Marine and Coastal

National Environmental Science Program

RESEARCH REPORT
Project 1.29a

Great Reef Census - a case study to integrate citizen science data into research output for marine habitat management

Dr Christopher L. Lawson, Prof Peter Mumby, Prof Chris Roelfsema, Dr Katie Chartrand, Aruna Kolluru, Heping Shi, Aloysius Tan, Zhi Yuan Toh, Chandramouleeswaran Ganesan, Navindra Viraj Kothalawal, Anand Iyengar, Nicole Senn, Ben Vozzo, Jasmina Uusitalo, Andy Ridley

Catalogue data

This publication (and any material sourced from it) should be attributed as:

Lawson C. L., Mumby P., Roelfsema C., Chartrand K., Kolluru A., Shi H., Tan A., Yuan Toh Z., Ganesan C., Navindra Kothalawal V., Iyengar A., Senn N., Vozzo B., Uusitalo J. and A. Ridley, 2023. Great Reef Census - a case study to integrate citizen science data into research output for marine habitat management. Report to the Reef and Rainforest Research Centre, Cairns, Queensland.

ISBN

978-1-922640-16-1

This publication is available on the NESP Marine and Coastal Hub website:

www.nespmarinecoastal.edu.au

Acknowledgment

This project is supported with funding from the Australian Government under the National Environmental Science Program. The Marine and Coastal Hub is funded by the Australian Government under the National Environmental Science Program.

Cover images

Right: Field collection of a snorkeller performing a survey of the reef for the Great Reef Census. Credit: Nicole McLachlan for Citizens of the Great Barrier Reef. Back: A volunteer analysing photos on the Great Reef Census online analysis platform. The platform uses a hybrid approach of AI and human eye to quickly analyse photos for major coral categories and coral cover. Credit: Nicole McLachlan for Citizens of the Great Barrier Reef.

Acknowledgement

The Marine and Coastal Hub acknowledges Aboriginal and Torres Strait Islander people as the first peoples and Traditional Owners and custodians of the land and waterways on which we live and work. We honour and pay our respects to Elders past, present and emerging.

Aboriginal and Torres Strait Islander peoples represent the world's oldest living culture. We celebrate and respect this continuing culture and strive to empower Aboriginal and Torres Strait Islander peoples.

Contents

List of figures

List of tables

Executive summary

To maximise our understanding of our marine and coastal environment, we need to take advantage of emerging technologies and approaches. This includes citizen science and community monitoring. Technology has greatly reduced the gap between mainstream science and citizen science to the point they may become almost identical in some integrated programs, especially when involving the collection of in-field information. The challenge for science is to integrate with the vast opportunities afforded by this congruence.

This project had two primary objectives. First, we compared and combined expert and citizen scientist analyses of geo-referenced images from a large-scale citizen science program, the Great Reef Census (GRC). Additionally, we examined the effectiveness of an updated online analysis platform (beta version) that incorporated both machine learning (artificial intelligence; AI) and citizen scientist validation steps for analysing the GRC Year 2 image collection.

To accomplish these goals, we 1) established a validation framework for GRC Year 1 & 2 analyses, using University of Queensland expert analysis as a reference for reliable data, to filter the citizen dataset, which may contain both accurate and erroneous results, 2) extracted insights to streamline image processing and training within the online platform, 3) evaluated the AI system's performance in identifying key coral groups in images from GRC Years 1 & 2, and 4) assessed a refined online platform, including AI integration, with a subset of citizen science users against a calibration expert pool's analyses from the GRC Year 2 image library to enhance data quality.

Here, we analyse the performance of a machine learning platform and a citizen science program deployed in schools across Australia in analysing the coral cover and coral type contained in images collected over two years by the Great Reef Census. The categories analysed were reef structure coverage (% of image), total coral cover (% of reef structure), branching coral cover (% of reef structure), table coral cover (% of reef structure), massive coral cover (% of reef structure), and other coral cover (% of reef structure).

The AI model was developed over three key versions. First, a basic Dell Technologies server hosted a rudimentary platform where AI-generated polygons required user labelling into categories. Despite being slow, buggy, and lacking design and training elements, it served as a starting point, tested by 200 corporate volunteers. Second, the software evolved into a more refined version, integrated into the Great Reef Census website and tested with around 100 corporate volunteers. Feedback revealed lingering issues, including overlapping polygons and small, hard-to-identify ones. Finally, the third platform iteration addressed feedback by eliminating small polygons for a smoother user experience.

The Great Reef Census School Program allowed us to compare the third AI platform iteration with the original non-AI analysis software used in 2021. The non-AI method involved 6,000 people analysing 30,000 images over 12 weeks, while the new AI method processed almost as many images (24,000) in just six weeks, with only 5% of the participants (300).

Image analysis was conducted by “experts”: individuals compensated for their time and who had experience in marine science and coral identification (e.g., holding a Bachelor’s degree in Marine Biology). These experts, affiliated with the University of Queensland, received training from the GRC science team to identify major coral groups of interest and calculate percent cover of each. By comparing the results of citizen science and AI analyses with those of the experts, we validated these methods as dependable tools for future reef health surveys.

The AI estimates were accurate and precise across almost all categories, in most cases reaching ± 5% accuracy while using just 20 images. However, the AI performed less accurately in general for images/reefs that had high (>70%) or low (< 20%) coral coverage; here, it may be beneficial to collect more images per location and develop improved analysis methodology, such as further training of the AI.

In general, the school program showed results that were reliable enough to inform research and management (i.e. within 10% of trained experts), however, it exhibited lower accuracy and was more variable compared to AI analysis. Further investigation may enable the school program system to fill in gaps where the AI performs poorly, e.g. images with very high coral cover. Combining the school program analysis with AI as a filtering mechanism enhanced AI accuracy when disparities between the two were observed, with the filter consistently working across diverse reef conditions. The mean accuracy of the AI performance was improved by up to ~3% if images with discrepancies (10% or more) between the school program and AI were removed. Consequently, citizen science may complement and improve AI analysis. However, the citizen science analysis may offer greater value for specific tasks, such as identifying coral health indicators that the AI isn't trained to detect.

The comparative approach of different techniques used in this project allowed us to assess their respective strengths, weaknesses, and suitability for diverse scenarios. The analysis tool and spatially explicit data validated from this project will be available to Commonwealth and regional management agencies as well as on-ground researchers, Traditional Owners and rangers to guide environmental decision-making and on-ground action.

1. Introduction

To maximise our understanding of our marine and coastal environment, we need to take advantage of emerging technologies and approaches. This includes citizen science and community monitoring. Technology has reduced the gap between mainstream science and citizen science, especially when involving the collection of in-field information (Green et al. 2020, Vohland et al. 2021). The challenge for science is to integrate with the vast opportunities afforded by this congruence.

The collaboration between technology and citizen science can enhance our capacity to address biodiversity loss and ecosystem challenges, issues that continue to escalate in the 'Anthropocene' (Johnson et al. 2017). The partnership of technology and community action can streamline data collection and analysis, promoting accurate and timely insights into ecological changes that can be used for evidence-based ecosystem management (Hauser et al. 2019) and achieving global targets for protected areas (Mammola et al. 2020). Meeting conservation goals hinges on the extensive monitoring programs that underpin informed and cooperative conservation efforts.

Given the sheer number of reefs on the Great Barrier Reef, less than 1% are regularly monitored by scientists. Monitoring of the Great Barrier Reef involves the Great Barrier Reef Marine Park Authority and other various programs and groups, such as the Australian Institute of Marine Science Long-term Monitoring Program, which surveys the same reefs regularly to track coral cover change over time. Various methods are used for collecting environmental and benthic field data across the Great Barrier Reef, with trade-offs in data resolution and spatial coverage (Roelfsema et al. 2020). While high-quality field data for every reef would be ideal, resource constraints make it unattainable. This leads to spatial and temporal data gaps, potentially affecting management. Despite the success of existing monitoring, more widespread data collection would help inform better decision-making. The intersection of technology and citizen science presents as a potentially useful tool here.

Citizen science engages the public in various aspects of scientific research and management. Over the past decade, there has been a global increase in citizen science projects, driven by growing public awareness and conservation concerns (McKinley et al. 2017). These initiatives address data gaps in professional science, yielding valuable insights for diverse species and ecosystems (McKinley et al. 2017). Technological advancements, such as the widespread proliferation of smartphones and underwater cameras, have fuelled the expansion of citizen science projects (e.g. Ruiz-Gutierrez et al. 2021). These projects can effectively provide extensive spatial data on species distributions, abundances, traits, or ecosystem functions (Chandler et al. 2017). Many of these initiatives have contributed reliable data, enhanced knowledge and reduced monitoring costs (McKinley et al. 2017).

