

Reflections from the Workshop on AI-Assisted Decision Making for Conservation

Abstract

In this white paper, we synthesize key points made during presentations and discussions from the AI-Assisted Decision Making for Conservation workshop, hosted by the Center for Research on Computation and Society at Harvard University on October 20–21, 2022. We identify key open research questions in resource allocation, planning, and interventions for biodiversity conservation, highlighting conservation challenges that not only require AI solutions, but also require novel methodological advances. In addition to providing a summary of the workshop talks and discussions, we hope this document serves as a call-to-action to orient the expansion of algorithmic decision-making approaches to prioritize real-world conservation challenges, through collaborative efforts of ecologists, conservation decision-makers, and AI researchers.

Lily Xu^{1,*}, Esther Rolf^{2,3,4,*}, Sara Beery⁵, Joseph R. Bennett⁶, Tanya Berger-Wolf^{7,8,9,10}, Tanya Birch^{11,†}, Elizabeth Bondi-Kelly^{5,12}, Justin Brashares¹³, Melissa Chapman¹³, Anthony Corso¹⁴, Andrew Davies¹⁵, Nikhil Garg¹⁶, Angela Gaylard¹⁷, Robert Heilmayr¹⁸, Hannah Kerner¹⁹, Konstantin Klemmer²⁰, Vipin Kumar²¹, Lester Mackey²⁰, Claire Monteleoni^{4,22}, Paul Moorcroft¹⁵, Jonathan Palmer²³, Andrew Perrault⁷, David Thau²⁴, and Milind Tambe^{1,2,11}

¹School of Engineering and Applied Sciences, Harvard University

²Center for Research on Computation and Society (CRCS), Harvard University

³Harvard Data Science Initiative (HDSI)

⁴Department of Computer Science, University of Colorado Boulder

⁵Electrical Engineering and Computer Science, Massachusetts Institute of Technology

⁶Department of Biology, Carleton University

⁷Department of Computer Science and Engineering, The Ohio State University

⁸Department of Electrical and Computer Engineering, The Ohio State University

⁹Department of Evolution, Ecology, and Organismal Biology, The Ohio State University

¹⁰Translational Data Analytics Institute, The Ohio State University

¹¹Google

¹²Electrical Engineering and Computer Science, University of Michigan

¹³Department of Environmental Science, Policy, and Management, University of California Berkeley

¹⁴Department of Aeronautics and Astronautics, Stanford University

¹⁵Department of Organismic and Evolutionary Biology, Harvard University

¹⁶Operations Research and Information Engineering, Cornell Tech

¹⁷African Parks

¹⁸Bren School of Environmental Science and Management, University of California, Santa Barbara

¹⁹School of Computing and Augmented Intelligence, Arizona State University

²⁰Microsoft Research

²¹Department of Computer Science and Engineering, University of Minnesota

²²INRIA Paris

²³Wildlife Conservation Society (WCS)

²⁴WWF Global Science

* These authors contributed equally

† The opinions are the authors' and don't necessarily reflect those of Google.

1 Introduction

Effective area-based conservation management is more imperative than ever before [1, 2, 3, 4]. The world has lost one-third of global forests [5] and one million species face extinction due to human activities [6]. To reverse these trends, conservation decision-makers are tasked with managing protected areas [7, 8], prioritizing endangered species [9, 10, 11], and navigating opportunity and cost trade-offs between conservation and economic development [12, 13, 14] — often in the face of uncertainty, complex constraints, and long planning horizons [15, 16, 17, 18]. Simultaneously, recent years have seen rapid advances in artificial intelligence (AI), including deep learning to uncover patterns in high-dimensional datasets [19], tractable approaches to navigate exponentially large problem spaces [20], and statistical inference to learn causal effects of actions [21].

The majority of conservation areas around the world lack funding, personnel, and equipment [22]. Without an increase in the resources available for conservation areas, it is critical to prioritize and scale the impact of the scarce resources we have. Much existing work at the intersection of AI and conservation focuses on extracting basic information from primary data, largely in the form of computer vision for biodiversity monitoring, such as supervised learning approaches to identify animals or classify land cover [23, 24, 25]. In addition to these applications, AI holds potential to contribute to *prioritization* of resources within protected areas and to improve the execution of conservation area *management* activities at scale, for example through optimization, sequential planning, and causal inference.

To this end, the AI-Assisted Decision Making for Conservation workshop convened to explore opportunities to expand the breadth of AI methods — particularly for optimization, planning, and evaluation — that are being developed for and applied to conservation. In this report, we summarize the discussions held and challenges raised during the workshop surrounding conservation area management; the role of AI, Earth observation, and causality in conservation decision making; and limitations in the application of AI for conservation.

First, we synthesize key conservation priorities. For effective conservation area management, the primary steps are to:


1. **Understand the world.** To inform conservation decision making, managers must first develop an understanding of the status of animal populations, species distributions [26, 27], land cover [28], animal threats [29], and human activity [30].
2. **Act in the world.** Using their best understanding of the world, managers are tasked with taking conservation actions like designating new conservation areas [11], allocating ranger patrols [31], implementing human–wildlife conflict deterrence programs [29], intervening in the illegal wildlife trade [32], securing sustainable harvest [33], supporting sustainable economies [34], and countless other decisions — almost always in the face of limited resources and uncertainty.
3. **Evaluate impact.** Managers want to evaluate the outcomes of their actions in order to learn whether certain actions are achieving desired impact [35], and to iteratively improve decision making within complex, interactive systems [36, 37].

We discuss these steps in detail in section 2; see figure 1 for a visual summary.

AI methods have the potential to aid in each of these three components of conservation management. Some approaches have already been in use for years; other approaches are not yet deployable but are growing in promise. Algorithmic approaches can help:

1. **Understand the world.** AI for Earth observation can help across multiple scales and modalities, from studying changes using remote sensing imagery [38, 39, 40, 41, 42, 43, 44, 45] to monitoring individual animal movement with photographs [46, 47], acoustic monitoring [48], and GPS tracking [49]. This is the focus of section 3.1.
2. **Act in the world.** AI methods can help plan across a range of contexts relevant to conservation, including sequential settings, under model uncertainty, and in the presence of multiple agents [50]. This is the focus of section 3.2.

AI + Conservation area management




1 Understand the world

Predictive models based on **proximal and remote sensing technologies** can help conservation managers monitor land cover, animal populations, and human activity.




2 Act in the world

Algorithmic decision making can help optimize and prioritize conservation efforts in the face of limited resources, uncertainty, and repeated decisions.




AI decision making can aid designing new conservation areas or allocating resources, accounting for uncertainty in species distribution and habitat value.



AI researchers ought to work closely with conservation managers to discuss resource constraints and priorities.

3 Evaluate impact

Causal inference can help understand the direct effect of management actions to improve conservation program impact, including the use of AI – particularly for adaptive management.



Evaluating impact of conservation efforts is challenging in an ever-evolving world.



Evaluating effectiveness of AI within conservation management programs is essential to building trust and long-term use.



Social, financial, and environmental impact of conservation decisions interact in complex ways.

Figure 1: Effective conservation management is a planning, implementation, and adaptation cycle consisting of three key steps. First, we must *understand the world*; proximal and remote sensing technologies can help monitor land cover and animal populations (section 3.1). Second, we must *act in the world*; algorithmic decision making can help optimize and prioritize conservation efforts (section 3.2). Third, we must *evaluate the impact* of our actions; causal inference can help understand the direct effect of our actions (section 3.3).

3. **Evaluate impact.** One important aspect of evaluating the impact of conservation decisions is isolating causal effects of these decisions. Causal inference techniques can help understand the *causal* effect of conservation actions, including scenarios where randomized controlled trials are too expensive or infeasible to conduct [51, 52, 53]. In addition, understanding causal relationships between elements of the environment, such as the effect of elephant populations on carbon stocks in savannas [54], can aid in designing comprehensive management decisions. This is the focus of section 3.3.

In this white paper, we recollect and synthesize key points made during presentations and discussions from the 2022 AI-Assisted Decision Making for Conservation workshop (figure 2). Where possible, we identify key open research questions in resource allocation, planning, and interventions for biodiversity conservation that arose in whole-group or partial-group conversations. In particular, we pinpoint conservation challenges that not only require AI solutions, but require novel methodological advances. While this synthesis document is not meant to serve as an entirely complete or comprehensive research agenda, we nonetheless hope that it can help **orient the expansion of algorithmic decision-making approaches so that they prioritize real-world conservation challenges**, through collaborative efforts of ecologists, conservation decision-makers, and AI researchers.



Figure 2: The AI-Assisted Decision Making for Conservation workshop was hosted by the Center for Research on Computation Society (CRCS) at Harvard University in October 2022. The co-authors of this paper were the workshop participants, coming from academia, industry, and non-profits with expertise in computation, conservation, and deploying research.

A full schedule of the two-day workshop, with PDF slides from all talks, is available on the workshop webpage: <https://crs.seas.harvard.edu/conservation-workshop>.

2 Primer to conservation area management

Conservation areas are defined geographical spaces on land or sea managed for biodiversity and climate mitigation. These include protected areas [7], Indigenous and Community Conserved Areas (ICCAs) [55], and Other Effective area-based Conservation Measures (OECMs) [56]. Initiatives like 30x30, which aims to protect 30% of land and sea by 2030, underscore the importance of protected areas for the sustainable development goals (SDGs) and for the future of life on Earth [57, 58].

For the scope of this document, we focus on field conservation, namely conservation area management and species management. Field conservation efforts include preventing human-wildlife conflict [59], processing multi-modal inputs from sensors including camera traps [60, 61, 62, 63], supporting Indigenous and local communities [64, 65], supporting women [66], and tracking and predicting animal movement [67]. The full scope of conservation initiatives is much broader, as documented by the Conservation Measures Partnership [68, 69]. Other activities include raising awareness, providing alternative livelihoods [70], and creating policy guidelines [71].

Ecosystems are highly complex dynamic systems, making conservation planning challenging. If we don't know the answer to simple questions of monitoring, such as: *what is the range of the western gorilla?*, then it is difficult to answer a host of more complex yet critically important questions such as: *how do we spatially prioritize land to protect? Does deforestation lead to increased human-wildlife conflict? Does a conservation program to protect elephants improve people's livelihoods?*

In this section, we provide an overview of management theory and current conservation practices while highlighting key challenges in monitoring, decision making, and impact evaluation.

Prioritization Conservation prioritization includes both deciding how best to allocate conservation resources in areas that are already protected as well as deciding where to expand existing conservation measures. The designation of new protected areas is part of the planning process, which requires taking into account the locations of existing protected areas, landscape connectivity, and biodiversity distributions.

Adaptive management Adaptive management is a planning, implementation, and adaptation cycle [72, 73]. The goal of adaptive management is to manage ecosystems using the best available information, accounting for feedback between decisions and outcomes.

Adaptive management approaches are complicated by exogenous changes, including climate change [74], which require repeated cycles of monitoring, planning, and action. Often, *sequences* of events are important: for example, availability of surface water impacts elephant movements, but also depends on previous years' rainfall. Other sources for dynamic change include species migration, seasonal weather patterns, and species interaction (including predator-prey interaction and habitat competition).

