The potential for AI to revolutionise conservation: a horizon scan

Sam A. Reynolds[1], Sara Beery[2], Neil Burgess[3,4], Mark Burgman[5], Stuart H. M. Butchart[6], Steven J. Cooke[7], David Coomes[8], Finn Danielsen[9], Enrico Di Minin [10,11,12], América Paz Durán[13], Francis Gassert[14], Amy Hinsley[15], Sadiq Jaffer[16], Julia P. G. Jones[17], Binbin V. Li[18,19], Oisin Mac Aodha[20], Anil Madhavapeddy[16], Stephanie A. L. O'Donnell[21], William M. Oxbury[22], Lloyd Peck[23], Nathalie Pettorelli[24], Jon Paul Rodríguez[25,26], Emily Shuckburgh[16], Bernardo Strassburg[27,28], Hiromi Yamashita[29], Zhongqi Miao[30], William J. Sutherland[1]

[1] Conservation Science Group, Department of Zoology, Cambridge University, The David Attenborough Building, Pembroke Street, Cambridge, CB2 3QZ, UK

[2] Faculty of AI and Decision Making, Massachusetts Institute of Technology, 32 Vassar St, Cambridge, MA, USA

[3] UN Environment Programme World Conservation Monitoring Centre (UNEP-WCMC), 219 Huntington Road, Cambridge, CB3 0DL, UK

[4] Center for Macroecology, Evolution and Climate, University of Copenhagen, DK-2100, Copenhagen, Denmark

[5] School of Life Sciences, University of Hawai'i at Mānoa, 3190 Maile Way, St. John 101, Honolulu, HI 96822, USA

[6] BirdLife International, David Attenborough Building, Pembroke Street, Cambridge CB2 3QZ, UK; Department of Zoology, University of Cambridge, Downing Street, Cambridge CB2 3EJ, UK.

[7] Department of Biology and Institute of Environmental and Interdisciplinary Science, Carleton University, 1125 Colonel By Dr., Ottawa, ON, K1S 5B6, Canada

[8] Conservation Research Institute and Department of Plant Sciences, The David Attenborough Building, Pembroke Street, Cambridge CB2 3QZ, UK

[9] Nordic Foundation for Development and Ecology (NORDECO), Skindergade 23, 3rd floor, Copenhagen DK-1159, Denmark

[10] Department of Geosciences and Geography, FI-00014, University of Helsinki, Finland
[11] Helsinki Institute of Sustainability Science (HELSUS), FI-00014, University of Helsinki, Finland

[12] School of Life Sciences, University of KwaZulu-Natal, Durban 4041, South Africa

[13] Instituto de Ecología y Biodiversidad (IEB-Chile), Santiago, Chile

[14] Vizzuality, Calle Fuencarral 123, 28010, Madrid, Spain

[15] Department of Biology, University of Oxford, Oxford, OX1 3SZ, UK

[16] Department of Computer Science & Technology, University of Cambridge, Cambridge, CB3 0FD, UK

[17] School of Environmental and Natural Sciences, Bangor University, UK

[18] Environmental Research Centre, Duke Kunshan University, Kunshan, Jiangsu, 215316, China

[19] Nicholas School of the Environment, Box 90328, Duke University, Durham, NC 27708, USA

[20] School of Informatics, University of Edinburgh, 10 Crichton St, Edinburgh, EH8 9AB, UK

[21] Fauna & Flora International, The David Attenborough Building, Pembroke Street, Cambridge, CB2 3QZ, UK

[22] School of Mathematical Sciences, Lancaster University, Lancaster LA1 4YQ

[23] British Antarctic Survey, NERC, High Cross, Cambridge CB3 0ET, United Kingdom

[24] Institute of Zoology, Zoological Society of London, Regent’s Park NW1 4RY London, UK

[25] IUCN Species Survival Commission, and Provita, Calle La Joya, Edificio Unidad Técnica del Este, Chacao, Caracas 1060, Venezuela

[26] Centro de Ecología, Instituto Venezolano de Investigaciones Científicas (IVIC), km 11 Carretera Panamericana, Altos de Pipe, Miranda 1204, Venezuela

[27] re.green, Rio de Janeiro 22470-060, Brazil.

[28] Rio Conservation and Sustainability Science Centre, Department of Geography and the Environment, Pontifícia Universidade Católica, Rio de Janeiro 22451-900, Brazil.

[29] Ritsumeikan Asia Pacific University, Jyumonji-baru, Beppu City, Oita, 874-8577, Japan

[30] AI for Good Lab, Microsoft, One Microsoft Way, Redmond, WA, USA, 98052

Corresponding author: Reynolds S.A. (sar87@cam.ac.uk)

Glossary

Artificial intelligence The broad field of creating systems that can execute tasks traditionally associated with human cognitive abilities.

Copilot An AI system that improves user productivity by assisting them, surfacing information, suggesting changes, and receiving corrections.

Digital twin A model coupled to and learning from data generated by a physical system.

Federated learning A field of machine learning whereby participants can keep their data private whilst still collaboratively training a model with other participants.

Foundation(al) models A type of machine learning model that is usually pre-trained via self-supervision on extremely large datasets to perform a range of tasks. This includes all data types, including text, images, audio and data from satellite sensors.

Human in the loop A system in which human experts provide evaluation and feedback during the training process to improve accuracy.

Large language models Foundation models pre-trained specifically on textual material

Machine learning A sub-field of Artificial Intelligence, in which a computer is algorithmically trained to perform a task, without following explicit instructions.

Model An abstraction of a system, process, or phenomenon, designed to capture its essential features and relationships. Models can take various forms such as mathematical equations, statistical distributions, algorithms, or neural networks, and are used to make predictions, explain patterns, or simulate behaviour based on input data or parameters.

Multimodal models AI models that are able to work with inputs of different modalities such as images, text, video or audio.

Neural network A computational deep-learning model inspired by the human brain's network of neurons, designed to recognize patterns and make decisions based on data.

Physics-inspired AI An AI approach that incorporates principles from physics to improve model performance, efficiency, and interpretability by leveraging known physical laws and constraints.

Reinforcement learning A type of Machine Learning whereby a system makes decisions in a potentially changing environment and may receive only intermittent signals as to the effectiveness of its decisions in reaching its intended goal.

Retrieval Augmented Generation A technique where an AI model retrieves relevant information from a large dataset to enhance the generation of more accurate and contextually appropriate responses.

Species distribution models Models that use environmental variables and species occurrence data to predict the distribution of a species across geographic space and time.

Supervised learning A method of learning in which a model is trained with both the inputs and corresponding desired outputs/labels.

Self-supervised learning Learning without explicit labels by exploiting some property or structure of the input data.

Transfer learning Where a model developed for one task is reused as the starting point for a model on a second, related task. Leveraging knowledge already gained from solving one problem to a new, similar problem.

Abstract

Artificial Intelligence (AI) is an emerging tool that could be leveraged to identify the effective conservation solutions demanded by the urgent biodiversity crisis. We present the results of our horizon scan of AI applications likely to significantly benefit biological conservation. An international panel of conservation scientists and AI experts identified 21 key ideas. These included species recognition to uncover ‘dark diversity’, multi-modal models to improve biodiversity loss predictions, monitoring wildlife trade, and addressing human-wildlife conflict. We consider potential negative impacts of AI adoption, such as AI colonialism and the loss of essential conservation skills, and suggest how the conservation field might adapt to harness the benefits of AI while mitigating its risks.

Keywords

Artificial Intelligence; Machine learning; Conservation; Biodiversity; Delphi

The intersection between Conservation and AI

Biological conservation science is concerned with understanding the causes and consequences of biodiversity loss, and developing and testing solutions to halt and reverse that loss. Given the urgency of the biodiversity crisis [1], more effective solutions are rapidly required. Biodiversity is complex and multifaceted [2], meaning that understanding status and trends often requires processing large quantities of data, which is time consuming to collect and analyse. In addition, conservation operates in a coupled human and natural system with complex feedbacks between social and ecological components [3] that can be difficult to fully understand and to model accurately [4], obstructing decision-making.

One hope of lifting these barriers is to look to methods of **Artificial Intelligence (AI)** (See Glossary), which have been transformative in many domains such as health care [5] and international security [6]. The key to this success has been the use of **machine learning**, which, given sufficient data, can demonstrate high predictive performance on tasks where good mathematical **models** are lacking [7]. Even in scientific areas that do rely on well understood mathematical models, such as weather forecasting, machine learning can often increase modelling performance at reduced computational cost [8].

Decades of systematic and opportunistic data collection by ecologists and conservation scientists (whether using citizen/community science, expert or remote sensed approaches) has provided a wealth of data, much of which is underutilised [9]. Moreover, increasing availability of novel sensors (e.g., for remote sensing, biotelemetry and biologging) and hardware (including smart phones) for conservation applications are rapidly accumulating new data, and we require analytical tools to keep pace with these advances. Combining these data with the power of AI represents a potentially revolutionising force to increase the effectiveness of conservation approaches, and thus accelerate efforts to the levels needed to meet the targets in the recently agreed Kunming-Montreal Global Biodiversity Framework [10].

