




Citizen science in European research infrastructures

Stephen Serjeant^a , James Pearson^b, Hugh Dickinson^c, Johanna Jarvis^d

School of Physical Sciences, The Open University, Walton Hall, Milton Keynes MK7 6AA, UK

Received: 2 February 2024 / Accepted: 29 April 2024
© The Author(s) 2024

Abstract Major European Union-funded research infrastructure and open science projects have traditionally included dissemination work, for mostly one-way communication of the research activities. Here, we present and review our radical re-envisioning of this work, by directly engaging citizen science volunteers into the research. We summarise the citizen science in the Horizon-funded projects ASTERICS (Astronomy ESFRI and Research Infrastructure Clusters) and ESCAPE (European Science Cluster of Astronomy and Particle Physics ESFRI Research Infrastructures), engaging hundreds of thousands of volunteers in providing millions of data mining classifications. Not only does this have enormously more scientific and societal impact than conventional dissemination, but it facilitates the direct research involvement of what is often arguably the most neglected stakeholder group in Horizon projects, the science-inclined public. We conclude with recommendations and opportunities for deploying crowdsourced data mining in the physical sciences, noting that the primary goal is always the fundamental research question; if public engagement is the primary goal to optimise, then other, more targeted approaches may be more effective.

1 Introduction

There has been an explosion in the volume and complexity of scientific data. Many scientific disciplines have begun a new data-rich era of discovery, from genomics to astronomical imaging to social sciences working with the firehose of data from social networks. But the data avalanche is so fast, so large and so complex that it often presents serious challenges for computing. Machine learning/artificial intelligence can alleviate these challenges, but this is usually contingent on having sufficient labelled data sets to act as training sets for machine learning. Unsupervised machine learning technologies are often insufficient to the task.

Humans are still often much better than machine learning at many classification tasks. This has led to a new way of doing science: crowdsourcing the data mining, with the help of citizen science volunteers. Indeed some open science data mining problems are at present intrinsically intractable even for machine learning, particularly where a subjective human assessment is an essential element of the problem at hand. Furthermore, even the most sophisticated machine learning/artificial intelligence (ML/AI) technologies are unable to respond to a classification problem with, in effect, “After seeing the data, I now believe this project is asking the wrong question, because the data appear to be showing us something unexpected”. For example, Serjeant [1] presented a hypothetical example of a facial recognition machine learning algorithm being presented with images of commuters. Such an algorithm would report the positions and possibly the identities of the faces, but if there were, for example, a circus clown working their way through the commuters, the algorithm would fail to identify the unusual occurrence. A human volunteer, however, would be much more likely to spot the outlier in the data. This capacity also makes human volunteers much better at classifying data where the subjects of interest are too rare, or the sample sizes too small, for machine learning. One of the best examples of this is from the Snapshot Serengeti citizen science project, which asked volunteers to identify wildlife from webcam images in a nature reserve [2]. This project identified over 100 thousand wildebeest, for which there are now machine learning algorithms that could replicate the effort, but the project also found 17 examples of the rare and elusive zorilla.

There is no realistic prospect that machine learning will be able to entirely remove the need for human classifications, for multiple reasons. Firstly, machine learning is only as good as its training data, and where classification subjects are rare (as in the zorilla problem above), by definition the machine learning will fail. Secondly, despite the name “artificial intelligence”, there is never any underlying intelligence of any sort. Large language models are now infamous for producing lucid text that is both confident and wrong, which is in turn due to there being no underlying knowledge model. Thirdly, the internal operations of machine learning are

^a e-mail: stephen.serjeant@open.ac.uk (corresponding author)

^b e-mail: jpearson42@outlook.com

^c e-mail: hugh.dickinson@open.ac.uk

^d e-mail: johanna@thejarvishouse.co.uk

often opaque, to the extent that entire computer science research themes are aimed at unpicking these internal operations. Pernicious biases have been discovered in many machine learning classification outputs, which can sometimes have serious consequences (e.g. facial recognition trained on biased training sets), while human volunteers are much more able to step outside the question posed (see the clown example above).

The science-inclined public is also by far the largest, but often most overlooked, set of stakeholders in open science. A central vision of the European Open Science Cloud (EOSC) is to make scientific data FAIR, that is, Findable, Accessible, Interoperable and Reusable. Implicit in this vision is that FAIR data should also be useful, but this is far from being guaranteed, especially given its inter- and multi-disciplinary remit. There are many examples of errors in FAIR data use by researchers outside their direct specialism. The farther from one's subject specialism, the more curated one's interaction with data needs to be. For example, Daylan et al. [3] reanalysed public data from the Fermi gamma-ray telescope taken in the direction of the Galactic centre, where the dark matter particle density is predicted to be the highest in the Milky Way halo. The authors found a gamma-ray excess, which they interpreted as a signature of dark matter particle annihilation, which in turn would represent a major insight into the nature of dark matter particles. However, the instrument team re-interpreted this excess as an observational systematic [4]. Without taking any view on this technical debate, it is clear that the usefulness of FAIR data will always be limited by the supporting contextual information. This is sometimes glibly incorporated into expectations on metadata, but it is clear from this example that this may even extend to user training. The further away that you are working from your subject specialism, the more curated your interaction must be with data. Perhaps the most extreme example of this is the science-inclined public, who need the most support in their interaction with FAIR data.

With a suitably carefully curated interaction with FAIR data, members of the public can make a genuine and valuable participation in scientific discovery. Furthermore, there is a huge public appetite for taking part. Therefore, our approach has been to involve the science-inclined public with the research infrastructures (and their precursors/pathfinders) and open science projects with data mining and data collection citizen science projects. In this paper, we will summarise the efforts taken towards these goals in several large European Union funded projects: ASTERICS (Astronomy ESFRI and Research Infrastructure Clusters), ESCAPE (European Science Cluster of Astronomy and Particle Physics ESFRI Research Infrastructures), and the EOSC-Future project that aimed to develop services for the European Open Science Cloud.

2 Methodology: existing infrastructure

There are a number of citizen science platforms suitable in principle for integration with research infrastructure and/or EOSC services. In this section, we will review a selection of the current marketplace of these platforms, and the use cases for which each platform is often deployed and perhaps best suited. We focus on data mining and data collection tasks where the majority of the end users are the science-inclined public, despite projects being managed by core teams of subject-specialist academics. For the sake of brevity, we do not discuss policy or training portals such as EU-Citizen.Science.

The Zooniverse has a very wide range of data mining capabilities for researchers to tailor for their non-specialist volunteers. Volunteer workflows can include polling among options, free text entry, image annotation of features or regions of interest, and more. Volunteers may be given options of reversing the colour balances of videos, and may pan or scroll through images. Short movies can be included as subjects using animated gifs, and the platform also supports audio clips. This is enough scope to facilitate a wide range of scientific capabilities, including transcription of handwritten texts, identification of objects or features in images, classification of audio and/or visual data, and more. The reach of the Zooniverse platform is enormous with over 2 million volunteers. In the 2019–20 year alone,¹ Zooniverse hosted 65 new projects, with volunteers contributing over 85 million classifications. In that year alone, the research teams published 35 papers. The platform grew from a desire to generalise the successes of the Galaxy Zoo project, and perhaps partly as a result of that heritage from astronomy and astrophysics, the Zooniverse has functionalities that are well-tested in ESCAPE science domains. The platform also supports project site translations of navigation interfaces in some languages (researchers running projects must still supply translations of project-specific text), making the platform suitable for the international contexts of EOSC. Therefore, the Zooniverse platform was ESCAPE's obvious first choice for integration into EOSC.

The Zooniverse is well suited to data mining problems, but a host of other platforms have specialised in the different requirements of data collection. For example, the nQuire site has capabilities for confidentially surveying volunteers themselves, and for uploading audio recordings, images or movies. Volunteer surveys can include free text entry, numerical entry, checkboxes or radio buttons, as well as supporting data from external sensors such as accelerometers, gyroscopes, orientation indicators, ambient light sensors, magnetometers, audio volumes, and location trackers.

Another platform specialising in data collection is CitSci.org, which now hosts 1249 projects, which in turn have contributed 1 825 928 data points. The platform offers volunteers the capability of uploading images and observations on a wide range of devices, plus some options for volunteers to visualise and mine the data.

A more specialised free platform is Treezilla, which aims to map, measure and monitor trees across the UK and the Republic of Ireland, with the help of public volunteers as well as local authorities, businesses and other organisations. Another UK-based

¹ Source: <https://blog.zooniverse.org/2020/11/17/into-the-zooniverse-vol-ii-now-available/>.

specialised biodiversity project enrolling non-specialists is the Big Butterfly Count. Both this and Treezilla are single bespoke solutions for individual projects, in contrast to nQuire or CitSci.org. SPOTTERON is a platform that specialises in environmental monitoring for a wide range of citizen science projects, but the platform is not free at the point of use for researchers.

