



Delft University of Technology

Empowering Local Communities using Artificial Intelligence

Hsu, Yen-Chia; Huang, Ting-Hao; Verma, Himanshu; Mauri, Andrea; Nourbakhsh, Illah; Bozzon, Alessandro

DOI

[10.1016/j.patter.2022.100449](https://doi.org/10.1016/j.patter.2022.100449)

Publication date

2022

Document Version

Final published version

Published in

Patterns

Citation (APA)

Hsu, Y.-C., Huang, T.-H., Verma, H., Mauri, A., Nourbakhsh, I., & Bozzon, A. (2022). Empowering Local Communities using Artificial Intelligence. *Patterns*, 3(3), Article 100449. <https://doi.org/10.1016/j.patter.2022.100449>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Perspective

Empowering local communities using artificial intelligence

Yen-Chia Hsu,^{1,*} Ting-Hao ‘Kenneth’ Huang,² Himanshu Verma,¹ Andrea Mauri,^{1,4} Illah Nourbakhsh,³ and Alessandro Bozzon¹

¹Faculty of Industrial Design Engineering, Delft University of Technology, Delft, the Netherlands

²College of Information Sciences and Technology, Pennsylvania State University, State College, PA, USA

³Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA

⁴Amsterdam Institute for Advanced Metropolitan Solutions, Amsterdam, the Netherlands

*Correspondence: y.hsu-1@tudelft.nl

<https://doi.org/10.1016/j.patter.2022.100449>

THE BIGGER PICTURE Artificial intelligence (AI) is increasingly used to analyze large amounts of data in various practices, such as object recognition. We are specifically interested in using AI-powered systems to engage local communities in developing plans or solutions for pressing societal and environmental concerns. Such local contexts often involve multiple stakeholders with different and even contradictory agendas, resulting in mismatched expectations of the behaviors and desired outcomes of these systems. There is a need to investigate whether AI models and pipelines can work as expected in different contexts through co-creation and field deployment. Based on case studies in co-creating AI-powered systems with local people, we explain challenges that require more attention and provide viable paths to bridge AI research with citizen needs. We advocate for developing new collaboration approaches and mindsets that are needed to co-create AI-powered systems in multi-stakeholder contexts to address local concerns.



Proof-of-Concept: Data science output has been formulated, implemented, and tested for one domain/problem

SUMMARY

Artificial intelligence (AI) applications can profoundly affect society. Recently, there has been extensive interest in studying how scientists design AI systems for general tasks. However, it remains an open question as to whether the AI systems developed in this way can work as expected in different regional contexts while simultaneously empowering local people. How can scientists co-create AI systems with local communities to address regional concerns? This article contributes new perspectives in this underexplored direction at the intersection of data science, AI, citizen science, and human-computer interaction. Through case studies, we discuss challenges in co-designing AI systems with local people, collecting and explaining community data using AI, and adapting AI systems to long-term social change. We also consolidate insights into bridging AI research and citizen needs, including evaluating the social impact of AI, curating community datasets for AI development, and building AI pipelines to explain data patterns to laypeople.

INTRODUCTION

Artificial intelligence (AI) techniques are typically engineered with the goals of high performance and accuracy. Recently, AI algorithms have been integrated into diverse and real-world applications. Exploring the impact of AI on society from a people-centered perspective has become an important topic.¹ Previous works in citizen science have identified methods of using AI to engage the public in research, such as sustaining participation, verifying data quality, classifying and labeling objects, predicting user interests, and explaining data patterns.^{2–5} These works

investigated the challenges regarding how scientists design AI systems for citizens to participate in research projects at a large geographic scale in a generalizable way, such as building applications for citizens globally to participate in completing tasks. In contrast, we are interested in another area that receives significantly less attention: *How can scientists co-create AI systems with local communities to address context-specific concerns and influence a particular geographic region?*

Our perspective is based on applying AI in Community Citizen Science^{6,7} (CCS), a framework to create social impact through community empowerment at an intensely place-based local



