Supplemental File 1: Appendix A: Project timeline

2025

- January - March
    - Administrative setup, planning and preparation
    - Began soliciting expressions of interest for workshops
- April - May
    - Completed draft of workshop participants based on criteria set
    - Established workshop and working session dates, times
    - Prepared pre-reading lists for workshops
    - Sent out initial workshop invitations, and additional invitations (if no response or unavailable from initial message)
- June
    - 2 weeks before W1, sent Zooniverse participant survey
    - Finalized W1 agenda using survey response data
    - **Workshop 1 (June 24)**
    - Refined notes from W1 and summarized key concepts
- July
    - 2 weeks before W2, sent Zooniverse participant survey
    - Finalized W2 agenda using survey response data
    - **Workshop 2 (July 24)**
    - Refined notes from W2 and summarized key concepts
- August
    - 2 weeks before W3, sent Zooniverse participant survey
    - Finalized W3 agenda using survey response data
    - *Working Session 1 (August 26)*
    - Refined notes from WS1 and summarized key concepts; shared W1+WS1 outcomes with W1 attendees for feedback
    - **Workshop 3 (August 29)**
    - Refined notes from W3 and summarized key concepts
- September
    - 2 weeks before W4, sent Zooniverse participant survey
    - Finalized W4 agenda using survey response data
    - *Working Session 2 (September 22)*
    - Refined notes from WS2 and summarized key concepts; shared W2+WS2 outcomes with W2 attendees for feedback
- October
    - **Workshop 4 (October 2)**
    - Refined notes from W4 and summarized key concepts
    - *Working Session 3 (October 22)*
    - Refined notes from WS3 and summarized key concepts; shared W3+WS3 outcomes with W3 attendees for feedback

    - *Working Session 4 (October 29)*
    - Refined notes from WS4 and summarized key concepts; shared W4+WS4 outcomes with W4 attendees for feedback
- November - December
    - Reconciled Workshop and Working Session outcomes with attendee feedback
    - Began drafting recommendations

2026

- January - July
    - Continued drafting and iterating recommendations (now called "Zooniverse AI Ethics Framework"), with input from Zooniverse leadership team
    - Finalized communications plan for Ethics Framework launch
- August
    - Finalized copy for all pages and launch plan (Framework, help.zooniverse.org guidance for project teams, blog post, email announcement, social media content)
- September
    - Launched Zooniverse AI Ethics Framework (September 9, 2026) and carried out all components of launch plan

Appendix B: Workshop participant solicitation and selection

Workshop participants were identified through a multi-step process that included emails sent to our existing Zooniverse researcher and volunteer community, and public requests for expressions of interest disseminated among the broader participatory science field.

This included the following:

- Identified potential participants at the AAAS and CAPS annual conferences in 2025, using a QR code linked to a Google Form to collect expressions of interest for people who wanted to participate as workshop participants
- Sent an email to Zooniverse project teams, requesting expressions of interest via a Google Form (including name and any professional title/affiliation, email address, and which of the 4 workshops they would be interested in)
- Asked our Advisory Board to share among their networks

We used a Google Sheet to track expressions of interest received via direct email, social media messages, Google Form sign-ups, or other. On the Google Sheet, we noted information about possible workshop participants to determine where they fit into our pre-selected categories, which included:

Stakeholder role: Zooniverse leadership, project lead, volunteer, disciplinary expert, platform maintainer (i.e. non-Zooniverse participatory research specialists)

- This information allowed us to make sure each workshop included the appropriate number of participants. We established 15 as our ideal number, to keep conversations intimate and ensure that everyone participating would have a chance to be heard within a 2.5-hour discussion.
- Each group of 15 would include:
    - Project Coordinator (Hillary Burgess)
    - Workshop Facilitator (Tanisha Tate Woodson, Catalyst Consulting)
    - 2 UC-Berkeley Kavli Center for Ethics, Science, and the Public Fellows (Chad Harper and Kaylene Caswell Stocking)
    - 3 topic-area experts (based on workshop theme)
    - 3 Zooniverse volunteers
    - 2-3 Zooniverse project leads
    - 2-3 members of the Zooniverse leadership team (Project PI Samantha Blickhan and 1-2 other leadership team members)

Discipline or research specialty: e.g. Anthropology, Artificial Intelligence, Astronomy, Biomedical, Biology, Citizen/Participatory Science (broadly), Communications, Copyright, Data Licensing, Digital Humanities, Ecology, Ethics, History, Law, Library Science, Linguistics, Machine Learning, Open Research/Open Science, Policy & Governance

Project affiliation (if any): This included Zooniverse projects as well as other participatory research platforms and websites, including iNaturalist, FromThePage, CitSci.org, as well as individual projects not affiliated with participatory science platforms.

