Using artificial intelligence to optimize ecological restoration for climate and biodiversity

DANIELE SILVESTRO[1,2,3,4], STEFANO GORIA[3,5], BEN GROOM[6,7], PIOTR JACOBSSON[8]
THOMAS STERNER[3,9], ALEXANDRE ANTONELLI[2,3,10,11,12]

[1] *Department of Biosystems Science and Engineering, ETH Zurich, Basel, Switzerland*
[2] *Gothenburg Global Biodiversity Centre, Department of Biological and Environmental Sciences, University of Gothenburg, Gothenburg, Sweden*
[3] *Captain Technologies Ltd, London, UK*
[4] *Swiss Institute of Bioinformatics, Basel, Switzerland*
[5] *Thymia Ltd, London, UK*
[6] *LEEP Institute, Department of Economics, University of Exeter Business School, Exeter University, Exeter, UK*
[7] *Grantham Research Institute on Climate Change and the Environment, London School of Economics, London, UK*
[8] *Department of Biology, University of Fribourg, Fribourg, Switzerland*
[9] *Department of Economics, University of Gothenburg, Gothenburg, Sweden*
[10] *Royal Botanic Gardens, Kew, Richmond, UK*
[11] *Antonelli Foundation for Biodiversity Research and Conservation, Nova Friburgo, Brazil*
[12] *Department of Biology, University of Oxford, Oxford, UK*
Correspondence: daniele[at]captaintech.ai, a.antonelli[at]kew.org

Abstract

The restoration of degraded ecosystems can provide critical contributions to help mitigate climate change and bend the curve of biodiversity loss. Depending on the primary objective – such as maximizing carbon storage or protecting threatened species – and within the boundaries of budget constraints, different spatial priorities have been identified at global and regional levels. Funding mechanisms to support such work comprise public sources, philanthropy, and the private sector, including the sales of carbon and biodiversity credits. However, effectively exploring tradeoffs between restoration objectives and estimating the price of biodiversity and carbon credits to design financially viable projects remain challenging. Here we harness the power of artificial intelligence in our software CAPTAIN, which we further develop to identify spatial priorities for ecological restoration that maximize multiple objectives at once and which allows a robust evaluation of biodiversity and climate outcomes. We find through a series of realistic simulations that even low-to-moderate consideration of biodiversity in restoration projects leads to the selection of restored areas that substantially improve the conservation of threatened species, although this can result in a decrease in total carbon captured. We propose a data-driven valuation of biodiversity credits in relation to carbon credits, enabling the design of a blended financial model that could support restoration efforts even in areas previously excluded for economic reasons. This study shows how the use of a robust methodological framework can lead to significant improvements in outcomes for climate and nature, while minimizing costs.

Keywords: *multi-objective reinforcement learning, Global Biodiversity Framework, Paris Agreement, Conservation Area Prioritization Through Artificial INtelligence (CAPTAIN)*

Introduction

Ecological restoration is one of the most powerful nature-based solutions to tackle the twin planetary crises of biodiversity loss and climate change (Seddon et al., 2020). This is why the United Nations has declared 2021-2030 the Decade on Ecosystem Restoration and it motivates the commitment under Target 2 of the Kunming-Montreal Global Biodiversity Framework – ratified by 196 countries – that "*by 2030 at least 30 per cent of areas of degraded terrestrial, inland water, and marine and coastal ecosystems are under effective restoration*". In addition to protecting natural ecosystems, the preferred option whenever there is an option, well-designed and implemented ecological restoration can capture and store carbon, increase populations of threatened species, improve ecosystem services and many other contributions to people and the natural environment (e.g., Di Sacco et al., 2021).

Despite the multiple benefits of restoration and its urgency in the face of global environmental change, two major challenges hamper its implementation. Firstly, much like prioritizing areas for protection (Jung et al., 2021), there are clear trade-offs when designing restoration projects aiming to maximize carbon capture or biodiversity recovery (Strassburg et al., 2020). For example, the planting of exotic trees across Africa, while growing fast and capturing large amounts of carbon (in principle helping to mitigate climate change, if the biomass produced is retained long-term), has inadvertently led to detrimental effects on native biodiversity (Parr et al., 2024). Trade-off effects have been widely discussed in the literature (e.g. Holl and Brancalion, 2020; Hua et al., 2022), but to our knowledge no tools have yet been developed to allow their quantification while taking into account the complexity of the real world: including the spatial heterogeneity in biodiversity, carbon storage potential, costs, and their changes over time.

Secondly, there is a major funding gap – estimated between USD 600 billion and over 800 billion annually (Deutz et al., 2020) – to implement the Global Biodiversity Framework, including its restoration target. Given that public funding from institutions such as the Global Environment Facility, which provides funding of roughly USD 1.3 billion per year, falls far short of reaching the required levels, many restoration projects today are supported instead by the sales of carbon credits, which can be voluntary or mandatory depending on the country. Alas, the simplicity of carbon credits, where one credit equals one tonne of atmospheric carbon captured, regardless of location, is also its pitfall. As market forces push trade towards low prices, this incentivizes restoration projects that minimize costs (such as reforestation on marginal and highly degraded land) and focus exclusively on maximum carbon recovery, even at the cost of biodiversity benefits. To improve this situation, carbon credits are being designed to

include additional benefits such as biodiversity, and biodiversity credits are emerging following strict principles for high integrity and governance (Antonelli et al., 2024). However, we are not aware of attempts to quantify the cost relationship between carbon and biodiversity credits to make them financially interchangeable (allowing, for instance, a restoration project to become viable by selling biodiversity credits rather than carbon credits), nor the economic opportunities for bundling carbon and biodiversity credits for the same site (while avoiding double claims) making it viable to restore ecosystems in expensive areas, which have so far been excluded due to insufficient profit margins.

Here, we propose an approach based on multi-objective reinforcement learning to address these shortcomings. Specifically, we i) demonstrate through simulations that this approach can help spatially guide restoration programs while monitoring their progress over time; ii) explore trade-offs and synergies that exist between restoration efforts focused on carbon, biodiversity and costs, through multi-objective optimization to identify socially optimal solutions; and iii) link our predictions to the emerging concept of biodiversity credits alongside carbon credits to design financially viable restoration programs. Our results show that identifying trade-offs between carbon and biodiversity is crucial for appropriately valuing biodiversity credits to incentivize local communities to develop restoration programs that maximize benefits for both climate and biodiversity.

