]. Crowdsourcing has been also applied in article evaluation using bookmarks of journal articles [14]. However, there is a lack of research on better software tools intended to support research tasks though crowdsourcing, as current systems do not provide easy access and dynamic task generation [16].

To fill this gap, researchers have coupled the outputs of humans with machine learning algorithms to improve academic knowledge discovery [19]. As a consequence, there has been an increasing interest in designing human-AI interfaces [17] and mixed-initiative systems [18] as iterative, intelligent approaches that combine the strengths of human interaction with the algorithmic power of AI to solve problems that could not be solved by either algorithms or crowds alone. In this scenario, humans can process, filter, classify, or simply validate machine-extracted data in order to provide evidence on demand using automated reasoning techniques [11]. As more scholars have become aware of hybrid algorithmic-crowd enabled systems, there has been some concern about the main limitations that impede us from realizing the great promise of IT-enabled crowd science. In view of this fact, there is a lack of systematic methods to handle false observations and low quality outputs. Palacin-Silva and Porras [20] go even further by identifying concerns related to the limited expertise of participants, standardization, data aggregation, and privacy. Following this line of thought, a key challenge for a system exploring scholarly data by means of human-AI interaction relies on attracting and sustaining participation over time.

III. SYSTEM DESCRIPTION

SciCrowd is a crowd-powered human-machine hybrid system for collecting and aggregating metadata from papers. The system combines crowdsourced human expertise and automated indexing while allowing users to label publication records and annotate them with contextual information. With this system we aim to find associations that were not explored before by means of shared conceptualization and collaborative conceptual modeling. In particular, each user can link entities/semantic resources (e.g., method, sample, findings, concepts) and explore their ramifications. Such environment also enables scientists to work collaboratively with crowd workers and volunteers in a distributed basis in order to extract relevant scientific facts locked within papers, a task that is usually insurmountable for researchers. Tchoua et al. [12] go even further by arguing that this issue "hinders the advancement of science" in several ways. According to the authors, this kind of system can be particularly useful to interpret data that is not machine accessible while building on existing findings and avoiding duplicated efforts.

A. System Design

1) Database

SciCrowd's initial reference database contains records crawled from DBLP[1] using a simple XML query API. Our system then gathers and aggregates the general outputs provided by the crowd. Given the demonstrated value of all the crowd-AI techniques in other domains, the search parameters are evolved based on the data granules provided by users. This is based on the evidence that looking at a more accurate candidate pool of contributions (e.g., labels, tags) reduces the number of false positives [22]. The data repository is responsible for facilitating data reuse by allowing researchers to store, manage, and make data available for other users [23]. In addition, we are interested in developing a dynamic indexing mechanism enriched with bibliometrics (including citation count and in-text citation contexts [8]) and alternative metrics such as downloads and impact of a certain paper on social media.

2) Data Analysis and Primary Sources

The application is deployed to fulfill the requirements of each research community by involving users as part of the entire cycle of system development. The crowd is encouraged to provide original labels and observations to supplement the original sources. For instance, a user begins the classification process by adding metadata about a publication record. Such metadata provide insights into the aspects reported in the record for improving the search filters. In other words, the inputs provided by users enrich the database for potentially identifying future actions. In its current form, the prototype comprises a limited set of features such as editing a publication record automatically extracted or added manually by users. Publication details can be visualized by pressing the "show" button, which opens an internal page that enables the user to insert data about a publication entry. Another feature relies on annotating excerpts of papers as *subtasks*, as done in other systems such as CommentSpace [26].

3) Faceted Search and Iterative Concept Discovery

The key motivating concept behind SciCrowd relies on providing researchers with a way to support their scientific knowledge discovery practices through a faceted search interface. These kinds of *search filters* allow users to filter results and retrieving important data by selecting details that would otherwise be unknown. The interface also scaffolds domain expertise to prevent users from applying search filters that might contradict each other. We aim to create a knowledge base for tracking the evolution of concepts through structured labeling, crowd-AI integration, and

[1] https://dblp.uni-trier.de/

iterative visualization. In this sense, SciCrowd uses an interface that enables a user to specify and refine search topics and taxonomies in a dynamic way through a tailored explanation of search results. This stresses the importance of deploying mechanisms to filter, discover, and define concepts and data relationships [24].

