ESSAY

Opportunities and Risks for Citizen Science in the Age of Artificial Intelligence

Luigi Ceccaroni*, James Bibby†, Erin Roger‡, Paul Flemons§, Katina Michael‖, Laura Fagan¶ and Jessica L. Oliver**

Members of the public are making substantial contributions to science as citizen scientists, and advances in technologies have enabled citizens to make even more substantial contributions. Technologies that allow computers and machines to function in an intelligent manner, often referred to as artificial intelligence (AI), are now being applied in citizen science. Discussions about guidelines, responsibilities, and ethics of AI usage are already happening outside the field of citizen science. We suggest such considerations should also be explored carefully in the context of citizen science applications. To start the conversation, we offer the citizen science community an essay to introduce the state-of-play for AI in citizen science and its potential uses in the future. We begin by presenting a systematic overview of AI technologies currently being applied, highlighting exemplary projects for each technology type described. We then discuss how AI is likely to be increasingly utilised in citizen science into the future, and, through scenarios, we explore both future opportunities and potential risks. Lastly, we conclude by providing recommendations that warrant consideration by the citizen science community, such as developing a data stewardship plan to inform citizens in advance of plans and expected outcomes of using data for AI training, or adopting good practice around anonymity. Our intent is for this essay to lead to further critical discussions among citizen science practitioners, which is needed for responsible, ethical, and useful use of AI in citizen science.

Introduction

Technologies that allow computers and machines to perform tasks normally requiring human intelligence are often referred to as artificial intelligence (AI). These technologies allow machines to complete tasks with traits or capabilities ordinarily associated with human cognition, such as reasoning, problem solving, common-sense knowledge management, planning, learning, translation, perception, vision, speech recognition, and social intelligence (Kaplan and Haenlein 2019). Research in AI is rapidly increasing, as indicated when comparing the annual publishing rate of papers focused on AI between 1996 and 2017 against the publishing rates of papers focused on any topic or against the pub-

lishing rates of papers in the field of computer science (see the growth of annually published papers by topic in Shoham et al. [2018; p. 9]). This growth in AI publications has prompted researchers to critically explore the potential promises and risks of AI (Scherer 2016; Webb 2019; Yudkowsky 2008) as well as ethics and responsibilities (Miller 2019; Cowls and Floridi 2018; Scherer 2016; Dawson et al. 2019).

AI has been used in citizen science projects for about 20 years. It was first used in this context in 2000, in collaborative AI databases such as the Generic Artificial Consciousness (GAC)/Mindpixel Digital Mind Modeling Project (McKinstry 2009) and the Open Mind Common Sense project (Singh et al. 2002). In these models, user-submitted propositions were meant to create a database of common-sense knowledge that could function as a kind of digital brain. This relationship between collective knowledge and algorithmic processing evolved in many directions and, in 2019, is predominantly represented by machine learning, especially applied to computer vision, which includes diverse methods of automatically identifying objects from digital photographs. For example, the iNaturalist platform, a citizen science project and online social network, is designed to enable citizen scientists and ecologists alike to upload observations from the natural world, such as images of animals and plants (Van Horn et al. 2018). The platform is one among many (Wäldchen et al. 2018) that include an automated

* Earthwatch, GB

† Centre for Advanced Analytics and Economics, Environment, Energy and Science, NSW Department of Planning, Industry and Environment, AU

‡ Australian Citizen Science Association, AU

§ Australian Museum, AU

‖ University of Wollongong, AU

¶ Department of Primary Industries and Regional Development, AU

** Queensland University of Technology, AU

Corresponding author: Luigi Ceccaroni (lceccaroni@earthwatch.org.uk)

species/taxon-identification machine-learning algorithm applied to computer vision (Weinstein 2018). Images can be identified via an AI model that has been trained on the large database of "research grade" observations on iNaturalist (Bowser et al. 2014; [https://www.inaturalist.org/pages/help#quality]).

The same types of machine-learning algorithms used by iNaturalist's community of users are also helping ecologists to classify millions of underwater snapshots of corals via the XL Catlin Global Reef Record project (Tollefson 2016). Currently, AI researchers, whether in citizen science or more broadly, tend to test their algorithms on a few standard data sets. For instance, image-recognition software is generally tested on ImageNet (for examples see Shoham et al. 2018; p. 47), a database of around 14 million photographs (Russakovsky et al. 2015) including people, scenes, and objects, as well as plants and animals. In the field of biodiversity, in 2017 iNaturalist made one of its data sets of 5,000 photographs of birds, mammals, amphibians, and other taxonomic groups available for attendees of the Computer Vision and Pattern Recognition Conference in Honolulu, Hawaii, to train and test computer-vision algorithms (Joppa 2017).

With the proliferation of connected devices and increased data collection, AI technology has the potential to dramatically impact society, including business and the workforce. The benefits of a prudent and planned approach to AI are manifold, from increasing user engagement in scientific activities to producing better scientific outcomes. As with any endeavour that could impact human well-being, it is important to examine the risks and opportunities of AI before developing citizen science projects that include it, in order to make informed decisions. For example, before we design and deploy computer-vision technology, we may want to ask the question: How do we acknowledge, respect, and reward the people whose data and expertise have helped to train the computer-vision algorithms? Data in citizen science are usually open and accessible to participants. However, to prevent the concentration of wealth and power in the hands of the AI companies controlling the data-processing technology, the regulation of data ownership requires more thought. If access to AI resources is restricted by commercial interests, data contributors (i.e., citizens) may be excluded from decisions about data use or from involvement in research that uses AI. Therefore, it is important that AI computing resources are openly accessible and available to all, creating opportunities for citizens to be involved in AI research and to understand how the data they collect are being used.

Intergovernmental agencies, technologists, and conservationists have identified the need to coordinate the creation and use of technologies to solve global problems (Campbell and Jensen 2019; Lahoz-Monfort et al. 2019). The citizen science community is well positioned to contribute in a variety of ways to global coordination initiatives, such as the United Nations Sustainable Development Goals (https://sustainabledevelopment.un.org/), whether through providing methodologies or contributing data not otherwise obtainable (See et al. 2019). Innovative solutions such as AI are required to make meaning of large datasets, and citizen science has a significant role to play in ensuring that data are collected, analysed, and interpreted in meaningful ways that benefit everyone. Here, we provide a systematic overview of AI technologies currently being implemented in citizen science. We then explore potential opportunities and risks that may arise as technologies evolve. Lastly, we provide recommendations to ensure that the opportunities and risks of AI use are adequately identified. It is our intention for this article to serve as a practical introduction to how AI is used in citizen science, and for it to elicit more in-depth discussions about AI use by members of the citizen science community.

Our Approach for This Essay

To explore the current use, opportunities, and risks of AI in citizen science, we elected to conduct a systematic overview (Grant and Booth 2009) of the use of AI in citizen science. Our overview is intended to provide readers with a broad understanding of AI and its applicability to citizen science, rather than providing an exhaustive list of citizen science projects applying AI. We did, however, want to ensure that we captured the diversity of AI technologies being included in citizen science. To develop a broad understanding of current AI use in citizen science, we queried two technology-focused academic literature databases, the Association of Computer Machinery Digital Library (ACM DL: [https://dl.acm.org/]) and the Institute of Electrical and Electronics Engineers (IEEE Xplore: [https://ieeexplore.ieee.org]) databases, using the terms "artificial intelligence" and "citizen science." The ACM DL and IEEE Xplore databases returned 92 and 8 articles respectively. We reviewed these articles to understand whether and how AI was being implemented, without making an assessment of the quality of the research, as this was not relevant to our aims. We identified that some form of AI used in citizen science was found in 50 and 6 articles from ACM DL and IEEE Xplore databases respectively. We identified the following types of AI in those papers: Automated reasoning and machine learning; computer vision and computer hearing; knowledge representation and ontologies; natural language processing; and robotic systems. These types are defined and described below. Given the interdisciplinary nature of citizen science research and associated publishing, we supplemented ACM DL and IEEE Xplore database query results with additional peer-reviewed literature by drawing from our collective knowledge. The authors are involved in citizen science globally, with particularly extensive knowledge of projects across Europe, Australia, and the United States. We decided that, for a specific AI technology to be considered currently applied in citizen science, some articles explicitly discussing its use in a citizen science project must be published in academic literature.

