Feedback mechanisms for image quality

I am developing a prototype web application to test various feedback mechanisms on users. The goal is to guide them to take better, more useful photos for fossil identification. To validate the results, there will be a baseline prototype and two more to test feedback mechanisms. The images taken by testers are collected and then compared based on five metrics. Sharpness, lighting, contrast, number of angles and scale present.

Created by: Megan
Created on: November 24, 2025 8:33 AM
Changed on: November 27, 2025 1:00 PM

Context of use: Education
Level of education: Master

Impact on society

What impact is expected from your technology?

What is exactly the problem? Is it really a problem? Are you sure?

This technology tackles poor data quality in citizen-submitted fossil images, which limits both scientific usefulness and AI performance in the LegaSea project. Submissions are often blurry or incomplete, feedback is slow, and users receive little guidance, resulting in data that experts struggle to validate and AI models struggle to learn from. By adding real-time capture guidance and a more natural, narrative input flow, the system aims to produce higher-quality fossil data and a smoother user experience.

More expanded version:
This technology addresses a challenge in AI-supported citizen science for paleontology, particularly within the LegaSea project at Naturalis. The main issue is that the scientific value of citizen-submitted fossil data is limited by inconsistent image quality and a slow, static submission workflow. Images are often blurry, poorly lit, or incomplete, and these issues can not be fixed afterward. As a result, AI models trained on these submissions perform worse than models trained on standardized images. The current system also relies on manual expert review, which slows down feedback. The platform provides little guidance for taking good photos, and the lack of tools, real-time feedback, and mobile support results in a poor user experience.
The project aims to solve this by adding (real-time) guidance during data collection and by replacing forms with a more narrative interface that helps users provide more useful information.
The people who benefit include citizen scientists, who will get immediate feedback and a more engaging experience. Researchers at Naturalis, who need large amounts of high-quality fossil data. Volunteer experts, who currently spend much time reviewing submissions, and AI developers, who need consistent, high-quality data to build accurate models.
This is a real problem for both science and AI development. Paleontological research depends on high-quality documentation, and the research goals of LegaSea are hard to reach with inconsistent submissions. AI accuracy drops when trained on citizen-generated images, showing that poor data quality directly limits technological progress.
Solving the problem improves both scientific output and public engagement. Better data support the reconstruction of Ice Age ecosystems, while improved tools help the public learn how to collect research-grade data and participate in science. Overall, the solution transforms unreliable, inconsistent submissions into usable, high-quality scientific data, improving both research and citizen participation.

Are you sure that this technology is solving the RIGHT problem?

Yes, I am confident that this technology is addressing the right problem. The

core issue is the poor quality of citizen-submitted fossil images, which directly impacts both scientific research and AI model performance. The technology is not just addressing surface-level symptoms, like adding a fancy interface or speeding up submission. It targets the root causes of low-quality data: blurry, poorly lit, or incomplete images, and the lack of immediate guidance for users.

I applied the Five Whys approach to confirm this:

1. Why are the submissions of low quality? Because images are often blurry, poorly lit, or incomplete.
2. Why are images blurry or poorly lit? Because users do not have guidance while capturing the photo.
3. Why dont users have guidance? The current system relies on post-submission expert review rather than real-time feedback.
4. Why is post-submission feedback insufficient? Because it is slow, discouraging, and cannot correct mistakes after the image is taken.
5. Why is real-time feedback not provided in the current system? Existing workflows use static forms and lack AI-assisted guidance during capture.

How is this technology going to solve the problem?

This technology solves the problem of poor citizen-submitted fossil images by giving users real-time guidance and a user-friendly submission process. It tackles the root issue by providing instant feedback on lighting, sharpness, framing, and scale, while replacing static forms with an interactive, step-by-step workflow that teaches best practices.
I am confident it will work because it is based on research in user interface design, citizen science, and image processing, which shows that guided input improves data quality and engagement. The prototype includes feedback loops and usability studies to compare submission quality before and after real-time guidance.