Results

A reinforcement learning framework for restoration

We developed an approach to optimize restoration planning building upon the reinforcement learning software CAPTAIN (Conservation Area Prioritization Through Artificial Intelligence) (Silvestro et al., 2022). The implementation includes a spatially explicit dynamic *environment* characterized by climatic features, species and their geographic ranges, anthropogenic disturbances, and restoration costs. The disturbance could be any kind of anthropogenic activity including settlements, roads, farming or logging on land, fishing or mining in the sea, while costs represent direct restoration costs and opportunity costs, e.g. if a restored area can no longer be used for farming.

The CAPTAIN framework also includes an *agent* (i.e. a manager or social planner) who monitors the environment and implements a restoration program pursuing one or multiple predefined restoration targets. We used a deep neural network to map the monitored environment to the selected areas for restoration. The environment evolves over time through a number of steps (e.g. years, here set to 30) at which multiple processes take place, including dispersal and mortality among species, changes to climate or disturbances, and the establishment of restoration areas selected by the agent, where the disturbance is reduced.

We implemented a multi-objective reinforcement learning algorithm to evaluate the ability of this approach to optimize restoration toward multiple goals, specifically targeting the net carbon sequestration and reduction in species extinction risks, while incorporating budget constraints. Here we measured extinction risks in a 5-tiered classification inspired by the International Union for Conservation of Nature (IUCN; Red List of Threatened Species, 2023) and ranging from Least Concern to Critically Endangered. The model is highly flexible and can be tuned to optimize other parameters, such as biodiversity metrics based on richness, abundance and function. We optimized five models, hereafter referred to as *policies* (Sutton et al., 1998), using differently weighted rewards that ranged from one focusing on carbon (Policy 1) to one targeting biodiversity outcomes (Policy 5), with three models exploring intermediate solutions (Table 2). All policies shared the same weight on costs, thus implying equal budget constraints. We then compared the predicted restoration outcomes of different policies, in terms of selected areas, resulting net carbon sequestration, and impact on species extinction risks. This allowed us to evaluate a gradient of restoration policies ranging from climate- to biodiversity-optimal.

Optimization of restoration policies and trade-offs

Our realistic simulations of biodiversity used to optimize restoration policies and predict their outcomes revealed that the five policies resulted in substantial spatial differences in the prioritization of restoration areas and led to widely different carbon and biodiversity outcomes. While model training and simulations were based on the same arbitrarily chosen geographic region (Figs. S1), all datasets differed in the distribution of simulated species, threats, and costs. Thus, the placement of the selected restored areas varied among them. However, at the level of each simulation, the policies clearly prioritized different areas (Fig. 1).

Policy 1 favored the restoration of areas that resulted in the highest net carbon gain (24% on average after 30 times steps), as shown for one simulated dataset in Fig. 2a. In contrast, Policy 5 prioritized areas that contributed the most to reducing species extinction risk, therefore considering threats, species richness, and complementarity of species composition (Fig. 1). The result was a more modest carbon gain (just above 11% on average; Fig. 2a) but a substantially stronger reduction of the extinction risk compared with Policy 1. For instance, the number of Critically Endangered species at the end of the simulation was 36% lower under Policy 5 compared to Policy 1 (Fig. 2b). Similarly, the number of Endangered species decreased on average by 73% when giving priority to biodiversity over carbon. These experiments also showed that while carbon accumulated through time following almost linear trajectories, changes in species extinction risks followed more complex pathways (Fig. 2b-f), reflecting the fact that species moved across risk categories and that their conservation status was determined by their ability to re-colonize restored areas based on their dispersal ability and connectivity and disturbance among areas.

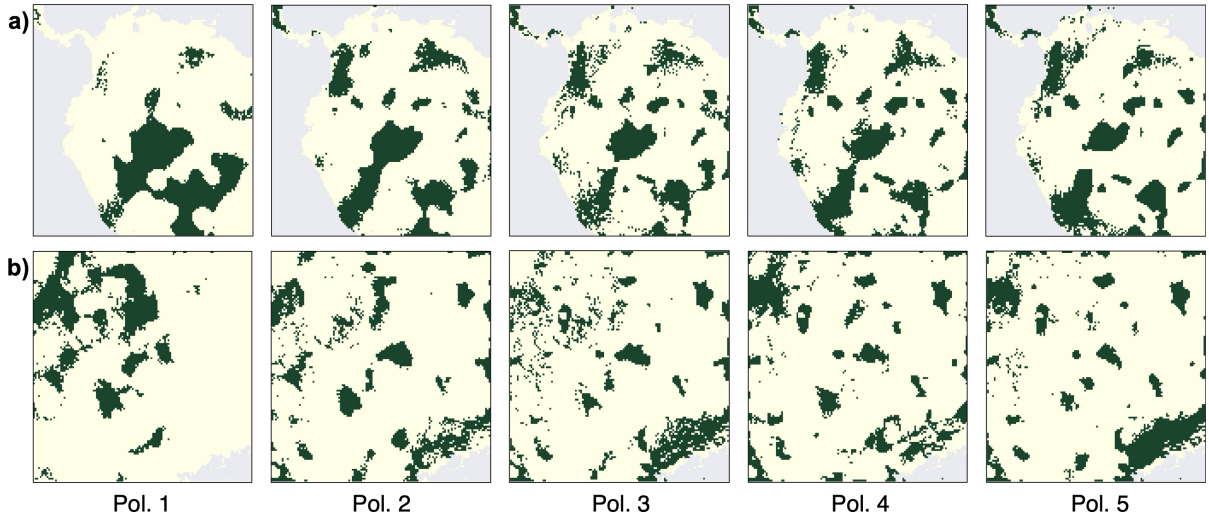

Figure 1: Changes in spatial restoration priorities depending on the relative weights of carbon and biodiversity. a) Example of the restored areas in one of the 100 simulations performed on the same region used to optimize the models (Fig. S1) as selected under our five policies (Table 2). Note that the placement of the restored areas depends on the specific conditions of the simulated dataset. Restored areas will therefore vary among simulations, reflecting the randomly initialized distribution of species, disturbance, and costs (the outcomes of this specific simulation in terms of resulting net carbon gain and changes in species extinction risks are shown in Fig. 2). b) Inferred restoration priorities obtained from simulations based on another, randomly selected region, showing similar discrepancies in spatial prioritization depending on the relative importance attributed to carbon vs biodiversity outcomes. (Fig. S2).

