

# Machine Learning in Citizen Science: Promises and Implications

Martina Franzen, Institute for Advanced Studies in the Humanities (KWI), Germany

Laure Kloetzer, University of Neuchâtel, Switzerland

Marisa Ponti, University of Gothenburg, Sweden

Jakub Trojan, Tomas Bata University in Zlin, Czech Republic

Julián Vicens, Data Science and Big Data Analytics, Eurecat; OpenSystems Universitat de Barcelona; Universitat de Barcelona Institute of Complex Systems, Spain

## Chapter Abstract

The chapter gives an account of both opportunities and challenges of human-machine collaboration in citizen science. In the age of big data, scientists are facing the overwhelming task of analysing massive amounts of data, and machine learning techniques are becoming a possible solution. Human and artificial intelligence can be recombined in citizen science in numerous ways. For example, citizen scientists can be involved in training machine learning algorithms in such a way that they perform certain tasks such as image recognition. To illustrate the possible applications in different areas, we discuss example projects of human-machine cooperation with regard to their underlying concepts of learning. The use of machine learning techniques creates lots of opportunities, such as reducing the time of classification and scaling expert decision-making to large data sets. However, algorithms often remain black boxes and data biases are not visible at first glance. Addressing the lack of transparency both in terms of machine action and in handling user-generated data, the chapter discusses how machine learning is actually compatible with the idea of active citizenship and what conditions need to be met in order to move forward – both in citizen science and beyond.

### Keywords:

Algorithms, Artificial Intelligence, Computer Vision, Machine Learning, Transparency, Sensor, Datafication

## Introduction

The combination of human and machine learning, wherever they complement one another, has a lot of potential applications in citizen science. Several projects have already integrated both forms of learning to perform data-centered tasks (Willi et al. 2019; Sullivan et al. 2018). While the term artificial intelligence (AI)

is generally used to refer to any kind of machine or algorithm able to observe the environment, learn, and make decisions, the term machine learning (ML) has been defined ‘as a subfield of artificial intelligence that includes software able to recognize patterns, make predictions, and apply newly discovered patterns to situations that were not included or covered by their initial design’ (Popenici and Kerr 2017, p. 2). ML algorithms are currently the most widely used and applied, for example, in image and speech recognition, fraud detection, and reproducing human abilities in playing Go or driving cars. In scientific research, they find many applications in different fields such as biology, astronomy, and social sciences, just to mention a few (Jordan and Mitchel 2015). Although AI is not new to citizen science (Ceccaroni et al. 2019), the convergence of advanced computing, availability of data, and learning algorithms can introduce something dramatically new in this area. The opportunities are many, and in some cases not yet foreseen, but so are the challenges, including the need to advance the explainability, accountability, and fairness of algorithms from the perspective of ML research and from that of citizen scientists using the applications.

We address two main questions here: (1) what tasks are citizens being invited to perform in citizen science projects through the use of ML, and (2) what are the main risks and opportunities of using ML in citizen science? The majority of citizen science projects is centered around data provided, for example, by satellites, cameras, or, more generally, sensors (Neal 2013). Collecting, analysing, and interpreting data are some of the most common activities that participants carry out, depending on their level of engagement in the scientific research process (Bonney et al. 2009). Similarly, ML makes sense in a variety of stages of the data–science life cycle through algorithms that perform tasks like classification, regression, clustering, and association, especially when dealing with huge amounts of data.

Many research problems are still considered computationally intractable and need human cognitive skills. For example, machines cannot yet match a person’s ability to identify certain objects and it is unclear to what extent they will ever succeed. Conversely, manual classification or identification of a large data set can be made more efficient in combination with ML approaches. Even so, the participation of citizens and the collective intelligence that emerges from it becomes fundamental to perform certain tasks, such as the creation of data sets with correctly tagged data to feed algorithms (Torney et al. 2019). The Galaxy Zoo project and the classification and identification of galaxy morphological shapes is a good case in point (Fortson et al. 2012; Walmsley et al. 2019). This procedure of combining cognitive skills and technical assignments is also called human computation. It is moreover an approach that has also been successfully tested in areas other than science (e.g. Google Maps). Human computation is used when it comes to the handling and classification of large, partly user-generated amounts of data.

Data has always been an intrinsic part of science and a rigorous methodology is needed to ensure data quality (see Balázs et al. in this volume), a topic extensively studied and discussed in the context of citizen science (Lukyanenko et al. 2019). With the advent of *big data*, not only is scientific resolution increasing, but so is

the ability to automate certain routine and repetitive tasks. The application of ML algorithms in the stage of data collection offers guidance in the subsequent analysis – identification and classification tasks – minimizing errors and maximizing data quality (Lukyanenko et al. 2019).