Evaluation and evidence-based management Adaptive management is closely related to evidence-based management [75], which centers around using repeated assessments to support decision making. Defining metrics of conservation assessment — what counts as conserved, and/or what defines successful protected area management — is critical to ensuring conservation strategies contribute to underlying objectives of protected areas. For example, land area being classified as protected does not necessarily mean it is well-resourced or effectively managed. The metrics we define to evaluate the outcomes of protected area management, such as changes in population abundance or land use, can then be used to update the management strategy (adaptive management).

In addressing these challenges, we need to think about simultaneously optimizing for

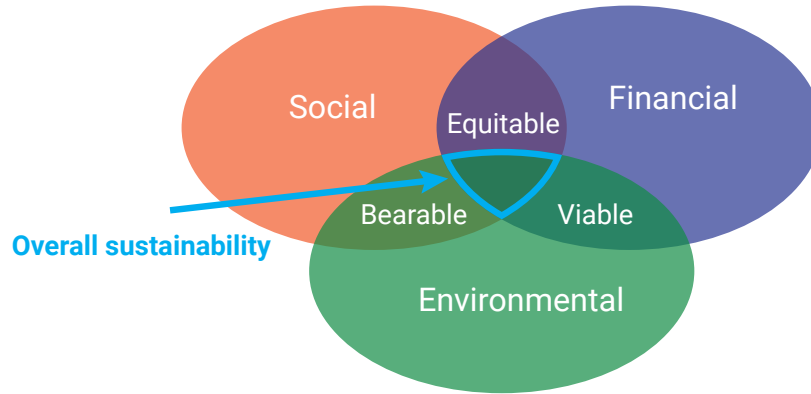


Figure 3: Overall sustainability must satisfy social, environmental, and financial considerations. These solutions ought to be viable to implement, equitable to impacted communities, and bearable to both humans and ecosystems. (Figure adopted from the triple-bottom line sustainability framework [86].)

environmental, social, and economic sustainability (see figure 3). Social and economic benefits and risks must be considered in tandem with environmental impacts [76, 77]. For example, increased lion conservation near human habitation may come with risks to those neighboring communities by increasing the risk of human–wildlife conflict [78]. Accounting for the many facets of sustainability can be an effort stemming from many communities working together [79]. The move towards truly participatory planning requires community engagement to empower Indigenous and local communities and integrate their expertise [80, 81, 82]. For computational scientists, this includes reducing barriers to access to conservation tools — including Wildlife Insights [62], eBird [83], iNaturalist [84], SMART [85], and decision-facilitation tools.

3 AI can assist conservation decision making

AI has the potential to aid in different stages of the conservation management cycle through data-driven decision support. In conservation area management, AI is already used to aid at every step of setting goals, monitoring resources (the step at which AI is most commonly involved), taking management action and measuring outcomes (some AI methods at present), analyzing trends and refining the approach (some AI methods at present) [87, 88, 63].

For monitoring, AI can be paired with Earth observation data to help understand the world. Examples include predicting areas of deforestation risk [89], estimating carbon stock [90], and learning and downsizing climate variables [91, 92]. AI can help extract relevant scientific information from passive and active sensor data, including images (satellites, camera traps, citizen science platforms, drones), acoustic data, motion tracking, on-body (e.g., radio collars), in-situ sensors, eDNA measurements, and even media sources, including social media, for alerts of environmental hazards and to monitor public sentiment of deployed programs. AI can also help overcome biases in data collection and missing data [93].

For decision making, AI can be used for resource allocation (e.g., patrols, sensor deployment, and prioritization), strategic reasoning, and multiobjective optimization. AI can also be a tool for increasing equity around language variation via internalization and auto-translation. AI can also enable participatory action, both through intelligent agents and by making humans more efficient and able to scale.

Lastly, for evaluating impact of actions, combining causal inference with AI has the potential to describe trends in conservation.

At the same time that AI can be powerful in reducing barriers and access to data for protected area managers and partnering communities, it can inadvertently serve to reinforce existing power structures. We discuss the interconnectedness and complexity of design goals and system impacts in section 4.

3.1 Earth observation: AI to help understand the world

The immense scale at which many conservation decisions are made requires comprehensive monitoring of the Earth over time and space. AI techniques can extract higher-order information from datasets of raw observations collected at different spatial and temporal resolutions across the globe. The outputs from AI models, such as machine learning predictions, may be useful for conservation decision making directly or may be combined as inputs to other downstream analyses to produce information needed for decision making.

The field of Earth observation involves the collection of physical, chemical, and biological data about planet Earth and its surface, often using sensors with variable resolution, sampling, and instrumentation from the ground up to space. Figure 4 shows an illustrative example of how AI methods can extract information from raw data to aid conservation management decisions.

3.1.1 Data collection and characteristics

We distinguish Earth observation sensors into two main categories: proximal sensing (e.g., field data from camera traps, acoustic monitors, temperature sensors) and remote sensing (i.e., satellites, UAVs).

Proximal data are collected “in situ” and close to the target of interest, using on-the-ground sensors. Throughout the last decade, the emergence of affordable IoT devices and robots has enabled the creation of vast sensor networks with pervasive spatial coverage, such as networks of camera traps for monitoring wildlife distribution or urban air quality sensors. The appropriate spatial and temporal allocation of sensors is a central determinant of the success of proximal sensing.

Remotely sensed data are captured using sensors that are distant from the target of interest, such as satellites. More pervasive than ever, satellites and other sensors provide near real-time data streams of optical imagery (RGB and multispectral), synthetic-aperture

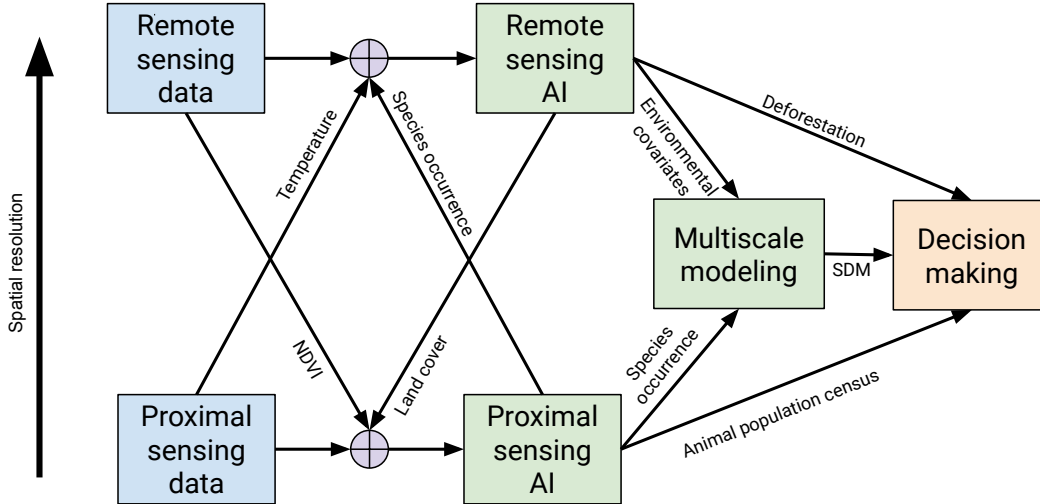


Figure 4: An illustrative example of how proximal and remote sensing approaches, the processing of the captured data, and AI can be integrated to inform decision making. The vertical axis roughly represents the spatial scale of the observations (from most proximal to most remote), the horizontal axis roughly represents the steps of the decision process, from data sourcing via modeling to decision making. Purple circles denote data (dis)aggregation and (pre)processing steps.

radar (SAR), and more. Remotely sensed data are characterized by their tremendous spatio-temporal heterogeneity, which leads to distinct modeling challenges. For example, resolution, frequency, and quality of observations may vary across time and space depending on sensor availability. In addition, specific spatial, temporal, and spectral resolutions (e.g., the resolution of a satellite camera, the fly-over frequency, or the measured wavelengths) may constrain modeling efforts with remotely sensed data.

For enabling evidence-based conservation management, proximal and remotely sensed data can be valuable individually or in combination. The proximal vs. remote sensing distinction is an oversimplification reflecting end-points of a range of spatial scales (e.g., sensors onboard drones and aircraft lie somewhere between proximal and remote sensing). Likewise, AI methods are not constrained to working with either remote or proximal sensing data individually, but can also combine multiple scales into one modeling framework.

Data from different sources can be complementary [94]: a camera trap (proximal sensor) might take a picture of a buffalo on the ground, while a satellite (remote sensor) with a high-resolution camera might take an overhead picture of the same animal. The two sources can be combined for a more comprehensive buffalo census than either source independently would allow.

3.1.2 Data analysis and modeling

Machine learning approaches, particularly computer vision, have challenged traditional image processing techniques in the analysis of remotely sensed data. The large scale and pervasiveness of the available data can synergize well with the “data-hungriness” of machine learning. However, labeled data are often scarce, especially in conservation-related domains. This scarcity of labeled data motivates dedicated learning approaches that are (1) resource efficient (e.g., zero- or few-shot learners [95]) and (2) able to learn in an unsupervised or weakly-supervised manner. Spatio-temporal dependencies in the data are another source of complexity [96, 97].

Machine learning methods for Earth observation data have begun to adapt to these particular challenges [44]. These developments are evidenced by emerging domains in method-

ological machine learning research, such as physics-informed deep learning [98], spatially-explicit learning [99, 100], weakly-supervised learning using spatial data [43, 101, 102] or spatial representation learning [103].

As such, dedicated methodological research is also required to generalize and adapt existing computer vision approaches to — or develop new customized approaches for — the remote sensing domain in use for conservation. For example, computer vision algorithms for natural images can take advantage of the abstractions of “foreground” and “background,” but there is not necessarily a notion of foreground or background in remote sensing images. Object detection tasks can also differ substantially in conservation images (like satellite or aerial imagery) compared to traditional computer vision for natural images, due to the existence of a large number possible objects across a variety of spatial scales.

3.1.3 Open challenges in AI for Earth observation

Concurrent revolutions in data (due to new satellite missions and other sensor data availability), compute (due to cloud infrastructure and cloud-hosted open datasets), and analysis methods have enabled tremendous advances in the development and application of AI for remote sensing and Earth observation. At the same time, these advances have exposed several challenges to be addressed. Here we outline several challenges in AI for Earth observation that are of key concern in conservation and decision-making settings.

Data availability and transparency The open availability of remote sensing imagery (Landsat, Sentinel, Planet/NICFI) has spurred the creation of derivative datasets (e.g., Global Forest Change, Dynamic World). At the same time, there are not always direct incentives for releasing source training data or validation data, or the models used to create the derivatives. In addition to availability, reporting the accuracy and applicability of measured or predicted datasets is particularly important [96, 104]: errors in ground-truth labels or features may propagate throughout model training and into deployment, and any downstream policies or analyses. Even in cases where ground truth data are less sparse, human errors can persist, requiring substantial data pre-processing to be suitable for model training [43].

Model interpretability and transparency As dynamic data streams are constantly changing, models — and the ways people use them — need to be dynamic as well. This introduces additional challenges for model interpretability, transparency, and explainability, toward answering questions such as: *When should the model be used? Who is the intended user? What is the model intended to do? What are the challenges and risks of using the model outputs? If the models themselves are changing constantly, does that make them usable in unseen environments or different spatial-temporal contexts?*

Missing labels While many labeled datasets have been published in recent years for training and evaluating AI models, existing labeled datasets for remote sensing are relatively small and sparse. This is in part because labels for Earth observation data can be difficult to annotate, for example due to overlapping ecotones between classes in remote sensing inputs. There is often a need for domain expertise for annotation, and many use cases require ground-truthing (e.g., on-the-ground observation) to determine labels.