The distributions of outcomes at the end of the 30 time-step restoration show a substantial level of consistency across 100 simulations (Fig. S3). The net carbon gain in Policy 1, which was trained without considering a biodiversity reward, resulted in an average 2.2-fold increase (standard deviation among simulations: ±0.26) in carbon storage compared to Policy 5, which did not include carbon in the reward (Fig. 3a). Policies with mixed rewards revealed a gradient of outcomes with net carbon gain gradually decreasing when setting the relative carbon reward to lower values (Figs. 2b, 3a). In contrast, the biodiversity outcomes revealed drastic improvements as soon as a non-zero weight was assigned to the biodiversity reward with Policies 2–5 all resulting in similar numbers of species with reduced extinction risk 3b), while Policy 1 resulted in significantly worse impacts on extinction risks compared to most other policies (Figs. 3, S3b). For instance, while on average 30% of the species decreased their extinction risk under Policy 1 (150 out of the 500 simulated species), the percentage raised to 71–76% (354–378 species) under all other policies (Fig. S3b). The fraction of Critically Endangered species remained virtually unchanged under Policy 1, while it decreased on average by 29% under Policy 5, and by 27% under all other policies (Fig. S3d). When combining all threatened species (CR, EN, VU), these decreased by about 25% under Policy 1 (i.e. 68 species) but dropped by 62–69% (i.e. 170–190 species) units under the others (Fig. S3i), thus showing a 2.6-fold improvement. These results show that even when biodiversity outcomes are given a smaller weight compared to carbon in

the optimization of the restoration policy (e.g., Policy 2) the impact on extinction risk reduction is large.

Finally, the costs of implementation of the different policies, i.e. the cumulative cost of all restored cells, were relatively similar, with changes between 4 and 17% among policies (Fig. 3c). This reflected the fact that all policies were run under the same budget constraints.

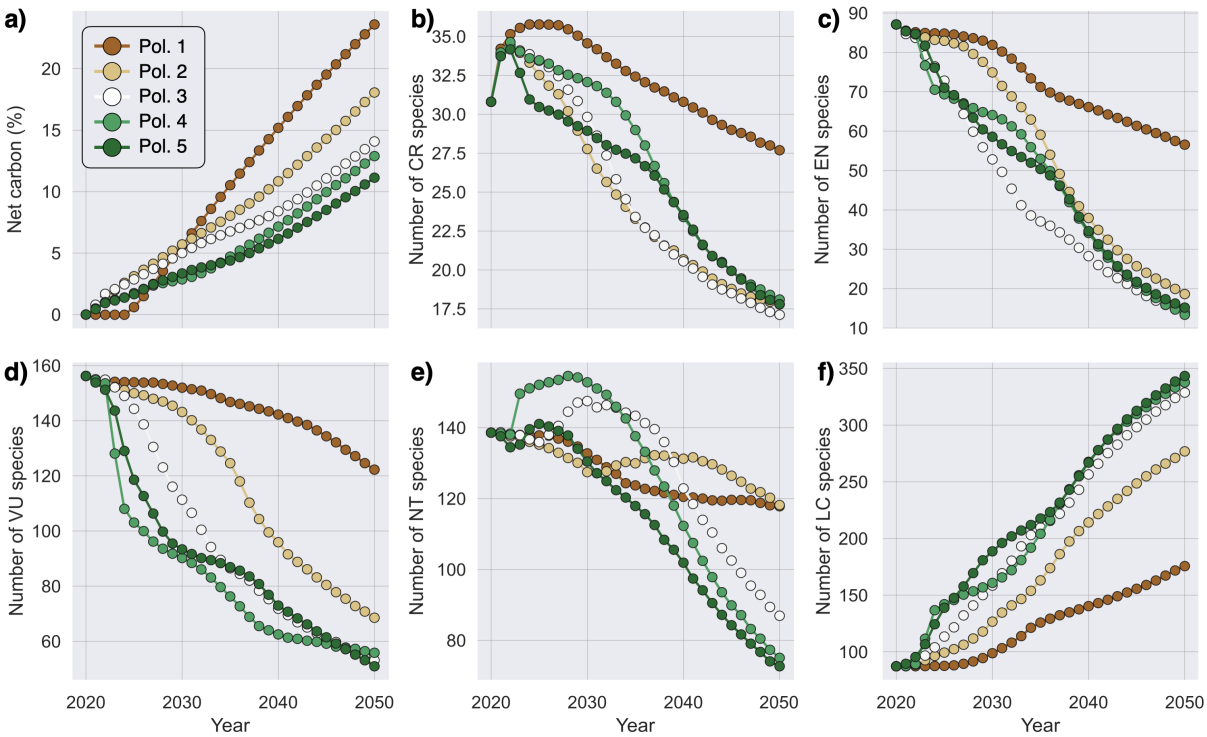

Figure 2: Accumulation of carbon and evolution of extinction risks across 30 time-steps averaged over the 100 simulations. The results are shown for Policies 1–5, which differ in the relative weight assigned to carbon sequestration (highest in Policy 1) and reduction of extinction risks (highest in Policy 5). The plots show that when biodiversity is factored in, the carbon accumulation curve reaches lower levels decreasing from 24% to 11% (a). This trend is counterbalanced by a substantial improvement of species conservation status, with a reduction of Critically Endangered and Endangered species of 36 and 73%, respectively, from Policy 1 to 5 (b, c). Meanwhile, the number of Least Concern species increased by 51% (f), highlighting the substantial impact of biodiversity-targeted restoration in our simulations. The time labels (years) used here are arbitrary.

