Optimising biodiversity protection through artificial intelligence

Alexandre Antonelli[1,2,3*§], Stefano Goria[4], Thomas Sterner[5*], Daniele Silvestro[2,6*§]

[1]Royal Botanic Gardens, Kew, TW9 3AE, United Kingdom
[2]Gothenburg Global Biodiversity Centre, Department of Biological and Environmental Sciences, University of Gothenburg, 405 30 Gothenburg, Sweden
[3]Department of Plant Sciences, University of Oxford, South Parks Road, Oxford, OX1 3RB, United Kingdom
[4]Thymia Limited, London, United Kingdom
[5]Department of Economics, University of Gothenburg, Gothenburg, Sweden
[6]Department of Biology, University of Fribourg, Switzerland

*Corresponding authors: a.antonelli@kew.org, thomas.sterner@economics.gu.se, daniele.silvestro@unifr.ch
[§]Equal contributions.

Abstract: Over a million species face extinction, carrying with them untold options for food, medicine, fibre, shelter, ecological resilience, aesthetic and cultural values. There is therefore an urgent need to design conservation policies that maximise the protection of biodiversity and its contributions to people, within the constraints of limited budgets. Here we propose a novel framework for spatial conservation prioritisation that combines simulation models, reinforcement learning and ground validation to identify optimal policies. Our model quantifies the trade-off between the costs and benefits of area and biodiversity protection, allowing the exploration of multiple biodiversity metrics. Under a fixed budget, our model protects substantially more species from extinction than the random or naively targeted protection of areas. We find that regular biodiversity monitoring, even if simple and with a degree of inaccuracy, substantially improves biodiversity outcomes. Given the complexity of people–nature interactions and wealth of associated data, artificial intelligence holds great promise for improving the conservation of biological and ecosystem values in a rapidly changing and resource-limited world.

Keywords: Artificial intelligence, sustainable development, systematic conservation planning.

Biodiversity is the variety of all life on Earth, from genes through to species, functions and ecosystems. It provides us with clean water, pollination, building materials, clothing, food and medicine, among many other physical and cultural contributions that species make to ecosystem services and people's lives[1,2]. The contradiction is that our endeavours to maximise short-term benefits are depleting biodiversity and threatening the life-sustaining foundations of modern societies in the long-run[3] (Box 1). This can help explain why, despite the risks, we are living in an age of mass extinction[4,5] and massive population declines of wild species[6]. One million species are currently at risk of disappearing[7], including nearly 40% (or c. 140,000) of all plant species[8,9], which could provide solutions to current and future challenges such as emerging diseases and climate-resilient crops[10-12]. The imperative to feed and house massively 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. Since not a single of the 20 Aichi biodiversity targets agreed by 196 nations for the period 2011–2020 has been fully met[13], there is now an urgent need to design more realistic and effective policies for a sustainable future[14].

There have been several theoretical and practical frameworks underlying biological conservation since the 1960s[15]. The field was initially focused on the conservation of nature for itself, without human interference, but gradually incorporated the bidirectional links to people – recognising our ubiquitous influence on nature, and the multi-faceted contributions we derive from it[1,6,15]. Throughout this progress, a critical step has been the identification of priority areas for targeted protection, restoration planning and impact avoidance – triggering the development of the field of spatial conservation prioritisation or systematic conservation planning[16-18].

Several tools and algorithms have been designed to facilitate systematic conservation planning[19]. They often allow the exploration and optimisation of trade-offs between variables, something not readily available in Geographic Information Systems[20], which can lead to substantial economic, social and environmental gains[21]. While the initial focus has been on maximising the protection of species while minimising costs, additional parameters can sometimes be modelled, such as species rarity and threat, total protected area and evolutionary diversity[22,23 19]. Among the two most widely used, Marxan[24] 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, while Zonation[19] produces a balanced ranking, from which arbitrary top or bottom priority areas can be selected. Despite their usefulness and popularity, these techniques are 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[25]). In addition, the optimized conservation planning does not explicitly incorporate climate change, variation in anthropogenic pressure, or species-specific sensitivities to such changes.

Here we tackle the challenge of optimising biodiversity protection in a complex and rapidly evolving world by harnessing the power of artificial intelligence (AI). We develop an entirely novel tool in systematic conservation planning[19] by creating a simulation system (Fig. 1) to explore multiple previously identified trade-offs in real-world conservation, and to evaluate the impact of data gathering on specific outcomes[26]. We also explore the impact of species-specific sensitivity to geographically varying local disturbances (e.g., as a consequence of new roads, mining, trawling or other forms of 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). Our framework is particularly applicable to sessile organisms, such as forest trees or reef corals, but implements mechanisms of offspring dispersal (e.g., seeds or larvae) that could be easily adaptable to many other systems.

Within AI, we implement a reinforcement learning (RL) framework based on our 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) and action (the effect of which is quantified by the outcome, also termed reward). It optimises 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 estimated within the RL framework to maximise the reward.

We combine these AI components with a simulation module that allows us to maximise various biodiversity metrics (including species richness and phylogenetic diversity) within a given budget constraint (Fig. 1). We investigate the importance of data gathering, focusing on

frequency (a single 'Initial', or multiple 'Recurrent' inputs of biodiversity data), level of detail ('Simple' species reports of species presence, or 'Full' monitoring of all species in targeted groups, and their population sizes) and accuracy (no errors in the recording, or allowing for a fraction of mistakes in species identifications). Our platform enables us to assess the influence of model assumptions on the reward, mimicking the use of counterfactual analyses[23].

