



OPEN

# Improving biodiversity protection through artificial intelligence

Daniele Silvestro<sup>1,2,3,8</sup> , Stefano Gorla<sup>4</sup>, Thomas Sterner<sup>5</sup>  and Alexandre Antonelli<sup>3,6,7,8</sup> 

**Over a million species face extinction, highlighting the urgent need for conservation policies that maximize the protection of biodiversity to sustain its manifold contributions to people's lives. Here we present a novel framework for spatial conservation prioritization based on reinforcement learning that consistently outperforms available state-of-the-art software using simulated and empirical data. Our methodology, conservation area prioritization through artificial intelligence (CAPTAIN), quantifies the trade-off between the costs and benefits of area and biodiversity protection, allowing the exploration of multiple biodiversity metrics. Under a limited budget, our model protects significantly more species from extinction than areas selected randomly or naively (such as based on species richness). CAPTAIN achieves substantially better solutions with empirical data than alternative software, meeting conservation targets more reliably and generating more interpretable prioritization maps. Regular biodiversity monitoring, even with a degree of inaccuracy characteristic of citizen science surveys, further improves biodiversity outcomes. Artificial intelligence holds great promise for improving the conservation and sustainable use of biological and ecosystem values in a rapidly changing and resource-limited world.**

Biodiversity is the variety of all life on Earth, from genes to populations, species, functions and ecosystems. Alongside its own intrinsic value and ecological roles, biodiversity provides us with clean water, pollination services, building materials, clothing, food and medicine, among many other physical and cultural contributions that species make to ecosystem services and people's lives<sup>1,2</sup>. The contradiction is that our endeavours to maximize short-term benefits have become unsustainable, depleting biodiversity and threatening the life-sustaining foundations of humanity in the long term<sup>3</sup> (Supplementary Box 1). This can help explain why, despite the risks, we are living in an age of mass extinction<sup>4,5</sup>. The imperative to feed and house rapidly growing human populations, with an estimated 2.4 billion more people by 2050, together with increasing disruptions from climate change, will put tremendous pressure on the world's last remaining native ecosystems and the species they contain. Because not a single one of the 20 Aichi Biodiversity Targets agreed by 196 nations for the period 2011–2020 has been fully met<sup>6</sup>, there is now an urgent need to design more realistic and effective policies for a sustainable future<sup>7</sup> that help deliver the conservation targets under the post-2020 Global Biodiversity Framework, the focus of the 15th Conference of the Parties in 2022.

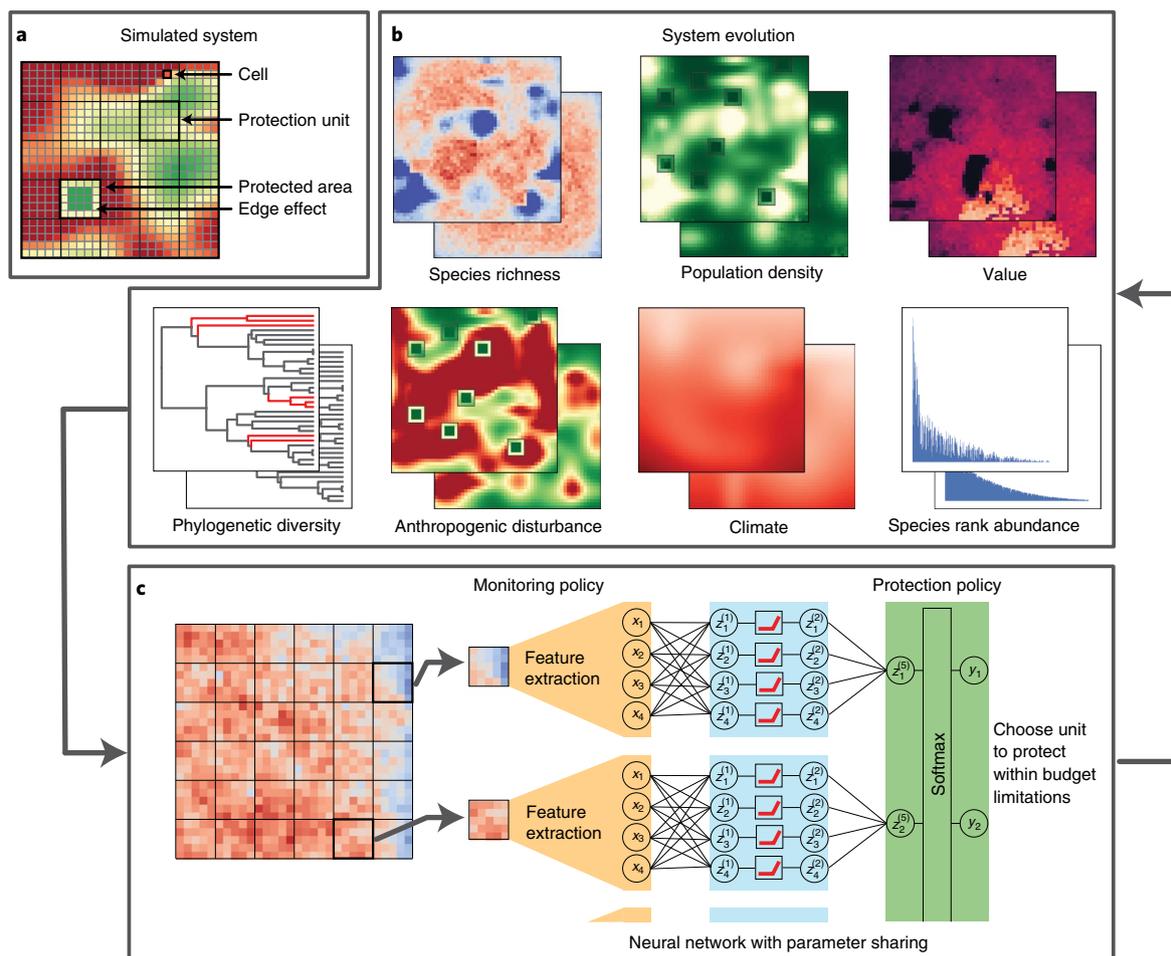
There have been several theoretical and practical frameworks underlying biological conservation<sup>8</sup> since the 1960s. The field was initially focused on the conservation of nature for itself, without human interference, but gradually incorporated the bidirectional links to people, recognizing our ubiquitous influence on nature and the multifaceted contributions we derive from it, including the sustainable use of species<sup>1,8,9</sup>. Throughout this progress, a critical step has been the identification of priority areas for targeted protection, restoration planning, impact avoidance and loss minimization, triggering the development of the fields of spatial conservation prioritization and systematic conservation planning<sup>10–16</sup>. While humans and wild species are increasingly sharing the same space<sup>17</sup>, the preservation

of largely intact nature remains critical for safeguarding many species and ecosystems, such as tropical rainforests.

Several tools and algorithms have been designed to facilitate systematic conservation planning<sup>18</sup>. They often allow the exploration and optimization of trade-offs between variables, something not readily available in Geographic Information Systems<sup>19</sup>, which can lead to substantial economic, social and environmental gains<sup>20</sup>. While the initial focus has been on maximizing the protection of species while minimizing costs, additional parameters can sometimes be modelled, such as species rarity and threat, total protected area and evolutionary diversity<sup>18,21,22</sup>. The most widely used method so far, Marxan<sup>23</sup>, seeks to identify a set of protected areas that collectively allow particular conservation targets to be met under minimal costs, using a simulated annealing optimization algorithm. Despite its usefulness and popularity, Marxan and similar methods<sup>18</sup> were designed to optimize a one-time policy, do not directly incorporate changes through time, and assume a single initial gathering of biodiversity and cost data (although temporal aspects can be explored by manually updating and re-running the models under various targets<sup>24</sup>). In addition, the optimized conservation planning does not explicitly incorporate climate change, variation in anthropogenic pressure (although varying threat probabilities are dealt with in recent software extensions of Marxan<sup>25,26</sup>), or species-specific sensitivities to such changes.

In this study we have tackled the challenge of optimizing biodiversity protection in a complex and rapidly evolving world by harnessing the power of artificial intelligence (AI). We have developed an entirely novel tool for systematic conservation planning (Fig. 1) that optimizes a conservation policy based on static or dynamic biodiversity monitoring towards user-defined targets (such as minimizing species loss) and within the constraints of a limited financial budget, and used it to explore, through simulations and empirical analyses, multiple previously identified trade-offs in

<sup>1</sup>Department of Biology, University of Fribourg, Fribourg, Switzerland. <sup>2</sup>Swiss Institute of Bioinformatics, Fribourg, Switzerland. <sup>3</sup>Gothenburg Global Biodiversity Centre, Department of Biological and Environmental Sciences, University of Gothenburg, Gothenburg, Sweden. <sup>4</sup>Thymia, London, UK. <sup>5</sup>Department of Economics, University of Gothenburg, Gothenburg, Sweden. <sup>6</sup>Department of Plant Sciences, University of Oxford, Oxford, UK. <sup>7</sup>Royal Botanic Gardens, Kew, UK. <sup>8</sup>These authors contributed equally: Daniele Silvestro, Alexandre Antonelli. ✉e-mail: [daniele.silvestro@unifr.ch](mailto:daniele.silvestro@unifr.ch); [a.antonelli@kew.org](mailto:a.antonelli@kew.org)

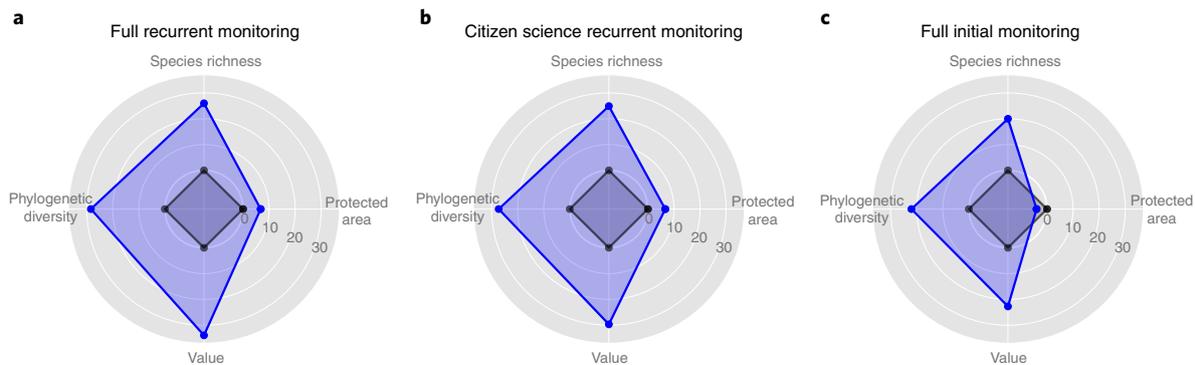


**Fig. 1 | The CAPTAIN reinforcement learning framework.** **a**, A simulated system, which could be equivalent to a country an island or a large coral reef, consists of a number of cells, each with a number of individuals of various species. Once a protection unit is identified and protected, its human-driven disturbance (for example, forest logging or sea trawling) will immediately reduce to an arbitrarily low level, except for the well-known edge effect<sup>47</sup> characterized by intermediate levels of disturbance. All simulation settings are provided with initial default values but are fully customizable (Supplementary Tables 1 and 2). Simulated systems evolve through time and are used to optimize a conservation policy using RL. After training the model, the optimized policy can be used to evaluate the model performance based on simulated or empirical data. Using empirical data, the simulated system is replaced with available biodiversity and disturbance data. **b, c**, Analysis flowchart integrating system evolution (**b**) with simulations and AI modules (**c**) to maximize selected outcomes (for example, species richness). The system evolves between two points in time, with several time-dependent variables considered (seven plotted here): species richness, population density, economic value, phylogenetic diversity, anthropogenic disturbance, climate and species rank abundance (see [www.captain-project.net](http://www.captain-project.net) for animations depicting these and additional variables). Biodiversity features (species presence per protection unit at a minimum, plus their abundance under full monitoring schemes as defined here; see Methods and Supplementary Box 2 for advances in data gathering approaches) are extracted from the system at regular steps, and are then fed into a neural network that learns from the system's evolution to identify conservation policies that maximize a reward, such as protection of the maximum species diversity within a fixed budget. The vectors of parameters  $x$ ,  $z$  and  $y$  represent the nodes of the input, hidden layers and output of the neural network, respectively.

real-world conservation and to evaluate the impact of data gathering on specific outcomes<sup>27</sup>. We have also explored the impact of species-specific sensitivity to geographically varying local disturbances (for example, as a consequence of new roads, mining, trawling or other forms of unsustainable economic activity with negative impacts on natural ecosystems) and climate change (overall temperature increases as well as short-term variations to reflect extreme weather events). We name our framework CAPTAIN, denoting Conservation Area Prioritization Through Artificial Intelligence.

Within AI, we implemented a reinforcement learning (RL) framework based on a spatially explicit simulation of biodiversity and its evolution through time in response to anthropogenic pressure and climate change. The RL algorithm is designed to find an optimal balance between data generation (learning from the

current state of a system, also termed 'exploration') and action (called 'exploitation', the effect of which is quantified by the outcome, also termed 'reward'). Our platform enables us to assess the influence of model assumptions on the reward, mimicking the use of counterfactual analyses<sup>22</sup>. CAPTAIN can optimize a static policy, where all the budget is spent at once, or (more in line with its primary objective) a conservation policy that develops over time, thus being particularly suitable for designing policies and testing their short- and long-term effects. Actions are decided based on the state of the system through a neural network, whose parameters are optimized within the RL framework to maximize the reward. Once a model is trained through RL, it can be used to identify conservation priorities in space and time using simulated or empirical data.