Appendix C: Survey preamble

Surveys were sent out to anyone subscribed to the global Zooniverse email list. We also announced the surveys through a blog post on June 10, 2026.[1]

The following text was included at the beginning of all surveys sent to the Zooniverse volunteer community, to explain the purpose of their participation and what would be done with the data collected.

> What is the purpose of this survey? Results from this and other surveys will help inform a set of recommendations for running Machine Learning and Artificial Intelligence-engaged projects on the Zooniverse platform (NASA definition of AI).
>
> Why have I been invited to participate? You have been invited to participate because you have contributed to one or more Zooniverse projects.
>
> Do I have to participate? You do not have to participate. Participation is voluntary.

[1] https://blog.zooniverse.org/2025/06/10/ethical-considerations-for-machine-learning-in-public-engaged-research/

What are the benefits of participating? There are no immediate benefits for participants, but your contributions will help us understand participant perspectives on ML/AI and will help ensure that the interests, priorities and concerns of Zooniverse project participants are reflected in recommendations for the platform and project leaders.

How will we use information about you? The survey has been designed so that you can answer anonymously if you wish and will not be identifiable by those with access to the data.

What will we do with the results, and who can access them? The survey results will be summarized and shared with Zooniverse leadership. Anonymized summaries of responses may appear in a report or peer-reviewed publication in support of the recommendations developed through this project.

Who is running this study? The Project Director is Dr. Samantha Blickhan, Zooniverse Co-Director and Digital Humanities Lead.

Who is funding this research? This research is funded by the Kavli Foundation.

How can I contact the team? Questions can be addressed to hillary@zooniverse.org or samantha@zooniverse.org

Appendix D: Preparatory materials for workshops

This section contains the pre-reading, questions, and agendas sent out to participants in advance of each workshop.

Workshop 1: Transparency and Communication Best Practices

Ceccaroni et al. (2019). Opportunities and Risks for Citizen Science in the Age of Artificial Intelligence. *Citizen Science Theory & Practice* 4(1), 29. https://doi.org/10.5334/cstp.241

- This pre-reading is intended to serve as a reminder of how ML/AI is used in participatory research projects (see especially Tables 1 and 2) and provide a general overview of potential opportunities and risks.

Gillespie, N., Lockey, S., Ward, T., Macdade, A., & Hassed, G. (2025). Trust, attitudes and use of artificial intelligence: A global study. The University of Melbourne and KPMG. https://doi.org/10.26188/28822919. Note: Executive Summary only (p. 4-10)

- This pre-reading is a helpful reminder of the reasons for hosting this workshop series, i.e. that public opinion around the concept of ML/AI is varied and often complex. This Executive Summary provides some data-driven insight into public opinion via the results of a recent global survey of 48,000 people.

Bhaskar Chakravorti (2024) AI's Trust Problem: Twelve persistent risks of AI that are driving skepticism. *Harvard Business Review*. https://hbr.org/2024/05/ais-trust-problem

- This pre-reading is intended to help us think about how 'AI' is being communicated to the general public, outside academic journals or in participatory research message boards (like Talk). It is not a reflection of the opinions of the workshop organizers, but rather an exercise to consider how participatory research communications must take into account the variety of messaging people may be encountering in their daily lives. We would specifically like you to consider what we might learn from a piece like this, as well as what you might critique or change about it from a communications perspective.

Review a few Zooniverse Projects that incorporate ML/AI:

- This pre-reading is a reminder of the variety of ways that ML/AI is used across the Zooniverse platform, as well as providing a range of examples of how teams are communicating their project's ML/AI use. Please read each project's About page, linked below. For active projects, you may also wish to submit a classification (or at least test the workflow). We will discuss these examples during the workshop.
    - Click-A-Coral https://www.zooniverse.org/projects/jordan-pierce/click-a-coral/about/research
    - Are we alone in the universe? https://www.zooniverse.org/projects/ucla-seti-group/are-we-alone-in-the-universe/about/research
    - Eye for Diabetes https://www.zooniverse.org/projects/oogvoordiabetes/eye-for-diabetes/about/research

Questions to consider before the workshop:

- In your experience, does the use of ML/AI impact motivation to participate in Zooniverse projects? Why?
- How do the terminology and language used to describe ML/AI influence public perception?
- What role does ML/AI literacy play in informing motivation? Informing consent?
- Do platforms for public-engaged research (such as Zooniverse) have a responsibility to facilitate literacy in ML/AI if they are hosting projects that incorporate this technology? If so, what could this look like?
- What information should platform participants have access to, in deciding whether or not to engage with a project?