Label biases Furthermore, both the sampling and authoring institutions of existing labeled datasets are mostly biased toward North America and Europe [105], which in turn results in disparate levels of attention and model performance across geographies (a phenomenon that has also been observed for traditional AI data modalities [106]). This challenge is particularly important in conservation, as ecological systems vary substantially throughout different geographical areas and models trained on data-rich environments in the Global North might not be transferable to the Global South [107]. Future efforts could

investigate how to more efficiently, scalably, and equitably collect labeled datasets for remote sensing and AI.

Generalization Generalization of models across spatial (e.g., from one landscape to another), temporal (e.g., from rainy season to dry season), and spectral dimensions (e.g., from RGB to thermal imagery) is an important open challenge for conservation that may be aided by methods for transfer learning or distribution shift.

Responsibility and governance While the increased availability of data, methods, and compute for AI and remote sensing is empowering for many researchers and practitioners, it also raises important ethical concerns. *Who is responsible for the output of a model (predicted map) and how it gets used? How do we prevent bad actors from misusing data and methods? Who has the authority to make maps, particularly given the history of extraction (local ownership, parachute science [108], and data colonialism)? How can a predicted map be used measure the impact of efforts on conservation or sustainable development goals?*

Increased collaboration The AI, Earth observation, and conservation research communities can benefit from more direct collaboration across fields [109]. For example, collaboration between the satellite remote sensing community and the biodiversity monitoring community (camera trap, acoustic, eDNA [110]) could lead to sharing of methods as well as development of new multi-modal methods leveraging the breadth of available data.

Computational costs The size (and constant growth) of both proximal and remotely sensed datasets also poses computational challenges to store, process, and analyze. The scale of many remote sensing datasets means that traditional methods can be computationally expensive; AI systems may provide faster modeling, particularly at inference time, but large-scale geospatial extrapolations can still be computationally intensive, even (or especially) using the latest ML models. As such, the development of large-scale cloud compute infrastructure over recent decades has benefited the Earth observation community substantially. We discuss computational costs in greater detail in section 4.

3.2 Algorithmic decision aids: AI for decision making

3.2.1 The role of AI in adaptive management

As discussed in section 2, adaptive management is a cycle consisting of planning, implementing, and adapting conservation management decisions. AI can provide a robust means of integrating evidence into adaptive management and has the potential to deal with complexity in ways that spatial statistics, simple heuristics, or general rules of thumb cannot. Management decisions in conservation are often complex, with trade-offs among multiple objectives [111] and limited information as a basis for decisions [112]. Such complexities amplify when the impacts of decisions must be predicted across large spatial and temporal scales, for example when planning for migratory bird conservation [10], building climate change resilience into protected area networks [113], and accounting for socio-political risk in conservation planning [114].

3.2.2 Sequential decision making

We can define a decision problem by an objective, a set of possible actions over time, and a model of the effects of those actions in the environment. Many decision problems in biodiversity conservation involve reasoning about a sequence of decisions for settings with complex dynamics and large amounts of uncertainty. AI methods can be applied to build models of the system dynamics, infer unobserved variables from observations, and to plan

with those models and observations to determine an optimal sequence of actions under uncertainty [115, 50]. Examples discussed over the course of the workshop included:

- **Patrolling to interdict poachers** [116, 117]: The objective is to deter poachers by deciding when and where to send rangers within a protected area. Poaching activity is only observed indirectly through the discovery of traps and encampments, and only in regions where the rangers patrol. The behavior of poachers and their response to changes in patrolling behavior poses a challenging modeling task, which may require robust planning [118].
- **Reintroduction of species** [119]: The objective is to ensure the long-term viability of the species by deciding when, where, and how many individuals should be reintroduced. The effects of reintroduction require ecological models, which can be informed by tracking the locations and movements of the individuals via collars or sensing modalities.
- **Protected area design** [120, 121]: The objective in protected area design is to increase species richness, preserve threatened species, or generally improve the connectedness and resilience of protected areas, while satisfying constraints of cost or total protected area. Decisions include where and when to purchase land, and making purchases now or delaying in order to gather more information or wait for opportunities. We have many tools to observe the land attributes we care about, but our observations are highly imperfect. Sometimes, lands are entirely inaccessible for field surveys. Time horizons for protected area design are often very long and ideally evaluation is done against a wide range of possible scenarios, including budgetary uncertainty.

Sequential decision problems can be modeled as Markov decision processes (MDPs), which account for uncertainty in the dynamics, and partially-observable Markov decision processes (POMDPs), which account for uncertainty in the state of the world. Recently, AI has yielded rapid advances in methods for solving large and complex MDPs and POMDPs such as deep reinforcement learning [122] and belief state planning [123]. These advances offer the potential to rigorously optimize sequential decisions for conservation [124, 125, 126]. The parallel advances in machine learning-enabled sensing methods (see section 3.1) accelerate this potential by providing data that can be used to create dynamic models and infer the system state, which in turn can inform more precisely targeted and useful actions.

3.2.3 Open challenges in sequential decision making for conservation

Imperfect or unknown models Previous successes of algorithmic decision making have been concentrated in domains that can be accurately represented by a known model, such as games [127], transportation systems [128], or subsurface applications [129]. However, in conservation, a key challenge to applying algorithmic decision making is building accurate models of the environment dynamics and the effect of actions. Data-driven approaches have been successful at modeling complex dynamical systems, but often require large amounts of data or additional expert knowledge in order to generalize well. Since no model is perfect, model uncertainty should be accurately accounted for so that plans derived from the model work well in environments that deviate from the model.

Challenging optimization Finding an optimal strategy in well-described environments may also be difficult. Decision problems in conservation often involve long time horizons or high-dimensional and continuous state, action, and observation spaces. Modern solvers have begun to scale to such problems [123], but more work needs to be done to improve these optimization techniques.

Integrating human expertise In addition to pure algorithmic development, approaches that integrate expert human knowledge in the planning loop may also be needed.

Mismatch between predictions and decisions Conservation decision making often involves machine learning as a substep to generate predictions, such as predicting species populations or identifying areas vulnerable to illegal mining. But then those predictions are ultimately used for some downstream decision, such as selecting areas to conserve as species habitat or add monitoring stations. Thus, aligning predictions with decision making is another priority [130].

Validation Once a plan has been generated, it is crucial to validate that plan for efficacy and safety [131]. Often, high quality plans can be developed from simple models, but we need extensive validation to check that the plan is indeed of high quality. Validation can be done by building additional validation models [128], stress-testing the plan to discover the most likely failure modes [132], and testing the plan on a small scale before larger scale deployment [128]. That said, no model accounts for all possible consequences of actions, so additional validation is needed to anticipate unforeseen consequences of the proposed actions, such as negatively affecting marginalized groups, or damaging other aspects of the environment.

Unknown unknowns A final challenge in sequential decision making involves so-called “unknown unknowns” that cannot be quantified at the time of an initial decision. These phenomena are relatively common in conservation practice [133, 134], and can be related to biological, environmental or sociopolitical factors [114]. Potential research avenues to help better account for such phenomena include designing adaptive management programs that can simultaneously build general knowledge of a system while informing more specific decisions, and explicitly including risk-tolerance scenarios into management prescriptions.

3.3 Causality: AI and economics to understand conservation impacts

3.3.1 The role of causal inference in conservation management

Applied challenges in conservation science and management often raise difficult questions of cause and effect. *Will additional anti-poaching patrols allow a threatened species to recover? What impact will a new conservation area have on deforestation?* Computer science, economics, and conservation biology have all made important advances in developing causal inference methods that enable researchers to use observational data to estimate causal effects [53, 52, 21]. AI can improve our understanding of these causal relationships in multiple ways by, for example, (a) extracting more useful information about the world from existing data; (b) improving predictions of counterfactual scenarios of what might have happened in the absence of some intervention; or (c) providing more data-driven approaches to evaluate heterogeneity in treatment effects [135].

However, data limitations often inhibit the direct use of off-the-shelf techniques from causal inference to study conservation impacts. Ecology settings commonly involve continuous real-time data, distribution shifts, and nonstationarity — settings that are all challenging to study because ground truth is not easily available for evaluating the models. AI offers opportunities to expand causal studies by generating data on outcomes and improving causal estimates.

While causal AI has made important contributions to all of these areas of research, our workshop’s discussions primarily focused upon the *potential* for AI to extract new information from data, as well as the *methodological considerations* that these data necessitate for the goal of accuracy of causal estimates.

Potential to uncover causal relationships in conservation Many rigorous research designs for observational, causal inference require consistent data spanning large geographic

areas, measured over multiple periods of time. Given these needs, AI-based Earth observation has proven to be a valuable source of information for causal inference in conservation.

However, these requirements have led the environmental economics and conservation biology communities to often focus on outcomes that are relatively easily measured. Rather than track a policy’s impacts on the primary outcomes conservation practitioners truly care about (e.g., biodiversity and species populations), impact studies often evaluate changes in more easily measured proxies such as land cover and land use change. As we develop longer time series and broader spatial coverage of biodiversity data (based on e-DNA, audio sensors, and camera traps; see section 3.1), there is growing opportunity to more directly estimate treatment effects on the real variables of interest.

Methodological considerations for causal inference for conservation The economics community often uses predictions from machine learning models as inputs into causal inference analyses, where an assumption might be made that derived data generated through AI can be integrated into existing causal inference research designs without significantly biasing estimated treatment effects. However, a growing body of research highlights that the structure of these data, and their underlying error distributions, require more careful reconsideration of causal inference techniques in order to generate unbiased estimates. For example, commonly-used, two-way fixed effects models may yield biased estimates of conservation impact when applied to pixel-level data detailing deforestation [136].

Careful consideration of the underlying data generating processes can help scientists structure their causal inference models to yield more accurate estimates of the impacts of conservation policies. More broadly, remotely sensed measures of environmental conditions exhibit both random and more structured measurement errors, which can bias estimates of causal relationships. In these settings, multiple imputation methods can generate useful error corrections to improve the accuracy of causal estimates [137]. This growing literature highlights the value of greater collaborations across disciplines — a broader understanding of both the ways that machine learning data are generated and used can help identify ways to more reliably estimate the causal impacts of conservation interventions.

Overcoming missing data Often in conservation, the outcomes managers and decision-makers care about are not directly observed in the data. Such missing data problems are studied in machine learning and statistics, including in areas such as positive-unlabeled learning. Conversely, statistical methods originally developed in ecology, such as missing species estimation, can be applied to other settings such as understanding where we might be missing data about urban incidents (fallen trees and sidewalk repairs) for responders to address [138]. These novel methods enable us to infer outcomes where there is no labeled outcome data.

3.3.2 Open challenges in causality for conservation decision making

In addition to the challenges of using predicted or remotely sensed variables in causal inference, we also highlight questions of external validity that arise in any causal inference. In particular, the following key challenges should come into consideration when studying causal relationships of conservation factors:

Shifts over time A correctly identified causal relationship that holds in past data may or may not hold in current or future data. This presents an opportunity to design experiments, data collection protocols, and algorithms that integrate monitoring and optimization, trading off improving model fidelity with taking impactful intervention action.