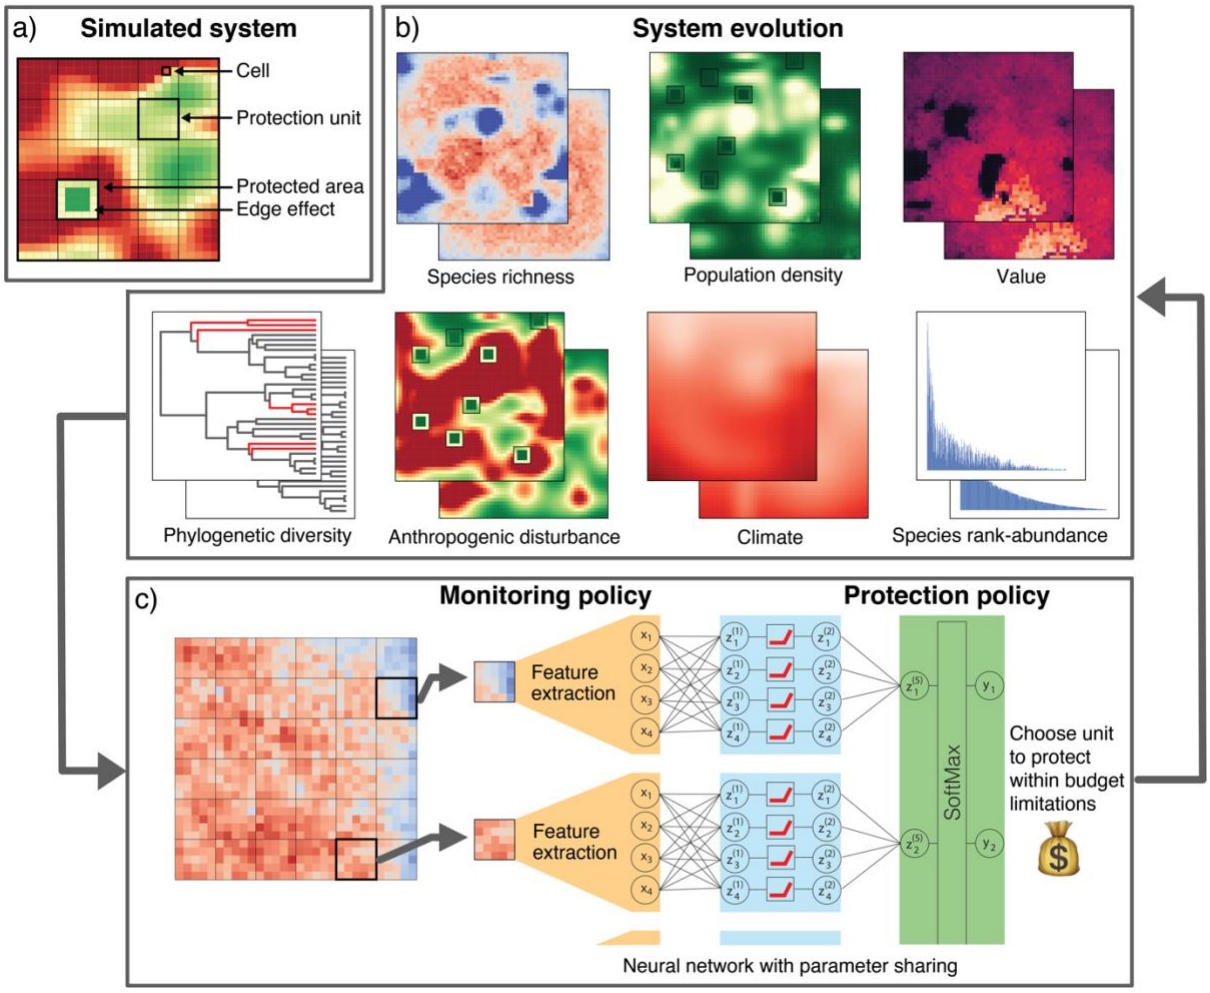

Fig. 1 | Simulation system and its evolution through time. a, The system – which could be the equivalent to a country state, an island or a large coral reef – consists of individual cells, each with a number of individuals of various species. Once a protection unit is identified and protected, its human-driven disturbance (e.g., forest loggings) will immediately reduce to an arbitrarily low level, except for the well-known edge effect[27] characterised by intermediate levels of disturbance. All 44 model parameters (see Table S1) are provided with initial default values but are fully customisable. **b-c,** Analysis flowchart integrating simulations and AI modules to maximise selected outcomes (e.g., species richness). **b,** System evolution between two points in time, in relation to six variables: species richness, population density, economic

value, phylogenetic diversity, anthropogenic disturbance, climate and species-rank abundance (see Animation S1 for a time-lapse video depicting these and additional attributes). Note that site protection does not protect species from climate change. **c,** Biodiversity features (species presence per cell at a minimum, plus their abundance under full monitoring schemes; see Methods and Box 2 for advances in data-gathering approaches) are extracted from the system at regular steps, which are then fed into a neural network that learns from the system's evolution to identify conservation policies that maximise a reward, such as protection of the maximum species diversity within a fixed budget.

We use our model to assess the following questions: i) What role does data gathering strategy play for effective conservation? ii) What trade-offs exist depending on the optimised variable, such as species richness, economic value or total area protected? iii) What can the simulation model teach us in terms of winners and losers – i.e., which traits characterise the species and areas protected over time? and iv) How does the size of protected areas influence biodiversity outcomes?

RESULTS AND DISCUSSION

Impact of data gathering strategy. We find 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 90% more species than a random protection policy (median values; Fig. 2a). A very similar outcome is generated by the Simple Recurrent Monitoring strategy (where only presence/absence of species are recorded in each cell) with a degree of error (Fig. 2b; see Methods). This similar outcome is noteworthy because the information it requires, although not currently available at high spatial resolution for most regions and taxonomic groups, could be obtained through modern technologies, such as remote sensing and environmental DNA, and for accessible locations also citizen science[28] in cost-efficient ways (Box 2).

The two monitoring strategies above outperform a Simple Initial Monitoring with no error (Fig. 2c), which protects an average of only 23% more species than a random policy. This is because the policy ignores the temporal dynamics caused by disturbances, population and range changes, and the costs of area protection – all of which are likely to change through time in real-world situations.

Our results differ from previous suggestions that area protection should not await prolonged data gathering[26]. This discrepancy can be explained by the nature of changes over time: while some systems may remain largely static over decades (e.g., tree species in old-growth forests), where a Simple Initial Monitoring could suffice, others may change drastically (e.g., 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 model.

Optimisation trade-offs. The policy objective, which determines the optimality criterion in our RL framework, significantly influences the outcome of the simulations. Minimizing species loss gives a median 37% gain in species, economic value and phylogenetic diversity (see Fig. S3) as compared to a random policy, and a 16% increase in protected area (Fig. 2a), despite an equal budgetary allocation. A policy maximizing species commercial values (such as timber price) tends to sacrifice more species in order to prioritise the protection of fewer, highly valuable ones. Under this scenario, species losses decrease by only 25% compared with the random baseline, while cumulative value losses decrease by about 50% (Fig. 2d).

A policy that maximises protected area results in a 50% increase in number of protected cells (selecting those cheapest to buy), but leads to substantial losses in species numbers, value and phylogenetic diversity (considerably worse than the random baseline; Fig. 2e). This is a key finding, given that total protected area has been at the core of previous international targets for biodiversity (such as Aichi), inadvertently supporting political pressure for 'residual' reservation[29,30] – the selection of new protected areas on land and at sea that are unsuitable for extractive activities (reducing costs and risks for conflicts), rather than key for biodiversity conservation.

Since the simulations are non-deterministic to better reflect natural stochasticity, their reliability in relation to species loss differ largely across simulations (Fig. 2f). Citizen Science Monitoring (Fig. 2b) is the most reliable policy – consistently yielding the same results, regardless of initial conditions of the system. This is probably because it depends on simpler data (presence/absence) with fewer confounding factors than the fuller data gathered for other simulations optimising the same policy (Fig. 2a,c). The citizen science policy is also more than three times more reliable than the Simple Initial Monitoring (Fig. 2c,f), because that scheme fails to adapt to changing conditions in the system over time (e.g., anthropogenic disturbance

and species richness). Since current methodologies for systematic conservation planning are static – relying on a similar initial data gathering as modelled here – this means their recommendations for area protection may be less reliable. As expected, the reliability for optimisations on economic value and total protected area (Fig. 2d,e) is even lower, meaning that those variables should not be used as surrogates for biodiversity protection unless economic values are amended to fully include ecosystem values.

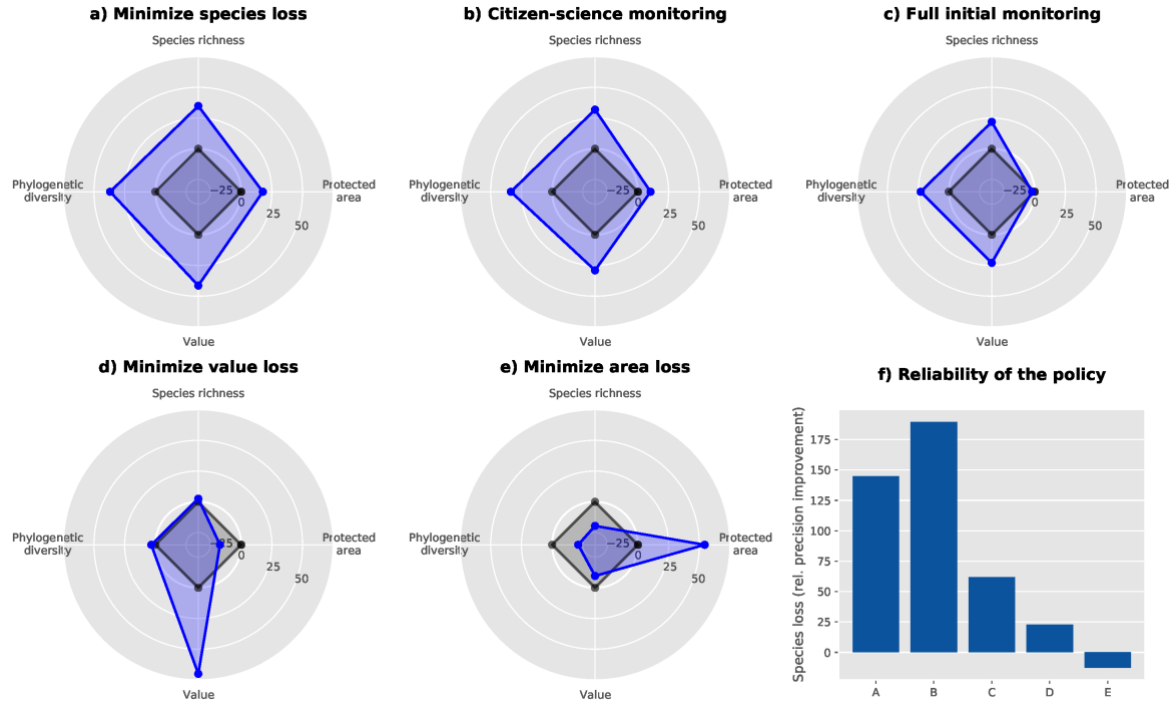

Fig. 2 | Improvement of policies optimised through reinforcement learning as compared with random area protection. In charts **a-e,** the blue polygons show the reduced loss of species, total protected area, accumulated species values and phylogenetic diversity between a random protection policy (grey shaded area) and AI-optimised simulations. **a-c,** Outcome of policies designed to minimise species loss based on different monitoring strategies: **a,** Full Recurrent Monitoring (of species presence and abundance at each time step); **b,** Citizen Science Monitoring (limited to species presence/absence with some error, at each time step); **c,** Full Initial Monitoring (species presence and abundance only at initial time). Plots **d-f** are based on Full Recurrent Monitoring, but using policies designed to (**d)** minimize value loss and (**e)** maximize the amount of protected area. **f,** reliability of the different policies (a-e) computed as the inverse of the variance in species loss across all simulations relative to that of a random policy (the dashed line at y = 0); high reliability values indicate high levels of consistency across simulations. Note that the model retained the full species data, from which dimension

indices were compiled for visualisation purposes. All results are averaged across 250 simulations. Each simulation was based on the same budget and resolution of the protection units (5x5) but differed in their initial natural system (species distributions, abundances, tolerances, phylogenetic relationships) and in the dynamics of the disturbance patterns.

