Human Computation (2021) 8:2:54-75

ISSN: 2330-8001, DOI: 10.15346/hc.v8i2.123

Crowdsourcing, computing, and conservation: how citizen science and artificial intelligence can improve the use of camera trap data to tackle large-scale ecological challenges

MEREDITH S. PALMER, Princeton University

SARAH E. HUEBNER, University of Minnesota

MARCO WILLI, University of Minnesota

LUCY FORTSON, University of Minnesota

CRAIG PACKER, University of Minnesota

ABSTRACT

Camera traps - remote cameras that capture images of passing wildlife - have become a ubiquitous tool in ecology and conservation. Systematic camera trap surveys generate 'Big Data' across broad spatial and temporal scales, providing valuable information on environmental and anthropogenic factors affecting vulnerable wildlife populations. However, the sheer number of images amassed can quickly outpace researchers' ability to manually extract data from these images (e.g., species identities, counts, and behaviors) in timeframes useful for making scientifically-guided conservation and management decisions. Here, we present 'Snapshot Safari' as a case study for merging citizen science and machine learning to rapidly generate highly accurate ecological data from camera trap surveys. Snapshot Safari is a collaborative cross-continental research and conservation effort with 1500+ cameras deployed at over 40 protected areas in eastern and southern Africa, generating millions of images per year. As one of the first and largest-scale camera trapping initiatives, Snapshot Safari spearheaded innovative developments in citizen science and machine learning. We highlight the advances made and discuss the challenges that arose using each of these methods to annotate camera trap data. We end by describing how we combined human and machine classification methods ('Crowd AI') to create an efficient integrated data pipeline. Ultimately, by using a feedback loop in which humans validate machine learning predictions and machine learning algorithms are iteratively retrained on new human classifications, we can capitalize on the strengths of both classification methods while

mitigating their weaknesses. Using Crowd AI to quickly and accurately 'unlock' ecological Big Data is revolutionizing the way we take on critical environmental issues in the Anthropocene era.

---

1. CAMERA TRAP 'BIG DATA' IN CONSERVATION - OPPORTUNITIES AND CHALLENGES

The Earth is currently undergoing extreme human-driven biodiversity loss, with devastating impacts on the functioning of ecological systems (Barnosky et al., 2011). In the current extinction crisis, collecting and analyzing Big Data is critical for identifying global biodiversity trends and diagnosing drivers of wildlife decline (Jetz et al., 2012; Stephenson et al., 2017; Kays et al., 2019). Technological advances are transforming our ability to rapidly gather massive quantities of high-resolution ecological data at unprecedented spatial and temporal scales (Rich et al., 2017; Steenweg et al., 2017). Remote (e.g., satellite) and in situ (e.g., camera traps, biologgers) sensors allow monitoring of entire wildlife communities across large areas and over long time periods (O'Connell et al., 2010; Wearn & Glover-Kapfer, 2017). While this information is vital for management and conservation, the rate of accumulation quickly outpaces researchers' abilities to extract the necessary data from collected information on timescales necessary to effectively understand and protect wildlife populations (Ahumada et al., 2019).

Camera traps have emerged as a popular tool in ecological research, enabling systematic collection of spatial and temporal information on wildlife community dynamics (O'Connell et al., 2010). These remote cameras are automatically triggered by passing animals, unobtrusively collecting data on the abundance, distribution, and behavior of medium- and large-bodied vertebrates. Each camera trap trigger generates a photographic record that includes the date, time, and location of the animal observation. Recent technological innovations have improved camera trap capabilities while lowering costs, leading to an exponential increase in the number of camera trap studies (Burton et al., 2015). While camera trap use has grown, a major obstacle to harnessing their full potential is the substantial time and effort involved for researchers to manually annotate each image (Ahumada et al. 2019). Even an average-sized survey (~78 camera traps; Steenweg et al. 2017) produces tens to hundreds of thousands of images over a 12-month period, which can take a research team months of dedicated work to classify (Ahumada et al., 2019; Glover-Kapfer et al., 2019). This image annotation bottleneck severely limits the usefulness of these data to address rapidly-changing local and global challenges.

Here, we present a case study of how citizen science and machine learning (ML), independently and, ultimately, combined created new opportunities for efficiently processing Big Data generated by one of the world's largest camera trapping initiatives. 'Snapshot Safari' (www.snapshotsafari.org) is a collaboration between dozens of standardized wildlife monitoring surveys across eastern and southern Africa. This network produces vast quantities of image data at a continent-wide scale, analysis of which drives scientific discovery and conservation policy. As one of the first large-scale camera trap projects, Snapshot Safari has been at the forefront of developments in citizen science (Swanson et al., 2015, 2016a,b; Hines et al., 2015; Kosmala et al., 2016), ML (Villa et al., 2017; Norouzzadeh et al., 2018; Tabak et al., 2019; Willi et al., 2019), and now, the powerful integration of human and machine classifications into a single pipeline, what we term 'Crowd AI'. We outline our successes using citizen science and ML to process camera trap images within the quick timeframe necessary to address critical conservation issues and describe the barriers we encountered employing each of these methods independently. We end by detailing the assimilation of both annotation techniques in our current pipeline and highlighting current challenges and new opportunities presented by using Crowd AI to process ecological data from camera trap images.

2. *SNAPSHOT SAFARI* - PROJECT DESCRIPTION AND AIMS

Substantive conservation science increasingly relies on trans-national collaborations to understand the drivers for global wildlife decline (O'Brien et al., 2010; Hampton et al., 2013) and standardized sampling over multiple sites allows direct comparison of the ecological and anthropogenic factors shaping animal communities at regional, continental, or even global scales (Rich et al., 2017; Steenweg et al., 2017). *Snapshot Safari* began as a single camera trap survey in Serengeti National Park, Tanzania, in 2010 (*'Snapshot Serengeti'*; Swanson et al., 2015, 2016a) and has grown into a collaborative effort managing surveys run in conjunction with 30 partner organizations across six African countries. Camera trap monitoring within each protected area provides continuous, fine-scale data on mammal and bird species >1 kg (O'Connell et al., 2010; Wearn & Glover-Kapfer, 2017). Data generated by *Snapshot Safari* are being used by government agencies, conservation organizations, protected area managers, and academic institutions to understand patterns of wildlife movement, behavior, and interactions and to examine how these trends change across environmental gradients and in the face of anthropogenic perturbations (e.g., Anderson et al., 2016; Swanson et al., 2016b; Palmer et al., 2017, 2019; Allen et al., 2018; Palmer & Packer, 2018; Muzena et al., 2019).

In the *Snapshot Safari* network, over 40 permanently-deployed camera trap surveys (totaling ~1500 individual cameras as of 2020) generate in excess of three million images every year. At each site, cameras are laid out in systematic 5-km^2 grids and positioned to maximize the likelihood of observing medium- and large-sized vertebrates (see Swanson et al., 2015 for details on survey design). Each camera is programmed to take a 'capture event' – a series of three quickfire images when motion and heat sensors are triggered during the day and a single image at night. These individual surveys operate continuously, monitoring expanses of up to 1400-km^2, and are intended to run for a decade or longer. In the following sections, we work through the successes and challenges encountered while developing the *Snapshot Safari* platform.