Public interest in AI has exploded in recent years with platforms such as ChatGPT and Midjourney capturing people's imaginations through their generation of human-like text and image content. However, in scientific research, AI is not a novel topic and has been adopted by researchers in some form for decades [11]. Nonetheless, the rapid investment in AI within academic and industry research has irrevocably altered the scientific landscape. The functions of AI in scientific understanding can be outlined as follows [12]:

(1) *Providing information through advanced simulation and data representation that cannot be obtained through experimentation*, to reveal properties of a physical

system that are otherwise difficult or even impossible to probe. In conservation, AI systems are making a rapidly growing contribution to providing better information and in some cases, more effective action such as ‘Skylight.global’, which identifies illegal and unreported fishing in near real time from satellite data [13].

(2) *Providing information that expands the scope of human imagination or creativity, such as identifying surprises in data, models and literature.* Recent projects allow researchers to automatically gain ecological and animal behavioural insights from audio [14] and image data [15]. AI technology is also helping to model species distributions from partial observations [16].

(3) *Providing insights to human experts by translating observations into new knowledge.* Platforms such as CAPTAIN [17] take the first steps to allow policy makers to identify optimal areas for protection to preserve species and habitats.

These advances not only improve current practices in conservation (by improving, for example, methods for monitoring trends in species, habitats and threats) but open completely novel areas for research and development. However, new technology also brings challenges. There is growing awareness of how the deployment of wildlife monitoring technologies might have unintended consequences, potentially undermining conservation efforts [18] and wider use of drones, audio recording, camera traps, and electronic tagging and tracking tools, all brings risks of misuse by bad actors. There are also justifiable concerns about the energy use requirements of AI [19], and the way in which use of AI could aggravate existing inequalities and biases [20]. Furthermore, current **large language models** can ‘hallucinate’, generating inaccurate or nonsensical outputs, which are presented as fact.

The aim of this paper is to identify the key areas where AI can help to improve the effectiveness of conservation, as revealed through a horizon scan methodology consulting both conservation and AI experts and also reflect on the challenges remaining to ensure this powerful emerging technology is a force for good in improving the effectiveness of conservation.

Recent developments in AI and machine learning

The potential for AI methods to process and analyse large datasets, identify subtle patterns, and generate novel insights offers promising opportunities for conservation and ecology [21]. An overview of AI and machine learning is available in **Box 1**.

Much of the media excitement in recent years has been the promise of large language models and generative models. These have been shown to be generalisable even with no additional training data for a given task (zero shot learning) [22] and can also reduce the cost of training new models across a wide range of domains through fine-tuning or **transfer learning**. Large language models in particular have seen rapid adoption through chat interfaces such as ChatGPT [23] and assistive "**copilots**" such as Github Copilot [24]. Further, their reasoning capabilities are starting to be used to drive AI Agents [25] that can sense their environment, make decisions and take action.

These **foundation models** [26] take unlabelled input and exploit structure from the data itself to learn. A common and general approach is to hide part of each input example and train the model to predict the missing parts. This process is part of the model's training, not a model evaluation method like cross-validation in conventional statistics. The next-word prediction task used to train Large Language Models such as ChatGPT is one example of **self-supervised learning**.

Freed of the need for labels, models can be trained with large amounts of data and through their self-supervised task can learn generic representations. For example, examination of large visual models trained through self-supervision reveals structures specialised for detecting different classes of objects or animals [27], Once trained, these models can be fine-tuned to a specific task. Consider a visual model that can detect various tree species and leaf shapes. It will require fewer new labelled examples of a specific endangered tree species of interest to be able to classify new examples compared with training a new visual model from scratch [28].

Current successful deployments of large language models often take the form of *copilots*, where AI systems improve user productivity by assisting decision-makers, surfacing information, suggesting changes, and receiving corrections. This approach leverages AI's strengths in synthesising large amounts of information while mitigating risks associated with potential errors or biases. However, challenges remain, including the need to address biases in training data and the importance of human oversight in critical applications.

Box 1 - ***An overview of AI and machine learning***

While AI covers the broad field of intelligent software systems, today the term is most often used to refer to systems that implement Machine Learning (ML) algorithms. An ML algorithm is one that incorporates a data structure (called a model) obtained as a compression of ’training’ data, usually with the aim of approximating probabilities of interest.

Neural networks are the most well-known example of such models, and the field of Deep Learning has shown outstanding success in applying large-scale neural networks to many previously intractable tasks in image and audio processing, Natural Language Processing (NLP) and other areas.

In its simplest form, the training of ML algorithms is supervised. This means that the training data comprises (usually human-labelled) pairs such as image pixels and type of object in the image or audio spectrogram to a bird species. Example applications include predicting taxonomic identity from a photo or audio recording of an animal, or predicting land cover class from remote sensing tiles.

The process of generating trustworthy labelled data for **supervised learning** is, and remains, labour-intensive. However, there are powerful variants of the supervised approach: the ’training signal’ on which model training depends need not always come from human-generated labels. It may come from structure observed within the data itself (for example, next-term prediction in sequential data such as text) or between data sets (for example, machine translation models learned from comparison of many parallel texts). This is commonly referred to as self-supervised learning. It may also come from **transfer learning**, where a model trained on one task (for example, recognising road traffic objects) is fine-tuned with limited data to solve another task (such as alerting to the presence of a pedestrian).

Reinforcement Learning solves decision problems (such as game playing) and takes its training signal from direct exploratory interaction with its environment. It learns a decision policy as an ML model, often a neural network, and can converge to super-human skill levels (for example in Chess, Go and some video games).

Sequential models, in particular, have reached a high degree of sophistication in Large Language Models (LLMs) (such as ChatGPT), which employ a combination of large-scale neural networks and transformer architectures and can be further fine tuned with reinforcement learning.

Finally, the idea of transfer learning can potentially be taken a long way with the proposal of Foundation Models (FMs) [26]. The idea here is to invest one-off (for an application domain of interest) in very large-scale, ideally self-supervised, models. The resulting FM can then be re-tuned at relatively low cost for specific predictive tasks in the domain. However, it is important to note that the current success of foundation models is derived from their access to large, high-quality datasets for training, which are often manually annotated in an expensive and time consuming process and therefore are not based purely on unsupervised learning.

Scanning the horizon for future applications of AI to improve conservation

To identify the areas where AI has considerable potential to revolutionise conservation, we applied a modified Delphi technique to select and rank suggestions from experts [29,30] (Box 2). This technique maintains the transparency, repeatability, and inclusivity of the process [31]. Participants in the Horizon Scan were selected based on consultation of the professional networks of the organising committee and internet searches, with an attempt to produce a mix of subject area expertise and geographical representation. The organising committee produced an initial document summarising recent advances in AI, the needs of conservation, and the main areas where AI is relevant to conservation. This provided a grounding in AI for conservation experts who may have had little exposure to AI advances, and to give AI experts an understanding of the interests of conservationists to facilitate discussion at the workshop. Authors were asked to suggest two to eight potential developments, each with a short explanation. In some cases, the authors gathered ideas from within their organisations, further expanding the sample of experts consulted and the geographical spread of idea generation.

We recognise there may be limitations to ideas generated in the horizon scanning process and that a different group of experts may identify a different set of ideas. Our methodology of inviting participants from a range of subject backgrounds and global regions and asking them to canvass their network of colleagues and collaborators, aims to identify as broad a set of issues as possible and limit bias towards a particular discipline or study area. We note also that Sutherland et al., 2019 [32] reported no significant correlation between participants' areas of research expertise and the top issues selected in a horizon scan conducted in 2009. Therefore, horizon scans do not necessarily represent ideas that reflect the expertise of participants.

Box 2 – ***Conducting the horizon scan***

In July 2024, a panel of 27 experts accepted the invitation to participate in the Horizon Scan, composed of 18 conservation scientists and 9 computer scientists and AI specialists. Some experts sit across both areas of expertise. The experts are authors of the paper.

Participants were required to submit between two and eight ideas detailing how AI could be leveraged in conservation across a range of applications, including image and audio recognition; eDNA analysis, modelling and data interpretation and integration. Participants canvassed their networks and colleagues to broaden perspectives captured in the Horizon Scan.

Figure I provides an overview of the process used to identify and score ideas. For this Horizon Scan, 104 ideas were submitted by participants for consideration (Table S1, Supplementary information). Participants then confidentially and independently reviewed each idea submitted, and ranked them by assigning each idea a score of between 1 (least significant) and 1000 (most significant). The large range of scores was simply to make it easy for scorers to rank items, the scores themselves were not used. Participants assigned a single score to each idea, subjectively combining two criteria: (i) its potential to have an important impact on the field of conservation science and/or practice, and (ii) its likelihood of coming to fruition. Notes were added by participants to provide further information as to whether ideas should be retained for the second round. To counter the effect of voter fatigue [33], participants were randomly assigned one of three lists of ideas presented in different orders. Ideas were clustered into broad topics (audio, images, reviewing literature, modelling, data, generating text, negative consequences, citizen science, eDNA, remote sensing, robotics, society and other) by the lead author. Participant scores were then converted into ranks (1 - 104), with the highest score denoting the highest rank. The median rank across all participants for each idea was calculated, with the top 30 median ranked ideas brought forward for discussion in the workshop. All rank calculations were conducted in R Statistical Software [34], the code is available via GitHub [35]. Two ideas in the top 30 were deemed by the organising committee to be similar to two other suggested ideas, and so were presented together in the shortlist. This led to 28 distinct ideas being shortlisted.