**Fig. 2 | Impact of monitoring strategies on biodiversity protection.** **a–c**, Outcome of policies designed to minimize species loss based on different monitoring strategies: full recurrent monitoring (of species presence and abundance at each time step; **a**), citizen science recurrent monitoring (limited to species presence/absence with some error at each time step; **b**) and full initial monitoring (species presence and abundance only at the initial time; **c**). The results show the percentage change in species loss, total protected area, accumulated species value and phylogenetic diversity between a random protection policy (black polygons) and models optimized by CAPTAIN (blue polygons). All results are averaged across 250 simulations, with more details shown in Supplementary Table 3. Each simulation was based on the same budget and resolution of the protection units ( $5 \times 5$  cells) but differed in their initial natural system (species distributions, abundances, tolerances and phylogenetic relationships) and in the dynamics of climate change and disturbance patterns.

Although AI solutions have been previously proposed and to some extent are already used in conservation science<sup>28,29</sup>, to our knowledge RL has only been advocated<sup>30</sup> and not yet implemented in practical conservation tools. In particular, CAPTAIN aims to tackle multidimensional problems of loss minimization considered by techniques such as stochastic dynamic programming but proven thus far intractable for large systems<sup>13</sup>. It thus fills an important space in conservation in a dynamic world<sup>31</sup>, characterized by heterogeneous and often unpredictable habitat loss<sup>14</sup>, which require iterative and regular conservation interventions.

We have used CAPTAIN to address the following questions: (1) What role does the data-gathering strategy have in effective conservation? (2) What trade-offs arise depending on the optimized variable, such as species richness, economic value or total area protected? (3) What can the simulation framework reveal in terms of winners and losers, that is, which traits characterize the species and areas protected over time? (4) How does our framework perform compared with the state-of-the-art model for conservation planning, Marxan<sup>32</sup>? Finally, we demonstrate here the usefulness of our framework and direct applicability of models trained through RL on an empirical dataset of endemic trees of Madagascar.

## Results

**Impact of data gathering strategy.** Using CAPTAIN we found that full recurrent monitoring (where the system is monitored at each time step, including species presence and abundance) results in the smallest species loss: it succeeds in protecting on average 26% more species than a random protection policy (Fig. 2a and Supplementary Table 3). A very similar outcome (24.9% improvement) is generated by the citizen science recurrent monitoring strategy (where only presence/absence of species are recorded in each cell), with a degree of error characteristic of citizen science efforts (Fig. 2b and Methods). These two monitoring strategies outperform a full initial monitoring with no error, which only saves from extinction an average of 20% more species than a random policy (Fig. 2c and Supplementary Table 3).

To thoroughly explore the parameter space of the simulations, each system was initialized with different species composition and distributions and different anthropogenic pressure and climate change patterns (Supplementary Figs. 1–4). Because of this stochasticity, the reliability of the protection policies in relation to species

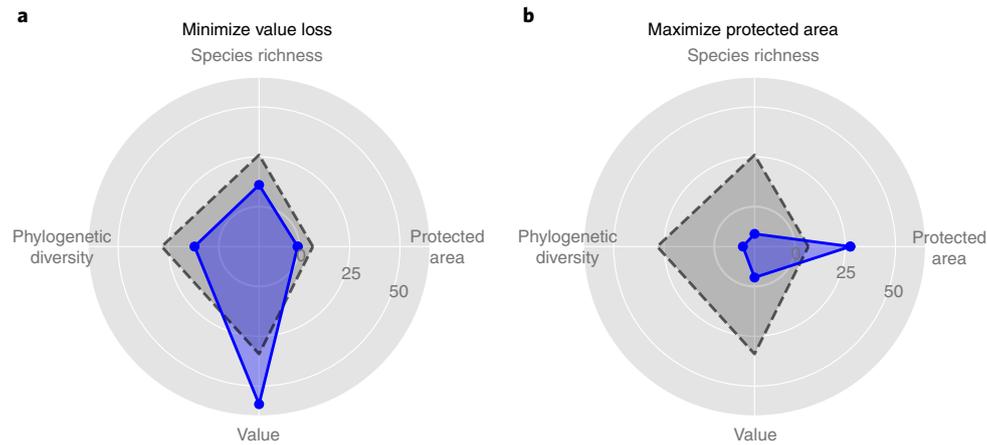
loss varies across simulations. The policies based on full recurrent monitoring and citizen science recurrent monitoring are the most reliable, outperforming the baseline random policy in 97.2% of the simulations. Both these policies are more reliable than the full initial monitoring, which in addition to protecting fewer species on average (Fig. 2) also results in a slightly lower reliability of the outcome, outperforming the random policy in 91.2% of the simulations.

**Optimization trade-offs.** The policy objective, which determines the optimality criterion in our RL framework, strongly influences the outcome of the simulations. A policy minimizing species loss based on their commercial value (such as timber price) tends to sacrifice more species to prioritize the protection of fewer, highly valuable ones. This policy, while efficiently reducing the loss of cumulative value, decreases species losses by only 10.9% compared with the random baseline (Supplementary Table 3). Thus, a policy targeting exclusively the preservation of species with high economic value may have a strongly negative impact on the total protected species richness, phylogenetic diversity and even amount of protected area compared with a policy minimizing species loss (Fig. 3a).

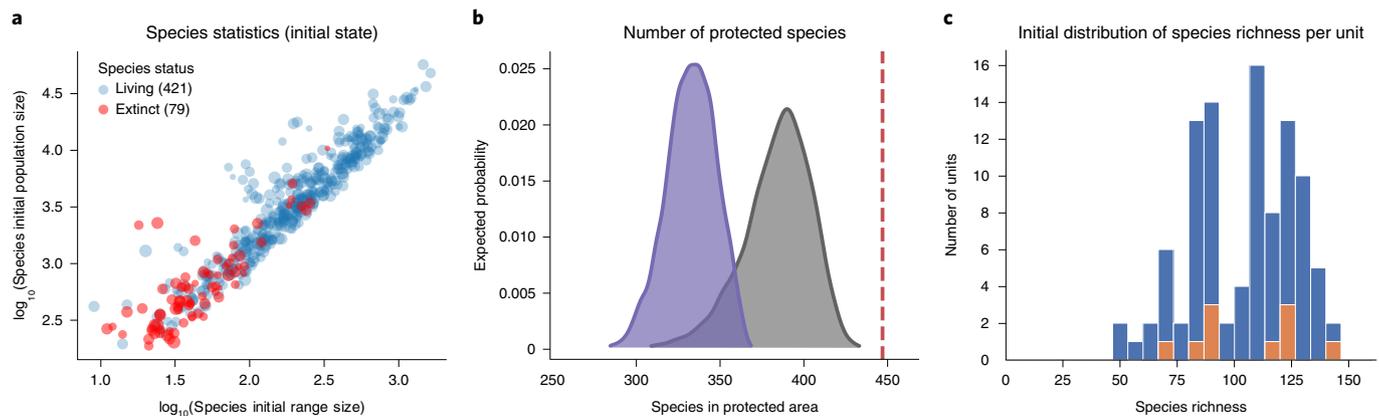
A policy that maximizes protected area results in a 27.6% increase in the number of protected cells by selecting those cheapest to buy; however, it leads to substantial losses in species numbers, value and phylogenetic diversity, which are considerably worse than the random baseline, with 13.6% more species losses on average (Supplementary Table 3). The decreased performance in terms of preventing extinctions is even more pronounced when compared with a policy minimizing species loss (Fig. 3b).

As expected, the reliability for optimizations based on economic value and total protected area is high for the respective policy objectives, but they result in highly inconsistent outcomes in terms of preventing species extinctions, with biodiversity losses not significantly different from those of the random baseline policy (Supplementary Table 3).

**Winners and losers.** Focusing on the policy developed under full recurrent monitoring and optimized on reducing species loss, we explored the properties of species that survived in comparison with those that went extinct, despite optimal area protection. Species that went extinct are characterized by relatively small initial ranges,



**Fig. 3 | Trade-offs in conservation outcomes in relation to policy objectives. a,b**, The impact of different policy objectives based on full recurrent monitoring: the policies were designed to minimize value loss (**a**) and maximize the amount of protected area (**b**). The plots show the outcome averaged across 250 simulations (blue polygons). The radial axis shows the percentage change compared with the baseline random policy, and the dashed grey polygons show the outcome of a policy with full recurrent monitoring optimized to minimize biodiversity loss (Fig. 2a).

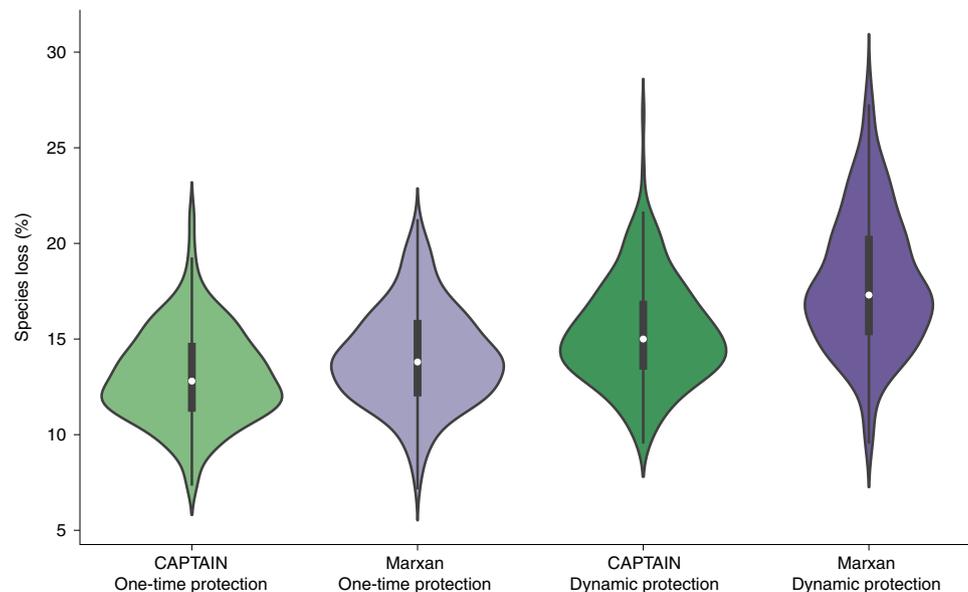


**Fig. 4 | Summary statistics for one simulation optimized to reduce species loss and informed through full recurrent monitoring. a**, Living (or surviving) and locally extinct species after a simulation of 30 time steps with increasing disturbance and climate change. The x and y axes show the initial range and population sizes of the species ( $\log_{10}$  transformed), respectively. The size of the circles is proportional to the resilience of each species to anthropogenic disturbance, with smaller circles representing more sensitive species. **b**, Cumulative number of species encompassed in the ten protected units ( $5 \times 5$  cells) selected on the basis of a policy optimized to minimize species loss. The grey density plot shows the expected distribution from 10,000 random draws, and the purple shaded area shows the expected distribution when protected units are selected 'naively' (here, randomly chosen from among the top 20 most diverse units). The dashed red line indicates the number of species included in the units selected by the optimized CAPTAIN policy, which is higher than in all the random draws. The optimized policy learned to maximize the total number of species included in protected units, thus accounting for their complementarity. Note that fewer species survived (421) in this simulation compared with how many were included in protected areas (447). This discrepancy is due to the effect of climate change, in which area protection does not play a role<sup>47</sup> (see also [www.captain-project.net](http://www.captain-project.net)). **c**, Species richness across the 100 protection units included in the area (blue), ten of which were selected to be protected (orange). The plot shows that the protection policy does not exclusively target units with the highest diversity.

small populations and intermediate or low resilience to disturbance (Fig. 4a). In contrast, species that survived have either low resilience but widespread ranges and high population sizes, or high resilience with small ranges and population sizes.

We further assessed what characterizes the grid cells that are selected for protection by the optimized policy. The cumulative number of species included in these cells is significantly higher than the cumulative species richness across a random set of cells of equal area (Fig. 4b). Thus, the model learns to protect a diversity of species assemblages to minimize species loss. Interestingly, the cells selected for protection did not include only areas with the highest species richness (Fig. 4c).

**Benchmarking through simulations.** We evaluated our simulation framework by comparing its performance in optimizing policies with that of the current state-of-the-art tool for conservation prioritization, Marxan<sup>23</sup>. The methods differ conceptually in that while CAPTAIN is explicitly designed to minimize loss (for example, local species extinction) within the constraints of a limited budget, Marxan's default algorithms minimize the cost of reaching a conservation target (for example, protecting at least 10% of all species ranges). Additionally, Marxan is typically used to optimize the placement of protected units in a single step, while CAPTAIN places the protection units across different time steps.



**Fig. 5 | Benchmarking of CAPTAIN through simulations.** The violin plots show the distribution of species loss outcomes across 250 simulations run under different monitoring policies based on the CAPTAIN framework developed here and on Marxan<sup>23</sup>. Species loss is expressed as a percentage of the total initial number of species in the simulated systems. See text for details of the analyses. The white circle in each box represents the median, the top and bottom of each box represent the 75th and 25th percentiles, respectively, and the whiskers are extended to 1.5× the interquartile range.

To compare the two models, we set up all Marxan analyses with an explicit budget constraint, following other tailored implementations<sup>15,32</sup> (see Methods). In a first comparison, we tested a protection policy in which all protection units (within a predefined budget) are established in one step. To this end, we trained an additional model in CAPTAIN based on a full initial monitoring and on a policy in which all budget for protection is spent in one step. The analysis of 250 simulations showed that CAPTAIN outperforms Marxan in 64% of the cases with an average improvement in terms of prevented species loss of 9.2% (Fig. 5).

In a second comparison, we used CAPTAIN with full recurrent monitoring and allowed the establishment of a single protection unit per time step for both programs (see Methods). Under this condition, our model outperforms Marxan in 77.2% of the simulations with an average reduction of species loss of 18.5% (Fig. 5).

**Empirical applications.** To demonstrate the applicability of our framework and its scalability to large, real-world tasks, we analysed a Madagascar biodiversity dataset recently used in a systematic conservation planning experiment<sup>33</sup> under Marxan<sup>23</sup>. The dataset included 22,394 protection units (5×5 km) and presence/absence data for 1,517 endemic tree species. The cost of area protection was set proportional to anthropogenic disturbance across cells, as in the original publication<sup>33</sup> (see Methods for more details and Supplementary Fig. 5).