AGENDA
Tuesday, June 24, 2025
1:00pm - 3:00pm Central Time

- Introductions
- Project overview (10 minutes)
- Session 1 (30 minutes; full-group discussion)
    - Experiences giving and receiving communication on AI/ML in participatory research
    - Reflections on examples from pre-reading

- Session 2 (20 minutes; breakout groups; end with 10-minute share-out)
    - Participant motivation and AI/ML
- Break (10 minutes)
- Session 3 (25 minutes; breakout groups; end with 10-minute share-out)
    - Transparency, AI literacy, agency, and consent
- Brainstorming actions
    - Add sticky notes to Miro board
    - Read and vote
    - Group discussion
- Recap and next steps
- Feedback from participants (brief survey via QR code)

Workshop 2: Ethical Approaches to AI

Pre-reading and Questions

Ceccaroni et al. (2019). Opportunities and Risks for Citizen Science in the Age of Artificial Intelligence. *Citizen Science Theory & Practice* 4(1), 29. https://doi.org/10.5334/cstp.241

- This pre-reading is intended to serve as a reminder of how AI/ML is used in participatory research projects (see especially Tables 1 and 2) and provide a general overview of potential opportunities and risks.

Principles for the Ethical Use of Artificial Intelligence in the United Nations System (2022)

- This pre-reading is intended to provide an example of how ethical information about AI can be operationalized at a high level but intended for a broad range of use cases.

Cooper, C. B., Rasmussen, L. M., and Jones, E. D. (2022). A Toolkit for Data Ethics in the Participatory Sciences. Citizen Science Association.

- Pages 5-15 only
- This pre-reading is intended to provide an example of how ethical information can be operationalized in a relevant context (i.e. participatory research), but that doesn't include the scope of our work (AI/ML).

Questions to consider before the workshop:

- What ethical questions and issues arise at the intersection of participatory science and AI/machine learning?
    - What harm could be caused if these aren't addressed? What scenarios do we want to avoid?
- What actions could be taken to address these?
- Who is responsible to carry out actions or ensure key principles are adhered to? (Zooniverse and/or Project Leaders and/or Participants and/or outside of this community's locus of control, e.g. need to use a legal framework)

- Are there existing practices, standards, frameworks that we could draw on during this process (even that aren't related to participatory research)?

AGENDA
Thursday, July 24, 2025
10:30am - 1:00pm Central Time

- Introductions
- Project overview (10 minutes)
- Session 1 (45 minutes; full-group discussion)
    - Ethical questions and issues; harm avoidance
- Session 2 (30 minutes; breakout groups; end with 10-minute share-out)
    - Identify actions and principles to address ethical issues
- Break (10 minutes)
- Session 3 (20 minutes; breakout groups; end with 10-minute share-out)
    - Roles, responsibility, and accountability
- Recap and next steps
- Feedback from participants (brief survey via QR code)

Workshop 3: Deepening Contextual Understanding

Pre-reading and Questions

Ceccaroni et al. (2019). Opportunities and Risks for Citizen Science in the Age of Artificial Intelligence. *Citizen Science Theory & Practice* 4(1), 29. https://doi.org/10.5334/cstp.241

- This pre-reading is intended to serve as a reminder of how AI/ML is used in participatory research projects (see especially Tables 1 and 2) and provide a general overview of potential opportunities and risks.

Review the following 3 scenarios and reflect on the questions that follow. These scenarios are based on real Zooniverse projects, and in some cases on real research studies. Each case has been adapted for the purposes of our speculative discussions in this workshop, with actual studies cited when possible.

Scenario 1: Using AI to alter volunteer experience and workflow
Project: Snapshot Serengeti
Purpose: Study how over 30 species are distributed across the Serengeti National park landscape – and how they interact with lions and one another. Insights guide strategies for species reintroduction, conservation, and ecosystem management around the world
Task: Wildlife identification from camera trap images
ML Application: Raw image library contains many "blanks" (no wildlife subjects within the image). When trained on crowd-labeled data, machines can be applied at a pre-processing stage, to remove

'blank' images from the workflow. In turn, participants are able to focus their efforts on identifying species in the images.

Variation Scenario 1: Using AI to alter volunteer experience and workflow
Project: Snapshot Serengeti
Task: Wildlife identification from camera trap images
ML Application: Data on individualized user engagement can be used to determine the kinds of images that participants find rewarding. Machine systems can insert rewarding images at a frequency optimized to retain participants, based on previous classification behavior patterns.