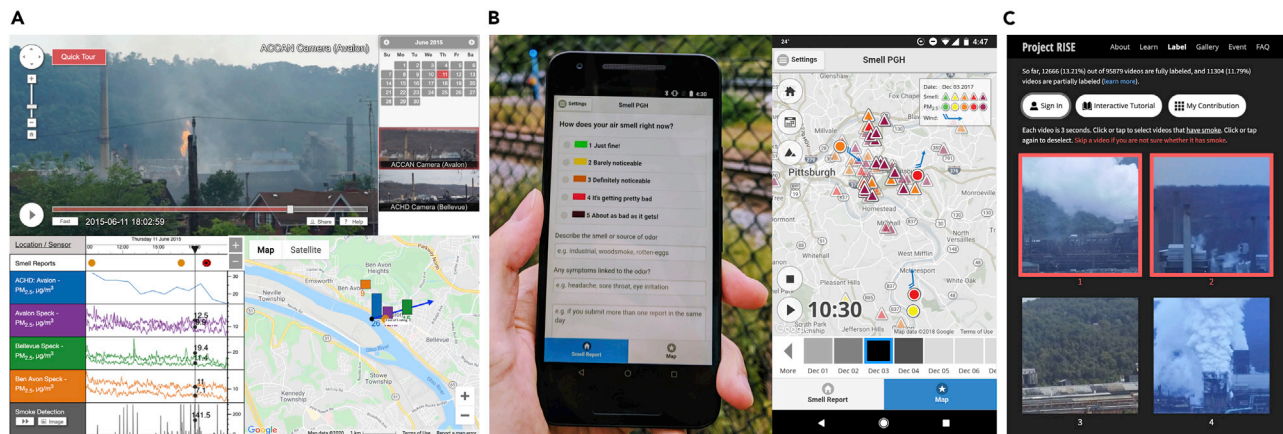


Figure 1. Case studies of community citizen science projects that involve co-designing AI tools with local communities
(A) The air pollution monitoring project²³ that empowered the Pittsburgh community to collect air pollution evidence in the local region for taking action.
(B) The Small Pittsburgh Project²⁴ invited citizens to report pollution odors and use the data as evidence to conduct air pollution studies.
(C) The RISE project²⁵ enabled citizens and scientists to annotate industrial smoke emissions and build an AI model to recognize pollution events.
These cases were approved by the ethics committee of the university that hosted the projects.

scale. We define *community* as a group of people who are indirectly or directly affected by issues in civil society and are dedicated to making sure that these issues are recognized and resolved. We define *social impact* as how a project influences the society and local communities that are affected by social or environmental issues. We define *community empowerment* as a process of yielding agency to communities so that they can use technology, data, and informed rhetoric to create and disseminate evidence to advocate for social and policy changes. The CCS framework, a branch of citizen science,^{8,9} is beneficial in co-creating solutions and driving social impact with communities that pursue the Sustainable Development Goals.¹⁰ Based on the literature and our experiences in co-creating AI systems with citizens, this article provides critical reflections regarding this underexplored topic for data science, AI, citizen science, and human-computer interaction fields. We discuss the challenges and insights in connecting AI research closely to social issues and citizen needs, using prior works as examples.

How CCS links to other frameworks

CCS emphasizes close collaborations among stakeholders when tackling local concerns. It is inspired by community-based participatory research¹¹ and popular epidemiology,¹² in which citizens directly engage in gathering data and extract knowledge from these data for advocacy and activism. Examples involve co-designing technology for local watershed management,¹³ understanding water quality with local communities,^{14,15} and using geo-information tools to monitor noise and earthquakes.¹⁶ CCS intends to extend the scope of previous frameworks to Sustainable Development Goals, especially the goal of sustainable cities and communities. This article discusses using CCS to integrate AI in-the-wild and local regions, which is different from those that conducted studies in living lab environments (e.g., the work by Alavi et al.¹⁷) or in online communities (e.g., the work by Brambilla et al.¹⁸).

In addition, CCS is related to action research,¹⁹ RtD (research through design),²⁰ service design,²¹ and the PACT (participatory

approach to enable capabilities in communities) framework.²² Extending action research, CCS encourages scientists to immerse themselves in the field by taking on a social role and conducting research from a first-person view. Complementing RtD that creates prototypes as proof-of-concept, CCS develops functional systems that can be deployed and used by local people. Unlike service design, citizens' roles extend beyond service consumers to co-designers who co-create knowledge and systems with scientists and other stakeholders. The PACT framework and CCS share the same goal of co-designing AI systems to address critical societal issues, while CCS has an additional goal that needs to be achieved simultaneously: empowering local communities to catalyze social impact.

Challenges

Due to its region-based characteristics, CCS often involves many local stakeholders—including communities, citizens, scientists, designers, and policy makers—with complex relationships. CCS creates the space for the stakeholders to reveal underlying difficulties and locally oriented action plans in tackling social concerns that are hard to uncover in traditional technology-oriented and researcher-centered approaches. However, stakeholders often have divergent and even contradicting values, which result in conflicts that pose challenges when designing, engineering, deploying, and evaluating AI systems.