In other cases, the use of ML algorithms is applied once the data has been modelled, with the objective of analysing the model, extracting information, and giving responses to research questions. Standard statistical analysis, but also supervised and unsupervised learning (see below), is used to find causal relationships in the observations or look for patterns in the data collected (Vicens et al. 2018; Poncela-Casasnovas et al. 2016). At this point it is also possible to implement ML algorithms to detect biases in the data, such as location biases (Chen and Gomes 2018), or to analyse the influence of different explanatory factors in the model (Bird et al. 2014).

These introductory remarks already indicate that the application of ML is associated with various functions for science and for citizen science in particular. The aim of this chapter is, first, to give an overview of the application of ML in citizen science and, building on this, to explore the relationship between humans and machines in knowledge production. In the next section we will therefore present the current learning paradigms associated with ML, illustrated by sample projects, followed by a discussion of the main ethical challenges for citizen science that arise from the opacity of the algorithm from outside and how this imbalance can possibly be overcome. In the discussion section, we use these recent developments to identify the opportunities and challenges arising from collaboration between humans and machines in citizen science in the long run.

## Learning Paradigms in ML

To examine the tasks citizens are being invited to perform in citizen science projects through the use of ML, we need to see the learning paradigms associated to it. Currently, in the field of machine learning three main learning paradigms can be distinguished: supervised, unsupervised, and reinforcement (cf. Sathya and Abraham 2013). *Supervised learning* is based on training or teaching an algorithm using sample data - also called training data - already correctly classified by an expert. After that, the machine is provided with a new set of examples (data) so that a supervised learning algorithm can analyse the training data (set of training examples) and produce a solution from labelled data. *Unsupervised learning* is based on training a machine using unclassified data and allowing the algorithm to act on that data without any guidance. Unlike supervised learning, no classifications are provided which means no training is given to the machine. Therefore, the machine itself has to derive the hidden structure in unlabelled data. *Reinforcement learning* entails taking a suitable action to maximize rewards in a particular situation. It is employed by a variety of software and machines to find the best possible behaviour or path to be taken in a specific situation. While in supervised learning the algorithm is trained on data containing the correct answers, in reinforcement learning there is no answer but the reinforcement agent

decides what to do to perform the given task. In the absence of a training data set, the algorithm has to learn from its own experience. A form of ML that can use either supervised or unsupervised algorithms is *deep learning*. Deep learning can help solve certain types of difficult computer problems, most notably in computer vision/computer hearing and natural language processing (NLP). Computer vision or hearing defines a subset of AI which automatically extracts information from image, video, and audio data using algorithms (see Ceccaroni et al. 2019). The ‘deep’ in deep learning refers to the many layers that are built into a model, which are typically neural networks. A convolutional neural network (CNN) can consist of many layers of models, where each layer takes input from the previous layer, processes it, and outputs it to the next layer, in a daisy chain fashion. The probably most famous example of CNN is the one developed by Google’s DeepMind team, which beat the human world champion of the ancient Chinese game of Go.

### Examples of ML in Citizen Science

To provide some examples to our conceptual discussion, we reviewed a small convenience sample of nine citizen science applications using ML (Table 8.1). While we do not consider these projects to be representative of the entire population of ML applications in citizen science, they still offer some interesting indications.

**Table 8.1** Examples of ML in citizen science projects

Project Name	Category of Task	Types of AI - Purpose	Field	Machine Learning Paradigm
Galaxy Zoo	Classification	Computer vision - recognize shape and structure of galaxies.	Astronomy	Supervised learning and use of CNN
Human Protein Atlas	Classification	Computer vision - combination of pattern recognition and classification of images for large-scale classification of protein localisation patterns in microscopy images.	Life sciences	Supervised learning
iNaturalist	Classification	Computer vision -	Environment	Supervised