4) Human Intelligence Tasks

Based on prior work on crowd-AI hybrids for scientific research [9], several components are being developed for supporting Human Intelligence Tasks (HITs). Such tasks are created by users and specified in the system as actions with a start and end date. The tasks will be available to be assigned to each crowd worker and the information about the actions performed by users will be updated accordingly. This includes a dashboard with contextual information about the progress of the tasks (e.g., percentage of the completed activities). The *task manager* guides the requester on the orchestration of the set of tasks to be executed and to assign them to crowd workers. Furthermore, a *consensus manager* examines the level of agreement over the crowd outputs, while the *reward manager* is responsible for updating information about the worker reputation and trustworthiness.

5) Quality Assessment

In software engineering, quality attributes such as reliability and performance have significant effects on the design of software systems [28]. Crowdsourcing can be expensive and unreliable, providing bad data and faulty observations resulting from crowd bias [29]. According to Dickinson and co-workers [32], misclassification and individual subjectivity are serious problems that can be mitigated through the aggregation of annotation data provided by crowd members for the same subject. We are investigating new forms of measuring performance indicators such as efficiency, quality, and success rate supported by a crowd-powered mechanism for error analysis [30]. The main idea relies on enabling crowd members to submit their asymmetric contributions to a peer assessment process (e.g., single/double blind, open peer-review). Likewise, we should not underestimate quality control mechanisms such as majority voting and ground truth [27].

6) Collaboration Functionalities

The cooperation model is responsible for aiding the execution of tasks, while the coordination between crowd workers and requesters is ensured via a collaborative environment where users registered in the system can contribute to the progress of scientific discovery. As argued by Xia and colleagues [36], this stigmergic form of crowd work "allows participants to build on one another's contributions without explicit coordination among collaborators or the division of labor into discrete steps and responsibilities." Perceiving the presence of other crowd members is also critical for diminishing the sense of distance and impersonality. More than that, the lack of awareness (a concept largely explored in the field of HCI as the understanding of the activities and individual contributions of other members in a cooperative work effort) can result in errors and misunderstandings while affecting the flow and naturalness of work [33]. In this sense, SciCrowd must be able to show the contributions provided by other users and provide targeted requests taking into account the different needs of intervention. In other words, a contributor must be aware of new papers in the system, see who is online in the platform, track the completion of a scientific workflow, and receive notifications about activity progress.

7) Hybrid Classifiers

At the start of the development of SciCrowd, it was recognized that we need an algorithm trained according to the crowd behavior to generate new hypotheses. In this scenario, a member of the crowd can validate the quality of the data processed by the machine, identify overlooked data regions, and provide explanations in order to justify each decision [35]. The proposed system's workflow is divided into three phases. In the first phase, requesters (e.g., researchers, institutions) design HITs that motivate crowd workers to participate. A requester who wants to register tasks into SciCrowd writes a project description. At this stage, the required skills are specified and the system coordinates the roles and relationships of each user within the research projects in which they are involved. In the second phase, crowd members can contribute to each project through the identification of relevant sources, data collection, analysis, and quality assessment. In the third (result-oriented) phase, the study results are generated and available for (re)use.

B. Scenario

We describe a scenario in which a user can login in the system and classify a publication record with different levels of granularity in order to make data reusable in further analyses [6]. Suppose that the user intends to perform a scientometric analysis of IEEE SMC conference proceedings from 1989 to 2018. After the user makes the choice for which type of task he/she wants to perform, an interface is presented for him/her to select the data items that will be considered in the analysis. The user can then edit the metadata of each included record and/or ask the crowd for intervention by publishing parts of the analysis as microtasks. For example, a paper can be classified taking into account the method(s) used, system features, concepts, and main findings, as shown in Figure 1.