3. FIRST HURDLE: OVERCOMING DATA PROCESSING BOTTLENECKS THROUGH CITIZEN SCIENCE

'Citizen science', the colloquial term for the involvement of non-professionals in scientific research, has played a key role in environmental research for centuries (Lepczyk et al., 2009). Technological advances over the past few decades have proliferated the number and diversity of projects citizen scientists can participate in, either as data contributors, processors, or analyzers (Shirk et al., 2012; Kosmala et al., 2016; Watson & Floridi, 2018). In particular, the advent of online citizen science now allows millions of participants to advance scientific research by classifying (predominantly) image, video, or audio information, enabling ecologists to gather and utilize data on a scale that would be infeasible without crowdsourcing (Bonney et al., 2014, 2016; McKinley et al., 2015; Jennett et al., 2016).

The world's largest online citizen science platform, Zooniverse (www.zooniverse.org), was created in 2007 with the goal of engaging volunteers in data classification tasks that would otherwise exceed the researchers' own capacity (Fortson et al., 2012; Prather et al., 2013; Watson & Floridi, 2018; Trouille et al. 2019). Data object annotation is decomposed into simple tasks and crowdsourced to volunteers (Rosser & Wiggins, 2018). While the Zooniverse platform was originally developed for astronomy projects, through the launch of the Zooniverse Project Builder in 2015 (www.zooniverse.org/lab) it has come to host a broad range of projects in biology, earth sciences, art, history, and social sciences (Simpson et al., 2014, Spiers et al., 2019).

The original *Snapshot Safari* initiative, *Snapshot Serengeti* (www.snapshotserengeti.org), was launched in 2012 as the first Zooniverse camera trap project. As such, *Snapshot Serengeti* was pivotal in developing the camera trap classifying interface and backend that are now used by >50 Zooniverse projects (Data Availability [2a]). In this system, volunteers view camera trap capture events along with a list of possible species, which can be filtered based on morphological traits such as coat color or horn shape (Figure 1). Instructional tools including tutorials, field guides, and FAQs help volunteers identify animals in the camera trap images. After selecting a possible species classification, users are presented with pictures of that species from various angles, a detailed description of the species' characteristics, and comparisons to 'commonly confused with' species to assist in verifying the classification. After identification of all species in the image, users count the number of individuals, classify the animals' behavior, and indicate the presence/absence of juveniles and the presence of horns in sexually dimorphic antelope (Figure 1). Captures are circulated to multiple users (originally, 10-25 volunteers for each image containing animals, 5-10 volunteers for 'empty' images) and resulting data are aggregated based on the plurality of classifications to form a dataset that is 98% accurate compared to expert classifications (see Swanson et al. 2015, 2016a for further information).

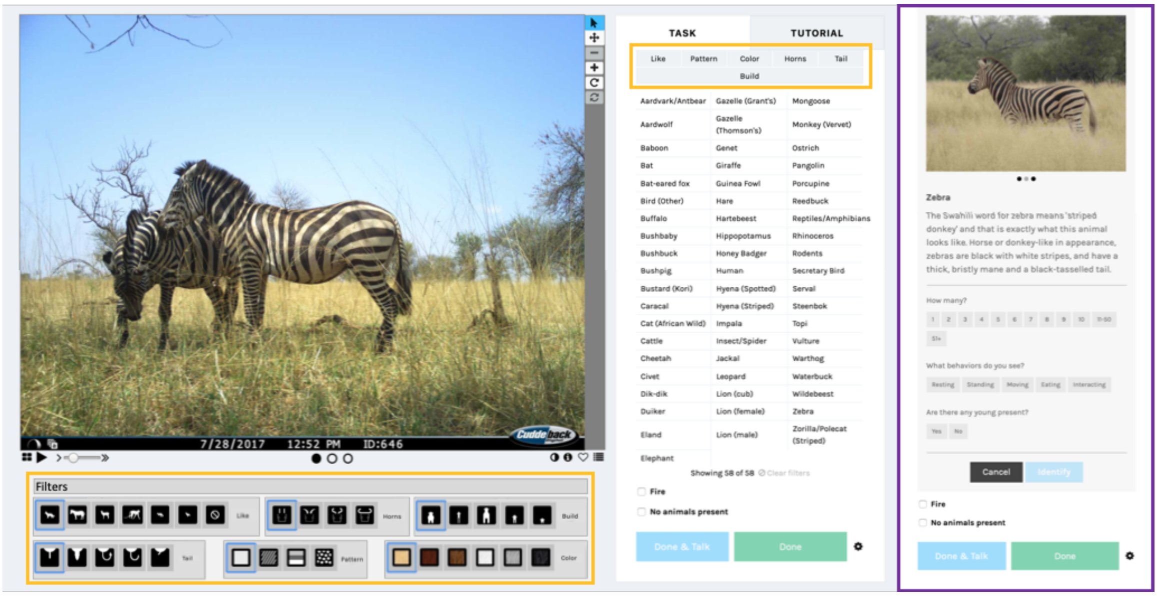

Figure 1. ***A typical Snapshot Safari project interface on the Zooniverse: the camera trap image is presented on the left-hand side and a species list on the right. Volunteers are able to select species directly or use the 'filter' tools (highlighted in orange; expanded in orange inset) to narrow down the selection. Once a species is chosen, additional information about the species and the opportunity to count individuals and tag behaviors and other additional annotations is presented (purple inset).***

3.1 Benefits of citizen science for processing camera trap imagery

The first three years of *Snapshot Serengeti* camera trap data were classified by over 30,000 volunteers who collectively donated a total of ~14.6 years of 40-hour-a-week effort (Swanson et al., 2015). This engagement has continued to grow: >138,000 users worldwide have contributed annotations on more than nine million images generated by *Snapshot Safari* projects between 2013-2020.

Beyond the benefits provided to researchers, the accessible nature of online citizen science enables novel forms of science education and outreach (Dickinson & Bonney, 2012; Shirk et al., 2012). We use the *Snapshot Safari* platform as a tool to disseminate information on African wildlife, share ongoing research and conservation efforts, and provide non-scientists with an intimate understanding of the scientific process. Participants are encouraged to interact with project researchers remotely via discussion boards and social media or in person at citizen science events. We have developed educational materials, including activities, multimedia, and course curricula, based on the *Snapshot Safari* platform to teach ecological principles and to engage students in authentic research experiences (Data Availability [1a-c]).

3.2 Challenges in the application of citizen science

Volunteers value learning, a sense of purpose, and the connections they develop with the research team (Rotman et al., 2012). As such, the researchers need to actively interact with volunteers to retain levels of participation sufficient to process data in the required timeframe (Cox et al., 2015). To draw in new volunteers and, of particular value, to retain experienced users, it is important to create and maintain engaging and informative websites, discussion boards, and social media platforms (Rotman et al. 2012; Reed et al., 2013). Project researchers can maximize participation by communicating frequently with volunteers and upticks in classifications are commonly observed after targeted engagement efforts. For example, classification rate approximately doubled after researchers conducted interactive educational webinars (42,639 classifications the day before vs. 94,625 the day after; April 2020) and podcast interviews in which volunteer contributions were explicitly highlighted (6,789 vs. 12,460; September 2020). We also found it highly beneficial to have a strong moderator corps that can amplify messages from the research team to volunteers and are knowledgeable enough to assist others with classifications. Overall, skills such as social media management, website design, and public communication are key for running a successful online citizen science project and the amount of effort that it takes to manage these platforms, particularly for the *Snapshot Safari* team which hosts dozens of projects, is not inconsequential.

Another complication is ensuring that 'sensitive' data are not viewed by the public. *Snapshot Safari* monitors certain protected areas that are subject to high rates of poaching and, as such, efforts are made to reduce the online circulation of images capturing wildlife crime and highly endangered species (e.g., rhinoceros, pangolin). Humans (e.g., protected area staff, tourists) are occasionally captured on camera and we act to remove these images from our platforms as quickly as possible. Although metadata such as geographic details are hidden from volunteers, minimizing exposure of these images reduces the possibility that contextual information from the photos could be used for unethical reasons.

