



OPEN ACCESS

EDITED AND REVIEWED BY

Claudia Wagner,
GESIS Leibniz Institute for the Social Sciences,
Germany

*CORRESPONDENCE

Jie Yang
✉ j.yang-3@tudelft.nl

SPECIALTY SECTION

This article was submitted to
Machine Learning and Artificial Intelligence,
a section of the journal
Frontiers in Artificial Intelligence

RECEIVED 07 February 2023

ACCEPTED 27 February 2023

PUBLISHED 15 March 2023

CITATION

Yang J, Bozzon A, Gadiraju U and Lease M
(2023) Editorial: Human-centered AI: Crowd
computing. *Front. Artif. Intell.* 6:1161006.
doi: 10.3389/frai.2023.1161006

COPYRIGHT

© 2023 Yang, Bozzon, Gadiraju and Lease. This
is an open-access article distributed under the
terms of the [Creative Commons Attribution
License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or
reproduction in other forums is permitted,
provided the original author(s) and the
copyright owner(s) are credited and that the
original publication in this journal is cited, in
accordance with accepted academic practice.
No use, distribution or reproduction is
permitted which does not comply with these
terms.

Editorial: Human-centered AI: Crowd computing

Jie Yang^{1*}, Alessandro Bozzon², Ujwal Gadiraju¹ and
Matthew Lease³

¹Web Information Systems, Delft University of Technology, Delft, Netherlands, ²Knowledge and
Intelligence Design, Delft University of Technology, Delft, Netherlands, ³School of Information, The
University of Texas at Austin, Austin, TX, United States

KEYWORDS

human-centered AI, human-in-the-loop AI, human-AI interaction, human computation,
crowdsourcing

Editorial on the Research Topic
[Human-centered AI: Crowd computing](#)

1. Introduction

Human computation (HCOMP) and crowdsourcing (Law and von Ahn, 2011; Quinn and Bederson, 2011; Kittur et al., 2013; Lease and Alonso, 2018) have been instrumental to advances seen in artificial intelligence (AI) and machine learning (ML) over the past 15+ years. AI/ML has an insatiable hunger for human labeled training to supervise models, with training data scale playing a significant (if not dominant) role in driving the predictive performance of models (Halevy et al., 2009). The centrality of such human-labeled data to the success and continuing advancement of AI/ML is thus at the heart of today's *data-centric AI* movement (Mazumder et al., 2022). Moreover, recent calls for *data excellence* (Aroyo et al., 2022) reflect growing recognition that AI/ML data scale alone does not suffice. The quality of human labeled data also plays a tremendous role in AI/ML success, and ignoring this can be perilous to deployed AI/ML systems (Sambasivan et al., 2021), as prominent, public failures have shown.

HCOMP and crowdsourcing have also enabled hybrid, human-in-the-loop, crowd-powered computing (Demartini et al., 2017). When state-of-the-art AI/ML cannot provide sufficient capabilities or predictive performance to meet practical needs for real-world deployment, hybrid systems utilize HCOMP at run-time to deliver last-mile capabilities where AI/ML fall short (Gadiraju and Yang, 2020). This has enabled a new class of innovative and more capable applications, systems, and companies to be built (Barr and Cabrera, 2006). While work in HCOMP is centuries old (Grier, 2013), access to an increasingly Internet-connected and well-educated world population led to the advent of crowdsourcing (Howe, 2006). This has allowed AI/ML systems to call on human help at run-time as "Human Processing Units (HPUs)" (Davis et al., 2010), "Remote Person Calls (RPCs)" (Bederson and Quinn, 2011), and "the Human API" (Irani and Silberman, 2013).

Across both data labeling and run-time HCOMP, crowdsourcing has enabled AI/ML builders to tap into the "wisdom of the crowd" (Surowiecki, 2005) and harness *collective intelligence* from large groups of people. As AI/ML systems have grown both more powerful and ubiquitous, appreciation of their capabilities has also been tempered by concerns of prevalence and propagation of biases, lack of robustness, fairness, and transparency as well as

ethical and societal implications. At the same time, crowdsourced access to a global, diverse set of contributors provides an incredible avenue to boost inclusivity and fairness in both AI/ML labeled datasets and hybrid, human-in-the-loop systems. However, important questions remain about the roles and treatment of AI/ML data workers, and the extent to which AI/ML advances are creating new economic opportunities for human workers (Paritosh et al., 2011) or exploiting hidden human labor (Bederson and Quinn, 2011; Fort et al., 2011; Irani and Silberman, 2013; Lease and Alonso, 2018; Gray and Suri, 2019). This has prompted the development of ethical principles for crowd work (Graham et al., 2020) and calls for *responsible sourcing* of AI/ML data (Partnership on AI, 2021).