Such data granules are particularly useful when a user is interested in answering research questions such as "What are the best practices for designing a crowdsourcing system intended to support users with cognitive impairment?" In order to answer this, a user can explore some concepts, observations, and semantic links obtained from the crowd. Moreover, a user can also provide his/her own contribution by logging into the system. An overview of the complete list of published tasks is then presented to the user. In addition, he/she is able to receive suggestions of other artifacts to be processed and each member is able to see the additions in response to his/her behavior. This gives users more opportunities to engage in others' insights and perspectives. Finally, the crowd worker finishes a task and submits the candidate answer according to the task requirements. The requester may now validate the answers based on the suggestion provided by the system, and a crowd worker can then get the reward credits for his/her contribution.

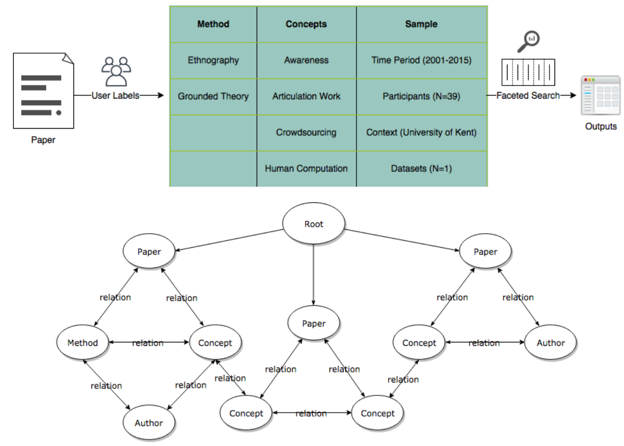

Figure 1. SciCrowd's basic workflow and ontological concept scheme.

C. Architecture

Figure 2 depicts an overview of the SciCrowd system's architecture and the relationships involved in performing crowd-powered research tasks. The 3-layer architecture of our system is composed by *Client*, *Server*, and *Database Abstraction* layers. The Client layer provides the interface and endpoints to standardize the access to the service layer. The Server layer provides a controlled access point to the other two layers in the architecture. This layer implements all services that are consumed by the Client layer and can be used to track the workflow execution. Moreover, a Database Abstraction layer is used to store the outputs from the crowd.

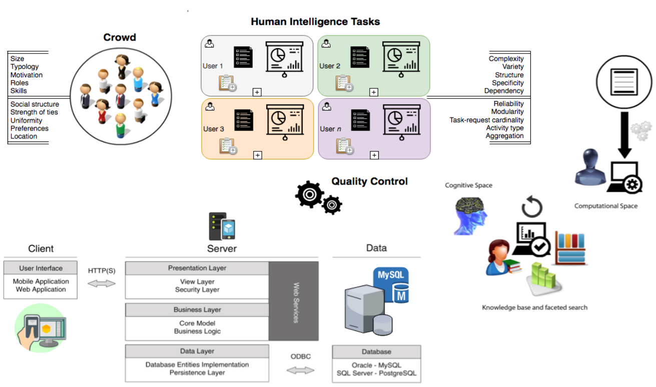

Figure 2. System architecture.

The research activities can be further decomposed into microtasks. However, non-decomposable macrotasks require high levels of coordination. At this stage, the system is only able to support microtasks such as classify a paper according to the categories of a taxonomy. The HITs vary in terms of complexity and modularity, structure, and dependencies. The crowd workforce is influenced by factors such as size and domain expertise, and a key challenge relies on motivating crowd workers. In order to be able to do this, each user is rewarded through gamification elements (badges, reputation points). Newcomers need to be supported with different levels of engagement and guided by seasoned users taking into account their roles in the community [25]. As argued by Sieber and Slonosky [31], such practices "provide crucial insight into the design and user experience of the system and to induce a sense of ownership of the project".

D. Implementation Details

The implementation refers to the design, installation and configuration of SciCrowd. The system is built using the PHP language and the Symfony framework[2]. The web app is hosted on a Linux server running Apache, and the metadata is stored in a MySQL relational database. Figure 3 presents a conceptual database schema representing the core structure (publication entity). The implemented classes are related to the publication metadata. This includes institution, authors, publisher, type of publication, ID, country, and additional (open) labels that constitute the folksonomy as a (free-form) classification scheme that enables each user to provide open categories into the system [14].