	tion	classification of images or acoustic data for species identification.	mental science	learning
Machine learning Assisted Image Annotation (MAIA)	Classification	Computer vision - image annotation method for environmental monitoring and exploration.	Oceanography	Supervised learning
Mindcontrol	Classification	Computer vision - scaling expertise in neuroimaging. Semi-automated replacement for manual coding of neuroimages through a mixture of crowdsourcing and machine learning. Citizen scientists are trained.	Neuroscience	Supervised learning and use of CNN
Nature through the eyes of many	Classification	Computer vision - pre-selection of animals trapped in camera trap databases and learning the species by citizen scientists' inputs.	Environmental science	Supervised learning
Observation.org	Classification	Computer vision - classification of images of species identification.	Environmental science	Supervised learning
SciStarter	Prediction	Automated reasoning - use of recommendation algorithms to predict users' interests based on their past interactions with SciStarter, a portal of 3,000 global citizen	Citizen Science Platform	Unsupervised learning in most cases, generally based on the project listed in the platform

		science projects, that aims to help users to find suitable projects.		
Snapshot Serengeti	Classification	Deep learning on structured data - automatic animal identification using deep learning to identify, count, and describe the behavior of animal species captured with camera traps.	Environmental science	Supervised learning

Most of the projects in Table 8.1 are examples of supervised learning in which algorithms are used that do not have a priori recognition abilities and, thus, need external training. Therefore, they usually start with a *golden set* of data labelled by domain human experts (e.g. Mindcontrol, A Shared Space and A Space for Sharing, and the Multiple Sclerosis course). Untrained citizens are then involved to use those labels for annotating a larger set of data, and, at the end, this larger data set is utilised to train a supervised machine learning model that automatically labels the entire data set.

Most projects in our list involve supervised learning by using image recognition software in the realm of computer vision. Computer vision is used on citizen science data and camera trap data to assist or replace citizen scientists in fine-grain image classification for taxon/species detection and identification (plant or animal) (Ceccaroni et al. 2019). A good example is the project Wildlife Insights (Ahumada et al. 2020) covering images from camera trap databases. Another example is a prototype called ‘Nature through the eyes of many’ as a utilised output from the project ‘National database of photo trap records’ (Lehejcek et al. 2019). Camera traps are commonly used in environmental monitoring, geography, and beyond (Trojan et al. 2019). Millions of pictures are collected throughout the extensive network of camera traps every day. This project combines pictures from various camera trap databases and serves as a management tool for the collected images. Like in the project Snapshot Serengeti (Swanson et al. 2016), machines are not always successful in identifying the proper animals in the collected pictures. In this case, citizen scientists in the role of spotters identify the animals and serve as teachers for the AI algorithms. Firstly, the AI will run the automatic classification of the picture. If the animal is detected with a certain probability, spotters come to the scene. AI offers a primary classification (animal recognition) to the spotter (also the trapper who uploaded records can pre-classify the image). A spotter validates/ invalidates the pre-classification and the image is not considered as validated until there is at least a 75% consensus (which can be adjusted in a certain project) among all the spotters involved. This is the input for the ML algorithms. The simplified schema of the whole process is visualised in Fig. 8.1, in which the parameters of the ML process are suppressed and

generalised. However, the example using a camera trap database could be analogically used in other ML in citizen science.

Although there are technical issues related to ML mechanisms, they can be utilised for gamification purposes increasing spotters' motivation levels and making participation for citizen scientists more attractive. Every spotter builds up their credibility. Spotters who tag pictures with greater consensus get higher weighting for their future votes, spotters who do not classify records well or who want to spoil the system are automatically weighted lower. The process of image classification involving AI/ML and citizens ends back with the trappers, who upload their camera trap data into the database. Trappers benefit from both AI and citizen science approaches and can easily manage their data within the database. The systems combining these methods in one place are in high demand, which can be substantiated by the support of big technology companies like Google (see the case of the Wildlife Insights project, Ahumada et al. 2020). For instance, national agencies for nature conservation and landscape protection using a significant amount of data from several remote camera trap repositories could manage the records in one place.

<FIG. 8.1 ABOUT HERE>

**Fig. 8.1** The interaction between spotters and ML processes during image classification within the camera trap database

## Challenges and Opportunities of Using ML

In the most applied ML paradigm – supervised learning – solutions are inferred directly from the data following the mathematical rules used to create such a paradigm (Sathya and Abraham 2013). Applications using this paradigm embed an idea of learning as acquisition or enhancement of knowledge to improve predictive accuracy or make more effective decisions (Blackwell 2015). This idea of learning builds on the strengths of machines, including performing tedious and repetitive tasks, fast processing of huge amounts of data, recognising complex patterns, and making predictions under uncertainty (Dellermann et al. 2018). Therefore, training ML models at high speed while maintaining accuracy and precision remains a vital goal for science. However, the application of ML in citizen science produces both epistemic and ethical challenges. Both have to do primarily with the opacity of the machine, whose operations and outcomes are largely obstructed by concrete human comprehension. Whether and how transparency can be created will be briefly discussed in the following section.