In 2020, a partnership of scientists, managers, and a citizen science organisation created The Great Reef Census (GRC): an innovative solution to scale collaborative conservation and research. Since then, GRC has activated a network of more than 100 vessels and operators, and hundreds of local volunteers working on or visiting the GBR to gather reconnaissance data on coral reef state across the 3000-reef 344,000km^2 system. GRC is an established citizen science innovation project, designed to pilot new ways of capturing reconnaissance citizen science data. By using citizen scientists to both collect and analyse reef images, as well as a team of professional scientists to ensure program rigour, the project is an innovative approach to assessing Great Barrier Reef health that complements and enhances existing monitoring programs.

A core deliverable of citizen science is to provide robust data at scales not achievable through classical funding streams and science programs while complementing ongoing monitoring and existing datasets. However, there is a risk that if trust in the analysis process and the integrity of the data are not established, the scientific community will not accept its significant value as a scalable solution to fill gaps in knowledge more broadly. Some citizen science programs contain higher levels of bias and human error than professional data collection (Moedas 2018). This can be seen with survey inconsistencies leading to temporal bias, as well as spatial bias where data is concentrated in more urban areas due to the larger number of people surveying (Millar et al. 2019). This can skew the results and provide inaccurate conclusions. There is also the potential for bias in the data collection, as people tend to collect data of things that are more appealing (e.g. charismatic megafauna). However, a well-designed program can reduce bias. For example, image collection in the GRC is directed according to prioritisation of reefs determined by the science advisory team based on management needs and further was shown to not be biased towards ‘feature’ corals (Taison Ka Tai Chang, unpublished thesis). If developed appropriately, such programs may provide data collection that is useful for applied management and research (Hughes et al. 2022).

The GRC field collection of images by citizen scientists amplified the spatial dataset used by researchers. However, expert analysis for an expanding program like GRC incurs significant time and costs. To achieve scale, GRC must rely on citizens as a virtual volunteer workforce to analyse the vastly growing image library. Whilst the expert analysis component of the previous platform has been effective, the citizen scientist component is an area that needs improvement. Here, we refine an analysis platform to ensure the “ask” is appropriately matched to the ability of a citizen scientist to retain data validity for end users.

Specifically, this project will test a simplified analysis platform that empowers citizen scientists to provide effective information from the images captured in GRC field campaigns and to build scalable capacity with a virtual volunteer workforce. Changes to the platform tool include machine learning capacity in order to remove more difficult image processing steps while amplifying the value of the individual to detect structures not achievable through machine learning (artificial intelligence; AI) alone.

The analysis tool and spatially explicit data validated from this project will be available to Commonwealth and regional management agencies as well as on-ground researchers, Traditional Owners and rangers to guide environmental decision-making and on-ground action.

Our study will:

- Explore and validate expert versus Artificial Intelligence analyses of geo-referenced images collected during the Great Reef Census Years 1 & 2 field campaign, and
- Explore and test the output of a refined online analysis platform (beta version) that integrates machine learning and citizen scientist validation steps to analyse the GRC Year 2 image library.

To achieve this, we will:

- Identify lessons learnt to simplify image processing and training through the online platform.
- Test the performance of the AI system alone in identifying key coral groups in images collected during the GRC Years 1 & 2.
- Test a refined online platform (with integrated AI) with a subset of citizen science users against a calibration expert pool of analyses from the GRC Year 2 image library to ensure data quality refinement.

2. Methods

2.1. Collection of images

The first task is for the Reef Census surveyor; their goal is to capture the most representative images of a particular reef slope or bommie and upload those to the Census Page directly, noting time and GPS position.

We are specifically interested in the slopes of reefs and larger bommies, rather than the shallow reef tops, because these are difficult to photograph with a wide angle (you cannot get sufficiently far away from the reef). Typically, the depth of the water will range from around 3 m to 20 m. The aim for the Reef Census surveyors is to look at the area of reef below them and imagine that they are describing a representative area of the reef, where each photograph might cover something like 5 m × 5 m of seabed or more. Typically, it will be necessary to take a minimum of 20 photographs to represent a section of reef. Each photograph should be approximately parallel to the reef orientation so that it serves as a 'vertical' shot looking directly down. In very clear water, the surveyor will likely be able to see 20 m on either side and will not need to swim far.

Photographs by snorkelers were mostly taken by duck-diving down as far as is comfortable, holding the camera steady to take one photograph per dive and holding the camera parallel to the reef. Divers generally kept a few metres away from the reef to obtain a good area of coverage per photograph (5 m or so within one shot). Images collected directly from the boat were obtained directly over a reef towards the seabed, ensuring it was not too deep and that the seabed was clearly visible in the image.

Once surveyors have finished their 20-photo survey of a stretch of reef, it was then recommended that they include two more zones on the same reef as far apart as is safe or logistically feasible. The minimum distance between sites was 200 m.

2.2. Expert analysis of images

The image analysis was carried out by experts who were compensated for their time and possessed expertise in marine science and coral identification, including qualifications such as a bachelor's degree in marine biology. These experts were affiliated with the University of Queensland and received training from the GRC science team to identify the key coral groups of interest. To perform the analysis, the experts outlined polygons around the primary coral types of interest within the images. Following that, the area of the polygons was summed by the software to calculate the percentage coverage of each variable in each image, rounding it to the nearest 1%. This analysis was conducted using an earlier 'expert version' of our GRC analysis platform, which had the necessary functionality for this task.

By comparing the outcomes of citizen science analysis and AI analysis with the results obtained by the experts, we can establish the reliability of these methods for future reef health surveys and integration with other conservation tools, such as ecosystem models (Bozec & Mumby 2020). Furthermore, this comparison allows us to gain insights into the strengths, weaknesses, and suitability of each approach for various scenarios. The following definitions are utilised to identify the key variables of interest (Mumby & Harborne 1999).

Table corals

Live table corals are typically circular and are elevated above the reef as they have a short ‘trunk’ at their centre. They can also hang off the side of a steep section of reef. Either way, the area beneath the table is often starved of sunlight making it difficult for competitors to grow. Table corals can grow fast and often exceed 1 m in diameter. The fine-grained lattice of living specimens is often clear and the edges of the branches may be paler in colour (although not always). Reefs can be dominated by this coral and recover rapidly quickly.

Dead table corals clearly have the same overall shape, but their colour tends to be drab and grey without any clear patterns in coral colouration (e.g., no pale edges). Dead specimens will often have dark patches of algal colonisation or even other corals growing on them in places.

Branching corals

Live branching corals can have a wide variety of shapes, ranging from fine branches (separated by less than 1 cm) to coarse branching, separated by tens of centimetres. The fine knobby structure of many corals can be seen on the surface and they may possess uniform patterns such as a consistent paling at the tips. Branching corals provide habitat for many fish.

Dead branching corals tend to be dull grey, have relatively uniform surface texture (i.e., without obvious small knobs), and are more likely to form damaged piles surrounded by rubble.

Massive corals

Live massive corals are dome-shaped and the most prominent are the large mushroom-shaped corals with a diameter ranging from 50 cm to several metres. They tend to have a uniform appearance and the colour is typically green or pale brown. Brain corals are also massive in structure and have a characteristic brain-like surface. In some cases, massive corals can take on bright colours, including green and magenta.

Dead massive corals make up much of the reef structure and like most dead corals have a dull grey appearance. Their surface tends to be less uniform in appearance because there will be patches colonised by other animals or plants including algae, other corals, sponges or soft corals.

Other corals

Other live corals encompass any live coral that does not fit one of the shapes we discriminate (i.e., table, branching or massive). They may have an encrusting, leaf-like, or flower-like shape, but their surface will tend to be of uniform appearance because the coral is living. Any pattern will tend to be consistent, from one surface to the next.

Other dead corals are reserved for corals whose shape is neither table-like, branching, or massive and where the coral is dead – i.e., dull grey. Note that any dead coral will also tend to have individual patches of other animals and plants.

Reef structure

This is reef without coral, or coral that is clearly dead or other reef organisms. It can include macroalgae (seaweed) in a variety of forms, but the most commonly seen on reefs are either (1) green plates/disks (Halimeda), (2) large leafy brown fronds that can be 50 cm to 2 m high (Sargassum), or (3) encrusting or leafy brown patches (Lobophora). Macroalgae grow on dead coral or the underlying reef substrate. Note that most of the underlying reef substrate is featureless limestone that was once coral but has died and been heavily eroded so that the nature of the coral type is no longer visible.

Water, sand, shadow

This is anything that doesn't belong to the reef. It could be simply the visible water in an image, wildlife, sand beneath a reef or a diver or snorkeller.

2.3. AI analysis of images

We used deep learning/computer vision to analyse the coral images. Deep learning is a type of machine learning that involves using algorithms to analyse and learn from data in a hierarchical manner, with the goal of learning from the data to make predictions or take actions (LeCun et al. 2015). It is called “deep” learning because it uses multiple layers of algorithms, each of which learns from the data and passes its findings on to the next layer. This hierarchical approach allows the system to learn complex patterns and relationships in the data, making deep learning a powerful tool for a variety of applications, such as image and speech recognition, natural language processing, and more. In simple terms, deep learning enables a computer to learn from data using many layers of algorithms (Rusk 2016).