Efforts have also been made to integrate data collection functionality into the Zooniverse. Rather than duplicating the functionalities of other platforms, Zooniverse partnered with the CitSci.org platform mentioned above. As with many technical developments in citizen science, this was driven by a direct scientific need, in this case from the Mountain Goat Molt Project. This project was crowdsourcing image data of goats at various locations and times, but wished to also crowdsource aspects of the image analysis. With the benefit of funding from the US National Science Foundation, projects on CitSci.org and Zooniverse can now be linked, and data from CitSci.org can be uploaded in near-real time for analysis in Zooniverse.²

There are other ways for non-specialist volunteers to participate in scientific discovery in open science, beyond data mining and data collection. One of the earliest examples of this is the SETI@Home project, in which volunteers downloaded a screensaver that mined the SETI data sets, so volunteers donate otherwise idle CPU time for the science objectives of the SETI institute. As the balance changed between crowdsourced computation and the costs of high performance computing, the focus has shifted to computationally-harder problems that involve user interaction. For example, the Foldit project gamifies the discovery of protein folding configurations. Protein folding problem is an NP-hard problem (non-deterministic polynomial hard), so it is not expected that there can be a general polynomial-time algorithm for finding the minimum energy configurations, although insights may be possible in individual cases. This is one of the few citizen science projects that involve elements of explicit competition between volunteers. At the time of writing, machine learning advances in protein structure prediction such as AlphaFold have not yet removed the need for testing against Foldit structure solutions,³ as well as against experimental data; meanwhile, Foldit has also been extending to other areas such as protein design.

3 Results

3.1 Overview

The EU-funded projects included deliverables to create a number of citizen science project demonstrators. This section describes the projects created during the funded lifetime of the ESCAPE and ASTERICS, and the extended work performed with the benefit of funding from the EOSC-Future project. Going beyond the baseline requirements of the minimum viable product of these demonstrators, we have also used several of them as exemplars in tutorial notebooks for creating and managing citizen science data mining projects, and integrating real-time machine learning into the projects; these tools will be described later in Sect. 4.1.

3.2 Muon Hunter and Muon Hunters 2.0—Return of the rings

The VERITAS team had been particularly keen to launch a major new citizen science experiment in the run-up to the Cherenkov Telescope Array. The original Muon Hunter experiment was therefore based on data from the VERITAS telescopes, used to detect some of the highest-energy photons in the Universe. These gamma-ray photons are generated in extreme astrophysical environments such as the expanding blast waves driven by supernova explosions, or from relativistic jets in active galaxies. Muons are not the target astrophysical signal, but rather are a prominent background contaminant when observing very-high-energy gamma rays on terrestrial facilities. Muons can be distinguished in principle from the gamma-ray signals, because the muons present a distinctive ring-like morphology on the detectors. These can be relatively easy for a human to discern or detect, but incomplete or truncated rings can appear as false-positive gamma-ray signals to automatic analysis algorithms. We therefore sought the help of the public to identify camera images that contain muon rings to generate a truth set for training machine learning algorithms to remove these false-positive signals. This project was a runaway success. This original Muon Hunters project welcomed 6107 citizen scientists who made 2 161 338 classifications of 135 000 objects. In its first five days, it attracted 1.3 million classifications, making it one of the most successful projects to run on the Zooniverse platform.

This machine learning led directly to the follow-up project, Muon Hunters 2.0, which tested the automated classifications, and sought to build a virtuous circle between human and machine classifications. The project also changed the manner in which the data was presented to volunteers, showing images in a grid pattern in order to bring additional efficiency to the classifications. Results from these projects are summarised in Feng et al. [5], Bird et al. [6], Laraia et al. [7], Bird et al. [8], and Flanagan et al. [9].

3.3 Classifying variable stars using SuperWASP data

This project aims to classify the many types of periodic variable stars found in SuperWASP survey data using the Zooniverse citizen science platform.

² Source: <https://blog.citsci.org/2021/11/17/crowdsource-images-with-citsci-analyze-them-with-the-zooniverse/>.

³ See e.g.: <https://fold.it/forum/blog/alphafold-machine-learning-for-protein-structure-prediction>.

The wide angle search for planets, SuperWASP, is a leading ground-based survey data set for transiting hot Jupiters (massive gaseous planets). The project accumulated 16 million images from 2800 nights, comprising 580 billion data points from 31 million unique objects outside the Galactic plane, over the courses of 2004–2013.

A re-analysis of all of the light curves led to the identification of 0.8 million unique objects have been identified as having statistically significant but unexplained variabilities. The volunteer tasks were to classify these variabilities, presented as folded light curves, as resembling that of an EA or EB type eclipsing binary star, an EW type eclipsing binary star, a pulsating star or a generic rotational modulation.

This project has been enormously successful. At the time of writing, it has attracted 14 052 volunteers who have contributed 5 192 629 classifications, covering 1 851 043 classification subjects. More details of the project results can be found in Norton [10] and Thiemann et al. [11].

3.4 CREDO—Involving the public in scientific research

The Cosmic Ray Extremely Distributed Observatory (CREDO) collaboration is an ongoing research project involving scientists and the public from around the world. Unlike most of the projects discussed in this paper, the CREDO project is primarily a data *collection exercise*. The “high risk, high return” goals include the determination of the nature of dark matter. The ASTERICS project partially supported the creation of the CREDO project in its early stages, and since that time the project has flourished.

The primary objective is to search for cascading products of the decay of supermassive particles, such as dark matter. These particle cascades occur elsewhere in the solar system so will be observed as being distributed over a very large international geographical area. CREDO is utilising the millions of small detectors throughout the globe in smartphone cameras, operating in a continuous readout of dark frames. These are supplemented by dark frames from worldwide astronomical observatories, and analysed by the public. The project platform development and science goals are described in more detail elsewhere [12–14].

3.5 Challenge the machines

The science goal of this citizen science experiment (available at <https://www.zooniverse.org/projects/hughdickinson/euclid-challenge-the-machines>) was to see whether human volunteer classifiers can find strong gravitational lens events better than machine learning. The Euclid consortium had recently ran a Strong Gravitational Lensing Challenge (with challenge data deposited in the public domain), in which teams contributed machine learning algorithms to find lensing events in simulated imaging data from single-band Euclid and the multiband ground-based KiDS (Kilo Degree Survey) project. One valiant expert volunteer visually classified the entire $\sim 100\,000$ image set. One of the great surprises of this challenge was that the valiant expert did not win the challenge [15]. Several machine learning algorithms were more effective than the human expert in finding lenses. But this is a feature recognition problem in which a non-expert can be very well trained, so the question is whether the artificial intelligence has exceeded human abilities or whether some humans can still beat the machine. Since we were more likely to engage volunteers with colourful images, we selected the simulated KiDS data for crowdsourced analysis, rather than that of Euclid.

Our tactic for educational resources associated with the mass participation experiments was to embed the material directly into the citizen science workflow. This means that the participant volunteers can be introduced gently to the science context as their involvement in the project extends.

We selected this challenge data for several reasons:

- It was to be the first of several similar challenges, so new versions of the crowdsourcing experiment could be run at later dates.
- There is a precedent for successful crowdsourcing in strong gravitational lensing, in the Space Warps experiment.
- Similar morphological searches for strong gravitational lenses can also be made in other ASTERICS and ASTERICS-related facilities, such as the SKA, the Large Synoptic Survey Telescope, the Hubble Space Telescope, and others.

Besides the particular science objectives of the experiment, we were also aiming to find experiments that are in some way aligned with the primary science goals of the facility, and which are likely to have some longer-term traction or application to other facilities. In the case of Euclid, the enormous hundred-fold increase in known strong gravitational lensed systems that Euclid will bring will provide independent statistical constraints on dark energy parameters.

The drafting of the experiment itself and its associated educational resources was done very quickly and efficiently on the Zooniverse’s Panoptes platform. We helped in the creation of the new mobile app functionality for Zooniverse, using a “Tinder”-like swiping left and right to indicate a binary classification. This citizen science experiment was one of the first to pioneer this mobile citizen science interface.

However, that is not to say that all aspects of this project were trivial to create. The classification data set was examined carefully by ourselves and by our collaborators in the Zooniverse Space Warps project. We discovered that while most of the simulated data sets were representative of the known gravitational lens systems discovered in the COSMOS survey by the Hubble Space Telescope, there was also a significant subset of outlier systems that appeared to be unphysical. Typically, these were systems in which there was an unusually faint foreground lensing galaxy with the background galaxy warped into a wide Einstein ring. Since the ring radius is a reliable proxy for the total mass of the system, this would mean that the simulated galaxies either have an unphysical mass-to-light ratio or have been placed at unphysically high redshifts (i.e. greater distances). We therefore took care to pre-filter the data set to be

classified, in order to better reflect the observed parameter distributions in the known Hubble Space Telescope gravitational lenses. This process took several months. It should be added that none of these anomalies that we discovered had been uncovered by any of the AI algorithms or human inspections of the data set prior to our work. Further detail of the results can be found in [16].