Winners and losers. Focusing on the policy developed under Full Recurrent Monitoring and optimised 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 were characterised by relatively small initial range, small populations and intermediate or low resilience to disturbance (Fig. 3a). In contrast, species that survived had either low resilience but widespread ranges and high population sizes, or high resilience with small ranges and population sizes.

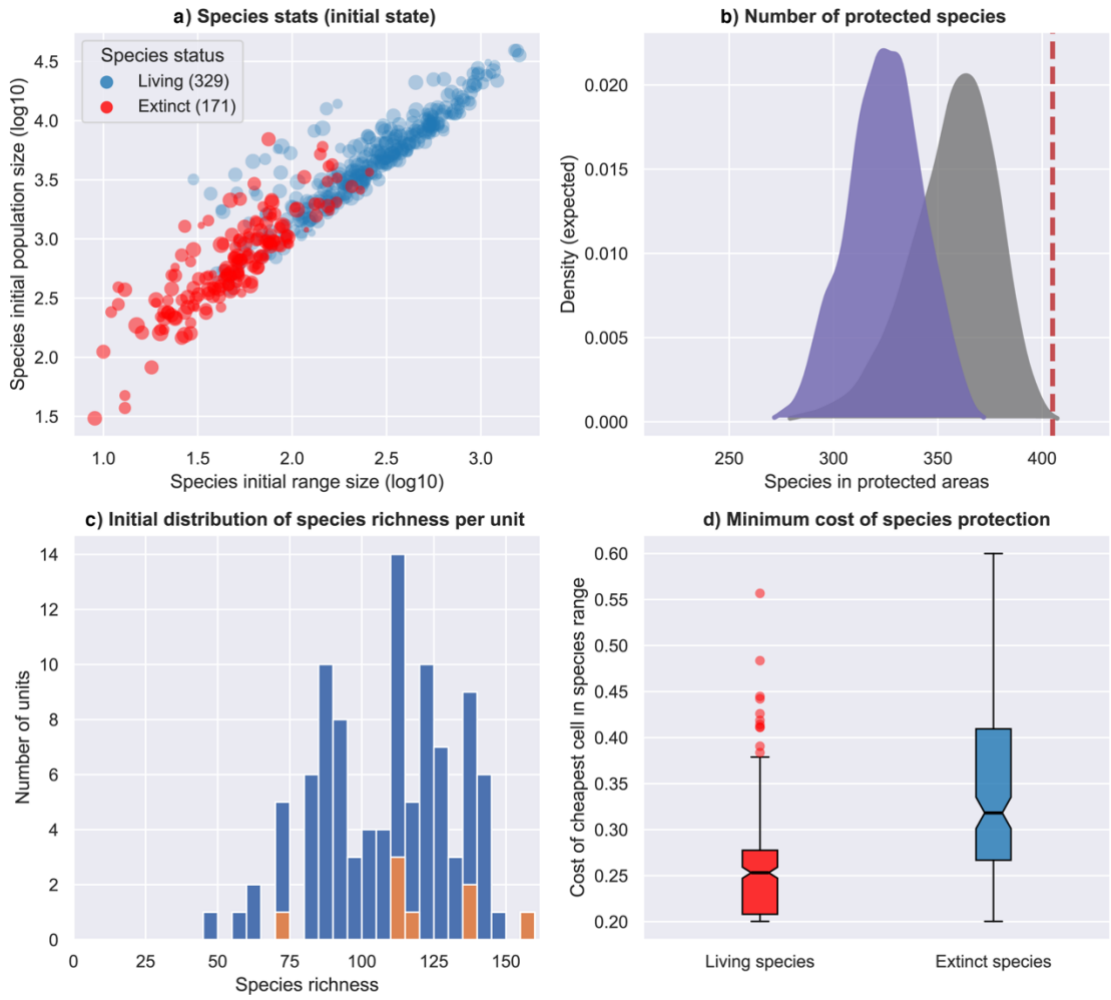

Fig. 3 | Summary statistics for one simulation optimised to reduce species loss and informed through Full Recurrent Monitoring. a, Living or surviving (blue circles) and

locally extinct (red circles) species after a simulation of 30 time steps with increasing disturbance and climate change. The X and Y axes show the initial range size and population size of the species (log10 transformed) and the size of the circles is proportional to the resilience of each species to anthropogenic disturbance and climate change, with smaller circles representing more sensitive species. **b,** Cumulative number of species encompassed in the eight protected units (5x5 cells) selected based on a policy optimised to minimise species loss. The grey density plot shows the expected distribution from 10,000 random draws, the purple shaded area shows the expected distribution when protected units are selected 'naively' (here, chosen among the top 20 most diverse ones), while the dashed red line indicates the number of species included in the units selected by the AI-optimised policy, which is higher than in 99.9% of the random draws. The optimised policy learned to maximise the total number of species included in protected units. Note that fewer species survived (329) in this simulation compared to how many were included in protected areas (405). This discrepancy is due to the effect of climate change, in which area protection does not play a role (Animation S1). Connectivity of protected areas could facilitate climate-induced range shifts. **c,** Species richness across the 100 protection units included in the area (blue), eight of which were selected to be protected (orange). The plot shows that the protection policy does not exclusively target units with the highest diversity. **d,** Distribution of the minimum cost of protection for living and extinct species. The plot shows the cost of the cell with the lowest cost of protection (as quantified at the last time step) within the initial range of each species.

We further assessed what characterised the grid cells that were selected for protection by the optimised policy. The cumulative number of species included in these cells was significantly higher than the cumulative species richness across a random set of cells of equal area (Fig. 3b). Thus, the model learns to protect a diversity of species assemblages to minimise species loss. Interestingly, the cells selected for protection did not include only areas with the highest species richness (Fig. 3c). This is important because areas with high richness constitute a 'naive' conservation target. Instead, protected cells spanned a range of areas with intermediate to high species richness, which reflects known differences among ecosystems or across environmental gradients and is more likely to increase protection complementarity[16]. Finally, extinct species tend to occur in more expensive regions within the system (Fig. 3d), corresponding to areas with high disturbance, which in turn might explain why they could not be effectively protected within the limited budget of the simulation.

If we instead examine which areas were protected under a policy maximizing total protected area, we find that they often include adjacent cells, for instance areas in a remote region, far from disturbance and therefore cheaper to buy and protect – typical of many current residual reserves[29,30], although less so for some newer conservation schemes[29].

Optimal size of protected areas. A long-standing debate in conservation biology centres on whether the optimal design is to protect "single large or several small" conservation areas – the so-called SLOSS debate[31]. Conservation effectiveness is related to both the size of protected areas and the dynamic pressures of the surrounding landscape[32]. To address this question, we estimated the biodiversity outcome in relation to protected areas constrained to sizes 2x1, 2x2, 2x5, 5x5 and 10x10 cells, with and without including climate change effects (Fig. 4 and Fig. S1a, respectively).

While in the short term (0-10 time steps) smaller protection units preserve more species, these scenarios lead to greater species loss in the longer term, whereas intermediate-sized areas are more resilient (in our simulations, the 5x5-cell area outperforms all others; Fig. 4a). Species losses are more linear and predictable for larger units, while smaller units show a rapid decline in response to anthropogenic disturbance and climate (Fig. 4b). The simulations without climate change show qualitatively similar results for the smaller areas, whereas for the larger sizes (5x5 and 10x10) they underestimate the rate and magnitude of species losses (Fig. S2a).