A meeting was held with 13 participants attending in person and 12 participants attending online. Two participants were unable to attend the workshop, but participated in the initial long-list scoring and preparation of the manuscript.

In order to account for participants attending from different time zones, the meeting was held in three sessions. Session 1 with 25 participants in attendance discussed audio, reviewing literature and generating text and Session 2 with 25 participants discussed images and remote sensing, including a focused discussion on the potential negative consequences and pitfalls of AI. The final session with 24 participants focused on modelling and data followed by a discussion on how the field of conservation may have to adapt to take full advantage of AI.

Each idea was discussed for 10 minutes. After each discussion, participants scored each idea on a scale of 1 - 1000 (low to high). High scores denote the most significant ideas, low scores the least significant. Each participant's scores were converted into ranks at the end of the workshop (the highest scoring idea being assigned rank 1), and the 20 ideas with the highest median ranks, and therefore the most significant ideas, were revealed.

Not all online participants could attend the full workshop. To incorporate their partially scored lists with the lists of participants who scored all ideas, we computed a ‘division factor’. The division factor is the ratio of total number of ideas to the number of ideas scored, plus one (Eq. 1):

$$DF = TI/(SI + 1)$$

Eq. 1

Where DF is the division factor, TI is the total number of ideas and SI is the number of ideas scored by the expert. The partially scored lists were then ranked (e.g. if only 10 ideas were scored, they were assigned ranks 1 through 10, with rank 1 being the most significant idea). To calculate the appropriate final rank for integration with the fully ranked score sheets, we used the following equation (Eq. 2):

$$FR = PR \times DF$$

Eq. 2

Where FR is the final rank, PR is the rank given when all partially scored ideas are ranked and DF denotes the division factor.

This creates a buffer in the ranking scale that allows for the possibility that ideas unranked by a participant could potentially be ranked higher than their partially ranked ideas if they were to be evaluated. This adjustment expands the range of the partial rankings to align with a full ranking scale, leaving space for potentially unranked items between them. Ideas that were not scored by a participant, did not have a rank imputed and were left blank. Therefore, participants only contributed to the ranking of items they were present to score.

Two ideas were related to identifying the expansion frontiers of human disturbance and it was proposed that these ideas should be amalgamated into a single idea. Additionally, there was a tied rank at 20, and it was decided both ideas should be included, therefore 21 ideas are included in the final list.

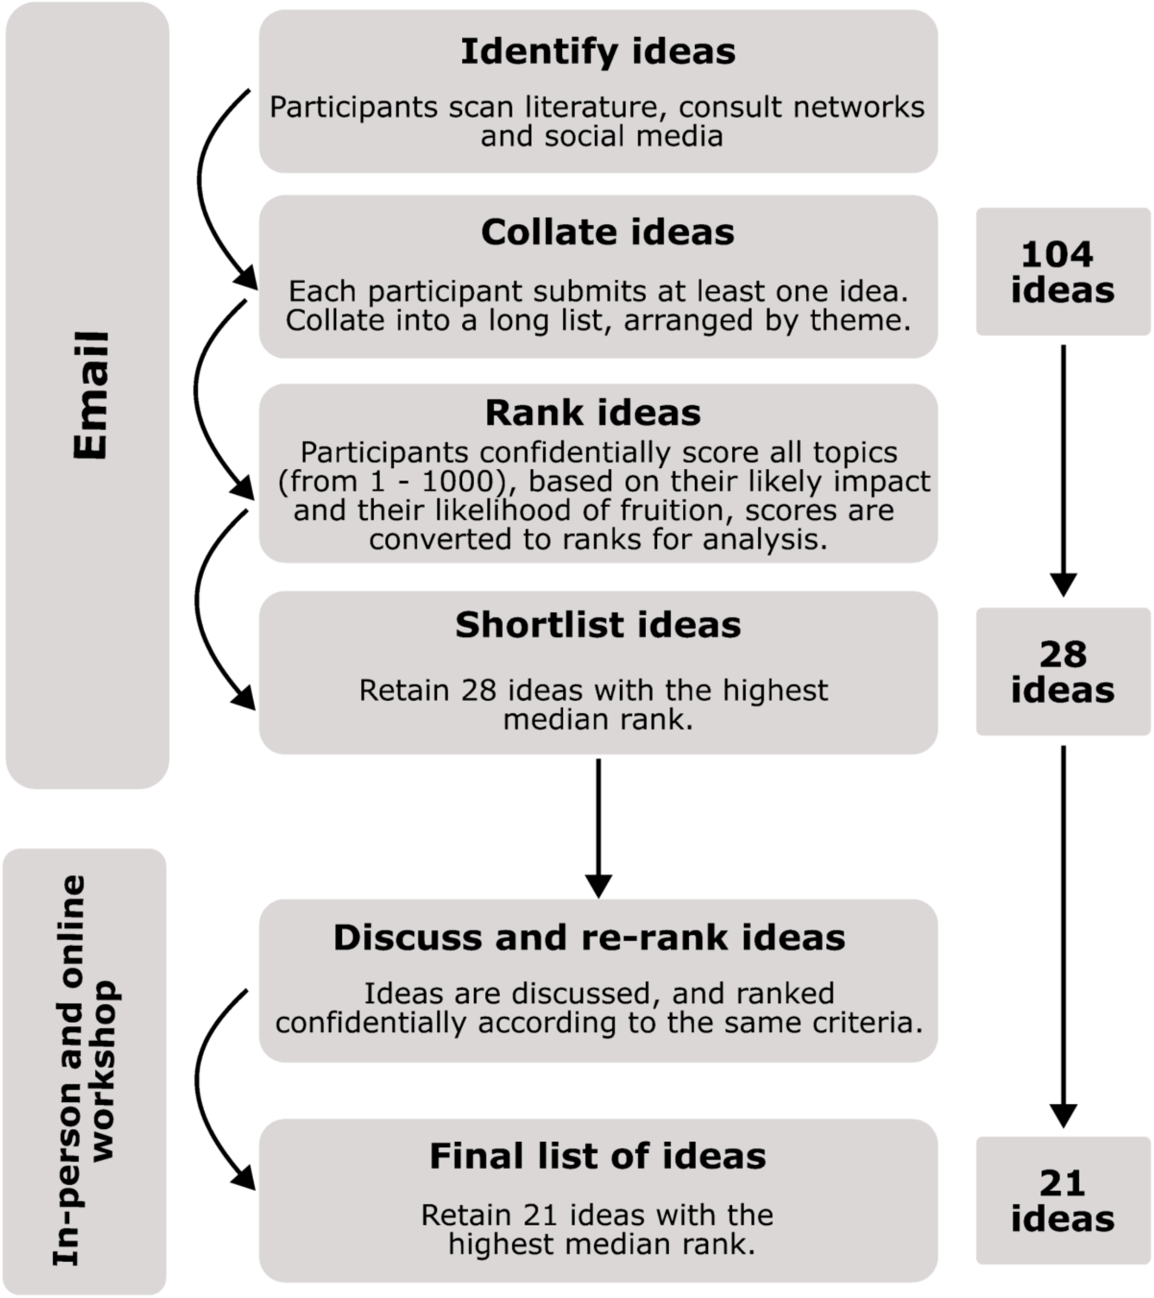

Figure 1. *The process used and the resulting filtering of ideas to produce a ranked list of how AI could revolutionise conservation effectiveness. Ideas were generated and ranked by a panel of 27 experts (18 conservation scientists and 9 computer scientists and AI specialists).*

Potential applications of AI in conservation

In the following we present the 21 ideas where AI has considerable potential to revolutionise biological conservation, (Figure 2). The ideas are presented in thematic groups, rather than in rank order. The full list of 104 ideas is included in the supplementary information (S1). Ideas within the full list vary in their detail, relevance and completeness, but may provide a useful insight for additional applications of AI within conservation. They include ideas relating to the use of robotics, citizen science and understanding peoples’ connections to nature.