Scenario 2: Training and integration of AI models for task optimization
Project: Civil War Bluejackets
Purpose: Create a database of every sailor involved in the American Civil War enabling exploration of historic naval records in the context of race, ethnicity, and class.
Task: In Phase 1, participants were shown full page scans and asked to transcribe different elements of Civil War musters. In Phase 2, participants are shown cropped sections of pages and asked to review the computer's transcriptions, verify their accuracy, and provide corrections.
ML Application: Participant-generated transcription data is used to train ML models on 19th-century handwriting. In Phase 2 of the project, the task shifts from volunteer transcription to volunteer review of machine-transcribed data, allowing for machine model improvement.

Scenario 3: Choosing not to use AI
Project: Everglades Wildlife Watch
Purpose: To document wildlife and the effects of the Florida Fish and Wildlife Commission's habitat management and restoration
Task: Wildlife identification from camera trap images
ML Application: When it first launched, this team did not use ML in their project for reasons including a lack of institutional guidance around the use of ML/AI in work products. Instead, the team chose to follow a 'traditional' participatory science approach where volunteers identify wildlife in camera trap photos. The team has indicated that the impact of their volunteer engagement has led to things like increased understanding among participants around the importance of conservation efforts in the region, as well as a spike in volunteering in local, in-person opportunities related to their research efforts.

Questions to consider about these scenarios:

- How might the following outcomes be impacted by the inclusion of Machine Learning (or not) in each scenario?
    - Learning
    - Skill building
    - Motivation
    - Engagement
    - Data quality
    - Trust
    - Efficiency

- How could participant agency, disclosure, and consent be addressed in this scenario? In this context, we define "participant agency" as the capacity of an individual to make choices, act, or exert power; "disclosure" as transparency around the application and methods of AI/ML use; and "consent" as giving permission to do something (e.g. to use personal data or contributions for a certain purpose).

- What if you swapped-in a different discipline or subject matter (e.g. astronomy or humanities for biology)? What if you swapped-in different tasks?

AGENDA
Friday, August 29, 2025
1:00pm - 3:30pm Central Time

- Introductions
- Project overview (10 minutes)
- Session 1 (45 minutes; breakout groups discussing pre-reading scenarios)
    - Explore how the use of AI/ML may impact outcomes, and how these impacts might vary by project context
- Break (10 minutes)
- Session 2 (50 minutes; full group discussion; start with share-outs from Session 1)
    - Explore how things that participants bring to the table (e.g. AI literacy, motivation, types of engagement) impact project outcomes
- Recap and next steps
- Feedback from participants (brief survey via QR code)

Workshop 4: Downstream Data Protection

Pre-reading and Questions

Ceccaroni et al. (2019). Opportunities and Risks for Citizen Science in the Age of Artificial Intelligence. *Citizen Science Theory & Practice* 4(1), 29. https://doi.org/10.5334/cstp.241

- This pre-reading is intended to serve as a reminder of how AI/ML is used in participatory research projects (see especially Tables 1 and 2) and provide a general overview of potential opportunities and risks.

Zooniverse's policy regarding data stewardship (see "Expectations of Approved Projects")
https://help.zooniverse.org/getting-started/lab-policies/

- This document provides important context about the platform's current policies regarding open data and the guidance that project builders receive

Contractor et al. (2022). Behavioral Use Licensing for Responsible AI.
ACM Conference on Fairness, FAccT '22, June 21–24, 2022, Seoul, Republic of Korea.
https://doi.org/10.1145/3531146.3533143

- This pre-reading articulates some of the tensions between the benefits of “open” work and concerns for potential uses if restrictions aren't in place or enforceable and explores how licenses might be applied to address this conundrum. The “Limitations and Open Questions” section is especially relevant for our discussion.

Questions to consider before the workshop:

- What are the benefits of open data?
- What are the risks of open data?
- Are these domain/subject specific?
- Are there mechanisms to maintain the benefits while reducing the risks?
- What is the responsibility of Zooniverse, Project leads, Participants, with respect to data protection?
- What rights do participants have to determine what happens to the data they contribute?

AGENDA
Thursday, October 2, 2025
3:00pm - 5:30pm Central Time

- Introductions
- Project overview (10 minutes)
- Session 1 (35 minutes; full group discussion)
    - Explore the benefits and risks of open data, how context matters, and existing mechanisms and frameworks
- Session 2 (40 minutes; breakout groups)
    - Deeper dive into mechanisms, limitations, and paths forward
- Break (10 minutes)
- Session 3 (30 minutes; full group discussion; start with share-outs from Session 2)
    - Considering next steps
- Recap and next steps
- Feedback from participants (brief survey via QR code)