Shifts over space Trends identified in case studies may not generalize to other areas. This presents an opportunity to clearly communicate regions where the inference is and is

not expected to be valid, and to design frameworks that increase the spatial extent that causal analysis can be applied to.

Generalizability How do we go from internal validity to external validity? This depends in part on the data we choose to make inferences on. Case studies may not generalize to other areas. It is essential to collect more datasets for distribution shifts in ecology settings. Creating systems to share data across geographic areas and time is a promising way to shed light on these challenges and provide a foundation for future methodological work.

4 Critical considerations in AI-assisted decision making

During our two-day workshop, several cross-cutting themes emerged throughout our sessions, highlighting critical considerations for AI-assisted decision making in conservation settings. In this section we synthesize those themes, grouped into: opportunities, limitations, and costs associated with using AI for conservation decision making.

4.1 Opportunities: What is AI-able?

Contexts where data are plentiful, such as remote sensing and camera traps, are amenable to machine learning. In more data sparse situations, additional methods may be needed, for example synthetic or model-derived data. AI has found limited use in predictive regimes in the context of conservation since the necessary data, both historical and ecological context-rich, were typically sparse, incomplete, or missing altogether. As those data become available (section 3.1) — in conjunction with the increasing availability of both current and historic data on human activity, weather, natural disasters, habitat, and more — there is an opportunity to develop AI approaches that can assist conservation efforts. Below we outline three considerations for assessing such opportunities.

Scale Conservation involves multiple scales (global to local) and different data are available at different scales. At a local scale there are clear challenges in data availability relating to the difficulty of data acquisition and data privacy issues. In particular, in conservation there are good reasons that data cannot be made FAIR [139]. These include risks around endangered species and data sovereignty issues at multiple levels including governmental restrictions on data sharing and restrictions on sharing Indigenous and local knowledge (see the CARE principles [140]).

Metrics Another challenge arises around the differences between outcome metrics (what real-world outcomes are desired) and AI performance metrics — what mathematical objectives the AI system is designed to optimize given the existing data — in a given conservation context. For example, biodiversity has a multifaceted definition which “includes diversity within species, between species and of ecosystems” [141]. In addition to biodiversity, conservation often focuses on ecosystem services such as pollination, water quality, carbon sequestration. And, for all aspects, benefits to people must be considered, including the health of communities on land being conserved or restored. Finally, the performance metrics used to evaluate the efficacy of AI approaches, such as species identification accuracy, are not necessarily well-aligned with outcome metrics. Performance metric improvement may or may not have a positive impact on the outcome of the entire conservation system, and this relationship is understudied.

Adaptability Conservation often occurs in non-stationary situations, necessitating AI solutions that can adapt under shifting conditions. AI methods frequently see large performance losses when used *out of distribution* — i.e., in a scenario that doesn’t match the data they were trained on [142, 143]. Understanding when and how much to trust the outcomes of an AI system in a changing world, and how to efficiently and cost effectively improve performance out of distribution is an open and impactful area of AI research paramount to the effectiveness of AI systems in conservation settings.

4.2 Limitations: What is not AI-able?

AI for conservation has the potential to address the challenge of scaling limited resources to significantly increase the area of well managed protected areas around the globe. For example, AI can help provide better predictive modeling capabilities than traditional spatial statistics. On the other hand, AI for conservation has the risk of promoting solutions that

fail to scale, are more complex and costly than is necessary for the task, and may distract conservation agency staff from applying more appropriate decision-making techniques.

As a community, we have a collective responsibility to identify solutions that will scale and have conservation impact, to be clear on our respective responsibilities and commitments, and to marshal the necessary resources to ensure these solutions are delivered. Conservation professions need guidance to allow them to better evaluate how to engage with conservation AI projects. Conversely, AI researchers and developers must invest in understanding practical challenges of conservation practitioners.

We recognize that while tech is often less than 20% of the overall solution, computation and AI have a potential role to play in every part of the conservation monitoring lifecycle. We emphasize that AI decision-making methods should be adopted as decision aids, but cannot replace the expertise of human practitioners. One way to frame this is to design AI to help invoke critical thinking — that is, AI to help enable decisions that then require conversations between human decision-makers to question the status quo of conservation initiatives, including defining what is beneficial in context, and potentially seeking alternative solutions to AI [82].

4.3 Costs: What are the costs of using AI?

There are several costs associated with the choice to use AI as a part of a management system that may not be readily apparent. We saw that across many of our talks, the hidden costs of AI were a large factor contributing to the potential impact or reduction of impact.

One factor that was brought up across many case studies was that the AI model was only 10% of the work of integrating AI in conservation management frameworks. Agreeing on where the responsibility lies for the remaining 90% of the work can be challenging, as there are differences of values and resources between academia, conservation land managers and organizations, and potential funding agencies or organizations. Further, in reality, AI will be only a component of a participatory system, as opposed to a fully-automated solution, and the use of AI will have ongoing computational and human costs.

We posit that it is important to assess all of the costs of AI use — not just the costs of data collection and model training — in deliberating the potential benefits of using of AI in conservation contexts. One key (and currently missing) component is low-cost assessment of AI feasibility for a new project, particularly in the case where there are existing off-the-shelf models that may be used. Making it cheap and easy to ask the question *can I use this model for my project, and will it save me time and/or money?* should be a key priority. For example, the satellite-image embedding and prediction model in MOSAIKS [41] prioritize model efficiency, accessibility, and speed of retraining to make this type of analysis lightweight. Projects like MOSAIKS evidence the value of re-evaluating AI performance metrics beyond accuracy to include usability metrics: *How expensive is the model to train (including hyperparameter tuning)? How expensive is the model to use? How do we help users determine model uncertainty? In what percentage of the time do humans need to verify the results to have a system that is reliable enough for use?*

4.3.1 Making the costs of AI transparent

Researchers and practitioners interested in applying AI to conservation decision making must account for the hidden costs of AI in deciding what to integrate into their practices. One important step forward is making the full costs of developing and maintaining AI solutions more transparent from the outset. There are hidden costs across the entire pipeline of AI development:

Data collection and analysis Data collection is costly and includes the purchase of sensors or data, acquisition of labels, management of the data, storage of the data, and

accessibility and standardization of the data. Additionally, cost of preliminary data analysis for ML readiness should be accounted for, including data science capacity cost and computational capacity cost.

Model development We must consider the cost of ML prototyping. This can be significantly more expensive if a new question is being asked and no existing labeled datasets, best practices for model training, or pretrained models are available for this task. Costs to be considered include GPU costs of training, which increase with data size, model size, and model robustness to hyperparameters. Note also that reinforcement learning and meta learning are much more expensive computationally than training a classification model. Smaller machine learning architectures, such as ResNet, are much cheaper to train and evaluate from both a size and robustness perspective than larger architectures.

System infrastructure Costly data infrastructure development is often necessary to scale the use of ML beyond a prototype. *How will the data get from your collection sites to the model? Where will the data be stored (locally or on the cloud), where will the model be stored, what are the costs of data movement, and what temporal latency do you need for your problem?* Larger models are also more costly to use for inference, either on the cloud or through the purchase of additional GPUs, and this increased cost should be weighed against performance gains that may be achieved with larger models.

Evaluation We must both plan for and account for the costs of persistent quality control. *How will a team verify and analyze performance of the model to assess bias? If the model is no longer working well, what will you do? Do you need to retrain your model periodically?* If this is the case, we must carefully consider what data and decision making infrastructure are needed to retrain models, and account for these costs in our planning.

Human resources Successful impact-driven, interdisciplinary collaboration requires extensive communication of both expectations and results. Human expertise ought to be integrated in the development of AI models [144]. If an AI partner pulls out of a collaboration before delivering a usable conservation technology tool to the end beneficiary, beyond breaching of trust, there is a reputational cost for other AI researchers. *If one of the players in the development of these models pulls out, what is the sunk cost for the other contributors and is that cost recoverable?* For example, if a conservation organization sinks considerable resources into ML data curation and then a usable model cannot be developed, *is that effort aligned with better data infrastructure and management for that organization in a way that is useful even without ML?*

Environmental Finally, all of the above costs are both financial and environmental. Both computational resources and data storage use power, which has an associated carbon and natural resource cost. We must be aware of the carbon and nature costs of the systems we build and account for these in our metrics and in our decision-making processes [145].

4.4 Guiding question: How to design AI that aids humans in conservation decision making?

Our talks and discussions illuminated a number of recurring themes on the ways in which the demands of human decision making require solutions that move beyond simply producing predictions or even suggesting future actions:

1. The need for understandable, computationally inexpensive, user-friendly, and easily deployed procedures.

2. The need for quantitative assessments of actual downstream conservation impact to judge the effectiveness of AI tools and interventions.
3. The need to determine which stakeholders or rights holders make which decisions and what form of information they need to make those decisions [146].
4. The importance of accounting for financial, social, and human effort costs in selecting which AI tools to deploy.
5. Model interpretability arose as an important factor for model usability and value in the hands of stakeholders [147].
6. Model bias, fairness, and trustworthiness all impact long term user adoption and satisfaction.

Each of these motifs in isolation highlights important questions, but, importantly, AI for conservation requires the *intersection* of these considerations. Moreover, important cultural changes are needed in the AI discipline to incentivize and reward the work needed to address these challenges.

AI must aid decision making in conservation, not replace the expert decision process with fully-automated systems. Due to the high risk associated with errors that AI for conservation systems often face, similar to AI for healthcare or self driving cars, there is a need for additional research and development of participatory systems that enable human-AI collaboration [148, 149, 150]. These systems and methods seek not to replace human expertise but instead to augment it, making efficient use of human intelligence and intervention and enabling a single expert to process far more — and far more complex — data than was previously possible.

One example of such a system is ElephantBook [151], a system which helps a small group of experts in the field monitor an elephant population of over 2,000 individuals by speeding up the visual re-identification process. Rangers and local experts photograph elephants when they are sighted, and the system combines expert individual codes with modern computer vision methods to find the best possible matches in the database of local known individuals (which includes over 50K images). Human experts sort through the possible matches and verify them, making the most of both machine efficiency and human ability to generalize and recognize patterns from only a few examples. This again points to a need to broaden our range of AI evaluation metrics, as a successful system should jointly meet accuracy goals needed for the task while minimizing the use of human time and resources.

Finally, one major remaining challenge is knowing when and how much to trust AI predictions or recommendations, particularly given that these models often have the most *potential* for added value when deployed outside of the distribution on which they were trained. We need to invest research cycles on determining both how best to estimate uncertainty and how to effectively communicate those uncertainty estimates. Here we can learn from adjacent domains: these types of uncertainty estimates and how they are best communicated have already been explored in weather forecasts [152] and in climate models [153].

5 Looking forward to the future of AI-assisted decision making for conservation

The 2022 AI-Assisted Decision Making for Conservation workshop convened to explore opportunities to expand the breadth of AI methods — particularly for optimization, planning, and evaluation — that are being developed for and applied to conservation. Our participants came from across academia, industry, and non-profits with expertise in computation, conservation, and deploying research. Together, we shared our expertise and hopes for the field of AI-assisted decision making in conservation, as it progresses.¹

We end by highlighting forward-looking considerations on the interconnectedness of different conservation management goals in 5.1 and summarizing prominent goals and ideas for moving forward in section 5.2.