Based on case studies (Figure 1) of collaborating with local people in building AI systems, we outline three major challenges:

- Co-designing AI systems with local communities
- Collecting and explaining community data using AI
- Adapting AI systems to long-term social changes

These challenges come mainly from the conflicts of interest between local communities (e.g., citizens, community groups) and university researchers (e.g., designers, scientists). AI researchers pursue knowledge to advance science, while local communities often desire social change. Such conflicts of interest among these two groups can lead to tensions,

socio-technical gaps, and mismatched expectations when co-designing and engineering AI systems. For instance, local communities need functional and reliable systems to collect evidence, but AI researchers may be interested in producing system prototypes only to prove concepts or answer their research questions. Local community concerns can be urgent and timely, and citizens need to take practical actions that can have an immediate and effective social impact, such as public policy change. However, scientists need to produce knowledge using rigorous methods and publish papers in the academic community, often requiring a long reviewing and publication cycle.

Co-designing AI systems with local communities

Challenges exist in community co-design, especially when translating multifaceted community needs to implementable AI system features without using research-centered methods. The current practice to design AI systems is mainly centered on researchers instead of local people. Popular research methods, such as participatory design workshops, interviews, and surveys, are normally used to help designers and scientists understand research questions. Although these methods enable researchers to better control the research process, essentially, university researchers are privileged and in charge of the conversations, leading to inappropriate power dynamics that can hinder trust and exacerbate inequality.^{26,27} For example, during our informal conversations with local people that suffer from environmental concerns in our air quality monitoring project,²³ many expressed feelings that scientists often treated them as experimental subjects (but not collaborators) in research studies. This imbalanced power relationship leads to difficulties in initiating conversations with citizens during our community outreach efforts.

Also, community data and knowledge are hyperlocal, which indicates that their underlying meanings ground closely to the local region and could be difficult to grasp for researchers who are not a part of the local community.²⁸ For example, citizen-organized meetings to discuss community actions are often dynamic and context-dependent, which is not designed nor structured for research purposes. To collect research data that represent community knowledge, the current intensive procedure, such as video or audio recording, can make citizens feel uncomfortable. One alternative is to be a part of the community, to design solutions with them, join their actions, and perform ethnographic observations. For instance, researchers can better understand local community needs by actively participating in regional citizen group meetings and daily conversations with citizens. However, such in-depth community outreach approaches take tremendous personal effort, which can be unmotivating or even infeasible due to the limited academic research cycle and research-oriented academic tenure awarding system.^{11,27}

Collecting and explaining community data using AI

Challenges arise in data collection, analysis, and interpretation due to conflicts of interest among scientists and citizens. Scientists look for rigorous procedures, but citizens seek evidence for action. Local communities are often frustrated by the formal scientific research procedure to prove the adverse impact of risk,¹² such as finding evidence of how pollution negatively affects

health. Traditional environmental risk assessment models require a causal link between the risk and the outcome with statistical significance before taking action, which can be very difficult to achieve due to complex relationships between local people and their environments.²⁹ As a result, citizens collect their own community data (as defined by Carroll et al.,³⁰ such as photographs of smoke emission from a nearby factory) as an alternative to prove their hypotheses. From scientists' point of view, however, such strong assumption-driven evidentiary collection can lead to biases since the collection, annotation, and analysis of community data are conducted in a manner that strongly favors the assumption. One example is confirmation bias, in which people are incentivized to search for information and provide data that confirm their prior beliefs,³¹ such as a high tendency to report odors related to pollution events.²⁴ Based on our experiences, it is extremely difficult to address or eliminate such biases when analyzing and interpreting how local social or environmental concerns affect communities.

Furthermore, researchers need to evaluate the social impact of AI systems to understand whether the community co-design approach is practical. However, it is hard to determine whether the intervention of AI systems actually influences the local people and leads to social changes by statistically analyzing community data. One difficulty is that people may have the implicit cognitive bias to overestimate and overstate the effect of the intervention since they are deeply involved in the co-creation of the AI system.³² Moreover, it can be infeasible to conduct randomized experiments to prove the effectiveness of the intervention of AI on local communities.⁷ Controlling volunteer demographics and participation levels can be unethical when analyzing impact among different groups of people. CCS treats local people as collaborators rather than as participants. Researchers in CCS take the supporting role to assist communities using technology, instead of supervising and overseeing the entire project.⁷ Therefore, citizens join the CCS project at will and are not recruited as in typical research studies. AI systems, in this case, are deployed in the wild with real consequences rather than a controlled test-bed environment that is designed for hypothesis testing. It remains an open research question as to how to integrate social science when studying the impact of AI systems.³³