We analysed the data assuming full initial monitoring in a static setting in which all protection units were placed in one step. We limited the budget to an amount that allows the protection of at most 10% of the units (or fewer if expensive units are chosen) and set the target of preserving at least 10% of the species' potential range within protected units. We repeated the Marxan analyses with a boundary length multiplier (BLM; which penalizes the placement of many isolated protection units in favour of larger contiguous areas; BLM=0.1, as in ref.<sup>33</sup>) and without it (BLM=0 for comparability with CAPTAIN, which does not include this feature).

The solutions found in CAPTAIN consistently outperform those obtained with Marxan. Within the budget constraints, CAPTAIN

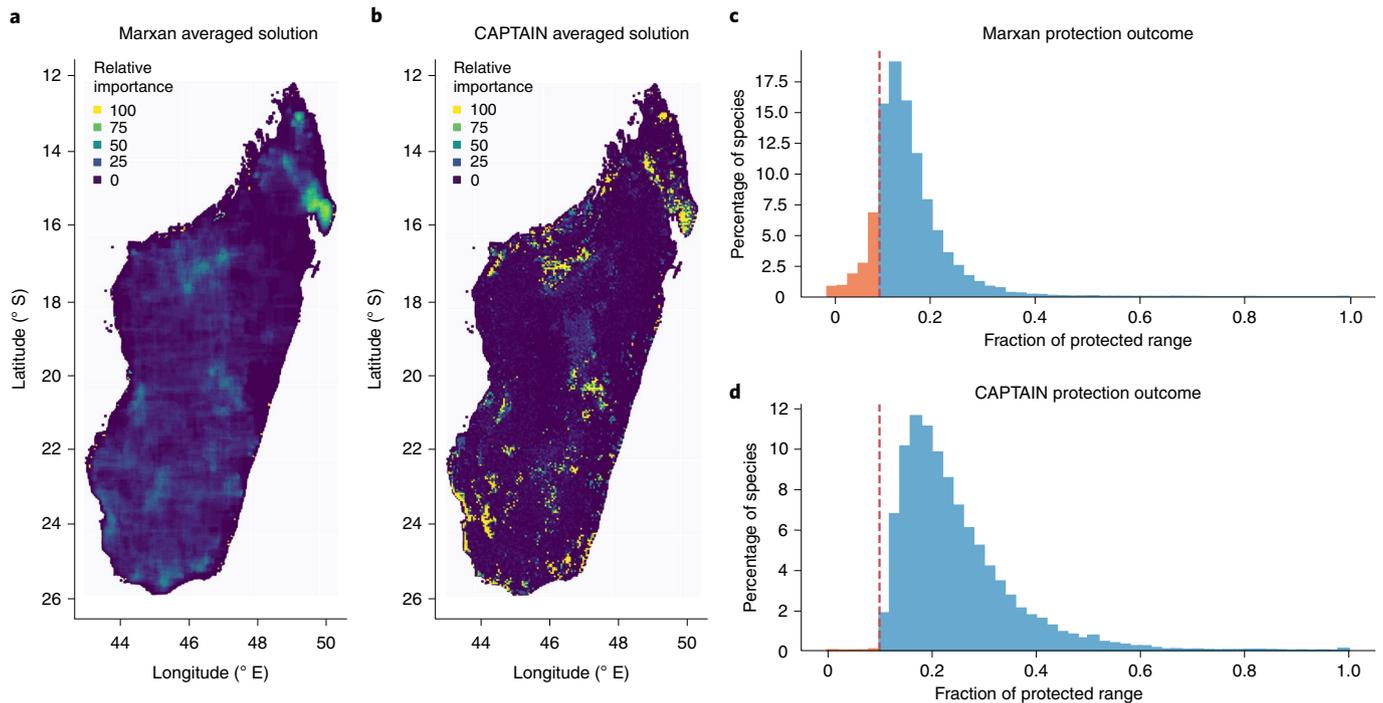
solutions meet the target of protecting 10% of the range for all species in 68% of the replicates, whereas only up to 2% of the Marxan results reach that target (Supplementary Table 4). Additionally, with CAPTAIN, a median of 22% of each species range is found within protected units, well above the set target of 10% and the 14% median protected range achieved with Marxan (Fig. 6c,d and Supplementary Fig. 6c,d). Importantly, CAPTAIN is able to identify priority areas for conservation at higher and therefore more interpretable spatial resolution (Fig. 6b and Supplementary Fig. 6b).

## Discussion

We have presented here a new framework to optimize dynamic conservation policies using RL and evaluate their biodiversity outcome through simulations.

**Data gathering and monitoring.** Our finding that even simple data (presence/absence of species) are sufficient to inform effective policies (Fig. 2 and Supplementary Table 3) is noteworthy because the information required is already available for many regions and taxonomic groups, and could be further complemented by modern technologies such as remote sensing and environmental DNA, and for accessible locations also citizen science<sup>34</sup> in cost-efficient ways (Supplementary Box 2).

The reason why single biodiversity assessments and area protection are often suboptimal is that they ignore the temporal dynamics caused by disturbances, population and range changes of species, all of which are likely to change through time in real-world situations. Although some systems may remain largely static over decades (for example, tree species in old-growth forests), others may change drastically (for example, alpine meadows or shallow-sea communities, where species shift their ranges rapidly in response to climatic and anthropogenic pressures); all such parameters can be tuned in our simulated system and accounted for in training the models through RL. Because current methodologies for systematic conservation planning are static, relying on a similar initial data gathering as modelled here, their recommendations for area protection may be less reliable.



**Fig. 6 | Empirical validation of CAPTAIN.** **a,b**, Maps showing the ranking of priority areas for protection across Madagascar based on the distribution of endemic trees and on a limited budget allowing for up to 10% of protected area overall: protection units identified through Marxan optimizations (with  $BLM = 0.1$ , as in Carrasco et al.<sup>33</sup>; **a**) and equivalent CAPTAIN results based on a full initial monitoring policy (**b**). The colour scale represents the relative importance of each protected unit, measured as the frequency at which it was included in the optimized solution. **c,d**, Histograms of the fraction of species ranges included in the protected units calculated for each species with Marxan (**c**) and CAPTAIN (**d**). The red dashed lines indicate the 10% threshold set as the target of the policy. Orange bars show species whose protection did not meet the target.

**Optimization trade-offs.** Our results indicate clear trade-offs, meaning that optimizing one value can be at the cost of another (Fig. 3 and Supplementary Table 3). In particular, our finding that maximizing total protected area can lead to substantial species loss is of urgent relevance, given that total protected area has been at the core of previous international targets for biodiversity (such as the Aichi Biodiversity Targets, <https://www.cbd.int/sp/targets>) and remains a key focus under the new post-2020 Global Biodiversity Framework under the Convention on Biological Diversity. Focusing on quantity (area protected) rather than quality (actual biodiversity protected) could inadvertently support political pressure for ‘residual’ reservation<sup>35,36</sup>, that is, the selection of new protected areas on land and at sea that are unsuitable for extractive activities, which may reduce costs and risk of conflicts, but are likely suboptimal for biodiversity conservation. Our trade-off analyses imply that economic value and total protected area should not be used as surrogates for biodiversity protection.

**Learning from the models.** Examination of our results reveals that, perhaps contrary to intuition, protected areas should not be primarily chosen based on high species richness (a ‘naive’ conservation target; Fig. 4). Instead, the simulations indicate that protected cells should span a range of areas with intermediate to high species richness, reflecting known differences between ecosystems or across environmental gradients. Such selection is more likely to increase protection complementarity for multiple species, a key factor incorporated in our software and some others<sup>10,23,37</sup>.

**Applications and prospects.** Our successful benchmarking against random, naive and Marxan-optimized solutions indicates that CAPTAIN has potential as a useful tool for informing

on-the-ground decisions by landowners and policymakers. Models trained through simulations in CAPTAIN can be readily applied to available empirical datasets.

In our experiments, CAPTAIN solutions outperform Marxan, even when based on the same input data, as in the example of Malagasy trees. Our simulations show that further improvement is expected when additional data describing the state of the system are used, and when the protection policy is developed over time rather than in a single step. These findings indicate that our AI parametric approach can (1) more efficiently use the available information on species distribution and (2) more easily integrate multidimensional and time-varying biodiversity data. As the number of standardized high-resolution biological datasets is increasing (for example, see ref. <sup>38</sup>), as a result of new and cost-effective monitoring technologies (Supplementary Box 2), our approach offers a future-proof tool for research, conservation and the sustainable use of natural resources. Our model can be easily expanded and adapted to almost any empirical dataset and to incorporate additional variables, such as functional diversity and more sophisticated measures of economic value. Similarly, the flexibility of our AI approach allows for the design of custom policy objectives, such as optimizing carbon sequestration and storage.

In contrast to many short-lived decisions by governments, the selection of which areas in a country’s territory should be protected will have long-term repercussions. Protecting the right areas, or developing sustainable models of using biodiversity without putting species at risk, will help safeguard natural assets and their contributions for the future. Choosing suboptimal areas for protection, by contrast, could not only waste public funding, but also lead to the loss of species, phylogenetic diversity, socioeconomic value and ecological functions. AI techniques should not replace human

judgement, and ultimately investment decisions will be based on more than just the parameters implemented in our models, including careful consideration of people's manifold interactions with nature<sup>1,8</sup>. It is also crucial to recognize the importance of ensuring the right conditions required for effective conservation of protected areas in the long term<sup>39,40</sup>. However, it is now time to acknowledge that the sheer complexity of sociobiological systems, multiplied by the increasing disturbances in a changing world, cannot be fully grasped by the human mind. As we progress in what many are calling the most decisive decade for nature<sup>9,41</sup>, we must take advantage of powerful tools that help us steward the planet's remaining ecosystems in sustainable ways—for the benefit of people and all life on Earth.

## Methods

**A biodiversity simulation framework.** We have developed a simulation framework modelling biodiversity loss to optimize and validate conservation policies (in this context, decisions about data gathering and area protection across a landscape) using an RL algorithm. We implemented a spatially explicit individual-based simulation to assess future biodiversity changes based on natural processes of mortality, replacement and dispersal. Our framework also incorporates anthropogenic processes such as habitat modifications, selective removal of a species, rapid climate change and existing conservation efforts. The simulation can include thousands of species and millions of individuals and track population sizes and species distributions and how they are affected by anthropogenic activity and climate change (for a detailed description of the model and its parameters see Supplementary Methods and Supplementary Table 1).

In our model, anthropogenic disturbance has the effect of altering the natural mortality rates on a species-specific level, which depends on the sensitivity of the species. It also affects the total number of individuals (the carrying capacity) of any species that can inhabit a spatial unit. Because sensitivity to disturbance differs among species, the relative abundance of species in each cell changes after adding disturbance and upon reaching the new equilibrium. The effect of climate change is modelled as locally affecting the mortality of individuals based on species-specific climatic tolerances. As a result, more tolerant or warmer-adapted species will tend to replace sensitive species in a warming environment, thus inducing range shifts, contraction or expansion across species depending on their climatic tolerance and dispersal ability.

We use time-forward simulations of biodiversity in time and space, with increasing anthropogenic disturbance through time, to optimize conservation policies and assess their performance. Along with a representation of the natural and anthropogenic evolution of the system, our framework includes an agent (that is, the policy maker) taking two types of actions: (1) monitoring, which provides information about the current state of biodiversity of the system, and (2) protecting, which uses that information to select areas for protection from anthropogenic disturbance. The monitoring policy defines the level of detail and temporal resolution of biodiversity surveys. At a minimal level, these include species lists for each cell, whereas more detailed surveys provide counts of population size for each species. The protection policy is informed by the results of monitoring and selects protected areas in which further anthropogenic disturbance is maintained at an arbitrarily low value (Fig. 1). Because the total number of areas that can be protected is limited by a finite budget, we use an RL algorithm<sup>42</sup> to optimize how to perform the protecting actions based on the information provided by monitoring, such that it minimizes species loss or other criteria depending on the policy.

We provide a full description of the simulation system in the Supplementary Methods. In the sections below we present the optimization algorithm, describe the experiments carried out to validate our framework and demonstrate its use with an empirical dataset.

**Conservation planning within a reinforcement learning framework.** In our model we use RL to optimize a conservation policy under a predefined policy objective (for example, to minimize the loss of biodiversity or maximize the extent of protected area). The CAPTAIN framework includes a space of actions, namely monitoring and protecting, that are optimized to maximize a reward  $R$ . The reward defines the optimality criterion of the simulation and can be quantified as the cumulative value of species that do not go extinct throughout the timeframe evaluated in the simulation. If the value is set equal across all species, the RL algorithm will minimize overall species extinctions. However, different definitions of value can be used to minimize loss based on evolutionary distinctiveness of species (for example, minimizing phylogenetic diversity loss), or their ecosystem or economic value. Alternatively, the reward can be set equal to the amount of protected area, in which case the RL algorithm maximizes the number of cells protected from disturbance, regardless of which species occur there. The amount of area that can be protected through the protecting action is determined by a budget  $B_t$  and by the cost of protection  $C_t^c$ , which can vary across cells  $c$  and through time  $t$ .

The granularity of monitoring and protecting actions is based on spatial units that may include one or more cells and which we define as the protection units. In our system, protection units are adjacent, non-overlapping areas of equal size (Fig. 1) that can be protected at a cost that cumulates the costs of all cells included in the unit.