Computer vision is the field of computer science that focuses on using algorithms and other techniques to allow computers to understand, analyse, and interpret visual data from the world around them (Stockman & Shapiro 2001). This can include things like analysing images or videos to identify objects, recognise faces, read text, and more. The goal of computer vision is to enable computers to understand and interpret visual data in the same way that humans do, and to be able to use that understanding to make decisions, take actions, and solve problems. Computer vision is essentially a way for computers to “see” and understand the world around them.

To identify the coral types in images, we used semantic segmentation, a type of image analysis that involves assigning labels or categories to different parts of an image (Guo et al. 2018). This allows a computer to understand and interpret the objects and scenes depicted in an image and to differentiate between different objects and their boundaries. For example, in an image of a street scene, semantic segmentation might be used to identify and label the buildings, the cars, the trees, the sidewalk, and other objects in the scene (Zhang et al. 2018). This can be useful for a variety of applications, such as autonomous vehicles, image search, and more.

Segformers is a type of algorithm that can be used for semantic segmentation of corals, which is the process of identifying and labelling the different types of corals in an image (Xie et al. 2021). Segformers were specifically used to handle the challenges of coral segmentation, such as the wide variety of coral shapes, sizes, and colours, as well as the complex backgrounds and lighting conditions often found in underwater images. Segformers use a hierarchical approach to segmentation, which allows them to learn from the data and make accurate predictions about the different types of corals in an image (Xie et al. 2021). As a result, Segformers are an effective tool for identifying and labelling different types of corals in images.

2.4. School program analysis of images

The Great Reef Census: Education and Analysis program (hereafter 'school program') has engaged students across Queensland to participate in citizen science to protect the Great Barrier Reef. During the pilot program, 302 students aged 10 - 17 took part and completed 24,000 image analyses over six weeks. These students were some of the first to use our Artificial Intelligence Deep Learning analysis software developed with Dell Technologies.

Five schools participated in the online analysis, with two schools attending in-person workshops and the other three completing online training with the GRC team prior to starting the analysis. Each class session incorporated both science and technology. Class time was split between:

1. A presentation on the importance of the Great Barrier Reef, its threats and how citizen science can help the Reef.
2. Training on how to identify key coral species and categories (Plate Corals, Branching Corals, Massive Corals, Other Corals, Reef Structure).
3. Image analysis on laptops/iPads using Artificial Intelligence software to label survey images.

In many cases, the program was run with students who were going to the reef as part of their curriculum - these sessions enhanced their reef experience by equipping students with the knowledge to identify key coral species while in the water. During each session, presenters gave detailed introductions including background information on their individual career pathways. Each session included a marine biologist, and allowed students to ask questions and find out more.

2.5. Comparing analysis outputs

2.5.1. AI alone

A total of 11,741 images were analysed by both the UQ experts and the AI program and could be included in the comparison. The expert data were amalgamated into the same six categories that the AI output provided: reef structure coverage (% of image), total coral cover (% of reef structure), branching coral cover (% of reef structure), table coral cover (% of reef structure), massive coral cover (% of reef structure), and other coral cover (% of reef structure). The difference between the expert analysis and the AI analysis of each category in each image was then calculated. All data wrangling and comparisons of data were completed in R Statistical Software (v4.2.1; R Core Team 2022).

The general performance of the AI against the expert analysis was assessed using the mean and standard error (SE) of the difference between AI and experts across all images in all categories (Andrew & Mapstone 1987). Our intention was not to engage in formal hypothesis testing to assess the AI results against those obtained from expert analysis; our primary focus was not on determining whether the AI results stem from a distinct 'population' compared to the expert results on an image-by-image basis. Instead, our goal was to delve into the AI's achievable level of accuracy and assess precision metrics. We aimed to understand how well the AI performs across various image types, including different reef states, and whether the associated uncertainty or variation reached unacceptable levels. The acceptability of this uncertainty or variation is contingent upon the specific management and research implications that arise from utilising the results.

Furthermore, from an operational standpoint, we sought to determine the optimal number of images to capture, which could influence the design of the Great Reef Census, in order to achieve varying levels of accuracy for different image types and relevant variables of interest. In essence, relying solely on the results of hypothesis testing or other statistical inferences may potentially misguide practitioners, managers, researchers, and other end-users. Our approach aims to provide a more comprehensive and context-specific assessment of the AI's performance and its practical implications.

2.5.2. Accuracy of AI

The accuracy of the AI output was assessed by how many images were required to reach a mean difference between AI and expert analysis of 20%, 15, 10% and 5% (95% confidence interval; CI) for each coral category, where ‘5% accuracy’ represents a mean difference of ± 5%.

A simulation approach was used to assess accuracy. To assess total coral coverage, a random image and its associated difference between AI and expert analysis was sampled until the absolute mean difference of sampled images was below the desired accuracy threshold (20, 15, 10 or 5%), with a minimum of 20 images sampled and a maximum of the total number of images available for analysis (11,741). This was repeated 1000 times to obtain a distribution of the number of images required to reach each accuracy threshold, and the mean and 95% CI were calculated. The upper 95% CI was used to conclude the required number of images to reach the target accuracy threshold in each category. This process was repeated for all six coral categories.

To understand the performance and variability of the AI in more detail, the AI vs expert results were also categorised according to reef state, i.e. the percentage of coral cover in the image in 10% bins (Table 1). The same simulation approach, used to analyse the performance of the AI using all images, was used in each coral category and 10% reef state bin.

Table 1*:* The number of images available per 10% reef state bins for the AI analysis.

Reef state bins (% coral cover of image)	Total images available for analysis
0 – 10	1382
11 – 20	1539
21 – 30	1415
31 – 40	1399
41 – 50	1316
51 – 60	1189
61 – 70	1066
71 - 80	835
81 - 90	682
91 - 100	918
All images	11,741

2.5.3. Dominant coral type

The dominant coral type in each image was determined as the coral with the highest percentage cover (branching, table/plating, massive or other). The performance of the AI in this regard was assessed by calculating the percentage of images in which the AI correctly identified the dominant coral type.

The dominant coral type at the reef scale was determined by the coral category that had the greatest number of associated images in which it was identified as the dominant coral. The performance of the AI was then assessed by calculating the percentage of reefs in which the dominant coral type was correctly identified. The same process was followed to assess the AI's performance at the site scale.

2.6. School Program analysis

The school program analysis was similar to the AI performance analysis. However, each image was processed by multiple school students; the values for each of the six coral categories of each image were taken as the mean value from all students who assessed the image. Only images that were assessed by the experts, the AI and the school program were included in the comparison of performance.

The accuracy of the school program in each category was assessed as the mean difference between the expert's value and the school program's value, and the precision was presented as the standard error of these mean values.

Secondly, to determine if the school program and the AI can be used in conjunction, the results of each image were compared between the school program and the AI output. The school program was used as a filter for the AI by removing images with discrepancies in reef substrate between the AI and school program. We then re-analysed the AI vs expert data using only the images that the AI and school program showed concordance on reef substrate.

To achieve this, we examined the difference in reef substrate coverage between AI and school program for each image and created multiple 'filters' to remove images that exceeded the filter threshold. The filter thresholds ranged from including images with 100% difference between AI and school program (i.e. all images) to 10% (i.e. most strict; only images that had less than 10% difference between AI and school program for reef substrate were included), in 10% increments. The assumption here is that the school program can identify individual images in which the AI does a poor job of parsing (based on reef substrate coverage results); reef substrate coverage is the only category in which the school program results were more accurate than the AI.

Note here there are different results (usually in non-consequential amounts) to the previous results for AI accuracy because a different set of images have been used (i.e. only the images that were analysed by experts, AI and school program).

There were an insufficient number of images that were analysed by both the school program and experts to check if the school program was more accurate than the AI for varying levels of coral cover. I.e., we were unable to check if citizen science analysis is more accurate for 10% reef state bins that the AI was unable to achieve 5% accuracy.

3. Results

3.1. Lessons learnt to simplify image processing and training through the online platform

The project resulted in three versions of the online analysis platform that combined AI and human analysis:

- A very basic platform housed on the Dell Technologies server. The software created AI-identified polygons for users to label as one of several categories. The platform was slow, had multiple bugs and featured no design or training elements. This version of the platform was tested with 200 corporate volunteers and their feedback was used to develop phase.
- A designed version of the software, housed on the Great Reef Census website, was tested with approximately 100 corporate volunteers. Feedback highlighted that there were still bugs within the software; for example, polygons overlapped, and many polygons were too small for users to identify and label.
- Based on feedback from corporate volunteers, the third evolution of the platform removed small polygons from the platform to make the user experience smoother.