3.6 Galaxy Zoo: Clump Scout

One of the main goals for modern observational cosmology is to discover and understand how galaxies and their constituent substructures have assembled and evolved throughout cosmic history. The diverse observed morphologies of individual galaxies are not only indicative of their current composition, but also encode a detailed record of their assembly histories, their past and ongoing star formation, and their interaction with local environments. Galaxies grow by forming stars. Today, the Hubble Space Telescope can detect distinct star-forming structures inside the galaxies that populated the Universe when it was less than a quarter of its current age. These early galaxies look very different to their modern-day counterparts. Their discs are thick, turbulent and violent environments, where hundreds of new stars are born every year. Many also exhibit giant regions of enhanced star formation that appear as bright clumps in telescope images. In contrast, today's star-forming galaxies are typically much more placid. Their discs are thin and well-ordered and clumpy star formation is much less common.

These profound differences raise obvious questions. Which physical mechanisms drove the observed evolution in star formation activity? Why are giant star-forming clumps so much more common in the early Universe?

To understand why clumpy galaxies became so rare, we need to find and investigate as many examples as possible. One potential approach involves training modern deep learning algorithms that use deep learning to identify galaxies with clumps. However, appropriately labelled training data for clump detection is scarce and laborious to generate. Moreover, automatic algorithms struggle to operate effectively if their limited training datasets underrepresent the diversity of the data being analysed. In contrast, human beings working in collaboration can extrapolate successfully from a handful of examples.

To benefit from this impressive human capability, we used the Zooniverse platform to develop a new citizen science project called Galaxy Zoo: Clump Scout. The project invites the general public to examine images of galaxies obtained by the Sloan Digital Sky Survey (SDSS) and annotate all the clumps they can see. By participating in Galaxy Zoo: Clump Scout, volunteer clicks will identify the locations of clumps within thousands of galaxies in the nearby Universe. The project uses a novel Bayesian aggregation algorithm that dynamically derives a consensus for the clump locations based on the annotations provided by multiple volunteers for the same image. The algorithm also estimates the reliability of the dynamic consensus, which helps to ensure completeness while avoiding spurious clump detections. Galaxy Zoo: Clump Scout represents one of the first large-scale studies of clumps in local galaxies.

In the future, new space telescopes like Euclid will image more than a billion galaxies. Using citizen science to manually check so many galaxies for clumps would take many years, even for the most dedicated Clump Scout volunteers. The speed of computer algorithms will be required to process such large volumes of data and so we have trained deep learning models for clump detection and localisation based on the faster RCNN architecture that exhibit very good performance on a wide range of galaxy images. However, there will always be galaxy images that confuse the computer algorithms and we will need the help of the Clump Scout volunteers to step in when deep learning fails. Even more importantly, human beings inspecting images are much better at spotting any unusual or unexpected phenomena that single-minded algorithms would just ignore. Indeed, the history of citizen science is full of examples when keen-eyed volunteers make serendipitous discoveries. Projects such as Clump Scout will help to maintain this tradition in the future.

The Galaxy Zoo: Clump Scout project has now concluded after collecting 1 738 822 from 13 762 volunteer citizen scientists who collectively annotated the locations of visible giant star-forming clumps in 85 286 low-redshift galaxies. To date, this project has generated two refereed journal papers. Dickinson et al. [17] presented a novel algorithm for aggregation of two-dimensional citizen science data, deployed on the clump annotations in this project. Adams et al. [18] presented the first results of the observed clump fraction in galaxies from this project, revealing a sharp decline in the clumpy fraction from redshift 0.5 to the present day. The aggregated annotations have since been used as truth sets for training machine learning algorithms, which will be the subject of future papers.

3.7 SuperWASP: black hole hunters

SuperWASP: Black Hole Hunters (SWBHH, <https://www.zooniverse.org/projects/hughdickinson/superwasp-black-hole-hunters>) was a new citizen science project that has volunteers search for an extremely rare event that of a hidden black hole orbiting a normal star in a binary system.

To fully understand the lifecycle of stars and how they contribute to the properties of a galaxy, astronomers not only look at regular stars, but also the remnants they become at the end of their lives: black holes and neutron stars. These compact objects form in the violent supernova explosions of massive dying stars. However, most stars have a companion (a binary system), so provided this explosion does not tear apart the system, then the result is the companion star and the compact object orbiting each other. A couple of hundred compact objects have so far been discovered because their orbit has decayed inwards to the point that they can

accrete matter from the companion star. However, simulations predict there should be hundreds of thousands of star-compact object binary systems, so another method is needed to identify these elusive objects.

One promising method is that of gravitational microlensing. When a foreground object such as a black hole or neutron star passes in front of a distant background star, the foreground object can act as a lens for the background star's light, causing the light to bend around it. Provided high enough resolution imaging, we observe this as the background star being distorted into a ring of light (called an Einstein ring). Unfortunately, for small objects like black holes and neutron stars, the angular resolution limits of telescopes prevent us from seeing these distortions directly. However, it is possible to infer these distortions by observing a temporary increase in brightness as the object passes between us and the background star. Such a microlensing event can even occur for the aforementioned star-compact object binary systems, with the compact object passing between us and its companion star. This produces a characteristic periodic increase in brightness that we can use to detect the compact object, and even learn about its properties, but to do this requires us to view the system almost perfectly edge-on.

These “self-lensing” signals are expected to be extremely rare—so far, nobody has seen conclusive evidence of a self-lensing signal from a black hole. Even though simulations predict numerous hidden black holes in our galaxy, only a handful will be orbiting in a plane that means they pass in front of their companion star when viewed from the Earth. To find such rare and infrequent events, we need to use survey data that covers large regions of the sky and spans long periods of time.

SuperWASP was an experiment to monitor the sky in a search for exoplanets. SuperWASP's exploration of the time domain parameter space in astronomy is a direct precursor to the Legacy Survey of Space and Time on the Vera Rubin Observatory, one of the forthcoming ESRIs supported by ESCAPE. The now-decommissioned SuperWASP telescopes are also an excellent instrument to use for self-lensing searches. During its operational lifetime, SuperWASP accumulated light curve data for millions of stars. It observed a very large fraction of the sky over a period of approximately 8 years with as little as 40 min between repeated observations. The primary shortcoming of SuperWASP data is that they are very noisy, which complicates automatic detection of subtle self-lensing signals and necessitates manual inspection of lightcurves.

To enable manual inspection of so many lightcurves, we designed and deployed SWBHH as a citizen science project on the Zooniverse web platform. SWBHH engaged 5 673 volunteers who have collectively completed a spectacular 2.1 million inspections of all its 208 700 lightcurves. The project launch was timed to coincide with the launch of a prime-time BBC science series “Universe”, which was presented by the popular UK-based science communicator, Professor Brian Cox. The SWBHH project was promoted to viewers of “Universe” using an accompanying poster designed by members of the ESCAPE collaboration. The SWBHH project was also promoted by Serjeant and Dickinson as part of an online panel discussion about black holes during British Science Week 2022 (<https://www.superwasp.org/britishscienceweek/>). During the event alone, the audience completed over 7000 classifications and the total for the day was over 14 000 classifications.

3.8 Galaxy zoo: Cosmic Dawn

Galaxy Zoo: Cosmic Dawn was a new citizen science project that formed the latest iteration of the longest running project on the Zooniverse, Galaxy Zoo, which aims to classify images of galaxies based on their visual morphologies.

A core aim of extragalactic astronomy is to study how galaxies form and evolve over cosmic time, including understanding the physical mechanisms that govern their structures and produce the wide range of galaxies we observe. Galaxies have a variety of shapes, from ball-like ellipticals to those with grand spiral arms, and different colours that indicate the presence and composition of dust and stars of different ages. These properties pertain to that galaxy's history, including through its rate of star formation, the activity of its central black hole, and the merging of it with other galaxies over cosmological time. Galaxies also typically live within extensive dark matter haloes that can, in rare circumstances, act as a (gravitational) lens to distort and magnify the light of distant background galaxies that would otherwise be hidden from us. Additionally, low surface brightness (LSB) galaxies contain a higher ratio of dark matter to baryonic matter, making them useful probes for studying the impact of dark matter, but are inherently faint and therefore difficult to detect. Hence, the appearance of a galaxy is the result of a combination of properties, and to study the impacts of these properties requires large samples of classified galaxies. Modern surveys are capable of collecting such large samples for fainter objects using deeper imaging, moving towards higher redshifts in order to examine galaxies from further back in time and their evolution from then to now. Any of those identified as a lens would additionally allow for the study of even more distant galaxies from the earliest epochs of galaxy formation.

The Cosmic Dawn survey aims to understand how galaxies, black holes and dark matter haloes co-evolve from the epoch of reionization (when the Universe was around 500 million years old) to the present. As a 50 square degree multi-wavelength survey of the Euclid deep fields, it covers some of the darkest areas of the sky that have been selected for study by multiple international observatories, partly in preparation for the ultra-deep photometry and spectroscopy to be produced by the upcoming Euclid mission. These deep fields allow for the study of large numbers of galaxies going back to when the first of them formed within the first billion years after the Big Bang.