5.1 Intersectionality of conservation management goals and its implications for useful AI

In discussing AI for conservation, it is important to simultaneously consider social and financial concerns rather than purely focusing on just environmental goals. Environmental goals should be socially viable. As a specific example, success in wildlife conservation efforts may lead to increased human-wildlife conflicts [154]; such costs must be considered and mitigated to the extent possible.

Similarly, environmental goals should be financially viable. Sustainable solutions in AI for conservation would allow us to simultaneously address all three concerns (see figure 3). Achieving these simultaneous goals requires deep partnership with local communities, non-profits, governments, and cannot be achieved by AI researchers in isolation, armed only with datasets.

It is unrealistic to expect to be able to assess all intersectional costs and interactions beforehand. Rather, we might develop and rely upon key partnerships to continually reassess the impact of the AI deployment. A staged approach to deployment with continuous assessment may help mitigate or address emerging financial or social challenges.

5.2 A vision for the future directions in AI-assisted conservation

To summarize, AI is being used widely for conservation [23, 155, 156], and holds promise for even more uses across a vast and intersectional range of conservation efforts. As we expand new frontiers of AI for conservation, it is useful to orient this vision around the following questions:

- What questions in conservation are answerable and what decisions are supportable with AI? Including:
 - When and how can AI approaches effectively assist conservation decision making and prioritization?
 - When and how can AI help assess conservation impacts and outcomes?
- What are the full costs associated with developing and deploying AI systems for conservation? How do these compare to the intended benefits?

¹Some ideas and goals were shared by individuals, some in small groups, and some were common to all participants in our group. Here we summarize and synthesize these points without necessarily differentiating between these cases.

Acknowledgments

We thank Carl Boettiger (UC Berkeley), Iadine Chadés (CSIRO), Ross Tyzack Pitman (Panthera), and Michael C. Runge (USGS) for helpful discussions prior to the workshop.

This workshop was funded by the Center for Research on Computation and Society at Harvard University. Thank you to Emma Hammack and Hila Bernstein of CRCS for all your support!

References

- [1] Bradley J. Cardinale, J. Emmett Duffy, Andrew Gonzalez, David U. Hooper, Charles Perrings, Patrick Venail, Anita Narwani, Georgina M. Mace, David Tilman, David A. Wardle, Ann P. Kinzig, Gretchen C. Daily, Michel Loreau, James B. Grace, Anne Larigauderie, Diane S. Srivastava, and Shahid Naeem. Biodiversity loss and its impact on humanity. *Nature*, 486(7401):59–67, 2012.
- [2] Gerardo Ceballos, Paul R Ehrlich, Anthony D Barnosky, Andrés García, Robert M Pringle, and Todd M Palmer. Accelerated modern human-induced species losses: Entering the sixth mass extinction. *Science Advances*, 1(5):e1400253, 2015.
- [3] Kendall R Jones, Oscar Venter, Richard A Fuller, James R Allan, Sean L Maxwell, Pablo Jose Negret, and James EM Watson. One-third of global protected land is under intense human pressure. *Science*, 360(6390):788–791, 2018.
- [4] James EM Watson, Jeremy S Simmonds, Divya Narain, Michelle Ward, Martine Maron, and Sean L Maxwell. Talk is cheap: Nations must act now to achieve long-term ambitions for biodiversity. *One Earth*, 4(7):897–900, 2021.
- [5] Michael Williams. *Deforesting the earth: from prehistory to global crisis*. University of Chicago Press, 2003.
- [6] Jeff Tollefson. Humans are driving one million species to extinction. *Nature*, 569(7755):171–172, 2019.
- [7] James EM Watson, Nigel Dudley, Daniel B Segan, and Marc Hockings. The performance and potential of protected areas. *Nature*, 515(7525):67–73, 2014.
- [8] Lauren Coad, Fiona Leverington, Kathryn Knights, Jonas Geldmann, April Eassom, Valerie Kapos, Naomi Kingston, Marcelo de Lima, Camilo Zamora, Ivon Cuadros, Christoph Nolte, Neil D. Burgess, and Marc Hockings. Measuring impact of protected area management interventions: Current and future use of the global database of protected area management effectiveness. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1681):20140281, 2015.
- [9] Iadine Chadès, Eve McDonald-Madden, Michael A McCarthy, Brendan Wintle, Matthew Linkie, and Hugh P Possingham. When to stop managing or surveying cryptic threatened species. *Proceedings of the National Academy of Sciences*, 105(37):13936–13940, 2008.
- [10] Richard Schuster, Scott Wilson, Amanda D. Rodewald, Peter Arcese, Daniel Fink, Tom Auer, and Joseph. R. Bennett. Optimizing the conservation of migratory species over their full annual cycle. *Nature Communications*, 10(1):1–8, 2019.
- [11] Ian R Ball, Hugh P Possingham, and Matthew Watts. Marxan and relatives: Software for spatial conservation prioritisation. *Spatial conservation prioritisation: Quantitative methods and computational tools*, pages 185–195, 2009.
- [12] Alan D Fox, David W Corne, C Gabriela Mayorga Adame, Jeff A Polton, Lea-Anne Henry, and J Murray Roberts. An efficient multi-objective optimization method for use in the design of marine protected area networks. *Frontiers in Marine Science*, 6:17, 2019.
- [13] Edmond Sacre, Rebecca Weeks, Michael Bode, and Robert L Pressey. The relative conservation impact of strategies that prioritize biodiversity representation, threats, and protection costs. *Conservation Science and Practice*, 2(8):e221, 2020.
- [14] Sally Dowd, Melissa Chapman, Laura E Koehn, and Porter Hoagland. The economic trade-offs and ecological impacts associated with a potential mesopelagic fishery in the California Current. *Ecological Applications*, page e2578, 2022.
- [15] Piero Visconti, Robert L Pressey, Daniel B Segan, and Brendan A Wintle. Conservation planning with dynamic threats: The role of spatial design and priority setting for species’ persistence. *Biological Conservation*, 143(3):756–767, 2010.

- [16] Silvia B Carvalho, Guillermo Velo-Anton, Pedro Tarroso, Ana Paula Portela, Mafalda Barata, Salvador Carranza, Craig Moritz, and Hugh P Possingham. Spatial conservation prioritization of biodiversity spanning the evolutionary continuum. *Nature Ecology & Evolution*, 1(6):1–8, 2017.
- [17] Robert L Pressey, Morena Mills, Rebecca Weeks, and Jon C Day. The plan of the day: Managing the dynamic transition from regional conservation designs to local conservation actions. *Biological Conservation*, 166:155–169, 2013.
- [18] Vivitskaia J Tulloch, Hugh P Possingham, Stacy D Jupiter, Chris Roelfsema, Ayesha IT Tulloch, and Carissa J Klein. Incorporating uncertainty associated with habitat data in marine reserve design. *Biological Conservation*, 162:41–51, 2013.
- [19] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [20] Cameron B Browne, Edward Powley, Daniel Whitehouse, Simon M Lucas, Peter I Cowl- ing, Philipp Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. A survey of Monte Carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games*, 4(1):1–43, 2012.
- [21] Judea Pearl. Causal inference in statistics: An overview. *Statistics Surveys*, 3:96–146, 2009.
- [22] Lauren Coad, James EM Watson, Jonas Geldmann, Neil D Burgess, Fiona Leverington, Marc Hockings, Kathryn Knights, and Moreno Di Marco. Widespread shortfalls in protected area resourcing undermine efforts to conserve biodiversity. *Frontiers in Ecology and the Environment*, 17(5):259–264, 2019.
- [23] Devis Tuia, Benjamin Kellenberger, Sara Beery, Blair R. Costelloe, Silvia Zuffi, Benjamin Risse, Alexander Mathis, Mackenzie W. Mathis, Frank van Langevelde, Tilo Burghardt, Roland Kays, Holger Klinck, Martin Wikelski, Iain D. Couzin, Grant van Horn, Margaret C. Crofoot, Charles V. Stewart, and Tanya Berger-Wolf. Perspectives in machine learning for wildlife conservation. *Nature Communications*, 13(1):1–15, 2022.
- [24] Ben G Weinstein. A computer vision for animal ecology. *Journal of Animal Ecology*, 87(3): 533–545, 2018.
- [25] Rongjun Qin and Tao Liu. A review of landcover classification with very-high resolution remotely sensed optical images—analysis unit, model scalability and transferability. *Remote Sensing*, 14(3):646, 2022.
- [26] Antoine Guisan, Reid Tingley, John B. Baumgartner, Ilona Naujokaitis-Lewis, Patricia R. Sutcliffe, Ayesha I. T. Tulloch, Tracey J. Regan, Lluís Brotons, Eve McDonald-Madden, Chrystal Mantyka-Pringle, Tara G. Martin, Jonathan R. Rhodes, Ramona Maggini, Samantha A. Setterfield, Jane Elith, Mark W. Schwartz, Brendan A. Wintle, Olivier Broennimann, Mike Austin, Simon Ferrier, Michael R. Kearney, Hugh P. Possingham, and Yvonne M. Buckley. Predicting species distributions for conservation decisions. *Ecology Letters*, 16(12): 1424–1435, 2013.
- [27] Colin M Beale and Jack J Lennon. Incorporating uncertainty in predictive species distribution modelling. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1586): 247–258, 2012.
- [28] Karen M Bailey, Robert A McCleery, Michael W Binford, and Christa Zweig. Land-cover change within and around protected areas in a biodiversity hotspot. *Journal of Land Use Science*, 11(2):154–176, 2016.
- [29] Timothy Kuiper, Andrew J Loveridge, and David W Macdonald. Robust mapping of human–wildlife conflict: Controlling for livestock distribution in carnivore depredation models. *Animal Conservation*, 25(2):195–207, 2022.
- [30] Andrew S Goudie. *Human impact on the natural environment*. John Wiley & Sons, 2018.