The monitoring action collects information within each protection unit about the state of the system  $S_t$ , which includes species abundances and geographic distribution:

$$S_t = \{H_t, D_t, F_t, T_t, C_t, P_t, B_t\} \quad (1)$$

where  $H_t$  is the matrix with the number of individuals across species and cells,  $D_t$  and  $F_t$  are matrices describing anthropogenic disturbance on the system,  $T_t$  is a matrix quantifying climate,  $C_t$  is the cost matrix,  $P_t$  is the current protection matrix and  $B_t$  is the available budget (for more details see Supplementary Methods and Supplementary Table 1). We define as feature extraction the result of a function  $X(S_t)$ , which returns for each protection unit a set of features summarizing the state of the system in the unit. The number and selection of features (Supplementary Methods and Supplementary Table 2) depends on the monitoring policy  $\pi_x$ , which is decided a priori in the simulation. A predefined monitoring policy also determines the temporal frequency of this action throughout the simulation, for example, only at the first time step or repeated at each time step. The features extracted for each unit represent the input upon which a protecting action can take place, if the budget allows for it, following a protection policy  $\pi_y$ . These features (listed in Supplementary Table 2) include the number of species that are not already protected in other units, the number of rare species and the cost of the unit relative to the remaining budget. Different subsets of these features are used depending on the monitoring policy and on the optimality criterion of the protection policy  $\pi_y$ .

We do not assume species-specific sensitivities to disturbance (parameters  $d_s$ ,  $f_s$  in Supplementary Table 1 and Supplementary Methods) to be known features, because a precise estimation of these parameters in an empirical case would require targeted experiments, which we consider unfeasible across a large number of species. Instead, species-specific sensitivities can be learned from the system through the observation of changes in the relative abundances of species ( $x_s$  in Supplementary Table 2). The features tested across different policies are specified in the subsection Experiments below and in the Supplementary Methods.

The protecting action selects a protection unit and resets the disturbance in the included cells to an arbitrarily low level. A protected unit is also immune from future anthropogenic disturbance increases, but protection does not prevent climate change in the unit. The model can include a buffer area along the perimeter of a protected unit, in which the level of protection is lower than in the centre, to mimic the generally negative edge effects in protected areas (for example, higher vulnerability to extreme weather). Although protecting a disturbed area theoretically allows it to return to its initial biodiversity levels, population growth and species composition of the protected area will still be controlled by the death–replacement–dispersal processes described above, as well as by the state of neighbouring areas. Thus, protecting an area that has already undergone biodiversity loss may not result in the restoration of its original biodiversity levels.

The protecting action has a cost determined by the cumulative cost of all cells in the selected protection unit. The cost of protection can be set equal across all cells and constant through time. Alternatively, it can be defined as a function of the current level of anthropogenic disturbance in the cell. The cost of each protecting action is taken from a predetermined finite budget and a unit can be protected only if the remaining budget allows it.

**Policy definition and optimization algorithm.** We frame the optimization problem as a stochastic control problem where the state of the system  $S_t$  evolves through time as described in the section above (see also Supplementary Methods), but it is also influenced by a set of discrete actions determined by the protection policy  $\pi_y$ . The protection policy is a probabilistic policy: for a given set of policy parameters and an input state, the policy outputs an array of probabilities associated with all possible protecting actions. While optimizing the model, we extract actions according to the probabilities produced by the policy to make sure that we explore the space of actions. When we run experiments with a fixed policy instead, we choose the action with highest probability. The input state is transformed by the feature extraction function  $X(S_t)$  defined by the monitoring policy, and the features are mapped to a probability through a neural network with the architecture described below.

In our simulations, we fix monitoring policy  $\pi_x$ , thus predefining the frequency of monitoring (for example, at each time step or only at the first time step) and the amount of information produced by  $X(S_t)$ , and we optimize  $\pi_y$ , which determines how to best use the available budget to maximize the reward. Each action  $A$  has a cost, defined by the function  $\text{Cost}(A, S_t)$ , which here we set to zero for the monitoring action ( $X$ ) across all monitoring policies. The cost of the protecting action ( $Y$ ) is instead set to the cumulative cost of all cells in the selected protection unit. In the simulations presented here, unless otherwise specified, the protection policy can only add one protected unit at each time step, if the budget allows, that is if  $\text{Cost}(Y, S_t) < B_t$ .

The protection policy is parametrized as a feed-forward neural network with a hidden layer using a rectified linear unit (ReLU) activation function (Eq. (3)) and

an output layer using a softmax function (Eq. (5)). The input of the neural network is a matrix  $x$  of  $J$  features extracted through the most recent monitoring across  $U$  protection units. The output, of size  $U$ , is a vector of probabilities, which provides the basis to select a unit for protection. Given a number of nodes  $L$ , the hidden layer  $h^{(1)}$  is a matrix  $U \times L$ :

$$h_{ul}^{(1)} = g \left( \sum_{j=1}^J x_{uj} W_{jl}^{(1)} \right) \quad (2)$$

where  $u \in \{1, \dots, U\}$  identifies the protection unit,  $l \in \{1, \dots, L\}$  indicates the hidden nodes and  $j \in \{1, \dots, J\}$  the features and where

$$g(x) = \max(0, x) \quad (3)$$

is the ReLU activation function. We indicate with  $W^{(1)}$  the matrix of  $J \times L$  coefficients (shared among all protection units) that we are optimizing. Additional hidden layers can be added to the model between the input and the output layer. The output layer takes  $h^{(1)}$  as input and gives an output vector of  $U$  variables:

$$h_u^{(2)} = \sigma \left( \sum_{l=1}^L h_{ul}^{(1)} W_l^{(2)} \right) \quad (4)$$

where  $\sigma$  is a softmax function:

$$\sigma(x_i) = \frac{\exp(x_i)}{\sum_u \exp(x_u)} \quad (5)$$

We interpret the output vector of  $U$  variables as the probability of protecting the unit  $u$ .

This architecture implements parameter sharing across all protection units when connecting the input nodes to the hidden layer; this reduces the dimensionality of the problem at the cost of losing some spatial information, which we encode in the feature extraction function. The natural next step would be to use a convolutional layer to discover relevant shape and space features instead of using a feature extraction function. To define a baseline for comparisons in the experiments described below, we also define a random protection policy  $\tilde{z}$ , which sets a uniform probability to protect units that have not yet been protected. This policy does not include any trainable parameter and relies on feature  $x_u$  (an indicator variable for protected units; Supplementary Table 2) to randomly select the proposed unit for protection.

The optimization algorithm implemented in CAPTAIN optimizes the parameters of a neural network such that they maximize the expected reward resulting from the protecting actions. With this aim, we implemented a combination of standard algorithms using a genetic strategies algorithm<sup>43</sup> and incorporating aspects of classical policy gradient methods such as an advantage function<sup>44</sup>. Specifically, our algorithm is an implementation of the Parallelized Evolution Strategies<sup>45</sup>, in which two phases are repeated across several iterations (hereafter, epochs) until convergence. In the first phase, the policy parameters are randomly perturbed and then evaluated by running one full episode of the environment, that is, a full simulation with the system evolving for a predefined number of steps. In the second phase, the results from different runs are combined and the parameters updated following a stochastic gradient estimate<sup>45</sup>. We performed several runs in parallel on different workers (for example, processing units) and aggregated the results before updating the parameters. To improve the convergence we followed the standard approach used in policy optimization algorithms<sup>45</sup>, where the parameter update is linked to an advantage function  $A$  as opposed to the return alone (Eq. (6)). Our advantage function measures the improvement of the running reward (weighted average of rewards across different epochs) with respect to the last reward. Thus, our algorithm optimizes a policy without the need to compute gradients and allowing for easy parallelization. Each epoch in our algorithm works as:

```

for every worker  $p$  do
   $\epsilon_p \leftarrow \mathcal{N}(0, \sigma)$ , with diagonal covariance and dimension  $W + M$ 
  for  $t = 1, \dots, T$  do
     $R_t \leftarrow R_{t-1} + r_t(\theta + \epsilon_p)$ 
  end for
end for
 $R \leftarrow$  average of  $R_t$  across workers
 $R_c \leftarrow \alpha R + (1 - \alpha)R_{c-1}$ 
for every coefficient  $\theta$  in  $W + M$  do
   $\theta \leftarrow \theta + \lambda A(R_c, R_t, \epsilon)$ 
end for

```

where  $\mathcal{N}$  is a normal distribution and  $W + M$  is the number of parameters in the model (following the notation in Supplementary Table 1). We indicate with  $r_t$  the reward at time  $t$ , with  $R$  the cumulative reward over  $T$  time steps.  $R_c$  is the running average reward calculated as an exponential moving average where  $\alpha = 0.25$  represents the degree of weighting decrease and  $R_{c-1}$  is the running average

reward at the previous epoch.  $\lambda = 0.1$  is a learning rate and  $A$  is an advantage function defined as the average of final reward increments with respect to the running average reward  $R_c$  on every worker  $p$  weighted by the corresponding noise  $\epsilon_p$ :

$$A(R_c, R_t, \epsilon) = \frac{1}{P} \sum_p (R_c - R_t^p) \epsilon_p. \quad (6)$$

**Experiments.** We used our CAPTAIN framework to explore the properties of our model and the effect of different policies through simulations. Specifically, we ran three sets of experiments. The first set aimed at assessing the effectiveness of different policies optimized to minimize species loss based on different monitoring strategies. We ran a second set of simulations to determine how policies optimized to minimize value loss or maximize the amount of protected area may impact species loss. Finally, we compared the performance of the CAPTAIN models against the state-of-the-art method for conservation planning (Marxan<sup>25</sup>). A detailed description of the settings we used in our experiments is provided in the Supplementary Methods. Additionally, all scripts used to run CAPTAIN and Marxan analyses are provided as Supplementary Information.

**Analysis of Madagascar endemic tree diversity.** We analysed a recently published<sup>33</sup> dataset of 1,517 tree species endemic to Madagascar, for which presence/absence data had been approximated through species distribution models across 22,394 units of  $5 \times 5$  km spanning the entire country (Supplementary Fig. 5a). Their analyses included a spatial quantification of threats affecting the local conservation of species and assumed the cost of each protection unit as proportional to its level of threat (Supplementary Fig. 5b), similarly to how our CAPTAIN framework models protection costs as proportional to anthropogenic disturbance.

We re-analysed these data within a limited budget, allowing for a maximum of 10% of the units with the lowest cost to be protected (that is, 2,239 units). This figure can actually be lower if the optimized solution includes units with higher cost. We did not include temporal dynamics in our analysis, instead choosing to simply monitor the system once to generate the features used by CAPTAIN and Marxan to place the protected units. Because the dataset did not include abundance data, the features only included species presence/absence information in each unit and the cost of the unit.

Because the presence of a species in the input data represents a theoretical expectation based on species distribution modelling, it does not consider the fact that strong anthropogenic pressure on a unit (for example, clearing a forest) might result in the local disappearance of some of the species. We therefore considered the potential effect of disturbance in the monitoring step. Specifically, in the absence of more detailed data about the actual presence or absence of species, we initialized the sensitivity of each species to anthropogenic disturbance as a random draw from a uniform distribution  $d_s \sim \mathcal{U}(0, 1)$  and we modelled the presence of a species  $s$  in a unit  $c$  as a random draw from a binomial distribution with a parameter set equal to  $p_s^c = 1 - d_s \times D^c$ , where  $D^c \in [0, 1]$  is the disturbance (or 'threat' sensu Carrasco et al.<sup>33</sup>) in the unit. Under this approach, most of the species expected to live in a unit are considered to be present if the unit is undisturbed. Conversely, many (especially sensitive) species are assumed to be absent from units with high anthropogenic disturbance. This resampled diversity was used for feature extraction in the monitoring steps (Fig. 1c). While this approach is an approximation of how species might respond to anthropogenic pressure, the use of additional empirical data on species-specific sensitivity to disturbance can provide a more realistic input in the CAPTAIN analysis.

We repeated this random resampling 50 times and analysed the resulting biodiversity data in CAPTAIN using the one-time protection model, trained through simulations in the experiments described in the previous section and in the Supplementary Methods. We note that it is possible, and perhaps desirable, in principle to train a new model specifically for this empirical dataset or at least fine-tune a model pretrained through simulations (a technique known as transfer learning), for instance, using historical time series and future projections of land use and climate change. Yet, our experiment shows that even a model trained solely using simulated datasets can be successfully applied to empirical data. Following Carrasco et al.<sup>33</sup>, we set as the target of our policy the protection of at least 10% of each species range. To achieve this in CAPTAIN, we modified the monitoring action such that a species is counted as protected only when at least 10% of its range falls within already protected units. We ran the CAPTAIN analysis for a single step, in which all protection units are established.

We analysed the same resampled datasets using Marxan with the initial budget used in the CAPTAIN analyses and under two configurations. First, we used a BLM (BLM = 0.1) to penalize the establishment of non-adjacent protected units following the settings used in Carrasco et al.<sup>33</sup>. After some testing, as suggested in Marxan's manual<sup>45</sup>, we set penalties on exceeding the budget, such that the cost of the optimized results indeed does not exceed the total budget (THRESHPEN1 = 500, THRESHPEN2 = 10). For each resampled dataset we ran 100 optimizations (with Marxan settings NUMITNS = 1,000,000, STARTTEMP = -1 and NUMTEMP = 10,000 (ref. 45) and used the best of them as the final result. Second, because the BLM adds a constraint that does not have a

direct equivalent in the CAPTAIN model, we also repeated the analyses without it (BLM=0) for comparison.

To assess the performance of CAPTAIN and compare it with that of Marxan, we computed the fraction of replicates in which the target was met for all species, the average number of species for which the target was missed and the number of protected units (Supplementary Table 4). We also calculated the fraction of each species range included in protected units to compare it with the target of 10% (Fig. 6c,d and Supplementary Fig. 6c,d). Finally, we calculated the frequency at which each unit was selected for protection across the 50 resampled datasets as a measure of its relative importance (priority) in the conservation plan.

**Reporting Summary.** Further information on research design is available in the Nature Research Reporting Summary linked to this article.

### Data availability

All data necessary to run the analyses presented here are available in a permanent repository on Zenodo<sup>46</sup>.

### Code availability

The CAPTAIN software is implemented in Python v.3 and is available at [captain-project.net](https://github.com/captain-project/captain). All scripts and data necessary to run the analyses presented here are available in a permanent repository on Zenodo<sup>46</sup>.