The school program was an opportunity to compare the third evolution of our AI analysis software to that of the original, non-AI analysis software which was used in 2021. This non-AI method saw 30,000 images analysed by 6,000 people across a 12-week period, whereas the new AI method saw almost as many analyses (24,000) in half the time (6 weeks) by 5% of the people (300).

These results show the fast pace and ease with which the new AI platform can be used by school students. Feedback from the program showed that students were interested in answering multiple-choice questions, e.g. what the most dominant coral is, what is the percentage of coral cover. For the next evolution of the platform, we will be incorporating these multiple-choice questions into the platform to provide further information to scientists and managers on the reef.

3.2. Evaluation of AI

The AI output, in general, was accurate and precise when compared to the image results derived by the experts, achieving 5% accuracy across all categories when all available images that were also analysed by experts were included (Table 2; Figures 1 - 2). Using all available images, the AI slightly overestimated the mean reef cover (calculated as all reef substrate excluding sand, including live coral) and the 'other' coral cover (any coral not classed as table, branching or massive; Table 2). Similarly, the AI output mean of all images was precise, with SE-to-mean ratios of less than 0.06 across all categories except 'other coral (0.21; Andrew & Mapstone 1987).

The AI estimates of coral cover were precise across most categories; in the images available for most categories, the AI output converged on the mean value within 20 images: the minimum number of images collected in a Great Reef Census survey.

Table 2*:* Mean difference between expert and AI output (as a proportion of the reef cover variable, which was measured as all reef substrate (not sand or background water; all images)).

	Reef cover	**Coral cover**	**Table cover**	**Branching cover**	**Massive cover**	**Other coral cover**
Mean	2.53	-3.42	-2.15	-1.13	-1.04	0.89
SE	0.15	0.19	0.06	0.09	0.06	0.19

3.2.1. Coral Cover

Number of images required to reach accuracy thresholds

Our simulation sampling approach demonstrated that the AI output consistently converged on ± 5% accuracy with few images required:

- 20% accuracy (± 20% absolute coral cover): 95% CI = **20 images**
- 15% accuracy (± 15% absolute coral cover): 95% CI = **21 images**
- 10% accuracy (± 10% absolute coral cover): 95% CI = **21 images**
- 5% accuracy (± 5% absolute coral cover): 95% CI = **22 images**

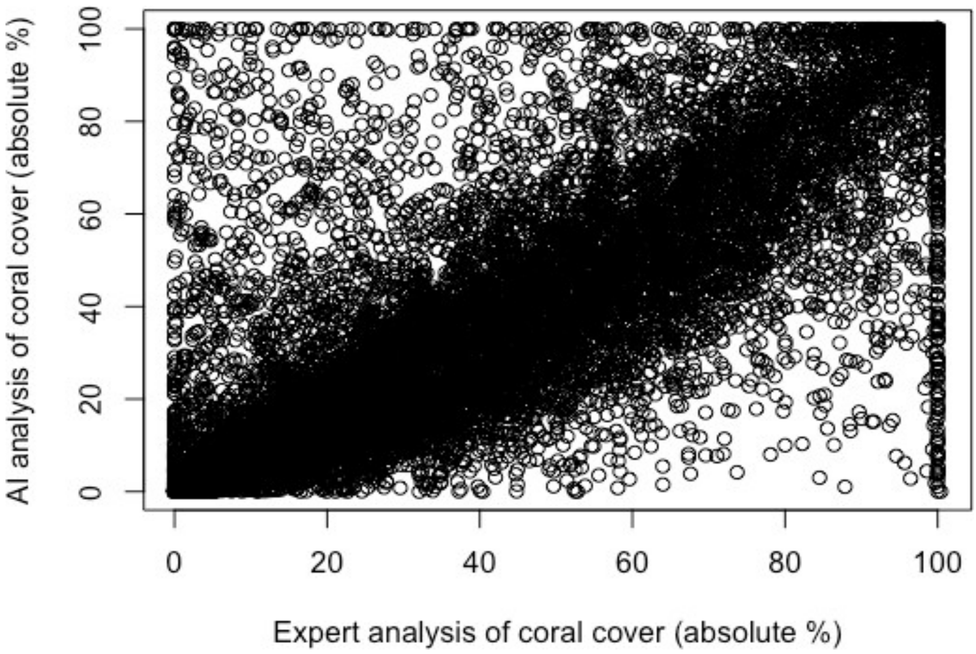

Figure 1*:* Expert vs AI analysis of total coral cover. Each point represents an image (n = 11,741 images).

Accuracy of coral cover as a function of reef state

When images were divided into bins based on reef state (coral coverage), the AI output could not reach a ± 5% accuracy threshold in some cases. Between all of the categories, the highest number of images in any reef state bin that was needed for the 95% confidence interval to reach the target accuracy was 127 (for other coral, followed by 58 for coral cover). This suggests that while in most cases 20 images at any given site or reef will be sufficient to reach the desired optimal accuracy (i.e. as high as possible for the AI), there may be a benefit to collecting up to ~127 images at each site.

The AI output was unable to meet some accuracy thresholds, especially at higher reef states (i.e. higher coral cover; Table 3) after using all available images for analysis.

Table 3*:* Coral cover: the number of images (minimum 20) required for the 95% confidence interval (CI) to meet accuracy thresholds according to the percentage of coral cover in an image (reef state bins of 10%).

Coral cover (% of image)	Mean difference using all images (% cover)	± 20% accuracy: images required (95% CI)	± 15% accuracy: images required (95% CI)	± 10% accuracy: images required (95% CI)	± 5% accuracy: images required (95% CI)
0 – 10	11.5	21	27	NA (>1382)	NA (>1382)
11 – 20	4.1	21	21	22	58
21 – 30	0.0	20	20	21	23
31 – 40	-3.1	20	20	21	36
41 – 50	-5.3	20	21	22	NA (>1316)
51 – 60	-6.8	20	21	25	NA (>1189)
61 – 70	-9.0	21	21	53	NA (>1066)
71 – 80	-10.1	21	22	NA (>835)	NA (>835)
81 – 90	-10.9	21	24	NA (>682)	NA (>682)
91 – 100	-19.2	NA (>918)	NA (>918)	NA (>918)	NA (>918)

3.2.2. Table, Branching, Boulder and Other Coral

For the remaining coral parameters (table, plating, massive and other coral), the AI converged on ± 5% accuracy within 25 images or less (Table 4). In general, the AI was more accurate for these individual coral groups than the total coral cover on average.

Number of images required to reach accuracy thresholds

Across all available images, the AI output consistently converged on ± 5% accuracy with few images required (Table 4; Figure 2).

Table 4*:* Results of the simulations that tested the number of images required for the AI to reach various levels of accuracy (compared to trained experts). The "% accuracy" refers to the ± % absolute coral cover. The number of images are the upper 95% confidence intervals of our simulation for each category.

	20% accuracy	15% accuracy	10% accuracy	5% accuracy
Table Coral	20	20	20	21
Branching Coral	20	20	20	21
Massive Coral	20	20	20	21
Other Coral	20	21	21	25