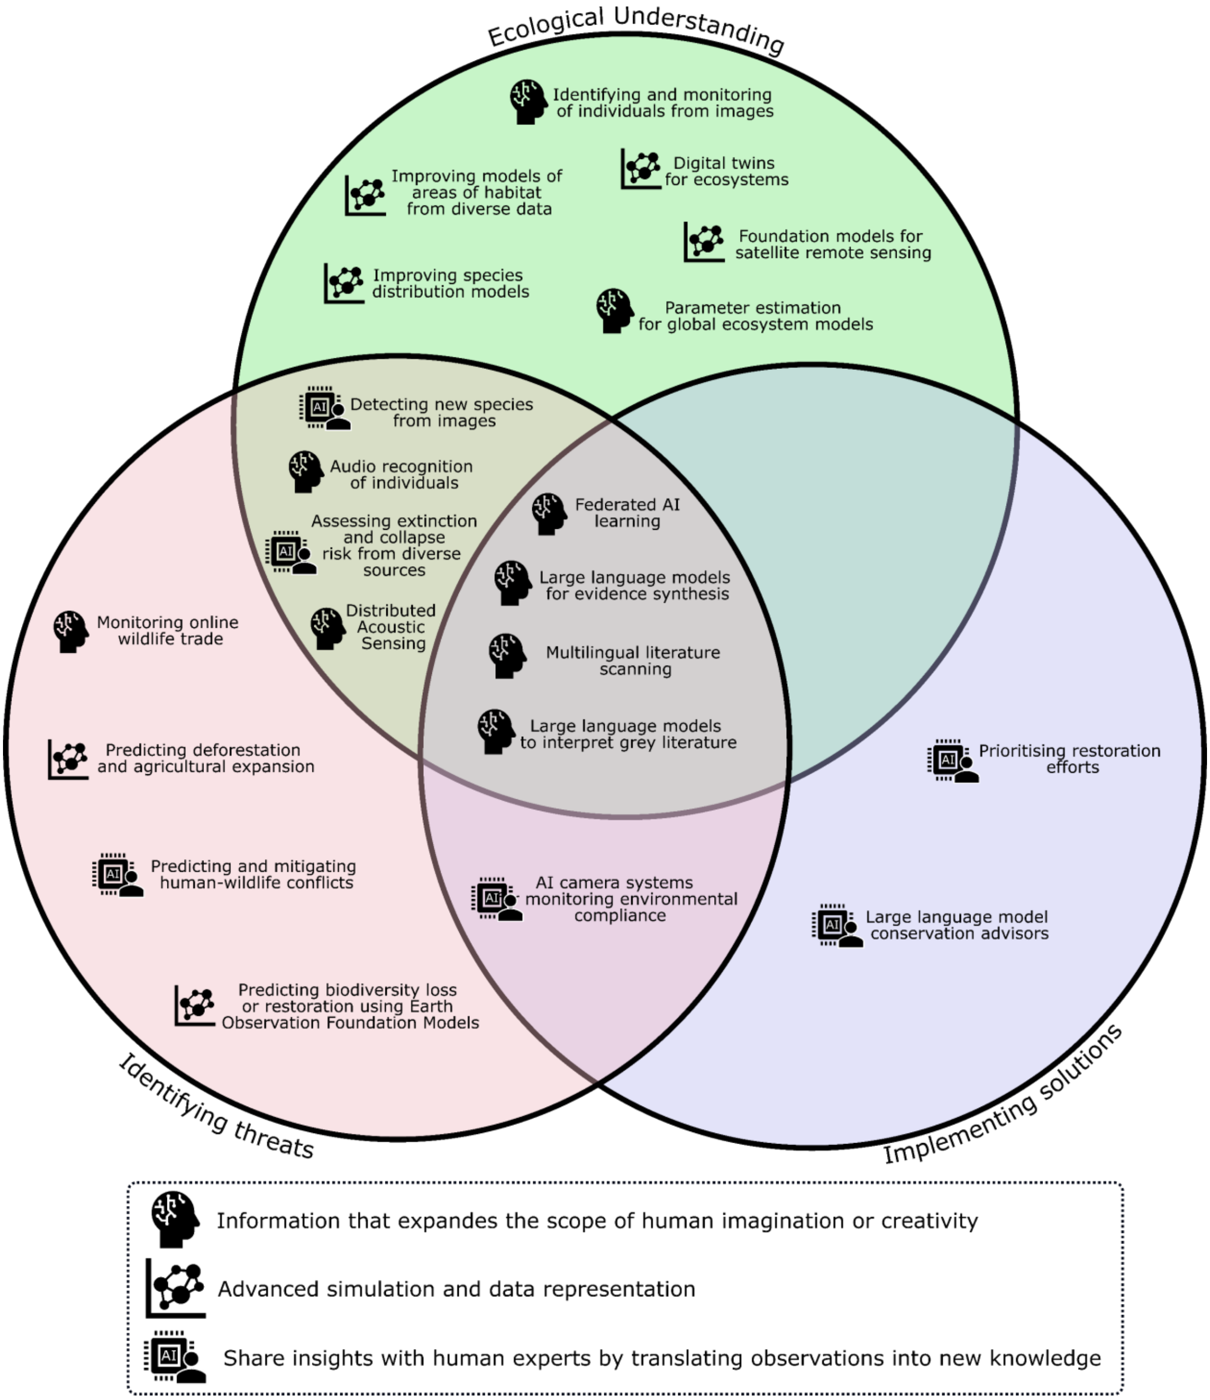

Figure 2. *Places the 21 AI applications identified in the Horizon Scan into a framework adapted from Krenn et al.[12]. Each circle represents a different area of study within conservation science: Ecological understanding, identifying threats and implementing solutions, while each icon represents a different way AI contributes to scientific understanding.*

Interpretation of images collected by ground-based sensors

Automated species recognition from images collected by devices like camera traps and mobile phones is well advanced and available through applications, such as iNaturalist [36], Pl@ntnet [37], ObsIdentify [38] ,Google Lens [39] and Merlin [40]. However, improvements in AI will enable substantial acceleration of its implementation and application, including real-time identification systems that can send alerts, for example, when specific species, large numbers of individuals, or threats are detected [41]. Scaling up image acquisition could be achieved through citizen science, community-based monitoring, harvesting images from social media, or repurposing existing datasets (e.g. Google Street View). This supports applications, such as mapping species distributions and range extensions, monitoring the establishment or spread of invasive alien species, and detecting and identifying illegal imports of traded species at customs [42].

***Identifying and monitoring of individuals from images*.** Scaling up automated recognition of individual animals from images could enable more widespread and accurate assessment of population sizes, for example by using mark-recapture [43] leading to more accurate assessment of status and trends and other opportunities such as quantifying home ranges and identifying movement patterns.

***Detecting new species from images*.** AI workflows have recently been developed for the detection and confirmation of unknown species identity from images, also known as novel category discovery [44]. This technique could accelerate the documentation of ‘dark diversity’, especially in combination with DNA analysis. These approaches could be used on both images captured from the field or from digitisation of museum collections [45].

AI camera systems for monitoring environmental compliance. AI systems for monitoring compliance could be as broad as measuring biodiversity gains promised by developers, analysing water quality measurements from treatment plants, or detecting illegal deforestation from satellite imagery. An obvious candidate area would be monitoring compliance with measures to mitigate commercial fishery bycatch. Bycatch of seabird species are results in mortality rates driving some species towards extinction [46]. Regional Fisheries Management Organisations already require implementation of measures to counter this [47]. However, detecting compliance relies on boat-based observers, which is resource intensive and dangerous. AI-enabled on-board camera systems can facilitate monitoring seabird presence, interactions with fishing gear (e.g. cable strikes), bycatch (i.e. ensnared birds) and use of mitigation measures, enabling safer, cheaper and more reliable monitoring of compliance.

Audio

AI is already being used extensively to identify species from audio recordings, principally for birds and bats. For example, using machine learning, BirdNET can currently identify about 3000 of the world's most common bird species [48]. However, there is considerable potential for scaling up to cover the remaining 75% of bird taxa, and other vocalising animals, including invertebrates. In combination with distribution of autonomous recording units or other devices running identification algorithms locally, this could transform our ability to monitor biodiversity [49]. AI is already being used to assess soundscapes to generate insights into the state of ecosystems, such as coral reefs [50] and tropical forests [51].

Audio recognition of individuals. In some cases, it is already possible to distinguish individuals within a species from their sounds [52]. There is likely to be considerable opportunity to extend this given that (a) voice recognition in humans is well-understood (using 'voiceprint' vectors) [53], and (b) many species are observed to recognise individuals from their vocalisations [52]. Potential applications include estimating population sizes, tracking individual movements, and quantifying dispersal.

Distributed acoustic sensing. Novel AI-enabled technologies for detecting sounds in the marine environment could transform monitoring of marine fauna. For example, Distributed Acoustic Sensing (DAS) can detect sounds in real time using the dense network of fibre-optic telecommunication cables covering both deep ocean and coastal areas worldwide. It could be used for audio recognition of marine fauna, such as cetaceans, seals and fish, as well as human activities (shipping and deep-sea mining) that may impact them [54,55]. This would enable monitoring even in the most remote areas or for oceanic features that are difficult to access, such as seamounts. Similarly, scaling up application of hydrophones with built-in low-powered AI audio classifiers could revolutionise our understanding of the distribution and abundance of marine animals [56,57].

Satellite and airborne remote sensing

Machine learning is commonly used to interpolate data from satellite sensors to classify ecosystems and track degradation and land use change [58]. There is also a rapidly advancing application of this technology to identify species from space [59,60]. Increasingly, deep learning approaches are also considered to detect patterns not easily picked up by conventional approaches [61]. The planet’s surface can be obscured by clouds and imagery quality is affected by illumination angle and atmospheric effects, requiring the inclusion of complex preprocessing steps; AI is more effective than conventional approaches at making these corrections [62].