When *Snapshot Serengeti* was released, it was the first project of its kind; over the last decade, the number of online camera trap citizen science projects has proliferated exponentially. There are now dozens of camera trap projects on the Zooniverse, with additional projects appearing on other platforms (e.g., *eMammal* www.emammal.si.edu, *SciStarter* www.scistarter.org, *MammalWeb* www.mammalweb.org). While participants continue to contribute millions of classifications every year, the available volunteer-base has grown more slowly than the number of projects relying on citizen science (Willi et al., 2019). It is increasingly challenging to recruit sufficient effort to process data, leading to substantial delays between data collection and analysis for real-world application. As these data are vital for practitioners who need to track population demographics, monitor threatened species, and assess outcomes of existing programs, the longer it takes for data to be processed, the less value they have for active, ongoing conservation initiatives.

4. SECOND HURDLE: EXPEDITING DATA PROCESSING WITH ARTIFICIAL INTELLIGENCE

We next investigated the use of artificial intelligence (AI) for camera trap image processing. New advances in deep learning, a sub-field of AI, allow computers to perform complex tasks with near human-level accuracy (Russakovsky et al., 2015). Particularly, convolutional neural networks (CNNs) have high accuracy in visual tasks such as image classification and object recognition, making them a powerful tool for ecology and conservation projects that rely on camera trapping (Norouzzadeh et al., 2018).

CNNs evaluate image data in a hierarchical fashion similar to the way mammalian brains process visual information (Eickenberg et al., 2017). They are trained to identify objects of interest by learning how to associate raw inputs (i.e., pixel values) to labelled image data (for illustrations on how CNNs identify animal species, see Miao et al., 2019). CNNs are a powerful class of models because they are able to learn and extract complex visual features without direct supervision of a human expert. However, learning to associate specific features with objects of interest - in this case, animal species - requires a substantial amount of labelled data. For example, the ImageNet computer vision challenge dataset contains 1,000 images per object class (Russakovsky, 2015). Real-world datasets are often highly imbalanced, dominated by images of common objects with few examples of rarer classes. As such, rare objects (species) are more difficult for algorithms to recognize (see results and discussion in Norouzzadeh, 2018). Camera trap images present additional challenges in the application of deep learning as animals in the images can be difficult to discern or distorted by occlusion, unusual perspectives, camouflage, and more (Beery, 2018).

When *Snapshot Serengeti* went online in 2012, CNNs had just demonstrated their superiority to other algorithms by significantly reducing error rates on image classification tasks (Krizhevsky, 2012). At that time, however, we lacked the necessary volume of labelled data to train a wildlife identification algorithm for camera trap images. Over the following years, the subsequent contributions of citizen scientists generated an unprecedented dataset of several million labelled images of 48 animal species (Swanson et al., 2015). This unique resource enabled ML researchers to significantly improve automated animal recognition (Chen et al., 2014, Villa et al., 2017, Tabak et al., 2018), reaching up to 96% accuracy (Norouzzadeh et al., 2018).

4.1 AI for camera trap image processing

A typical *Snapshot Safari* survey is composed of ~80% empty capture events; this is a common camera trapping issue, as moving vegetation can trigger a camera's sensor and lead to long sequences of images containing only shrubs or grass (Green et al. 2020). Removing empty images is tedious and time-consuming for humans but can be quickly and efficiently accomplished by AI. We decomposed the classification task into two components: first, eliminating empty images (e.g., misfires that do not contain animals) and second, extracting several types of data from images of wildlife. This separation required us to train two separate models, the first which identifies wildlife presence (*'Empty or Not'*) and the second (*'Species'*) that extracts information from capture events flagged as non-empty. This approach is closely aligned with the citizen science workflow and has additional advantages as discussed in Willi et al. (2019) and Norouzzadeh et al. (2018) (Data Availability [2b]).

We trained our models on a dataset of 3.66 million labelled capture events on 85 species from nine *Snapshot Safari* surveys (three in Tanzania, one in Mozambique, and five in South Africa; Data Availability [3]) and images collected across South Africa by a similar camera trapping project, *Camera CATalogue* (www.panthera.org/cameracatalogue). While there was large overlap in species community composition across these surveys, including data from a diverse array of sites allowed us to generate an algorithm that could perform across substantially different environments (e.g., image backgrounds; Beery, 2018).

Our modelling approach is based on the description in Willi et al. (2019), with a few notable changes. We used a more complex model architecture for the *Species* model (Chollet, 2017) and added additional outputs that classify animal behaviors, detect the presence of young animals, and estimate animal counts. While the *Empty or Not* model architecture remains unchanged, we increased the input image resolution from 224x224 to 448x448 pixels to better detect small or obscured animals. The most significant difference is the performance improvements resulting from using a substantially larger and more diverse training dataset (~3x larger than that used in Willi et al., 2019). Each model outputs a probability distribution over all objects it was trained to recognize, with the sum over all probabilities (over all classes) on any given image equaling 100. The class with the highest probability is considered the model's classification for that image and the associated probability value can be interpreted as the model's confidence in its classification. With appropriate models, confidence thresholding (i.e., ignoring model classifications below a certain confidence value) can result in accurate predictions on a large subset of a dataset and is an effective way to identify images that need manual review (see Norouzzadeh et al., 2018; Willi et al., 2019).

Both models were evaluated on several test datasets which were not used during model training. One dataset was composed of capture events from sites that the model was trained on (in-sample) and three datasets were from new sites (out-of-sample). The latter provide a more realistic estimate of model performance if applied to new sites. The following figures refer to the much larger in-sample dataset.

The *Empty or Not* model achieved an overall accuracy of 96.0% at identifying empty images. Accuracy rose to 99.5% on images where the AI indicated greater than 95.0% ('high') confidence. The share of such high confidence images in the test dataset is 74.4%, meaning the accuracy of that part of the dataset can be accepted without additional input from researchers, while the remaining 25.6% may need to be reviewed.

The *Species* model extracts annotations comparable to those supplied by citizen scientists, specifically, species identity, count of animals, behavior (e.g., resting, standing, moving, eating, interacting), and presence of offspring. This model performed well on the species category, returning accurate responses of 89.3% across the board and 97.2% accuracy on high confidence images, roughly equal to the volunteers' aggregated accuracy of 97.9% (Swanson et al., 2016). Identification of behaviors and juveniles achieved 89.22% overall and count accuracy varies by binned count category, with low (count = 1, accuracy = 94.7%) and high (count = 11-50, accuracy = 89.0%) values achieving particularly high accuracy. See Supplemental Materials for breakdown of model accuracy and recall by attribute.

4.2 AI application issues

AI presents a substantial advance in how ecologists manage Big Data by significantly reducing the classification workload, but AI is not yet a panacea for camera trappers. In our project, AI algorithms achieve high accuracy on empty images and perform adequately on species identification if trained on large labelled image libraries. However, these models have yet to achieve confident predictions for rare species with few labelled images, images collected from areas unknown to the model, and, due to image-level classification of the training set (still the prevailing form of image data classification on the Zooniverse), multiple species in a single image or counts of large groups (Figure 2A-E). Every additional variable used by ecologists to assess abundance, species interactions, or demographic parameters (e.g., count, behavior, age, sex) introduces new layers of complexity in training a model (Figure 2F-I). This could have conservation consequences if AI fails to identify the presence of new, rare, or endangered species or if AI does not annotate information at a resolution necessary to develop accurate predictive models of wildlife responses to management interventions or environmental change (Wearn et al., 2019). Lastly, AI cannot yet compete with the human ability to recognize and flag intriguing interactions or behaviors which may lead to unexpected scientific discoveries.