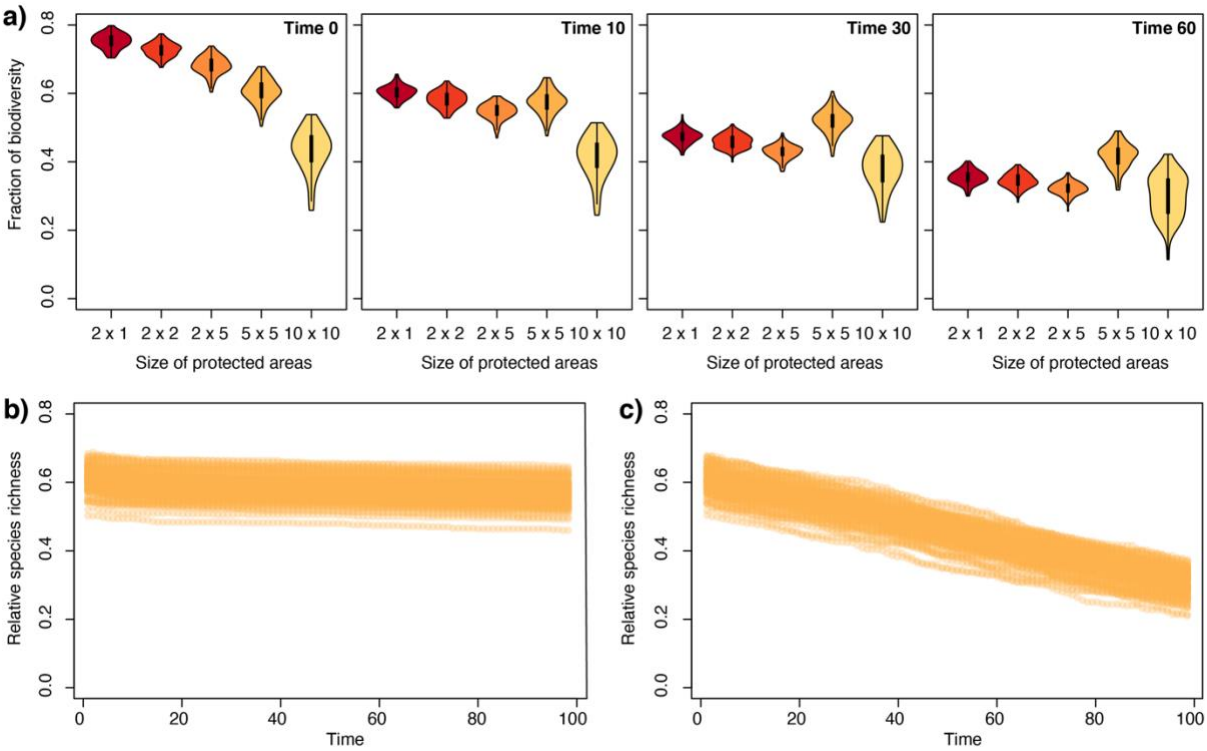

Fig. 4 | Effects of size of protected units on biodiversity outcomes. a, Violin plots summarising the fraction of surviving species at times steps 0, 10, 30, 60 across 100 simulations with climate change and full anthropogenetic disturbance in all non-protected cells. **b-c,** Rate and magnitude of species loss through time based on protection units of 5x5 cells in the absence of climate change (**b**) and with climate change (**c;** 2.5°C increase per 100 years). All simulations were performed under a random protection policy to establish a base line across 5 settings, varying the size of the protection units (redder colours indicate smaller protection units). They were based on the same budget, therefore protecting the same amount of area (8% of the total), with all protected areas established in the first-time step. The protected areas included a buffer zone (1-cell thick) where the level of protection is 50%.

Conclusions

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 will help safeguard natural assets and their contributions for the future. Choosing suboptimal areas, by contrast, will lead to the loss of species, phylogenetic diversity, socio-economic value and ecological functions. As public funding for nature conservation is scarce and needs to compete with many other important environmental issues such as climate change or the spread of toxins, there is little room for mistakes.

The framework presented here provides a spatially explicit way to explore the complex decisions facing policymakers and conservationists today, focusing on data gathering, placement and size of protected areas, and policy objective – all within the context of a rapidly changing world affected by local anthropogenic disturbances and large-scale climate change. Our results highlight the importance of regularly gathering basic biodiversity information, and demonstrate substantial trade-offs in biodiversity outcome depending on the policy being optimised – such as maximizing the protection of species or areas.

We recognise that few real-world systems are available for a direct application of our theoretical models and simulations. However, an increasing number of high-resolution biological datasets are being compiled and standardised (e.g.,[33]), which in addition to the use

of new and cost-effective monitoring technologies (Box 2) could provide empirical data for further exploration of our approach, and added complexity. The flexibility of our model, offered by its nearly 40 variables (see Methods), could be easily expanded and adapted to almost any empirical dataset. Future implementations could model additional variables, such as functional diversity and other biodiversity and socio-economic metrics.

Artificial intelligence techniques should not replace human judgment, and ultimately investment decisions will be based on more than just the simple parameters implemented in our models, including careful consideration of people's manifold interactions with nature[1,15]. More sophisticated measures of value could model this in future work. It is also crucial to recognise the importance of ensuring the right conditions required for effective conservation of protected areas in the long term[34,35]. However, it is now time to acknowledge that the sheer complexity of biological systems, multiplied by the increasing disturbances in a changing world, cannot be fully grasped by the human mind. As we are at the dawn of what many are calling the most decisive decade for biodiversity[6,36], 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 developed a simulation framework to model biodiversity loss and optimize conservation policies using a reinforcement learning algorithm. We implemented a spatially explicit individual-based simulation to assess future biodiversity changes based on natural processes of mortality, replacement, and dispersal and on anthropogenic processes such as habitat modifications, selective removal of a species, rapid climate change, as well as 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.

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 modeled 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. Along with a representation of the natural and anthropogenic evolution of the system, our simulation framework includes an agent (i.e., 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 sizes for

each species, their temporal trends and geographic ranges. 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 a reinforcement learning algorithm to optimize how to perform the protect actions based on the information provided by monitoring, such that it minimizes species loss or other criteria depending on the policy (see below).

We use this framework to: 1) assess the effect of different monitoring policies on the protection of species, phylogenetic diversity, and other measures of biodiversity; 2) explore trade-offs when attempting to minimize losses in biodiversity, economic value, and fraction of disturbed area; and 3) assess the effects of modulating number and size of protected areas on short- and long-term biodiversity loss.

We provide a full description of the simulation system in the Supplementary Methods. In the sections below we present our optimization methods and algorithms and describe the experiments carried out within this framework.

Optimization of conservation policy using reinforcement learning

We use a reinforcement learning framework to optimize a conservation policy under a pre-defined criterion (e.g., aiming to minimize the loss of various biodiversity metrics). The 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 time frame evaluated in the simulation. If the value is set equal across all species, the reinforcement learning algorithm will minimize overall species extinctions. However, different definitions of value can be used to minimize loss based on evolutionary distinctiveness of species (e.g., 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 and through time.

The granularity of monitoring and protection 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), which 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, which includes species abundances and geographic distribution,

$$S_t = \{\mathbf{H_t}, \mathbf{D_t}, \mathbf{F_t}, \mathbf{T_t}, \mathbf{C_t}, \mathbf{P_t}, B_t\} \tag{1}$$

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 (Table S2) depends on the monitoring policy (π_X), which is decided *a priori* in the simulation. A predefined monitoring policy also determines the temporal frequency of this action throughout the simulation, e.g. only at the first time step or repeated at each time step. The features extracted for each unit represent the input upon which a protect action can take place, if the budget allows for it, following a protection policy π_Y.

Monitoring action

The monitoring action extracts features from the simulated system within each protection unit. These features (listed in Table S2) include the number of species which 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 (π_Y).

We do not assume species-specific sensitivities to disturbance (d_s, f_s) to be a known feature, since 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 relative abundances of species (x_3 in Tables S2). The features tested across different policies are specified in the subsection *Experiments* (see below).

Protect action

The protect 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, while 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 compared to the centre, to mimic the generally negative edge effects in protected areas (e.g., higher vulnerability to extreme weather). While 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 the state of neighboring areas. Thus, protecting an area that has already undergone biodiversity loss may not result in the restoration of its original biodiversity levels.