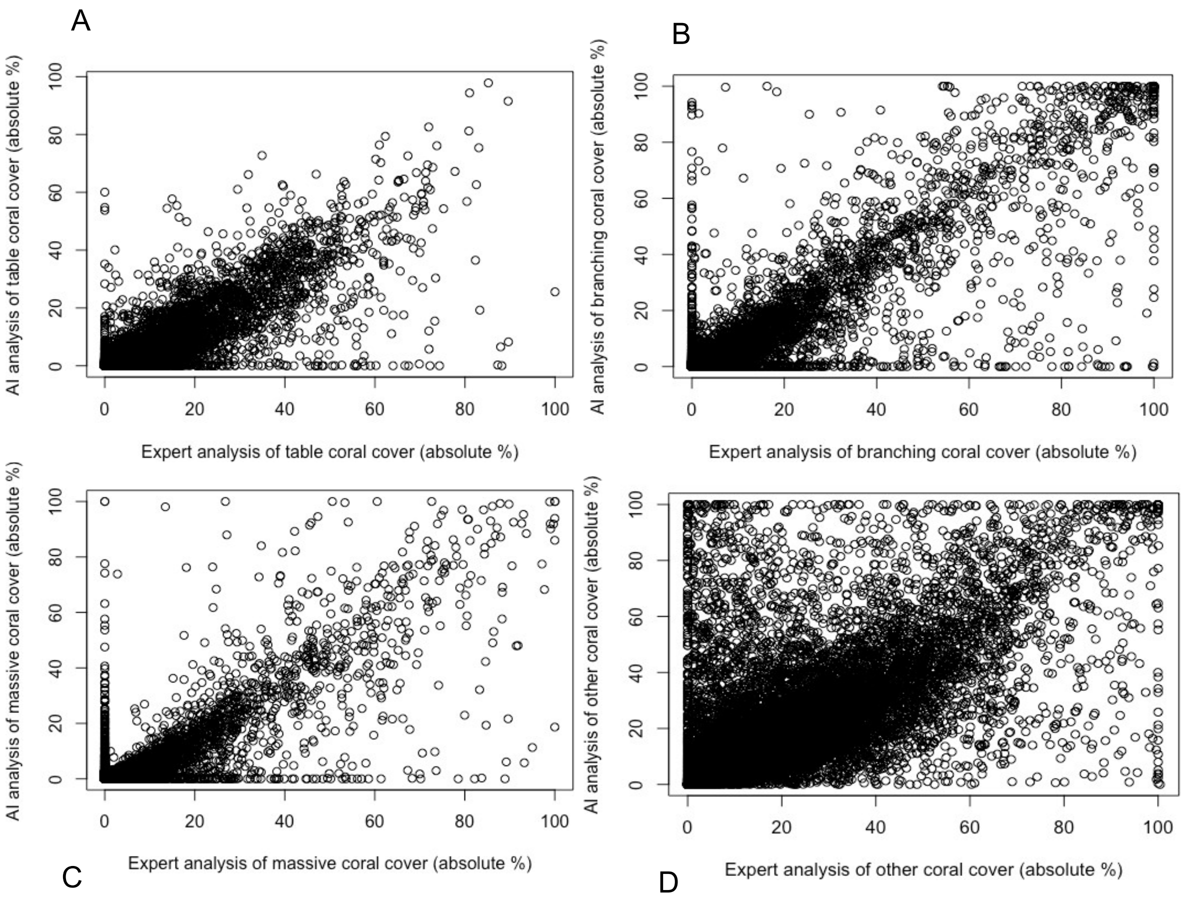

Figure 2*:* Expert vs AI analysis of A) table Acropora coral cover; B) branching Acropora coral cover; C) massive coral cover; D) 'other' coral cover. Each point represents an image (n = 11,741 images).

Accuracy of table, branching, massive and other coral cover as a function of reef state

Our simulation showed similar patterns between total coral cover and individual coral groups when examining the performance of the AI when categorising images according to reef state. Most notably, the AI was able to achieve ± 5% accuracy for images with very low or high reef states when examining individual coral groups rather than total coral cover, or required fewer images to do so (Tables 5 – 8).

Table 5*:* Table coral: the number of images (minimum 20) required for the 95% confidence interval (CI) to meet accuracy thresholds according to the percentage of coral cover in an image (reef state bins of 10%).

Coral cover (% of image)	Mean difference using all images (% cover)	± 20% accuracy: images required (95% CI)	± 15% accuracy: images required (95% CI)	± 10% accuracy: images required (95% CI)	± 5% accuracy: images required (95% CI)
0 – 10	0.02	20	20	20	20
11 – 20	-6.2	20	20	20	NA (>1539)
21 – 30	-1.0	20	20	20	20
31 – 40	-1.5	20	20	20	20
41 – 50	-2.0	20	20	20	21
51 – 60	-2.6	20	20	21	22
61 – 70	-3.6	20	20	20	25
71 – 80	-4.4	20	20	21	44
81 – 90	-4.9	20	20	21	NA (>682)
91 – 100	-4.6	20	21	21	NA (>918)

Table 6*:* Branching coral: the number of images (minimum 20) required for the 95% confidence interval (CI) to meet accuracy thresholds according to the percentage of coral cover in an image (reef state bins of 10%).

Coral cover (% of image)	Mean difference using all images (% cover)	± 20% accuracy: images required (95% CI)	± 15% accuracy: images required (95% CI)	± 10% accuracy: images required (95% CI)	± 5% accuracy: images required (95% CI)
0 – 10	0.9	20	20	21	21
11 – 20	-0.1	20	20	20	20
21 – 30	-0.5	20	20	20	20
31 – 40	-0.6	20	20	20	21
41 – 50	-0.5	20	20	20	21
51 – 60	-0.9	20	20	20	21
61 – 70	-1.3	20	20	20	21
71 – 80	-2.3	20	20	21	23
81 – 90	-3.6	20	21	21	35
91 – 100	-5.9	21	21	24	NA (>918)

Table 7*:* Massive coral: the number of images (minimum 20) required for the 95% confidence interval (CI) to meet accuracy thresholds according to the percentage of coral cover in an image (reef state bins of 10%).

Coral cover (% of image)	Mean difference using all images (% cover)	± 20% accuracy: images required (95% CI)	± 15% accuracy: images required (95% CI)	± 10% accuracy: images required (95% CI)	± 5% accuracy: images required (95% CI)
0 – 10	0.1	20	20	20	20
11 – 20	-0.1	20	20	20	20
21 – 30	-0.3	20	20	20	21
31 – 40	-0.7	20	20	20	20
41 – 50	-1.2	20	20	20	21
51 – 60	-1.4	20	20	20	21
61 – 70	-1.9	20	20	21	21
71 – 80	-2.1	20	20	20	21
81 – 90	-2.0	20	20	20	21
91 – 100	-2.4	20	20	21	23

Table 8*:* Other coral: The number of images (minimum 20) required for the 95% confidence interval (CI) to meet accuracy thresholds according to the percentage of coral cover in an image (reef state bins of 10%).

Reef substrate (% of image)	Mean difference using all images (% cover)	20% accuracy: images required (95% CI)	15% accuracy: images required (95% CI)	10% accuracy: images required (95% CI)	5% accuracy: images required (95% CI)
0 – 10	10.5	21	25	NA (>1382)	NA (>1382)
11 – 20	4.9	21	21	22	127
21 – 30	1.8	20	21	21	26
31 – 40	-0.3	20	21	21	23
41 – 50	-1.6	20	20	21	25
51 – 60	-1.9	20	20	21	25
61 – 70	-2.2	20	20	21	28
71 – 80	-1.2	20	20	21	24
81 – 90	-0.5	20	21	21	26
91 – 100	-6.3	21	22	31	NA (>918)

3.2.3. Coral dominance

Across all images, the AI selected the correct dominant coral type in 9,735 out of 11,741 images (83%). The AI selecting the correct coral dominance was reliant on the AI performance of differentiating and correctly identifying the coral types in each image. Hence, the performance of correct coral dominance of each coral type is dependent on the accuracy of coral types per image (Table 2).

Reef scale as a function of which coral dominates

The AI correctly predicted the dominant coral at the reef level in 175 of 189 reefs (93%).

In general, the AI correctly predicted the dominant coral type at reefs with more images available (correct reefs: mean = 63 images, median = 40 images. Incorrect reefs: mean = 57 images, median = 20 images; Figure 3).

Dominant coral types at reef scale (according to expert analysis):

- Table: 4.8% (9/189)
- Branching: 9.5% (18/189)
- Massive: 4.8% (9/189)
- Other: 81.0% (153/189)

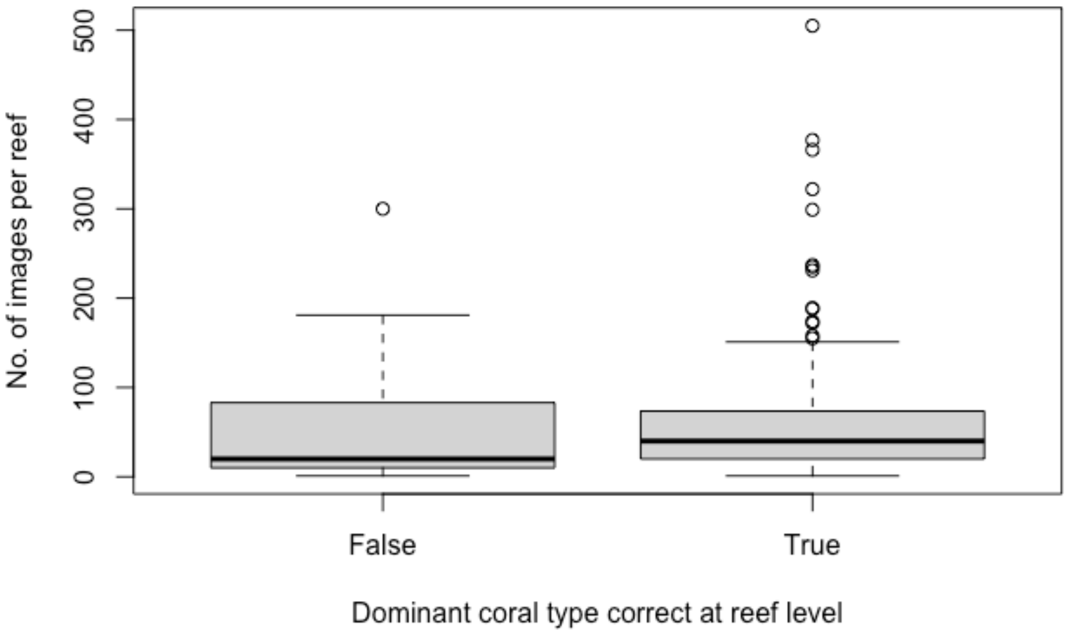

Figure 3: The number of images available for each reef according to if the AI correctly identified the dominant coral type at the reef level.