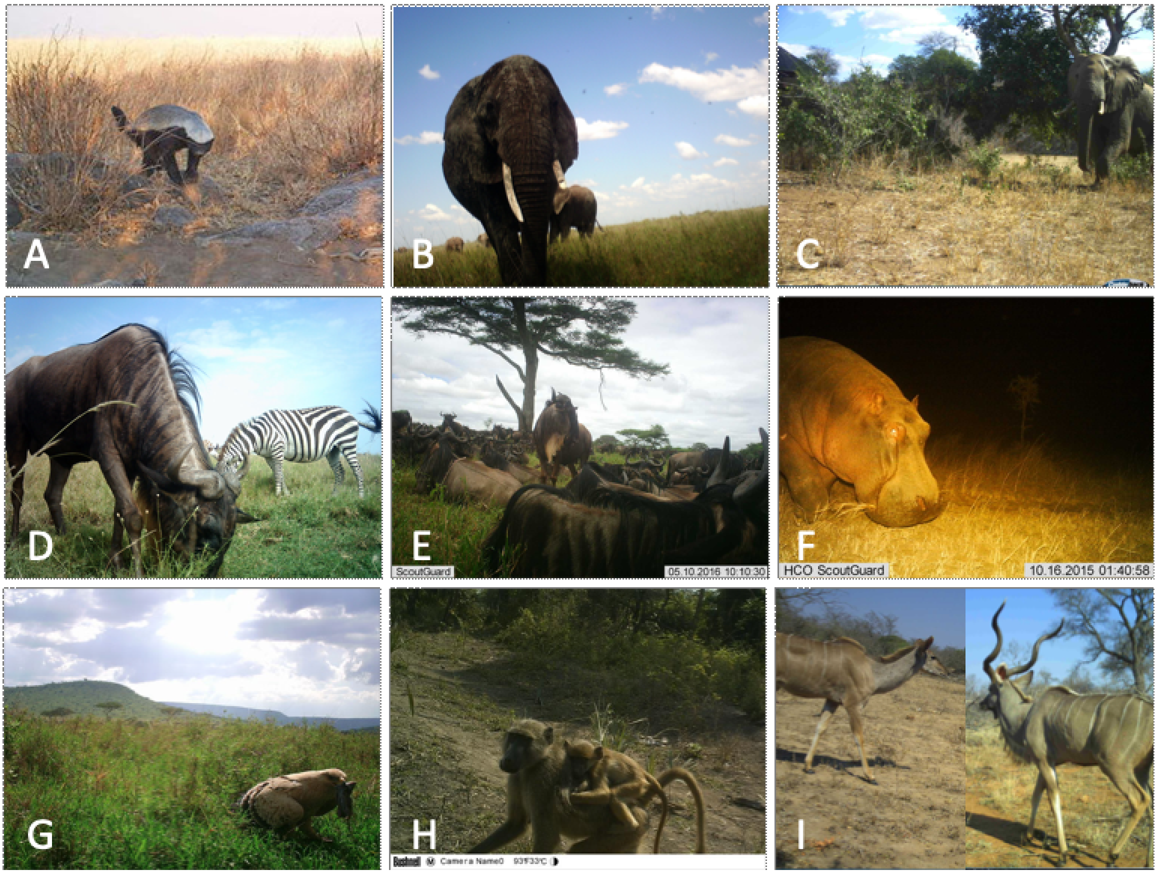

Figure 2. ***Examples of camera trap images that are difficult for our current ML algorithm to annotate: (A) rare species - honey badger compose less than 0.03% of the overall Snapshot Serengeti dataset; (B, C) different backgrounds - the savanna habitat in Serengeti National Park, Tanzania, differs significantly from the bushveldt habitat of Kruger National Park, South Africa; (D) multiple species - wildebeest and zebra commonly herd together and can consequently be captured in the same images; (E) large counts - animals in groups of >10 are difficult for humans and computers to resolve. Furthermore, additional image attributes are collected by human volunteers that our algorithm struggles to predict with reasonable accuracy: (F, G) behavior - activities such as 'feeding' can look different across species; (G, H) age and sex - extracting data on demographic parameters is important for understanding population dynamics.***

5. INTEGRATION OF AI AND CROWDSOURCING TO OVERCOME CURRENT CHALLENGES IN DATA PROCESSING

Humans have the pattern recognition capability to easily extract data, but it can become difficult to marshal enough volunteer effort to return information in a timely manner. In *Snapshot Safari*, issues still remain with using AI alone for data processing in terms of model accuracy, the types of data that can currently be extracted, and the loss of broader outreach benefits from engaging volunteers in authentic research experiences. The combination of ML and citizen science offers numerous advantages over using either technique independently (Green et al. 2020). Foremost is drastically improved classification rates without sacrificing accuracy or the potential to discover new, unanticipated information (Branson, 2017; Topol et al., 2019). Secondly, human input and verification of ML classifications is used to iteratively retrain the algorithms, enabling us to further advance ML accuracy and capability to handle difficult classification tasks. Furthermore, we continue to involve the public in authentic research, a mutually beneficial relationship that serves to increase scientific literacy and raise awareness of pressing conservation issues.

5.1 Crowd AI for classifying camera trap images

In the integrated *Snapshot Safari* data pipeline (Figure 3), the task of identifying empty images is primarily delegated to the ML algorithms, allowing us to concentrate volunteer effort on annotation tasks where the algorithm is unable to generate confident predictions (e.g., captures that contain data-deficient species, captures that are of poor quality due to glare, obstructions, or only portions of the animal are in view). Throughout this process, human classifications confirm and expand on the computer annotations. This ensures that important discoveries related to behavior, species interactions, and novel species are not lost due to current limitations of the AI. Citizen scientists' efforts then provide new labels for iteratively training the classifier between data batches to improve the accuracy of the first phase.