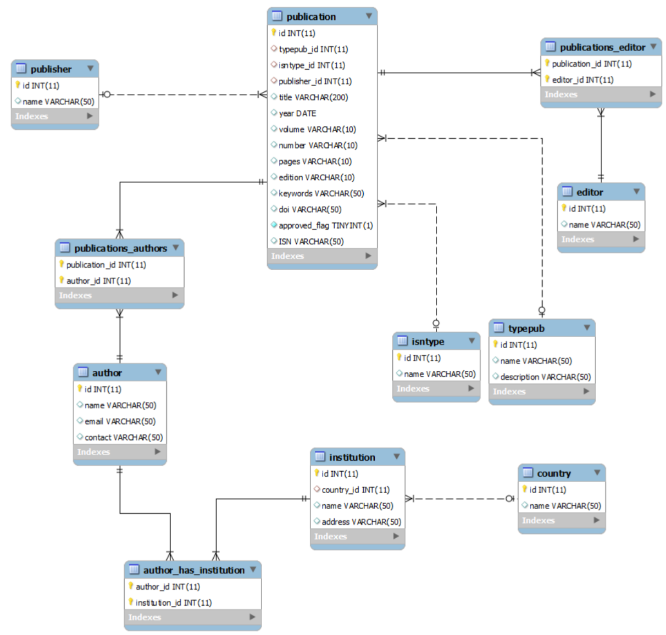

Figure 3. Excerpt of SciCrowd metamodel.

An authentication mechanism ensures that only authorized users can enter into the system and insert or modify content using different roles. Such authentication levels are divided into groups (common user, moderator, administrator) and permissions such as edit, delete, and add data. When designing the authentication mechanism, it was necessary to create three classes: *Users*, *Roles*, and *Groups*. These classes are responsible for storing information about each user (or group) and their permissions. The connection between classes is as follows: a user belongs to one or more groups and these groups have one or more permissions.

E. Runtime Analysis

The SciCrowd platform tests consisted of a performance evaluation to measure its efficiency. The performance tests were conducted to verify the response time between the

[2] https://symfony.com/

request and the reply, and later presentation of the data in the database. Two illustrative graphs are presented (Figure 4). The first graph (a) shows the system behavior in processing simple queries involving few or no connections between tables, while the second graph (b) involves complex queries.

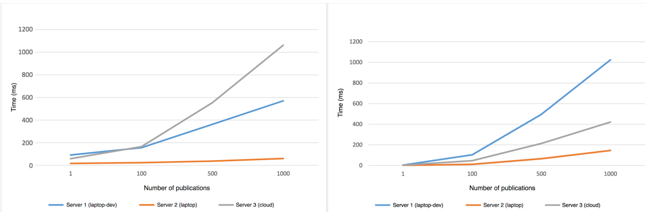

(a) Response time for *n* publications (simple query) (b) Response time for *n* publications (concurrent queries)

Machine	Server 1 (laptop-dev)	Server 2 (laptop)	Server 3 (cloud)
Processor	Intel® Core™ i5 CPU M 480	Intel® Core™ T7200	AMD Opteron
Memory	8GB	2GB	2GB
Hard Disk	WD 320GB 7200	WD 80GB 5400	-
Bundle	XAMPP	Isolated	XAMPP
Operating System	Windows 8	Linux Debian 7	Windows 2008 Server

Figure 4. Performance of the SciCrowd's system prototype.

To carry out the tests on three different machines, it was necessary to implement methods to insert data in the database. The system was installed using a XAMPP Apache distribution following the standard installation procedure with all available components. The number of authors was determined by the system (in a maximum of five per paper). Moreover, it is worth noting that the underlying hardware type also influenced the efficiency of the system.

A closer look at the results further highlight lower levels for Server 1 and Server 3 compared to Server 2. As outlined above, these servers were executing on a computer with an user interface and hosted in a bundle. The analysis of the two types of tests accentuates the differences in results between the other servers. It is important to note that the Debian Linux operating system without an user interface is faster and the results make the system of high performance and reliable for extracting data from publications. Our tests also suggest that the database query and presentation are both satisfactory. Given the characteristics of the computers where the tests were performed, it can be hypothesized a better performance in a more equipped and dedicated infrastructure. Furthermore, it should also be noted that these tests were carried out with several resources in parallel, such as print server, file server, and user account management.