### Epistemological and Ethical Challenges

The extended use of AI, particularly ML, has initiated a general debate on the different forms of opacity (Burrell 2016) and bias (Mehrabi et al. 2019) that it promotes. Drawing from Burrell (2016), we use opacity to describe the difficulties encountered by a user of the output of an algorithm (for example, a classification decision) to make sense of how or why that particular classification has been arrived at from inputs. We use the term bias to refer to any prejudice or favoritism toward an individual or a group based on certain characteristics (Mehrabi et al. 2019).

The issues connected to opacity and bias in ML have brought to light the need for more transparency in the designing of algorithms and the data used for training in order to prevent or mitigate adverse effects. This consideration transcends citizen science and unequivocally affects every area in which the ML algorithms are applied. However, the very nature of citizen science projects and their possible biases mean that citizen science researchers devote much attention to ensuring data quality, a task which is even more important when we are using ML approaches.

Opacity in ML takes many forms but one of the most recently scrutinised is the black box effect. In general, a black box is a system in which we can observe the inputs and outputs but not the internal process. ML algorithms like neural networks and deep learning are so intrinsically complex that it is virtually unworkable to get to the bottom of their operations and internal decision-making processes. Those algorithms are designed to achieve the best performance possible given particular metrics, thus they are very useful when the cost of an error is low (Rudin 2019). This, for instance, happens when the consequences of unacceptable results are not significant or when the results are studied and validated in real applications (Doshi-Velez and Kim 2017). Nevertheless, the black box effect can cause biases and unfairness that impact human lives deeply. In those cases, it is advisable not to use opaque systems in high-stakes decisions regarding justice, healthcare, and employment, to mention just a few (Rudin 2019).

### **Making ML More Transparent**

To achieve further and sustained progress by the implementation of ML, explainable, interpretable, and comprehensible algorithms are needed to reduce biases (such as gender and racial biases), produced in both the design of the algorithms and the data used to train them

Concepts such as explainability, interpretability, and transparency are widely used in the AI literature and, in some cases, they have even been used interchangeably. Gilpin et al. (2018) state that an *explanation* can be evaluated in two ways: ‘according to its interpretability, and according to its completeness’ (p. 2), where *interpretability* describes the internal mechanisms of a system in a way that is understandable to humans, and *completeness* describes the operation of a system in an accurate way. Doshi-Velez and Kim (2017) define interpretability in machine learning as ‘the ability to explain or to present in understandable terms to a human’.

Therefore, the idea behind explainable AI radiates from the implementation of algorithms that are understandable to a human expert who can discern the internal mechanisms and understand what is happening. This idea is in contrast to black boxes. In a similar way, interpretable algorithms are the ones that allow the observation of the cause and effect in a system and predicting what is going to happen if there are changes in the input or in the algorithmic parameters.

In some citizen science projects, an explanation may not be required unless there is a decision-making process envisaged by the outcome. For instance, if we want to classify images that contain a whale and images which do not, in principle it is acceptable to use black box models. However, if from this outcome we need to make critical decisions, or we want to know how the decision-making process to detect whales works, then we would need an interpretable model. This is particularly critical in the context of citizen social science where we work with sensible social data, thus inferences in the analysis can have a direct impact on societal concerns. In this case, the generalisation bias means not only that data does not represent the whole context but, moreover, that data represents and reproduces situations of social injustice or prejudices. Transparent systems to avoid intentional bias (Burrell et al. 2016) in projects that involve sensitive data, such as biometric and genetic information, political opinions, and sexual orientation, need adopting mechanisms to ensure principles and guidelines regarding ethics: transparency, justice and fairness, non-maleficence, responsibility and privacy (Jobin et al. 2019; Floridi and Cowls 2019).

From these concerns about unfairness in ML emerges the ‘right to an explanation’, which basically states that a decision should not be based solely on automated decision-making, but also provide an explanation about the outcome of the decision-making process (Edwards and Veale 2018). The application of this principle in science is very interesting in the sense that the scientific understanding needs not only the outcome of the systems, but also the process leading to this outcome in order to extract knowledge from the procedure and be able to interpret it (Doshi-Velez and Kim 2017), let alone the possibility of replicating the results.