One of the programmes forming Cosmic Dawn is the Hawaii Two-0 (H20) survey of the Euclid deep calibration fields, which involves deep multiband imaging of the 10 square degree Euclid Deep Field North (EDF-N) using the Hyper Suprime-Cam (HSC) on the 8.2m Subaru Telescope on the summit of Mauna Kea in Hawaii. The ultra-deep HSC imaging contains around a million galaxies per square degree down to their magnitude limit, with data processing still ongoing. Through combining this with Spitzer

Space Telescope infrared imaging and Keck DEIMOS (DEep Imaging Multi-Object Spectrograph) spectroscopy, H20 will advance the study of galaxy evolution and co-evolution with dark matter haloes out to a high redshift of $z = 7$ (<800 million years since the Big Bang), pushing the boundaries of extragalactic astronomy.

Citizen science helps play a crucial role in the examination of such large data sets, with volunteers able to classify huge amounts of data much faster than a small team of researchers. While machine learning methods has developed rapidly over the past few years, volunteers still outmatch them when it comes to classifying complex features or the serendipitous discovery of rare objects. As such, the Galaxy Zoo: Cosmic Dawn project has volunteers classify the hundreds of thousands of galaxies detected in H20's HSC imaging of the EDF-N, exploring a region of intense study with deep multiband imaging. These classified images extend to higher redshift sources and provide a means of statistically studying objects such as galaxies with LSB features while expanding the list of interesting, rarer objects found through serendipitous discovery. In addition, mapping the EDF-N is important for future surveys such as the Euclid mission, with classifications providing a basis for rapid follow-up of the most interesting objects. HSC is also a precursor for the Legacy Survey of Space and Time (LSST) Vera C. Rubin Observatory, with this classification work providing multiband ground truth sets for use in training deep learning models, such as for strong gravitational lens finding.

The project workflow involved presenting colour postage stamp cutout images of each galaxy to volunteers to classify following the set of questions developed by Galaxy Zoo, modified to suit these HSC images. Initially focusing on 2 square degrees in the EDF-N, it can be extended to a wider area as the processed H20 data becomes available. Due to the images featuring sources at higher redshifts than is typical for Galaxy Zoo, a new colour scaling had been implemented. Additionally, the added confusion noise from the increased number of background galaxies in these deeper images proved challenging for the H20 team's source detection and modelling, hence requiring re-sizing of their cutouts to correctly display sources. The Zooniverse tutorial and help text for the project were updated to reflect this and explain the project's new data, and the "star/artefact" option normally presented in the Galaxy Zoo workflow was also expanded in order to aid the H20 team with improving their source detection and modelling methods going forward.

When the project had its public launch, a blog post was announced alongside it informing volunteers of the launch and to keep an eye out for and tag any extremely red sources or gravitational lenses due to their rarity. This also featured in the tutorial and help sections of the project, and, as volunteers can select lensing features as part of the workflow, this additionally allowed the Galaxy Zoo team to investigate how tagging rare sources compares to asking via the workflow questions. The Galaxy Zoo: Cosmic Dawn project was completed in June 2023, and a general data release paper is currently in production. The volunteers' classifications will also serve to generate multiple publications, including the detection of strong gravitational lenses, statistics of clumpy galaxies, and examinations of galaxies with low surface brightness features.

3.9 Radio galaxy zoo: LOFAR

The low-frequency array (LOFAR) is a large interferometric array of radio telescopes located primarily in the Netherlands, but with outlying antennae dispersed across Europe. LOFAR is also a recognised science and technology pathfinder facility for the next-generation radio telescope, the Square Kilometre Array (SKA).

Radio Galaxy Zoo: LOFAR (<https://www.zooniverse.org/projects/chrisrnp/radio-galaxy-zoo-lofar>) was a new citizen science project led by ASTRON in the Netherlands with substantial ESCAPE-funded support provided by the Zooniverse platform and the Open University (OU). The project invites volunteers to classify radio images extracted from the first data release of the LOFAR Two-metre Sky Survey (LoTSS) which covers 424 square degrees in the region of the HETDEX Spring Field. In this release, 325 694 individual radio sources were detected with a signal five times greater than a typical background noise fluctuation.

Classification entails attribution of distinct regions of radio emission to a single origin and (where possible) identifying an optical counterpart for the radio emission's source. By Zooniverse standards this is a very complicated analysis task, which requires consideration of multiple images, representing radio and optical data. Moreover, the degree of scientific comprehension that volunteers require to successfully provide the required classifications is more than typical Zooniverse projects, which often rely on somewhat mechanical "microtasks" that can be performed without complete understanding.

To render such complex classifications tractable for citizen scientists, the OU and Zooniverse teams have developed an advanced volunteer training and feedback system. The project uses a tutorial video paired with a dedicated training workflow that allows volunteers to mimic the classification process as demonstrated by one of the LOFAR project scientists. The training workflow presents subjects in the same order as they appear in the video (unlike the normal random ordering employed by the Zooniverse platform) and volunteers receive real-time feedback in response to the annotations they provide. This is the most advanced training infrastructure that has been deployed using the Zooniverse project builder platform and the upgrades that have been developed by the OU and Zooniverse with ESCAPE support will be available for future CS projects to leverage. It has been shown that volunteers' confidence is a critical factor in citizen science projects, which improves classification accuracy and volunteer retention.

The project is now complete, with 11 583 volunteers having contributed 0.95 million classifications on a total of 189 375 subjects. A refereed journal paper is now in preparation on the catalogue of optical identifications generated in this project.

3.10 Knitting patterns

Our ambition has always been to extend the EOSC citizen science demonstrators to beyond the subject-specialist areas of the ESCAPE EOSC cell. Therefore, with the support of the EOSC-Future project, we have assisted in the development of the Knitting Leaflets citizen science project (<https://www.zooniverse.org/projects/elliereed185/knitting-leaflet-project>). This project aims to discover more about how knitting and knitwear developed in Britain during the twentieth century, by carefully characterising information on the front covers of knitting pattern leaflets and magazines, to yield knowledge of how versions of femininity connected to fashion and consumption, and how they changed over time. Three workflows were developed requiring increasing levels of knowledge in this subject area. The project had its public launch on 27th September 2023, and proved exceedingly popular as it was completed within a week, with 1 327 volunteers making a total of 52 928 classifications of 9 624 subjects.

The Knitting Leaflets project was promoted to the Zooniverse community of over 2.6 million registered users through the Zooniverse front page as well as through the Zooniverse email newsletter (announcements@lists.zooniverse.org). It was also promoted in an article on the Knitting Industry website (<https://www.knittingindustry.com/creative/the-knitting-leaflet-project/>; 400+ listed views) as well as the Knitting History Forum Facebook page (https://www.facebook.com/KnittingHistoryForum/?locale=en_GB; 1 600 followers), and the project also featured in the forums of the Gardeners' World website (<https://forum.gardenersworld.com/discussion/1077606/anybody-interested-in-the-history-of-knitting-patterns>).

3.11 African indigenous knowledge

African Indigenous Knowledge (AIK; <https://www.zooniverse.org/projects/xsr/african-indigenous-knowledge-aik-m>) is the second citizen science project we have helped to develop to extend the EOSC citizen science demonstrators beyond the ESCAPE remit, with the project also using the Zooniverse platform. This project is still under development, but aims to study traditional African indigenous food system knowledge for better food production, processing and consumption in Africa. The project asks volunteers to classify the types of food and tools present in photographs from traditional farmers in Sierra Leone, and is currently paused while data collection is underway.

4 Discussion

4.1 Tools and data

Subject-specialist researchers must often go to substantial lengths to convert experimental data into high-level “subjects” (images, videos etc.) that can be successfully interpreted and analysed by non-expert volunteers. Therefore, in order to streamline this process and reduce the workload of running a successful citizen science project, we have created tools and template workflows that generate attractive subject data, upload them to the Zooniverse servers and manage them effectively as the volunteer classifications accumulate.

Firstly, we created a Jupyter notebook and documentary materials demonstrating web-interface based and programmatic (scriptable) Zooniverse project management including project and workflow creation, subject creation and upload, improving the volunteer experience through extra metadata, retrieving a list of subject sets, downloading classification and subject data, setting up volunteer feedback and training. This tutorial makes use of example material (subjects, metadata, classifications) from the SuperWASP Variable Stars Zooniverse project, which in turn was created as an EOSC demonstrator, and which is also discussed in Sect. 3.3. The notebook is available at <https://git.astron.nl/astron-sdc/escape-wp5/workflows/zooniverse-advanced-project-building>, and a recorded walkthrough of this advanced tutorial is available at <https://youtu.be/o9SzgsZvOCg>.

The ESAP Archives (accessible via the ESAP GUI) include data retrieval from the Zooniverse Classification Database using the ESAP Shopping Basket. We therefore also created a tutorial on loading Zooniverse data from a saved shopping basket into a notebook and performing simple aggregation of the classification results (also available as an interactive analysis workflow). The tutorial notebook includes importing the Panoptes Python Client and ESAP client Zooniverse connector, retrieving items added to an ESAP Shopping Basket, and an example of retrieval and analysis of data from the Muon Hunter Zooniverse project, which in turn was created as an EOSC demonstrator, and which is also discussed in Sect. 3.2 above. The notebook covers the downloading and processing of classification data, as well as plotting the data and performing simple aggregation of the results. The tutorial is available at <https://git.astron.nl/astron-sdc/escape-wp5/workflows/muon-hunters-example/>.