Site scale as a function of which coral dominates

The AI correctly predicted the dominant coral at the site level in 309 of 339 sites (91.2%).

In general, the AI correctly predicted the dominant coral type at sites with more images available (correct sites: mean = 35 images, median = 24 images. Incorrect sites: mean = 40 images, median = 14 images; Figure 4).

Dominant coral types at site scale (according to expert analysis):

- Table: 5.0% (17/339)
- Branching: 9.4% (32/339)
- Massive: 4.7% (16/339)
- Other: 80.8% (274/339)

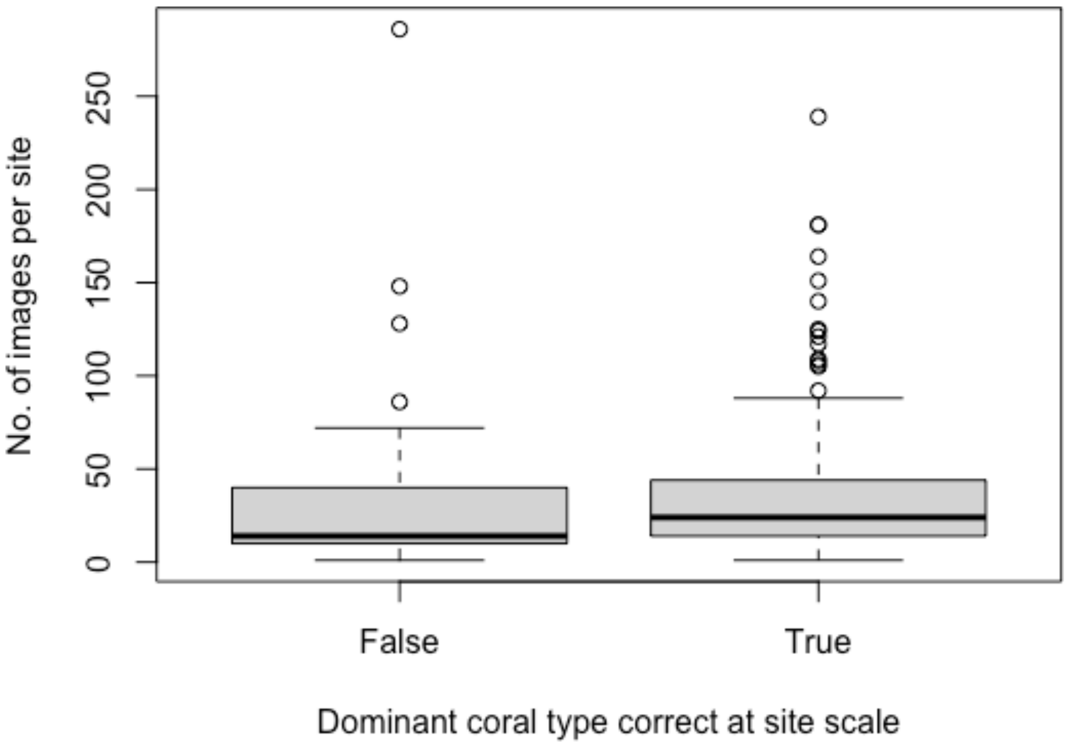

Figure 4*:* The number of images available for each reef according to if the AI correctly identified the dominant coral type at the site scale.

3.3. Evaluation of School Program data

The school program saw 302 school students complete 24,000 individual image analyses. The top 5 students who participated in the program analysed over 500 images each, with the top student analysing 926 images.

3.3.1. Comparison of accuracy of AI-supported school program vs AI alone

The school program analysis was less accurate than the AI analysis in all categories except reef cover and less precise in all categories, with a greater standard error (Table 9). Hence, the school program performance was more variable compared to the AI results.

Table 9*:* Performance of school program against Expert and AI analysis. Numbers in parentheses are the percent differences observed between AI and the expert analysis taken from Table 2.

	Reef cover	Coral cover	Table cover	Branchin g cover	Massiv e cover	Other coral cover
Mean	0.69 (2.53)	8.70 (-3.42)	3.96 (-2.15)	4.16 (-1.13)	3.41 (-1.04)	-2.84 (0.89)
SE	0.43 (0.15)	0.50 (0.19)	0.41 (0.06)	0.43 (0.09)	0.37 (0.06)	0.49 (0.19)

3.3.2. Utilising discrepancies between the school program and AI

The results from the school program data may be used in conjunction with the AI output by acting as a filter, i.e. any images that have a large discrepancy between the AI and school program outputs can be removed to improve the accuracy of the AI interpretation (Figure 5).

For example, when using only the images that were used in all three analyses (experts, AI and school program), by removing images with more than 10% difference between AI and the school program reef substrate results (547 images; i.e. the most strict filter):

- Reef structure mean accuracy: decreased from -0.4% to -4.0%.
- Coral cover mean accuracy: improved from -4.08% to -1.15%.
- Table coral mean accuracy: improved from -3.5% to -1.6%.
- Branching coral mean accuracy: improved from -0.8% to -0.4%
- Massive coral mean accuracy: improved from -0.97% to 0.1%.
- Other coral mean accuracy: improved from 1.3% to 0.9%, although with slightly less linear improvement with increasing filter strictness compared to the other variables.

The use of such filters will be largely dependent on the goals of the individual management or research plan. In general, using the filter improved the results of the AI performance, especially for total coral cover (arguably the most commonly needed result from the Great Reef Census), which saw a ~3% absolute improvement in accuracy. However, the accuracy of reef structure marginally declined.

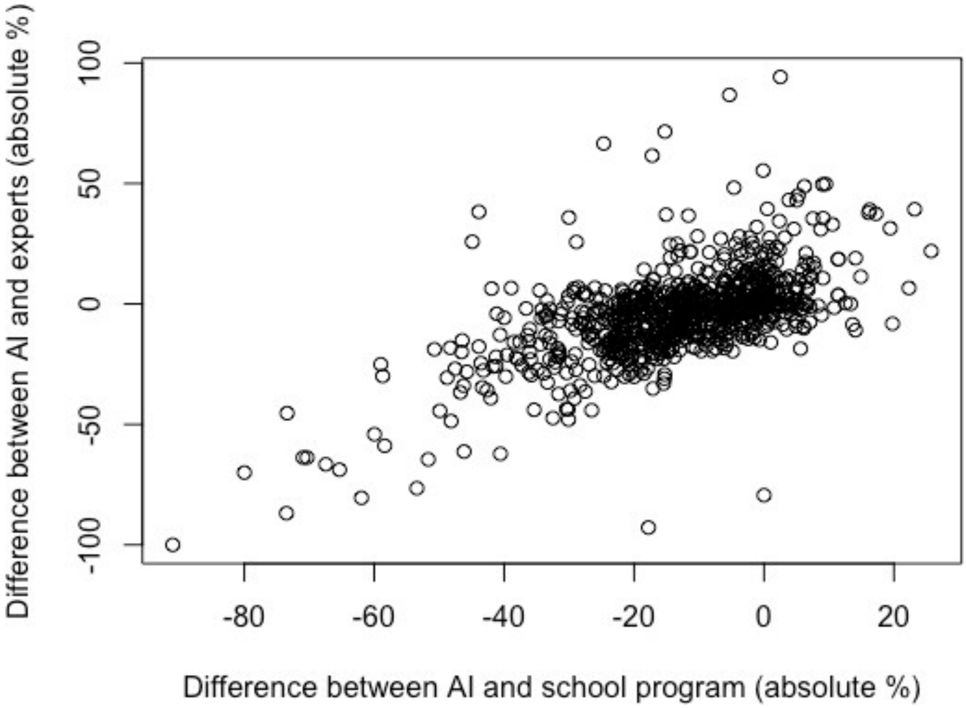

Figure 5: Differences between AI and school program outputs approximately correlate with the difference between the AI and experts. Each point represents one image.

Figure 6 shows the results of applying the filter of each threshold to each variable. As the filter applied becomes stricter, based on the difference between the school program and AI results of reef substrate, the AI is more accurate across all categories except for reef coverage (i.e. each category approaches the zero line, which indicates perfect accuracy). For example, total coral cover accuracy improves from -4.0% to -1.16%.