***Foundation models for satellite remote sensing*.** Pre-trained foundation models (see Box 1) can facilitate the monitoring of areas with restricted availability of training and validation data [63]. However, the fine tuning of these models for specialised tasks can be highly challenging [64]. AI can facilitate the fusing of multiple data modalities, such as optical and radar imagery, as well as the use of time series analyses for improved land cover classification and threat detection [65] and monitoring [66,67].

Predicting biodiversity loss or restoration using Earth Observation Foundation Models. Global biodiversity monitoring uses a series of metrics, but there are often challenges in keeping agreed metrics up to date, and as close to real time as possible. AI could facilitate the automated production of earth observation-based datasets used in policy and business applications, such as the human footprint index or the Biodiversity Intactness Index, which assesses the degree to which terrestrial ecosystems are affected by land-use change and intensification. However, applying this will require some care, to avoid error propagation and its associated impact on biodiversity science [68].

Combining Datasets for New Insights

Primary observational datasets are the “fuel” for machine learning algorithms, but AI also creates opportunities to improve higher-level datasets by combining diverse input data. For example, training models on photo catalogues together with remote sensing images has facilitated the development and validation of land cover change products like Dynamic World [69]. Gains are likely when bringing multiple datasets together. For example, the CAPTAIN project [17] uses reinforcement learning to train models to identify areas for conservation prioritisation by integrating species distribution data, anthropogenic disturbance data, human population densities, phylogenetic diversity, land use data and climate change projections, while another project combines similar datasets to estimate extinction risks for 89% of known tree species [70].

Federated AI learning. Several biodiversity data systems contain "data islands" that cannot be directly shared for the purposes of training AI models (e.g. critically endangered species' locations, or illegal hunting and trade data). Using expertise in managing sensitive data from other fields, such as healthcare [71], it may be possible to train a distributed model for biodiversity characteristics which is an accurate representation of the underlying raw sensitive data and delivers real time insights on the status and trends in biodiversity, pressures, responses and benefits. If possible, this would allow the unification of numerous distributed databases into models that facilitate decision-making without revealing sensitive raw data. However, much work remains to ensure these models can be made robust to attacks that leak the original sensitive training data [72].

Monitor online wildlife trade. Computer vision (i.e., automated processing and understanding of images), natural language processing (i.e., automatic processing of textual content), and **multimodal** machine learning can be used to understand where, when, how, why, and what species and wildlife products are being traded on which online platforms [73]. For example, a recent study of global trade in chameleons, used multiple lines of evidence to understand trade patterns and the impacts of trade bans [74], but automation and AI-assisted insights could have dramatically speeded up the process. Research using these methods must prioritise adherence to the highest standards of data privacy and protection [75].

***Improving models of areas of habitat from diverse data**.* Our understanding of species' areas of habitat (AoH) currently relies on extremely patchy sampling, estimated range maps, and incomplete knowledge of species' ecology and habitat preferences. AI could improve AoH maps by integrating occurrence data on specimen locations, habitat and cohabitation preferences, environmental and ecological data, remote sensing data and spatial datasets on human impacts. These approaches are already being tested [76] and could move from relatively simple correlative AoH models (e.g. masking estimated species’ range maps with elevation and climate) to more sophisticated covariates describing species habitat selection across diverse ecosystems, with deep learning making more comprehensive predictions of species’ ranges and, perhaps, also populations [77,78].

Predicting and mitigating human-wildlife conflicts. Human-wildlife conflict, where people themselves, or their property, crops or livestock, are harmed by wildlife is a serious problem for some human communities and, when resulting in retaliatory killing, is an important threat to some species [79]. AI-powered cameras are already used in various countries to warn communities when particular mammal species are nearby, for example elephants in Africa or tigers in Asia [80]. Similar technology can also track humans within landscapes set aside for animals [81]. However, there are important ethical and privacy concerns with such technology which potentially may exacerbate conflict between conservation and local communities [18]. AI-enabled systems combining short- and long-term data on factors correlated with human-wildlife conflicts will provide more accurate predictions and opportunities to take action earlier.

Modelling and causal inference

Using data sets, usually in combination, there is the potential to apply AI to build system models that can be used to test hypotheses or to build credible counterfactuals for the evaluation of conservation policy and practice. Uses could include simulating the outcomes of policy decisions and predicting projections based on a range of data sources. Large-scale modelling of biodiversity outcomes from different policy interventions, using General Ecosystem Models, are increasingly available [82,83]. However, they are time consuming and resource intensive, and are typically used to inform international policy formulation - for example the Global Biodiversity Framework [84], or the EU new green deal [85]. AI may offer opportunities to accelerate these large-scale modelling efforts and observe emerging patterns important for policy decision-making.

Digital twins for ecosystems. **Digital twins** are used in engineering and earth sciences [86], but they can also be applied to the problem of predicting outcomes of conservation interventions. Digital twins for ecosystem modelling and prediction could be developed, potentially paired across multiple domains (e.g. physics, fish demography and seabird ecology) to understand and model the likely outcomes of different interventions [87,88]. Moreover, where we seek to apply AI to optimise intervention policies (as in CAPTAIN), simulation from a digital twin can provide an additional training signal.

Prioritising restoration efforts. Restoring degraded habitats is essential in order to avert extinctions, recover populations and minimise climate change. AI could help us direct such efforts more efficiently, identifying the most important locations for restoration given habitat loss and degradation to date, the distribution of species and ecosystems, projected climate change, shifts in energy production, food production and human population distributions. Artificial neural networks have been previously explored for prioritising areas for wetland restoration (utilising multiple inputs including elevation); soil texture and permeability maps; protected areas for waterbirds; and for likelihood of dust storms and urbanisation [89].

Improving species' distribution models. **Species distribution models** (SDMs) are valuable for conservation decision making such as assessing land use change impacts. Using species' distribution records and remotely sensed data, SDMs produce maps of the (relative) probability of occurrences. Machine learning models are well-suited to integrate heterogeneous and multi-fidelity datasets efficiently, handling missing data, noise, and non-linear relationships; they can automatically learn relevant features from raw spatial data and uncover complex patterns and interactions. However, current limitations include the scarcity of labelled data (especially for plants and non-vertebrate animals), spatial biases in data, temporal mismatches between field and remotely sensed data, data accuracy, over reliance on abiotic conditions, limited considerations for biotic interactions and overfitting. New AI approaches (e.g. **physics-inspired AI** and xAI) and rapidly increasing citizen science datasets may help address these problems.

Predicting deforestation and agricultural expansion. Deep learning can be used to analyse satellite imagery and other geospatial data to predict where land use change is most likely, learning spatio-temporal patterns from the remote sensing data [90] rather than the mechanistic modelling approaches currently used. Such models could be used to predict where deforestation would have occurred under business-as-usual scenarios to allow better estimates of the effectiveness of interventions aimed at slowing land conversion (such as zero deforestation commitments or REDD+ projects). Similar techniques can also be applied to training predictive models from time-series maps of agricultural expansion, which has seen rapid acceleration in the past two decades [91].

Parameter estimation for global ecosystem models. Global ecosystem models (GEMs) simulate the dynamics of life on Earth, including the interactions between plants, animals, and the environment. GEMs are valuable for modelling the impacts of climate change on nature. However, predictions of dynamic GEMs do not align well with field measurements (e.g. of fluxes of greenhouse gas); they fail to make reliable predictions of large-scale fires and droughts, and they rarely include plant-animal interactions. Many processes in dynamic GEMs are represented by semi-empirical semi-theoretical equations. AI optimisation techniques could automate the calibration of the parameters in these models across multiple scales. Physics-inspired AI techniques, such as reduced-order modelling, could potentially find low-dimensional representations that capture the essential dynamics while being much faster to compute [88].

Reviewing literature

There is considerable potential for using AI to improve the efficiency of extracting, screening, and collating literature for conservation [92] as is already underway to some extent within medicine [93]. For example, creation of systematic reviews could be automated so they can be carried out in days rather than months or years, enabling more timely guidance drawing on the full range of evidence.

Large language models for evidence synthesis. Evidence syntheses for biodiversity conservation are key for effective decision-making, yet are challenged by increasingly time-consuming tasks, a broad evidence base, and persistent underfunding. Moreover, most evidence syntheses produced in the conservation space are static and fail to incorporate new evidence as it is generated. AI has the potential to be harnessed to identify relevant (new) evidence and integrate it into the existing evidence base and synthesise key messages rapidly (i.e., living evidence syntheses; [94]). Doing so will ensure decision makers have access to the best available evidence to guide them.

Multilingual literature scanning. Large Language Models (LLM) enable multilingual literature searches for non-English evidence [95], which is key given that non-english languages typically dominate in areas of most conservation concern [96], not to mention the growing quantity of non-English evidence [97].

Large language models to interpret grey literature. LLMs can be used for identifying grey literature and text matching to cluster relevant evidence. They can accelerate evidence synthesis and make it more timely, equitable and inclusive in terms of the evidence base and perspectives considered.