Current Applications of AI in Citizen Science

In this section we provide an overview of citizen science, AI, and how the two currently interplay. To set the stage, we begin by broadly describing citizen science and AI. Then we describe the types of AI already being applied in citizen science and highlight the use of these technologies by describing associated exemplary projects.

Citizen science can be described as work undertaken by civic educators and scientists together with citizen communities to advance science, foster a broad scientific mentality, and/or encourage democratic engagement, which allows society to deal rationally with complex modern problems (Ceccaroni et al. 2017). Put simply, it involves public participation and collaboration in scientific research with the aim of increasing scientific knowledge. The citizen science community occasionally uses supporting technologies that allow computers and machines to function in an intelligent manner, to achieve particular traits or capabilities often associated with AI.

AI can be described as intelligence demonstrated by machines, in contrast to the natural intelligence displayed by humans and other animals. In computer science, AI research is defined as the study of "intelligent agents," which are any devices that perceive their environment and take actions that maximise the chance of successfully achieving goals (Poole et al. 1998). Colloquially, the expression artificial intelligence is applied when a machine mimics cognitive functions that people associate with human minds, such as learning and problem solving (Russell and Norvig 2016). AI can be a challenging concept for humans (Sterne 2017). Intrinsically, humans want to believe that the wonders of the mind (for example, in identifying species or sounds) are inaccessible by material processes—that minds are, if not literally miraculous, then mysterious in ways that defy natural science. This is, among other motives, because of something truly unsettling to a human mind: Competence without comprehension (Dennett 2017).

Below we provide a description of the technologies commonly used in citizen science that allow machines to complete tasks and achieve particular traits or capabilities that are often referred to as AI, such as machine learning. Real-world examples are provided, with references, so that people less familiar with the AI technologies will have a way to conceptualise use of these AI types and their impacts.

Automated reasoning and machine learning

Automated reasoning is an area of computer science and mathematical logic dedicated to understanding different aspects of reasoning. Automated reasoning helps to produce computer programs that allow computers to reason semiautomatically, or entirely automatically. Machine learning uses statistical techniques to give computers the ability to "learn" (i.e., progressively improve performance on a specific task) with data. With machine learning, programs can be designed to learn things on their own. One program, for example, can learn to detect a specimen of a specific taxon in a picture. It is not necessary to tell the program whether each picture has a specimen of that specific taxon in it or not; the program will learn that itself using machine learning. A motivation for research in this area, for example, is the desire to design programs that simulate empathy and improve the program's understanding of human nature (Kido and Swan 2015). The machine interprets the emotional state of humans and adapts its behaviour to them, in an attempt to give an appropriate response to the human's emotional state (Picard 1995; Jaques et al. 2016; Herzig et al. 2017; Feffer et al. 2018). One common machine-learning approach involves the application of deep-learning techniques (or artificial neural networks), which have been shown to be effective and efficient in addressing classification-type problems such as identifying objects or categorising digital imagery.

Computer vision and computer hearing

Computer vision and hearing are interdisciplinary fields that explore how computer algorithms and systems can classify and/or identify content and achieve high-level understanding from digital images, videos, or audio recordings. They could broadly be called a subfield of AI and machine learning, which may involve the use of specialised methods and make use of general learning algorithms. We distinguish computer vision from machine learning because of the high number of applications using computer vision specifically, but we would like to make clear that they are not separate fields of research. Computer vision and computer hearing are 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). A good example of this is iNaturalist (discussed above), built on the concept of mapping and sharing observations of biodiversity across the globe. As of July 2018, iNaturalist users have contributed more than 14,000,000 observations of plants, animals, and other organisms worldwide. In addition to observations being identified by the user community, iNaturalist includes an automated species identification tool based on computer vision. Images can be identified via an AI model, which has been trained on the large database of "research grade" observations on iNaturalist (Bowser et al. 2014). A broader taxon such as a genus or family is typically provided if the model cannot decide what the species is. If the image has poor lighting, is blurry, or contains multiple subjects, it can be difficult for the model to determine the species and it may decide incorrectly. Multiple species suggestions are typically provided, with the species that the algorithm believes the image to be most closely aligned placed at the top of the list of suggested matches. iNaturalist still relies on experts to validate users' recordings, but deep convolutional neural networks are reducing the amount of repetitive expert-input required. Currently, limited availability of experts remains one of the biggest bottlenecks in the growth of validated user observations (Joppa 2017). Computer vision and computer hearing also can be used to automatically annotate previously collected data on undescribed or undiscovered species (Le et al. 2013; Sun et al. 2017).

Knowledge representation and ontologies

Knowledge representation is the field of AI dedicated to representing information about the world in a form that a computer system can utilize to solve complex tasks such as assessing environmental impact or having a dialog in a natural language.

"Ontology," in philosophy, refers to the set of "things" that a person believes to exist. In AI, it has proven more than convenient to extend the term "ontology" beyond this primary meaning and use it for the set of "things" that a computer program must be able to deal with to do its job (Dennett 2017). An ontology then encompasses a representation of the categories, properties, and relations among the concepts, data, and entities of a domain (Ceccaroni et al. 2017). Several organisations work on the development of a recommendation on how to represent data and metadata in citizen science. This work is based on previous efforts by the US Citizen Science Association's (CSA) international Data and Metadata Working Group. The Group's aim is to promote collaboration in citizen science through the development and/or improvement of international standards for citizen science data and metadata. This working group collaborates on citizen science at the international level, and became a coordinating and umbrella group crossing many thematic and geographically distributed organisations that provide relevant complementary work. Contributions have been provided by the European Citizen Science Association (ECSA), the COST Action 15212's Working Group 5 ("Improve data standardization and interoperability"), and the Australian Citizen Science Association (ACSA). These organisations also address the definition of interoperability standards for data exchange, reusability, and compatibility in citizen science. They contributed to defining core building blocks of these interoperability standards, and outlined the way ahead based on the CSA Data and Metadata Working Group's previous work. Providing guidance on how to use standards across communities with varying knowledge and technical expertise will support uptake of project results and improve project sustainability.

Natural language processing

Natural language processing (NLP) is an area of computer science and AI concerned with the interactions between computers and human (natural) languages. In particular, NLP considers how to program computers to process and analyse large amounts of natural language data (Deng et al. 2012).

Robotic systems

Robotics is an interdisciplinary branch of engineering and science that includes mechanical engineering, electronic engineering, information engineering, computer science, and others. Robotics deals with the design, construction, operation, and use of robots and computer systems for their control, sensory feedback, and information processing abilities (Joshi et al. 2018).

Categorisation of Applied Uses of AI

As discussed above, there are a number of types of AI techniques, and a number of ways in which each type can be applied across science disciplines (e.g., Rogers and Aikawa 2019; Hecht 2018; Korot et al. 2019). To better understand how AI is currently used in citizen science, and the possible extensions of its current use into the future, we divided uses into three broad and overlapping categories (**Tables 1** and **2**). These categories are arbitrary and group an otherwise long list of uses. The first category is *assisting or replacing humans in completing tasks*, which means that AI is enabling tasks traditionally done by people to be partly or completely automated. The second category of AI use is associated with *influencing human behaviour*. Human behaviour is a major source of data in the current digital economy and in the training of AI. At the same time, it is also one of the main objects of data science in the sense that many data science, AI, and citizen science models are aimed at influencing human behaviour (e.g., through personalisation and behavioural segmentation, or providing people a means to be comfortable with citizen science and get involved). The third category of AI use relates to having *improved insights* as a result of using AI to enhance data analysis. For example, AI can now offer greater insights from data to inform research and policies, thanks to the training of computer-vision and computer-hearing algorithms using citizen science data. AI also can facilitate sharing the meaning of terms among machines thanks to the use of ontologies.