Carbon credits and valuation of biodiversity credits

As expected, different policies led to vastly different carbon and biodiversity outcomes, while maintaining similar implementation costs. To address the question of whether they are equally feasible in financial terms, we first explored their feasibility in a context in which carbon credits were used to cover the restoration costs. To this end, we used an arbitrary valuation of carbon credits such that the cumulative value of carbon credits accumulated with Policy 1 matched exactly its restoration costs (Eq. 14; Fig. 4a), Under this definition of carbon credits, all other

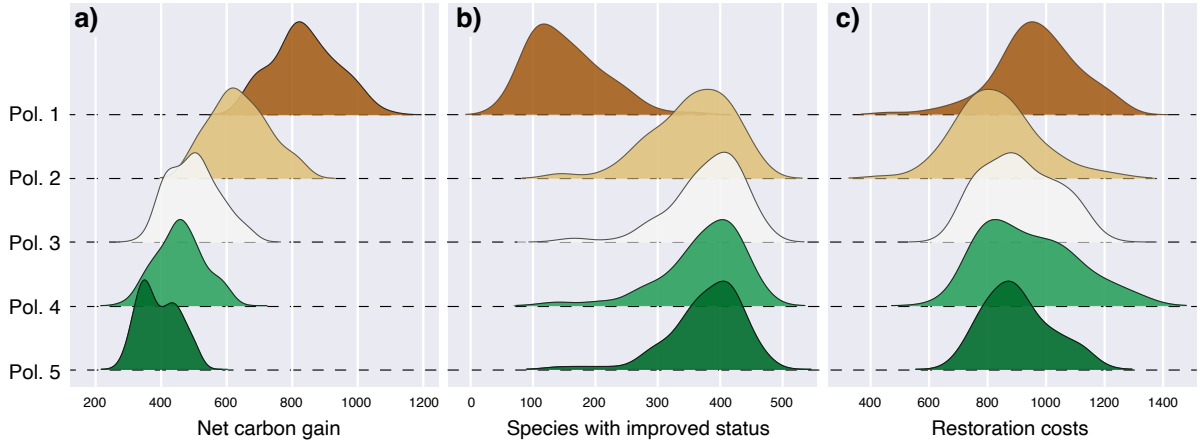

Figure 3: Summary statistics of key outcomes from 100 simulations run under five trained restoration models, Policies 1–5. The net carbon (a) represents the net gain in amount of carbon (here expressed in arbitrary units – see Methods) in the system after 30 time-steps during which 3,000 cells were restored. As the weight assigned to the carbon reward was lowered from 0.5 to 0 in favor of the biodiversity reward, the net carbon gain gradually decreased, although all restoration policies resulted in a net gain. In contrast, the number of species with improved conservation status, i.e. lower extinction risk, drastically increased as soon as a non-zero weight was assigned to the biodiversity reward (Policies 2–5; b). The implementation cost was similar across all policies (c), reflecting the fact that all models were trained with the same budget constraint.

policies resulted in a net financial loss that ranged from -10% on average in Policy 2 to about -50% in Policy 5 (Fig. 4a).

The results changed substantially when adding the contribution of biodiversity credits to the value of carbon credits. This required defining the unit for a biodiversity credit, a concept around which there is significant ongoing discussion and currently no consensus. Here we adopted a simple definition, for demonstration, where one credit corresponds to one species moving from a threatened category before restoration to a not-threatened one at the end of the simulation (see *Methods)*. Without implying that this should be preferred over alternative definitions, we used this measure to evaluate the potential effects of biodiversity credits on the financial viability of different restoration policies.

Next, we needed to define the value of a biodiversity credit. We therefore explored different valuations relative to a carbon credit. When setting the value of one biodiversity credit equal to one carbon credit (i.e. $\beta = \chi$ based on the notation used in the *Methods*), both Policy 1 and 2 became profitable with financial gains above 10% for the latter. Under this valuation, the implementation of Policies 3–5 continued to lead to significant financial loss (Fig. 4b). Increasing the relative value of biodiversity credits ($\beta = 2\chi$ and $\beta = 4\chi$) led to improved financial outcomes of all policies (Fig. 4c,d). Yet, we found that Policy 2 consistently led to the best outcomes, reflecting the non-linear improvement in biodiversity outcomes observed in Fig. S3b. These experiments showed that biodiversity credits may be used to level out financial

barriers impeding the implementation of restoration policies that would not be otherwise viable with carbon credits alone.

We also derived the relative value of a biodiversity credit such that a biodiversity-aware policy (i.e. Policies 2–5) reached a net monetary profit matching that of Policy 1 (Eq. 17). We found that policies 3–5 could reach financial gains equal to Policy 1 if the value of a biodiversity credit was set to about $\beta = 2.6 - -3.5\chi$, while for Policy 2, a value of $\beta = 0.8\chi$ was sufficient (Table S1). This means that if a biodiversity credit is valued at 1.2 times the value of a carbon credit, then Policies 1 and 2 result in the same net relative profit based on the cumulative carbon and biodiversity credits. These results indicate that the value of biodiversity credits converge to a relatively constrained range across different optimal policies in which biodiversity outcomes are assigned different weights. Consequently, for a given valuation of biodiversity credits there is some room for favoring carbon or biodiversity in restoration programs with similar economic incentives. In contrast, assigning a null value to biodiversity credits ($\beta = 0$) leaves limited options to target biodiversity in restoration programs without drastically reducing the financial viability of the policy (Fig 4a).