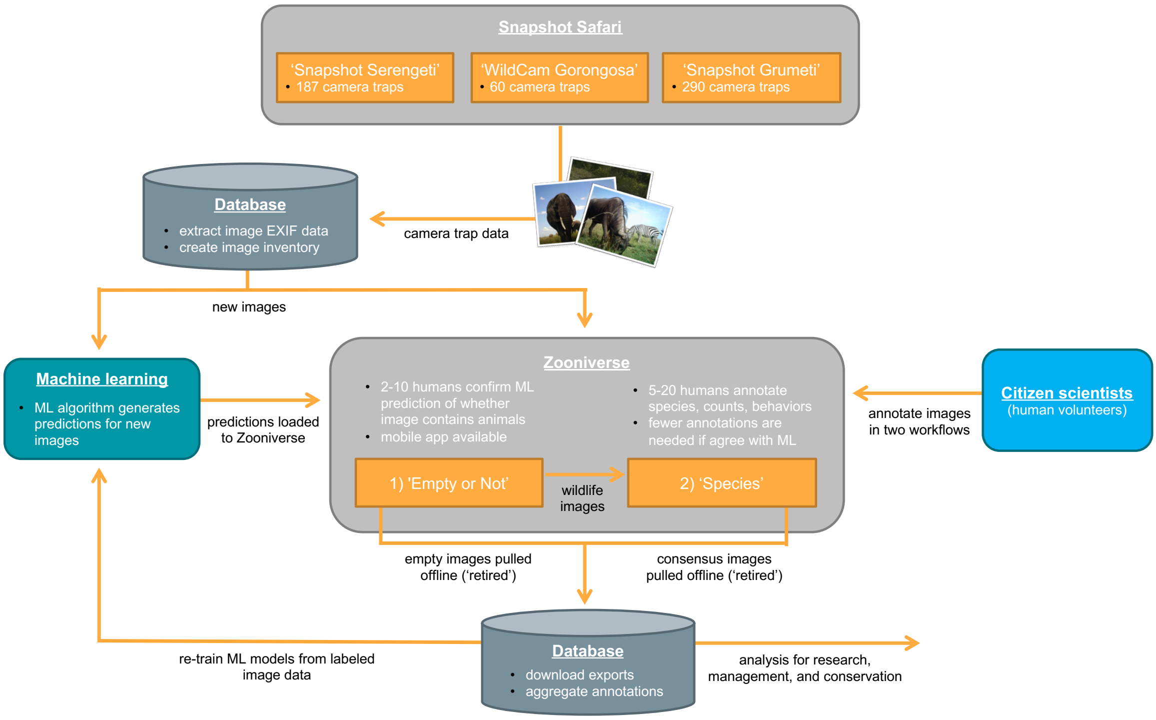

***Figure 3.* Workflow outline of the integrated 'Crowd AI' data pipeline.**

This integration is made possible by a key Zooniverse functionality, the '*Caesar*' decision engine, which allows us to determine whether to keep circulating an image for additional human annotation by comparing incoming volunteer classifications to ML-generated predictions (Data Availability [2a]; Fortson et al. 2018). With *Caesar*, researchers can set logical 'retirement rules' for a capture event (i.e., when we are confident that we can remove the capture event from online circulation having achieved highly accurate annotations) based on the ML predictions, the number of volunteers who have annotated the event, and the level of agreement among ML and volunteers. While Type II errors (i.e., model identifies an empty image as containing an animal) result in a slight increase in human effort, Type I errors (i.e., model identifies a wildlife-containing image as empty) can lead to potentially costly data loss. Although the *Empty or Not* CNN has a high degree of accuracy, we have adopted a conservative set of rules to ensure all images containing wildlife are properly identified (see Supplemental Materials for rules and background on their development and justification).

Images identified as non-empty are then moved to the *Species* workflow to annotate the identity, counts, behaviors, age, and sex of animals in the capture. When using citizen science alone, capture events were circulated to an average of 26 users to achieve an overall 98.6% accuracy on animal identifications and 75.2% accuracy on counts (Swanson et al., 2016a). With Crowd AI, we have reduced the number of human volunteers needed by half to produce comparable levels of accuracy (for example, average users to retire captures before/after implementing Crowd AI: *Snapshot APNR* 18.9/8.9; *Snapshot Serengeti* 16.0/8.1). As in the preceding workflow, dynamic retirement rules implemented through *Caesar* allow us to retire images sooner that exhibit high agreement among volunteers and/or between human and ML annotations. Distinct rule sets are applied to sensitive (i.e., humans or threatened species), common, and rare animals, allowing us to move quickly though the classification process for species the algorithm performs well on while ensuring maximum volunteer input on difficult images and minimum online circulation for sensitive species (see Supplemental Materials). The labelled images for rare species and for images initially mis-classified by the computer are then used to retrain the ML model to improve accuracy on future data.

The Zooniverse continues to develop new tools for integrating ML and citizen science classifications (Data Availability [2a]; Fortson et al. 2018, Trouille et al. 2019). In addition to *Caesar*, we have been able to use the Zooniverse's mobile application to rapidly process high volumes of images on the order of several million annually. The mobile app is ideal for binary classifications like the *Empty or Not* workflow, as volunteers are able to swipe left or right to indicate whether a capture is empty or contains wildlife. This improved accessibility on mobile devices expands the way that we interact with volunteers and the way volunteers engage with us.

Depending on the size of the dataset, removal of empty images can take as little as ~3-7 days, while the more complex species identification task can be completed in 4-6 weeks for small protected areas (~20 cameras) and several months for large national parks (~150-200 cameras). Classification times for the flagship *Snapshot Serengeti* project (currently 187 camera traps) have been significantly reduced with Crowd AI: prior to implementing the combined workflow, a batch of data collected over six months in the field took 452 days (~15 months) to process, while a similarly-sized batch with integrated ML predictions was processed in less than half the time (209 days or ~7 months).

6. CROWD AI: UNLOCKING NEW OPPORTUNITIES IN ECOLOGICAL RESEARCH

The *Snapshot Safari* Crowd AI pipeline is constantly being refined as new labelled image datasets become available. In our current system, volunteers are used to confirm computer classifications for empty images and identities of common species. The next challenge is improving our models to generate more accurate predictions for complex annotation tasks, such as counting animals, identifying multiple species in the same image, and classifying animal behavior (Figure 2). With improved classification models, the majority of volunteer effort can be directed towards tackling the much smaller percentage of the data that requires advanced human pattern recognition skills.

6.1 Current challenges in the use of Crowd AI

Determining when conservation efforts need to be implemented and tracking their success requires a precise understanding of animal population sizes. Counts of animals from camera trap images can be used to generate abundance estimates (e.g., Morellet et al. 2007, Palmer et al. 2018), but *Snapshot Safari* citizen science and, consequently ML, annotations are currently binned into categories ('11-50', '51+') too broad to accurately calculate population size. One issue is that certain species travel in enormous herds numbering hundreds or thousands of individuals, making it nearly impossible for humans to distinguish and count each animal (Figure 2E). Additionally, some volunteers only count animals near the camera (i.e., those that triggered the capture), while others attempt to also count animals in the far distance (i.e., those that would not have otherwise been photographed), creating a bimodal distribution that further complicates our ability to train an accurate counting model.

Due to the image-level labeling of the camera trap captures, another drawback with our current approach is that our AI cannot identify more than one species per capture event. While <2% of Snapshot Safari images overall contain multiple species (Norouddezzah et al., 2018; Willi et al., 2019; Figure 2D), these images contain valuable information about species interactions such as mutualisms, facilitation, competition, and predator-prey relationships (e.g., Palmer et al., 2017; Beaudrot et al., 2020).

At this stage, neither citizen scientists nor machine learning algorithms provide robust data on additional image attributes such as animal behavior. This issue is driven by lack of human agreement (and therefore, labelling) on what behavioral performance looks like within and across species. For example, to different volunteers, 'moving' could be interpreted as an animal going from place to place or a stationary animal moving parts of its body (e.g., tail swinging, grooming). Significantly different-looking actions can be associated with the same behavioral label across species: a wildebeest grazing on grass, a lion eating a carcass, and a bird catching insects would all be labelled as 'feeding' (Figure 2F-G). The problem is exacerbated, as with species classification, by image-level labeling of behaviors, possibility for multiple behaviors to be tagged in a single image, and highly imbalanced examples of each behavior. Currently, the best classifiers are only 76.2% accurate compared to human classifications (Norouzzadeh et al., 2018).

6.2 On-going advances in Crowd AI

On-going technological advances allow for increasingly sophisticated forms of data extraction and Crowd AI integration through our hybrid approach and are an active area of investigation.

CNNs with object detection capabilities offer a potential solution to improve count estimates, handle images containing multiple species, and address the issue of classifying complex behavioral data. While image classification models by definition only predict the presence of one species for a given image (i.e., are trained on and provide labels at the image-level), object detector models explicitly identify each individual animal in an image and output its location and species class (Beery et al., 2018; Schneider et. al, 2019; Parham et al., 2018; Data Availability [2c]). Object detection models are trained with localized annotations which are typically tight rectangular boxes ('bounding boxes') around each object of interest ('Data Availability [3]). Bounding boxes additionally allow for developing standardized cut-off points to determine which individuals to consider a valid observation. Counts could be limited to those animals that exceed a minimum detection size limit, thus preventing the inclusion of background animals that would never have been captured in the absence of a foreground animal triggering the camera trap. The size limit would likely vary by species as a distant elephant will be larger than a distant gazelle.