Adapting AI systems to long-term social changes

Conflicts of interest in the diverged values of citizens and scientists can lead to challenges in adapting AI systems to long-term social changes. The relationship between local people and AI systems is a feedback loop, which is similar to the concept that human interactions with architectural infrastructure are a continuous adaptation process that spans over long periods of time.³⁴ When embedded in the social context, AI systems interact with citizens daily as community infrastructure. Communities are dynamic and frequently evolve their agenda to adapt to social context changes. This means that the AI systems also need to adjust to such changes in local communities continuously. For instance, as we understand more about the real-life effects of the deployed AI systems on local people, we may need to fine-tune the underlying machine learning model using local community data. We may also need to improve the data analysis pipeline and strategies for interpreting results to fit local community needs in taking action. We may even need to stop the AI

system from intervening in the local community under certain conditions. Such adaptation at scale requires ongoing commitment from researchers, designers, and developers to continuously maintain the infrastructure, involve local people in assessing the impact of AI, adjust the behavior of AI systems, and support communities in taking action to advocate for social changes.³⁵

However, it is very challenging to estimate and obtain the required resources to sustain such long-term university-community engagement with local people,³⁶ especially in financially supporting local community members for their efforts. Typical research procedures can be laborious in data collection and analysis, and engineering AI systems with local people requires a tremendous community outreach effort to establish mutual trust. In our experiences, applying and evaluating AI in CCS relies heavily on an environment that has a sustainable fundraising mechanism in community organizations and universities. For example, funding is needed to hire software engineers that can maintain AI systems as community infrastructure in the long term, which can be hard to achieve in the current academic grant instruments and funding cycles.

The success of CCS also depends on sustainable participation, which requires high levels of altruism, high awareness of local issues, and sufficient self-efficacy among local people. However, the complexity of the underlying machine learning techniques can affect the willingness to participate. On one side is whether the automation technique is trustworthy. In our experience, local communities often perceive AI as a mysterious box that can be questionable and is not always guaranteed to work. Hence, the willingness of citizens to provide data can be low, but AI systems that use machine learning and computer vision need data to be functional. On the other side, “what citizens think the AI system can do” does not match “what the AI system can actually do,” resulting in socio-technical gaps and pitfalls for actual usage.³⁷ In our experience, local communities often have high expectations about what AI techniques can do for them—for example, automatically determining whether an industrial site is violating environmental regulations. However, in practice, the AI system may only identify whether a factory emits smoke and degrades the quality of the air through sensors and cameras, which requires additional human efforts to verify whether the pollution event is indeed a violation.

Bridge AI research and citizen needs

University researchers typically lead the development of AI systems using a researcher-centered approach, in which they often have more power over local communities (especially underserved ones) in terms of scientific authority and available resources. This unequal power relationship can result in a lack of trust and cause harm to underserved communities.^{26,36} An underlying assumption of this researcher-centered approach is that designers and scientists can put themselves in the situation of citizens and empathize with local people’s perspectives. However, university researchers are in a privileged situation in terms of socio-economic status and may come from other geographical regions or cultures, which means it can be very challenging for researchers to fully and authentically understand local people’s experiences.¹² Only by admitting this weakness and recognizing the power inequality can researchers truly respect

community knowledge and be sincerely open-minded in involving local communities—especially those affected by the problems the most—in the center of the design process when creating AI systems. Beyond being like the local people and designing solutions for them, researchers need to be with people who are affected by local concerns to co-create historicity and ensure that the AI systems are created to be valuable and beneficial to them.^{35,38}

The critical role of creating social impact lies in local people and their long-term perseverance in advocating for changes. We believe that scientists need to collaborate with local people to address pressing social concerns genuinely, and even further, to immerse themselves into the local context and become citizens, hence “scientific citizens” (as defined by Irwin⁸). However, pursuing academic research and addressing citizen concerns require different (even contradictory) efforts and can be difficult to achieve at the same time. Academic research requires contributing papers with scientific knowledge primarily to the research community, while citizen concerns typically involve many other stakeholders in a large and regional socio-technical system. It remains an open question how scientists and citizens can collaborate effectively under such dynamic, hyperlocal, and place-based conditions.³⁹

To move forward, we propose three viable CCS approaches about how AI designers and scientists can conduct research and co-create social impact with local communities:

- Evaluate social impact of AI as empirical contributions
- Curate community data as dataset contributions
- Build AI pipelines as methodological contributions

These approaches produce empirical, dataset, and methodological contributions, respectively, to the research community, as defined by Wobbrock and Kientz.⁴⁰ To the local people, these approaches establish a long-term fair university-community partnership in addressing community concerns, increase literacy in collecting community data, and equip communities with AI tools to interpret data. CCS projects will succeed when designers and scientists see themselves as citizens, and in turn, when local communities and citizens see themselves as innovators. It is essential for all parties to collaborate around the lived experiences of one another and listen to each other’s voices with humility and respect.