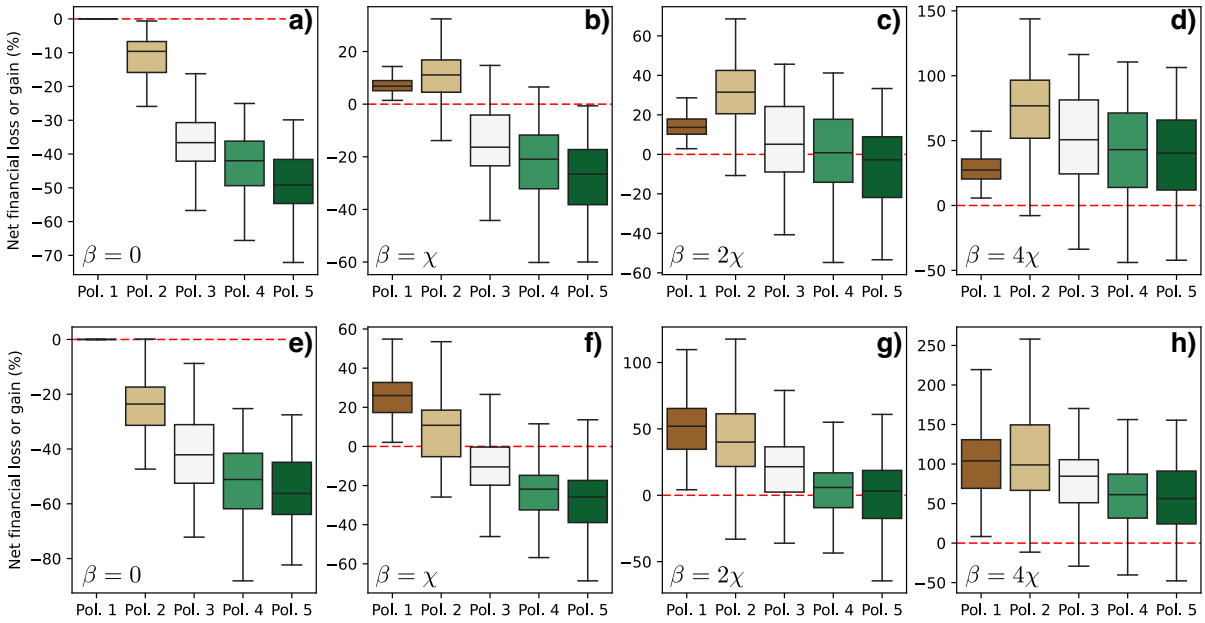

Figure 4: Financial outcomes (Φ, Equation 15) of restoration policies (net percentage loss or gain) under different valuations of biodiversity credits. Negative values indicate a net economic loss where restoration costs exceed the cumulative value of carbon and biodiversity credits, while positive values indicate a net gain. The results are shown when only carbon credits (χ) are included (i.e. $\beta = 0$; a), which we used as a baseline of net 0 financial outcome to show how biodiversity credits could contribute value to the simulated restoration programs. The results in b–d) show that when biodiversity credits are valuated at $\beta = \chi, 2\chi$ and 4χ, respectively, all policies become financially viable ($\Phi > 0$). Similar outcomes were obtained from running the simulations across 100 random regions with different spatial location, extent and resolution (e–h).

Generality of carbon and biodiversity tradeoffs

We tested to what extent our policies trained on a single arbitrary region could be applied to other geographic and climate contexts, by running restoration experiments on a set of 100 additional datasets of varying geographic location, extent, and resolution (Fig. S2). We found that the policies still resulted in general trends reflecting the objectives used to train them. Accordingly, the net carbon gain decreased from Policy 1 to 5, the number of species with improved conservation status was higher in Policies 2–5, and costs were roughly constant across all policies (Fig. S4a–c). However the patterns are less consistent when looking at changes in species number across risk category (Fig. S4d–i). These results suggest that while models trained on a region can be applied with reasonable performance to other regions they would almost certainly benefit from a region-specific re-training or fine-tuning.

Consistently with our previous experiments we found that the restoration areas selected under different policies only partly overlap (Fig. 1b), indicating that trade-offs between carbon and biodiversity priorities are common at different geographic locations and scales. These experiments also confirmed the crucial potential of biodiversity credits in contributing to the financial feasibility restoration policies that target both carbon and extinction risk reduction (Fig. 4e–h; Table S2).

Discussion

Reinforcement learning as a tool for guiding restoration programs

We demonstrated the use of multi-objective reinforcement learning to optimize ecological restoration strategies targeting climate and biodiversity positive outcomes while accounting for costs of implementation. Using weighted rewards in the training phase allowed us to calibrate models that explore the boundary of optimality (Lin et al., 2019) between net carbon gain and reduction in extinction risks obtained through restoration. This simulation and optimization framework thus allows us to evaluate the overlaps and trade-offs among restoration policies with different targets.

As our models are trained within a dynamic environment, the solutions incorporate the effects of connectivity among restored areas, that favor migration processes and range expansion underlying the simulation of species recovery. They also account for species-specific features such as climatic niche, growth rates, population density, and sensitivity to anthropogenic disturbance. The simulations show that the spatial distribution of biodiversity and carbon priorities differ, as observed in real-world scenarios (Jung et al., 2021; Alejo et al., 2025). This mismatch means that a restoration program focusing on the maximization of carbon will likely not result in the best biodiversity outcome, and vice versa.

Toward a data-driven valuation of biodiversity credits

Our results show that simulation experiments coupled with multi-objective reinforcement learning optimization allow us to explore the trade-offs between carbon and biodiversity outcomes within the constraints of a limited restoration budget. We propose that this systematic evaluation of the predicted restoration outcomes and associated costs can be used as a data-driven approach to valuate biodiversity credits relative to carbon credits, such that restoration policies targeting both carbon fixation and extinction risk reduction become financially sustainable. We showed that even a relatively small value assigned to biodiversity outcomes can create incentives favoring a multi-objective restoration policy. While this could in principle mean that "high-integrity" carbon credits, sometimes marketed as "carbon credits with co-benefits" (i.e., carbon credits sold at a higher premium cost), could in principle suffice to finance restoration projects that meet adequate climate and biodiversity goals, in practice ecological restoration companies frequently report that such added costs remain prohibitively high to be competitive in the carbon market – pushing restoration projects to marginal, cheap land.