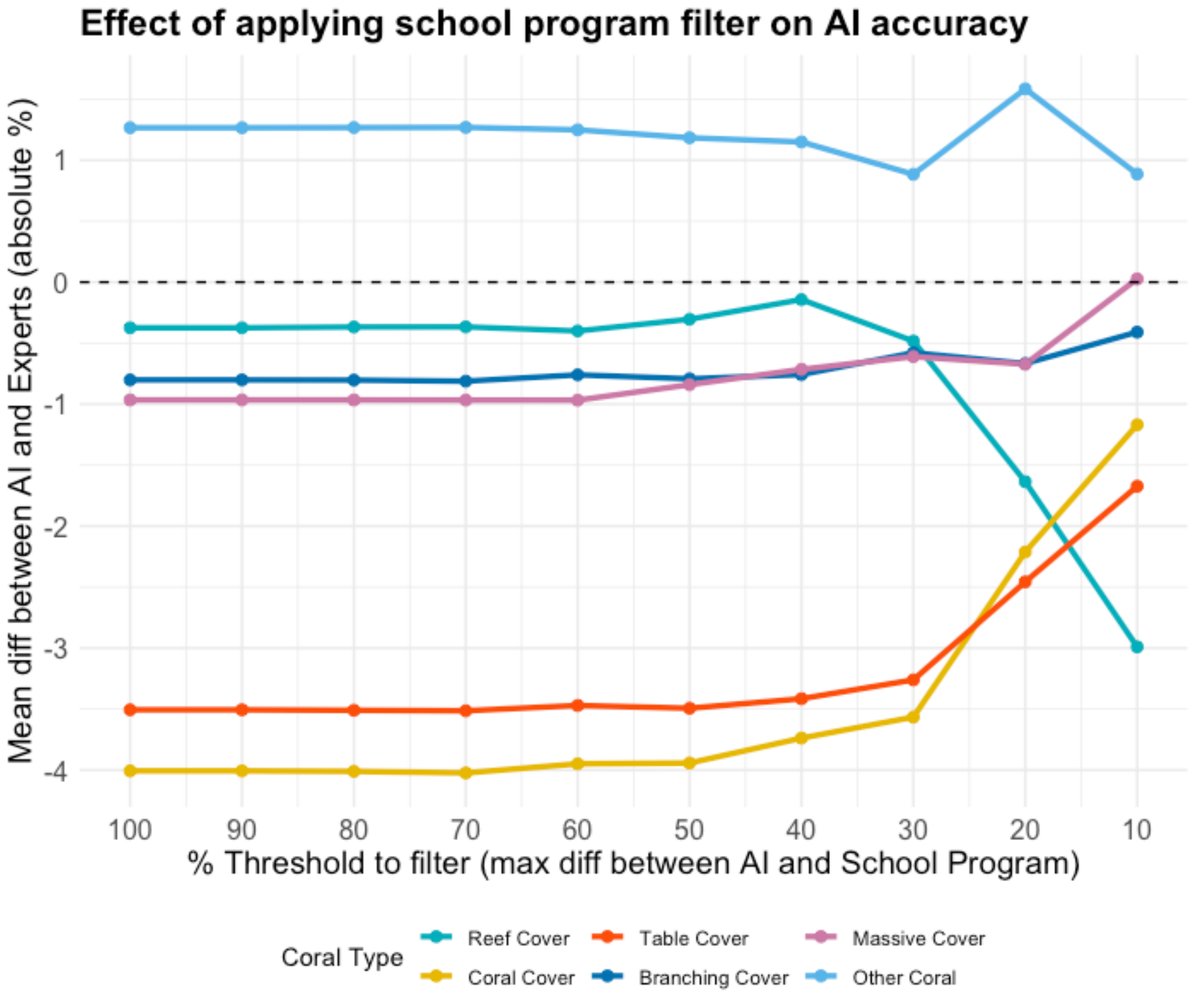

Figure 6*: T*he effect of applying a progressively strict filter that removes images based on the difference in reef substrate between AI and the school program. The left-hand side of the x-axis (100%) represents no filter applied, and on the right-hand side, the filter becomes progressively stricter (ultimately only keeping images with less than 10% difference between AI and school program results). The y-axis represents the mean difference between the AI and experts after the filter has been applied. The dashed dotted line at 0% represents no difference in mean results between the AI and experts.

To better investigate which images the filters were removing and potentially introducing bias in the results, we categorised the images in 10% bins according to coral cover (% of reef; Figure 7, Table 10). The filter removed an approximately equivalent number of images from high and low bins, i.e. it was not selecting predominately low or high reef state images. However, the question of bias created by applying such a filter needs to be assessed more thoroughly with a greater number of images from varying environmental and coral reef conditions.

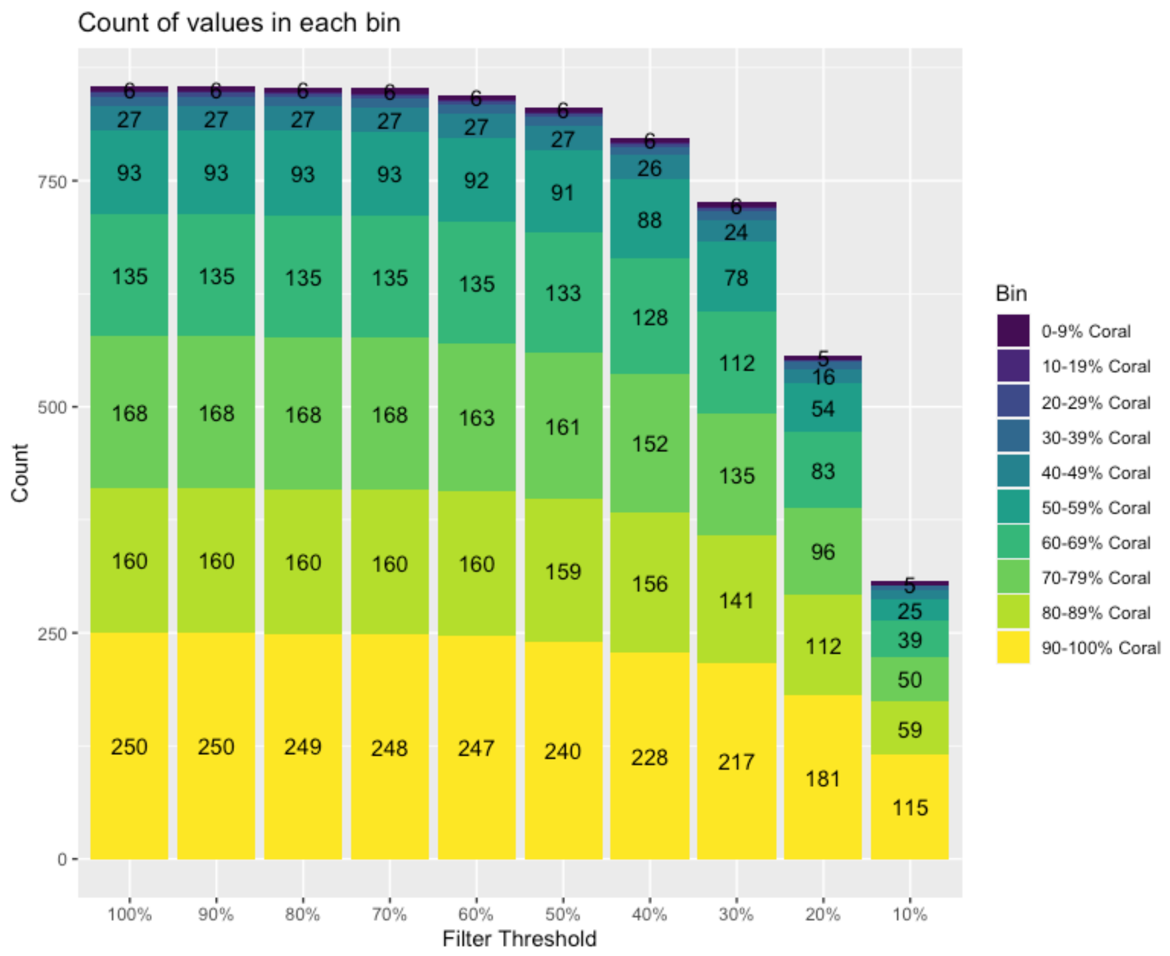

Figure 7*:* The number of images included in each level of coral cover at each threshold of filter (100% filter = all images, 10% filter = most strict filtering and fewest images based on difference between AI and school program results).

Table 10*:* Tabulated view of data presented in Figure 7 for clarity.

Threshold	**0-9% Coral**	**10-19% Coral**	**20-29% Coral**	**30-39% Coral**	**40-49% Coral**	**50-59% Coral**	**60-69% Coral**	**70-79% Coral**	**80-89% Coral**	**90-100% Coral**
10%	5	0	0	4	10	25	39	50	59	115
20%	5	0	2	8	16	54	83	96	112	181
30%	6	1	3	9	24	78	112	135	141	217
40%	6	1	4	9	26	88	128	152	156	228
50%	6	1	4	9	27	91	133	161	159	240
60%	6	1	4	10	27	92	135	163	160	247
70%	6	1	4	10	27	93	135	168	160	248
80%	6	1	4	10	27	93	135	168	160	249
90%	6	1	4	10	27	93	135	168	160	250
100%	6	1	4	10	27	93	135	168	160	250

Caveats to the filter results

Applying the filter introduces three important considerations. Firstly, the accuracy of reef substrate identification tends to decrease, necessitating further investigation. Secondly, the strictest filter, which prioritises quality, has the effect of removing 64% of images from the dataset, signifying a trade-off between quantity and quality. Lastly, there is a potential risk that the filter may inadvertently eliminate entire categories of images, such as those with low visibility, which could have broader implications. For instance, it might lead to the exclusion of entire sites or reefs from the analysis, particularly if all photos were taken within a condensed time frame. These caveats underscore the need for further investigation into the optimal balance between AI and citizen science.