Future Applications of AI in Citizen Science

In addition to more people integrating AI into a wider diversity of projects and improvement of existing methods, we foresee a wider array of AI technologies being applied to citizen science, which we explore in the section below. We have created two scenarios relating to different potentials of AI to impact citizen science and potentially society more broadly. The first scenario describes a future in which AI competence is inferior to human competence in relation to citizen science tasks. The second scenario describes a future in which AI competence equals or surpasses human capability in relation to citizen science (Barrat 2013).

Scenario one: AI for engaging citizens

Imagine we have a project with a large dataset of images, and computer scientists apply computer vision to identify objects of interest from images. Citizen scientists can be engaged to identify objects and train algorithms to improve their accuracy rates. Apart from improving its automated image classification, AI proves a very effective tool for engaging and connecting people to science. AI benefits the amateur participants and creates a more inclusive, inspiring, and impactful scientific practice.

Scenario two: AI for engaging citizens and as basis for new applications

Imagine a similar scenario as the one outlined above, though a key difference is that AI computer-vision techniques can identify objects in images with a competence equal to or superior to human competence: AI tools can instantly analyse and identify animals

Table 1: Summary of the categories of AI used in citizen science and their applications. (The list of categories is not ranked in terms of importance.)

Description of instances where AI is applied	Types of AI	Examples of citizen-science software-applications
Applied use and impact: Assisting or replacing humans in completing tasks		
Improving image or audio classification	Computer vision and computer hearing	Computer vision and computer hearing can be applied to photographic images (e.g., from cameras that are triggered by motion detection) or acoustic data, to assist or replace citizen scientists in classifying images or sounds for species detection and identification (Parham et al. 2018). Examples include citizen science biodiversity project iNaturalist (Joppa 2017; Van Hon et al. 2018); improvement of species monitoring and automatic annotation of previously collected data on undescribed or undiscovered species (Sun et al. 2017; Sullivan et al. 2018); and automatic detection of acoustic events such as bat vocalisations from audio recordings (Mac Aodha et al. 2018).
Accelerating the digitization of biodiversity research specimens	Computer vision and computer hearing	In digitising museum specimens, computer vision can assist citizens with tasks related to identifying labels, sorting handwritten versus typed labels, capturing label data, parsing information into field notes, normalising data, and minimising duplication. Examples include Leafsnap, for the identification of tree species in the North-eastern United States (Kumar et al. 2012); SPIDA, for the identification of one family of Australasian ground spiders (Russell et al. 2007).
Verifying the accuracy and consistency of contributors' submissions	Automated reasoning and machine learning	The citizen-science biodiversity projects eBird (Sullivan et al. 2014) and iNaturalist.
Providing more rapid response to complex modern problems	Automated reasoning and machine learning	The citizen-science monitoring project Citclops for early warning of harmful algal blooms (Ceccaroni et al. 2018).
Applied use and impact: Influencing human behaviour		
Extend social impact of citizen science	Robotic systems	A community-oriented robotic system designed to extend the social, educational, economic, and health benefits from citizen science to a more general public (Joshi et al. 2018).
Using social media for collaborative species identification and occurrence	Natural language processing, Knowledge representation and ontologies	Using specific social media to engage participants in contributing their observations over a long time-period (Deng et al. 2012).
Applied use and impact: Improving insights		
Training of computer-vision and computer-hearing algorithms using citizen-science data	Computer vision and computer hearing	Data collected by citizens are used by knowledge engineers, people who integrate knowledge into computer systems to solve complex problems normally requiring a high level of human expertise, to train AIs. Examples include citizen-science biodiversity projects iNaturalist (Van Horn et al. 2018), Leafsnap and Pl@ntNet (as discussed in Bonnet et al. 2016).
Facilitating sharing the meaning of terms	Knowledge representation and ontologies	Citizen-science associations and projects based in the US, Europe, and Australia working together to design an ontology to represent knowledge in the domain of citizen science (Storksdieck et al. 2016).
Mining social-network data	Natural language processing	Citizen science projects can collect and analyse Twitter/Google data about health or the environment. An example is Aurorasaurus, a project to collect auroral observations (MacDonald et al. 2015).

and plants in our environment, without the need for human-based methods of classification. In this case, not only is AI a tool to engage citizens, it also opens the possibility of creating new applications based on automatic nature classification.

Opportunity exploration
The positive impact of AI is clear from Scenario one, with AI proving an effective tool to engage and connect people to science. Positive impact related to Scenario two is potentially less clear, if the "human training AI" relationship is

Table 2: Summary of new applications of AI in citizen science likely to appear in the near future.

Description of instances where AI is likely to be applied	Types of AI	Examples of citizen science software applications
Applied use and impact: Assisting or replacing humans in completing tasks		
Filtering out hard, repetitive, routine, or mundane tasks	Automated reasoning and machine learning	Software applications that allow citizen scientists to focus on more engaging tasks, for example, focusing on observations of interactions, or developing/contributing to innovative projects in the field.
Providing training/support	Automated reasoning and machine learning	AI systems that can be used in regions where citizen science training/support by humans is limited, such as when direct access to people with expertise is limited and/or human-language barriers exist.
Identifying species	Computer vision and computer hearing	AI tools that can instantly classify species based on images or sounds.
Applied use and impact: Influencing human behaviour		
Describing and formally representing the domain of citizen science in all languages	Knowledge representation and ontologies	An ontology that can facilitate the creation of new citizen science applications in any language and the translation of existing applications into any language.
Making information and data more accessible in citizen science applications	Automated reasoning and machine learning; Natural language processing	Applications using machine learning and natural language processing to overcome information overload in citizen science platforms.
Providing an easy, engaging, and enjoyable citizen scientist experience with AI-based virtual assistance	Automated reasoning and machine learning	Virtual/simulated environments, in which citizens interact with AI to test tasks before real-world deployment.
Notifying citizens about what is likely to occur near them or what/when they could observe	Automated reasoning and machine learning	Mobile apps providing satellite-based information to citizen scientists (e.g., satellite-overpass maps). Applications that provide contextual information to citizens: What is measured, why, when, and where.
Adaptively managing and changing citizen science activities	Automated reasoning and machine learning	Trigger service for citizens to measure at certain times/frequencies (e.g., measuring at a satellite overpass or triggering a measurement for a certain monitoring request). Environmental data can be used to change the frequency or moment of monitoring by citizens, for example when an AI detects that there will be no satellite coverage due to cloud presence and alerts citizens to provide more observations in that particular time and location. AI models that benefit from information theory and statistics to help to prioritise effort in field work.
Motivating citizen scientists to participate	Automated reasoning and machine learning	Applications providing personalised reward models for making tools appealing to users. AI that optimises reward models to reflect the personality of the individual. Applications introducing context, information requirements, and gamification aspects.
Providing personalised notifications to increase engagement	Automated reasoning and machine learning	Notifications about collecting or analysing data, which are provided when and where appropriate and with personalised frequency.
Applied use and impact: Improving insights		
Improving data quality control	Automated reasoning and machine learning	Applications that provide means to quality control data using cross checks between citizen science and other in-situ methods to address issues in the data that cannot be addressed by internal quality control (e.g., combining citizen data with satellite data).
Validating outputs through automatic procedures	Automated reasoning and machine learning	Machine-learning algorithms trained to filter out irrelevant data.

removed. However, imagine being in a forest and encountering a rare type of mushroom, wondering whether it would be advisable to pick it up and add it to your dinner plans or if this might lead to serious food poisoning. A tool for nature classification would come in handy. You could then point your phone to the mushroom, snap, and it would instantly tell you everything there is to know about it, including whether cooking it is a wise choice. Some organisations are working on exactly this (Bonnet et al. 2016), training their AI algorithms on the huge amounts of past data and observations collected by scientists and citizens worldwide. AI tools that can instantly classify species could be valuable in other ways. For instance, plant-recognition software and other similar tools appear to be awakening botanical interest among much of the general population, sparking their curiosity about the natural world. Furthermore, computer algorithms trained on classifying dried plants could help researchers to process herb samples, a process that often requires hours of human work (Carranza-Rojas et al. 2017).