IV. Implications for Design

Hybrid machine-crowd interaction may exploit the best of both human and machine abilities for performing complex research tasks. This line of work reinforces a recent interest in training algorithms and developing hybrid intelligence systems and architectures to solving difficult problems and coordinating work with high levels of structural complexity [34]. Although some studies explore how crowd-AI hybrids can be used to improve scientific discovery [12,19,21], they address only partially the problems posed by AI concerning its intrusive actions and unpredictable behaviors [37]. When studying crowd science and AI as an integrated system, there are also concerns related to the errors that can result from insufficient training data. Thus, the uncertainty that advent from computer decisions affect the reliability of the system as a whole. An important concept is what Sheridan and Verplank [38] called by 'levels of automation'. These levels must be negotiated and articulated taking into account the different needs of the environment. In addition, there are also challenges related to task modularity and coordination, lack of scientific domain expertise, and quality control [31].

Understanding science software and data-intensive socio-technical infrastructures can inform the design of human-centered systems by means of transferrable insights [6]. A lens into the contribution patterns of crowd-AI systems can avoid problems such as the highest amount of crowd work often performed by only a small portion of community members [2] and the system must be able to avoid cheating and manipulation of task outputs by crowd workers [27]. To mitigate possible errors, a crowd-powered system must be able to provide feedback and incorporate explanations about how crowd participants and machines make their observations and classifications [30]. Paine and colleagues [6] go even further by claiming that "understanding the context and process of the creation of datasets is necessary and important for researchers to be able to analyze, share, or reuse data in research work". The authors also argue that data must be transformed into a common format to be used in subsequent activities or as input to computational models.

V. Conclusion

In this paper we presented SciCrowd as a human-centered, crowd-powered system based on an open model for aiding discovery-based science. The solution is being developed through a series of observations retrieved from literature studies and a survey with experienced researchers with diverse background in the field of crowdsourcing. The system has addressed part of the main functionalities presented in previous section and the next stage of development will comprise a set of usability studies. We elaborate on some requirements identified previously to assume that SciCrowd will need to include quality control, aggregation, and reputation mechanisms. A missing piece relies on the development of a HIT-management and task assistant module to provide and manage tasks for users. In addition, a way to customize how the data is presented and visualized is also desirable. There is a need for data synchronization in the database and support for multi-dimensional data fields through an evolutionary taxonomy.

Another challenge is to design a computational model that learns from crowd behavior while providing a self-adapting environment in which the user can validate the outputs provided by the machine in a highly symbiotic process. However, some complexities (e.g., dependency, scalability) remain unsolved regarding the lack of machine-crowd integration when considering macrotasks. In this sense, there is a large opportunity for enhancing scientific work through crowd-AI hybrids and this has been under-explored due to the scarcity of large-scale deployments in research settings. Finally, releasing SciCrowd as an open source project is an ongoing debate that, although outside the scope of the present paper, would benefit the whole scientific community while harnessing crowd contributions for developing and maintaining the system code.

ACKNOWLEDGEMENTS

This work is financed by National Funds through FCT - Portuguese Foundation for Science and Technology, in the scope of the projects UID/EEA/50014/2019 and CMU/CS/0012/2017. The PhD thesis of António Correia is funded by FCT grant SFRH/BD/136211/2018. The authors also wish to thank Jorge Santos for doing the implementation work.

REFERENCES

[1] K. Li, J. Greenberg, and J. Dunic, "Data objects and documenting scientific processes: An analysis of data events in biodiversity data papers," *J Assoc Inf Sci Technol*, 2019.

[2] A. Basiri, M. Haklay, and Z. Gardner, *The impact of biases in the crowdsourced trajectories on the output of data mining processes*. in *Proc. of the AGILE Conference on Geo-information science*, 2018.

[3] D. Paine and C. P. Lee, "Who has plots?: Contextualizing scientific software, practice, and visualizations," in *Proceedings of the ACM Conf. Human-Computer Interaction*, CSCW, vol. 1, article 85, 2017.