What negative effects do you expect from this technology?

Some negative effects I anticipate include user frustration, overreliance on the technology, and exclusion of certain participants. Because the system provides constant real-time guidance, some users (casual contributors) might find it intrusive, distracting, or annoying, which could lead to drop-off. Users with older or low-power devices may also experience slower performance or incompatibility, creating a technological bias.
Overreliance is another risk: users might follow AI guidance blindly without learning the reasoning behind it, reducing their ability to independently assess fossil quality. Theres also the potential for the AI feedback to misinterpret unusual fossils or field conditions, giving misleading guidance.
I accept some of these risks because the benefits (higher-quality, scientifically useful submissions) outweigh them in this prototype. I plan to mitigate negative effects through usability testing, user feedback, and

iterative improvements, such as refining the guidance interface and ensuring accessibility across devices.

In what way is this technology contributing to a world you want to live in?

This technology contributes to a world I want to live in by making science more accessible and engaging. In the short term, it helps people engage with fossil research, learn best practices, and submit higher-quality data. In the long term, it fosters a culture of careful, responsible citizen science, encourages learning, and strengthens the connection between the public and professional researchers. It aligns with my values of transparency, inclusivity, education, and ethical data handling, supporting a society where knowledge is shared, participation is valued, and scientific contributions are meaningful.

Now that you have thought hard about the impact of this technology on society (by filling out the questions above), what improvements would you like to make to the technology? List them below.

Based on my reflections on the problem, its impact, and potential negative effects, I would make the following improvements to the technology:

1. Improve data quality and reduce bias:

- Incorporate multi-expert validation for quality metrics to reduce individual bias.
- Develop AI modules to assess more quality metrics, like scale inclusion and multiple angles.

2. Enhance user experience and accessibility:

- Simplify real-time feedback with visual cues and icons to reduce cognitive load.
- Integrate image capture and metadata collection into a seamless workflow.
- Test and optimize the app for older or lower-powered devices to include more participants.

3. Strengthen privacy and transparency:

- Provide clear, jargon-free explanations of how location and image data are used, stored, and shared.

Hateful and criminal actors

What can bad actors do with your technology?

In which way can the technology be used to break the law or avoid the consequences of breaking the law?

This technology will be used within an AI-supported citizen science app for improving fossil documentation, user engagement, and data quality. Although designed for scientific and educational purposes, its use of mobile cameras, timestamps, and location results in risks involving privacy, data handling, and information integrity.
Because the app collects images and location data, users could accidentally or intentionally capture people, private property, or sensitive locations, creating privacy concerns. Its offline functionality could also allow someone to delay reporting the location of protected finds, potentially avoiding regulations. The system must follow GDPR rules, and any failure in consent handling, storage security, or data protection could lead to privacy violations.
There is also a risk of scientific deception. Users could exploit the AI guidance to create high-quality images of non-fossil objects or fabricate submissions, misleading experts and corrupting the dataset. They could also enter false metadata, such as changing discovery dates or locations, which would undermine scientific accuracy.

Can fakers, thieves or scammers abuse the technology?

Yes, even a harmless citizenscience tool can be misused in unexpected ways. Some potential risks include:

1. Fake or malicious submissions
People could intentionally upload fake fossils, misleading photos, or irrelevant images to disrupt the platform or skew scientific data.

2. Location misuse
If users share GPS coordinates, someone could use the platform to track fossilhunting spots or private locations, leading to theft of valuable fossils or trespassing.

3. Harassment or community misuse
If the final platform includes community features (comments or discussions), users could misuse them to insult others, spread misinformation, or discourage beginners.

4. Data poisoning attacks
Bad actors could upload large amounts of low-quality or intentionally misleading data to degrade the AI model on purpose.

5. Impersonation or fraud

Users could pretend to be experts or create fake profiles to influence others or gain trust within the community.

Most of these risks are manageable with moderation tools, clear reporting mechanisms, validation checks, and strong privacy protections.