Deep learning can be combined with massive-scale citizen science to improve large-scale image classification. Sullivan et al. (2018) showed how citizens and AI excel at different types of classifications and that citizen output can be used to augment and improve deep-learning models. These authors speculated that the integration of scientific tasks into established computer games will be a commonly used approach in the future to harness the brain processing power of humans. They concluded that intricate designs of citizen science games that feed directly into machine-learning models through techniques such as reinforcement learning have the power to rapidly leverage the output of large-scale science efforts. Other examples of data annotated by citizens that have the potential to inform AI in the future are projects administered on websites such as Zooniverse (https://www.zooniverse.org/) and DigiVol (http://digivol.org), and citizens transcribing and annotating museum collection information (Ellwood et al. 2015). Apart from extending current use, new applications of AI in citizen science are likely to appear in the near future as summarised above (**Table 2**). We believe that a wide array of AI applications have the potential to provide new opportunities and positive impact.

Risks exploration
The exploration of risks related to the use of AI in citizen science is driven, at least in part, by the recognition of an existential risk from artificial general intelligence (AGI) (Müller 2016; Yampolskiy and Fox 2013; Ramamoorthy and Yampolskiy 2018), which is the hypothesis that substantial progress in AGI could someday, among other impacts, result in human extinction or some other unrecoverable global catastrophe. Even if this risk is small and the use of AI in citizen science is limited, the potential significant negative consequences for humanity should be reason enough to highlight concerns about the possible impact of AGI (Müller and Bostrom 2016).

In relation to the use of AGI, Dennett highlights the importance of distinguishing between peripheral and central intellectual powers, and of not prematurely ceding authority to AI. "So far, there is a fairly sharp boundary between machines that enhance our 'peripheral' intellectual powers (of perception, algorithmic calculation, and memory) and machines that at least purport to replace our 'central' intellectual powers of comprehension (including imagination), planning, and decision-making" (2017; p. 402). Citizen science's use of AI can contribute to the danger of overestimating AI tools, "prematurely ceding authority to them far beyond their competence."

Ethical concerns commonly associated with robots and other artificially intelligent systems programmed with AI are typically divided into two groups: (1) the moral behaviour of humans as they design, construct, use, and treat artificially intelligent beings, and (2) the moral behaviour of artificial moral agents/machines (AMAs), or machine ethics (The Future of Life Institute 2017; IEEE 2018; Pichai 2018; Shaw 2019; European Group on Ethics in Science and New Technologies 2018; House of Lords Select Committee on Artificial Intelligence 2018; Cowls and Floridi 2018; Winfield et al. 2019; Université de Montreal 2018). In this paper we focus on the first group, given that the presence of AMAs in citizen science is currently very limited.

As the use of AI grows and humans increasingly rely on machines to complete tasks, it is important that the citizen science community gathers data on how AI is used and on the ethical considerations that arise. In contemplating this scenario, we give an overview of AI risks specific to citizen science (and sometimes broader), and are important to consider into the future.

With respect to citizen engagement in citizen science, risks exist that citizens disengage if:

- when contributing expertise to develop and train AI, they are not properly and fairly acknowledged, respected, and rewarded;
- they think that new technologies could be driven more by short-term commercial necessity than longer-term social good;
- they are not comfortable sharing their data because of concerns that their data might be unfairly appropriated (especially for commercial purposes);
- they are forced (because of ethical considerations) to provide too-frequent re-confirmation of their willingness to share their data openly. (See GDPR (2016) as an example of where good intention can sometimes become burdensome.)

Technology giants like Google and Facebook (Webb 2019) are emerging as likely oligopolists in the new world of digital advertising (Mims 2018; Pedemonte 2016), monetising personal data by offering target-oriented advertising services (Krombholz et al. 2012; Teece 2018). Their competitive advantage is largely due to their exclusive access to personal data used to train their algorithms (Mims 2018; Sivinski et al. 2017). Himel and Seamans noted that "*Artificial intelligence ("AI") relies on the use of large datasets to train AI algorithms. Access to such data is therefore a critical resource, the lack of which may create barriers to entry for both AI startups and established firms developing AI technologies*" (2017). It is now recognised

that the existing regulatory frameworks for anti-competitive behaviour have not adequately evaluated the risk nor intervened to prevent data oligopoly, due to lack of recognition of the critical value of data (Pedemonte 2016; Stucke and Grunes 2016). This is a key lesson for citizen science: There is a risk that, as AI-based services arise in the field of citizen science, the same restrictive data policies used by technology giants could be used to create similar oligopolies.

It is possible that citizen science AI startups which lack a long-term funding model will adopt revenue models to monetise their "value-added" services, i.e., algorithmic intellectual property (Brownlow et al. 2015; Hartmann et al. 2014; Schüritz et al. 2017). Where citizens indeed value such services, the market should be left to determine the viability of such revenue models. Citizens engage in citizen science and contribute data for a number of reasons, including public good, curiosity, fun, prestige, and the desire to name their own species (Roger et al. 2019). When citizens contribute data for public good (to mitigate against the risk of creating new oligopolies where they have no choice but to pay for services created from data they contributed), we recommend that an open-data policy is adopted by default. That is, in partnering with technology startups, it should be agreed up front that all data contributed by citizen scientists should be made openly available via Creative Commons licensing. We also recommend exploring whether fragmenting solutions hinders effectiveness in delivering outcomes that users want. It is much easier to contribute expertise in the context of one large well-connected system than through dozens of discrete systems, each with their own quirks.

One of the drawbacks of using some AI approaches, for example deep-learning techniques, is that they are opaque. Specifically, the limitation is the difficulty of explaining, in human terms, the results of large and complex models, such as why a certain decision was reached. The risk is to treat AI as a final authority. For example, validation mechanisms could be established for the automatic verification, by AI, of the accuracy of submissions of data. If this becomes the case, the lack of transparency in reasoning, coupled with our tendency to trust in technology, will inhibit a critical debate in the decisions reached by AI. Among other constraints, regulators will need rules and choice criteria to be clearly explainable to meet transparency and accountability requirements. Some nascent approaches to increasing model transparency, including local-interpretable-model-agnostic explanations (LIME), which attempt to identify which parts of input data a trained model relies on most to make predictions, may help to resolve this explanation challenge in many cases (Henke et al. 2016; Chui et al. 2018). The general recommendation, at least in the short term, is to treat AI as a tool that ideally may be further validated or overturned by human experts. With respect to human relationship with machines, recommendations should be provided about which processes and tasks should be carried out by humans and which ones by machines as well as about how to best manage the replacement or augmentation of humans by machines.

Even if open-source machine-learning toolsets are becoming increasingly available for all to use, an issue with current Google, Microsoft, Amazon, Facebook, IBM, Apple, Baidu, Alibaba, and Tencent ethics policies (Google's DeepMind, for instance), is that we hardly know what the ethics panels are all about (Webb 2019); they are not transparent to public observers. Publicly accountable ethics panels should supervise the processes of AI augmenting the way that people think or taking over certain cognitive tasks. Also, in a data economy where AI algorithms often tend to use personal data as training sets, the ability of AI algorithms to spot patterns makes them very effective at re-identifying personal data in "anonymised" data sets, causing significant concerns about individual and group privacy. The risks related to AI industry are not limited to ethics; a separate risk exists of the AI industry dictating the general direction of citizen science.

Finally, there is an emerging issue of gender and racial bias in AI. Leavy (2018) highlighted the over-representation of white men in the design of technologies. Also, machines largely reflect values of their creators, which can be deeply embedded in machine algorithms. For example, facial recognition software works best for those who are white and male (Buolamwini and Gebru 2018). These gender and racial biases can be reflected in naming, ordering, and descriptions. The risk is that technologies developed for use by citizen scientists (applications and platforms, for example) may alienate users if not tailored to their needs. In addition there is the risk of embedding western views of science and taxonomy into AI, which may preclude ways of grouping organisms according to indigenous knowledge frameworks or alternative cultures. Citizen science presents a special opportunity to engage a wider cohort in training algorithms, which would help in not extending to algorithms the existing biases that are entrenching gender and racial discrimination in modern society.