[4] E. Law, K. Z. Gajos, A. Wiggins, M. L. Gray, and A. Williams, "Crowdsourcing as a tool for research: Implications of uncertainty," in *Proceedings of the ACM Conference on Computer Supported Cooperative Work and Social Computing*, 2017, pp. 1544-1561.

[5] M. Vukovic and C. Bartolini, "Crowd-driven processes: state of the art and research challenges," in *Proceedings of the International Conference on Service-Oriented Computing*, 2010, pp. 733-733.

[6] D. Paine, E. Sy, R. Piell, and C. P. Lee, "Examining data processing work as part of the scientific data lifecycle: Comparing practices across four scientific research groups," in *Proc. iConference*, 2015.

[7] M. Weiss, "Crowdsourcing literature reviews in new domains," *Technology Innovation Management Review*, 6, 2, 2016, pp. 5-14.

[8] M. Färber, A. Sampath, and A. Jatowt, "PaperHunter: A system for exploring papers and citation contexts," in *Proceedings of the European Conference on Information Retrieval*, 2019, pp. 246-250.

[9] A. Correia, D. Schneider, H. Paredes, and B. Fonseca, "SciCrowd: Towards a hybrid, crowd-computing system for supporting research groups in academic settings," in *Proc. of the International Conference on Collaboration and Technology*, 2018, pp. 34-41.

[10] H. Sauermann and C. Franzoni, "Crowd science user contribution patterns and their implications," *National Academy of Sciences*, vol. 112, no. 3, 2015, pp. 679-684.

[11] A. Correia, D. Schneider, B. Fonseca, and H. Paredes, "Crowdsourcing and massively collaborative science: Systematic literature review and mapping study," in *Proc. of the International Conference on Collaboration and Technology*, 2018, pp. 133-154.

[12] R. Tchoua, A. Ajith, Z. Hong, L. Ward, K. Chard, D. Audus, S. Patel, J. de Pablo, and I. Foster, "Towards hybrid human-machine scientific information extraction," *New York Scientific Data Summit*, 2018,1-3.

[13] P. André, H. Zhang, J. Kim, L. Chilton, S. P. Dow, and R. C. Miller, "Community clustering: Leveraging an academic crowd to form coherent conference sessions," in *Proceedings of the First AAAI Conference on Human Computation and Crowdsourcing*, 2013.

[14] I. Peters, S. Haustein, and J. Terliesner, "Crowdsourcing in article evaluation," in *Proceedings of the ACM International Conference on Web Science*, 2011.

[15] K. M. Gregory, H. Cousijn, P. Groth, A. Scharnhorst, and S. Wyatt, "Understanding data search as a socio-technical practice," *Journal of Information Science*, 2019.

[16] R. Englund, S. Kottravel, and T. Ropinski, "A crowdsourcing system for integrated and reproducible evaluation in scientific visualization," in *Proc. IEEE Pacific Visualization Symposium*, 2016, pp. 40-47.

[17] A. T. Nguyen, A. Kharosekar, S. Krishnan, S. Krishnan, E. Tate, B. C. Wallace, and M. Lease, "Believe it or not: Designing a human-AI partnership for mixed-initiative fact-checking," in *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, 2018, pp. 189-199.

[18] E. Horvitz, "Principles of mixed-initiative user interfaces," in *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*, 1999, pp. 159-166.

[19] Z. Dong, J. Lu, T. W. Ling, J. Fan, and Y. Chen, "Using hybrid algorithmic-crowdsourcing methods for academic knowledge acquisition," *Cluster Computing*, vol. 20, no. 4, 2017, pp. 3629-3641.

[20] M. Palacin-Silva and J. Porras, "Shut up and take my environmental data! A study on ICT enabled citizen science practices, participation approaches and challenges," *EPiC Series in Computing*, vol. 52, 2018, pp. 270-288.

[21] E. Krivosheev, F. Casati, and B. Benatallah, "Crowd-based multi-predicate screening of papers in literature reviews," *arXiv:1803.09814*, 2018.