The protect 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 thorough time. Alternatively, it can be defined as a function of the current level of anthropogenic disturbance in the cell. The cost of each protection 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

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, but it is also influenced by a set of discrete actions determined by the protection policy π_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 to all possible *protect* actions. While calibrating our 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 architecture described below. The parameters of the neural network are optimized such that they maximize the expected reward resulting from the *protect* actions. We optimize the model using a genetic strategies algorithm[45] incorporating aspects of classical Policy Gradient methods such as an Advantage function[46].

In our simulations we fix π_X, thus pre-defining the frequency of monitoring (e.g. at each time step or only at the first time step) and the amount of information produced by $X(S_t)$ and we optimize π_Y, which determines how to best use the available budget to maximize the reward. Each action has cost, defined by the function $\texttt{Cost}(A, S_t)$, which here we set to a constant for monitoring and equal across all monitoring policies. The cost of the *protect* action 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, i.e. if $\texttt{Cost}(Y, S_t) < B_t$.

The protection policy is parametrized as a feed forward neural network with 2 hidden layers using a ReLU activation function and an output layer using a softMax function. The input of the neural network is a matrix 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. The first hidden layer $h^{(1)}$ is a matrix $J \times L_1$

$$h^{(1)}_{nl_1} = g\left(\sum_f x_{nf} W^{(1)}_{fl_1}\right) \tag{2}$$

where

$$g(x) = \max(0, x) \tag{3}$$

is the ReLU activation function and the matrix $W^{(1)}$ is the matrix of coefficients we are

optimizing. The second hidden layer takes the $h^{(1)}$ as input and gives a N vector as output:

$$h_n^{(2)} = \sigma\left(\sum_{l_1} h_{nl_1}^{(1)} W_{l_1}^{(2)}\right) \tag{4}$$

With σ a softMax function

$$\sigma(x)_i = \frac{\exp(x_i)}{\sum_j \exp(x_j)} \tag{5}$$

We interpret the output vector of N variables as the probability of protecting the cell n.

This architecture implements parameter sharing across all protection units when connecting the input nodes to the first 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 protect policy, $\hat{\pi}_P$, 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_6 (an indicator variable for protected units) to randomly select the proposed unit for protection.

The optimization algorithm is a variation of the Evolution Strategies algorithm[45] , which optimizes a policy without the need of computing gradients and allowing for easy parallelization. In particular one epoch of our algorithm works as

for every worker p **do**
  $\epsilon_p \leftarrow \mathcal{N}(0, \sigma)$ with diagonal covariance and dimension $W + M$
  for $t = 1, ..., T$ **do**
    $R_t \leftarrow R_{t-1} + r_t(\theta + \epsilon_p)$
  end for
end for
$R \leftarrow$ average of R_T across workers

$R_e \leftarrow \alpha R + (1-\alpha) R_{e-1}$

for every coefficient θ in $W + M$ **do**

$\theta \leftarrow \theta + \lambda A(R_e, R_T, \epsilon)$

end for

where R is the cumulative reward over T time steps, $\lambda = 0.1$ 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_e on every worker p weighted by the corresponding noise ϵ_p:

$$A(R_e, R_T, \epsilon) = \frac{1}{P} \sum_p (R_e - R_T^p)\epsilon_p. \tag{6}$$

Experiments

We used our simulation framework to explore the properties of our model and the effect of different policies. 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 tested the effect of varying the size of the protection units on species loss in a context of disturbance and or climate change.

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[47] , totaling 2.5 million individuals (Fig. S3a). Species ranges were generated based on a random diffusion process, with the number of individuals in a cell constrained by the carrying capacity.

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 ($f_s \sim \mathcal{B}[0.2, 0.7]$) (Fig. S3c,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[48] 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 harbor more valuable species for timber then 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 (Fig. S3b).

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 modeled a dynamic system in which an initial area is subject to a level of disturbance that intensifies and expands spatially over time (Fig. S5). 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 5×5 neighboring 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, species with a declining population size, and budget and cost of each unit (features x_2, x_3, x_5, x_6 in Table S2). The protection policy included one meta-parameter (defining the threshold for rare species; Table S2) and a neural network with 4×3 and 3 parameters for the first and second layers, respectively.

- *Simple 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. The protection policy included a neural network with 3×3 and 3 parameters.

- *One-time 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 4×4 and 4 parameters.

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.

We summarized the outcome of the simulations by measuring at the last time step the biodiversity 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. We also computed the reliability of the policy as the inverse of the variance (i.e., precision) in biodiversity loss across all simulations relative to that of a random policy. Thus a higher reliability value indicates that the policy provides a lower variance in biodiversity loss across stochastically different systems and simulations. We interpret high reliability of a policy

as a measure of how well it adapts to different systems.

Testing the trade-offs among biodiversity, 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 5×5 and 5 parameters.

- *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 2×2 and 2 parameters.

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

Testing the effect of the size of protection units

We ran additional simulations to assess the impact of the size and number of protection units on short- and long-term biodiversity loss in the face of 1) high anthropogenic disturbance and 2) high anthropogenic disturbance coupled with climate change. Rather than testing policies optimized under different settings, we established baseline expectations under a random policy where the only difference among sets of simulations was the size and number of protection units. Disturbance was set to increase from 0 to 1 uniformly across the system in the first 4 steps and remain constant afterward, while at the third step all protected units allowed by the

budget were randomly placed in the system. After that, the simulations were run until reaching 100 steps. We tested protection units of five sizes: 2×1, 2×2, 2×5, 5×5, 10×10 and specified a budget such that a total of 200 cells were protected in all simulations. Thus the number of individual protected units varied with size: 100, 50, 20, 8, 2, respectively. The protected units were assumed to include a 1-cell thick buffer zone along their perimeter, in which the protection level was set to 0.5 to mimic the commonly recognised edge effect of protected areas, whereas internal cells had full protection (equal to 1). We ran 100 analyses under each setting and compared the levels of biodiversity loss at different time steps to assess the effects on short, medium and long terms.

We repeated these simulations with the addition of climate change, which was modeled as a stochastic process with an average temperature increase of 0.025 degrees per time step and random spatially heterogeneous anomalies varying at each step (Fig. S6). Under these settings species in protected units are still exposed to climate-driven mortality affecting more strongly species with more limited climatic tolerances and species with restricted spatial range, which makes them more sensitive to local climate anomalies.

Implementation and code availability

The simulation framework and algorithms described here were implemented in Python v.3. All the codes and scripts will be made available along with all simulations as an open access repository upon acceptance of the paper.

Box 1: Socio-economic valuation of biodiversity

The field of economics sees nature as an asset, analogous to diversity in portfolio management[3,37]. It provides valuable insurance that may prove vital[38]. Risks are reduced if two assets are negatively correlated. Weitzman[39] incorporates the idea of redundancy when species share genes, and shows that optimal policy under a budget constraint will involve strict prioritisation of some species.

The value of ecosystem services hinges on ease of substitution by man-made capital. In some respects, man 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[40].

We see massive destruction 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[41]. Still, private incentives to preserve biodiversity (e.g. research and development in biodiversity-based solutions) remain weak[42], and strong policies are needed[14].

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 model 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[43]. Spatial targeting of sites that increase biodiversity protected per dollar is cost effective, and may also include recreational and other values[21,44] that we hope to develop more in future model versions.

<End of Box 1>

Box 2: New techniques for biodiversity monitoring

Our simulations show that the regular monitoring of species is critical for prioritising areas for conservation under changing anthropogenic and climate pressure (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[45]. New approaches are now speeding up and popularising 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 analysed. This approach reveals the biotic composition of whole communities, including plants, animals, fungi and bacteria. Standardised sampling also allows direct comparison of taxonomic and phylogenetic diversity across multiple sites[46], 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[47], creating unseen possibilities of linking millions of people to nature through species logging (**b**), using platforms such as iNaturalist[48] which increasingly use image-recognition software and geolocations for automated species identification. In our simulations, even identification errors of up to 20% 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 (Fig. 3c). [Photo credits: a-b: A.A.; c: Justin Moat].

Finally, remote sensing technologies can rapidly scan and characterise relatively large areas. Using 3-D sensing technology, they now hold the promise of automated identification of species using multi-spectrum images[49]. 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.

<End of Box 2>

Acknowledgements

We thank Kerstin Johannesson for the initial discussions that sparked this project; Ben Groom, Atte Moilanen, Bob Pressey and Ian Bateman for invaluable feedback on an earlier version of this manuscript; Justin Moat for input on remote sensing technologies; Rhian Smith for editorial support; and many colleagues at our respective research groups for discussions. AA is funded by the Swedish Foundation for Strategic Research, the Swedish Research Council and the Royal Botanic Gardens, Kew. TS and AA acknowledge funding from the Strategic Research Area Biodiversity and Ecosystem Services in a Changing Climate, BECC, funded by the Swedish government. DS received funding from the Swiss National Science Foundation (PCEFP3_187012) and from 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.