- [31] Rob Critchlow, Andrew J Plumptre, Bazil Alidria, Mustapha Nsubuga, Margaret Driciru, Aggrey Rwetsiba, F Wanyama, and Colin M Beale. Improving law-enforcement effectiveness and efficiency in protected areas using ranger-collected monitoring data. *Conservation Letters*, 10(5):572–580, 2017.
- [32] Meredith L Gore, Emily Griffin, Bistra Dilkina, Aaron Ferber, Stanley E Griffis, Burcu B Keskin, and John Macdonald. Advancing interdisciplinary science for disrupting wildlife trafficking networks. *Proceedings of the National Academy of Sciences*, 120(10):e2208268120, 2023.
- [33] Milad Memarzadeh, Gregory L Britten, Boris Worm, and Carl Boettiger. Rebuilding global fisheries under uncertainty. *Proceedings of the National Academy of Sciences*, 116(32):15985–15990, 2019.
- [34] Alexander Teytelboym. Natural capital market design. *Oxford Review of Economic Policy*, 35(1):138–161, 2019.
- [35] Seema Jayachandran, Joost De Laat, Eric F Lambin, Charlotte Y Stanton, Robin Audy, and Nancy E Thomas. Cash for carbon: A randomized trial of payments for ecosystem services to reduce deforestation. *Science*, 357(6348):267–273, 2017.
- [36] Simon A. Levin, Tasos Xepapadeas, Anne-Sophie Crépin, Jon Norberg, Aart de Zeeuw, Carl Folke, Terry P. Hughes, Kenneth J. Arrow, Scott Barrett, Gretchen C. Daily, Paul R. Ehrlich, Nils Kautsky, K-G. Mäler, Stephen Polasky, Max Troell, Jeffrey R. Vincent, and Brian Walker. Social-ecological systems as complex adaptive systems: Modeling and policy implications. *Environment and Development Economics*, 18(2):111–132, 2013.
- [37] Steven F Wilson, Thomas D Nudds, and Andrew de Vries. A causal modelling approach to informing woodland caribou conservation policy from observational studies. *Biological Conservation*, 264:109370, 2021.
- [38] Philip G Brodrick, Andrew B Davies, and Gregory P Asner. Uncovering ecological patterns with convolutional neural networks. *Trends in Ecology & Evolution*, 34(8):734–745, 2019.
- [39] David B Lobell, David Thau, Christopher Seifert, Eric Engle, and Bertis Little. A scalable satellite-based crop yield mapper. *Remote Sensing of Environment*, 164:324–333, 2015.
- [40] Matthew C Hansen, Peter V Potapov, Rebecca Moore, Matt Hancher, Svetlana A Turubanova, Alexandra Tyukavina, David Thau, Stephen V Stehman, Scott J Goetz, Thomas R Loveland, Anil Kommareddy, Alexey Egorov, L Chini, C.O. Justice, and J Townshend. High-resolution global maps of 21st-century forest cover change. *Science*, 342(6160):850–853, 2013.
- [41] Esther Rolf, Jonathan Proctor, Tamma Carleton, Ian Bolliger, Vaishaal Shankar, Miyabi Ishihara, Benjamin Recht, and Solomon Hsiang. A generalizable and accessible approach to machine learning with global satellite imagery. *Nature Communications*, 12(1):1–11, 2021.
- [42] Ankush Khandelwal, Anuj Karpatne, Praveen Ravirathinam, Rahul Ghosh, Zhihao Wei, Hilary A Dugan, Paul C Hanson, and Vipin Kumar. ReaLSAT, a global dataset of reservoir and lake surface area variations. *Scientific Data*, 9(1):1–12, 2022.
- [43] Christopher F. Brown, Steven P. Brumby, Brookie Guzder-Williams, Tanya Birch, Samantha Brooks Hyde, Joseph Mazzariello, Wanda Czerwinski, Valerie J. Pasquarella, Robert Haertel, Simon Ilyushchenko, Kurt Schwehr, Mikaela Weisse, Fred Stolle, Craig Hanson, Oliver Guinan, Rebecca Moore, and Alexander M. Tait. Dynamic World, near real-time global 10 m land use land cover mapping. *Scientific Data*, 9(1):1–17, 2022.
- [44] Markus Reichstein, Gustau Camps-Valls, Bjorn Stevens, Martin Jung, Joachim Denzler, Nuno Carvalhais, and Prabhat. Deep learning and process understanding for data-driven Earth system science. *Nature*, 566(7743):195–204, 2019.
- [45] Benjamin Kellenberger, Diego Marcos, and Devis Tuia. When a few clicks make all the difference: Improving weakly-supervised wildlife detection in UAV images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2019.

- [46] Mayank Lahiri, Chayant Tantipathananandh, Rosemary Warungu, Daniel I Rubenstein, and Tanya Y Berger-Wolf. Biometric animal databases from field photographs: Identification of individual zebra in the wild. In *Proceedings of the 1st ACM International Conference on Multimedia Retrieval*, pages 1–8, 2011.
- [47] Ross T. Pitman, Julien Fattebert, Samuel T. Williams, Kathryn S. Williams, Russell A. Hill, Luke T. B. Hunter, Hugh Robinson, John Power, Lourens Swanepoel, Rob Slotow, and Guy A. Balme. Cats, connectivity and conservation: Incorporating data sets and integrating scales for wildlife management. *Journal of Applied Ecology*, 54(6):1687–1698, 2017.
- [48] Larissa Sayuri Moreira Sugai, Thiago Sanna Freire Silva, Jose Wagner Ribeiro Jr, and Diego Llusia. Terrestrial passive acoustic monitoring: Review and perspectives. *BioScience*, 69(1): 15–25, 2019.
- [49] Ella Browning, Mark Bolton, Ellie Owen, Akiko Shoji, Tim Guilford, and Robin Freeman. Predicting animal behaviour using deep learning: GPS data alone accurately predict diving in seabirds. *Methods in Ecology and Evolution*, 9(3):681–692, 2018.
- [50] Mykel J Kochenderfer, Tim A Wheeler, and Kyle H Wray. *Algorithms for Decision Making*. MIT Press, 2022.
- [51] Daniel W Beyers. Causal inference in environmental impact studies. *Journal of the North American Benthological Society*, 17(3):367–373, 1998.
- [52] Paul J Ferraro, James N Sanchirico, and Martin D Smith. Causal inference in coupled human and natural systems. *Proceedings of the National Academy of Sciences*, 116(12):5311–5318, 2019.
- [53] Van Butsic, David J Lewis, Volker C Radeloff, Matthias Baumann, and Tobias Kuemmerle. Quasi-experimental methods enable stronger inferences from observational data in ecology. *Basic and Applied Ecology*, 19:1–10, 2017.
- [54] Andrew B Davies and Gregory P Asner. Elephants limit aboveground carbon gains in African savannas. *Global Change Biology*, 25(4):1368–1382, 2019.
- [55] Fikret Berkes. Community conserved areas: Policy issues in historic and contemporary context. *Conservation Letters*, 2(1):20–25, 2009.
- [56] Nigel Dudley, Holly Jonas, Fred Nelson, Jeffrey Parrish, Aili Pyhälä, Sue Stolton, and James EM Watson. The essential role of other effective area-based conservation measures in achieving big bold conservation targets. *Global Ecology and Conservation*, 15:e00424, 2018.
- [57] Eric Dinerstein, Carly Vynne, Enric Sala, Anup R Joshi, Sanjiv Fernando, Thomas E Lovejoy, Juan Mayorga, David Olson, Gregory P Asner, Jonathan EM Baillie, Neil D Burgess, Karl Burkart, Reed F Noss, Yaping P Zhang, Alessandro Baccini, Tanya Birch, Nathan Hahn, Lucas N Joppa, and Eric Wikramanayake. A global deal for nature: Guiding principles, milestones, and targets. *Science Advances*, 5(4), 2019.
- [58] Sean L. Maxwell, Victor Cazalis, Nigel Dudley, Michael Hoffmann, Ana S. L. Rodrigues, Sue Stolton, Piero Visconti, Stephen Woodley, Naomi Kingston, Edward Lewis, Martine Maron, Bernardo B. N. Strassburg, Amelia Wenger, Harry D. Jonas, Oscar Venter, and James E. M. Watson. Area-based conservation in the twenty-first century. *Nature*, 586(7828):217–227, 2020.
- [59] Christine E. Wilkinson, Alex McInturff, Jennifer R. B. Miller, Veronica Yovovich, Kaitlyn M. Gaynor, Kendall Calhoun, Harshad Karandikar, Jeff Vance Martin, Phoebe Parker-Shames, Avery Shawler, Amy Van Scoyoc, and Justin S. Brashares. An ecological framework for contextualizing carnivore–livestock conflict. *Conservation Biology*, 34(4):854–867, 2020.
- [60] David C Marvin, Lian Pin Koh, Antony J Lynam, Serge Wich, Andrew B Davies, Ramesh Krishnamurthy, Emma Stokes, Ruth Starkey, and Gregory P Asner. Integrating technologies for scalable ecology and conservation. *Global Ecology and Conservation*, 7:262–275, 2016.

- [61] David Thau, Jorge A Ahumada, Tanya Birch, Eric Fegraus, Nicole Flores, Walter Jetz, Roland Kays, Margaret Kinnaird, Sayali Kulkarni, Arnaud Lyet, Timothy G O'Brien, Jonathan Palmer, Stephanie Schuttler, Michelle Duong, Ruth Oliver, Jennifer Y Zhao, and William J McShea. Artificial intelligence's role in global camera trap data management and analytics via wildlife insights. *Biodiversity Information Science and Standards*, 3:e38233, 2019.
- [62] Jorge Ahumada, Eric Fegraus, Tanya Birch, Nicole Flores, Roland Kays, Timothy O'Brien, Jonathan Palmer, Stephanie Schuttler, Jennifer Zhao, Walter Jetz, Margaret Kinnaird, Sayali Kulkarni, Arnaud Lyet, David Thau, Michelle Duong, Ruth Oliver, and Anthony Dancer. Wildlife Insights: A platform to maximize the potential of camera trap and other passive sensor wildlife data for the planet. *Environmental Conservation*, 47(1):1–6, 2020.
- [63] Talia Speaker, Stephanie O'Donnell, George Wittemyer, Brett Bruyere, Colby Loucks, Anthony Dancer, Marianne Carter, Eric Fegraus, Jonathan Palmer, Ellie Warren, and Jennifer Solomon. A global community-sourced assessment of the state of conservation technology. *Conservation Biology*, 2022.
- [64] Richard Schuster, Ryan R Germain, Joseph R Bennett, Nicholas J Reo, and Peter Arcese. Vertebrate biodiversity on indigenous-managed lands in Australia, Brazil, and Canada equals that in protected areas. *Environmental Science & Policy*, 101:1–6, 2019.
- [65] Meredith L Gore, Michelle L Lute, Jonah H Ratsimbazafy, and Andry Rajaonson. Local perspectives on environmental insecurity and its influence on illegal biodiversity exploitation. *PLoS One*, 11(4):e0150337, 2016.
- [66] Helen U Agu and Meredith L Gore. Women in wildlife trafficking in Africa: A synthesis of literature. *Global Ecology and Conservation*, 23, 2020.
- [67] Peter E Smouse, Stefano Focardi, Paul R Moorcroft, John G Kie, James D Forester, and Juan M Morales. Stochastic modelling of animal movement. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1550):2201–2211, 2010.
- [68] Lou Ann Dietz, Marcia Brown, and Vinaya Swaminathan. Increasing the impact of conservation projects. *American Journal of Primatology*, 72(5):425–440, 2010.
- [69] Kent H Redford, Kristin B Hulvey, Matthew A Williamson, and Mark W Schwartz. Assessment of the Conservation Measures Partnership's effort to improve conservation outcomes through adaptive management. *Conservation Biology*, 32(4):926–937, 2018.
- [70] Rosie Cooney, Dilys Roe, Holly Dublin, Jacob Phelps, David Wilkie, Aidan Keane, Henry Travers, Diane Skinner, Daniel WS Challender, James R Allan, and Duan Biggs. From poachers to protectors: Engaging local communities in solutions to illegal wildlife trade. *Conservation Letters*, 10(3):367–374, 2017.
- [71] Charles A Emogor, Daniel J Ingram, Lauren Coad, Thomas A Worthington, Andrew Dunn, Inaoyom Imong, and Andrew Balmford. The scale of Nigeria's involvement in the transnational illegal pangolin trade: Temporal and spatial patterns and the effectiveness of wildlife trade regulations. *Biological Conservation*, 264:109365, 2021.
- [72] Angela Gaylard and Sam Ferreira. Advances and challenges in the implementation of strategic adaptive management beyond the Kruger National Park-making linkages between science and biodiversity management. *Koedoe: African Protected Area Conservation and Science*, 53(2): 1–8, 2011.
- [73] Joseph R Bennett, Justine D Shaw, Aleks Terauds, John P Smol, Rien Aerts, Dana M Bergstrom, Jules M Blais, William WL Cheung, Steven L Chown, Mary-Anne Lea, Uffe N Nielsen, Daniel Pauly, Kenneth J Reimer, Martin J Riddle, Ian Snape, Jonathan S Stark, Vivitskaia J Tulloch, and Hugh P Possingham. Polar lessons learned: Long-term management based on shared threats in Arctic and Antarctic environments. *Frontiers in Ecology and the Environment*, 13(6):316–324, 2015.
- [74] Lynn H Kaack, Priya L Donti, Emma Strubell, George Kamiya, Felix Creutzig, and David Rolnick. Aligning artificial intelligence with climate change mitigation. *Nature Climate Change*, 12(6):518–527, 2022.