[22] V. Mohanty, D. Thames, S. Mehta, and K. Luther, "Photo Sleuth: Combining human expertise and face recognition to identify historical portraits," in *Proceedings of the 24th International Conference on Intelligent User Interfaces*, 2019, pp. 547-557.

[23] Y. Kim and A. Yoon, "Scientists' data reuse behaviors: A multilevel analysis," *Journal of the Association for Information Science and Technology*, vol. 68, no. 12, 2017, pp. 2709-2719.

[24] J. C. Chang, N. Hahn, A. Perer, and A. Kittur, "SearchLens: Composing and capturing complex user interests for exploratory search," in *Proceedings of the 24th International Conference on Intelligent User Interfaces*, 2019, pp. 498-509.

[25] D. Gurzick and W. G. Lutters, "Towards a design theory for online communities," in *Proc. of the International Conf. on Design Science Research in Information Systems and Technology*, 2009, p. 11.

[26] W. Willett, J. Heer, J. Hellerstein, and M. Agrawala, "CommentSpace: Structured support for collaborative visual analysis," in *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*, 2011, pp. 3131-3140.

[27] F. Daniel, P. Kucherbaev, C. Cappiello, B. Benatallah, and M. Allahbakhsh, "Quality control in crowdsourcing: A survey of quality attributes, assessment techniques, and assurance actions," *ACM Computing Surveys*, vol. 51, no. 1, 2018.

[28] D. Ameller, M. Galster, P. Avgeriou, and X. Franch, "A survey on quality attributes in service-based systems," *Software Quality Journal*, vol. 24, no. 2, 2016, pp. 271-299.

[29] S. Greengard, "Following the crowd," *Communications of the ACM*, vol. 54, no. 2, 2011, pp. 20-22.

[30] S. Fu, A. Xu, X. Liu, H. Zhou, and R. Akkiraju, "Challenge AI mind: A crowd system for proactive AI testing," *arXiv:1810.09030*, 2018.

[31] R. Sieber and V. Slonosky, "Developing a flexible platform for crowdsourcing historical weather records," *Historical Methods: A Journal of Quantitative and Interdisciplinary History*, 2019, pp. 1-14.

[32] H. Dickinson, L. Fortson, C. Scarlata, M. Beck, and M. Walmsley, "Modeling with the crowd: Optimizing the human-machine partnership with Zooniverse," *arXiv:1903.07776*, 2019.

[33] I. Steinmacher, A. P. Chaves, and M. A. Gerosa, "Awareness support in distributed software development: A systematic review and mapping of the literature," *Computer Supported Cooperative Work*, vol. 22, nos. 2-3, 2013, pp. 113-158.

[34] J. Cheng and M. S. Bernstein, "Flock: Hybrid crowd-machine learning classifiers," in *Proc. of the ACM Conference on Computer Supported Cooperative Work and Social Computing*, 2015, pp. 600-611.

[35] J. Yang, T. Drake, A. Damianou, and Y. Maarek, "Leveraging crowdsourcing data for deep active learning an application: Learning intents in Alexa," in *Proceedings of the International World Wide Web Conference*, 2018, pp. 23-32.

[36] H. Xia, C. Østerlund, B. McKernan, J. Folkestad, P. Rossini, O. Boichak, J. Robinson, K. Kenski, R. Myers, B. A. Clegg, and J. Stromer-Galley, "TRACE: A stigmergic crowdsourcing platform for intelligence analysis," in *Proceedings of the Hawaii International Conference on System Sciences*, 2019, pp. 440-448.

[37] S. Amershi, D. Weld, M. Vorvoreanu, A. Fourney, B. Nushi, P. Collisson, J. Suh, S. Iqbal, P. N. Bennett, K. Inkpen, J. Teevan, R. Kikin-Gil, and E. Horvitz, "Guidelines for human-AI Interaction," in *Proc. ACM Conf. on Human Factors in Computing Systems*, 2019.

[38] T. B. Sheridan and W. L. Verplank, *Human and computer control of undersea teleoperators*. Massachusetts Institute of Technology Cambridge Man-Machine Systems Lab, 1978.