Can the technology be used against certain (ethnic) groups or (social) classes?

The prototype is not designed to identify people, so the risks are smaller, but some indirect harms are still possible.

1. Unequal access based on device quality
Because the prototype currently works best on modern smartphones, people with older or lowend devices (often linked to lower-income groups) may be excluded from participating fully.

2. Location-based disadvantages
If GPS data is required, people who cannot or prefer not to share location (for cultural, safety, or socioeconomic reasons) might be limited in how they can contribute.

3. Language and digital literacy barriers
If the interface is mainly in one language or requires high digital literacy, it could indirectly exclude non-native speakers or people with less technological experience.

4. Community misuse
If users can comment or interact in the final product, bias or discrimination could emerge through harmful messages, gatekeeping, or dismissive behavior toward certain groups.

These issues can be prevented by inclusive design choices: device-friendly performance, multilingual support, optional metadata, strong moderation, and clear community guidelines.

In which way can bad actors use this technology to pit certain groups against each other? These groups can be, but are not constrained to, ethnic, social, political or religious groups.

Because this technology focuses on fossil documentation rather than people, the risk is relatively low, but some forms of indirect misuse are still possible:

1. Using the platform to spread misinformation or polarizing narratives
Bad actors could use community features in the final app (comments, shared observations, messaging) to push conspiracies, politically charged narratives, or targeted misinformation disguised as scientific discussion. This could create division between groups with different beliefs or backgrounds.

2. Excluding or ridiculing certain groups within the community
Users could gatekeep or mock contributions from beginners, young participants, non-native speakers, or people from certain regions. Creating an in-group vs. out-group dynamic that causes division.

3. Manipulating scientific data to spread agenda-driven claims
Someone could intentionally upload misleading or fake fossil observations to push ideological narratives (e.g., denying established science or promoting pseudoscience). If targeted at certain educational or cultural groups, this could widen existing divides.

4. Device- or accessbased inequality
If modern devices are needed for best use, bad actors could frame this as a science for the privileged situation, using it to deepen resentment between socioeconomic groups.

These risks can be reduced with strong moderation, transparent data policies, inclusive design, and clear scientific communication to prevent the platform from becoming a tool for division.

How could bad actors use this technology to subvert or attack the truth?

Bad actors could misuse this technology to attack the truth by fabricating evidence, mass-producing misinformation, presenting misleading visuals, poisoning training data, claiming false authority, hiding or distorting results, or overwhelming information channels. Even objective AI outputs could be used to manipulate perception, create false narratives, and undermine trust in reality.

Now that you have thought hard about how bad actors can impact this technology, what improvements would you like to make? List them below.

To reduce risks from bad actors, I would add verification for submissions to prevent fake or misleading data. I would limit the exposure of sensitive information. I would make AI guidance and outputs transparent. Community interactions should be monitored and moderated. The system should be designed to make it hard to manipulate results or misrepresent findings. Clear instructions and user education would help. Multi-expert validation and audit logs would further protect the technology from abuse.

Privacy

Are you considering the privacy & personal data of the users of your technology?

This category is only partial filled.

Does the technology register personal data? If yes, what personal data?

This prototype also processes personal data, even though it is not a full application. Its purpose is to test feedback mechanisms for improving image quality, but doing so still needs the handling of information that falls under GDPR. The prototype collects fossil images, a title for each submission, and the location of the find. If the submission needs to be linked to a survey, it will likely also collect a user name so responses can be matched correctly.
Even in this limited form, the combination of images, names, and location data counts as personal data. An uploaded photo could also contain unintended details, and location details and timestamps can show where a user was at a given moment. Together, these elements form identifiable information connected to a specific person, even if the system is only meant for testing.
Because of this, the prototype still needs careful handling of data. This means secure storage, clear consent, minimal retention, and avoidance of unnecessary identifiers. Anonymization, where possible, and transparent communication with users help ensure that the prototype remains GDPR-compliant while gathering the data needed to study and improve real-time image feedback.