Received: 12 April 2021; Accepted: 17 January 2022;

Published online: 24 March 2022

### References

- Diaz, S. et al. Assessing nature's contributions to people. *Science* **359**, 270–272 (2018).
- Chaplin-Kramer, R. et al. Global modeling of nature's contributions to people. *Science* **366**, 255–258 (2019).
- Dasgupta, P. *The Economics of Biodiversity: The Dasgupta Review* (HM Treasury, 2021).
- Barnosky, A. D. et al. Has the Earth's sixth mass extinction already arrived? *Nature* **471**, 51–57 (2011).
- Andermann, T., Faury, S., Turvey, S. T., Antonelli, A. & Silvestro, D. The past and future human impact on mammalian diversity. *Sci. Adv.* **6**, eabb2313 (2020).
- Global Biodiversity Outlook 5 – Summary for Policy Makers* (Secretariat of the Convention on Biological Diversity, 2020).
- Sterner, T. et al. Policy design for the Anthropocene. *Nat. Sustain.* **2**, 14–21 (2019).
- Mace, G. M. Whose conservation? *Science* **345**, 1558–1560 (2014).
- Leclère, D. et al. Bending the curve of terrestrial biodiversity needs an integrated strategy. *Nature* **585**, 551–556 (2020).
- Moilanen, A. Generalized complementarity and mapping of the concepts of systematic conservation planning. *Conserv. Biol.* **22**, 1655–1658 (2008).
- Margules, C. & Sarkar, S. *Systematic Conservation Planning* (Cambridge Univ. Press, 2007).
- Margules, C. R. & Pressey, R. L. Systematic conservation planning. *Nature* **405**, 243–253 (2000).
- Wilson, K. A., McBride, M. F., Bode, M. & Possingham, H. P. Prioritizing global conservation efforts. *Nature* **440**, 337–340 (2006).
- Visconti, P., Pressey, R. L., Bode, M. & Segan, D. B. Habitat vulnerability in conservation planning—when it matters and how much. *Conserv. Lett.* **3**, 404–414 (2010).
- Visconti, P., Pressey, R. L., Segan, D. B. & Wintle, B. A. Conservation planning with dynamic threats: the role of spatial design and priority setting for species' persistence. *Biol. Conserv.* **143**, 756–767 (2010).
- Wilson, K. A. et al. Prioritizing conservation investments for mammal species globally. *Phil. Trans. R. Soc. B* **366**, 2670–2680 (2011).
- Obura, D. O. et al. Integrate biodiversity targets from local to global levels. *Science* **373**, 746–748 (2021).
- Moilanen, A., Wilson, K. & Possingham, H. (eds) *Spatial Conservation Prioritization: Quantitative Methods and Computational Tools* (Oxford Univ. Press, 2009).
- Honeck, E., Sanguet, A., Schlaepfer, M. A., Wyler, N. & Lehmann, A. Methods for identifying green infrastructure. *SN Appl. Sci.* **2**, 1916 (2020).
- Bateman, I. J. et al. Bringing ecosystem services into economic decision-making: land use in the United Kingdom. *Science* **341**, 45–50 (2013).
- Carvalho, S. et al. Spatial conservation prioritization of biodiversity spanning the evolutionary continuum. *Nat. Ecol. Evol.* **1**, 0151 (2017).
- Sacre, E., Weeks, R., Bode, M. & Pressey, R. L. The relative conservation impact of strategies that prioritize biodiversity representation, threats, and protection costs. *Conserv. Sci. Pract.* **2**, e221 (2020).
- Ball, I. R., Possingham, H. P. & Watts, M. in *Spatial Conservation Prioritization: Quantitative Methods and Computational Tools* (eds Moilanen, A. et al.) 185–195 (Oxford Univ. Press, 2009).
- Pressey, R. L., Mills, M., Weeks, R. & Day, J. C. The plan of the day: managing the dynamic transition from regional conservation designs to local conservation actions. *Biol. Conserv.* **166**, 155–169 (2013).
- Watts, M., Klein, C. J., Tulloch, V. J. D., Carvalho, S. B. & Possingham, H. P. Software for prioritizing conservation actions based on probabilistic information. *Conserv. Biol.* **35**, 1299–1308 (2021).
- Tulloch, V. J. et al. Incorporating uncertainty associated with habitat data in marine reserve design. *Biol. Conserv.* **162**, 41–51 (2013).
- Grantham, H. S., Wilson, K. A., Moilanen, A., Rebelo, T. & Possingham, H. P. Delaying conservation actions for improved knowledge: how long should we wait? *Ecol. Lett.* **12**, 293–301 (2009).
- Zizka, A., Silvestro, D., Vitt, P. & Knight, T. M. Automated conservation assessment of the orchid family with deep learning. *Conserv. Biol.* <https://doi.org/10.1111/cobi.13616> (2020).
- Gomes, C. et al. Computational sustainability: computing for a better world and a sustainable future. *Commun. ACM* **62**, 56–65 (2019).
- Spring, D. A., Cacho, O., MacNally, R. & Sabbadin, R. Pre-emptive conservation versus “fire-fighting”: a decision theoretic approach. *Biol. Conserv.* **136**, 531–540 (2007).
- Meir, E., Andelman, S. & Possingham, H. P. Does conservation planning matter in a dynamic and uncertain world? *Ecol. Lett.* **7**, 615–622 (2004).
- Adams, V. M. & Setterfield, S. A. Optimal dynamic control of invasions: applying a systematic conservation approach. *Ecol. Appl.* **25**, 1131–1141 (2015).
- Carrasco, J., Price, V., Tulloch, V. & Mills, M. Selecting priority areas for the conservation of endemic trees species and their ecosystems in Madagascar considering both conservation value and vulnerability to human pressure. *Biodivers. Conserv.* **29**, 1841–1854 (2020).
- Kosmala, M., Wiggins, A., Swanson, A. & Simmons, B. Assessing data quality in citizen science. *Front. Ecol. Environ.* **14**, 551–560 (2016).
- Devillers, R. et al. Reinventing residual reserves in the sea: are we favouring ease of establishment over need for protection? *Aquat. Conserv.* **25**, 480–504 (2015).
- Vieira, R. R. S., Pressey, R. L. & Loyola, R. The residual nature of protected areas in Brazil. *Biol. Conserv.* **233**, 152–161 (2019).
- Dilkina, B. et al. Trade-offs and efficiencies in optimal budget-constrained multispecies corridor networks. *Conserv. Biol.* **31**, 192–202 (2017).
- Bruehlheide, H. et al. sPlot – a new tool for global vegetation analyses. *J. Veg. Sci.* **30**, 161–186 (2019).
- Geldmann, J., Manica, A., Burgess, N. D., Coad, L. & Balmford, A. A global-level assessment of the effectiveness of protected areas at resisting anthropogenic pressures. *Proc. Natl Acad. Sci. USA* **116**, 23209–23215 (2019).
- Geldmann, J. et al. A global analysis of management capacity and ecological outcomes in terrestrial protected areas. *Conserv. Lett.* **11**, e12434 (2018).
- Diaz, S. et al. Set ambitious goals for biodiversity and sustainability. *Science* **370**, 411–413 (2020).
- Barto, A. G. Reinforcement learning and dynamic programming. *IFAC Proc. Volumes* **28**, 407–412 (1995).
- Salimans, T., Ho, J., Chen, X., Sidor, S. & Sutskever, I. Evolution strategies as a scalable alternative to reinforcement learning. Preprint at <https://arxiv.org/abs/1703.03864> (2017).
- Sutton, R. S., McAllester, D. A., Singh, S. P. & Mansour, Y. in *Advances in Neural Information Processing Systems* (eds Solla, S. et al.) 1057–1063 (MIT Press, 1999).
- Serra-Sogas et al. *Marxan User Manual: For Marxan Version 2.43 and Above* (The Nature Conservancy and PacMARA, 2021).
- Silvestro, D., Gorla, S., Sterner, T. & Antonelli, A. Supplementary data for: Improving biodiversity protection through artificial intelligence. *Zenodo* <https://doi.org/10.5281/zenodo.5643665> (2021).
- Murcia, C. Edge effects in fragmented forests: implications for conservation. *Trends Ecol. Evol.* **10**, 58–62 (1995).

### Acknowledgements

We thank K. Johannesson for the initial discussions that sparked this project, B. Groom, A. Moilanen, B. Pressey, K. Dhanjal-Adams and I. Bateman for invaluable feedback on an earlier version of this manuscript, J. Carrasco for providing the Madagascar tree distribution data, J. Moat for input on remote sensing technologies, W. Testo for suggesting the acronym CAPTAIN, R. Smith for editorial support, G. Dudas for coding support and the many colleagues in our respective research groups for discussions. A.A. is funded by the Swedish Foundation for Strategic Research (FFL15-0196), the Swedish Research Council (VR 2019-05191) and the Royal Botanic Gardens, Kew. T.S. and A.A. acknowledge funding from the Strategic Research Area Biodiversity and Ecosystem Services in a Changing Climate, BECC, funded by the Swedish government. D.S. received funding from the Swiss National Science Foundation (PCEFP3\_187012) and the Swedish Research Council (VR: 2019-04739).

### Author contributions

T.S. initiated the project with A.A. D.S. and S.G. developed the AI model and performed all analyses. A.A., D.S., T.S. and S.G. wrote the paper.

### Competing interests

The authors declare no competing interests.

**Additional information**

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41893-022-00851-6>.

**Correspondence and requests for materials** should be addressed to Daniele Silvestro or Alexandre Antonelli.

**Peer review information** *Nature Sustainability* thanks Stephen Newbold and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022

**Supplementary information**

---

**Improving biodiversity protection through artificial intelligence**

---

In the format provided by the authors and unedited

# Supplementary information

## Improving biodiversity protection through artificial intelligence

Daniele Silvestro<sup>1,2,\*,§</sup>, Stefano Gorla<sup>3</sup>, Thomas Sterner<sup>4</sup> Alexandre Antonelli<sup>2,5,6,\*,§</sup>

<sup>1</sup> *Department of Biology, University of Fribourg and Swiss Institute of Bioinformatics, 1700 Fribourg, Switzerland*

<sup>2</sup> *Gothenburg Global Biodiversity Centre, Department of Biological and Environmental Sciences, University of Gothenburg, 40530 Gothenburg, Sweden*

<sup>3</sup> *Thymia Limited, London, United Kingdom*

<sup>4</sup> *Department of Economics, University of Gothenburg, 40530 Gothenburg, Sweden*

<sup>5</sup> *Department of Plant Sciences, University of Oxford, Oxford, OX1 3RB, United Kingdom*

<sup>6</sup> *Royal Botanic Gardens, Kew, TW9 3AE, United Kingdom*

\*Corresponding authors: daniele.silvestro@unifr.ch, a.antonelli@kew.org

§ These authors contributed equally.

### Supplementary methods

**Natural evolution of the system.** We define as *system* a spatially explicit simulated set of  $N$  species, each with  $Q^s$  individuals distributed across a 2D grid of  $Z$  cells. The phylogenetic relationships among species are described by a time-calibrated phylogeny  $\phi$ . The initialization of the system used in our simulations is described in the *Experiments* section. The evolution of the system is discretized into time steps (e.g., years). We model natural processes that include the death of a fraction of individuals across all cells and their replacement. In the notation used in the following paragraphs, we indicate the vector including the number of individuals for each species in a cell at time  $t$  as  $H_t^c = \{H_t^{1,c}, \dots, H_t^{N,c}\}$  and as  $H_t$  the matrix of size  $Z \times N$  including the individuals counts across all species and all cells. The notation and all main parameters in the framework are summarized in Supplementary Table 1.

At each time step, a fraction of individuals die and are replaced by a number of new individuals sampled proportionally to their abundance in each cell, their dispersal ability (determined by a dispersal rate  $\lambda$ ), and their distance. Each cell is characterized by a carrying

capacity  $K_c$  determining the maximum number of individuals from any species that can occur in the cell. In our simulations we make a number of simplistic assumptions, which can be easily relaxed or modified in the program. In particular, we set the natural mortality and dispersal rates as equal among species and we assume that individuals of each species contribute equally toward the carrying capacity. Under this system, species are unlikely to naturally go extinct in the short term even if stochastic variation in death rates per species is applied. This is because the individuals dying at each time step are replaced by as many individuals (within the carrying capacity limits) sampled proportionally across species.

Each species  $s$  in cell  $i$  disperses a number of offspring to cell  $j$  equal to

$$n_{sij} = H_t^{s,i} \times g_s \times \exp\left(-\frac{1}{\lambda_s} d_{ij}\right) \quad (1)$$

where  $d_{ij}$  is the distance between the cells  $i$  and  $j$ , and  $g_s$  is a growth rate, indicating the number of dispersing offspring when the distance is  $d_{ij} = 0$  (i.e. when  $i = j$ ). The number of new individuals added to a cell at each time step is determined by the cumulative number offspring of all individuals and all  $N$  species and from all  $Z$  cells in the systems, within the limits allowed by the cell's carrying capacity,  $K_j$ :

$$n_j = \frac{K_j - \sum_s^N H_t^{s,j}}{K_j} \left( \sum_i^Z \sum_s^N n_{sij} \right) \quad (2)$$

Under this formulation and with the simulations settings used in our experiments (low mortality rates compared with growth and dispersal rates), all individuals dying in a time step are replaced by new ones, thus maintaining the density of individuals at carrying capacity. However, the growth rate will limit the speed of recovery if an event of higher mortality strikes a cell, for instance due to an episode of extreme climate or to a sudden increase in anthropogenic disturbance. While in our simulations we used deterministic counts of how many individuals die at each time step, stochastic variation can be introduced in the model.