A prerequisite for training such models is collection of new forms of citizen science data. Most pressingly, we require specialized workflows to annotate animals within localized bounding-boxes within camera trap images. We also are working towards enhancing the specificity of volunteer responses to complex image annotation questions. One solution for standardizing behavioral annotation, for example, is to train volunteers using ethograms (a method of systematically cataloging types of behavior) for particular species of research or conservation interest. This would then allow for improved training of species-behavior-specific AI. Designing these novel workflows is an area of on-going development by *Snapshot Safari* in collaboration with the Zooniverse and the Microsoft AI for Earth team.

The Zooniverse continues to innovate advances to better facilitate the dynamic integration of machine learning with volunteer classifications. One area of improvement is in bringing to bear differing levels of expertise within the volunteer corps. For example, in the 'leveling up' strategy employed by the '*Gravity Spy*' citizen science project (Zevin et al. 2017), volunteers label 'easy' images first and as their experience and expertise grows, access subsequent 'levels' with progressively harder images. Using this infrastructure, camera trap projects could present to qualified individuals sets of images that need more detailed classification. *Gravity Spy* has also demonstrated that this leveling up mechanism provides a way for citizen scientists to make serendipitous discoveries: at each level volunteers have the opportunity to flag something as 'odd' or 'not known', which then promotes that image to the next level.

6.3 Zooming out: Applying lessons from the Snapshot Safari case study to other monitoring efforts

Our work with *Snapshot Safari* demonstrates that combining citizen science and ML capacities in an integrated workflow allows us to rapidly generate highly accurate ecological Big Data from camera trap surveys. We have been able to break new ground in citizen science, AI, and Crowd AI due to the exceptional scale of our projects, and we have made the results of our efforts freely available to facilitate the work of other researchers who may lack sufficient resources, data, or capacity to develop their own classification pipelines (Data Availability [2a-c, 3]). Our pipeline will continue to improve through iterative feedback between human volunteers and new algorithms.

The foundations laid by *Snapshot Safari* have opened up new frontiers in how ecologists can monitor, understand, and protect our natural world (Green et al. 2020). Long-term, large-scale camera trap monitoring programs have proliferated over the last few years, enabling direct comparison of ecological and anthropogenic drivers of community dynamics at regional, continental, or even global scales (Rich et al., 2017; Steenweg et al., 2017). Other programs currently generating large camera trap image datasets that benefit from citizen science, AI, or Crowd AI include *Snapshot USA* (www.emammal.si.edu/snapshot-usa), *Tropical Ecology Assessment and Monitoring (TEAM*; www.wildlifeinsights.org/team-network), and the Zoological Society of London (www.cloud.google.com/customers/zsl). Many efforts exist to synthesize and process existing camera trap datasets, such as *eMammal*, *MammalWeb*, *Wildlife Insights* (www.wildlifeinsights.org), and the Global Biodiversity Information Facility (www.gbif.org).

Our project is not alone in integrating citizen science and machine learning capabilities for camera trap image processing, though much ongoing work advancing AI and Crowd AI relies on the labelled image libraries generated by *Snapshot Safari* (Data Availability [3]). Collaborations include those led by Microsoft's *AI for Earth* initiative (www.microsoft.com/en-us/ai/ai-for-earth) and the American Museum of Natural History (www.biodiversityinformatics.amnh.org/ml4conservation/animal-detection-network/species-identification-localization). More broadly, the pairing of citizen science and machine learning is now being used to process or collect other types of ecological data (e.g., *WildMe* www.wildme.org and *iNaturalist* www.inaturalist.org for volunteer-gathered wildlife images; *Satellites over Seals* for satellite imagery, LaRue et al., 2020; 'biophony' audio data, Sueuer et al., 2019).

Rapid access to these data presents unprecedented opportunities for us to discover fundamental information on the structuring of wildlife communities and to explore and evaluate new measures for preserving biodiversity and community functioning. The rate and scale of data collection and now, data processing, allows us to use this information to make near-real time decisions that can guide scientifically-driven conservation decisions. Modern conservation requires modern tools, and our new capacity to close the Big Data analysis gap is better enabling us to take on new and increasingly urgent environmental challenges.

DATA AVAILABILITY

[1] Educational materials:

- [1a] Middle school-level activities and multimedia: https://doi.org/10.13020/5r00-8c56
- [1b] High school-level lesson plans and multimedia: https://classroom.zooniverse.org/#/wildcam-gorongosa-lab
- [1c] Undergraduate-level lab materials for conducting authentic research inquiries: https://doi.org/10.24918/cs.2020.49 (Palmer et al. 2020)

[2] Code repositories:

- [2a] Zooniverse's code repository, including the *Panoptes* platform and *Caesar* decision engine: https://github.com/zooniverse
- [2b] CNN for classifying camera trap images (code and models, updated from Willi et al., 2019): https://github.com/marco-willi/camera-trap-classifier
- [2c] Microsoft's code repository for training, running, and evaluating detectors and classifiers for camera trap images, developed in part using *Snapshot Safari* image data (maintained by AI for Earth): https://github.com/microsoft/CameraTraps

[3] Camera trap data:

- Labelled Image Library of Alexandria - Biology and Conservation (LILA-BC) repository, containing camera trap datasets from *Snapshot Safari* projects including data collected in Serengeti National Park (Tanzania), Camdeboo National Park, Karoo National Park, Kgalagadi Transfrontier Park, Kruger National Park, and Mountain Zebra National Park (South Africa) and the Enonkishu Conservancy (Kenya). New data is uploaded as classifications are provided by citizen scientists (maintained by Microsoft's AI for Earth): http://lila.science/datasets

ACKNOWLEDGEMENTS

We thank our partners at the Zooniverse and the thousands of citizen scientists who have participated in our research, particularly those who help moderate *Snapshot Safari* platforms. We acknowledge the Minnesota Supercomputing Institute (www.msi.umn.org) for providing resources contributing to data storage and processing. We would like to thank our partner organizations, many of whom provided funding for this work: Cincinnati Zoo & Botanical Gardens, Detroit Zoological Society, Living Desert Zoo, Nelson Mandela University, Ruaha Carnivore Project, Seneca Park Zoo, Singita Grumeti Fund, South African National Biodiversity Institution, South African National Parks, University of Wyoming, and Zoo Miami. This work was supported by AZA's Saving Animals from Extinction, Disney's 'Protect the Pride', Microsoft's 'AI for Earth', National Geographic Explorer grants, U.S. National Science Foundation grants: PRFB #1810586: IIS #1619177: OAC #1835530, and National Research Foundation of South Africa. South African research permits were provided by South African National Parks PACC1403.

REFERENCES

Ahumada, J. A., Silva, C. E., Gajapersad, K., Hallam, C., Hurtado, J., Martin, E., ... & Sheil, D. (2011). Community structure and diversity of tropical forest mammals: data from a global camera trap network. Philosophical Transactions of the Royal Society B: Biological Sciences, 366(1578), 2703-2711.

Allen, M. L., Peterson, B., & Krofel, M. (2018). No respect for apex carnivores: distribution and activity patterns of honey badgers in the Serengeti. Mammalian Biology, 89(1), 90-94.

Anderson, T. M., White, S., Davis, B., Erhardt, R., Palmer, M., Swanson, A., ... & Packer, C. (2016). The spatial distribution of African savannah herbivores: species associations and habitat occupancy in a landscape context. Philosophical Transactions of the Royal Society B: Biological Sciences, 371(1703), 20150314.