- [75] Keith R Hayes, Geoffrey R Hosack, Emma Lawrence, Paul Hedge, Neville S Barrett, Rachel Przeslawski, M Julian Caley, and Scott D Foster. Designing monitoring programs for marine protected areas within an evidence based decision making paradigm. *Frontiers in Marine Science*, 6:746, 2019.
- [76] Nathan J. Bennett, Robin Roth, Sarah C. Klain, Kai Chan, Patrick Christie, Douglas A. Clark, Georgina Cullman, Deborah Curran, Trevor J. Durbin, Graham Epstein, Alison Greenberg, Michael P Nelson, John Sandlos, Richard Stedman, Tara L Teel, Rebecca Thomas, Diogo Veríssimo, and Carina Wyborn. Conservation social science: Understanding and integrating human dimensions to improve conservation. *Biological Conservation*, 205:93–108, 2017.
- [77] William M Adams, Ros Aveling, Dan Brockington, Barney Dickson, Jo Elliott, Jon Hutton, Dilys Roe, Bhaskar Vira, and William Wolmer. Biodiversity conservation and the eradication of poverty. *Science*, 306(5699):1146–1149, 2004.
- [78] Adrian Treves, Robert B Wallace, Lisa Naughton-Treves, and Andrea Morales. Co-managing human–wildlife conflicts: A review. *Human Dimensions of Wildlife*, 11(6):383–396, 2006.
- [79] Justin S Brashares, Peter Arcese, Moses K Sam, Peter B Coppolillo, and Andrew Balmford. Bushmeat hunting, wildlife declines, and fish supply in West Africa. *Science*, 306(5699):1180–1183, 2004.
- [80] Peggy M Shepard. Advancing environmental justice through community-based participatory research. *Environmental Health Perspectives*, 110(suppl 2):139–139, 2002.
- [81] Michalis Vitos, Jerome Lewis, Matthias Stevens, and Muki Haklay. Making local knowledge matter: Supporting non-literate people to monitor poaching in Congo. In *Proceedings of the 3rd ACM Symposium on Computing for Development*, pages 1–10, 2013.
- [82] Elizabeth Bondi, Lily Xu, Diana Acosta-Navas, and Jackson A Killian. Envisioning communities: A participatory approach towards AI for social good. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 425–436, 2021.
- [83] Brian L Sullivan, Christopher L Wood, Marshall J Iliff, Rick E Bonney, Daniel Fink, and Steve Kelling. eBird: A citizen-based bird observation network in the biological sciences. *Biological Conservation*, 142(10):2282–2292, 2009.
- [84] Grant Van Horn, Oisín Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The iNaturalist species classification and detection dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8769–8778, 2018.
- [85] Hugo Rainey, Antony Lynam, Richard Bergl, Drew Cronin, Chris Gordon, Jonathan Palmer, and Emma Stokes. Integrated technology for improving effectiveness of management of poaching in protected areas. 2021.
- [86] John Elkington. Accounting for the triple bottom line. *Measuring Business Excellence*, 1998.
- [87] Fei Fang, Milind Tambe, Bistra Dilikina, and Andrew J Plumptre. *Artificial Intelligence and Conservation*. Cambridge University Press, 2019.
- [88] Roberta Kwok. AI empowers conservation biology. *Nature*, 567(7746):133–134, 2019.
- [89] Helen J Mayfield, Carl Smith, Marcus Gallagher, and Marc Hockings. Considerations for selecting a machine learning technique for predicting deforestation. *Environmental Modelling & Software*, 131:104741, 2020.
- [90] Gyri Reiersen, David Dao, Björn Lütjens, Konstantin Klemmer, Kenza Amara, Attila Steinegger, Ce Zhang, and Xiaoxiang Zhu. ReforesTree: A dataset for estimating tropical forest carbon stock with deep learning and aerial imagery. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 12119–12125, 2022.

- [91] Caleb Scoville, Melissa Chapman, Razvan Amironesei, and Carl Boettiger. Algorithmic conservation in a changing climate. *Current Opinion in Environmental Sustainability*, 51:30–35, 2021.
- [92] Brian Groenke, Luke Madaus, and Claire Monteleoni. Climalign: Unsupervised statistical downscaling of climate variables via normalizing flows. In *Proceedings of the 10th International Conference on Climate Informatics, CI2020*, page 60–66, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450388481.
- [93] A.D.M. Dobson, E.J. Milner-Gulland, Nicholas J. Aebischer, Colin M. Beale, Robert Brozovic, Peter Coals, Rob Critchlow, Anthony Dancer, Michelle Greve, Amy Hinsley, Harriet Ibbett, Alison Johnston, Timothy Kuiper, Steven Le Comber, Simon P. Mahood, Jennifer F. Moore, Erlend B. Nilsen, Michael J.O. Pocock, Anthony Quinn, Henry Travers, Paulo Wilfred, Joss Wright, and Aidan Keane. Making messy data work for conservation. *One Earth*, 2(5):455–465, 2020.
- [94] JF Moore, WE Pine, F Mulindahabi, P Niyigaba, G Gatorano, MK Masozera, and L Beaudrot. Comparison of species richness and detection between line transects, ground camera traps, and arboreal camera traps. *Animal Conservation*, 23(5):561–572, 2020.
- [95] Gencer Sumbul, Ramazan Gokberk Cinbis, and Selim Aksoy. Fine-grained object recognition and zero-shot learning in remote sensing imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 56(2):770–779, 2017.
- [96] Hanna Meyer and Edzer Pebesma. Machine learning-based global maps of ecological variables and the challenge of assessing them. *Nature Communications*, 13(1):2208, 2022.
- [97] Esther Rolf. Evaluation challenges for geospatial ML. In *ICLR 2023 Workshop on Machine Learning for Remote Sensing*, 2023.
- [98] Rui Wang, Karthik Kashinath, Mustafa Mustafa, Adrian Albert, and Rose Yu. Towards physics-informed deep learning for turbulent flow prediction. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1457–1466, 2020.
- [99] Konstantin Klemmer, Tianlin Xu, Beatrice Acciaio, and Daniel B Neill. SPATE-GAN: Improved generative modeling of dynamic spatio-temporal patterns with an autoregressive embedding loss. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 4523–4531, 2022.
- [100] Gabriel Tseng, Ivan Zvonkov, Mirali Purohit, David Rolnick, and Hannah Kerner. Lightweight, pre-trained transformers for remote sensing timeseries. *arXiv preprint arXiv:2304.14065*, 2023.
- [101] Saumya Sinha, Sophie Giffard-Roisin, Fatima Karbou, Michael Deschatres, Anna Karas, Nicolas Eckert, Cécile Coléou, and Claire Monteleoni. Variational autoencoder anomaly-detection of avalanche deposits in satellite sar imagery. In *Proceedings of the 10th International Conference on Climate Informatics, CI2020*, page 113–119, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450388481.
- [102] Esther Rolf, Nikolay Malkin, Alexandros Graikos, Ana Jojic, Caleb Robinson, and Nebojsa Jojic. Resolving label uncertainty with implicit posterior models. In *Proceedings of the 38th Conference on Uncertainty in Artificial Intelligence*, 2022.
- [103] Yifang Yin, Ying Zhang, Zhenguang Liu, Sheng Wang, Rajiv Ratn Shah, and Roger Zimmermann. GPS2Vec: Pre-trained semantic embeddings for worldwide GPS coordinates. *IEEE Transactions on Multimedia*, 24:890–903, 2021.
- [104] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé Iii, and Kate Crawford. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92, 2021.
- [105] Alexandra Sasha Luccioni and David Rolnick. Bugs in the data: How ImageNet misrepresents biodiversity. *arXiv preprint arXiv:2208.11695*, 2022.

- [106] Shreya Shankar, Yoni Halpern, Eric Breck, James Atwood, Jimbo Wilson, and D Sculley. No classification without representation: Assessing geodiversity issues in open data sets for the developing world. *arXiv preprint arXiv:1711.08536*, 2017.
- [107] Miao Zhang, Harvineet Singh, Lazarus Chok, and Rumi Chunara. Segmenting across places: The need for fair transfer learning with satellite imagery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2916–2925, 2022.
- [108] Alex Asase, Tiwonge I Mzumara-Gawa, Jesse O Owino, Andrew T Peterson, and Erin Saupe. Replacing “parachute science” with “global science” in ecology and conservation biology. *Conservation Science and Practice*, 4(5), 2022.
- [109] Catherine Nakalembe and Hannah Kerner. Considerations for AI-EO for agriculture in Sub-Saharan Africa. *Environmental Research Letters*, 18(4):041002, 2023.
- [110] José Luis Mena, Hiromi Yagui, Vania Tejada, Emilio Bonifaz, Eva Bellemain, Alice Valentini, Mathias W Tobler, Pamela Sánchez-Vendizú, and Arnaud Lyet. Environmental DNA metabarcoding as a useful tool for evaluating terrestrial mammal diversity in tropical forests. *Ecological Applications*, 31(5):e02335, 2021.
- [111] Lily Xu, Arpita Biswas, Fei Fang, and Milind Tambe. Ranked prioritization of groups in combinatorial bandit allocation. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence*, 2022.
- [112] Edward T Game, Erik Meijaard, Douglas Sheil, and Eve McDonald-Madden. Conservation in a wicked complex world; challenges and solutions. *Conservation Letters*, 7(3):271–277, 2014.
- [113] Diogo Alagador, Jorge Orestes Cerdeira, and Miguel Bastos Araújo. Shifting protected areas: Scheduling spatial priorities under climate change. *Journal of Applied Ecology*, 51(3):703–713, 2014.
- [114] Richard Schuster, Rachel Buxton, Jeffrey O. Hanson, Allison D. Binley, Jeremy Pittman, Vivitskaia Tulloch, Frank A. La Sorte, Patrick R. Roehrdanz, Peter H. Verburg, Amanda D. Rodewald, Scott Wilson, Hugh P. Possingham, and Joseph R. Bennett. Protected area planning to conserve biodiversity in an uncertain future. *Conservation Biology*, 2022.
- [115] Lily Xu, Elizabeth Bondi, Fei Fang, Andrew Perrault, Kai Wang, and Milind Tambe. Dual-mandate patrols: Multi-armed bandits for green security. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14974–14982, 2021.
- [116] Fei Fang, Thanh Nguyen, Rob Pickles, Wai Lam, Gopalasamy Clements, Bo An, Amandeep Singh, Milind Tambe, and Andrew Lemieux. Deploying PAWS: Field optimization of the Protection Assistant for Wildlife Security. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, pages 3966–3973, 2016.
- [117] Lily Xu, Shahrzad Gholami, Sara Marie Mc Carthy, Bistra Dilkina, Andrew Plumptre, Milind Tambe, Rohit Singh, Mustapha Nsubuga, Joshua Mabonga, Margaret Driciru, Fred Wanyama, Aggrey Rwetsiba, Tom Okello, and Eric Enyel. Stay ahead of poachers: Illegal wildlife poaching prediction and patrol planning under uncertainty with field test evaluations. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*, pages 1898–1901. IEEE, 2020.
- [118] Lily Xu, Andrew Perrault, Fei Fang, Haipeng Chen, and Milind Tambe. Robust reinforcement learning under minimax regret for green security. In *Proceedings of the 37th Conference on Uncertainty in Artificial Intelligence*, pages 257–267. PMLR, 2021.
- [119] Eric W Sanderson, Kent H Redford, Cheryl-Lesley B Chetkiewicz, Rodrigo A Medellin, Alan R Rabinowitz, John G Robinson, and Andrew B Taber. Planning to save a species: The jaguar as a model. *Conservation Biology*, 16(1):58–72, 2002.
- [120] Jeffrey O Hanson, Richard Schuster, Nina Morrell, Matthew Strimas-Mackey, ME Watts, P Arcese, J Bennett, and HP Possingham. prioritizr: Systematic conservation prioritization in R. *R package version*, 3(1.1), 2017.