**Evolution of anthropogenic effects.** Anthropogenic impact in the system is modelled through different processes. We use one layer of disturbance, which characterizes the general impact of human activities on each cell. The disturbance  $D_t^c$  increases the mortality rate in a cell  $c$  at time  $t$  and reduces the carrying capacity of a cell such to  $(1 - D_t^c)K_i$  so that, e.g. with

$D_c = 0.75$ , only 25% of the theoretical maximum number of individuals can inhabit the cell. This layer captures land-use changes, such as urbanization, that reduces the available space for species. Importantly, disturbance affects the mortality rate of individuals differently as a function of a species-specific sensitivity  $d_s$ , such that the disturbance-driven mortality rate of an individual of species  $s$  in cell  $c$  is:

$$\delta_t^{s,c} = \max \left[ -\log(1 - d_s) \left( \frac{H_t^c}{(1 - D_t^c)K_c} - 1 \right), 0 \right] \quad (3)$$

Thus, death rates increase with higher disturbance as this reduces the carrying capacity of a cell and with increasing species-specific sensitivity, such that  $\delta_t^{s,c}$  tends to infinity as  $d_s$  approaches 1. Based on this definition, the fraction of individuals of species  $s$  dying in a cell  $c$  due to disturbance in one time step is  $1 - \exp(-\delta_t^{s,c})$ . Once reached the new carrying capacity  $(1 - D_t^c)K_c$ , the system returns to equilibrium, as  $H_t^c / (1 - D_t^c)K_c - 1 = 0$ , and only natural mortality and replacement occur. However, because sensitivity to disturbance differs among species, the relative abundances of species in each cell is likely to have changed upon reaching the new equilibrium, and these changes may involve the disappearance of rare and/or sensitive species from the cell.

We also model a second layer of selective disturbance, indicated with  $F_t$  where  $F_t^c \in (0, 1)$ , affects the mortality of individuals in cell  $c$  based on a species-specific sensitivity  $f_s \in (0, 1)$ , but does not alter the cell's carrying capacity. This layer captures some land-use changes (e.g., turning a forest into pasture), activities specifically targeting particular species (e.g., timber logging). In a species  $s$ , a fraction  $f_s F_t^c$  of individuals will die due to selective disturbance at each step, in cell  $c$ . As this does not affect the carrying capacity of a cell, growth and migration rates will replace these individuals with others, increasingly representative of species with relatively low sensitivity ( $f_s$ ). Thus, selective disturbance can cause the (local) disappearance of species and alter their natural rank abundance.

**Climate change.** We included a third layer of disturbance due to climate change, which affects the mortality of individuals in a cell based on a species-specific sensitivity. Sensitivity to climate change can be interpreted as drought tolerance, resistance to extreme weather or increasing temperatures, among others. Conceptually, climate change disturbance differs from the other two forms of disturbance as it cannot be mitigated by a protection policy. While the

climate layer can potentially represent different climatic variables, we simplify this here by referring to it as mean annual temperature.

In our experiments, we initialize a climate layer as a latitudinal gradient with cooler temperatures in the top cells in the system, which increase linearly towards the bottom of the grid. We define empirically the temperature tolerance of each species  $\tau_s$  as the difference between the maximum and the minimum temperatures at which the species occurs in the initial system. As the system evolves, individuals of a species can occur outside their temperature range because they disperse outside of their initial geographic range, or because they track their preferred temperatures during climate change.

For individuals occurring outside of their climatic tolerance, we define as  $\Delta T_{sc}$  the absolute distance between the mean annual temperature in the cell  $T_c$  and the species temperature range  $\tau_s$ . We set  $\Delta T_{sc} = 0$  for individuals occurring within their tolerance range and define the death rate of species  $s$  driven by climate as:

$$\delta(s, T_c) = \frac{1}{\tau_s} \Delta T_{sc}$$

Thus, climate has the effect of linearly increasing the death rates of individuals occurring outside of their climate tolerance. Additionally, as the death rate is an inverse function of  $\tau_s$ , species with a wider tolerance are subject to lower death rates than species with narrower tolerance.

## Experiments settings.

**Initialization of the systems.** Across all our simulations, we used systems of  $50 \times 50$  cells with uniform carrying capacity across all cells set at  $K = 1,000$ . Each simulated system was initialized with 500 species with random Weibull-distributed population sizes to reflect empirical rank-abundance plots<sup>1</sup>, totalling 2.5 million individuals (Supplementary Fig. 1a). Species ranges were generated based on a random diffusion process, with the number of individuals in a cell constrained by the carrying capacity (Supplementary Fig. 2).

The natural mortality was set constant across species and equal to 0.01, implying that 1% of the individuals of each species die at each time step. We set the growth rate to 1, indicating that, on average, each individual gives origin to one offspring. In our simulations we also

assumed dispersal rates to equal across all species ( $\lambda = 0.1$ ), with the distance between adjacent cells set to 1. Species were assigned a random beta sensitivity to disturbance ( $d_s \sim \mathcal{B}[1, 1]$ ) and sensitivity to selective disturbance ( $f_s \sim \mathcal{B}[0.2, 0.7]$ ) (Supplementary Fig. 1c,d).

To track the effects of biodiversity loss on phylogenetic diversity, we simulated phylogenetic trees describing the relationships and divergence times among all species. Phylogenies with 500 extant species were simulated based on a birth-death process using the TreePar package<sup>2</sup> with a speciation rate set to 1 and extinction rate sampled for each tree from a uniform distribution  $\mathcal{U}[0, 0.9]$ .

We assigned an economic value to each species using an approach that generates a high variation among species but also a geographic pattern, such that species occurring in some regions tend to be more valuable than species occurring in others (e.g., some types of forest might harbour more valuable species for timber than others, depending on moisture and soil gradients). We first randomly initialized species values such that 20% of the species had a value 100 times higher than each of the remaining 80% of the species. We then randomly selected a cell in the system and divided the value of each species by the distance between the centroid of the species' geographic range and the selected cell. The values were then re-scaled so that that mean value across all species became 1 (Supplementary Fig. 1b).

In our simulations we defined the cost of protection ( $C_c$ ) for a cell as a function of the current level ( $D_c$ ) of anthropogenic disturbance:  $C_c = 0.2 + 0.4D_c$ , where 0.2 is an arbitrary baseline cost. Under this assumption, the cost of protecting a cell with full disturbance is three times higher than that of a cell with no disturbance. The rationale for this disturbance-dependent cost is that enforcing protection implies removing (or relocating) the source of disturbance. If there is a strong disturbance this can be interpreted as a considerable amount of economic activity (e.g. large town or mine) which will be expensive to remove or relocate.

**Monitoring policies to minimize species loss.** We performed simulations to assess the effect of different monitoring policies on biodiversity loss. We modelled a dynamic system in which an initial area is subject to a level of disturbance that intensifies and expands spatially over time (Supplementary Fig. 3). Disturbance values are initialized at 0 across all cells and then increased based on random diffusion processes to create spatial heterogeneity in the disturbance intensity and a trend toward increased total disturbance through time.

We divided the system into 100 protection units of  $5 \times 5$  neighbouring cells and set the initial budget to 55. Given a baseline cost of 0.2 per cell, the cost of protection for one unit is 5 in the absence of disturbance, thus a maximum of 11 units can be protected under these settings. However, since the cost of protection is also a function of disturbance, the effective number of protected units is lower. The protection policy in these simulations allowed for the selection of a single protection unit at each time step.

We tested three monitoring policies which differ in the frequency of monitoring and the quality and quantity of information obtained through the feature extraction function.

- *Full Recurrent Monitoring*. This policy involves monitoring at each time step with the feature extraction function tracking the number of non-protected rare species, the mean relative change in populations size compared to previous monitoring, and budget and cost of each unit (features  $x_1, x_2, x_3, x_5, x_6$  in Supplementary Table 2). The protection policy included one meta-parameter (defining the threshold for rare species; Supplementary Table 2) and a neural network with two hidden layers of 4 and 2 nodes, respectively.
- *Citizen science Recurrent Monitoring*. This policy involves the monitoring of the system at each time step, but with a lower amount and quality of information compared to the full monitoring policy (features  $x_1, x_5, x_6$ ). Monitoring only tracks the number of non protected species per unit without information about population sizes. The monitoring also includes the available budget and cost of protection. We additionally introduced a fraction of stochastic error in the feature extraction function. Specifically, we assumed that an average of 5% of the species in a unit might go undetected (with probability inversely proportional to their abundance in the unit) and that another 5% of the species can be misidentified. To assess the effect of monitoring error in the outcome of the model we also tested the model with the error set to 25%. The protection policy included a neural network with two hidden layers of 4 and 2 nodes, respectively.
- *Full Initial Monitoring*. This policy involves monitoring only at the first time step, with the feature extraction function tracking the number of non protected species, rare species, and budget and cost of each unit (features  $x_1, x_2, x_5, x_6$ ). Since monitoring only takes place once, no information about trends in population sizes is included. We consider the cost feature ( $x_5$ ) and the protection indicator ( $x_6$ ) as known even without monitoring, and

therefore update them at each step. The protection policy included one meta-parameter (defining the threshold for rare species) and a neural network with two hidden layers of 4 and 2 nodes, respectively.

The models were optimized using the algorithm described above, running 6 simulations in parallel for 1,000 epochs. Each simulation was based on a different randomly initialized system and ran for 25 time steps. We then used the estimated parameters to assess the effect of different monitoring policies in 250 additional simulations. Each simulation was run for 25 time steps and we fixed the seeds of the random number generators used in the runs to ensure the simulations were comparable among monitoring policies (i.e. based on the same underlying system and disturbance processes). For comparison, we repeated the 250 simulations using a random protection policy, in which non-protected units are assigned a uniform probability of being selected for protection. Additionally, we repeated the analyses using a greedy algorithm that at each step selects the protection unit with the highest ratio between the number of non-protected species and its cost.

We summarized the outcome of the simulations by measuring at the last time step the species loss, value loss, phylogenetic diversity loss and the amount of protected area. We report all measurements as percentage changes compared to the baseline outcome of the random policy.

**Testing the trade-offs among species, value, and area losses.** We ran simulations in which the policy aimed at minimizing value loss or at maximizing the number of protected units, using the same settings as in the simulations described above.

- *Minimizing value loss.* In this policy the negative reward was set equal to the cumulative value of extinct species. We applied a full recurrent monitoring, with the feature extraction function tracking the cumulative value of non protected species, that of rare species with declining population sizes, as well as budget and cost of each unit (features  $x_4, x_5, x_6, x_7, x_8$ ). The protection policy included one meta-parameter (defining the threshold for rare species) and a neural network with two hidden layers of 4 and 2 nodes, respectively.
- *Maximizing protected area.* In this policy the positive reward was set equal to the amount of protected area, regardless of which and how many species occurred in it. We applied

recurrent monitoring, but the feature extraction function here only returned the cost of each unit and the available budget (features  $x_5, x_6$ ). The protection policy included a neural network with one hidden layer of 4 nodes.

We optimized the models, performed simulations and summarized the results as described in the section above.

**Comparison with Marxan.** We ran additional simulations to compare the performance of our model optimized through reinforcement learning framework with Marxan<sup>3</sup> (v. 4.0.6), one of the state-of-the-art and most widely used programs for systematic conservation planning. By default Marxan optimizes the placement of protection units to meet a predefined conservation target while minimizing the cost. A Marxan analysis can be however set up to run within a limited budget<sup>4</sup> by setting a maximum budget (parameter COSTTHRESH; here set to 55 matching our CAPTAIN settings and allowing for a maximum of 11 protected units) and penalty scores that prevent the algorithm from exceeding the budget (parameters THRESHPEN1 and THRESHPEN2 here set both equal to 10). Given the relatively small number of protection units (100 in our simulations), we set the boundary length multiplier (BLM) to 0, thus not applying any penalties for choosing non-adjacent protection units. We applied the same cost of protection across units and the same budget used in the CAPTAIN analyses, such that it allowed the protection of at most 10% of the protection units in the system.

Because Marxan is designed to identify all protection units in one step, we performed a Full Initial Monitoring of the system to generate Marxan input files containing: a definition of the protection units (as in Fig. 1a) and their cost, the presence and absence of species across units and their abundance (number of individuals), and the conservation target for each species. Since in all our CAPTAIN simulations we considered species as extinct when their total population size is lower than 10 individuals, we set a conservation target of 10 individuals per species to be included within the designated protected units in the Marxan setup. We performed two sets of experiments, each including 250 simulations: in the first (one-time protection policy), all budget is spent in one step after monitoring the system for three steps, while in the second (dynamic protection policy) the establishment of protected units is done dynamically as the system evolves.

- *One-time protection policy.* In CAPTAIN, we trained a model (using the optimization al-

gorithm described above) based on recurrent full monitoring performed in the first three time steps followed by the placement of all protected units. The protection policy included one meta-parameter (defining the threshold for rare species; Supplementary Table 2) and a neural network with one hidden layer of 8 nodes. In Marxan, the system was monitored as described above at step three. We then performed 10 optimizations using the simulated annealing algorithm followed by iterative improvement<sup>4</sup> (with settings: NUMITNS = 1,000,000; STARTTEMP = -1; NUMTEMP = 10000) and chose the output of the best among them to select the protected units. For both CAPTAIN and Marxan policies, the system evolved until reaching 25 steps with increasing and spatially inhomogeneous disturbance (Supplementary Fig. 3).

- *Dynamic protection policy.* In CAPTAIN, this corresponds to the Full Recurrent Monitoring experiment described in the section *Monitoring policies to minimize species loss*. In Marxan, the system was monitored and the model optimized as in the One-time protection policy, but the selected units were established one at a time, i.e. placing at most one unit at each step. Following the settings used in our previous experiments, the system evolved through time with increasing and spatially heterogeneous disturbance.

For comparison, as in the other experiments, we repeated the simulations using a random protection policy and a greedy algorithm that at each step selects the protection unit with the highest ratio between the number of non-protected species and its cost. We summarized the outcome of the simulations by measuring at the last time step the species loss (measured as number of species with a remaining population size greater than 10 individuals), value loss, phylogenetic diversity loss and the amount of protected area. We report all measurements as percentage changes compared to the baseline outcome of the random policy.

### **Supplementary Box 1: Socio-economic valuation of biodiversity**

The field of economics sees nature as an asset, analogous to diversity in portfolio management<sup>5,6</sup>. It provides valuable insurance that may prove vital<sup>7</sup>. Risks are reduced if two assets are negatively correlated. Weitzman<sup>8</sup> incorporates the idea of redundancy when species share genes, and shows that optimal policy under a budget constraint will involve strict prioritization of some species.