Assessing extinction and collapse risk from diverse sources. Assessments of extinction risk for species and collapse risk for ecosystems currently rely on manual compilation of information on species’ population sizes, distributions and trends, as well as trends in area, biotic and abiotic factors for ecosystems, to produce parameter estimates that are applied to IUCN Red Lists criteria. AI could accelerate and expand the process by scanning the scientific literature and other online materials to locate relevant information in published and unpublished sources in all languages, including real-time land-cover change. This would substantially accelerate and improve the process of assessing extinction and collapse risk while significantly reducing costs. More ambitiously, extinction and collapse risk could be estimated directly in the future by combining relevant literature-derived parameter estimates with spatially explicit predictive models informed by remote sensing.

Generating text

The success of Large Language Models (LLMs) in creating novel text, images and videos is well recognised. There has long been a distinction between the use of generative AI for error tolerant commercial applications, such as marketing and social media, and for safety-critical or politically sensitive applications in specialist use cases. However, their recent success has led to their application being debated in medicine [98]. Their ability to generate personalised outputs and recommendations based on a synthesis of available evidence has led to suggestions that they could play a role as policy advisors [99]. This could include synthesis of multiple data sources for a specific location to provide bespoke management recommendations, or application to complex numerical datasets to provide non-specialist text summaries for decision-makers.

Large language model conservation advisors. LLMs could synthesise evidence of impact for different conservation management options [92] and combine this with data on species, land-use, or socioeconomic factors for a specific context, to produce easily understandable and evidence-based information for practitioners and policy-makers. The inclusion of **human in the loop**, whereby human experts are involved in providing feedback and evaluation of advisor outputs during model training, would be essential to limit biases and hallucinations [99], as would fine-tuning and **retrieval augmented generation** (RAG) based on curated, robust evidence and data.

Negative consequences of AI

AI will undoubtedly lead to great changes in conservation in the coming years. While many will be positive there is also the potential for unintended negative outcomes. These need to be understood and mitigated.

There is a danger that AI could have a polarising effect on conservation and conservation funding. If AI-supported conservation becomes the ‘gold standard’, as people are impressed by the novelty, or claims of large impacts or the promise to revolutionise conservation then we may risk seeing a shift in funding and leadership in conservation. Distorting away from conventional experimentation and on the ground practice, which already struggles to attract resources, towards financially wealthy institutions which are able to undertake AI work. This could be especially true in conservation research where the sort of grounded field-based and participatory studies, which have a role in advancing understanding and local ownership, may become ever more difficult to fund. This could undermine efforts to improve the diversity of voices, knowledge and approaches in conservation. Hence it is important funders recognise the importance of supporting a spectrum of conservation research and practice embracing both conventional and AI approaches.

There is a fear that we could see a loss of essential skills in conservation if people in the field pivot towards implementing AI over conventional techniques. Retaining species, ecosystem and community experts will be integral to creating reliable AI technology. Data is the fuel of AI, and data collected by conservation experts will be essential to producing better models, and this critical data-collection work must be appropriately recognised. Moreover, information itself, however it is obtained, does not lead to better conservation, it is important that recommendations are designed to work in the real world and are not detached from the social and ecological reality on the ground.

AI colonialism is a central concern, with data potentially being extracted from the Global South to data centres in the Global North, where it is used to train AI models to issue AI-driven mandates to the Global South on how land and resources should be managed. This would undermine efforts of the conservation community to address the colonial legacy of contemporary conservation and to recognise the importance of indigenous rights and voices. Furthermore, there is a risk that AI contributes to a militarization of conservation, with computer systems, developed far from the area concerned, identifying infractions and triggering enforcement without understanding of local context. Given the importance of local perceived legitimacy to compliance with conservation rules [100], this could create or exacerbate conflict. To address digital inequalities and injustices and to produce less biased, fairer and more robust information for conservation actions there is a need to integrate epistemic feedback loops into black box models. This can be achieved by leveraging human-in-the-loop designs along with political agencies and democratic decision-making [101].

The coupling of AI models containing inherent biases, propagates bias and error. This can be exacerbated by ‘AI pollution’, whereby the outputs of biased models are used to train new models. This may lead to increasingly poor representation of understudied species or ecosystems, potentially pushing people away from considering these understudied areas if we over-rely on AI for decision making and needs to be explicitly considered in any AI-based approaches.

A lot of the promise of AI is bigger, better models. However, the bigger the models get, the more expensive they are to run, in terms of computation, bandwidth requirements, power and expertise. Already there is significant concern about environmental externalities of the power consumption and cooling requirements required to run AI infrastructure, which look likely to increase. The cost of using these models may push AI beyond the current resources (financial and human) of conservation. However, in conservation there is less call for large generalist AI models, like Chat GPT, rather smaller more specialised models for specific use cases. It would be helpful for the conservation community to adopt a code-of-practice to address sustainability considerations associated with AI-research.

Finally, although AI models to further the effectiveness of conservation are built with good intentions, it must be remembered that they could also be used by bad actors. For example, image and audio tools used by conservationists to track and locate endangered or protected species can equally be exploited by poachers or be co-opted by government regimes to monitor the movement of people. Similarly, remote sensing data could be contaminated or poisoned by malicious private companies to exaggerate restoration efforts in order to attract greater revenue from carbon or biodiversity credit schemes. It is important that significant steps are taken to protect data, and ensure tools are used only for their intended purpose.

How might conservation be organised to take advantage of AI and reduce problems?

To take advantage of the potential capability of AI to identify new patterns, generate accurate predictions, run more accurate simulations drawing on diverse data sources, and help decision makers better assess the efficacy of potential interventions, society and the field of conservation will have to adapt.

Reproducibility of results, one of the pillars of the scientific method , becomes challenging when there is insufficient documentation, limited access to underlying code and data and lack of understanding of how AI tools reach conclusions, which makes it difficult to scrutinise, verify and replicate experiments [11]. Improved literacy around AI will help to counteract its misinterpretation and recognise its limitations. Improved AI literacy will also be essential to ensure sufficient peer review of papers using AI.

To counteract the influx of outputs from generative AI models it is important society identifies a way to segregate human-produced content from that produced by AI. One such example could be digital kite marks, where individuals, organisations and institutions are assigned digital signatures that can be applied to authenticate a document or dataset’s human-generated origin.

Interdisciplinary collaboration between AI researchers and experts and domain subject experts within natural science, social science, and Indigenous and local knowledge, will be central to building accurate AI models. There is a danger of a decoupling between the creators and users of AI tools. Hallucinations in outputs are more likely to be detected by subject experts. This also extends to ensuring there is consultation between developers of AI tools and the experts who manage underlying datasets to ensure data limitations are understood and are not used inappropriately, as inaccurate use of data can lead to misleading advice for decision makers [102].

Moreover, researchers should not incorporate advanced AI techniques at the cost of appropriate conventional methodologies. It should be remembered that AI is only one of the tools in the toolbox.

In conservation there are additional barriers to broad uptake, which will need to be considered. There is a large inequality of data availability between the Global South and the Global North. This needs to be acknowledged and understood to avoid embedding biases in AI models predicting outputs on a global scale. This inequality extends also to ecosystems and species where there is a disparity in data availability, for example, data on birds in the Amazonian rainforest are more numerous than for marine bryozoans in Antarctica.

Access to infrastructure and resources to be able to train and run these models is currently limited and is disproportionally located in the Global North and held by elite academic institutions, governments and technology companies. Training foundation AI models is expensive and requires access to a large amount of computing power, as well as IT expertise, high internet bandwidth and reliable power supply. Methods of access to these resources, as well as overcoming language barriers, need to be thoughtfully considered to widen participation in AI research by conservation practitioners, researchers and governance institutions including civil society associations, NGOs and government conservation managers based in the Global South.

The role of expert data labellers will be crucial to maintaining data quality, preventing biased outputs, and addressing data fabrication. Many species and ecosystem experts are based in the Global South, and it is imperative that their skills and labour are equitably employed and recognised. It also highlights the value of conventional conservation expertise in the production of these models. Moreover, concerted effort is needed to substantially strengthen the connection between the outputs of AI models and the ecological and social realities of implementing and sustaining conservation actions on-the-ground [103].

While AI systems have potential to make monitoring of infractions of conservation rules cheaper and therefore more widespread, it cannot by itself overcome social and governance challenges which strongly influence compliance. For example, just because it is possible for regulators to put AI-enabled cameras on board ships to monitor bycatch, how useful that technology is will depend on compliance and follow-through.

Concluding remarks

Following a horizon scan methodology, we identified 21 areas where AI can help to revolutionise conservation. These represent both a cross-section of research areas within conservation science and practice, as well as methods through which AI contributes to scientific understanding.

AI stands to rapidly improve our ability to understand distributions of species; locate rare or yet unknown species; identify, model and monitor threats; identify priority areas for conservation; model effectiveness of planned conservation actions; monitor adherence to environmental legislation; and assimilate and interact with scientific evidence. However, we need to ensure AI-supported conversation does not replace valuable established conservation techniques, education and on the ground research.