- [121] Daniele Silvestro, Stefano Gorla, Thomas Sterner, and Alexandre Antonelli. Improving biodiversity protection through artificial intelligence. *Nature Sustainability*, 5(5):415–424, 2022.
- [122] Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. A brief survey of deep reinforcement learning. *arXiv preprint arXiv:1708.05866*, 2017.
- [123] Zachary Sunberg and Mykel Kochenderfer. Online algorithms for POMDPs with continuous state, action, and observation spaces. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 28, pages 259–263, 2018.
- [124] Samuel Nicol, Olivier Buffet, Takuya Iwamura, and Iadine Chades. Adaptive management of migratory birds under sea level rise. In *23rd International Joint Conference on Artificial Intelligence*, pages 2955–2957, 2013.
- [125] Iadine Chadès, Luz V Pascal, Sam Nicol, Cameron S Fletcher, and Jonathan Ferrer-Mestres. A primer on partially observable Markov decision processes (POMDPs). *Methods in Ecology and Evolution*, 12(11):2058–2072, 2021.
- [126] Marcus Lapeyrolerie, Melissa S Chapman, Kari EA Norman, and Carl Boettiger. Deep reinforcement learning for conservation decisions. *Methods in Ecology and Evolution*, 13(11):2649–2662, 2022.
- [127] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharmashan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419):1140–1144, 2018.
- [128] Mykel J Kochenderfer, Jessica E Holland, and James P Chryssanthacopoulos. Next-generation airborne collision avoidance system. Technical report, Massachusetts Institute of Technology-Lincoln Laboratory Lexington United States, 2012.
- [129] Yizheng Wang, Markus Zechner, John Michael Mern, Mykel J Kochenderfer, and Jef Karel Caers. A sequential decision-making framework with uncertainty quantification for ground-water management. *Advances in Water Resources*, 166:104266, 2022.
- [130] Priya Donti, Brandon Amos, and J Zico Kolter. Task-based end-to-end model learning in stochastic optimization. *Advances in Neural Information Processing Systems*, 30, 2017.
- [131] Anthony Corso, Robert Moss, Mark Koren, Ritchie Lee, and Mykel Kochenderfer. A survey of algorithms for black-box safety validation of cyber-physical systems. *Journal of Artificial Intelligence Research*, 72:377–428, 2021.
- [132] Ritchie Lee, Ole J Mengshoel, Anshu Saksena, Ryan W Gardner, Daniel Genin, Joshua Silbermann, Michael Owen, and Mykel J Kochenderfer. Adaptive stress testing: Finding likely failure events with reinforcement learning. *Journal of Artificial Intelligence Research*, 69:1165–1201, 2020.
- [133] Daniel F. Doak, James A. Estes, Benjamin S. Halpern, Ute Jacob, David R. Lindberg, James Lovvorn, Daniel H. Monson, M. Timothy Tinker, Terrie M. Williams, J. Timothy Wootton, Ian Carroll, Mark Emmerson, Fiorenza Micheli, and Mark Novak. Understanding and predicting ecological dynamics: Are major surprises inevitable? *Ecology*, 89(4):952–961, 2008.
- [134] Brendan A Wintle, Michael C Runge, and Sarah A Bekessy. Allocating monitoring effort in the face of unknown unknowns. *Ecology Letters*, 13(11):1325–1337, 2010.
- [135] Hugo Storm, Kathy Baylis, and Thomas Heckeley. Machine learning in agricultural and applied economics. *European Review of Agricultural Economics*, 47(3):849–892, 2020.
- [136] Alberto Garcia and Robert Heilmayr. Conservation impact evaluation using remotely sensed data. Available at SSRN 4179782, 2022.
- [137] Jonathan Proctor, Tamma Carleton, and Sandy Sum. Parameter recovery using remotely sensed variables. Technical report, National Bureau of Economic Research, 2023.

- [138] Zhi Liu and Nikhil Garg. Equity in resident crowdsourcing: Measuring under-reporting without ground truth data. *arXiv preprint arXiv:2204.08620*, 2022.
- [139] Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E. Bourne, Jildau Bouwman, Anthony J. Brookes, Tim Clark, Mercè Crosas, Ingrid Dillo, Olivier Dumon, Scott Edmunds, Chris T. Evelo, Richard Finkers, Alejandra Gonzalez-Beltran, Alasdair J.G. Gray, Paul Groth, Carole Goble, Jeffrey S. Grethe, Jaap Heringa, Peter A.C 't Hoen, Rob Hooft, Tobias Kuhn, Ruben Kok, Joost Kok, Scott J. Lusher, Maryann E. Martone, Albert Mons, Abel L. Packer, Bengt Persson, Philippe Rocca-Serra, Marco Roos, Rene van Schaik, Susanna-Assunta Sansone, Erik Schultes, Thierry Sengstag, Ted Slater, George Strawn, Morris A. Swertz, Mark Thompson, Johan van der Lei, Erik van Mulligen, Jan Velterop, Andra Waagmeester, Peter Wittenburg, Katherine Wolstencroft, Jun Zhao, and Barend Mons. The FAIR guiding principles for scientific data management and stewardship. *Scientific Data*, 3(1):1–9, 2016.
- [140] Stephanie Russo Carroll, Ibrahim Garba, Oscar L. Figueroa-Rodríguez, Jarita Holbrook, Raymond Lovett, Simeon Materechera, Mark Parsons, Kay Raseroka, Desi Rodriguez-Lonebear, Robyn Rowe, Rodrigo Sara, Jennifer D. Walker, Jane Anderson, and Maui Hudson. The CARE principles for indigenous data governance. *Data Science Journal*, Nov 2020.
- [141] IUCN. *IUCN guidelines for the prevention of biodiversity loss caused by alien invasive species*. Species Survival Commission, 2000.
- [142] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, Tony Lee, Etienne David, Ian Stavness, Wei Guo, Berton Earnshaw, Imran Haque, Sara M Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. WILDS: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, pages 5637–5664. PMLR, 2021.
- [143] Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 456–473, 2018.
- [144] Petra M Kuhnert, Tara G Martin, and Shane P Griffiths. A guide to eliciting and using expert knowledge in Bayesian ecological models. *Ecology Letters*, 13(7):900–914, 2010.
- [145] Roy Schwartz, Jesse Dodge, Noah A Smith, and Oren Etzioni. Green AI. *Communications of the ACM*, 63(12):54–63, 2020.
- [146] William D Moreto. Poaching in Uganda: Perspectives of law enforcement rangers. *Deviant Behavior*, 36(11):853–873, 2015.
- [147] Timothy Kuiper, Blessing Kavhu, Nobesuthu A Ngwenya, Roseline Mandisodza-Chikerema, and EJ Milner-Gulland. Rangers and modellers collaborate to build and evaluate spatial models of African elephant poaching. *Biological Conservation*, 243:108486, 2020.
- [148] Elizabeth Bondi, Raphael Koster, Hannah Sheahan, Martin Chadwick, Yoram Bachrach, Taylan Cemgil, Ulrich Paquet, and Krishnamurthy Dvijotham. Role of human-AI interaction in selective prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 5286–5294, 2022.
- [149] Mohammad Sadegh Norouzzadeh, Dan Morris, Sara Beery, Neel Joshi, Nebojsa Jojic, and Jeff Clune. A deep active learning system for species identification and counting in camera trap images. *Methods in ecology and evolution*, 12(1):150–161, 2021.
- [150] Zhongqi Miao, Ziwei Liu, Kaitlyn M Gaynor, Meredith S Palmer, Stella X Yu, and Wayne M Getz. Iterative human and automated identification of wildlife images. *Nature Machine Intelligence*, 3(10):885–895, 2021.
- [151] Peter Kulits, Jake Wall, Anka Bedetti, Michelle Henley, and Sara Beery. Elephantbook: A semi-automated human-in-the-loop system for elephant re-identification. In *ACM SIGCAS Conference on Computing and Sustainable Societies*, pages 88–98, 2021.

- [152] Rebecca E Morss, Julie L Demuth, and Jeffrey K Lazo. Communicating uncertainty in weather forecasts: A survey of the US public. *Weather and Forecasting*, 23(5):974–991, 2008.
- [153] Chris Huntingford, Elizabeth S Jeffers, Michael B Bonsall, Hannah M Christensen, Thomas Lees, and Hui Yang. Machine learning and artificial intelligence to aid climate change research and preparedness. *Environmental Research Letters*, 14(12):124007, 2019.
- [154] Amy J Dickman. Complexities of conflict: The importance of considering social factors for effectively resolving human–wildlife conflict. *Animal Conservation*, 13(5):458–466, 2010.
- [155] Carla Gomes, Thomas Dietterich, Christopher Barrett, Jon Conrad, Bistra Dilkina, Stefano Ermon, Fei Fang, Andrew Farnsworth, Alan Fern, Xiaoli Fern, Daniel Fink, Douglas Fisher, Alexander Flecker, Daniel Freund, Angela Fuller, John Gregoire, John Hopcroft, Steve Kelling, Zico Kolter, Warren Powell, Nicole Sintov, John Selker, Bart Selman, Daniel Sheldon, David Shmoys, Milind Tambe, Weng-Keen Wong, Christopher Wood, Xiaojian Wu, Yexiang Xue, Amulya Yadav, Abdul-Aziz Yakubu, and Mary Lou Zeeman. Computational sustainability: Computing for a better world and a sustainable future. *Communications of the ACM*, 62(9):56–65, 2019.
- [156] Karen Bakker. Smart oceans: Artificial intelligence and marine protected area governance. *Earth System Governance*, 13:100141, 2022.