Do you think the technology invades the privacy of the stakeholders? If yes, in what way?

This question has not been answered yet.

Is the technology is compliant with prevailing privacy and data protection law? Can you indicate why?

This question has not been answered yet.

Does the technology mitigate privacy and data protection risks/ concerns (privacy by design)? Please indicate how.

This question has not been answered yet.

In which way can you imagine a future impact of the collection of personal data?

This question has not been answered yet.

Now that you have thought hard about privacy and data protection, what improvements would you like to make? List them below.

This question has not been answered yet.

Human values

How does the technology affect your human values?

This category is only partial filled.

How is the identity of the (intended) users affected by the technology?

In the prototype version of this technology, the focus is on testing AI-supported feedback for taking better fossil images, but even in this limited form, it still affects how users see themselves and how they participate in the process. Instead of simply uploading photos, users learn how to capture clearer, more useful images through real-time guidance. This helps fossil enthusiasts feel more capable and confident, turning them from casual contributors into users who understand what good scientific documentation looks like. Even without the full educational or community features of the future app, the prototype already gives people a sense of producing higher-quality, more meaningful submissions.

How does the technology influence the users' autonomy?

This question has not been answered yet.

What is the effect of the technology on the health and/or well-being of users?

This question has not been answered yet.

Now that you have thought hard about the impact of your technology on human values, what improvements would you like to make to the technology? List them below.

This question has not been answered yet.

Stakeholders

Have you considered all stakeholders?

Who are the main users/targetgroups/stakeholders for this technology? Think about the intended context by answering these questions.

Name of the stakeholder
Fossil hunter / Citizen scientist

How is this stakeholder affected?
The fossil hunter's experience is changed from static and discouraging to interactive and educational. The technology provides a user-friendly data collection process that streamlines the workflow. They receive real-time AI feedback (e.g., "move closer," "adjust lighting," or "include a scale") during image capture, which helps them collect research-grade data. This directly addresses their previously reported limitation of the absence of real-time image guidance.
I consulted an already done user analysis and user survey from the last semester.

Did you consult the stakeholder?
Yes

Are you going to take this stakeholder into account?
Yes

Name of the stakeholder
Paleontologists / fossil validators

How is this stakeholder affected?
The validators' workload is made more efficient and focused on high-value tasks. They currently spend time manually classifying submissions, a process that is slowed by inconsistent data quality. The new technology improves image and metadata quality before submission, which means they receive a higher proportion of scientifically useful records. The technology supports them in verifying and correcting AI or user-submitted identifications. They also directly influence the system, as paleontologist Isaak Eijkelboom helps define the standards for fossil documentation. I consulted experts in online meetings and in person at Leiden Biodiversity Center.

Did you consult the stakeholder?
Yes

Are you going to take this stakeholder into account?

Yes

Name of the stakeholder
AI developer at Naturalis

How is this stakeholder affected?
The AI developer will receive higher-quality, more consistent training data, improving the model performance. The current inconsistent image quality leads to considerably lower classification performance in existing AI models (Top-1 accuracy of 47% to 54% on citizen data vs. 88% to 91% on standardized data).

Did you consult the stakeholder?
Yes

Are you going to take this stakeholder into account?
Yes

Name of the stakeholder
Fontys side product owner

How is this stakeholder affected?
This stakeholder is responsible for coordinating the project and ensuring its strategic alignment. The Product Owner defines the research goals and expectations, and their primary role is to ensure that the project aligns with the larger LegaSea initiative and provides access to resources. The project deliverables, including the research-supporting prototype and the research paper, contribute directly to their overall project goals. I have weekly meetings with the PO and paleontologists and the AI developer from Naturalis.

Did you consult the stakeholder?
Yes

Are you going to take this stakeholder into account?
Yes

Did you consider all stakeholders, even the ones that might not be a user or target group, but still might be of interest?