We also created a Jupyter notebook tutorial and documentary materials demonstrating how to set up an active learning framework to continuously train machine learning models using volunteer classifications of optimally selected subjects. Zooniverse's Caesar advanced retirement and aggregation engine allows for the setup of more advanced rules for retiring subjects (as opposed to the default method of retiring after that subject has been classified a certain number of times). Caesar also provides a powerful way of collecting and analysing volunteer classifications (aggregation). This tutorial makes use of example material (subjects, metadata, classifications) from the Penguin Watch Zooniverse project, which involves counting the numbers of penguin adults, chicks and eggs in images to help understand their lives and environment. A recorded walkthrough of this advanced tutorial is available at <https://>

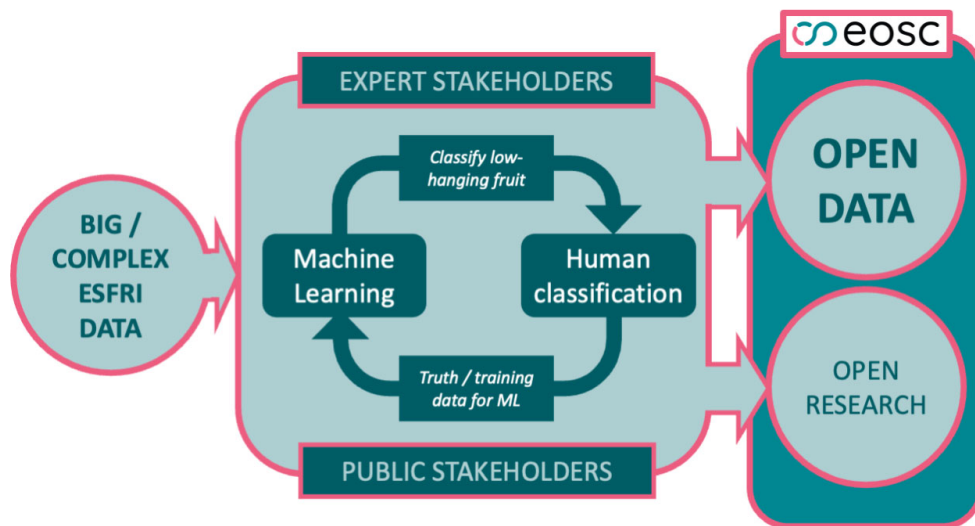


Fig. 1 Schematic PERT diagram showing selected workflows of big/complex data from the European Strategic Forum for Research Infrastructures (ESFRIs), and the external context of the European Open Science Cloud (EOSC). This illustrative workflow shows the virtuous circle between human and machine learning

youtu.be/o9SzgsZvOCg?t=3840 while the notebook itself is at <https://git.astron.nl/astron-sdc/escape-wp5/workflows/zooniverse-advanced-aggregation-with-caesar>.

Finally, we created a notebook demonstrating how to integrate Zooniverse projects with existing machine learning frameworks and combine volunteer classifications with machine learning predictions. The tutorial introduces advanced retirement rules using machine learning, with choices of either pre-classifying with machine learning in preparation for volunteer classifications, or performing “on the fly” retirement decisions made after both machine learning and volunteer classifications. The tutorial shows how to use machine learning to filter out uninteresting subjects prior to volunteer classification and demonstrates setting up active learning in which volunteer classifications train the machine learning model, which in turn handles the “boring” subjects and leaves the more challenging/interesting subjects for volunteers. The notebook is available at <https://git.astron.nl/astron-sdc/escape-wp5/workflows/zooniverse-integrating-machine-learning> and a recorded walkthrough is at <https://youtu.be/o9SzgsZvOCg?t=8218>, conducted as part of our first ESCAPE Citizen Science Workshop.

4.2 Ambition for the future

Our vision is for the community to have access to clear exemplars of planning, creating and managing crowdsourced data mining on EOSC, implementing machine learning in real time, across a wide range of scientific domains, which they can use as templates to deploy with ease. Through these activities, our vision is to increase the size of the community making real scientific engagement with EOSC by orders of magnitude, solving the difficult problem of usefulness of FAIR data by giving non-specialists a carefully curated and educationally supportive experience of EOSC. The workflow is illustrated schematically in Fig. 1.

In realising this EOSC citizen science vision, there are some important lessons from ESCAPE and other projects. One is that interoperability between citizen science services is not automatic; rather, it requires (funded) developer resource to create the interoperability, as was the case in the Zooniverse/CitSci.org cross-integration. Secondly, tutorials, notebook resources and other training materials also need funding for their creation. Thirdly, even with abundant online training materials, the most effective intervention is to have a “support scientist” on hand to assist in the creation, operation and ultimately the science exploitation of citizen science projects. Finally, all the citizen science demonstrators created in ESCAPE were driven by fundamental science questions, and had active commitments from the subject-specialist science teams who were motivated to find answers to these fundamental science questions. Science team engagement with volunteers is one of the key factors influencing success in citizen science projects, and all the ESCAPE citizen science demonstrators have been objectively extremely successful, but the lesson in EOSC is deeper: technical developments are, and must be, precipitated bottom-up by the science use cases. Features that are “nice to have” are not built unless they also have science drivers. This is consistently true in citizen science contexts such as the Zooniverse, and arguably also should apply more widely to EOSC development in general.

With those lessons in mind, we aspire to closer cooperation of citizen science platforms with EOSC. Our vision is to increase citizen science integration with EOSC services, expanding citizen science domains through EOSC. Partly, this will require advocates skilled in using citizen science platforms to engage with wider subject-specialist communities, possibly (though not necessarily) through the creation of multi-disciplinary demonstrator projects driven by specific science needs from those communities. It may

also help to train, or to simply provide as a service, the data visualisation tools that are widely used in the ESCAPE science domain areas but less common elsewhere. We would also support the creation of a Citizen Science Task Force within EOSC communities such as the EOSC Association.

We also aspire to deeper technology integration. Citizen science volunteers in the astrophysics domain already routinely ask or expect the functionality to navigate from an astronomical image subject in Zooniverse, out to a virtual observatory interface to interrogate other multi-wavelength data or to explore the data to be classified with a greater flexibility that is available on the Zooniverse platform alone, all in aid of open-ended and curiosity-driven enquiry by volunteers. For technical reasons relating to the underlying technologies, it is not simple to integrate virtual observatory tools directly into the Zooniverse platform itself, but the links out and in again can be made as frictionless as possible for particular science cases. Again, developments must be driven by community need, whether those communities are subject specialists or the science-inclined public.

One important caveat to this deeper technology integration is that astronomy, astroparticle physics and particle physics are relatively benign subject areas for non-specialists to explore data, in that there are not many serious consequences for the misunderstanding and misuse of data. This is not always the case in other domains, such as healthcare or climate science, where the policy stakes are much higher and there are even efforts by malicious actors to disseminate disinformation. The wider integration of EOSC technologies with citizen science must therefore account for the wider policy and ethics landscapes of each science domain.⁴ During the course of the ESCAPE project there has been a growing interest in science analysis platforms as a whole, not just in ESCAPE's own ESFRIs Science Analysis Platform (ESAP), but also in the Rubin observatory science analysis platform, and in the European Space Agency's new Datalabs platform. As part of the efforts of these platforms towards interoperability within the EOSC ecosystem, the citizen science use cases need to be included explicitly. With a suitable funding stream for the server and for compute, it may also be beneficial to run a dedicated ESAP instance for citizen science, if the user communities do not have routine access to these other science analysis platforms.

We would also like to see greater development of open data standards in citizen science, including creating a registry of terminology, models/formats, defining reporting guidelines, describing data policies, and creating models for identifier schemas related to citizen science.

Finally, as machine learning technologies improve, the synergies and opportunities with crowdsourced data mining will evolve. This may require more science-driven multi-disciplinary worked examples of plug-and-play citizen science notebooks with embedded machine learning integration, to facilitate citizen science uptake throughout the EOSC science domains.