Social impact of AI as empirical contributions

Lessons learned from previously deployed AI systems in other contexts cannot be simply applied in the current one, as local communities have various cultures, behaviors, beliefs, values, and characteristics.³³ Hence, it is essential to understand and document how scientists can co-design AI systems with local communities and co-create long-term social impact in diverse contexts. It is also important to study the effectiveness and impact of various AI interventions with different design criteria in sustaining participation, affecting community attitude, and empowering people. The CCS framework provides a promising path toward these goals. Implications of collaborating with local people in co-designing AI interventions, creating long-term effects, and tackling the conflicts of interest among stakeholders can be strong empirical contributions to the academic community.³⁵

The data-driven evidence and the interventions that are produced by AI systems can affect the local region in various ways, including increasing residents' confidence in addressing concerns, providing convincing evidence, or rebalancing power relationships among stakeholders.

For instance, our air pollution monitoring project documented the co-design process in regard to how designers translated citizen needs and local knowledge into implementable AI system features, as recognized by the human-computer interaction community and published in the proceedings of the Association for Computing Machinery Conference on Human Factors in Computing Systems (ACM CHI).²³ This work shows researchers how we co-created an AI system to support citizens in collecting air pollution evidence and how local communities used the evidence to take action. For example, in the computer vision model for finding industrial smoke emissions in videos, the feature vectors are handcrafted according to the behaviors and characteristics of smoke, which are provided by community knowledge. Also, the communities decide the areas in the video that require image processing. The decision of having high precision in the prediction (instead of high recall) is also a design choice by local people for quickly determining severe environmental violations. Another example is our study of push notifications, which are generated by an AI model to predict the presence of bad odors in the city.²⁴ The finding from the study explains how sending certain types of push notifications to local citizens is related to the increase of their engagement level, such as contributing more odor reports or browsing more data.

Although one may not simply duplicate the collaboration ecosystem in these contexts due to unique characteristics in the local communities, our projects can be seen as case studies in specific settings. Our air quality monitoring case provides insights to researchers working on similar problems in other contexts about integrating technology reliably into their settings, as cited by Ottinger.⁴¹ Moreover, the case also helps researchers understand and categorize different modes of community empowerment, as cited by Schneider et al.⁴²

Community data as dataset contributions

Data work is critical in building and maintaining AI systems,⁴³ as modern AI models are powered by large and constantly changing datasets. When addressing local concerns with the support of AI systems, researchers often need to fine-tune existing models or build new pipelines to fit local needs. This requires collecting data in a specific regional context and may introduce new tasks to the AI research field. CCS provides a sustainable way to co-create high-quality regional datasets while simultaneously increasing citizens' self-efficacy in addressing local problems. Based on our experiences, co-creating publicly available community data can also facilitate citizens' sense of ownership of the collaborative work. This value of community empowerment links AI research closely to social impact and the public good.

Besides the value of increasing citizens' data literacy, the collected real-world data, the data collection approach, and the data processing pipeline can be combined into a significant dataset contribution to the academic community in creating robust AI models. Such community datasets are gathered in the wild with local populations over a long period to reflect the regional context, which complements the datasets obtained us-

ing crowdsourcing approaches (e.g., Amazon Mechanical Turk) in a broader context. In this way, community datasets provide values for AI researchers to validate whether AI models trained on general datasets can work as expected in different regional contexts. Also, the accompanying software for data labeling can contribute reusable computational tools to the research community that investigates data annotation strategies.

For example, our RISE (Recognizing Industrial Smoke Emissions) project presented a novel video dataset for the smoke recognition task, which can help other researchers develop better computer vision models for similar tasks, as recognized by the AI community and published by the Association for the Advancement of Artificial Intelligence (AAAI).²⁵ Our project demonstrated the approach of collaborating with citizens affected by air pollution to annotate videos with industrial smoke emissions on a large scale under various weather and lighting conditions. The dataset was used to train a computer vision model to recognize industrial smoke emissions, which allowed community activists to curate a list of severe pollution events as evidence to conduct studies and advocate for enforcement action. Another example is the Mosquito Alert project, which curates and labels a large mosquito image dataset with local people using a mobile application.⁴⁴ The dataset is built with local community knowledge and is used to train a mosquito recognition model to support the local public health agency in disease management. In addition to its social impact, the Mosquito Alert project advances science by providing a real-world dataset for researching different mosquito recognition models, as cited by Adhane et al.⁴⁵

AI pipelines as methodological contributions

In CCS, there is a need to unite expertise from the local communities and scientists to build AI pipelines using machine learning to assist data labeling, predict trends, or interpret patterns in the data. An example is to forecast pollution and find evidence of how pollution affects the quality of living in a local region. Although the concept of machine learning is common among computer scientists, it can seem mysterious to citizens.