Discussion and Recommendations

Writing about opportunities and risks of AI in citizen science is difficult. Citizen science is not settled science, despite the growing body of research. AI is not settled science either; it inherently belongs to the frontier, not to the textbook, therefore referencing AI literature, in particular in relation to the human social context, has clear limits. In this paper, we did not write about the AI field in general, but confined ourselves to the field of its application to citizen science, where we can knowingly or unknowingly encounter AI. At times the very terminology can be alienating, and terms such as "AI" should be carefully chosen and well defined. The expression "machine learning" can often be a useful alternative. For example, machine learning applied to computer vision, which is the most common AI technology in citizen science:

· is used by biodiversity projects to verify the accuracy/consistency of contributors' submissions (coming, for example, from iNaturalist, which has created one of the world's largest network of citizen scientists, who have collected over 25 million records of

rare and common species around the world);
· supports citizen science monitoring projects in early warning of harmful algal blooms; and
· identifies the taxon of a species in a photo so that it can be monitored more easily.

Even in the reduced domain of citizen science, rapid advances in AI and the development of improved sensing systems offer the chance to introduce something dramatically new. Many people now engage in citizen science apps on their smartphones daily. As the list of applications grows, so too does awareness of AI in our lives. As a result, technologists pushing for the next big thing in automation now face more questions about what the public really wants. The small group of companies that are investing billions of dollars in using machine learning find themselves having to address the question of how to deal with the public's perception of AI (Hecht 2018).

A big part of citizen science is about connecting people to science, nature, and discovery, and about empowering human minds, mainly through education. Many established citizen science programmes see AI as having a role in this, and some of the biggest names in technology are now entering the citizen science sector through these programmes. Advocates of AI say that technology can make people's lives easier by filtering out hard/repetitive/mundane tasks, so that volunteer efforts can focus on more engaging tasks.

Let us consider projects where AI can streamline the identification of user observations, thus increasing the total number of records being identified. On the one hand we can see risks associated with the unnecessary use of AI. While AI may provide identification help for projects where citizen scientists contribute data and may increase validated users' recordings, there are other ways, apart from AI, to increase the expertise of a citizen science system. These include increasing expertise amongst users, improving the connectivity between experts, and providing more incentive for experts to participate. Moreover, it is not clear that increasing the validated users' recordings through AI helps in progressing citizen science or connecting more people to nature. Connecting and incentivising more human expertise, instead, is likely to progress citizen science and connect more people to nature. According to this vision, AI doesn't necessarily increase users' overall experience (e.g., their general interest, knowledge, or capability to recognise the same organism next time).

On the other hand, we can see opportunities associated with the ability to tackle global-scale challenges. There is little prospect of experts and new citizen scientists by themselves delivering the volumes of data that we need to monitor and understand earth systems, including biodiversity. We need this information for conservation, food security, and many other aspects, for example those related to the Sustainable Development Goals. We should be evaluating the risks associated with the introduction of AI, but we also should consider the risk of ignoring the tools we have to deliver much more data, in a much more usable form, much more quickly.

Since the turn of the millennium, a brute-force approach has been applied to the technology of machine learning, in which huge volumes of data are analysed to look for patterns (Mayer-Schönberger and Cukier 2013). Thanks to increasing citizen engagement and technological improvement, larger repositories of citizen-collected data are now available. As highlighted earlier, larger data repositories available for training AI are a potential risk. To address this, we recommend following the below practices whenever using people's data for AI training:

· An ethics framework about AI use should be created and applied (e.g., The Future of Life Institute 2017; Wehn et al. 2019; Williams et al. 2019).
· A data stewardship plan (e.g., Wilkinson et al. 2016) should inform citizens about plans for and expected outcomes of using data for AI training.
· Good anonymity practices should be adopted. It is important to evaluate to what extent the patterns of information captured may reveal personal information even if names or personal details are not retained. For example, all of the observations in certain areas may derive from a single individual. If any information about their movements is incorporated into the AI training, there is a risk (albeit very small) of revealing personal information about that individual. Anonymity management should be part of the documentation/information provided beforehand to citizens.
· Citizens should be given a standard opt-in/opt-out option (opt-in being best practice).
· Designers should be diverse in ethnicity, gender, and disciplines. This addresses issues such as "data bias".
· Measures of success should be clear. Saying that AI is "successful" in engaging citizens is not enough. Measurements should exist to determine whether citizen science is helping people to engage with nature.
· It should be possible to delete one's data from an AI system (untrain the system).
· It should be possible to challenge the AI. For example, if the number one expert in nudibranchs finds that an AI incorrectly identifies the image of a nudibranch on their phone, who do they call? Who do they talk to? Is there a phone number? A feedback link? How is that handled?

Conclusion

Most people today are only somewhat aware of the rise of AI and its potential impact on their lives. In this paper we discuss this impact in relation to the use of AI in citizen science. It is true that, for all their potential, AI technologies still have many limitations. Current AI limitations include not just issues related to data requirements, but also: (1) regulatory obstacles; (2) lack of social and user acceptance; (3) the challenge of labelling training data (which often must be done manually by citizens and is necessary for supervised learning); (4) the difficulty of obtaining data sets that are sufficiently large and comprehensive to be used for training; (5) the difficulty of

explaining in human terms the results from large and complex models (Why was a certain decision reached?); (6) the generalisability of learning (AI models continue to have difficulties in carrying their experiences from one set of circumstances to another); and (7) the risk of bias in data and algorithms (Chui et al. 2018). Societal concern and regulation, for example about safety, privacy, and use of personal data, can constrain AI use in the public and social sectors if these issues are not properly addressed.

At the same time, the scale of the potential economic and societal impact of AI creates an incentive for all the participants (AI innovators, AI-using organisations, citizens, scientists, and policy-makers) to ensure an AI environment that is friendly and can effectively and safely achieve economic and societal benefits. The potential value that could be harnessed provides the incentive for technology developers, companies, policy makers, and users to try to tackle current AI issues (Chui et al. 2018).

At present, the impact of AI on citizen science is limited, but it is indubitable that technological developments will gather momentum in the next few decades. We anticipate that the result will be all the applications of AI described in this paper and many more. If citizen science is to continue to make meaningful contributions to society and science in the near future, it will not only need to make sense of AI, it also will need to incorporate AI in a meaningful and considered way in future projects.

There is no question that AI potentially introduces significant risks for society and democracy, and ethical considerations regarding how we might retain some control in "central" intellectual powers should be carefully considered by policymakers and legislators.

However, at the same time, we are facing tremendous global-scale challenges across areas of human and planetary health. This means we have a moral obligation to make benign use of AI and every other appropriate and sustainable technology at our disposal to accelerate collection of the data needed to understand our environment, and to use this greater understanding to push for evidence-based decision making to put appropriate mitigation and safeguards in place. Therefore, the authors urge the citizen science community to implement AI, but in a careful way (i.e., only to enhance our "peripheral" intellectual powers). If carefully used, AI is an important tool for accelerating citizen science to ultimately massively scale scientific research.

Acknowledgements

Competing Interests
The authors have no competing interests to declare.

Authors Contributions
All authors made substantial contributions to the conception of the work, contributed to drafting the work or revising it critically for important intellectual content, provided final approval of the version to be published, agreed to be accountable for all aspects of the work in ensuring that questions related to the accuracy or integrity of any part of the work are appropriately investigated and resolved, agreed to be named on the author list, and approved of the full author list.

References

Australian Citizen Science Association. n.d. Who are we? What is citizen science? Available at https://citizenscience.org.au/who-we-are/ [Last accessed 28 September 2019].

Barrat, J. 2013. Our final invention: Artificial intelligence and the end of the human era. St. Martin's Press.