We demonstrated our approach with a fixed definition of biodiversity credits, based on species extinction risks. However, our approach could also be used to evaluate the effect of different metrics as the basis for biodiversity crediting. For instance, other approaches could be based on population recovery or geographic range expansion (Farnsworth et al., 2015; Schipper et al., 2020). In a simulated setting, virtually any metric can be tested as we can afford to quantify the evolution of an environment through restoration at any level of demographic and spatial detail. In practice, the reference metric will have to be chosen so that it can be monitored accurately and continuously in the real world, for instance, through the re-assessment of extinction risks of selected species (Durán et al., 2020; Antonelli et al., 2023).

Biodiversity credits can be defined such that an optimal solution for both carbon and biodiversity is preferable compared to a policy targeting carbon alone. Unlike carbon, a biodiversity outcome (and valuation) is inherently context dependent (Farnsworth et al., 2015). That is because a ton of fixated carbon provide the same contribution to climate wherever it is sequestrated (except for potentially losing value through time if the sequestration is not permanent (Groom and Venmans, 2023)), while the value of the restoration and recovery of one species is necessarily linked to where the species occurs and its environmental context. Estimating trade-offs can help us defining actionable and context-dependent valuation for biodiversity credits. In our experiments, we found that setting the value of one species recovering from threatened to a non-threatened status to between 1 and 4 carbon credits was required to make a biodiversity-oriented policy financially competitive against a carbon-focused one. The value will likely differ in different geographic contexts and depending on the spatial extent considered and the selection of reference species. Yet, this approach to value biodiversity credits relative to

carbon credits can be in principle applied to different contexts and different quantities used to measure a biodiversity credit.

Applicability to real-world scenarios

The current framework could be applied to real world scenarios by replacing some variables with actual data. For example, species distribution models based on species occurrence data and biodiversity monitoring can be used to replace the simulated biodiversity data and habitat suitability in our model. The objectives and needs of indigenous and local communities, which are essential for a fair and effective protection of nature (Erbaugh et al., 2020), can be integrated into the decision making process, for instance by restricting the range of candidate sites available for restoration, or imposing particular constraints on the optimization objectives or species to be utilized for restoration. Additionally, carbon estimates in both degraded and restored areas can be used to quantify the amount of fixated carbon in disturbed habitats and its potential amount after restoration of natural vegetation. Further, real land use and predicted restoration costs can replace the respective simulated variables in our reinforcement learning framework.

Extinction risks can be initialized based on existent or predicted species extinction risks evaluated at a global or regional level (Red List of Threatened Species, 2023; Zizka et al., 2022; Estopinan et al., 2024). In our simulations, we used simple rules based on simulated population recovery to model the changes in extinction risks after implementing restoration policies. These rules can be adapted to reflect the actual criteria used to evaluate the conservation status of species, for instance, based on IUCN guidelines for global assessments or other criteria implemented in regional Red Lists (IUCN Standards and Petitions Committee, 2025).

Simulation experiments carried out outside of the region used for model training showed that there is some ability of the optimized policies to generalize to different contexts with reasonable performance. That is, climate-focused policies still achieved higher carbon fixation, while biodiversity-focused policies led to more reduction of extinction risks. However, our results also suggest that models achieve higher and more consistent performance when applied within the region used in their optimization. Thus, the application of this approach to a real-world context will benefit from area-specific model training or fine-tuning of existing policies, to best adapt to the specific properties of the focal region and its biodiversity patterns.

Our approach can therefore be informed by actual data and provide restoration priorities while leveraging potentially conflicting targets for carbon sequestration and biodiversity outcomes. Its application to real-world scenarios can provide a data-driven approach to predict not only the carbon and biodiversity outcomes of different restoration policies, but also their financial viability based on the balance between costs, carbon credits, and the potential valuation of biodiversity credits (Antonelli et al., 2024). This in turn can serve as a tool to estimate how a biodiversity crediting system could help achieving both climate and biodiversity targets.

Conclusion

The controlled simulations performed here indicate that it is possible to select areas for ecological restoration that substantially benefit both climate and biodiversity, while incorporating the needs and objectives of local and indigenous communities. Focusing on a single restoration objective would result in missed opportunities (Brancalion et al., 2025), which a multi-objective prioritization can seize integrating both carbon capture and biodiversity protection. Ecosystem restoration is increasingly recognized as a crucial step in the conservation of nature and stabilization of Earth's climate (Strassburg et al., 2020; Seddon et al., 2020). The use of robust models and AI-based optimization can help us deal with the complexity of biodiversity and its spatial and temporal dynamics, thus playing a crucial role in guiding our efforts to tackle the twin crisis of biodiversity loss and climate change.

Methods

Reinforcement learning optimization of restoration

We optimized spatial restoration planning developing new functionalities within the CAPTAIN software (Silvestro et al., 2022). We used environments and species generated through simulations to assess the performance and results of different restorations models under fully controlled settings. However, the framework can be used with empirical data, for instance by using species distribution models, climate data, and land use and other indexes of human footprint, and real costs of restoration. Similarly empirical data, when available, can be used to inform the model about species dispersal ability, sensitivity to disturbances, and -in the case of plants- net sequestrated carbon. In the following sections we first present how the updated CAPTAIN framework works for restoration with multi-objective reinforcement learning and then provide the details how how we set up the specific experiments presented in this study.

Environment – The environment is discretized into spatial units of arbitrary size that represent the cells of a grid and are used as units for the description of environmental variables, species ranges, and restoration actions. Each cell features environmental variables distinguishing land from sea and providing elevation and climate data (such as temperature and precipitation – e.g. their mean, monthly and seasonal variations; for ocean environments, specific variables such as water depth, oxygen levels, and salinity could be used). Cells are also characterized by a value of anthropogenic disturbance, indicated with $\delta \in [0, 1]$, that alters the ability of species to live in the cell and their abundance (see below) and by a restoration cost (ξ), quantifying the relative cost of restoring the cell to a low level of disturbance. This could include, for instance, the cost of buying the cell, displacing or removing the cause of disturbance, and potentially taking action to reintroduce native species. Note that in real conservation work, restoration costs may be complex depending on the legal and socio-economic setting. For example, buying land

is not always sufficient to stop all logging, but we assume that we can model appropriate total costs here. All these features can be set to change through time, e.g., capturing the dynamics of predicted climate change, variation in disturbance patterns, and evolving costs per cell.