The intersection within conservation between scientists, practitioners, governments, local communities and indigenous peoples is nuanced and complex. It is important that AI technologies are developed and deployed with understanding of local contexts. Moreover, the fact that many of the planet's intact ecosystems reside in the Global South, poses many significant challenges to the equitable implementation of AI technologies, which are currently predominantly trained and developed in the Global North.

AI will be an invaluable tool to support conservation, however it is not a panacea. Proactive and creative efforts to embrace AI, while also ensuring proper protections and attention to equitable and just conservation practices are in place, will be needed for AI to reach its transformative potential (see Outstanding Questions).

Acknowledgements

This was undertaken as part of the AI@cam project, which we thank for funding. Neil Burgess was funded by UK Research and Innovation's Global Challenges Research Fund (UKRI GCRF) through the Trade, Development and the Environment Hub project (project number ES/S008160/1). Finn Danielsen was supported by the EC projects FRAMEwork, ECS, BESTLIFE2030 and more4nature (grants 862731, 101058509, 101096887, 101133983), UArctic, and Danida (Maarifa). Support for Amy Hinsley was provided by Schmidt Sciences, LLC. Nathalie Pettorelli was supported by Research England.

Declaration of interests

The authors declare no competing interests.

References

1. Ripple, W.J. *et al.* (2017) World Scientists' Warning to Humanity: A Second Notice. *BioScience* 67, 1026–1028

2. Pollock, L.J. *et al.* (2020) Protecting Biodiversity (in All Its Complexity): New Models and Methods. *Trends Ecol. Evol.* 35, 1119–1128

3. Ferraro, P.J. *et al.* (2019) Causal inference in coupled human and natural systems. *Proc. Natl. Acad. Sci.* 116, 5311–5318

4. Nuno, A. *et al.* (2014) Managing social–ecological systems under uncertainty: implementation in the real world. *Ecol. Soc.* 19

5. Jiang, F. *et al.* (2017) Artificial intelligence in healthcare: past, present and future. *Stroke Vasc. Neurol.* 2

6. Gill, A.S. (2019) Artificial Intelligence and International Security: The Long View. *Ethics Int. Aff.* 33, 169–179

7. Jordan, M.I. and Mitchell, T.M. (2015) Machine learning: Trends, perspectives, and prospects*Science*, 349, 255–260

8. Xu, L. *et al.* (2023) Enhanced Persistent Orientation of Asymmetric-Top Molecules Induced by Cross-Polarized Terahertz Pulses. arXiv. https://doi.org/10.48550/arXiv.2312.01579

9. Binley, A.D. *et al.* (2024) Making the most of existing data in conservation research. *Perspect. Ecol. Conserv.* 22, 122–128

10. CBD (2022) *15/4. Kunming-Montreal Global Biodiversity Framework*, Kunming-montreal global biodiversity framework.

11. Noble, A. (2024) *Science in the age of AI: How artificial intelligence is changing the nature and method of scientific research*, The Royal Society

12. Krenn, M. *et al.* (2022) On scientific understanding with artificial intelligence. *Nat. Rev. Phys.* 4, 761–769

13. Skylight.Global (2024) [Online]. Available: https://www.skylight.global/. [Accessed: 01-Aug-2024]

14. Arbimon (2024) Arbimon - Empower your wildlife research [Online]. Available: https://arbimon.org/. [Accessed: 01-Aug-2024]

15. Sentinel (2024) [Online]. Available: https://sentinel.conservationxlabs.com/. [Accessed: 01-Aug-2024]

16. Cole, E. *et al.* (2023) Spatial Implicit Neural Representations for Global-Scale Species Mapping. arXiv. https://doi.org/10.48550/arXiv.2306.02564

17. Silvestro, D. *et al.* (2022) Improving biodiversity protection through artificial intelligence. *Nat. Sustain.* 5, 415–424

18. Simlai, T. and Sandbrook, C. (2021) Digital surveillance technologies in conservation and their social implications. In *Conservation Technology* (Wich, S. A. and Piel, A. K., eds), pp. 0, Oxford University Press

19. Spelda, P. and Stritecky, V. (2020) The future of human-artificial intelligence nexus and its environmental costs. *Futures* 117, 102531

20. García-Roselló, E. *et al.* (2023) The biased distribution of existing information on biodiversity hinders its use in conservation, and we need an integrative approach to act urgently. *Biol. Conserv.* 283, 110118

21. Christin, S. *et al.* (2019) Applications for deep learning in ecology. *Methods Ecol. Evol.* 10, 1632–1644

22. Xian, Y. *et al.* (2017) Zero-Shot Learning — The Good, the Bad and the Ugly. in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3077–3086

23. Hu, K. and Hu, K. (2023) ChatGPT sets record for fastest-growing user base - analyst note*Reuters*

24. Dohmke, T. (2022) GitHub Copilot is generally available to all developers. *The GitHub Blog*.

25. Xi, Z. *et al.* (2023) The Rise and Potential of Large Language Model Based Agents: A Survey. arXiv. https://doi.org/10.48550/arXiv.2309.07864

26. Bommasani, R. *et al.* (2022) On the Opportunities and Risks of Foundation Models. arXiv. https://doi.org/10.48550/arXiv.2108.07258

27. Oquab, M. *et al.* (2024) DINOv2: Learning Robust Visual Features without Supervision. arXiv. https://doi.org/10.48550/arXiv.2304.07193

28. Chen, T. *et al.* (2020) Big self-supervised models are strong semi-supervised learners. in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, pp. 22243–22255

29. Sutherland, W.J. *et al.* (2014) A horizon scan of global conservation issues for 2014. *Trends Ecol. Evol.* 29, 15–22

30. Sutherland, W.J. *et al.* (2024) A horizon scan of global biological conservation issues for 2024. *Trends Ecol. Evol.* 39, 89–100

31. Mukherjee, N. *et al.* (2015) The Delphi technique in ecology and biological conservation: applications and guidelines. *Methods Ecol. Evol.* 6, 1097–1109

32. Sutherland, W.J. *et al.* (2019) Ten Years On: A Review of the First Global Conservation Horizon Scan. *Trends Ecol. Evol.* 34, 139–153

33. Danziger, S. *et al.* (2011) Extraneous factors in judicial decisions. *Proc. Natl. Acad. Sci.* 108, 6889–6892

34. R Core Team (2022) R: A language and environment for statistical computing.

35. Reynolds, S. and Tew, E. (2024) The potential for AI to revolutionise conservation: a horizon scan. *Github*. [Online]. Available: https://zenodo.org/records/13170385

36. Van Horn, G. *et al.* (2018) The iNaturalist Species Classification and Detection Dataset. in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8769–8778

37. Goëau, H. *et al.* (2013) Pl@ntNet mobile app. in *Proceedings of the 21st ACM international conference on Multimedia*, New York, NY, USA, pp. 423–424

38. Observation International and local partners (2022). *Observation.org*. [Online]. Available: https://observation.org/. [Accessed: 01-Aug-2024]

39. Google Lens (2024) *Google Lens*. [Online]. Available: https://lens.google/. [Accessed: 01-Aug-2024]

40. Merlin Bird ID (2024) *Merlin Bird ID - Free, instant bird identification help and guide for thousands of birds*. [Online]. Available: https://merlin.allaboutbirds.org. [Accessed: 01-Aug-2024]

41. Chalmers, C. *et al.* (2019) Conservation AI: Live Stream Analysis for the Detection of Endangered Species Using Convolutional Neural Networks and Drone Technology. arXiv. https://doi.org/10.48550/arXiv.1910.07360

42. Pirotta, V. *et al.* (2022) Detecting illegal wildlife trafficking via real time tomography 3D X-ray imaging and automated algorithms. *Front. Conserv. Sci.* 3

43. de Lorm, T.A. *et al.* (2023) Optimizing the automated recognition of individual animals to support population monitoring. *Ecol. Evol.* 13, e10260

44. Zhao, B. *et al.* (2024) Labeled Data Selection for Category Discovery. presented at the The 18th European Conference on Computer Vision (ECCV)

45. Lu, Y. *et al.* (2024) Accelerating the discovery of biodiversity by detecting “new” species based on machine learning method. *Research square*. https://doi.org/10.21203/rs.3.rs-3832815/v1

46. Dias, M.P. *et al.* (2019) Threats to seabirds: A global assessment. *Biol. Conserv.* 237, 525–537

47. Da Rocha, N. *et al.* (2021) Reduction in seabird mortality in Namibian fisheries following the introduction of bycatch regulation. *Biol. Conserv.* 253, 108915

48. Kahl, S. *et al.* (2021) BirdNET: A deep learning solution for avian diversity monitoring. *Ecol. Inform.* 61, 101236

49. Manzano-Rubio, R. *et al.* (2022) Low-cost open-source recorders and ready-to-use machine learning approaches provide effective monitoring of threatened species. *Ecol. Inform.* 72, 101910

50. Williams, B. *et al.* (2024) Unlocking the soundscape of coral reefs with artificial intelligence: pretrained networks and unsupervised learning win outbioRxiv, 2024.02.02.578582