Bonnet, P, Joly, A, Goëau, H, Champ, J, Vignau, C, Molino JF, Barthélémy, D and **Boujemaa, N.** 2016. Plant identification: man vs. machine. *Multimedia Tools and Applications*, 75(3): 1647. DOI: https://doi.org/10.1007/s11042-015-2607-4

Bowser, A, Wiggins, A, Shanley, L, Preece, J and **Henderson, S.** 2014. Sharing data while protecting privacy in citizen science. *Interactions*, 21(1): 70. DOI: https://doi.org/10.1145/2540032

Brownlow, J, Zaki, M, Neely, A and **Urmetzer, F.** 2015. Data and analytics-data-driven business models: A blueprint for innovation. University of Cambridge: Cambridge Service Alliance. Available at https://cambridgeservicealliance.eng.cam.ac.uk/resources/Downloads/Monthly%20Papers/2015FebruaryPaperTheDDBMInnovationBlueprint.pdf [Last accessed 28 September 2019].

Buolamwini, J and **Gebru, T.** 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency: Proceedings of Machine Learning Research*, 81: 77. Available at http://proceedings.mlr.press/v81/buolamwini18a.html?mod=article_inline [Last accessed 28 September 2019].

Campbell, J and **Jensen, DE.** 2019. The Promise and Peril of a Digital Ecosystem for the Planet: Key Decisions Are Needed in the Next 12 Months to Set in Motion a Robust Architecture and Governance Framework. Available at https://medium.com/@davidedjensen_99356/building-a-digital-ecosystem-for-the-planet-557c41225dc2 [Last accessed 29 September 2019].

Carranza-Rojas, J, Goëau, H, Bonnet, P, Mata-Montero, E and **Joly, A.** 2017. Going deeper in the automated identification of Herbarium specimens. *BMC Evolutionary Biology*, 17(1): 181. DOI: https://doi.org/10.1186/s12862-017-1014-z

Ceccaroni, L, Bowser, A and **Brenton, P.** 2017. Civic education and citizen science: Definitions, categories, knowledge representation. In *Analyzing the Role of Citizen Science in Modern Research*,1–23. Hershey, PA, USA: IGI Global. DOI: https://doi.org/10.4018/978-1-5225-0962-2

Ceccaroni, L, Velickovski, F, Blaas, M, Wernand, MR, Blauw, A and **Subirats, L**. 2018. Artificial Intelligence and Earth Observation to Explore Water Quality in the Wadden Sea. In: Mathieu, PP and Aubrecht, C (eds.), *Earth Observation Open Science and Innovation. ISSI Scientific Report Series*, 15, 311–320. Cham, Switzerland: Springer. DOI: https://doi.org/10.1007/978-3-319-65633-5_18

Chui, M, Manyika, J, Miremadi, M, Henke, N, Chung, R, Nel, P and **Malhotra, S.** 2018. Notes from the AI frontier: Insights from hundreds of use cases: Discussion Paper. McKinsey Global Institute. Available at https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-applications-and-value-of-deep-learning [Last accessed 28 September 2019].

Cowls, J and **Floridi, L.** 2018. Prolegomena to a White Paper on an Ethical Framework for a Good AI Society. Available at SSRN: https://ssrn.com/abstract=3198732. DOI: https://doi.org/10.2139/ssrn.3198732

Dawson, D, Schleiger, E, Horton, J, McLaughlin, J, Robinson, C, Quezada, G, Scowcroft, J and **Hajkowicz, S.** 2019. Artificial Intelligence: Australia's Ethics Framework. Australia: Data61 CSIRO. Available at https://consult.industry.gov.au/strategic-policy/artificial-intelligence-ethics-framework/ [Last accessed 29 September 2019].

Deng, DP, Chuang, TR, Shao, KT, Mai, GS, Lin, TE, Lemmens, R, Hsu, CH, Lin, HH and **Kraak, MJ.** 2012. Using social media for collaborative species identification and occurrence: issues, methods, and tools. In *Proceedings of the 1st ACM SIGSPATIAL International Workshop on Crowdsourced and Volunteered Geographic Information*, 22–29. Association for Computing Machinery. DOI: https://doi.org/10.1145/2442952.2442957

Dennett, DC. 2017. From bacteria to Bach and back: The evolution of minds. Great Britain.

Ellwood, ER, Dunckel, BA, Flemons, P, Guralnick, R, Nelson, G, Newman, G, Newman, S, Paul, D, Riccardi, G, Rios, N and **Seltmann, KC.** 2015. Accelerating the digitization of biodiversity research specimens through online public participation. *BioScience*, 65(4): 383–96. DOI: https://doi.org/10.1093/biosci/biv005

European Group on Ethics in Science and New Technologies. 2018. Statement on Artificial Intelligence, Robotics and 'Autonomous' Systems. *European Commission, European Union*. DOI: http://doi.org/10.2777/531856

Feffer, M, Rudovic, O and **Picard, RW.** 2018. A Mixture of Personalized Experts for Human Affect Estimation. In: *Machine Learning and Data Mining in Pattern Recognition*, Perner, P. (ed.). *MLDM 2018. Lecture Notes in Computer Science*, 10935. Cham: Springer. DOI: https://doi.org/10.1007/978-3-319-96133-0_24

GDPR. 2016. General Data Protection Regulation. *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC.*

Grant, MJ and **Booth, A.** 2009. A typology of reviews: an analysis of 14 review types and associated methodologies. *Health Information & Libraries Journal*, 26(2): 91–108. DOI: https://doi.org/10.1111/j.1471-1842.2009.00848.x

Hartmann, PM, Zaki, M, Feldmann, N and **Neely, A.** 2014. *Big data for big business? A taxonomy of data-driven business models used by start-up firms*. Cambridge: Cambridge Service Alliance, University of Cambridge. Available at https://cambridgeservicealliance.eng.cam.ac.uk/resources/Downloads/Monthly%20Papers/2014_March_DataDrivenBusinessModels.pdf [Last accessed 29 September 2019].

Hecht, J. 2018. Managing expectations of artificial intelligence. *Nature*, 563: S141–S143. DOI: https://doi.org/10.1038/d41586-018-07504-9

Henke, N, Bughin, J, Chui, M, Manyika, J, Saleh, T, Wiseman, B and **Sethupathy, G.** 2016. The age of analytics: Competing in a data-driven world. *McKinsey Global Institute*. Available at https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/the-age-of-analytics-competing-in-a-data-driven-world [Last accessed 29 September 2019]

Herzig, A, Lorini, E and **Pearce, D.** 2017. Social Intelligence. *AI & Soc.* London: Springer. DOI: https://doi.org/10.1007/s00146-017-0782-8

Himel, S and **Seamans, R.** 2017. Artificial Intelligence, Incentives to Innovate, and Competition Policy. *Competition Policy International Antitrust Chronicle*. Available at https://www.competitionpolicyinternational.com/artificial-intelligence-incentives-to-innovate-and-competition-policy [Last accessed 29 September 2019].

House of Lords Select Committee on Artificial Intelligence. 2018. AI in the UK: ready, willing and able? *Authority of the House of Lords*. Available at https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf [Last accessed 29 September 2019].

IEEE. 2018. Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems. *IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems*. Version 2 – For Public Discussion. Available at https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/other/ead_v2.pdf [Last accessed 29 September 2019].