The environment is also characterized by the species that live (or can live) in it. Specifically, for each cell and each species in the system, there is a value describing the number of individuals that can occupy the cell in the absence of disturbance (i.e. the natural state). This is a cell- and species-specific carrying capacity $K_{c,s}$ which depends on both the habitat suitability of the cell for that species (H_s) and a species-specific maximum abundance per cell (K_s):

$$K_{c,s} = H_{c,s} \times K_s, \text{ where } H \in [0,1] \text{ and } K_s \in \mathbb{R}^+ \tag{1}$$

The environment also includes the actual number of individuals of a species at any given time t, $N_{c,s,t}$, which results from the interaction between the species and the environment, including the effects of disturbance, growth rates, and dispersal.

For each species, we define a specific value of sensitivity $S_s \in [0,1]$ that determines individual mortality fractions in response to a disturbance in a cell, $\delta_{c,t}$. Species are also assigned a growth rate $G_s > 1$, which describes the demographic growth in one time step, and a dispersal rate $D_s \geq 0$, which determines the number of individuals that disperse from one cell to another. In the current implementation, we assume individuals to be static (e.g., plants). Therefore, dispersal does not move individuals from one cell to another. Instead, dispersal only adds individuals to the receiving cell (replicating for instance the dispersal of seeds). The number of individuals dispersing from cell i to j follows a truncated exponential distribution:

$$I_{i \to j,s,t} = \begin{cases} N_{i,s,t} \times \exp(-d_{ij} D_s^{-1}) \text{ if } d_{ij} \leq T \\ 0 \text{ if } d_{ij} > T \end{cases} \tag{2}$$

where d_{ij} is the distance between the two cells and T is the maximum distance that can be reached in a single time step. Species distributions and abundances change over time based on growth, dispersal, and mortality. Specifically, at each time step, the distribution and abundance of species is updated according to the following equation:

$$N_{c,s,t} = \min\left(N_{c,s,t-1} \times (1 - S_s \delta_{c,t}) \times G_s + I_{c,s,t}, K_{c,s}\right) \tag{3}$$

where $N_{c,s,t-1}$ is the number of individuals of species s in cell c at the previous time step, $S_s \delta_{c,t}$ is the product between the disturbance in the cell at time t and the species sensitivity, with $(1 - S_s \delta_{c,t})$ representing the mortality driven by disturbance. The term G_s represents the contribution of demographic growth, while

$$I_{c,s,t} = \sum_{\forall j \neq i} I_{j \to i,s,t-1} \tag{4}$$

is the sum of the individuals dispersing from all other cells and represents the contribution of dispersal to the cell's population.

Each species in the environment is assigned a value representing its expected carbon value $C_s \in \mathbb{R}^+$. This is a single value and can be interpreted as the net amount of carbon sequestered by one individual of species s. As our simulated system in its current implementation does not track the age of each individual, we define the total carbon in a cell as the sum of carbon storage from all individuals and all species in the cell:

$$C_{c,t} = \sum_s C_s N_{c,s,t} \tag{5}$$

This quantity could in principle extended to include carbon in the soil. We note that, as the age of individuals is not modeled here, the simulated accumulation of carbon in the current implementation might occur more rapidly than in empirical systems.

Finally, we also include a 5-tiered classification of species extinction risk, inspired by the International Union for Conservation of Nature (IUCN) Red List of Threatened Species (2023). When working on empirical data, this can be effectively replaced by actual IUCN Red List assessments or model-based predictions (Zizka et al., 2022; Antonelli et al., 2023). Hereafter, we will use the Red List-inspired labels to describe increasing extinction risks: Least Concern (LC), Near Threatened (NT), Vulnerable (VU), Endangered (EN), and Critically Endangered or (locally) extinct (CR). As part of the evaluation of the implemented restoration action, the initial extinction risk of a species can be updated based on the simulated dynamics of the species demography or geographic range, for instance, increasing or decreasing the extinction risk depending on whether the census population size of a species is decreasing or increasing, respectively. We indicate the extinction risk of each species s at time t as:

$$X_{s,t} \in \{1, ..., 5\}$$

Agent – The agent in the system observes the environment at each time step extracting features for each grid cell about its disturbance level, its cost for restoration, the species that occur there, those that could occur there naturally in the absence of disturbance, their current extinction risk and carbon contribution (a detailed list of the features used in our experiments is provided in the next section). Following the modeling framework developed in the initial implementation of CAPTAIN (Silvestro et al., 2022), the features of each cell feed into a fully connected deep neural network with `tanh` activation function (Apicella et al., 2021) in the hidden layers and a single output node. The parameters of the neural networks are shared

across all cells and their output values are pooled into a single `softmax` function providing a restoration rank of each cell. The agent will then select deterministically or probabilistically a pre-defined number of cells based on their ranking, within which restoration occurs. Here we assume restoration to result in the instantaneous removal of disturbance from the cells, although different dynamics such as a gradual decrease or partial reduction in disturbance can be implemented. The restoration action is associated with the cumulative cost of the selected cells. The agent is given an initial budget that can be spent for restoration, and only cells with a cost lower than the current budget can be selected for restoration. Once the disturbance is removed from a cell, its populations will continue to evolve according to Equation 3 with reduced (or removed) disturbance.

Rewards and optimization – The actions of the agent result in changes in the environment that are quantified through a *total reward*, defined based on the objective of the restoration policy. The parameters of the neural networks determining the restoration action are chosen to maximize the cumulative net reward at the end of a simulation of a fixed number of steps. The net reward or payoff is here defined as the weighted sum of three components quantifying net carbon gain (carbon reward), net biodiversity outcomes (biodiversity reward), and costs (budget reward). These quantities allow us to optimize policies with one or multiple objectives, generating socially optimal solutions when multiple objectives are used. At each time step, the carbon reward of a restoration action is computed as the fraction between the net change in total carbon stored in the system compared with its initial state (at time $t = 0$) and the amount of carbon lost between the natural state (in the absence of any human disturbance) and the initial state:

$$\mathcal{R}_{C,t} = \frac{\sum_c C_{c,t} - \sum_c C_{c,t=0}}{\sum_c C_{c,\text{natural}} - \sum_c C_{c,t=0}} - \mathcal{R}_{C,t-1} \tag{6}$$

We note that with simulated datasets, like the ones used in this study, the natural carbon amount C_c, natural is known *a priori*. When working with empirical data, this quantity might have to be approximated based on measurements, such as through survey data (for trees, this can be derived indirectly from measurements of height and stem diameter at breast height (dbh), or more directly from biovolumes estimated from lidar scans).