As the above discussion reflects, HCOMP and crowdsourcing reflects a rich amalgamation of interdisciplinary research. In particular, the confluence of two key research communities—AI/ML and human-computer interaction (HCI)—has been central to founding and advancing HCOMP and crowdsourcing. Beyond this, related work draws upon a wide and rich body of diverse areas, including (but not limited to) computational social science, digital humanities, economics, ethics, law / policy / regulation, and social computing. More broadly, the HCOMP and crowdsourcing community promotes the exchange of advances in the state-of-the-art and best practices not only among researchers but also engineers and practitioners, to encourage dialogue across disciplines and communities of practice.

2. Call for papers: Aim and scope

Our organization of this *Frontiers Research Topic* called for new and high-impact contributions in HCOMP and crowdsourcing. We especially encouraged work that generates new insights into the collaboration and interaction between humans and AI, enlarging understanding about hybrid human-in-the-loop and algorithm-in-the-loop systems (Green and Chen, 2020). This includes human-AI interaction, algorithmic and interface techniques for augmenting human abilities to AI systems. It also spans issues that affect how humans collaborate and interact with AI systems such as bias, interpretability, usability, and trustworthiness. We welcomed both system-centered and human-centered approaches to human+AI systems, considering humans as users and stakeholders, or as active contributors and an integral part of the system.

Our call for papers invited submissions relevant to theory, studies, tools and/or applications that present novel, interesting, impactful interactions between people and computational systems. These cover a broad range of scenarios across human computation, wisdom of the crowds, crowdsourcing, and people-centric AI methods, systems and applications.

The scope of the Research Topics included the following themes:

- Crowdsourcing applications and techniques.
- Techniques that enable and enhance human-in-the-loop systems, making them more efficient, accurate, and human-friendly.
- Studies about how people perform tasks individually, in groups, or as a crowd.

- Approaches to make crowd science FAIR (Findable, Accessible, Interoperable, Reproducible) and studies assessing and commenting on the FAIRness of human computation and crowdsourcing practice.
- Studies into fairness, accountability, transparency, ethics, and policy implications for crowdsourcing and human computation.
- Methods that use human computation and crowdsourcing to build people-centric AI systems and applications, including topics such as reliability, interpretability, usability, and trustworthiness.
- Studies into the reliability and other quality aspects of human-annotated and -curated datasets, especially for AI systems.
- Studies about how people and intelligent systems interact and collaborate with each other and studies revealing the influences and impact of intelligent systems on society.
- Crowdsourcing studies into the socio-technical aspects of AI systems: privacy, bias, and trust.

3. Partnership with AAAI HCOMP

For over a decade, the premier venue for disseminating the latest research findings on HCOMP and crowdsourcing has been the Association for the Advancement of Artificial Intelligence (AAAI) Conference on Human Computation and Crowdsourcing (AAAI HCOMP).¹ Early HCOMP workshops at KDD and AAAI conferences (2009-2012) led to the genesis of the AAAI HCOMP conference in 2013. To further strengthen this *Frontiers Research Topic*, we partnered with AAAI HCOMP to invite submissions; papers accepted to HCOMP 2021 were offered a streamlined process for publication in this topic (e.g., maintaining the same reviewers when possible). We accepted four such submissions that extend earlier HCOMP 2021 papers (Samiotis et al., 2021; Welty et al., 2021; Yamanaka, 2021; Yasmin et al., 2021).

4. Managing conflicts-of-interest (COI)

“The *Frontiers* review system is designed to guarantee the most transparent and objective editorial and review process, and because the handling editor’s and reviewers’ names are made public upon the publication of articles, conflicts of interest will be widely apparent” (*Frontiers*, 2023). For this *Frontiers Research Topic*, two submissions from topic editors were routed by *Frontiers* staff to other editors not otherwise associated with this *Research Topic* and had no COI with the topic editors. Both submissions were ultimately accepted (Pradhan et al.; Samiotis et al.), after which the identity of each handling editor became publicly available. We thank these additional editors for their contributions to this *Research Topic*.