There are useful methods for explaining black box models (Guidotti et al. 2018) that can be applied to citizen science projects (for a more generalised way of differences between black box projects and transparent projects, see Fig. 8.2). The open research culture of citizen science is a perfect context to promote transparency within AI. The first step should be transparency of the forms of collaboration between humans and AI in citizen science. The main reason most humans are willing to give time and money (through energy consumption and use of their computers, for example, in online citizen science) is to help science and scientists build knowledge and therefore act for a better world. The implicit contract of citizen science builds on the premises of collaboration with scientists - not with artificial agents programmed to make use of data provided by human volunteers. However, citizen science projects do not always communicate clearly which use they make of the inputs of the volunteers. A minimum ethical requirement of online citizen science is therefore to make the process of human-AI collaboration explicit.

<FIG. 8.2 ABOUT HERE>

**Fig. 8.2** Differences between black box and transparent projects

Conversely, this ethical issue is also an opportunity to introduce AI to the participants of citizen science projects: through participating in hybrid intelligence activities and platforms that connect human intelligence and artificial intelligence to advance scientific knowledge, volunteers might get first-hand experience and better understanding of how AI works, and what are its requirements and limits, especially regarding the quality of structured data needed for the algorithms to be useful, the corresponding lack of relevance of AI to ambiguous problems, and the complexity of the black box and need to control it if AI is to contribute to decision-making on important social and medical matters. This opportunity to learn about AI is not widely shared, therefore citizen science can play a limited but critical role in helping citizens learn about algorithms. We could even imagine citizen science projects which would explicitly aim to help citizens learn about AI. In some projects (for example in Eyewire), volunteers can contribute as ML experts not only to use but also to design the project. The combination of a concrete experience of collaboration with AI, measuring its benefits and limits, and opportunities for social learning in the field of AI through citizen science projects is a promising path to spreading a democratic understanding of AI.

## Lessons Learnt

The use of ML continues to grow, but scholarly reflection and discussion on the role of ML in citizen science are still in their infancy. In other words, a solid research overview on this topic is complicated by the fact that not only is citizen science research not settled science, as Ceccaroni et al. (2019) argue in their review essay, but also by the fact that “AI is not settled science either; it inherently belongs to the frontier, not to the textbook” (ibid., p. 8).

To further explore this topic, we have therefore selected the aspects that appear most relevant to us from the recent research literature, without claiming to be exhaustive. These include, firstly, the *approach to machine learning*, which is inscribed in the various citizen science projects in different ways. Secondly, it was important to explore more closely the way in which humans and machines work together, as established through the use of ML. From a technical–scientific point of view, this is expressed by the term *human computation* based on the concept of distributed intelligence. The central question therefore is: what is the *division of labour between humans and machines*? In contrast to conventional cooperative relationships in research, however, it means that the actions of the machine remain invisible, from which special challenges are derived for citizen science that aims to increase transparency, algorithmic de-biasing, and fairness. One of these is the approach of *fair machine learning*, which we have discussed in the section on making ML more transparent as a possible solution to ML opacity. This example moreover shows that the plea for more transparency in algorithms is neither

limited to citizen science nor to science as such. Algorithms today affect all areas of social life and here lies a sociopolitical challenge as to how the interplay between humans and machines will be shaped in the future.

However, the main issue that arises when citizen science is considered in the context of ML is whether the machine should actually be regarded as a cooperative partner or rather as a competitor to human research activities. Currently, it is emphasised, and this is often part of the initial call for participation in citizen science projects, that certain tasks can be performed better and more efficiently by humans than by computers. Above all, however, in projects built around monitoring issues, in which the role of the citizen is primarily that of the human sensor (Haklay 2013) or in projects based on classifier-based models (Haklay 2013; Lintott and Reed 2013), the question arises to what extent these activities cannot be completely automated sooner or later, thus rendering superfluous not only citizens but also professional scientists (Franzen 2019). If the machine learns to perform more and more tasks reliably, the question is where to look for the role of the citizen (and the human) in the future.

One possible response would be to involve citizens in scientific research for even more demanding activities, which is in line with the normative expectations of citizen science. In Haklay's typology of citizen participation, ranging from participatory sensing to collaborative science (2013), this would mean allowing non-scientists to participate in higher levels of citizen science other than crowdsourcing, up to the generation of having their own research projects by defining research problems (Haklay 2013). Particularly in view of today's increasingly data-driven research landscape, participation in citizen science would then not only depend on the *digital literacy* of the participants (with regard to the use of smartphones and apps, such as many of the citizen science projects provide), but would also require *code literacy* in order to actually exploit ML for this type of bottom-up research in citizen science.