Jaques, N, Kim, YL and **Picard, R.** 2016. Personality, Attitudes, and Bonding in Conversations. In: Traum, D, Swartout, W, Khooshabeh, P, Kopp, S, Scherer, S and Leuski, A (eds.), *International Conference on Intelligent Virtual Agents*, 378–382. Cham: Springer. DOI: https://doi.org/10.1007/978-3-319-47665-0_37

Joppa, LN. 2017. The case for technology investments in the environment. *Nature*, 552(7685): 325–328. DOI: https://doi.org/10.1038/d41586-017-08675-7

Joshi, S, Randall, N, Chiplunkar, S, Wattimena, T and **Stavrianakis, K.** 2018. 'We'-A Robotic System to Extend Social Impact of Community Garden. In *Proceeding HRI '18 Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction,*

349–350. Association for Computing Machinery. DOI: https://doi.org/10.1145/3173386.3177817

Kaplan, A and **Haenlein, M.** 2019. Siri, Siri, in my hand: Who's the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence. *Business Horizons*, 62(1): 15–25. DOI: https://doi.org/10.1016/j.bushor.2018.08.004

Kido, T and **Swan, M.** 2015. Ambient Intelligence and Crowdsourced Genetics for Understanding Loss Aversion in Decision Making. *2015 Association for the Advancement of Artificial Intelligence (AAAI) Spring Symposium Series. Ambient Intelligence for Health and Cognitive Enhancement.* Available at https://www.aaai.org/ocs/index.php/SSS/SSS15/paper/viewPaper/10338 [Last accessed 29 September 2019].

Korot, E, Wood, E, Weiner, A, Sim, DA and **Trese, M.** 2019. A renaissance of teleophthalmology through artificial intelligence. *Eye*, 33: 861–863. DOI: https://doi.org/10.1038/s41433-018-0324-8

Krombholz, K, Merkl, D and **Weippl, E.** 2012. Fake identities in social media: A case study on the sustainability of the Facebook business model. *Journal of Service Science Research*, 4(2): 175–212. DOI: https://doi.org/10.1007/s12927-012-0008-z

Kumar, N, Belhumeur, PN, Biswas, A, Jacobs, DW, Kress, WJ, Lopez, IC and **Soares, JV.** 2012. Leafsnap: A computer vision system for automatic plant species identification. In: Fitzgibbon, A, Lazebnik, S, Perona, P, Sato, Y and Schmid, C (eds.), *European Conference on Computer Vision*, 502–516. Berlin, Heidelberg: Springer. DOI: https://doi.org/10.1007/978-3-642-33709-3_36

Lahoz-Monfort, J, Chadès, I, Davies, A, Fegraus, E, Game, E, Guillera-Arroita, G, Harcourt, R, Indraswari, K, McGowan, J, Oliver, JL, Refisch, J, Rhodes, J, Roe, P, Rogers, A, Ward, A, Watson, D, Watson, J, Wintle, B and **Joppa, L.** 2019. A Call for International Leadership and Coordination to Realize the Potential of Conservation Technology. *BioScience*, 69(10): 823–832. DOI: https://doi.org/10.1093/biosci/biz090

Le, QV, Monga, R, Devin, M, Chen, K, Corrado, GS, Dean, J and **Ng, AY.** 2013. Building high-level features using large scale unsupervised learning. In: *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 8595–8598. Vancouver, BC, Canada: IEEE. DOI: https://doi.org/10.1109/ICASSP.2013.6639343

Leavy, S. 2018. Gender bias in Artificial Intelligence: The Need for Diversity and Gender Theory in Machine Learning. *Proceedings of the 1st International Workshop on Gender Equality in Software Engineering*, 14–16. DOI: https://doi.org/10.1145/3195570.3195580

Mac Aodha, O, Gibb, R, Barlow, KE, Browning, E, Firman, M, Freeman, R, Harder, B, Kinsey, L, Mead, GR, Newson, SE, Pandourski, I, Parsons, S, Russ J, Szodoray-Paradi, A, Szodoray-Paradi, F, Tilova, E, Girolami, M, Brostow, G and **Jones, KE.** 2018. Bat Detective—Deep Learning Tools for Bat Acoustic Signal Detection. *PLOS Computational Biology*, 14(3): e1005995. DOI: https://doi.org/10.1371/journal.pcbi.1005995

MacDonald, EA, Case, NA, Clayton, JH, Hall, MK, Heavner, M, Lalone, N, Patel, KG and **Tapia, A.** 2015. Aurorasaurus: A citizen science platform for viewing and reporting the aurora. *Space Weather*, 13(9): 548–559. DOI: https://doi.org/10.1002/2015SW001214

Mayer-Schönberger, V and **Cukier, K.** 2013. Big data: A revolution that will transform how we live, work, and think. Houghton Mifflin Harcourt.

McKinstry, C. 2009. Mind as Space. In: Epstein, R, Roberts, G and Beber, G (eds.), *Parsing the Turing Test: Philosophical and Methodological Issues in the Quest for the Thinking Computer*, 283–299. Dordrecht: Springer. DOI: https://doi.org/10.1007/978-1-4020-6710-5

Miller, T. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267: 1–38. DOI: https://doi.org/10.1016/j.artint.2018.07.007

Mims, C. 2018. Tech's Titans Tiptoe Toward Monopoly: Amazon, Facebook and Google may be repeating the history of steel, utility, rail and telegraph empires past—while Apple appears vulnerable. *Wall Street Journal.* Available at https://www.wsj.com/articles/techs-titans-tiptoe-toward-monopoly-1527783845 [Last accessed 29 September 2019].

Müller, VC (ed.). 2016. Risks of artificial intelligence. CRC Press. DOI: https://doi.org/10.1201/b19187

Müller, VC and **Bostrom, N.** 2016. Future progress in artificial intelligence: A survey of expert opinion. In: *Fundamental issues of artificial intelligence*, 555–572. Cham: Springer. DOI: https://doi.org/10.1007/978-3-319-26485-1_33

Parham, J, Stewart, C, Crall, J, Rubenstein, D, Holmberg, J and **Berger-Wolf, T.** 2018. An Animal Detection Pipeline for Identification. In: *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 1075–1083. IEEE. DOI: https://doi.org/10.1109/WACV.2018.00123

Pedemonte, E. 2016. Google, Facebook, the New Monopolies and Silicon Valley Ideologues. *DigitCult − Scientific Journal on Digital Cultures*, 1(2): 27–34. DOI: https://doi.org/10.4399/97888548960933

Picard, RW. 1995. Affective Computing. *M.I.T Media Laboratory Perceptual Computing Section Technical Report No. 321*.

Pichai, S. 2018. AI at Google: our principles, 7 June 2018. The Keyword. Available at https://www.blog.google/technology/ai/ai-principles/ [Last accessed 29 September 2019].

Poole, D, Mackworth, A and **Goebel, R.** 1998. *Computational intelligence: A logical approach.* New York: Oxford University Press.

Ramamoorthy, A and **Yampolskiy, R.** 2018. Beyond mad? the race for artificial general intelligence. *ITU Journal: ICT Discoveries*, Special Issue No. 1, 2 Feb. 2018.

Roger, E, Tegart, P, Dowsett, R, Kinsela, MA, Harley, MD and **Ortac, G.** 2019. Maximising the potential of citizen science in New South Wales. *Australian Zoologist*. DOI: https://doi.org/10.7882/AZ.2019.023

Rogers, MA and **Aikawa, E.** 2019. Cardiovascular calcification: Artificial intelligence and big data accelerate mechanistic discovery. *Nature Reviews Car-*

diology, 16(5): 261–274. DOI: https://doi.org/10.1038/s41569-018-0123-8

Russakovsky, O, Deng, J, Su, H, Krause, J, Satheesh, S, Ma, S, Huang, Z, Karpathy, A, Khosla, A, Bernstein, M, Berg, AC and **Fei-Fei, L.** 2015. ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, *115*(3): 211–252. DOI: https://doi.org/10.1007/s11263-015-0816-y

Russell, KN, Do, MT, Huff, JC and **Platnick, NI.** 2007. Introducing SPIDA-web: Wavelets, neural networks and Internet accessibility in an image-based automated identification system. In: *Automated Taxon Identification in Systematics: Theory, Approaches and Applications*, 131–149. DOI: https://doi.org/10.1201/9781420008074

Russell, SJ and **Norvig, P.** 2016. *Artificial intelligence: A modern approach*. USA: Pearson Education Limited.

Scherer, MU. 2016. Regulating artificial intelligence systems: Risks, challenges, competencies, and strategies. *Harvard Journal of Law Technology*, 29(2): 353–400. Available at https://heinonline.org/HOL/P?h=hein.journals/hjlt29&i=365 [Last accessed 29 September 2019].