The value of ecosystem services hinges on ease of substitution by human-made capital. In some respects, humans can “replace” a coral reef by artificial structures. Fish still find protection and food, but the artificial reef does not fully “replace” nature. When nature is irreplaceable, we speak of strong (rather than weak) sustainability; meaning more care must be taken to protect natural capital assets<sup>9</sup>.

We see massive degradation of natural assets that could have yielded a high return if properly managed by effective owners or stewards. Institutional mechanisms (including secure property rights) are vital to protect nature<sup>10</sup>. Still, private incentives to preserve biodiversity (e.g. research and development in biodiversity-based solutions) remain weak<sup>11</sup>, and strong policies are needed<sup>12</sup>.

Both benefits and costs of conservation vary considerably across space. Protection is typically expensive for farmland or real estate close to cities. Some land is cheap but protection enforcement expensive precisely because of remoteness. Some areas may be valuable because they have high biodiversity or endemism. What the framework we developed here strives for is high total biodiversity benefits with low land and opportunity costs, at the same time as it reduces risks associated with climate change uncertainties<sup>13</sup>. Spatial targeting of sites that increase biodiversity protected per monetary unit is cost effective, and may also include recreational and other values<sup>14,15</sup> that we hope to develop more in future model versions.

## Supplementary Box 2: New techniques for biodiversity monitoring

Our simulations show that the regular monitoring of species is critical for prioritizing areas for conservation under changing anthropogenic and climate pressure (main text, Fig. 2). But given scarce resources and vast areas to monitor, how can this be most efficiently accomplished?

For centuries, biodiversity inventories have been carried out by biodiversity experts. This work has comprised the collection, recording and eventual identification of biological samples, such as fertile tree branches or whole animals – a time- and resource-consuming process, resulting in taxonomically and spatially biased and incomplete biodiversity data<sup>16</sup>. New approaches are now speeding up and popularizing biodiversity monitoring, including environmental DNA, remote sensing and citizen science.



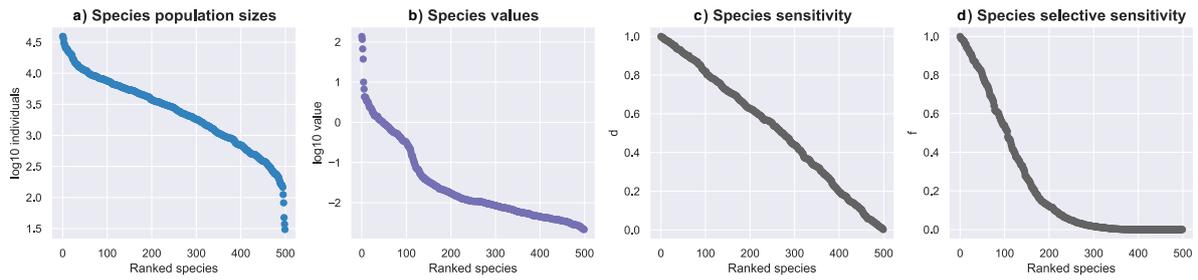
Rather than locating and recording each individual species directly, soil samples can be taken by non-specialists (a) and their environmental DNA (eDNA) contents subsequently analyzed. This approach reveals the biotic composition of whole communities, including plants, animals, fungi and bacteria. Standardized sampling also allows direct comparison of taxonomic and phylogenetic diversity across multiple sites<sup>17</sup>, helping to inform on presence and absence of species. Challenges remain in completing reference databases for matching sequences to names, and in separating local versus transported DNA from nearby areas.

Citizen science initiatives such as 'bio-blitzes' (an intense period of biological surveying, often by amateurs) can help fill up critical data gaps, reduce monitoring costs and

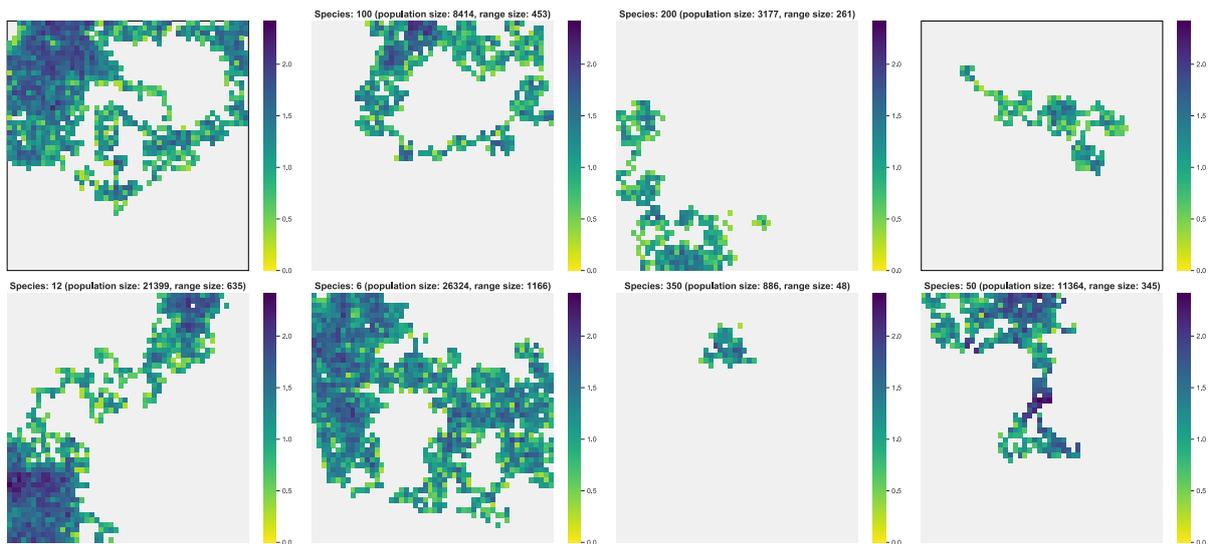
increase local awareness and engagement for biodiversity. An impressive 3.5 billion smartphones are now in use around the world<sup>18</sup>, creating unseen possibilities of linking millions of people to nature through species logging (b), using platforms such as iNaturalist<sup>19</sup> which increasingly use image-recognition software and geolocations for automated species identification. In our simulations, even identification errors of up to 25% during monitoring steps (a conservative rate for many organism groups, such as birds, mammals or trees) have only a marginal impact on the biodiversity outcomes (Supplementary Table 3). [Photo credits: a-b: A.A.; c: Justin Moat].

Finally, remote sensing technologies can rapidly scan and characterize relatively large areas. Using 3-D sensing technology, they now hold the promise of automated identification of species using multi-spectrum images<sup>20</sup>. The spectral signatures of species can be produced by combining on-the-ground (visible to NIR spectrophotometry of plant material and terrestrial LiDAR) with from-the-air (hyperspectral and LiDAR UAV survey) data. This allows the isolation of individual objects to record the presence and abundance of species (c). The technology can also provide measures of vegetation structure (informing on an areas' functional diversity) and ecosystem services, such as carbon storage, soil water content, and other parameters that could be easily implemented in our model to help the valuation and ranking of areas for protection, besides greatly improving the quality of habitat data publicly available.

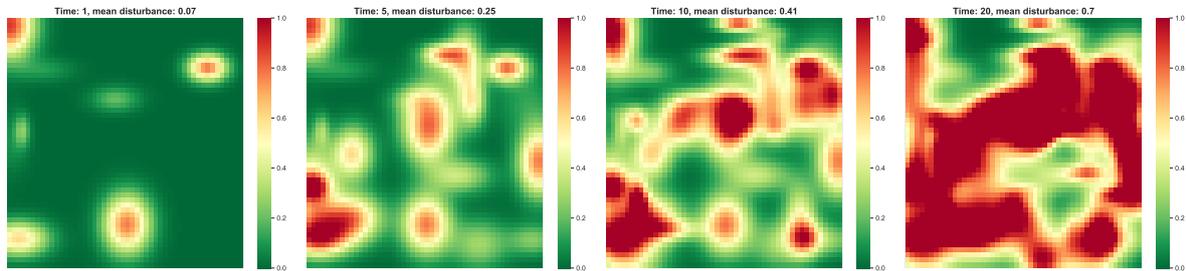
## Supplementary Figures



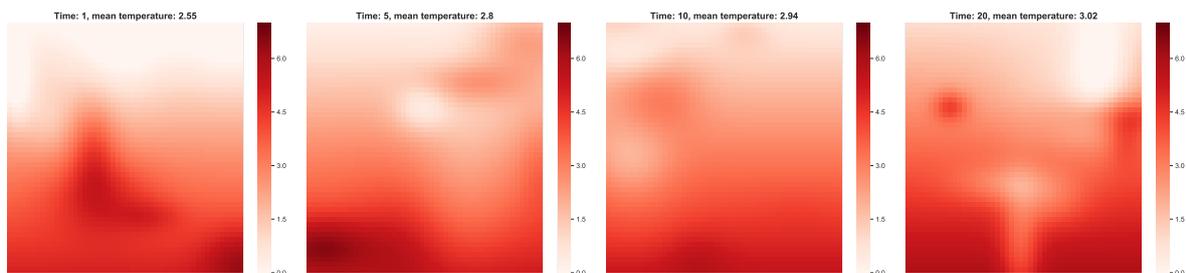
Supplementary Figure 1. Value distribution across four key species-specific model parameters. For a randomly selected simulation, the graphs show a, initial species population sizes, b, species economic values, c, species sensitivity to anthropogenic disturbance and d, species selective sensitivity (affecting them disproportionately, due e.g. to selective logging of trees). In each plot, species are ranked by the plotted metric. Population size and species value are log<sub>10</sub> transformed.



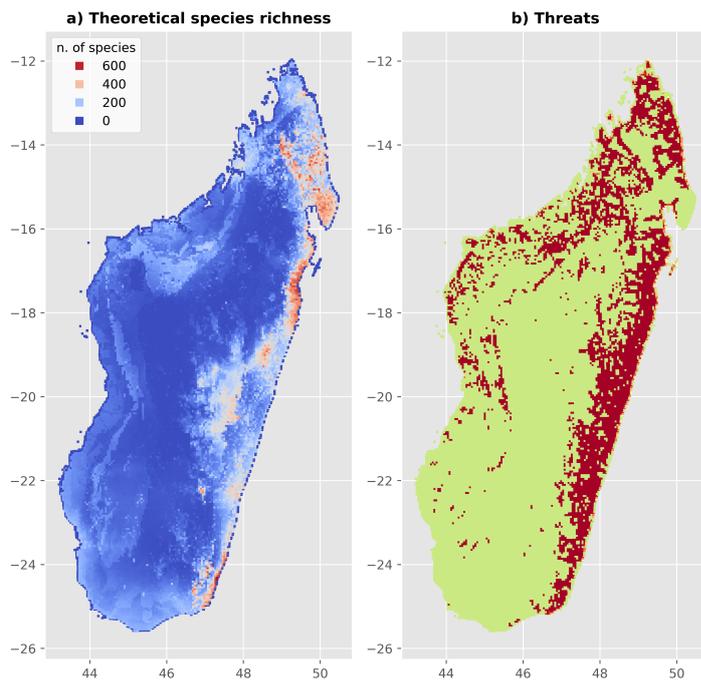
Supplementary Figure 2. Examples of species ranges. For eight randomly selected species, the plots show their total distribution (non-grey cells; summed to a value of range size) as well as total population sizes (i.e., the number of individuals in the system, or abundance). Cells are coloured based on the log<sub>10</sub>-transformed number of individuals.



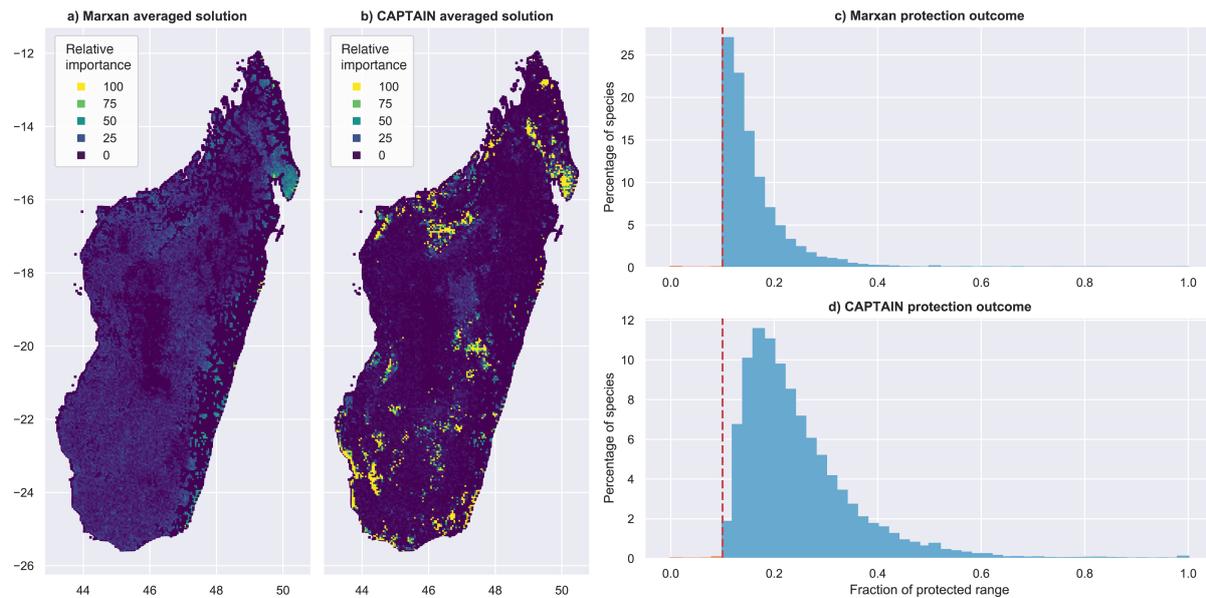
Supplementary Figure 3. Evolution of anthropogenic disturbance across the simulated system. For a randomly selected simulation, the colours show changes in human disturbance across individual cells at time steps 1, 5, 10, and 20 in a system where no protected areas were set. The mean overall disturbance increases 10-fold from 0.07 to 0.7.