Name of the stakeholder
Educator

How is this stakeholder affected?
Educators can use the technology as a teaching tool, showing students how

citizen science and real-time data collection work. It provides an example of AI-assisted research and encourages engagement with scientific practices. They benefit from clear instructions, educational materials, and structured workflows that they can integrate into lessons or workshops.

Did you consult the stakeholder?
No

Are you going to take this stakeholder into account?
No

Name of the stakeholder
Casual user

How is this stakeholder affected?
Casual users, such as hobbyists or beachgoers, could interact with the app without deep scientific knowledge. The technology could feel overwhelming if the guidance is too technical/scientific. When simplified, it can make participation easy, educational, and enjoyable, allowing them to contribute without frustration.

Did you consult the stakeholder?
No

Are you going to take this stakeholder into account?
No

Now that you have thought hard about all stakeholders, what improvements would you like to make? List them below.
I would improve the technology by making it more inclusive and accessible to all users, including those with older devices or limited technical experience. For fossil hunters, I would refine the real-time feedback to be clearer and less distracting, using icons or simple visual cues. For paleontologists, I would add multi-expert validation and quality checks to reduce bias and ensure high-value tasks are prioritized. For AI developers, I would provide standardized, high-quality datasets with complete metadata to improve model performance. For the Product Owner, I would ensure clear reporting so they can track progress and outcomes efficiently.
I would also consider less obvious stakeholders, such as educators or casual users, by providing tutorials, optional guidance modes, and transparency about data use.

Data

Is data in your technology properly used?

This category is only partial filled.

Are you familiar with the fundamental shortcomings and pitfalls of data and do you take this sufficiently into account in the technology?

In the prototype, I address the weaknesses of citizen science fossil data by testing real-time AI feedback to reduce poor and inconsistent image quality. LegaSea studies show that user-submitted photos are often blurry or poorly lit, and AI models trained on them perform worse than on controlled images. The prototype provides instant feedback during image capture, helping users avoid common mistakes before making a submission.
I also recognize that human judgment is subjective. Initial quality labels came from a single expert, and simple image metrics do not fully capture what experts consider research-grade. While the prototype does not yet include full metadata or multimodal features, it is designed to explore combining image feedback with contextual information to come close to expert reasoning.

How does the technology organize continuous improvement when it comes to the use of data?

This question has not been answered yet.

How will the technology keep the insights that it identifies with data sustainable over time?

This question has not been answered yet.

In what way do you consider the fact that data is collected from the users?

This question has not been answered yet.

Now that you have thought hard about the impact of data on this technology, what improvements would you like to make? List them below.

This question has not been answered yet.

Inclusivity

Is your technology fair for everyone?

This category is only partial filled.

Will everyone have access to the technology?

This question has not been answered yet.

Does this technology have a built-in bias?

In my prototype, while I aim to reduce the bias of low-quality citizen data, the system itself introduces several built-in biases related to reference data, technical requirements, and prescriptive guidance.

The most obvious bias comes from the quality benchmark. Real-time feedback is based on 100 fossil images graded by a single expert, creating a gold-standard bias toward one persons judgment. The feedback also enforces a museum-like aesthetic: white backgrounds, consistent lighting, and a scale. This can be difficult to achieve in the field. The computational metrics I use, like mean intensity or Laplacian variance, only moderately match expert judgment and do not fully capture aspects like fossil type, size, or preservation.

Other structural biases exist. It requires a modern smartphone, excluding some users, and supports only English, limiting participation by speakers of other languages.

The design is also prescriptive. By giving constant real-time guidance, the system assumes users want detailed instructions on producing research-grade data. For casual users, this could feel intrusive or slow down the submission process.

I am aware of these limitations and plan future mitigations, such as randomization in prototype testing. In essence, the prototype acts like an algorithmic museum curator: it improves data quality but biases the dataset toward contributions that fit its specific standards.

Does this technology make automatic decisions and how do you account for them?

This question has not been answered yet.

Is everyone benefitting from the technology or only a a small group? Do you see this as a problem? Why/why not?