4.3 Could you use citizen science for your physical science research?

We conclude this review with some brief reflections on the types of physical systems and problems that are amenable to citizen science approaches. The most important lesson from the examples above is that every case has been driven by the fundamental science question. In no cases has the activity been driven primarily by a desire for public engagement or outreach. In fact, it is logically impossible to optimise against more than one parameter. For example, one can seek to climb the highest mountain on Earth (Mt. Everest), or the hardest to climb (arguably K2), but it is not possible to simultaneously choose the highest *and* hardest. One can optimise subject to a constraint, such as maximising the science return given the use of crowdsourcing from volunteers or getting to the highest point in a country subject to the constraint that it is accessible by road, but there must always be only one parameter to optimise. Therefore, we do not make recommendations for public engagement or outreach in this review, even though there is abundant evidence that citizen science can demonstrably benefit public engagement, outreach and societal impact (see the outreach and societal impact studies in e.g. [19–22]). If public engagement is the primary goal, then rather than using techniques optimised for something else, in many cases, it may be more effective to make more targeted interventions to support that primary goal; in other words, if the most important factor to you is that you want to do outreach, then you may well be better off optimising your efforts to that specific goal, rather than deploying citizen science activities optimised for something else. We also do not address techniques of co-creation and the use of volunteers throughout the research cycle (e.g. [23]) since, with very few exceptions, this methodology is unsuitable to the physical sciences because of the depth and breadth of background information necessary. Even in the exceptions that exist, the volunteer involvement has to be carefully managed (e.g. [24]).

As a useful heuristic, it is worth considering the motivations for Zooniverse volunteers⁵:

- **Is it beautiful?** In other words, is the project visually engaging? For example, part of the success of the Muon Hunter project (Sect. 3.2) can be attributed to the strikingly pretty colourful arcs to be identified.
- **How easy is it?** Simpler workflows will engage larger cohorts of volunteers. Conversely, complex analysis will draw smaller but more dedicated volunteer teams. At the simplest end, binary classifications can be deployed on the Zooniverse mobile platform, as in Sect. 3.5.
- **How important is it?** The volunteers may be spending a large fraction of their free time on the project, so the onus is on the science team to convince the volunteers that the task is worthwhile and important.

⁴ Further discussion of the policy environment for citizen science can be found in Angelidakis et al., this volume.

⁵ These five insights are due originally to Becky Smethurst.

- **How much am I learning?** The in-built tutorial material supports volunteer learning, and although the learning journey is not the primary objective of the Zooniverse research projects, the projects that provide a sense of personal enrichments are more likely to succeed in practice. Moreover, direct science team engagement on online talk forums is essential not just for fostering volunteer engagement, but also for assessing the progress of the project and discussing volunteer discoveries. Exactly this science team engagement with volunteers led to the successful but unexpected discoveries of the new classes of Voorwerpjes [25] and green pea galaxies [26], identified first by volunteers in Galaxy Zoo.
- **How famous could I get?** Some citizen science projects have a tantalising possibility of a spectacular breakthrough, and this on its own can sometimes be enough to motivate volunteers. An analogy might be made with a lottery: time or money is spent on entertainment, with the tantalising possibility of a large return. In the context of citizen science volunteers it is also worth considering (or just asking) what the volunteers would value as a recognition, which may vary from project to project. For example, while the science team might place the highest value on a co-authorship of a high impact paper, a volunteer might place a higher value on a name appearing on a NASA website.

This is not an exhaustive list of volunteer motivations. Indeed for some citizen science projects, competition may be a factor. We will discuss gamification use cases below, which we use to refer to short-term competition for a particular task or for more long-term achievements on e.g. ranking on a leaderboard.

This review has focused mainly on imaging use cases in which subjective human assessments are used, whether for image annotation or image classification, which in turn can be used as training sets for machine learning technologies. These machines can then be used to “mop up” the easiest classification subjects, allowing human effort to be refocused on more difficult cases, ultimately creating a virtuous circle between human and machine learning (Fig. 1). This on its own captures a very wide variety of research use cases in the physical sciences, but there are many other ways that human effort can be deployed in the physical sciences.

One immediate generalisation is that human volunteers can parse more than visual data. Several efforts have already been made at sonifying astronomical data (e.g. [27, 28]) and there have been several successful Zooniverse projects that require volunteers to parse audio, including assessing the maturity of language development in babies (e.g. [29, 30]) or identifying urban noise in New York with the help of a spectrogram linked to audio (e.g. [31]). In the latter, the spectrograms presented to volunteers are qualitatively similar to those in gravitational wave Zooniverse projects. Perhaps because the professional community is not yet accustomed to parsing their own data through audio, this currently remains an under-explored use case for crowdsourced data mining in the physical sciences.

Going beyond the use of volunteers as distributed, massively parallel biological computers for data mining, there are many other ways in which non-specialists can contribute to research in the physical sciences. Astronomy in particular has a distinguished history of pioneering crowdsourced data collection (e.g. [32]), from the discovery of supernovae and comets to the identification of high energy cosmic rays (Sect. 3.4) to the discovery of meteors at radio and optical wavelengths (e.g. [33]) and of meteorites (e.g. [34, 35]). The Search for Extraterrestrial Intelligence (SETI) pioneered the deployment of otherwise idle CPU on volunteer computers for data analysis, in the SETI@Home screensaver. More opportunities in the physical sciences may arise as the balance evolves between the availability of high performance computing for research, and the potential crowdsourced capabilities of volunteer home computing.

“Amateur” astronomy includes volunteers who make very significant personal financial investments in hardware for observational astronomy, to the extent that there is an overlap between the capabilities of the smallest professional facilities and the largest amateur ones, at around ~ 0.5 m diameter primary mirrors for optical telescopes. This makes it possible for the non-professional community to contribute directly to monitoring campaigns of bright targets, such as measuring transiting exoplanets or monitoring nearby supernova light curves. Moreover, the extremely distributed nature of the data collection can be a decisive advantage for some use cases, such as in multi-messenger astronomy, where the detection of an optical transient may be prevented at one location by local weather conditions but possible at another. Even with the advent of large sky monitoring projects such as the Legacy Survey of Space and Time at the Vera Rubin Observatory, there will always be multi-messenger events inaccessible to any particular facility and so some potential for non-professional contributions will continue. The International Astronomical Union has a working group dedicated to promoting and facilitating research between professional and non-professional communities, at https://www.iau.org/science/scientific_bodies/working_groups/professional-amateur/.

Some citizen science projects have explored other approaches to motivating volunteer effort, such as gamification and contests. The FoldIt project seeks volunteer effort to compete to find novel protein folding configurations [36], a problem well-known to be NP-hard [37]. Another computational problem that is NP-hard is to find the longest common subsequence among sequences; this is relevant to the analysis of DNA, so there are other gamified citizen science applications in bioinformatics such as the Phylo project (<https://phylo.cs.mcgill.ca/>) and Borderlands (<https://borderlands.2k.com/news/borderlands-science/>). There are arguably comparably difficult NP-hard problems in the physical sciences, though so far gamification has been attempted for computationally simpler problems, such as exoplanet discovery in Project Discover, with volunteers competing for rankings in accuracy. There may be under-explored opportunities in the physical sciences for volunteer competitions in harder problems.

Acknowledgements We thank the anonymous referee for helpful suggestions. The authors were part-funded by the ASTERICS, ESCAPE and EOSC-Future projects, through a variety of European Commission Horizon calls. ASTERICS is a project supported by the European Commission Framework Programme Horizon 2020 Research and Innovation action under grant agreement no. 653477. ESCAPE—The European Science Cluster of Astronomy and Particle

Physics ESFRI Research Infrastructures has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement no. 824064. The EOSC-Future project is co-funded by the European Union Horizon Programme call INFRAEOSC-03-2020—Grant Agreement Number 101017536.

Funding Funding sources are declared above. All of the software notebooks described in this paper are part of the ESCAPE Open Source Software Repository (OSSR). All of the ESCAPE OSSR content is hosted in the Zenodo ESCAPE2020 community, at <https://zenodo.org/communities/escape2020>

Data Availability Statement No data associated in the manuscript.