4. Conclusions and recommendations

- Overall, the results of the AI performance demonstrated it to be capable of achieving accuracy of analysis of coral reef images that is comparable to trained human experts, and therefore, presents a useful tool for expanding the spatial scales of photo surveys of the Great Barrier Reef. The school program also demonstrated that citizen science can act as a useful tool for online analysis of coral reef survey images. Combining these two tools together demonstrated greater accuracy than either one alone.
- The AI estimates of coral cover were precise across most categories; in the images available for most categories, the AI output converged on the mean value with the minimum number of images (20).
- When using all available images, the AI output was accurate to ± 5% in all categories, including branching and table Acropora, which have been raised as important groups for management (e.g. informing crown-of-thorns starfish control and which reefs are likely to recover most quickly following disturbance).
- When images were divided into bins based on reef state (coral coverage), the AI output could not reach a ± 5% accuracy threshold in some cases. The highest number of images in any category or reef state bin that was needed for the 95% confidence interval to reach the target accuracy was 127 (followed by 58). This suggests that while in most cases 20 images at any given site or reef will be sufficient to reach the desired optimal accuracy (i.e. as high as possible for the AI), there may be a benefit to collecting up to ~127 images at each site.
- The school program output was generally less accurate than the AI analysis, however this may be improved by altering the process for the school students to follow (e.g. further training or adjusting the pre-set AI polygons).
- The school program analysis may be beneficial to use in conjunction with the AI to act as a filter to remove images in which the AI performs poorly. The mean accuracy of the AI output improved when removing images for which there was a discrepancy between the school program results and the AI results. The filter appears to remove images in an approximately equal manner across reef states (from 0 – 100% coral coverage).
- While the citizen science aspect of the analysis demonstrated capability in improving the results of an analysis system relying solely on AI, the AI performed adequately for most management needs in many image categories (i.e. coral cover ranges). Hence, citizen science analysis, such as the school program here, may provide more value by analysing categorically different questions, such as identifying if the coral is dead or bleached, which currently the AI has not been trained to detect.

4.1. **Future Directions**

The results of this study will be used to improve the methodology and expand outcomes across multiple areas:

- The school program approach has demonstrated that the current approach is a feasible method to gain buy-in and use by citizen scientists. Previously, the approach appeared too convoluted for lay people to participate; the current approach is simplified and presented in a more intuitive user interface, enabling significant participation by citizen scientists (e.g. some participants analysed several hundred images).
- While the AI output was generally highly effective, there are certain categories that did not meet desired accuracy thresholds (e.g. at low reef states). Using the school program as a filter for the AI appears to improve accuracy. Similarly, the school program (i.e. citizen scientists) effort may be directed towards the categories in which the AI performs poorly. A combined approach may optimise prediction accuracy and will continue to be tested.
- Continual optimisation and expansion of the citizen science analysis: the school program has acted as an effective pilot study for the citizen science image analysis and can now be used to expand more broadly outside of schools to potentially tens of thousands of citizen scientists globally. The results and feedback will be used to continually improve the user interface while focusing on the same six categories of coral presented here and that the AI has been trained on. This expansion will be rolled out to the general public over 2023 Q1 and Q2. For example:
    - Enable a human filter option by citizen scientists before they analyse an image to flag it as a “bad image”, e.g. blurry or poor water visibility.
    - Enable boundary adjustments of the polygons by citizen scientists. This may enable citizens to significantly improve accuracy without being fundamentally limited by the AI’s polygon boundaries.
    - Test citizen scientists’ ability on categorically different questions to the AI, such as detecting dead or bleached coral.

5. References

Andrew, N. L., & Mapstone, B. D. (1987). Sampling and the description of spatial pattern in marine ecology. *Oceanography and Marine Biology*, *25*, 39-90.

Bozec, Y. M., & Mumby, P. J. (2020). *Supplementary report to the final report of the coral reef expert group: S7. Coral reef models as assessment and reporting tools for the Reef 2050 Integrated Monitoring and Reporting Program–a review*. Great Barrier Reef Marine Park Authority.

Chandler, M., See, L., Copas, K., Bonde, A. M., López, B. C., Danielsen, F., ... & Turak, E. (2017). Contribution of citizen science towards international biodiversity monitoring. *Biological conservation*, *213*, 280-294.

Green, S. E., Rees, J. P., Stephens, P. A., Hill, R. A., & Giordano, A. J. (2020). Innovations in camera trapping technology and approaches: The integration of citizen science and artificial intelligence. *Animals*, *10*(1), 132.

Guo, Y., Liu, Y., Georgiou, T., & Lew, M. S. (2018). A review of semantic segmentation using deep neural networks. *International journal of multimedia information retrieval*, *7*, 87-93.

Hauser, C. E., Southwell, D., Lahoz-Monfort, J. J., Rumpff, L., Benshemesh, J., Burnard, T., ... & Bode, M. (2019). Adaptive management informs conservation and monitoring of Australia's threatened malleefowl. *Biological Conservation*, *233*, 31-40.

Hughes, C., Tremblett, K., Kummer, J., Lee, T. S., & Duke, D. (2022). How Can We Do Citizen Science Better? A Case Study Evaluating Grizzly Bear Citizen Science Using Principles of Good Practice in Alberta, Canada. *Animals*, *12*(9), 1068.

Johnson, C. N., Balmford, A., Brook, B. W., Buettel, J. C., Galetti, M., Guangchun, L., & Wilmshurst, J. M. (2017). Biodiversity losses and conservation responses in the Anthropocene. *Science*, *356*(6335), 270-275.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*(7553), 436-444.

Mammola, S., Riccardi, N., Prié, V., Correia, R., Cardoso, P., Lopes-Lima, M., & Sousa, R. (2020). Towards a taxonomically unbiased European Union biodiversity strategy for 2030. *Proceedings of the Royal Society B*, *287*(1940), 20202166.

McKinley, D. C., Miller-Rushing, A. J., Ballard, H. L., Bonney, R., Brown, H., Cook-Patton, S. C., ... & Soukup, M. A. (2017). Citizen science can improve conservation science, natural resource management, and environmental protection. *Biological Conservation*, *208*, 15-28.

Millar, E. E., Hazell, E. C., & Melles, S. J. (2019). The 'cottage effect'in citizen science? Spatial bias in aquatic monitoring programs. *International Journal of Geographical Information Science*, *33*(8), 1612-1632.

Moedas, C. (2018). *Citizen science: innovation in open science, society and policy*. UCL Press.

Mumby, P. J., & Harborne, A. R. (1999). Development of a systematic classification scheme of marine habitats to facilitate regional management and mapping of Caribbean coral reefs. *Biological conservation*, *88*(2), 155-163.

R Core Team (2022). R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. https://www.R-project.org/

Roelfsema, C. M., Kovacs, E. M., Ortiz, J. C., Callaghan, D. P., Hock, K., Mongin, M., ... & Phinn, S. R. (2020). Habitat maps to enhance monitoring and management of the Great Barrier Reef. *Coral Reefs*, *39*, 1039-1054.

Ruiz-Gutierrez, V., Bjerre, E. R., Otto, M. C., Zimmerman, G. S., Millsap, B. A., Fink, D., ... & Robinson, O. J. (2021). A pathway for citizen science data to inform policy: A case study using eBird data for defining low-risk collision areas for wind energy development. *Journal of Applied Ecology*, *58*(6), 1104-1111.

Rusk, N. (2016). Deep learning. *Nature Methods*, *13*(1), 35-35.

Stockman, G., & Shapiro, L. G. (2001). *Computer vision*. Prentice Hall PTR.

Vohland, K., Land-Zandstra, A., Ceccaroni, L., Lemmens, R., Perelló, J., Ponti, M., ... & Wagenknecht, K. (2021). *The science of citizen science* (p. 529). Springer Nature.

Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J. M., & Luo, P. (2021). SegFormer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, *34*, 12077-12090.

Zhang, H., Dana, K., Shi, J., Zhang, Z., Wang, X., Tyagi, A., & Agrawal, A. (2018). Context encoding for semantic segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition* (pp. 7151-7160).

CONTACT

Andy Ridley
andy@citizensgbr.org

citizensgbr.org

Marine and Coastal

National Environmental Science Program

CITIZENS OF THE GREAT BARRIER REEF

National Environmental Science Program

This project is supported with funding from the Australian Government under the National Environmental Science Program.