Barnosky, A. D., Matzke, N., Tomiya, S., Wogan, G. O., Swartz, B., Quental, T. B., ... & Mersey, B. (2011). Has the Earth's sixth mass extinction already arrived? Nature, 471(7336), 51-57.

Beaudrot, L., Palmer, M. S., Anderson, T. Michael, Packer, C. (2020). Mixed-species groups of Serengeti grazers: a test of the stress gradient hypothesis. Ecology, (In Press).

Beery, S., Van Horn, G., & Perona, P. (2018). Recognition in terra incognita. In Proceedings of the European Conference on Computer Vision (ECCV) (pp. 456-473).

Bonney, R., Shirk, J. L., Phillips, T. B., Wiggins, A., Ballard, H. L., Miller-Rushing, A. J., & Parrish, J. K. (2014). Next steps for citizen science. Science, 343(6178), 1436-1437.

Bonney, R., Phillips, T. B., Ballard, H. L., & Enck, J. W. (2016). Can citizen science enhance public understanding of science? Public Understanding of Science, 25(1), 2-16.

Branson, S., Van Horn, G., & Perona, P. (2017). Lean crowdsourcing: Combining humans and machines in an online system. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 7474-7483).

Burton, A. C., Neilson, E., Moreira, D., Ladle, A., Steenweg, R., Fisher, J. T., ... & Boutin, S. (2015). Wildlife camera trapping: a review and recommendations for linking surveys to ecological processes. Journal of Applied Ecology, 52(3), 675-685.

Cadman, M., & González-Talaván, A. (2014). Publishing camera trap data, a best practice guide. Global Biodiversity Information Facility.

Chen, G., Han, T. X., He, Z., Kays, R., & Forrester, T. (2014, October). Deep convolutional neural network based species recognition for wild animal monitoring. In 2014 IEEE International Conference on Image Processing (ICIP) (pp. 858-862). IEEE.

Chollet, F. (2017). Xception: Deep learning with depthwise separable convolutions. Proceedings – 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017-Janua, 1800-1807. https://doi.org/10.1109/CVPR.2017.195

Cox, J., Young, E. O., Simmons, B., Lintott, C., Masters, K., Greenhill, A., Graham, G., & Holmes, K. (2015). Defining and measuring success in online citizen science: A case study of Zooniverse projects. Computing in Science & Engineering, 17(4) (pp. 28–41).

Dickinson, J. L., & Bonney, R. (2012). Introduction: Why citizen science. In Citizen science: public participation in environmental research (pp. 1-14).

Dredze, M., & Crammer, K. (2008, October). Online methods for multi-domain learning and adaptation. In Proceedings of the Conference on Empirical Methods in Natural Language Processing (pp. 689-697). Association for Computational Linguistics.

Eickenberg, M., Gramfort, A., Varoquaux, G., & Thirion, B. (2017). Seeing it all: Convolutional network layers map the function of the human visual system. NeuroImage, 152, 184-194.

Fegraus, E. H., Lin, K., Ahumada, J. A., Baru, C., Chandra, S., & Youn, C. (2011). Data acquisition and management software for camera trap data: A case study from the TEAM Network. Ecological Informatics, 6(6), 345-353.

Foody, G. M., & Arora, M. K. (1997). An evaluation of some factors affecting the accuracy of classification by an artificial neural network. International Journal of Remote Sensing, 18(4), 799-810.

Forrester, T., O'Brien, T., Fegraus, E., Jansen, P. A., Palmer, J., Kays, R., ... & McShea, W. (2016). An open standard for camera trap data. Biodiversity Data Journal, (4).

Fortson, L. F., Masters, K., Nichol, R., Borne, K. D., Edmondson, E. M., Lintott, C., Raddick, J., Schawinski, K., Wallin, J. (2012). Galaxy Zoo: Morphological Classification and Citizen Science. Book Chapter: Advances in Machine Learning and Data Mining for Astronomy, Chapman & Hall/CRC editors Way, M. J., Scargle, D. J., Ali, K. M., Ashok, N. S., arXiv:1104.5513.

Fortson, L. F., Wright, D. E., Lintott, C. J., & Trouille, L. (2018). Optimizing the Human-Machine Partnership with Zooniverse. Keynote: 6th AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2018). arXiv e-prints, arXiv:1809.09738.

Glover-Kapfer, P., Soto-Navarro, C. A., & Wearn, O. R. (2019). Camera-trapping version 3.0: current constraints and future priorities for development. Remote Sensing in Ecology and Conservation, 5(3), 209-223.

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.

Hampton, S. E., Strasser, C. A., Tewksbury, J. J., Gram, W. K., Budden, A. E., Batcheller, A. L., ... & Porter, J. H. (2013). Big data and the future of ecology. Frontiers in Ecology and the Environment, 11(3), 156-162.

Harris, G., Thompson, R., Childs, J. L., & Sanderson, J. G. (2010). Automatic storage and analysis of camera trap data. The Bulletin of the Ecological Society of America, 91(3), 352-360.

Hines, G., Swanson, A., Kosmala, M., & Lintott, C. (2015, March). Aggregating user input in ecology citizen science projects. In Twenty-Seventh IAAI Conference.

Jennett, C., Kloetzer, L., Schneider, D., Iacovides, I., Cox, A., Gold, M., ... & Talsi, Y. (2016). Motivations, learning and creativity in online citizen science. Journal of Science Communication, 15(3).

Jetz, W., McPherson, J. M., & Guralnick, R. P. (2012). Integrating biodiversity distribution knowledge: toward a global map of life. Trends in Ecology and Evolution, 27(3), 151-159.

Kays, R., McShea, W. J., & Wikelski, M. (2020). Born-digital biodiversity data: Millions and billions. Diversity and Distributions.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems (pp. 1097-1105).

Kosmala, M., Wiggins, A., Swanson, A., & Simmons, B. (2016). Assessing data quality in citizen science. Frontiers in Ecology and the Environment, 14(10), 551-560.

LaRue, M. A., Ainley, D. G., Pennycook, J., Stamatiou, K., Salas, L., Nur, N., ... & Barrington, L. (2020). Engaging 'the crowd'in remote sensing to learn about habitat affinity of the Weddell seal in Antarctica. Remote Sensing in Ecology and Conservation, 6(1), 70-78.

Lepczyk, C. A., Boyle, O. D., Vargo, T. L., Gould, P., Jordan, R., Liebenberg, L., ... & Vaughan, H. (2009). Symposium 18: Citizen science in ecology: the intersection of research and education. The Bulletin of the Ecological Society of America, 90(3), 308-317.

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

Miao, Z., Gaynor, K. M., Wang, J., Liu, Z., Muellerklein, O., Norouzzadeh, M. S., ... & Getz, W. M. (2019). Insights and approaches using deep learning to classify wildlife. Scientific Reports, 9(1), 1-9.

Muneza, A. B., Ortiz-Calo, W., Packer, C., Cusack, J. J., Jones, T., Palmer, M. S., ... & Montgomery, R. A. (2019). Quantifying the severity of giraffe skin disease via photogrammetry analysis of camera trap data. Journal of Wildlife Diseases, 55(4), 770-781.

Norouzzadeh, M. S., Nguyen, A., Kosmala, M., Swanson, A., Palmer, M. S., Packer, C., & Clune, J. (2018). Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. Proceedings of the National Academy of Sciences, 115(25), E5716-E5725.