Thus, during public communication and community outreach, researchers often need to visualize analysis results and explain the statistical evidence for local residents, which is highly related to the explainable AI (XAI) and interpretable machine learning research.⁴⁶ However, current XAI research focuses mainly on making AI understandable for experts rather than laypeople and local communities.⁴⁷ This creates a unique research opportunity to study co-design methods and software engineering workflows of translating the predictions of AI models and their internal decision-making process into human-intelligible insights in the hyperlocal context.^{1,48,49} We believe the pipeline of such translation into explainable evidence can be a methodological contribution to the academic community, which provides a way to address the challenge of predicting future trends and interpreting similar types of real-world data. Also, the implemented machine learning pipeline and the design insights of developing the pipeline can contribute reusable computational tools and novel software engineering workflows to the research communities that study XAI and its user interfaces.

For instance, our Smell Pittsburgh project used machine learning to explain relationships between citizen-contributed odor reports and air quality sensor measurements, contributing

to a methodological pipeline of translating AI predictions, as recognized by the intelligent user interface community and published in the ACM Transactions on Interactive Intelligent Systems (TiiS).²⁴ In this way, the pollution patterns became visible to public scrutiny and debate. Another example is the xAire project, which co-designed solutions with local schools and communities to collect nitrogen dioxide data.⁵⁰ Air measurements were analyzed with asthma cases in children using a statistical machine learning model. The community outreach and public communication enabled laypeople to make sense of how nitrogen dioxide posed a risk to local community health. The pipelines in these two examples produced meaningful patterns for citizens to understand and communicate about how pollution affects the local region. They also informed researchers about how to process, wrangle, analyze, and interpret urban data to explain insights to laypeople.

Next steps

We have explained major challenges in co-designing AI systems with local people and empowering them to create a broader social impact. We also proposed CCS approaches to simultaneously addressing local societal issues and advancing science. Computing research communities have taken steps to recognize the impact of technology and AI on society, such as establishing a separate track for paper evaluation (e.g., the AAAI Special Track on AI for Social Impact). We urge the computing research communities to go further and acknowledge social impact as a type of formal contribution in scientific inquiry and paper publication. Promoting this kind of contribution can be a turning point to encourage scientists to link research to society and ultimately make university research socially responsible for the public good. Co-creating AI systems and developing reusable tools with local communities in the long term allows scientists and designers to explore real-world challenges and solution spaces for various AI techniques, including machine learning, computer vision, and natural language processing.

We also urge universities to integrate social impact into the evaluation criteria of the tenure roadmap of the academic professorship as the “service” pillar of the university that contributes to the public good. We envision that applying CCS when co-designing AI systems can advance science, build public trust in AI research through genuine reciprocal university-community partnership, and directly support community action to affect society. In this way, we may fundamentally change how universities, organizations, and companies partner with their neighbors to pursue shared prosperity in the future of community-empowered AI.

ACKNOWLEDGMENTS

We greatly appreciate the support of this work from the European Commission under the EU Horizon 2020 framework (grant number 101016233), within project PERISCOPE (Pan-European Response to the Impacts of COVID-19 and Future Pandemics and Epidemics).

AUTHOR CONTRIBUTIONS

Conceptualization, Y.-C.H., T.-H.K.H., H.V., A.M., I.N., and A.B.; methodology, Y.-C.H. and T.-H.K.H.; investigation, Y.-C.H. and T.-H.K.H.; writing – original draft, Y.-C.H.; writing – review & editing, Y.-C.H., T.-H.K.H., H.V., A.M., I.N.,

and A.B.; supervision, A.B.; project administration, Y.-C.H.; funding acquisition, A.B.

DECLARATION OF INTERESTS

The authors declare no competing interests.