Declarations

Conflict of interest The authors declare no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. S. Serjeant, Citizen science in the European open science cloud. *Europhys. News* **54**(2), 20–23 (2023). <https://doi.org/10.1051/epn/2023204>. arXiv: 2307.06896 [astro-ph.IM]
2. A. Swanson, M. Kosmala, C. Lintott, R. Simpson, A. Smith, C. Packer, Snapshot Serengeti, high-frequency annotated camera trap images of 40 mammalian species in an African savanna. *Sci. Data* **2**, 150026 (2015). <https://doi.org/10.1038/sdata.2015.26>
3. T. Daylan, D.P. Finkbeiner, D. Hooper, T. Linden, S.K.N. Portillo, N.L. Rodd, T.R. Slatyer, The characterization of the gamma-ray signal from the central milky way: a case for annihilating dark matter. *Phys. Dark Univ.* **12**, 1–23 (2016). <https://doi.org/10.1016/j.dark.2015.12.005>
4. M. Ackermann, M. Ajello, A. Albert, W.B. Atwood, L. Baldini, J. Ballet, G. Barbiellini, D. Bastieri, R. Bellazzini, E. Bissaldi, R.D. Blandford, E.D. Bloom, R. Bonino, E. Bottacini, T.J. Brandt, J. Bregeon, P. Bruel, R. Buehler, T.H. Burnett, R.A. Cameron, R. Caputo, M. Caragiulo, P.A. Caraveo, E. Cavazzuti, C. Cecchi, E. Charles, A. Chekhtman, J. Chiang, A. Chiappo, G. Chiaro, S. Ciprini, J. Conrad, F. Costanza, A. Cuoco, S. Cutini, F. D'Ammando, F. de Palma, R. Desiante, S.W. Digel, N. Di Lalla, M. Di Mauro, L. Di Venere, P.S. Drell, C. Favuzzi, S.J. Fegan, E.C. Ferrara, W.B. Focke, A. Franckowiak, Y. Fukazawa, S. Funk, P. Fusco, F. Gargano, D. Gasparri, N. Giglietto, F. Giordano, M. Giroletti, T. Glanzman, G.A. Gomez-Vargas, D. Green, I.A. Grenier, J.E. Grove, L. Guillemot, S. Guiriec, M. Gustafsson, A.K. Harding, E. Hays, J.W. Hewitt, D. Horan, T. Jogler, A.S. Johnson, T. Kamae, D. Kocovski, M. Kuss, G. La Mura, S. Larsson, L. Latronico, J. Li, F. Longo, F. Loparco, M.N. Lovellette, P. Lubrano, J.D. Magill, S. Maldera, D. Malyshev, A. Manfreda, P. Martin, M.N. Mazziotta, P.F. Michelson, N. Mirabal, W. Mitthumsiri, T. Mizuno, A.A. Moiseev, M.E. Monzani, A. Morselli, M. Negro, E. Nuss, T. Ohsugi, M. Orienti, E. Orlando, J.F. Ormes, D. Paneque, J.S. Perkins, M. Persic, M. Pesce-Rollins, F. Piron, G. Principe, S. Rainò, R. Rando, M. Razzano, S. Razzaque, A. Reimer, O. Reimer, M. Sánchez-Conde, C. Sgrò, D. Simone, E.J. Siskind, F. Spada, G. Spandre, P. Spinelli, D.J. Suson, H. Tajima, K. Tanaka, J.B. Thayer, L. Tibaldo, D.F. Torres, E. Troja, Y. Uchiyama, G. Vianello, K.S. Wood, M. Wood, G. Zaharijas, S. Zimmer, Fermi LAT collaboration: the fermi galactic center gev excess and implications for dark matter. *ApJ* **840**(1), 43 (2017). <https://doi.org/10.3847/1538-4357/aa6cab>. arXiv:1704.03910 [astro-ph.HE]
5. Q. Feng, J. Jarvis, VERITAS Collaboration: A citizen-science approach to muon events in imaging atmospheric Cherenkov telescope data: the Muon Hunter, in 35th International Cosmic Ray Conference (ICRC2017). International Cosmic Ray Conference, 301, 826 (2017). <https://doi.org/10.22323/1.301.0826>
6. R. Bird, M.K. Daniel, H. Dickinson, Q. Feng, L. Fortson, A. Furniss, J. Jarvis, R. Mukherjee, R. Ong, I. Sadeh, D. Williams, Muon Hunter: a Zooniverse project. arXiv e-prints, 1802–08907 (2018). <https://doi.org/10.48550/arXiv.1802.08907> [astro-ph.IM]
7. M. Laraia, D. Wright, H. Dickinson, A. Simenstad, K. Flanagan, S. Serjeant, L. Fortson, VERITAS Collaboration: Muon Hunter 2.0: efficient crowd-sourcing of labels for IACT image analysis, in 36th International Cosmic Ray Conference (ICRC2019). International Cosmic Ray Conference, vol. 36, p. 678 (2019). <https://doi.org/10.22323/1.358.0678>
8. R. Bird, M.K. Daniel, H. Dickinson, Q. Feng, L. Fortson, A. Furniss, J. Jarvis, R. Mukherjee, R. Ong, I. Sadeh, D. Williams, Muon Hunter: a Zooniverse project. *J. Phys. Conf. Ser.* **1342**, 012103 (2020). <https://doi.org/10.1088/1742-6596/1342/1/012103>
9. K. Flanagan, D. Wright, H. Dickinson, P. Wilcox, M. Laraia, S. Serjeant, M. Capasso, R. Ong, I. Sadeh, P. Kaaret, W. Jin, W. Benbow, R. Mukherjee, R. Prado, M. Lundy, S. Patel, P. Moriarty, G. Maier, A. Furniss, K. Ragan, D. Williams, J. Buckley, L. Fortson, J. Quinn, J. Holder, C. Giuri, E. Pueschel, D. Nieto, C. Adams, S. O'Brien, D. Ribeiro, K. Pfrang, O. Gueta, G. Foote, A. Weinstein, S. Kumar, T. Williamson, D. Tak, C. McGrath, T. Kleiner, M. Pohl, P. Reynolds, B. Hona, D. Hanna, M. Santander, G. Sembroski, S.R. Patel, M. Errando, M. Kertzman, O. Hervet, M. Nievas-Rosillo, M. Lang, E. Roache, T.B. Humensky, R.Y. Shang, V. Vassiliev, A. Chromey, A. Falcone, J. Christiansen, A. Otte, A.E. Gent, A. Brill, J. Ryan, K. Farrell, G. Gillanders, Q. Feng, A. Archer, D. Kieda, Identifying muon rings in VERITAS data using convolutional neural networks trained on images classified with Muon Hunters 2, in 37th International Cosmic Ray Conference, p. 766 (2022). <https://doi.org/10.22323/1.395.0766>
10. A.J. Norton, A Zooniverse project to classify periodic variable stars from superWASP. *Res. Notes Am. Astron. Soc.* **2**(4), 216 (2018). <https://doi.org/10.3847/2515-5172/aaf291>
11. H.B. Thiemann, A.J. Norton, H.J. Dickinson, A. McMaster, U.C. Kolb, SuperWASP variable stars: classifying light curves using citizen science. *MNRAS* **502**(1), 1299–1311 (2021). <https://doi.org/10.1093/mnras/stab140>
12. K.W. Woźniak, Detection of cosmic-ray ensembles with CREDO. *Eur. Phys. J. Web Conf.* **208**, 15006 (2019). <https://doi.org/10.1051/epjconf/201920815006>
13. T. Wibig, M. Karbowski, D. Alvarez-Castillo, O. Bar, Ł. Bibrzycki, D. Gora, P. Homola, P. Kovacs, M. Piekarczyk, J. Stasielak, S. Stuglik, O. Sushchov, A. Tursunov, Determination of Zenith Angle Dependence of Incoherent Cosmic Ray Muon Flux Using Smartphones of the CREDO Project, in 37th International Cosmic Ray Conference, p. 199 (2022). <https://doi.org/10.22323/1.395.0199>