REFERENCES

[Inserted using reference manager; they will be carefully formatted upon final submission. These references also include the Methods section[50-53]]

1 Díaz, S. *et al.* Assessing nature's contributions to people. *Science* **359**, 270-272, doi:10.1126/science.aap8826 (2018).

2 Chaplin-Kramer, R. *et al.* Global modeling of nature's contributions to people. *Science* **366**, 255-258, doi:10.1126/science.aaw3372 (2019).

3 Dasgupta, P. The Economics of Biodiversity: The Dasgupta Review. *London: HM Treasury* (2021).

4 Barnosky, A. D. *et al.* Has the Earth's sixth mass extinction already arrived? *Nature* **471**, 51-57, doi:10.1038/nature09678 (2011).

5 Andermann, T., Faurby, S., Turvey, S. T., Antonelli, A. & Silvestro, D. The past and future human impact on mammalian diversity. *Science Advances* **6**, eabb2313, doi:10.1126/sciadv.abb2313 (2020).

6 Leclère, D. *et al.* Bending the curve of terrestrial biodiversity needs an integrated strategy. *Nature* **585**, 551-556, doi:10.1038/s41586-020-2705-y (2020).

7 IPBES. IPBES Global Assessment on Biodiversity and Ecosystem Services. (2019).

8 Nic Lughadha, E. *et al.* Extinction risk and threats to plants and fungi. *PLANTS, PEOPLE, PLANET* **2**, 389-408, doi:10.1002/ppp3.10146 (2020).

9 Antonelli, A. *et al.* State of the World's Plants and Fungi 2020. (Royal Botanic Gardens, Kew, 2020).

10 Howes, M.-J. R. *et al.* Molecules from nature: Reconciling biodiversity conservation and global healthcare imperatives for sustainable use of medicinal plants and fungi. *PLANTS, PEOPLE, PLANET* **2**, 463-481, doi:10.1002/ppp3.10138 (2020).

11 Kersey, P. J. *et al.* Selecting for useful properties of plants and fungi – Novel approaches, opportunities, and challenges. *PLANTS, PEOPLE, PLANET* **2**, 409-420, doi:10.1002/ppp3.10136 (2020).

12 Seddon, N. *et al.* Understanding the value and limits of nature-based solutions to climate change and other global challenges. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **375**, 20190120, doi:doi:10.1098/rstb.2019.0120 (2020).

13 Diversity, S. o. t. C. o. B. Global Biodiversity Outlook 5 – Summary for Policy Makers. (2020).

14 Sterner, T. *et al.* Policy design for the Anthropocene. *Nature Sustainability* **2**, 14-21, doi:10.1038/s41893-018-0194-x (2019).

15 Mace, G. M. Whose conservation? *Science* **345**, 1558-1560, doi:10.1126/science.1254704 (2014).

16 Moilanen, A. Generalized Complementarity and Mapping of the Concepts of Systematic Conservation Planning. *Conservation biology : the journal of the Society for Conservation Biology* **22**, 1655-1658, doi:10.1111/j.1523-1739.2008.01043.x (2008).

17 Margules, C. & Sarkar, S. *Systematic conservation planning*. (Cambridge University Press, 2007).

18 Margules, C. R. & Pressey, R. L. Systematic conservation planning. *Nature* **405**, 243-253, doi:10.1038/35012251 (2000).

19 Moilanen, A., Wilson, K. & Possingham, H. *Spatial conservation prioritization: quantitative methods and computational tools*. (Oxford University Press, 2009).

20 Honeck, E., Sanguet, A., Schlaepfer, M. A., Wyler, N. & Lehmann, A. Methods for identifying green infrastructure. *SN Applied Sciences* **2**, 1916, doi:10.1007/s42452-020-03575-4 (2020).

21 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).

22 Carvalho, S. *et al.* Spatial conservation prioritization of biodiversity spanning the evolutionary continuum. *Nature Ecology & Evolution* **1**, 0151, doi:10.1038/s41559-017-0151 (2017).

23 Sacre, E., Weeks, R., Bode, M. & Pressey, R. L. The relative conservation impact of strategies that prioritize biodiversity representation, threats, and protection costs. *Conservation Science and Practice* **2**, e221, doi:https://doi.org/10.1111/csp2.221 (2020).

24 Ball, I. R., Possingham, H. P. & Watts, M. Marxan and relatives: software for spatial conservation prioritisation. *Spatial conservation prioritisation: Quantitative methods and computational tools*, 185-195 (2009).

25 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, doi:https://doi.org/10.1016/j.biocon.2013.06.025 (2013).

26 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, doi:https://doi.org/10.1111/j.1461-0248.2009.01287.x (2009).

27 Murcia, C. Edge effects in fragmented forests: implications for conservation. *Trends Ecol. Evol.* **10**, 58-62 (1995).

28 Kosmala, M., Wiggins, A., Swanson, A. & Simmons, B. Assessing data quality in citizen science. *Front. Ecol. Environ.* **14**, 551-560, doi:https://doi.org/10.1002/fee.1436 (2016).

29 Devillers, R. *et al.* Reinventing residual reserves in the sea: are we favouring ease of establishment over need for protection? *Aquat. Conserv.: Mar. Freshwat. Ecosyst.* **25**, 480-504, doi:https://doi.org/10.1002/aqc.2445 (2015).

30 Vieira, R. R. S., Pressey, R. L. & Loyola, R. The residual nature of protected areas in Brazil. *Biol. Conserv.* **233**, 152-161, doi:https://doi.org/10.1016/j.biocon.2019.02.010 (2019).

31 Järvinen, O. Conservation of endangered plant populations: single large or several small reserves? *Oikos*, 301-307 (1982).

32 Saunders, D. A., Hobbs, R. J. & Margules, C. R. Biological consequences of ecosystem fragmentation: a review. *Conserv. Biol.* **5**, 18-32 (1991).

33 Bruelheide, H. *et al.* sPlot – A new tool for global vegetation analyses. *Journal of Vegetation Science* **30**, 161-186, doi:10.1111/jvs.12710 (2019).

34 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, doi:10.1073/pnas.1908221116 (2019).

35 Geldmann, J. *et al.* A global analysis of management capacity and ecological outcomes in terrestrial protected areas. *Conservation Letters* **11**, e12434, doi:https://doi.org/10.1111/conl.12434 (2018).

36 Díaz, S. *et al.* Set ambitious goals for biodiversity and sustainability. *Science* **370**, 411-413, doi:10.1126/science.abe1530 (2020).

37 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).

38 Goeschl, T. & Swanson, T. The social value of biodiversity for R&D. *Environmental and Resource Economics* **22**, 477-504 (2002).

39 Weitzman, M. L. The Noah's ark problem. *Econometrica*, 1279-1298 (1998).

40 Barbier, E. B. The concept of natural capital. *Oxford Review of Economic Policy* **35**, 14-36, doi:10.1093/oxrep/gry028 (2019).

41 Dasgupta, P. The Dasgupta Review –Independent Review on the Economics of Biodiversity. Interim Report. (2020).

42 Simpson, R. D., Sedjo, R. A. & Reid, J. W. Valuing biodiversity for use in pharmaceutical research. *Journal of Political Economy* **104**, 163-185 (1996).

43 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).

44 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).

45 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).

46 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).

47 Deyan, G. 61+ Revealing Smartphone Statistics For 2020. (2019). <https://techjury.net/stats-about/smartphone-usage/#gref>.

48 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).

49 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).

50 Salimans, T., Ho, J., Chen, X., Sidor, S. & Sutskever, I. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864* (2017).

51 Sutton, R. S., McAllester, D. A., Singh, S. P. & Mansour, Y. in *NIPs.* 1057-1063 (Citeseer).

52 ter Steege, H. *et al.* Hyperdominance in the Amazonian Tree Flora. *Science* **342**, doi:10.1126/science.1243092 (2013).

53 Stadler, T. Simulating trees with a fixed number of extant species. *Syst. Biol.* **60**, 676-684 (2011).

Supplementary information

Optimising biodiversity protection through artificial intelligence

Alexandre Antonelli, Stefano Goria, Thomas Sterner, Daniele Silvestro

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 ϕ. 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 $H_t^c = \{H_t^{1,c}, \ldots, H_t^{N,c}\}$ and as $\mathbf{H}_t$ 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 Table S1.