REFERENCES

- Shneiderman, B. (2020). Bridging the gap between ethics and practice: Guidelines for reliable, safe, and trustworthy human-centered AI systems. *ACM Trans. Interactive Intell. Syst. (TiiS)* 10, 1–31.
- Ceccaroni, L., Bibby, J., Roger, E., Flemons, P., Michael, K., Fagan, L., and Oliver, J.L. (2019). Opportunities and risks for citizen science in the age of Artificial Intelligence. *Citizen Sci. Theor. Pract.* 4, 1.
- Franzen, M., Kloetzer, L., Ponti, M., Trojan, J., and Vicens, J. (2021). *Machine Learning in Citizen Science: Promises and Implications (The Science of Citizen Science)*, p. 183.
- Lotfian, M., Ingensand, J., and Brovelli, M.A. (2021). The partnership of citizen science and machine learning: benefits, risks, and future challenges for engagement, data collection, and data quality. *Sustainability* 13, 8087.
- McClure, E.C., Sievers, M., Brown, C.J., Buelow, C.A., Ditria, E.M., Hayes, M.A., Pearson, R.M., Tulloch, V.J., Unsworth, R.K., and Connolly, R.M. (2020). Artificial intelligence meets citizen science to supercharge ecological monitoring. *Patterns* 1, 100109.
- Chari, R., Matthews, L.J., Blumenthal, M., Edelman, A.F., and Jones, T. (2017). *The Promise of Community Citizen Science (RAND)*.
- Hsu, Y.-C., and Nourbakhsh, I. (2020). When human-computer interaction meets community citizen science. *Commun. ACM* 63, 31–34.
- Irwin, A. (2001). Constructing the scientific citizen: science and democracy in the biosciences. *Public Understanding Sci.* 10, 1–18.
- Shirk, J.L., Ballard, H.L., Wilderman, C.C., Phillips, T., Wiggins, A., Jordan, R., McCallie, E., Minarchek, M., Lewenstein, B.V., Krasny, M.E., et al. (2012). Public participation in scientific research: a framework for deliberate design. *Ecol. Soc.* 17, 29.
- Fritz, S., See, L., Carlson, T., Haklay, M.M., Oliver, J.L., Fraisl, D., Mondardini, R., Brocklehurst, M., Shanley, L.A., Schade, S., et al. (2019). Citizen science and the United Nations sustainable development goals. *Nat. Sustainability* 2, 922–930.
- Wallerstein, N.B., and Duran, B. (2006). Using community-based participatory research to address health disparities. *Health Promot. Pract.* 7, 312–323.
- Brown, P. (1993). When the public knows better: popular epidemiology challenges the system. *Environ. Sci. Policy Sustainable Development* 35, 16–41.
- Preece, J., Pauw, D., and Clegg, T. (2019). Interaction design of community-driven environmental projects (CDEPS): a case study from the anastasia watershed. *Proc. Natl. Acad. Sci.* 116, 1886–1893.
- Carroll, J.M., Beck, J., Boyer, E.W., Dhanorkar, S., and Gupta, S. (2019). Empowering community water data stakeholders. *Interacting Comput.* 31, 492–506.
- Jollymore, A., Haines, M.J., Satterfield, T., and Johnson, M.S. (2017). Citizen science for water quality monitoring: data implications of citizen perspectives. *J. Environ. Manage.* 200, 456–467.
- Carton, L., and Ache, P. (2017). Citizen-sensor-networks to confront government decision-makers: two lessons from The Netherlands. *J. Environ. Manage.* 196, 234–251.
- Alavi, H.S., Verma, H., Mlynar, J., and Lalanne, D. (2018). *The Hide and Seek of Workspace: Towards Human-Centric Sustainable Architecture (Association for Computing Machinery)*, pp. 1–12.
- Brambilla, M., Ceri, S., Mauri, A., and Volonterio, R. (2014). Community-based crowdsourcing. In *Proceedings of the 23rd International Conference on World Wide Web, WWW'14 Companion (Association for Computing Machinery)*, pp. 891–896.