O'Connell, A. F., Nichols, J. D., & Karanth, K. U. (Eds.). (2010). Camera traps in animal ecology: methods and analyses. Springer Science & Business Media.

O'Brien, T. G., Baillie, J. E. M., Krueger, L., & Cuke, M. (2010). The Wildlife Picture Index: monitoring top trophic levels. Animal Conservation, 13(4), 335-343.

Palmer, M. S., Fieberg, J., Swanson, A., Kosmala, M., & Packer, C. (2017). A 'dynamic' landscape of fear: prey responses to spatiotemporal variations in predation risk across the lunar cycle. Ecology Letters, 20(11), 1364-1373.

Palmer, M. S., Swanson, A., Kosmala, M., Arnold, T., & Packer, C. (2018). Evaluating relative abundance indices for terrestrial herbivores from large-scale camera trap surveys. African Journal of Ecology, 56(4), 791-803.

Palmer, M. S., & Packer, C. (2018). Giraffe bed and breakfast: Camera traps reveal Tanzanian yellow-billed oxpeckers roosting on their large mammalian hosts. African Journal of Ecology, 00, 1–3.

Palmer, M. S., Willis, C., Barry, K., Packer, C., Moe, A., & Wassenberg, D. (2020). Exploring species interactions with "Snapshot Serengeti." CourseSource https://doi.org/10.24918/cs.2020.49.

Parham, J., Stewart, C., Crall, J., Rubenstein, D., Holmberg, J., & Berger-Wolf, T. (2018). An animal detection pipeline for identification. In 2018 IEEE Winter Conference on Applications of Computer Vision (WACV) (pp. 1075-1083). IEEE.

Piper, J. (1992). Variability and bias in experimentally measured classifier error rates. Pattern Recognition Letters, 13(10), 685-692.

Prather, E. E., Cormier, S., Wallace, C. S., Lintott, C., Raddick, M. J., & Smith, A. (2013). Measuring the conceptual understandings of citizen scientists participating in zooniverse projects: A first approach. Astronomy Education Review, 12(1).

Reed, J., Raddick, M. J., Lardner, A., & Carney, K. (2013) An exploratory factor analysis of motivations for participating in Zooniverse, a collection of virtual citizen science projects. In 46th Hawaii International Conference on System Sciences (pp. 610-619).

Rich, L. N., Davis, C. L., Farris, Z. J., Miller, D. A., Tucker, J. M., Hamel, S., ... & Kane, M. D. (2017). Assessing global patterns in mammalian carnivore occupancy and richness by integrating local camera trap surveys. Global Ecology and Biogeography, 26(8), 918-929.

Rosser, H., & Wiggins, A. (2018, October). Tutorial Designs and Task Types in Zooniverse. In Companion of the 2018 ACM Conference on Computer Supported Cooperative Work and Social Computing (pp. 177-180).

Rotman, D., Preece, J., Hammock, J., Procita, K., Hansen, D., Parr, C., ... & Jacobs, D. (2012). Dynamic changes in motivation in collaborative citizen-science projects. In Proceedings of the ACM 2012 conference on computer supported cooperative work (pp. 217-226).

Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., … Fei-Fei, L. (2015). ImageNet Large Scale Visual Recognition Challenge. International Journal of Computer Vision.

Schneider, S., Taylor, G. W., Linquist, S., & Kremer, S. C. (2019). Past, present and future approaches using computer vision for animal re-identification from camera trap data. Methods in Ecology and Evolution, 10(4), 461-470.

Shirk, J. L., Ballard, H. L., Wilderman, C. C., Phillips, T., Wiggins, A., Jordan, R., ... & Bonney, R. (2012). Public participation in scientific research: a framework for deliberate design. Ecology and Society, 17(2).

Simpson, R., Page, K. R., & De Roure, D. (2014, April). Zooniverse: observing the world's largest citizen science platform. In Proceedings of the 23rd international conference on world wide web (pp. 1049-1054).

Spiers, H., Swanson, A., Fortson, L., Simmons, B. D., Trouille, L., Blickhan, S. and Lintott, C. (2019). Everyone counts? design considerations in online citizen science. Journal of Science Communication, 18(1), A04.

Steenweg, R., Hebblewhite, M., Kays, R., Ahumada, J., Fisher, J. T., Burton, C., ... & Brodie, J. (2017). Scaling-up camera traps: Monitoring the planet's biodiversity with networks of remote sensors. Frontiers in Ecology and the Environment, 15(1), 26-34.

Stephenson, P. J., Brooks, T., Butchart, S., Fegraus, E., Geller, G., Hoft, R., ... & McRae, L. (2017). Priorities for big biodiversity data. Frontiers in Ecology and the Environment, 15(3), 124-125.

Sueur, J., Krause, B., & Farina, A. (2019). Climate change is breaking Earth's beat. Trends in Ecology & Evolution, 34(11), 971-973.

Swanson, A., Kosmala, M., Lintott, C., Simpson, R., Smith, A., & Packer, C. (2015). Snapshot Serengeti, high-frequency annotated camera trap images of 40 mammalian species in an African savanna. Scientific Data, 2, 150026.

Swanson, A., Kosmala, M., Lintott, C., & Packer, C. (2016a). A generalized approach for producing, quantifying, and validating citizen science data from wildlife images. Conservation Biology, 30(3), 520-531.

Swanson, A., Arnold, T., Kosmala, M., Forester, J., & Packer, C. (2016b). In the absence of a "landscape of fear": How lions, hyenas, and cheetahs coexist. Ecology and Evolution, 6(23), 8534-8545.

Tabak, M. A., Norouzzadeh, M. S., Wolfson, D. W., Sweeney, S. J., VerCauteren, K. C., Snow, N. P., ... & Teton, B. (2019). Machine learning to classify animal species in camera trap images: applications in ecology. Methods in Ecology and Evolution, 10(4), 585-590.

Topol, E. J. (2019). High-performance medicine: the convergence of human and artificial intelligence. Nature Medicine, 25(1), 44-56.

Trouille, L., Lintott, C. J., & Fortson, L. F. (2019). Citizen science frontiers: Efficiency, engagement, and serendipitous discovery with human–machine systems. Proceedings of the National Academy of Sciences, 116(6), 1902–1909.

Villa, A. G., Salazar, A., & Vargas, F. (2017). Towards automatic wild animal monitoring: Identification of animal species in camera-trap images using very deep convolutional neural networks. Ecological Informatics, 41, 24-32.

Watson, D., & Floridi, L. (2018). Crowdsourced science: sociotechnical epistemology in the e-research paradigm. Synthese, 195(2), 741-764.

Wearn, O. R., Freeman, R., & Jacoby, D. M. (2019). Responsible AI for conservation. Nature Machine Intelligence, 1(2), 72-73.

Wearn, O. R., & Glover-Kapfer, P. (2017). Camera-trapping for conservation: a guide to best-practices. WWF Conservation Technology Series, 1(1), 181.

Willi, M., Pitman, R. T., Cardoso, A. W., Locke, C., Swanson, A., Boyer, A., ... & Fortson, L. (2019). Identifying animal species in camera trap images using deep learning and citizen science. Methods in Ecology and Evolution, 10(1), 80-91.

Zevin, M., Coughlin, S., Bahaadini, S., Besler, E., Rohani, N., Allen, S., Cabero, M., Crowston, K., Katsaggelos,

A. K., Larson, S. L., Lee, T. K., Lintott, C. J., Littenberg, T., Lundgren, A., Oesterlund, C., Smith, J., Trouille, L., Kalogera, V. (2017). Gravity Spy: integrating advanced ligo detector characterization, machine learning, and citizen science. Classical and Quantum Gravity, 34(6):064003.