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 λ), 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

$$n_{sij} = H_t^{s,i} \times g_s \times \exp\left(-\frac{1}{\lambda_s} d_{ij}\right) \tag{7}$$

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:

$$n_j = \frac{K_j - \sum_s^N H_t^{s,j}}{K_j} \left(\sum_i^Z \sum_s^N n_{sij}\right) \tag{8}$$

Under this formulation and with the simulations settings utilized 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 modeled 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:

$$\delta_t^{s,c} = \max\left[-\log(1-d_s)\left(\frac{H_t^c}{(1-D_t^c)K_c} - 1\right), 0\right] \tag{9}$$

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 $\mathbf{F}_t$ 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 $f_s F_c$ 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 τ_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 ΔT_{sc} the absolute distance between the mean annual temperature in the cell T_c and the species temperature range τ_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 τ_s, species with a wider tolerance are subject to lower death rates than species with narrower tolerance.

Supplementary Tables

S1. 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
$\mathbf{H}_t$	Matrix with number of individuals for each species across each cell

Parameter	Explanation
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)
λ^s	Dispersal rate
ϕ	Phylogenetic tree connecting all species in the system
$\mathbf{D}_t$	Disturbance matrix
D_t^c	Disturbance at cell c
d_s	Sensitivity to disturbance of species s
$\mathbf{F}_t$	Selective disturbance matrix
F_t^c	Selective disturbance at cell c
f_s	Sensitivity to selective disturbance of species s
$\mathbf{T}_t$	Climate matrix
T_t^c	Climate in cell c
τ_s	Climatic tolerance of species s
$\mathbf{P}_t$	Protection matrix
P_t^c	Protection level of cell c
B_t	Available budget
$\mathbf{C}_t$	Matrix of cost of protection
C_t^c	Cost of protection of cell c
S_t	State of the system. This includes $\{\mathbf{H}_t, \mathbf{D}_t, \mathbf{F}_t, \mathbf{T}_t, \mathbf{C}_t, \mathbf{P}_t, B_t\}$
π_X	Monitoring policy setting the frequency of monitoring and defining the function $X(S_t)$
π_Y	Protect policy
U	Number of protection units
$X(S_t)$	Function extracting features from last monitored state of the system
$\{x_1, x_2, \ldots x_n\}$	List of features extracted for each protection unit
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

Parameter	Explanation
r_t	Reward in one time step t. This depends on the policy, e.g. $r_t = -1$ if the policy aims to minimize biodiversity loss and one species went extinct at time t.

S2. 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	Number of non-protected species with a relative decrease in population size. A species is considered to have undergone a relative decrease in population size if $$\frac{E_t^s}{E_{t-1}^s} - 1 < 0 \text{ where } E_t^s = \frac{Q_t^s}{\sum_{i=1}^{N} Q_t^i}$$ is the population size of species s at time t relative to the total number of individuals across all species. Time $t-1$ indicates the time of the previous monitoring. We note that to generate this feature the feature extraction function requires information from a previous state of the system $X(S_t, S_{t-1})$.
x_4	Available budget (B_t)
x_5	Available budget minus current 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 figures

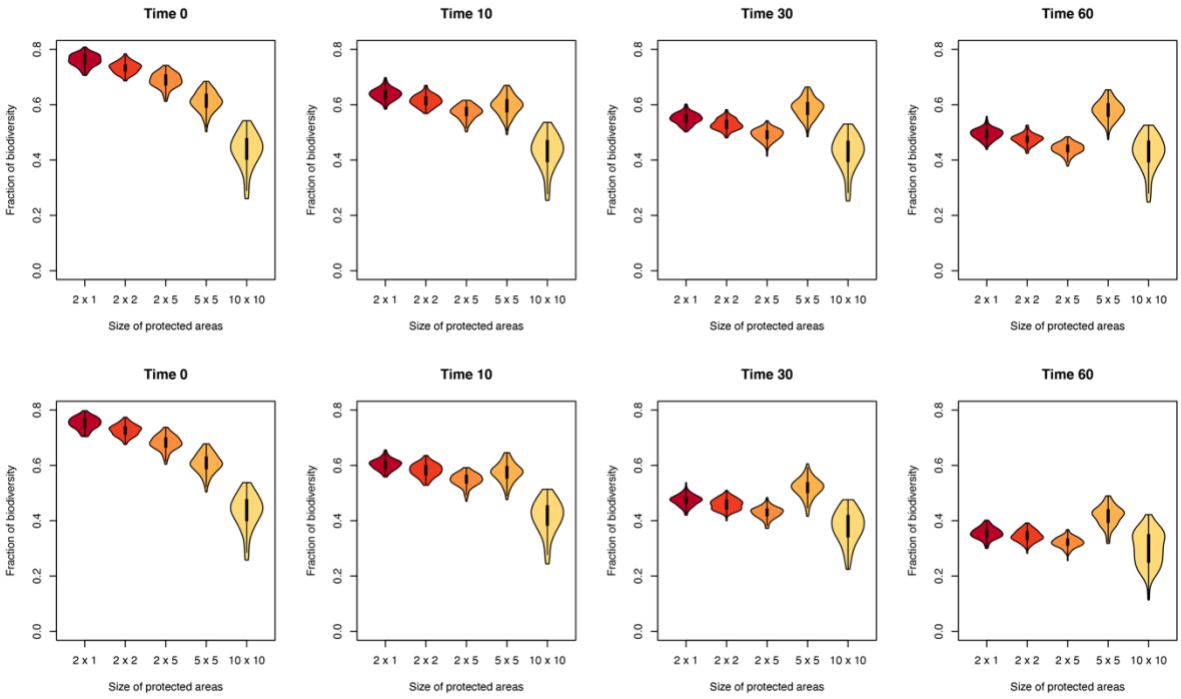

Fig. S1 | Impact of climate change on species losses in relation to size of protected unit. The violin plots show the fraction of surviving species at times steps 0, 10, 30, 60 across 100 simulations. Upper plots: only anthropogenic disturbance is included in the model. Lower plots: climate change is added (2.5°C increase per 100 years; as in Fig. 5a). All simulations were performed under a random protection policy to establish a baseline across 5 settings, varying the size of the protection units (redder colours indicate smaller protection units). They were based on the same initial budget, therefore protecting the same amount of area (8% of the total), with all protected areas established at the first-time step. The protected areas included a buffer zone (1-cell thick) where the level of protection is 50%.

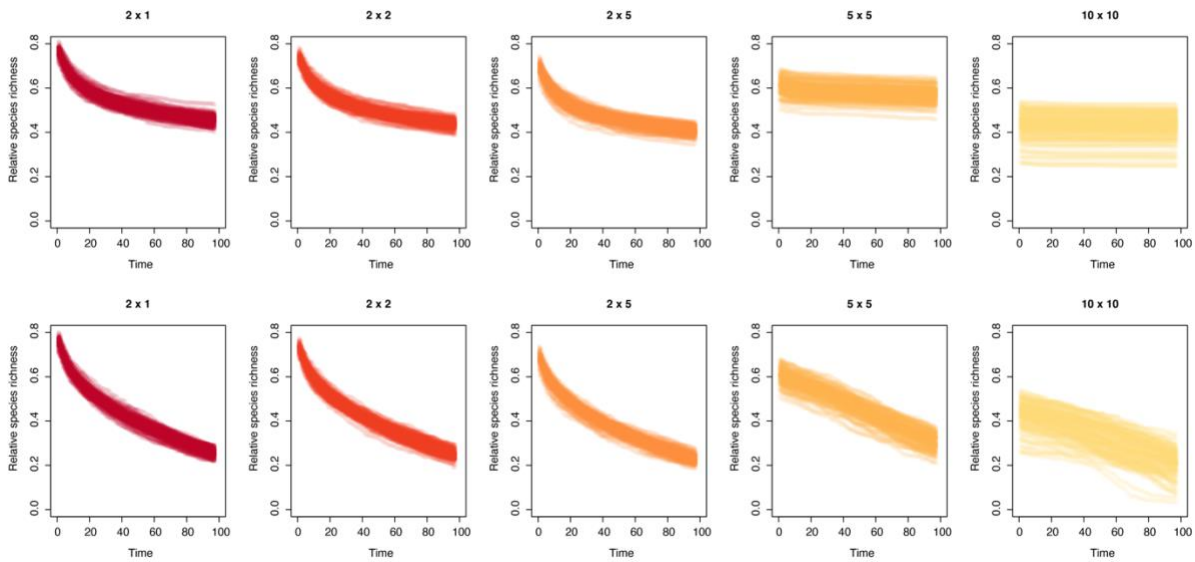

Fig. S2 | Impact of climate change on the rate and magnitude of species losses in relation to size of protected unit. Upper plots: only anthropogenic disturbance is included in the model. Lower plots: climate change is added (2.5°C increase per 100 years; as in Fig. 5b). All simulations as in Fig. 5.