This question has not been answered yet.

Does the team that creates the technology represent the diversity of our society?

This question has not been answered yet.

Now that you have thought hard about the inclusivity of the technology, what improvements would you like to make? List them below.
This question has not been answered yet.

Transparency

Are you transparent about how your technology works?

This category is only partial filled.

Is it explained to the users/stakeholders how the technology works and how the business model works?

In my project, transparency about how the prototype works and its goals is a very important aspect. The technology is designed as an academic, non-commercial initiative, so no business model is involved.
The goals are clearly communicated: the prototype supports the LegaSea project by helping reconstruct Dutch Ice Age biomes. Its immediate purpose is to improve the quality and quantity of fossil images and metadata, while also educating fossil enthusiasts on collecting research-grade data. The prototype replaces the slow, static workflow of the existing platform with AI-assisted real-time guidance for data collection.
Users can understand how the system works through its live feedback, which gives clear instructions such as move closer, adjust lighting, or include a scale. Guidance is linked to quantifiable metrics from expert grading, so users can see exactly why the system provides certain instructions.

If the technology makes an (algorithmic) decision, is it explained to the users/stakeholders how the decision was reached?

This question has not been answered yet.

Is it possible to file a complaint or ask questions/get answers about this technology?

This question has not been answered yet.

Is the technology (company) clear about possible negative consequences or shortcomings of the technology?

This question has not been answered yet.

Now that you have thought hard about the transparency of this technology, what improvements would you like to make? List them below.

This question has not been answered yet.

Sustainability

Is your technology environmentally sustainable?

In what way is the direct and indirect energy use of this technology taken into account?

In my prototype, energy use is addressed through resource-efficient AI, focusing on minimizing power consumption on user devices. The AI components are optimized to run on low-power smartphones using techniques like quantization and on-device inference. By performing real-time processing locally, the system eliminates network latency and does not use another cloud server to run the backend.
Currently, the web app requires an internet connection both to access the app and to submit data. Offline capability is planned for the final app, allowing users to capture and store images in the field without internet, which will further reduce energy spent on data transfer.

Do you think alternative materials could have been considered in the technology?

No, this is a software product.

Do you think the lifespan of the technology is realistic?

Its a software product. The prototypes will only exist and be used until their purpose is fulfilled. About one semester.

What is the hidden impact of the technology in the whole chain?

?? Not applicable

Now that you have thought hard about the sustainability of this technology, what improvements would you like to make? List them below.

Backend server costs have already been eliminated
The final application will run offline on the user's device, so no additional server cost.
AI will be designed to be resource-efficient, so as not to drain the user's phone battery.

Future

Did you consider future impact?

This category is only partial filled.

What could possibly happen with this technology in the future?

In my project, the technology is designed to help citizen science data collection and has the potential for a larger impact. If scaled to a million users, it could change data standards, user habits, and the interaction between amateur and professional scientists.
With real-time AI guidance, users would learn to consistently produce research-grade data, making high-quality, standardized documentation the norm. This would greatly improve AI model accuracy, which currently suffers when trained on inconsistent citizen data. The large volume of contributions would allow not only regional but potentially global ecological and biome reconstruction, far beyond what scientists could achieve alone.
Expert roles would also shift. AI could handle routine validation of common finds, letting experts focus on rare or complex specimens.

Sketch a or some future scenario (s) (20-50 years up front) regarding the technology with the help of storytelling. Start with at least one utopian scenario.

This question has not been answered yet.

Sketch a or some future scenario (s) (20-50 years up front) regarding the technology with the help of storytelling. Start with at least one dystopian scenario.

This question has not been answered yet.

Would you like to live in one of this scenario's? Why? Why not?

This question has not been answered yet.

What happens if the technology (which you have thought of as ethically well-considered) is bought or taken over by another party?

This question has not been answered yet.

Impact Improvement: Now that you have thought hard about the future impact of the technology, what improvements would you like to make? List them below.

This question has not been answered yet.