5. Research Topic contributions

A total of nine articles were accepted, contributing studies into factors of human computation and crowdsourcing, to their

¹ <https://humancomputation.com/>

applications to human-AI collaborative systems and large-scale behavioral studies. In the following, we very briefly summarize these works.

5.1. Quality in crowdsourced data annotation

Annotation quality is often a key concern in crowdsourced labeling. Pradhan et al. introduce a three-stage FIND-RESOLVE-LABEL workflow to reduce ambiguity in annotation task instructions. Their workflow allows workers to provide feedback on ambiguous task instructions to a requester. Another aspect of annotation quality is worker disagreement, for which a number of methods have been developed. Drawing from the observation that the effectiveness of annotation depends on the level of noise in the data, Uma et al. investigate the use of temperature scaling to estimate noise. Yasmin et al. investigates the effect of different forms of input elicitation to improve the quality of inferred labels in image classification, suggesting that more accurate results can be achieved when labels and self-reported confidence are used as features for classifiers.

5.2. Human-centered computation and interaction in AI

Tocchetti et al. study the effect of gamified activities to improve crowds' understanding of black-box models, addressing the intelligibility issue of explainable AI. They consider gamified activities to educate crowds by AI researchers. Yamanaka investigates the effectiveness of crowdsourcing for validating user performance models, especially the error-rate prediction model in target pointing tasks, which requires data from many repetitive experiments by participants for each task condition to measure the central tendency of the error rate. Welty et al. studies crowd knowledge creation for curating class-level knowledge graphs. Their three-tier crowd approach to elicit class-level attributes addresses the label sparsity problem faced by AI/ML systems.

5.3. Human factors in human computation

Vinella, Hu et al. focuses the effect of human agency in team formation on team performance. They found that in open collaboration scenarios, e.g., hackathon, teams formed by workers

themselves are more competitive, compared to those formed by algorithms. Samiotis et al. explore the possession of musical skills in the worker population. Their study shows that untrained workers possess high perception skills that can be useful in many music annotation tasks. Vinella, Odo et al. study the effect of personality on task performance by ad-hoc teams composed of strangers, especially in solving critical tasks that are often time-bounded and high-stress, e.g., incident response. Their results identify personality traits that affect team performance and in addition to that, relevant communication patterns used by winning teams.

Author contributions

All authors listed have made a substantial, direct, and intellectual contribution to the work and approved it for publication.

Acknowledgments

We would like to thank the AAAI HCOMP Steering Committee, as well as the organizers of the 2021 conference <https://www.humancomputation.com/2021/>, for their partnership with our Research Topic on Human-Centered AI: Crowd Computing. ML was supported in part by Good Systems <https://goodsystems.utexas.edu>, a UT Austin Grand Challenge to develop responsible AI technologies. Any opinions expressed in this article are entirely our own.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Aroyo, L., Lease, M., Paritosh, P., and Schaeckermann, M. (2022). Data excellence for AI: why should you care? *Interactions* 29, 66–69. doi: 10.1145/3517337
- Barr, J., and Cabrera, L. F. (2006). AI gets a brain: new technology allows software to tap real human intelligence. *Queue* 4, 24–29. doi: 10.1145/1142055.1142067
- Bederson, B. B., and Quinn, A. J. (2011). "Web workers unite! addressing challenges of online laborers," in *CHI Workshop on Crowdsourcing and Human Computation* (Vancouver, BC: ACM).
- Davis, J., Arderiu, J., Lin, H., Nevins, Z., Schuon, S., Gallo, O., et al. (2010). "The HPU," in *Computer Vision and Pattern Recognition Workshops (CVPRW)* (San Francisco, CA), 9–16.
- Demartini, G., Difallah, D. E., Gadiraju, U., and Catasta, M. (2017). An introduction to hybrid human-machine information systems. *Found. Trends® Web Sci.* 7, 1–87. doi: 10.1561/18000000025
- Fort, K., Adda, G., and Cohen, K. B. (2011). Amazon mechanical turk: gold mine or coal mine? *Comput. Linguist.* 37, 413–420. doi: 10.1162/COLL_a_00057