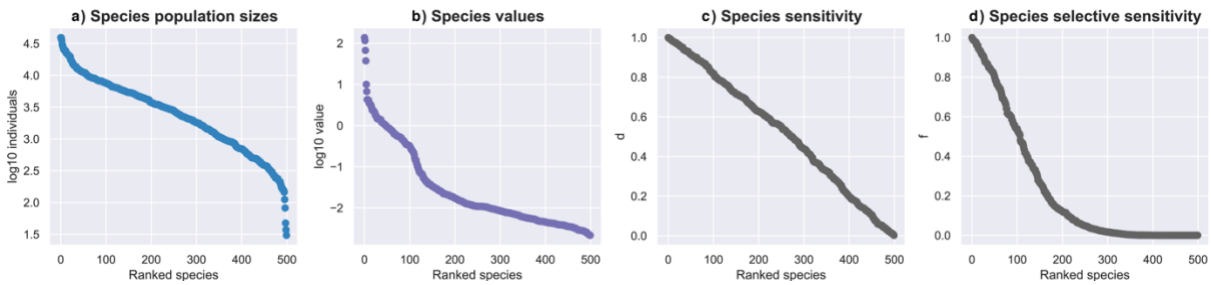

Fig. S3 | Value distribution across four key species-specific model parameters. For a randomly selected simulation, the graphs show **a,** initial species population sizes, **b,** species economic values, **c,** species sensitivity to anthropogenic disturbance and **d,** species selective sensitivity (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 log10 transformed.

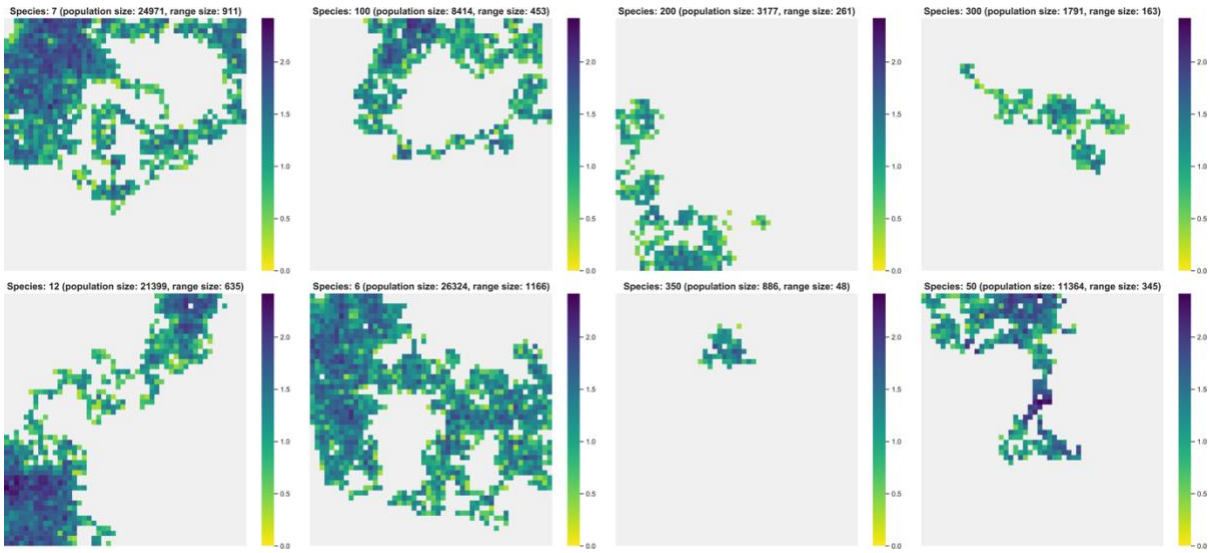

Fig. S4 | Examples of species ranges. For eight randomly selected species, the plots show their total distribution (non-grey cells; summed to a value of range size) as well as total population sizes (i.e., the number of individuals in the system, or abundance). Cells are coloured based on the log10-transformed number of individuals.

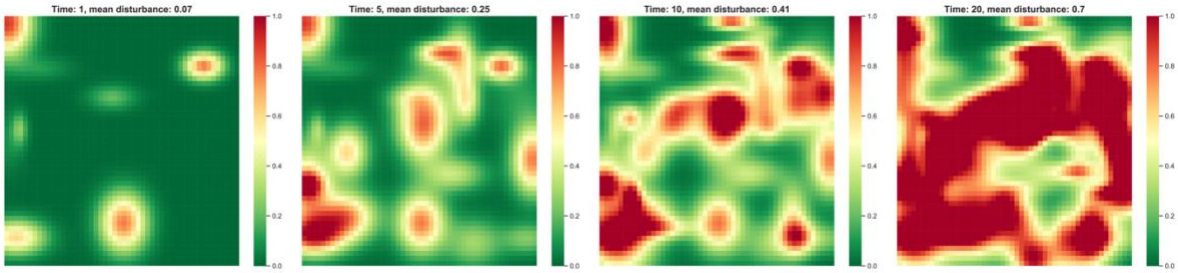

Fig. S5 | Evolution of anthropogenic disturbance across the simulated system. For a randomly selected simulation, the colours show changes in human disturbance across individual cells at time steps 1, 5, 10, and 20 in a system where no protected areas were set. The mean overall disturbance increases 10-fold from 0.07 to 0.7.

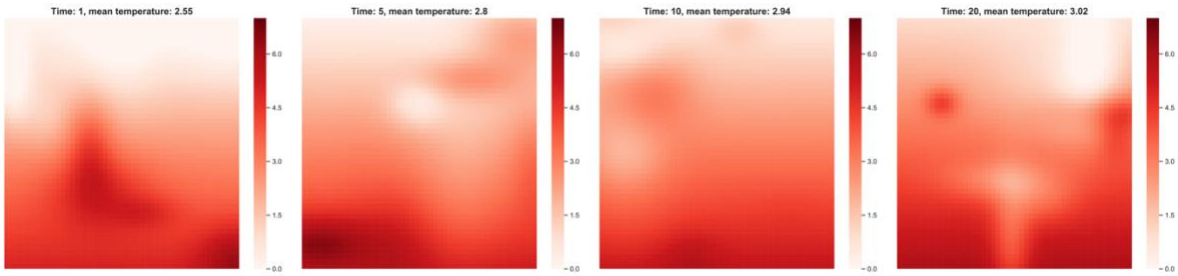

Fig. S6 | Evolution of climate change across the simulated system. For a randomly selected simulation, the colours show temperature changes across individual cells at time steps 1, 5, 10, and 20 resulting in an average increase across the entire system of 0.47 degrees with random local deviations at each step, and the overall increase in mean temperature across the entire system.

Multimedia

Animation S1 | Time-lapse animation depicting the evolution of the simulation system over 30 time steps (such as years in real-life situations). Each plot depicts spatial changes through time for one model variable, departing from a randomly chosen simulation optimised to maximise species richness. The variables shown are: species richness (initial state: 500 species), mean population density (initial state: 1000 individuals per cell), total population size (initial state: 2.5 million individuals), phylogenetic diversity, anthropogenic disturbance, mean annual temperature and economic loss. The system evolves under increasing anthropogenic disturbance and climate change (global warming plus local yearly anomalies) and with eight protection units of 5 by 5 cells, established in the first 10 time steps based on our reinforcement-learning optimised policy. To facilitate comparison, the trend curves show the quantitative changes over time for economic value, phylogenetic diversity and species richness, expressed as percentages of the initial values. The bottom plots show range maps and abundance for 4 different species. While all of them decline in abundance by the end of the simulation, the first species actually experiences an initial small increase in population size from ca. 26,000 to 27,000, likely linked to the species' high tolerance toward low levels of disturbance, before declining to 17,000 individuals on step 30. The third species only survives within protected

areas at the end, while the fourth species goes extinct at time step 18, due to climate change, even though it was included in a protected area. Available at https://vimeo.com/512631818 (password: Biodiversity*2021).