Schüritz, R, Seebacher, S and **Dorner, R.** 2017. Capturing value from data: Revenue models for data-driven services. In: *Proceedings of the 50th Hawaii International Conference on System Sciences*. DOI: https://doi.org/10.24251/HICSS.2017.648

See, L, Carlson, T, Haklay, M, Oliver, JL, Fraisl, D, Mondardini, R, Brocklehurst, M, Shanley, L, Schade, S, Wehn, U, Abrate, T, Anstee, J, Arnold, S, Billot, M, Campbell, J, Espey, J, Gold, M, Hager, G, He, S, Hepburn, L, Hsu, A, Long, D, Maso, J, McCallum, I, Muniafu, M, Moorthy, I, Obersteiner, M, Parker, A, Weisspflug, M and **West, S.** 2019. Citizen Science and the United Nations Sustainable Development Goals. *Nature Sustainability*. (in press).

Shaw, G. 2019. The Future Computed: Artificial Intelligence and its Role in Society. Redmond, WA: Microsoft Corporation. Available at https://news.microsoft.com/futurecomputed/ [Last accessed 29 September 2019].

Shoham, Y, Perrault, R, Brynjolfsson, E, Clark, J, Manyika, J, Niebles, JC, Lyons, T, Etchemendy, J, Grosz, B and **Bauer, Z.** 2018. *The AI Index 2018 Annual Report*. Stanford, CA, US: AI Index Steering Committee, Human-Centered AI Initiative, Stanford University. Available at http://cdn.aiindex.org/2018/AI%20Index%202018%20Annual%20Report.pdf [Last accessed 29 September 2019].

Singh, P, Lin, T, Mueller, ET, Lim, G, Perkins, T and **Li Zhu, W**. 2002. Open Mind Common Sense: Knowledge Acquisition from the General Public. In: Meersman, R and Tari Z (eds.), *On the Move to Meaningful Internet Systems 2002: CoopIS, DOA, and ODBASE. Lecture Notes in Computer Science*, 2519. Berlin, Heidelberg: Springer. DOI: https://doi.org/10.1007/3-540-36124-3_77

Sivinski, G, Okuliar, A and **Kjolbye, L.** 2017. Is big data a big deal? A competition law approach to big data. *European Competition Journal*, 13(2–3): 199–227.

DOI: https://doi.org/10.1080/17441056.2017.1362866

Sterne, J. 2017. *Artificial intelligence for marketing: Practical applications*. Hoboken, New Jersey, US: John Wiley & Sons, Inc. DOI: https://doi.org/10.1002/9781119406341

Storksdieck, M, Shirk, JL, Cappadonna, JL, Domroese, M, Göbel, C, Haklay, M, Miller-Rushing, AJ, Roetman, P, Sbrocchi, C and **Vohland, K.** 2016. Associations for citizen science: Regional knowledge, global collaboration. *Citizen Science: Theory and Practice*, 1(2). DOI: https://doi.org/10.5334/cstp.55

Stucke, ME and **Grunes, AP.** 2016. Introduction: Big Data and Competition Policy. *Big Data and Competition Policy*. Oxford University Press. Available at https://ssrn.com/abstract=2849074 [Last accessed 29 September 2019].

Sullivan, BL, Aycrigg, JL, Barry, JH, Bonney, RE, Bruns, N, Cooper, CB, Damoulas, T, Dhondt, AA, Dieterich, T, Farnsworth, A, Fink, D, Fitzpatrick, JW, Fredericks, T, Gerbracht, J, Gomes, C, Hochachka, WM, Iliff, MJ, Lagoze, C, La Sorte, FA, Merrifield, M, Morris, W, Phillips, TB, Reynolds, M, Rodewald, AD, Rosenberg, KV, Trautmann, NM, Wiggins, A, Winkler, DW, Wong, WK, Wood, CL, Yu, J and **Kelling, S.** 2014. The eBird Enterprise: An Integrated Approach to Development and Application of Citizen Science. *Biological Conservation*, 169: 31–40. DOI: https://doi.org/10.1016/j.biocon.2013.11.003

Sullivan, DP, Winsnes, CF, Åkesson, I, Hjelmare, M, Wiking, M, Schutten, R, Campbell, L, Leifsson, H, Rhodes, S, Nordgren, A, Smith, K, Revaz, B, Finnbogason, B, Szantner, A and **Lundberg, E.** 2018. Deep Learning Is Combined with Massive-Scale Citizen Science to Improve Large-Scale Image Classification. *Nature Biotechnology*, 36: 820–828. DOI: https://doi.org/10.1038/nbt.4225

Sun, Y, Liu, Y, Wang, G and **Zhang, H.** 2017. Deep learning for plant identification in natural environment. *Computational Intelligence and Neuroscience*. 2017. DOI: https://doi.org/10.1155/2017/7361042

Teece, DJ. 2018. Business models and dynamic capabilities. *Long Range Planning*, 51(1): 40–49. DOI: https://doi.org/10.1016/j.lrp.2017.06.007

The Future of Life Institute. 2017. *Asilomar AI principles [Principles developed in conjunction with the 2017 Asilomar conference]*. Available at https://futureoflife.org/ai-principles/ [Last accessed 28 September 2019].

Tollefson, J. 2016. Computers on the reef: Software tools that digitize and annotate underwater images are transforming marine ecology. *Nature*, 537: 123–124. DOI: https://doi.org/10.1038/537123a

Universite de Montreal. 2018. Montreal Declaration for a Responsible Development of Artificial Intelligence 2018. Available at https://docs.wixstatic.com/ugd/ebc3a3_c5c1c196fc164756afb92466c081d7ae.pdf [Last accessed 1 October 2019].

Van Horn, G, Mac Aodha, O, Song, Y, Cui, Y, Sun, C, Shepard, A, Adam, H, Perona, P and **Belongie, S.** 2018. The iNaturalist Species Classification and Detection Dataset. In: *Proceedings of the IEEE conference on*

computer vision and pattern recognition, 8769–8778. DOI: https://doi.org/10.1109/CVPR.2018.00914

Wäldchen, J, Rzanny, M, Seeland, M and **Mäder, P.** 2018. Automated plant species identification–Trends and future directions. *PLoS Computational Biology*, 14(4): e1005993. DOI: https://doi.org/10.1371/journal.pcbi.1005993

Webb, A. 2019. *The Big Nine: How the Tech Titans and Their Thinking Machines Could Warp Humanity*. Hachette, UK.

Wehn, U, Williams, C and **Ceccaroni, L.** 2019. D6.1: Ethics H – Requirement No. 1. Deliverable report of project H2020 MICS (grant agreement No 824711).

Weinstein, BG. 2018. A computer vision for animal ecology. *Journal of Animal Ecology*, 87(3): 533–545. DOI: https://doi.org/10.1111/1365-2656.12780

Wilkinson, MD, Dumontier, M, Aalbersberg, IJ, Appleton, G, Axton, M, Baak, A, Bouwman, J, et al. 2016. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific data*, 3. DOI: https://doi.org/10.1038/sdata.2016.18

Williams, C, Wehn, U and **Ceccaroni, L.** 2019. D6.2: Ethics POPD – Requirement No. 2. Deliverable report of project H2020 MICS (grant agreement No 824711).

Winfield, AF, Michael, K, Pitt, J and **Evers, V.** 2019. Machine ethics: the design and governance of ethical AI and autonomous systems. *Proceedings of the IEEE*, 107(3): 509–517. DOI: https://doi.org/10.1109/JPROC.2019.2900622

Yampolskiy, R and **Fox, J.** 2013. Safety engineering for artificial general intelligence. *Topoi*, 32(2): 217–226. DOI: https://doi.org/10.1007/s11245-012-9128-9

Yudkowsky, E. 2008. Artificial intelligence as a positive and negative factor in global risk. In: Rees, M, Bostrom, N and Ćirković, M (eds.), *Global Catastrophic Risks*, 308–345. Oxford: Oxford University Press.