14. A. Tursunov, P. Homola, D. Alvarez-Castillo, N. Budnev, A. Gupta, B. Hnatyk, M. Kasztelan, P. Kovacs, B. Łozowski, M. Medvedev, A. Mozgova, M. Niedzwiecki, M. Pawlik, M. Rosas, K. Rzecki, K. Smelcerz, K. Smolek, J. Stasielak, S. Stuglik, M. Svanidze, O. Sushchov, Y. Verbetsky, T. Wibig, J. Zamora-Saa.: Credo Collaboration: Probing UHECR and cosmic ray ensemble scenarios with a global CREDO network, in: 37th International Cosmic Ray Conference, p. 472 (2022). <https://doi.org/10.22323/1.395.0472>
15. R.B. Metcalf, M. Meneghetti, C. Aveztruz, F. Bellagamba, C.R. Bom, E. Bertin, R. Cabanac, F. Courbin, A. Davies, E. Decenci re, R. Flamary, R. Gavazzi, M. Geiger, P. Hartley, M. Huertas-Company, N. Jackson, C. Jacobs, E. Jullo, J.-P. Kneib, L.V.E. Koopmans, F. Lanusse, C.-L. Li, Q. Ma, M. Makler, N. Li, M. Lightman, C.E. Petrillo, S. Serjeant, C. Sch fer, A. Sonnenfeld, A. Tagore, C. Tortora, D. Tuccillo, M.B. Valent n, S. Velasco-Forero, G.A. Verdoes Kleijn, G. VERNARDOS, The strong gravitational lens finding challenge. *A & A* 625, 119 (2019). <https://doi.org/10.1051/0004-6361/201832797>
16. A. Davies, Using machine learning techniques to detect, classify and separate strong gravitational lensing systems from astronomical images. PhD thesis, The Open University, The Open University, Milton Keynes, MK7 6AA, UK (2022). <https://doi.org/10.21954/ou.ro.000166ad>
17. H. Dickinson, D. Adams, V. Mehta, C. Scarlata, L. Fortson, S. Serjeant, C. Krawczyk, S. Kruk, C. Lintott, K.B. Mantha, B.D. Simmons, M. Walmsley, Galaxy Zoo: Clump Scout—Design and first application of a two-dimensional aggregation tool for citizen science. *MNRAS* 517(4), 5882–5911 (2022). <https://doi.org/10.1093/mnras/stac2919>. arXiv:2210.03684 [astro-ph.GA]
18. D. Adams, V. Mehta, H. Dickinson, C. Scarlata, L. Fortson, S. Kruk, B. Simmons, C. Lintott, Galaxy zoo: clump scout: surveying the local universe for giant star-forming clumps. *ApJ* 931(1), 16 (2022). <https://doi.org/10.3847/1538-4357/ac6512>. arXiv:2201.06581 [astro-ph.GA]
19. E.E. Prather, S. Cormier, C.S. Wallace, C. Lintott, M. Jordan Raddick, A. Smith, Measuring the conceptual understandings of citizen scientists participating in zooniverse projects: a first approach. *Astron. Edu. Rev.* (2013). <https://doi.org/10.3847/aer2013002>
20. M. Luczak-Roesch, R. Tinati, E. Simperl, M. Van Kleek, N. Shadbolt, R. Simpson, Why won't aliens talk to us? content and community dynamics in online citizen science. *Proc. Int. AAAI Conf. Web Soc. Media* 8(1), 315–324 (2014). <https://doi.org/10.1609/icwsm.v8i1.14539>
21. K. Masters, E.Y. Oh, J. Cox, B. Simmons, C. Lintott, G. Graham, A. Greenhill, K. Holmes, Science learning via participation in online citizen science. arXiv (2016). <https://doi.org/10.48550/ARXIV.1601.05973>. <https://arxiv.org/abs/1601.05973>
22. L. Trouille, T. Nelson, J. Feldt, J. Keller, M. Buie, C. Cardamone, B.C. Kung, K. Masters, K. Meredith, K. Borden, Citizen science in astronomy education. *Astron. Edu.* 2514–3433, 8–1824 (2019). <https://doi.org/10.1088/2514-3433/ab2b42ch8>
23. E. Senabre Hidalgo, J. Perell , F. Becker, I. Bonhoure, M. Legris, A. Cigarini, In: Vohland, K., Land-Zandstra, A., Ceccaroni, L., Lemmens, R., Perell , J., Ponti, M., Samson, R., Wagenknecht, K. (eds.) Participation and Co-creation in Citizen Science, pp. 199–218. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-58278-4_11
24. W.C. Keel, J. Tate, O.I. Wong, J.K. Banfield, C.J. Lintott, K.L. Masters, B.D. Simmons, C. Scarlata, C. Cardamone, R. Smethurst, L. Fortson, J. Shanahan, S. Kruk, I.L. Garland, C. Hancock, D. O'Ryan, Gems of the galaxy zoos—a wide-ranging hubble space telescope gap-filler program*. *Astron. J.* 163(4), 150 (2022). <https://doi.org/10.3847/1538-3881/ac517d>
25. C.J. Lintott, S.P. Schawinski, W. Keel, H. van Arkel, N. Bennert, E. Edmondson, D. Thomas, D.J.B. Smith, P.D. Herbert, M.J. Jarvis, S. Virani, D. Andreescu, S.P. Bamford, K. Land, P. Murray, R.C. Nichol, M.J. Raddick, A. Slosar, A. Szalay, J. Vandenberg, Galaxy Zoo: 'Hanny's Voorwerp', a quasar light echo? *MNRAS* 399(1), 129–140 (2009). <https://doi.org/10.1111/j.1365-2966.2009.15299.x>
26. C. Cardamone, K. Schawinski, M. Sarzi, S.P. Bamford, N. Bennert, C.M. Urry, C. Lintott, W.C. Keel, J. Parejko, R.C. Nichol, D. Thomas, D. Andreescu, P. Murray, M.J. Raddick, A. Slosar, A. Szalay, J. Vandenberg, Galaxy zoo green peas: discovery of a class of compact extremely star-forming galaxies. *MNRAS* 399(3), 1191–1205 (2009). <https://doi.org/10.1111/j.1365-2966.2009.15383.x>
27. G. De La Vega, L.M. Exequiel Dominguez, J. Casado, B. Garc a, SonoUno web: an innovative user centred web interface. arXiv e-prints, 2302–00081 (2023) <https://doi.org/10.48550/arXiv.2302.00081>
28. A. Zanella, C.M. Harrison, S. Lenzi, J. Cooke, P. Damsma, S.W. Fleming, Sonification and sound design for astronomy research, education and public engagement. *Nat. Astron.* 6, 1241–1248 (2022). <https://doi.org/10.1038/s41550-022-01721-z>
29. C. Semenzin, L. Hamrick, A. Seidl, B.L. Kelleher, A. Cristia, Describing vocalizations in young children: A big data approach through citizen science annotation. *OSF Preprints* (2020). <https://doi.org/10.31219/osf.io/z6exv>
30. C. Semenzin, L. Hamrick, A. Seidl, B.L. Kelleher, A. Cristia, Towards large-scale data annotation of audio from wearables: validating zooniverse annotations of infant vocalization types. *OSF Preprints* (2020). <https://doi.org/10.31219/osf.io/gpxf5>
31. M. Cartwright, G. Dove, A.E. M endez M endez, J.P. Bello, O. Nov, Crowdsourcing multi-label audio annotation tasks with citizen scientists, in Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. CHI '19, pp. 1–11. Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3290605.3300522>. <https://doi.org/10.1145/3290605.3300522>
32. A. Chapman, Victorian Amateur Astronomer: Independent Astronomical Research in Britain 1820 - 1920, (2017)
33. J. Rendtel, Review of amateur meteor research. *Planet. Space Sci.* 143, 7–11 (2017). <https://doi.org/10.1016/j.pss.2017.01.007>
34. A.J. King, L. Daly, J. Rowe, K.H. Joy, R.C. Greenwood, H.A.R. Devillepoix, M.D. Suttle, Q.H.S. Chan, S.S. Russell, H.C. Bates, J.F.J. Bryson, P.L. Clay, D. Vida, M.R. Lee,  . O'Brien, L.J. Hallis, N.R. Stephen, R. Tart se, E.K. Sansom, M.C. Towner, M. Cupak, P.M. Shober, P.A. Bland, R. Findlay, I.A. Franchi, A.B. Verchovsky, F.A.J. Abernethy, M.M. Grady, C.J. Floyd, M. Van Ginneken, J. Bridges, L.J. Hicks, R.H. Jones, J.T. Mitchell, M.J. Genge, L. Jenkins, P.-E. Martin, M.A. Sephton, J.S. Watson, T. Salge, K.A. Shirley, R.J. Curtis, T.J. Warren, N.E. Bowles, F.M. Stuart, L. Di Nicola, D. Gy re, A.J. Boyce, K.M.M. Shaw, T. Elliott, R.C.J. Steele, P. Povinec, M. Laubenstein, D. Sanderson, A. Cresswell, A.J.T. Jull, I. S ykora, S. Sridhar, R.J. Harrison, F.M. Willcocks, C.S. Harrison, D. Hallatt, P.J. Wozniakiewicz, M.J. Burchell, L.S. Alesbrook, A. Dignam, N.V. Almeida, C.L. Smith, B. Clark, E.R. Humphreys-Williams, P.F. Schofield, L.T. Cornwell, V. Spathis, G.H. Morgan, M.J. Perkins, R. Kacerek, P. Campbell-Burns, F. Colas, B. Zanda, P. Vernazza, S. Bouley, S. Jeanne, M. Hankey, G.S. Collins, J.S. Young, C. Shaw, J. Horak, D. Jones, N. James, S. Bosley, A. Shuttleworth, P. Dickinson, I. McMullan, D. Robson, A.R.D. Smedley, B. Stanley, R. Bassom, M. McIntyre, A.A. Suttle, R. Fleet, L. Bastiaens, M.B. Ih sz, S. McMullan, S.J. Boazman, Z.I. Dickeson, P.M. Grindrod, A.E. Pickersgill, C.J. Weir, F.M. Suttle, S. Farrelly, I. Spencer, S. Naqvi, B. Mayne, D. Skilton, D. Kirk, A. Mounsey, S.E. Mounsey, S. Mounsey, P. Godfrey, L. Bond, V. Bond, C. Wilcock, H. Wilcock, R. Wilcock, The Winchcombe meteorite, a unique and pristine witness from the outer solar system. *Sci. Adv.* 8(46), 3925 (2022). <https://doi.org/10.1126/sciadv.abq3925>
35.  .C. O'Brien, A. Pickersgill, L. Daly, L. Jenkins, C. Floyd, P.-E. Martin, L.J. Hallis, A. King, M. Lee, The Winchcombe Meteorite: one year on. *Astron. Geophys.* 63(1), 1–21123 (2022). <https://doi.org/10.1093/astrogeo/atac009>
36. B. Koeppnick, J. Flatten, T. Husain, A. Ford, D.-A. Silva, M.J. Bick, A. Bauer, G. Liu, Y. Ishida, A. Boykov, R.D. Estep, S. Kleinfelter, T. N rg rd-Solano, L. Wei, F. Players, G.T. Montelione, F. DiMaio, Z. Popovi , F. Khatib, S. Cooper, D. Baker, De novo protein design by citizen scientists. *Nature* 570(7761), 390–394 (2019). <https://doi.org/10.1038/s41586-019-1274-4>
37. A.S. Fraenkel, Complexity of protein folding. *Bull. Math. Biol.* 55(6), 1199–1210 (1993). <https://doi.org/10.1007/bf02460704>