In the context of human computation, however, there is another reflexive component, which is named analogously as *data literacy* but is rarely discussed in the discourse on citizen science. Volunteers should at least be aware that as soon as they participate in data-driven citizen science projects, they themselves become data that might be processed further. For the purpose of increasing data quality, user performance is not only recorded automatically in the systems, but is also partly used as a weighting factor in classification projects (e.g. Galaxy Zoo) or as information about the participant, in order to keep him or her on their toes, depending on the required commitment profile (cf. Lintott and Reed 2013). Since citizen science is primarily designed to advance collective knowledge, it is important to enlighten potential participants about the handling of user-generated data as it is demanded in all other areas of an increasingly datafied society (see e.g. the 'manifesto' for the 'public understanding of big data', posted by Michael and Lupton 2015).

We should therefore remember, that learning has a double meaning in this context: through the classification activities mostly carried out in large-scale citizen science projects, not only can the participants possibly learn something

about science but the machine also learns something about human actions in order to imitate them first and possibly exceed them sooner or later.

## **Future Trends, Recommendations, and Conclusions**

The processing power and sophistication of algorithms have improved at previously unimaginable levels and some ML techniques have already outperformed or at least paralleled human capabilities. Google-owned AI specialist, DeepMind, claimed a new milestone in being able to demonstrate the usefulness of AI to help with the task of predicting 3D structures of proteins based solely on their genetic sequence. Google's new algorithm AlphaFold showed at the last biannual protein-folding olympics that it is more efficient than humans in predicting protein structure based on amino acids (Sample 2018). In Galaxy Zoo, the use of CNN to classify galaxies led to impressive results in a task previously considered performed better by humans. Despite these remarkable achievements, there are still problems that machines cannot solve alone, such as those involving creative tasks or using expertise in decision-making (Dellerman et al. 2019).

As Watson and Floridi (2018) pointed out, 'We cannot be certain just what scientific developments the future holds in store, but we can be confident that many of our next great discoveries will be made thanks to some complex partnership of minds and machines' (p. 760). We must not forget that we are thus dealing with the question of development of science and society as a whole, even if we discuss the question of ML here using the example of citizen science. Recourse to the normative foundations of citizen science is, then, helpful in providing concrete indications for the thrust and democratic design of a socially desirable socio-technical development in the age of AI. Precisely because citizen scientists are volunteers who donate their time to 'help science', compliance with research ethics guidelines in the handling of personal data is a top priority. At the same time, however, citizen science projects might become forerunners in the drive to break down the opacity of algorithms as far as possible in favour of education and enlightenment. Whether approaches like local-interpretable-model-agnostic explanations (see Ceccarino et al. 2019, p. 8) are already sufficient to increase model transparency should be further discussed, not only in academia but also with citizens. For developers of future citizen science projects in the context of AI, the crucial question is therefore how to involve and motivate citizens not only in the processing of data, but also how to educate them and reward them for their work. With regard to volunteer monitoring projects, Ceccarino et al. have summarised the concern as follows: 'How do we acknowledge, respect, and reward the people whose data and expertise have helped to train the computer-vision algorithms?' (2019, p. 2). While these dimensions correspond to the normative structure of science, the question arises to what extent the principles for dealing with data and self-learning algorithms can or should be applied to other social sectors, especially if the aims are to be used for commercial or political/regulatory goals.

This kind of reasoning is part of a sociopolitical debate that needs to be conducted on a broad scale, because, with all the promises associated with AI, an informed view of the risks must not be neglected in order to shape socio-technical development for the common good. This addresses the scientific responsibility, not only of computer scientists and IT developers, but also of other (social) scientists and citizen science researchers, when it comes to applying ML to scientific knowledge production.

The rapid progress in the development of computing capacities – see Google’s major breakthrough with its new quantum computer (Arute et al. 2019) – means that we run the risk of being unable to keep up with the reflexive consideration of its significance and impact on science and society. This brings us to our last point: the success of citizen science and ML in citizen science therefore depends on the technical and financial resources available now and in the future for this type of research.