51. Müller, J. *et al.* (2023) Soundscapes and deep learning enable tracking biodiversity recovery in tropical forests. *Nat. Commun.* 14, 6191

52. Linhart, P. *et al.* (2022) The potential for acoustic individual identification in mammals. *Mamm. Biol.* 102, 667–683

53. Snyder, D. *et al.* (2018) X-Vectors: Robust DNN Embeddings for Speaker Recognition. in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5329–5333

54. Bouffaut, L. *et al.* (2022) Eavesdropping at the Speed of Light: Distributed Acoustic Sensing of Baleen Whales in the Arctic. *Front. Mar. Sci.* 9

55. Wilcock, W.S.D. *et al.* (2023) Distributed acoustic sensing recordings of low-frequency whale calls and ship noise offshore Central Oregon. *JASA Express Lett.* 3, 026002

56. Parijs, S.M.V. *et al.* (2009) Management and research applications of real-time and archival passive acoustic sensors over varying temporal and spatial scales. *Mar. Ecol. Prog. Ser.* 395, 21–36

57. Sanguineti, M. *et al.* (2021) Real-Time Continuous Acoustic Monitoring of Marine Mammals in the Mediterranean Sea. *J. Mar. Sci. Eng.* 9, 1389

58. Helber, P. *et al.* (2019) EuroSAT: A Novel Dataset and Deep Learning Benchmark for Land Use and Land Cover Classification. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* 12, 2217–2226

59. Duporge, I. *et al.* (2021) Using very-high-resolution satellite imagery and deep learning to detect and count African elephants in heterogeneous landscapes. *Remote Sens. Ecol. Conserv.* 7, 369–381

60. Cubaynes, H.C. and Fretwell, P.T. (2022) Whales from space dataset, an annotated satellite image dataset of whales for training machine learning models. *Sci. Data* 9, 245

61. Pettorelli, N. *et al.* (2024) Deep learning and satellite remote sensing for biodiversity monitoring and conservation. *Remote Sens. Ecol. Conserv.* DOI: 10.1002/rse2.415

62. Lisaius, M.C. *et al.* (2024) Using Barlow Twins to Create Representations from Cloud-corrupted Remote Sensing Time Series. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* DOI: 10.1109/JSTARS.2024.3426044

63. Jakubik, J. *et al.* (2023) Foundation Models for Generalist Geospatial Artificial Intelligence. arXiv. https://doi.org/10.48550/arXiv.2310.18660

64. Diao, S. *et al.* (2024) LMFlow: An Extensible Toolkit for Finetuning and Inference of Large Foundation Models. arXiv. https://doi.org/ 10.48550/arXiv.2306.12420

65. Ortega Adarme, M. *et al.* (2022) Improving Deforestation Detection on Tropical Rainforests Using Sentinel-1 Data and Convolutional Neural Networks. *Remote Sens.* 14, 3290

66. Ye, W. *et al.* (2022) Pine pest detection using remote sensing satellite images combined with a multi-scale attention-UNet model. *Ecol. Inform.* 72, 101906

67. López, S. *et al.* (2020) Landscape change in Southern Ecuador: An indicator-based and multi-temporal evaluation of land use and land cover in a mixed-use protected area. *Ecol. Indic.* 115, 106357

68. Mittelstadt, B. *et al.* (2023) To protect science, we must use LLMs as zero-shot translators. *Nat. Hum. Behav.* 7, 1830–1832

69. Brown, C.F. *et al.* (2022) Dynamic World, Near real-time global 10 m land use land cover mapping. *Sci. Data* 9, 251

70. Silva, S.V. *et al.* (2022) Global Estimation and Mapping of the Conservation Status of Tree Species Using Artificial Intelligence. *Front. Plant Sci.* 13

71. Rieke, N. *et al.* (2020) The future of digital health with federated learning. *Npj Digit. Med.* 3, 1–7

72. Nasr, M. *et al.* (2023) Scalable Extraction of Training Data from (Production) Language Models. arXiv. https://doi.org/10.48550/arXiv.2311.17035

73. Kulkarni, R. and Di Minin, E. (2023) Towards automatic detection of wildlife trade using machine vision models. *Biol. Conserv.* 279, 109924

74. Isaac, M.C. *et al.* (2024) Status and trends in the international wildlife trade in Chameleons with a focus on Tanzania. *PLOS ONE* 19, e0300371

75. Di Minin, E. *et al.* (2021) How to address data privacy concerns when using social media data in conservation science. *Conserv. Biol.* 35, 437–446

76. O'Malley, W.C. *et al.* (2024) Machine learning allows for large-scale habitat prediction of a wide-ranging carnivore across diverse ecoregions. *Landsc. Ecol.* 39, 106

77. Estopinan, J. *et al.* (2024) Modelling Species Distributions with Deep Learning to Predict Plant Extinction Risk and Assess Climate Change Impacts. arXiv. https://doi.org/10.48550/arXiv.2401.05470

78. Zhang, H.-T. *et al.* (2024) A novel hybrid model for species distribution prediction using neural networks and Grey Wolf Optimizer algorithm. *Sci. Rep.* 14, 11505

79. Gross, E. *et al.* (2021) *A future for all: the need for human-wildlife coexistence*, WWF, Gland, Switzerland

80. Dertien, J.S. *et al.* (2023) Mitigating human–wildlife conflict and monitoring endangered tigers using a real-time camera-based alert system. *BioScience* 73, 748–757

81. Dasgupta, S. (2018) This tiny camera aims to catch poachers — before they kill. *Mongabay Environmental News*. [Online]. Available: https://news.mongabay.com/2018/06/this-tiny-camera-aims-to-catch-poachers-before-they-kill/. [Accessed: 01-Aug-2024]

82. Leclère, D. *et al.* (2020) Bending the curve of terrestrial biodiversity needs an integrated strategy. *Nature* 585, 551–556

83. Newbold, T. *et al.* (2020) Non-linear changes in modelled terrestrial ecosystems subjected to perturbations. *Sci. Rep.* 10, 14051

84. Díaz, S.M. *et al.* (2019) *The global assessment report on biodiversity and ecosystem services: Summary for policy makers*, Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services

85. The European Green Deal (2021) *European Commission* [Online]. Available: https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/european-green-deal_en. [Accessed: 01-Aug-2024]

86. Tzachor, A. *et al.* (2023) Digital twins: a stepping stone to achieve ocean sustainability? *Npj Ocean Sustain.* 2, 1–8

87. Keith, D.A. *et al.* (2022) A function-based typology for Earth’s ecosystems. *Nature* 610, 513–518

88. Kochkov, D. *et al.* (2024) Neural general circulation models for weather and climate. *Nature* DOI: 10.1038/s41586-024-07744-y

89. Maleki, S. *et al.* (2020) Wetland Restoration Prioritization Using Artificial Neural Networks. *Wetlands* 40, 179–192

90. Ball, J.G.C. *et al.* (2022) Using deep convolutional neural networks to forecast spatial patterns of Amazonian deforestation. *Methods Ecol. Evol.* 13, 2622–2634

91. Potapov, P. *et al.* (2022) Global maps of cropland extent and change show accelerated cropland expansion in the twenty-first century. *Nat. Food* 3, 19–28

92. Berger-Tal, O. *et al.* (2024) Leveraging AI to improve evidence synthesis in conservation. *Trends Ecol. Evol.* 39, 548–557

93. Blaizot, A. *et al.* (2022) Using artificial intelligence methods for systematic review in health sciences: A systematic review. *Res. Synth. Methods* 13, 353–362

94. Elliott, J.H. *et al.* (2014) Living Systematic Reviews: An Emerging Opportunity to Narrow the Evidence-Practice Gap. *PLOS Med.* 11, e1001603

95. Zenni, R.D. *et al.* (2023) Multi-lingual literature searches are needed to unveil global knowledge. *J. Appl. Ecol.* 60, 380–383

96. Christie, A.P. *et al.* (2021) The challenge of biased evidence in conservation. *Conserv. Biol.* 35, 249–262

97. Chowdhury, S. *et al.* (2022) Growth of non-English-language literature on biodiversity conservation. *Conserv. Biol. J. Soc. Conserv. Biol.* 36, e13883

98. Thirunavukarasu, A.J. *et al.* (2023) Large language models in medicine. *Nat. Med.* 29, 1930–1940

99. Tyler, C. *et al.* (2023) AI tools as science policy advisers? The potential and the pitfalls. *Nature* 622, 27–30

100. Duffy, R. *et al.* (2019) Why we must question the militarisation of conservation. *Biol. Conserv.* 232, 66–73

101. Debnath, R. *et al.* (2023) Harnessing human and machine intelligence for planetary-level climate action. *Npj Clim. Action* 2, 1–11

102. Challender, D.W.S. *et al.* (2022) Mischaracterizing wildlife trade and its impacts may mislead policy processes. *Conserv. Lett.* 15, e12832

103. Danielsen, F. *et al.* (2021) The Concept, Practice, Application, and Results of Locally Based Monitoring of the Environment. *BioScience* 71, 484–502