Supplementary Figure 4. Evolution of climate change across the simulated system. For a randomly selected simulation, the colours show temperature changes across individual cells at time steps 1, 5, 10, and 20 resulting in an average increase across the entire system of 0.47 degrees with random local deviations at each step, and the overall increase in mean temperature across the entire system.



Supplementary Figure 5. Madagascar tree diversity and threats. The maps show data obtained from Carrasco et al.<sup>Ref. 34, main text</sup>: a) estimated diversity of endemic trees as predicted by species distribution modelling; b) threats to tree conservation, with redder colours showing higher levels of threat (*sensu* Carrasco et al.<sup>Ref. 34, main text</sup>).



Supplementary Figure 6. Ranking of potential protection units across Madagascar: The maps show a ranking of priority areas for protection across Madagascar based on the distribution of endemic trees based on a limited budget allowing for up to 10% of protected area overall. The colour scale is proportional to the relative importance of each protected unit measured as the frequency at which it was included in the optimised solution across 50 replicated analyses randomizing species sensitivities to disturbance a), protection unit ranks resulting from Marxan optimizations (with BLM = 0); b), equivalent CAPTAIN results based on a full initial monitoring policy; c-d) histograms of the fraction species range included in the protected units calculated for each species. The red dashed lines indicate the 10% threshold set as the target of the policy.

## Supplementary Tables

Supplementary Table 1. List of variables used in our model and their notation. Notation  $t$  always indicates “at time  $t$ ”.

Parameter	Explanation
$Z$	Number of cells in the system
$N$	Number of species
$K_c$	Carrying capacity of a cell (total number of individuals it can contain) $c$
$H_t$	Matrix with number of individuals for each species across each cell
$H_t^c$	Number of individuals for each species in cell $c$
$H_t^{c,s}$	Number of individuals of species $s$ in cell $c$
$Q_t^s$	Total number of individuals of species $s$ , i.e. $\sum_c^Z (H_t^{c,s})$
$E_t^s$	Share of individuals of species $s$ , i.e. $\sum_c^Z (H_t^{c,s}) / \sum_s^N \sum_c^Z (H_t^{c,s})$
$g_s$	Growth rate of species $s$
$\delta_t^{c,s}$	Death rate (cumulating the natural and disturbance-induced mortalities)
$\lambda^s$	Dispersal rate
$\phi$	Phylogenetic tree connecting all species in the system
$D_t$	Disturbance matrix
$D_t^c$	Disturbance at cell $c$
$d_s$	Sensitivity to disturbance of species $s$
$F_t$	Selective disturbance matrix
$F_t^c$	Selective disturbance at cell $c$
$f_s$	Sensitivity to selective disturbance of species $s$
$T_t$	Climate matrix
$T_t^c$	Climate in cell $c$
$\tau_s$	Climatic tolerance of species $s$
$P_t$	Protection matrix
$P_t^c$	Protection level of cell $c$
$B_t$	Available budget
$C_t$	Matrix of cost of protection
$C_t^c$	Cost of protection of cell $c$
$S_t$	State of the system. This includes $\{H_t, D_t, F_t, T_t, C_t, P_t, B_t\}$
$\pi_X$	Monitoring policy setting the frequency of monitoring and defining the function $X(S_t)$
$\pi_Y$	Protect policy
$U$	Number of protection units
$X(S_t)$	Function extracting features from last monitored state of the system
$\{x_1, x_2, \dots, x_n\}$	List of features extracted for each protection unit

Parameter	Explanation
$W$	Number of coefficients in the neural network determining the protection policy
$M$	Number of metaparameters (in our models $M = 1$ )
$q$	Metaparameter defining the threshold for rare species when extracting features
$Y(X(S_t), W, q)$	Function choosing the unit to protect based on the extracted features and estimated parameters $W$ and $q$
$r_t$	Reward in one time step $t$ . This depends on the policy, e.g. $r_t = -1$ if the policy aims to minimize species loss and one species went extinct at time $t$ .

Supplementary Table 2. List of features that can be extracted through monitoring, using the function  $X(S_t)$ .

Parameter	Explanation
$x_1$	Number of non-protected species, i.e. species not included in other already protected units. In systematic conservation planning, this value is critical to evaluate complementarity of protected areas
$x_2$	Number of non-protected rare species. This is a subset of the count ( $x_1$ ) in which only species with a population size smaller than a threshold $t$ are included. The threshold is defined as: $\tau = \exp(q \log(Q_{\max}))$ , where $Q_{\max}$ is the largest population size among all species and $q \in (0, 1)$ is an estimated parameter, such that if $q = 1$ all species are counted, whereas with $q < 1$ only a fraction of species will be included in the count. Note that, depending on the monitoring policy, $Q_{\max}$ can be updated at the time of monitoring, thus reflecting the largest population size at that time
$x_3$	Mean relative change in populations size <div style="text-align: center; margin: 10px 0;"> <math display="block">\frac{1}{N_u} \sum_{s \in u} \left( \frac{Q_t^s}{Q_{t-1}^s} - 1 \right)</math> </div> <p>where the population size of species <math>s</math> at time <math>t</math> (<math>Q_t^s</math>) is divided by a previous measure (<math>Q_{t-1}^s</math>; time <math>t - 1</math> indicates the time of the previous monitoring) and averaged across all <math>N_u</math> species present in a protection unit <math>u</math>. We note that to generate this feature the feature extraction function requires information from a previous state of the system <math>X(S_t, S_{t-1})</math>.</p>
$x_4$	Available budget ( $B_t$ )
$x_5$	Cost of the unit
$x_6$	An indicator defining whether the unit is already protected (1) or not (0)
$x_7$	Cumulative value of non-protected species
$x_8$	Cumulative value of non-protected rare species (as defined in $x_2$ ) with declining relative population sizes (as defined in $x_3$ )

**Supplementary Table 3. Results from 250 simulations based on models trained with different and monitoring policy, policy objectives, and protection policy.**

The results are shown in terms of mean percentage change compared to simulations using a random protection policy. Values in parentheses show the standard deviation (sd) of the percentage change across simulations and the proportion of simulations (P) yielding a positive change compared to the respective simulations under a random policy.

Monitoring policy	Policy objective	Protection policy	Change relative to a random policy			
			Protected area	Specie loss	PD loss	Value loss
Full recurrent monitoring	species loss	dynamic protection	6.8 (sd = 13.4; P = 0.464)	-26.0 (sd = 16.0; P = 0.972)	-28.4 (sd = 28.9; P = 0.832)	-33.9 (sd = 75.2; P = 0.684)
Citizen science recurrent monitoring (5% error)	species loss	dynamic protection	6.6 (sd = 13.2; P = 0.476)	-24.9 (sd = 14.9; P = 0.972)	-27.3 (sd = 27.9; P = 0.852)	-29.6 (sd = 75.7; P = 0.652)
Citizen science recurrent monitoring (25% error)	species loss	dynamic protection	8.8 (sd = 13.3; P = 0.504)	-21.9 (sd = 16.1; P = 0.924)	-23.9 (sd = 28.4; P = 0.792)	-28.3 (sd = 80.2; P = 0.66)
Full initial monitoring	species loss	dynamic protection	-4.0 (sd = 13.6; P = 0.184)	-20.0 (sd = 15.6; P = 0.912)	-22.0 (sd = 28.4; P = 0.772)	-22.7 (sd = 79.3; P = 0.596)
Marxan	species loss	dynamic protection	10.3 (sd = 14.3; P = 0.568)	-13.4 (sd = 17.0; P = 0.752)	-15.1 (sd = 20.7; P = 0.756)	-10.8 (sd = 54.5; P = 0.628)
Greedy optimization	species loss	dynamic protection	28.4 (sd = 12.2; P = 0.972)	-23.0 (sd = 15.3; P = 0.956)	-24.8 (sd = 19.4; P = 0.944)	-23.0 (sd = 52.4; P = 0.78)
Full recurrent monitoring	value loss	dynamic protection	-0.8 (sd = 14.0; P = 0.256)	-10.9 (sd = 16.4; P = 0.736)	-12.1 (sd = 29.9; P = 0.648)	-59.1 (sd = 77.0; P = 0.832)
Full recurrent monitoring	area loss	dynamic protection	27.6 (sd = 11.4; P = 0.98)	13.6 (sd = 18.9; P = 0.212)	14.2 (sd = 32.6; P = 0.308)	4.5 (sd = 85.9; P = 0.448)
Full monitoring	species loss	all budget spent at once	1.3 (sd = 12.8; P = 0.336)	-26.7 (sd = 17.8; P = 0.968)	-29.2 (sd = 30.4; P = 0.84)	-29.9 (sd = 90.3; P = 0.656)
Marxan	species loss	all budget spent at once	14.0 (sd = 12.6; P = 0.72)	-21.2 (sd = 17.4; P = 0.9)	-23.3 (sd = 21.5; P = 0.888)	-27.1 (sd = 62.1; P = 0.7)
Greedy optimization	species loss	all budget spent at once	21.5 (sd = 11.8; P = 0.888)	-28.4 (sd = 15.3; P = 0.98)	-29.9 (sd = 21.0; P = 0.944)	-27.7 (sd = 59.0; P = 0.768)

**Supplementary Table 4. Summary of the output of CAPTAIN and Marxan policies applied to the endemic trees of Madagascar datasets.**

The budget was set to allow for a maximum of 2239 protected units (10% of the total), although note that number of protection units might be smaller (depending on their cost). In Marxan, the optimization might result in a higher number of protected units but the program settings we tuned to return on average the correct number of units. Values are averaged across 50 resampled datasets (see Methods).

	<b>CAPTAIN</b>	<b>Marxan (no BLM)</b>	<b>Marxan (BLM 0.1)</b>
% met target for all species	0.68	0.02	0
% missed the target in fewer than 10 species	0.94	1	0
% missed the target in fewer than 100 species	1	1	0
Percentage protected species	99.98	99.93	99.35
Median n. species not meeting the target	0 (range: 0–14)	1 (range: 0–3)	202 (range: 159–237)
Median n. species outside PUs	0 (range: 0–3)	0 (range: 0–1)	10 (range: 2–14)
Median fraction of protected sp. range	0.22	0.14	0.14
Median n. protected units (max: 2239)	1876 (range: 1616–2116)	1949 (range: 1891–2010)	2196 (range: 1863–3085)

## Supplementary References

1. ter Steege, H. et al. Hyperdominance in the Amazonian Tree Flora. *Science* 342, doi: 10.1126/science.1243092 (2013).
2. Stadler, T. Simulating trees with a fixed number of extant species. *Syst. Biol.* 60, 676684 (2011).
3. Watts, M., Klein, C. J., Tulloch, V. J. D., Carvalho, S. B. and Possingham, H. P. Software for prioritizing conservation actions based on probabilistic information. *Conserv. Biol.* 35, 12991308, doi: 10.1111/cobi.13681 (2021).
4. SerraSogas et al. Marxan User Manual. For Marxan version 2.43 and above (2021).
5. Dasgupta, P. *The Economics of Biodiversity: The Dasgupta Review*. London: HM Treasury (2021).
6. Brock, W. A. & Xepapadeas, A. Valuing Biodiversity from an Economic Perspective: A Unified Economic, Ecological, and Genetic Approach. *Am Econ Rev* 93, 1597-1614, doi:10.1257/000282803322655464 (2003).
7. Goeschl, T. & Swanson, T. The social value of biodiversity for R&D. *Environmental and Resource Economics* 22, 477-504 (2002).
8. Weitzman, M. L. The Noah's ark problem. *Econometrica*, 1279-1298 (1998).
9. Barbier, E. B. The concept of natural capital. *Oxford Review of Economic Policy* 35, 14-36, doi:10.1093/oxrep/gry028 (2019).
10. Dasgupta, P. *The Dasgupta Review –Independent Review on the Economics of Biodiversity*. Interim Report. (2020).
11. Simpson, R. D., Sedjo, R. A. & Reid, J. W. Valuing biodiversity for use in pharmaceutical research. *Journal of Political Economy* 104, 163-185 (1996).

12. Sterner, T. et al. Policy design for the Anthropocene. *Nature Sustainability* 2, 14-21, doi:10.1038/s41893-018-0194-x (2019).
13. Ando, A. W. & Mallory, M. L. Optimal portfolio design to reduce climate-related conservation uncertainty in the Prairie Pothole Region. *Proc. Natl. Acad. Sci. USA* 109, 6484-6489, doi:10.1073/pnas.1114653109 (2012).
14. Bateman, I. J. et al. Conserving tropical biodiversity via market forces and spatial targeting. *Proceedings of the National Academy of Sciences* 112, 7408-7413 (2015).
15. Bateman, I. J. et al. Bringing Ecosystem Services into Economic Decision-Making: Land Use in the United Kingdom. *Science* 341, 45-50, doi:10.1126/science.1234379 (2013).
16. Meyer, C., Weigelt, P. & Kreft, H. Multidimensional biases, gaps and uncertainties in global plant occurrence information. *Ecology Letters* 19, 992-1006, doi:10.1111/ele.12624 (2016).
17. Ritter, C. D. et al. The pitfalls of biodiversity proxies: Differences in richness patterns of birds, trees and understudied diversity across Amazonia. *Scientific reports* 9, 19205, doi:10.1038/s41598-019-55490-3 (2019).
18. Deyan, G. 61+ Revealing Smartphone Statistics For 2020. (2019). <https://techjury.net/>.
19. Wäldchen, J., Rzanny, M., Seeland, M. & Mäder, P. Automated plant species identification—Trends and future directions. *PLoS Comp. Biol.* 14, e1005993, doi: 10.1371/journal.pcbi.1005993 (2018).
20. Baena, S., Moat, J., Whaley, O. & Boyd, D. S. Identifying species from the air: UAVs and the very high resolution challenge for plant conservation. *PLOS ONE* 12, e0188714, doi:10.1371/journal.pone.0188714 (2017).