19. Susman, G.I., and Evered, R.D. (1978). An assessment of the scientific merits of action research. *Administrative Sci. Q.* 582–603.
20. Zimmerman, J., Forlizzi, J., and Evenson, S. (2007). Research through design as a method for interaction design research in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI'07* (Association for Computing Machinery), pp. 493–502.
21. Zomerdiijk, L.G., and Voss, C.A. (2010). Service design for experience-centric services. *J. Serv. Res.* 13, 67–82.
22. Bondi, E., Xu, L., Acosta-Navas, D., and Killian, J.A. (2021). Envisioning communities: a participatory approach towards AI for social good. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society, AIES'21* (Association for Computing Machinery), pp. 425–436.
23. Hsu, Y.-C., Dille, P., Cross, J., Dias, B., Sargent, R., and Nourbakhsh, I. (2017). Community-empowered air quality monitoring system. In *Proceedings of the 2017 ACM Conference on Human Factors in Computing Systems*, pp. 1607–1619.
24. Hsu, Y.-C., Cross, J., Dille, P., Tasota, M., Dias, B., Sargent, R., Huang, T.-H., and Nourbakhsh, I. (2020). Smell Pittsburgh: engaging community citizen science for air quality. *ACM Trans. Interactive Intell. Syst.* 10, 1–49.
25. Hsu, Y.-C., Huang, T.-H.K., Hu, T.-Y., Dille, P., Prendi, S., Hoffman, R., Tshlaries, A., Pachuta, J., Sargent, R., and Nourbakhsh, I. (2021). Project RISE: recognizing industrial smoke emissions. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35, pp. 14813–14821.
26. Harrington, C., Erete, S., and Piper, A.M. (2019). Deconstructing community-based collaborative design: towards more equitable participatory design engagements. *Proc. ACM Human-Computer Interaction* 3, 1–25.
27. Klein, P., Fatima, M., McEwen, L., Moser, S.C., Schmidt, D., and Zupan, S. (2011). Dismantling the ivory tower: engaging geographers in university-community partnerships. *J. Geogr. Higher Education* 35, 425–444.
28. Carroll, J.M., Hoffman, B., Han, K., and Rosson, M.B. (2015). Reviving community networks: hyperlocality and suprathresholding in web 2.0 designs. *Personal. Ubiquitous Comput.* 19, 477–491.
29. Bidwell, D. (2009). Is community-based participatory research postnormal science? *Sci. Technol. Hum. Values* 34, 741–761.
30. Carroll, J.M., Beck, J., Dhanorkar, S., Binda, J., Gupta, S., and Zhu, H. (2018). Strengthening community data: towards pervasive participation. In *Proceedings of the 19th Annual International Conference on Digital Government Research: Governance in the Data Age*, pp. 1–9.
31. Nickerson, R.S. (1998). Confirmation bias: a ubiquitous phenomenon in many guises. *Rev. Gen. Psychol.* 2, 175–220.
32. Norton, M.I., Mochon, D., and Ariely, D. (2012). The IKEA effect: when labor leads to love. *J. Consumer Psychol.* 22, 453–460.
33. Sloane, M., and Moss, E. (2019). AI's social sciences deficit. *Nat. Machine Intelligence* 1, 330–331.
34. Alavi, H.S., Verma, H., Mlynar, J., and Lalanne, D. (2019). On the temporality of adaptive built environments. In *People, Personal Data and the Built Environment* (Springer), pp. 13–40.
35. Sloane, M., Moss, E., Awomolo, O., and Forlano, L. (2020). Participation is not a design fix for machine learning. In *Proceedings of the International Conference on Machine Learning*.
36. Koekoek, A., Van Ham, M., and Kleinhans, R. (2021). Unraveling university-community engagement: a literature review. *J. Higher Education Outreach Engagement* 25, 3.
37. Roberts, M., Driggs, D., Thorpe, M., Gilbey, J., Yeung, M., Ursprung, S., Aviles-Rivero, A.I., Etmann, C., McCague, C., Beer, L., et al. (2021). Common pitfalls and recommendations for using machine learning to detect and prognosticate for COVID-19 using chest radiographs and CT scans. *Nat. Machine Intelligence* 3, 199–217.
38. Bennett, C.L., and Rosner, D.K. (2019). The promise of empathy: design, disability, and knowing the “other”. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–13.
39. Bozzon, A., Houtkamp, J., Kresin, F., De Sena, N., and de Weerd, M. (2015). From Needs to Knowledge: A Reference Framework for Smart Citizens Initiatives, Tech. rep. (Amsterdam Institute for Advanced Metropolitan Solutions (AMS))
40. Wobbrock, J.O., and Kientz, J.A. (2016). Research contributions in human-computer interaction. *Interactions* 23, 38–44.
41. Ottinger, G. (2017). Crowdsourcing undone science. *Engaging Sci. Technol. Soc.* 3, 560–574.
42. Schneider, H., Eiband, M., Ullrich, D., and Butz, A. (2018). Empowerment in HCI - a survey and framework. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–14.
43. Sambasivan, N., Kapania, S., Highfill, H., Akrong, D., Paritosh, P., and Arayo, 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*, pp. 1–15.
44. Pataki, B.A., Garriga, J., Eritja, R., Palmer, J.R., Bartumeus, F., and Csabai, I. (2021). Deep learning identification for citizen science surveillance of tiger mosquitoes. *Scientific Rep.* 11, 1–12.
45. Adhane, G., Dehshibi, M.M., and Masip, D. (2021). A deep convolutional neural network for classification of aedes albopictus mosquitoes. *IEEE Access* 9, 72681–72690.
46. Doshi-Velez, F., and Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv*, arXiv:1702.08608.
47. Cheng, H.-F., Wang, R., Zhang, Z., O'Connell, F., Gray, T., Harper, F.M., and Zhu, H. (2019). Explaining decision-making algorithms through UI: strategies to help non-expert stakeholders. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–12.
48. Burrell, J. (2016). How the machine “thinks”: understanding opacity in machine learning algorithms. *Big Data Soc.* 3, 2053951715622512.
49. Miller, T. (2019). Explanation in artificial intelligence: insights from the social sciences. *Artif. Intelligence* 267, 1–38.
50. Perelló, J., Cigarini, A., Vicens, J., Bonhoure, I., Rojas-Rueda, D., Nieuwenhuijsen, M.J., Cirach, M., Daher, C., Targa, J., and Ripoll, A. (2021). Large-scale citizen science provides high-resolution nitrogen dioxide values and health impact while enhancing community knowledge and collective action. *Sci. Total Environ.* 789, 147750.