## References

- Ahumada, J.A., Fegraus, E., Birch, T., Fores, N., Kays, R., O’Brien, T.G. et al. (2020). Wildlife insights: A platform to maximize the potential of camera trap and other passive sensor wildlife data for the planet. *Environmental Conservation*, 47(1), 1–6.
- Arute, F., Arya, K., Babbush, R., Bacon, D., Bardin, J.C., Barends, R., et al. (2019). Quantum supremacy using a programmable superconducting processor, *Nature*, 574(7779), 505–510. <https://doi.org/10.1038/s41586-019-1666-5>.
- Bird, T.J., Bates, A.E., Lefcheck, J.S., Hill, N.A., Thomson, R.J., Edgar, G.J., et al. (2014). Statistical solutions for error and bias in global citizen science datasets. *Biological Conservation*, 173, 144–154.
- Blackwell, A. (2015). Interacting with an inferred world: the challenge of machine learning for humane computer interaction. *Aarhus Series on Human Centered Computing*, 1(1), 12. <https://doi.org/10.7146/aahcc.v1i1.21197>.
- Bonney, R., Ballard, H., Jordan, R., McCallie, E., Phillips, T., Shirk, J., Wilderman, C.C. (2009). *Public participation in scientific research: Defining the field and assessing its potential for informal science education. A CAISE Inquiry Group Report*. Washington, DC: Center for Advancement of Informal Science Education (CAISE). [https://safmc.net/wp-content/uploads/2016/06/Bonneyetal2009\\_PPSRExecutiveSummary.pdf](https://safmc.net/wp-content/uploads/2016/06/Bonneyetal2009_PPSRExecutiveSummary.pdf). Accessed 21 February 2020.
- Burrell, J. (2016). How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1). <https://doi.org/10.1177/2053951715622512>.
- Ceccaroni, L., Bibby, J., Roger, E., Flemons, P., Michael, K., Fagan, L., Oliver, J.L. (2019). Opportunities and risks for citizen science in the age of artificial intelligence. *Citizen Science: Theory and Practice*, 4(1), 29. <https://doi.org/10.5334/cstp.241>.

Chen, D., & Gomes, C.P. (2018). Bias reduction via end-to-end shift learning: Application to citizen science. <https://arxiv.org/pdf/1811.00458v4.pdf>. Accessed 21 February 2020.

Dellermann, D., Calma, A., Lipusch, N., Weber, T., Weigel, S., & Ebel, P. (2019). The future of human-AI collaboration: A taxonomy of design knowledge for hybrid intelligence systems. In T. Bui (Ed.), *Proceedings of the Hawaii International Conference on System Sciences (HICSS)* (pp. 1–10). ScholarSpace/AIS Electronic Library (AISeL). [https://www.alexandria.unisg.ch/254994/1/JML\\_706.pdf](https://www.alexandria.unisg.ch/254994/1/JML_706.pdf) Accessed 21 February 2020.

Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. <http://arxiv.org/abs/1702.08608>. Accessed 21 February 2020.

Edwards, L., & Veale, M. (2018). Enslaving the algorithm: From a ‘right to an explanation’ to a ‘right to better decisions’? *IEEE Security and Privacy*, 16(3), 46–54. <https://doi.org/10.1109/MSP.2018.2701152>.

Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1). <https://doi.org/10.1162/99608f92.8cd550d1>

Fortson, L., Masters, K., Nichol, R., Borne, K., Edmondson, E., Lintott, C., et al. (2012). Galaxy Zoo: Morphological classification and citizen science. <https://arxiv.org/pdf/1104.5513.pdf>. Accessed 21 February 2020.

Franzen, M. (2019). Changing science-society relations in the digital age: The citizen science movement and its broader implications. In D. Simon, S. Kuhlmann, J. Stamm, & W. Canzler (Eds.), *Handbook on science and public policy* (pp. 336–356). Cheltenham: Edward Elgar.

Gilpin, L. H., Bau, D., Yuan, B. Z., Bajwa, A., Specter, M., & Kagal, L. (2019). Explaining explanations: An overview of interpretability of machine learning. In *Proceedings - 2018 IEEE 5<sup>th</sup> International Conference on Data Science and Advanced Analytics - DSAA 2018* (pp. 80–89). <https://doi.org/10.1109/DSAA.2018.00018>.

Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5), Article 93, 1–42. <https://doi.org/10.1145/3236009>.

Haklay, M. (2013). Citizen science and volunteered geographic information: overview and typology of participation. In D. Z. Sui, S. Elwood, & M. Goodchild (Eds.), *Crowdsourcing geographic knowledge* (pp. 105–122). Dordrecht: Springer.

Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399.

Jordan, M.I., & Mitchell, T.M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260.

Lehejcek, J., Adam, M., Tomasek, P. & Trojan, J. (2019). Informacni system pro spravu fotopasti [National database of photo trap records]. <http://okovprirode.cz/>. Accessed 21 February 2020.