- Frontiers (2023). *Policies and Publication Ethics*. Available online at: <https://www.frontiersin.org/guidelines/policies-and-publication-ethics/>.
- Gadiraju, U., and Yang, J. (2020). "What can crowd computing do for the next generation of ai systems?" in *2020 Crowd Science Workshop: Remoteness, Fairness, and Mechanisms as Challenges of Data Supply by Humans for Automation* (CEUR), 7–13.
- Graham, M., Woodcock, J., Heeks, R., Mungai, P., Van Belle, J.-P., du Toit, D., et al. (2020). The fairwork foundation: strategies for improving platform work in a global context. *Geoforum* 112, 100–103. doi: 10.1016/j.geoforum.2020.01.023
- Gray, M. L., and Suri, S. (2019). *Ghost Work: How to Stop Silicon Valley From Building a New Global Underclass*. Eamon Dolan Books.
- Green, B., and Chen, Y. (2020). "Algorithm-in-the-loop decision making," in *Proceedings of the AAAI Conference on Artificial Intelligence* (New York, NY), Vol. 34, 13663–13664.
- Grier, D. A. (2013). *When Computers Were Human*. Princeton University Press.
- Halevy, A., Norvig, P., and Pereira, F. (2009). The unreasonable effectiveness of data. *IEEE Intellig. Syst.* 24, 8–12. doi: 10.1109/MIS.2009.36
- Howe, J. (2006). The rise of crowdsourcing. *Wired Magaz.* 14, 1–4.
- Irani, L., and Silberman, M. (2013). "Turkopticon: interrupting worker invisibility in amazon mechanical turk" in *Proceeding of the ACM SIGCHI Conference on Human Factors in Computing Systems* (Paris).
- Kittur, A., Nickerson, J. V., Bernstein, M., Gerber, E., Shaw, A., Zimmerman, J., et al. (2013). "The future of crowd work," in *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW)* (San Antonio, TX), 1301–1318.
- Law, E., and von Ahn, L. (2011). Human computation. *Synth. Lectur. Artif. Intellig. Mach. Learn.* 5, 1–121. doi: 10.1007/978-3-031-01555-7
- Lease, M., and Alonso, O. (2018). "Crowdsourcing and human computation: introduction," in *Encyclopedia of Social Network Analysis and Mining*, eds R. Alhajj and J. Rokne (New York, NY: Springer), 499–510.
- Mazumder, M., Banbury, C., Yao, X., Karlaš, B., Rojas, W. G., Damos, S., et al. (2022). Dataperf: benchmarks for data-centric AI development. *arXiv preprint arXiv:2207.10062*.
- Paritosh, P., Ipeiritos, P., Cooper, M., and Suri, S. (2011). "The computer is the new sewing machine: benefits and perils of crowdsourcing," in *Proceedings of the 20th International Conference Companion on World Wide Web* (Hyderabad: ACM), 325–326.
- Partnership on AI (2021). *Responsible Sourcing Across the Data Supply Line*. Available online at: <https://partnershiponai.org/workstream/responsible-sourcing/>.
- Quinn, A. J., and Bederson, B. B. (2011). "Human computation: a survey and taxonomy of a growing field," in *2011 Annual ACM SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC), 1403–1412.
- Sambasivan, N., Kapania, S., Highfill, H., Akrong, D., Paritosh, P., and Aroyo, L. M. (2021). "“Everyone wants to do the model work, not the data work”: Data cascades in high-stakes AI," in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–15.
- Samiotis, I. P., Qiu, S., Lofi, C., Yang, J., Gadiraju, U., and Bozzon, A. (2021). "Exploring the music perception skills of crowd workers," in *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 9, 108–119.
- Surowiecki, J. (2005). *The Wisdom of Crowds*. Anchor.
- Welty, C., Aroyo, L., Korn, F., McCarthy, S. M., and Zhao, S. (2021). "Rapid instance-level knowledge acquisition for google maps from class-level common sense," in *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, vol. 9, 143–154.
- Yamanaka, S. (2021). "Utility of crowdsourced user experiments for measuring the central tendency of user performance to evaluate error-rate models on guis," in *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 9, 155–165.
- Yasmin, R., Grassel, J. T., Hassan, M. M., Fuentes, O., and Escobedo, A. R. (2021). "Enhancing image classification capabilities of crowdsourcing-based methods through expanded input elicitation," in *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 9, 166–178.