Lintott, C., & Reed, J. (2013). Human computation in citizen science. In P. Michelucci (Ed.), *Handbook of human computation* (pp. 153–162). New York, NY: Springer.

Lukyanenko, R., Wiggins, A., & Rosser, H.K. (2019). Citizen science: An information quality research frontier. *Information Systems Frontiers*. <https://doi.org/10.1007/s10796-019-09915-z>.

Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2019). A survey on bias and fairness in machine learning. <https://arxiv.org/pdf/1908.09635.pdf>. Accessed 21 February 2020

Michael, M., & Lupton, D. (2015). Toward a manifesto for the ‘public understanding of big data’. *Public Understanding of Science*, 25, 104–116. <http://dx.doi.org/10.1177/0963662515609005>.

Neal, L. (2013). The ‘human sensor’. Bridging between human data and services. In P. Michelucci (Ed.), *Handbook of human computation* (pp. 581–593). New York, NY: Springer.

Poncela-Casasnovas, J., Gutiérrez-Roig, M., Gracia-Lázaro, C., Vicens, J., Gómez-Gardeñes, J., Perelló, J., et al. (2016). Humans display a reduced set of consistent behavioral phenotypes in dyadic games. *Science Advances*, 2(8), 1–9. <https://doi.org/10.1126/sciadv.1600451>.

Popenici, S., & Kerr, S. (2017). Exploring the impact of artificial intelligence on teaching and learning in higher education. *Research and Practice in Technology Enhanced Learning*, 12(22). <https://doi.org/10.1186/s41039-017-0062-8>.

Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215. <https://doi.org/10.1038/s42256-019-0048-x>.

Sample, Ian (2018). Guardian online Dec 2, 2018: <https://www.theguardian.com/science/2018/dec/02/google-deepminds-ai-program-alpha-fold-predicts-3d-shapes-of-proteins>

Sathya, R., & Abraham, A. (2013). Comparison of supervised and unsupervised learning algorithms for pattern classification. *International Journal of Advanced Research in Artificial Intelligence*, 2(2), 34–38.

Sullivan, D.P., Winsnes, C.F., Åkesson, L., Hjelmare, M., Wiking, M., Schutten, R., et al. (2018). Deep learning is combined with massive-scale citizen science to improve large-scale image classification. *Nature Biotechnology*, 36(9), 820–832. <https://doi.org/10.1038/nbt.4225>.

Swanson, A., Kosmala, M., Lintott, C., Packer, C. (2016). A generalized approach for producing, quantifying, and validating citizen science data from wildlife images. *Conservation Biology*, 30(3), 520–531.

Torney, C.J., Lloyd-Jones, D.J., Chevallier, M., Moyer, D.C., Maliti, H.T., Mwita, M., et al. (2019). A comparison of deep learning and citizen science techniques for counting wildlife in aerial survey images. *Methods in Ecology and Evolution*, 10(6), 779–787. <https://doi.org/10.1111/2041-210X.13165>.

Trojan, J., Schade, S., Lemmens, R., & Frantál, B. (2019). Citizen science as a new approach in geography and beyond: Review and reflections. *Moravian Geographical Reports*, 27(4), 254–264. <https://doi.org/10.2478/mgr-2019-0020>.

Vicens, J., Bueno-Guerra, N., Gutiérrez-Roig, M., Gracia-Lázaro, C., Gómez-Gardeñes, J., Perelló, J., et al. (2018). Resource heterogeneity leads to unjust effort

distribution in climate change mitigation. *PLoS ONE*, *13*(10), 1–17. <https://doi.org/10.1371/journal.pone.0204369>.

Walmsley, M., Smith, L., Lintott, C., Gal, Y., Bamford, S., Dickinson, H., et al. (2019). Galaxy Zoo: Probabilistic morphology through Bayesian CNNs and active learning. *Monthly Notices of the Royal Astronomical Society*, *491*(2), 1554–1574. <https://doi.org/10.1093/mnras/stz2816>.

Watson, D., & Floridi, L. (2018). Crowdsourced science: Sociotechnical epistemology in the e-research paradigm. *Synthese*, *195*, 741–764. <https://doi.org/10.1007/s11229-016-1238-2>.

Willi, M., Pitman, R.T., Cardoso, A.W., Locke, C., Swanson, A., Boyer, A., et al. (2019). Identifying animal species in camera trap images using deep learning and citizen science. *Methods in Ecology and Evolution*, *10*(1), 80–91. <https://doi.org/10.1111/2041-210X.13099>.