The total reward of a restoration action also includes a measurement of the resulting biodiversity outcome, which can be defined in many different ways, as shown by a recent compendium (Burgess et al., 2024) currently listing over 570 biodiversity metrics. In the current implementation, we exemplify the calculation of biodiversity gains in a similar way as recently calculated in Silvestro et al. (2023), who used the number of extinct or threatened species as a negative reward. The focus on extinction risk is also aligned with other recently proposed metrics for conservation priorities and threat abatement (Durán et al., 2020; Mair et al., 2021; Eyres et al.,

2024). Here we use an average of the restored species population weighted by their current conservation status. Specifically, for any given extinction risk status $X = i$, we identify the subset of species S_i such that:

$$X_{s,t} = i, \forall s \in S_i \tag{7}$$

and compute the respective reward as

$$\mathcal{R}_{X=i,t} = \begin{cases} \frac{w_i}{\sum_j w_j} \text{ if } S_i = \emptyset \\ \sum_{s \in S_i} \frac{w_i R_{s,t}}{\sum_j w_j} \text{ otherwise} \end{cases} \tag{8}$$

where w_i is a weight associated with extinction risk status i, and $R_{s,t}$ is the ratio between the number of individuals of species s at time t found within restored cells (indicated with M) and the total population of the species in the environment:

$$R_{s,t} = \frac{\sum_{\varkappa \in M} N_{\varkappa,s,t}}{\sum_c N_{c,s,t}} \tag{9}$$

Thus, $R_{s,t} = 0$ when the species is not found in any restored areas and $R_{s,t} = 1$ when all its population falls inside restored areas. The weights are set to favor scenarios where the highest average fraction of population is restored in species classified in a high-risk category:

$$\begin{cases} w_1 = 1 \\ w_2 = 8 \\ w_3 = 16 \\ w_4 = 32 \\ w_5 = 64 \end{cases} \tag{10}$$

so that, e.g., protecting 30% of the individuals in a Least Concern species ($X = 1$) contributes 0.0025 to the reward, while the same fraction in a Critically Endangered species ($X = 5$) contributes 0.1587. The overall biodiversity reward is set to the sum of $R_{X,t}$ over all extinction risk categories. The use of increasing weights associated with extinction risk is consistent with several metrics developed to quantify conservation priorities (Pimiento et al., 2020; Mair et al., 2021; Gumbs et al., 2023).

Finally, we include restoration costs ε_c as part of the total reward. Based on an initial budget $B_{t=0}$ we compute the budget reward $R_{\varepsilon,t}$ as the ratio between the current budget and the initial one. Here we set the initial budget as a function of the restoration target (30%) and the total cost of restoration. Specifically, given a restoration cost ε_c (defined in the next section) for cell c

and a restoration target $r \in [0,1]$ specifying the fraction of cells to be restored, we set the initial budget to a value arbitrarily larger than the total average cost a number of cells required to meet the target:

$$B_{t=0} > r \sum_{c} \varepsilon_c \tag{11}$$

The rewards can be used alone or in combination to optimize policies that pursue one or more targets. Thus, we defined the total reward as:

$$\mathcal{R}_t = w_C \mathcal{R}_{C,t} + w_X \mathcal{R}_{X,t} + w_\varepsilon \mathcal{R}_{\varepsilon,t} \tag{12}$$

were the pre-defined weights (w) quantify the proportional importance assigned to carbon, biodiversity, and budget. Here we set the weight associated with the budget (w_ε) to a constant and set the others to express their relative importance. Under this formulation, a policy targeting a single objective (e.g., focusing exclusively on budget) can be optimized by setting the respective weight to 1 and the others to 0 (e.g., $w_C, w_X = 0$, and $w_\varepsilon = 1$). The weights can also be used to estimate different solutions where the total reward includes two or more potentially competing optimality criteria.

We use the parallel genetic strategies algorithm used in Silvestro et al. (2022) to estimate the parameters of the neural networks in the model that maximize the reward. Here, multiple environments are run in parallel with an optimization step occurring at the end of each training epoch.

Simulations

We use simulations of biodiversity to evaluate the ability of restoration policies optimized through reinforcement learning to achieve different targets, focusing on carbon and biodiversity outcomes. Although the same policies could be applied to empirical data, simulations allow us to benchmark and compare different restoration actions across many simulated scenarios in which we can accurately measure all outcomes.

Parameters of the simulated environment – We simulated environments based on a raster of 128×128 cells, of 25 Km resolution at the equator, inspired from physical properties from an arbitrary region spanning a portion of Central and South America (Fig. S1a). The properties include its coastline, elevation and cell-specific climate average values obtained from CHELSA (Karger et al., 2023), here including the mean annual temperature, mean annual precipitation and the minimum temperature of the coldest month. The selection of variables is not intended to provide an exhaustive account of the climate variables determining a species niche but were used to generate artificial species distributions with realistically constrained geographic ranges.

After z-transforming these topographic and climate variables we indicate them as $V = V_1, ..., V_v$, where $V_{i,x,y}$ represents the value of the ith variable corresponding to the longitude and latitude coordinates x, y.

We simulated species geographic range by first drawing a random pair of coordinates (m, n), based on which we extracted the corresponding elevation and climate data. We use a normal distribution to describe the environmental and climatic niche of the species $\text{MVN}(V_{m,n}, \Sigma)$, where the diagonal of the covariance matrix is a vector $\sigma = \{\sigma_1, ..., \sigma_V\}$ is randomly drawn for each species from a uniform distribution $\mathcal{U}(0.05, 1)$ and off-diagonal values (covariances) are set to 0. The σ vector will therefore generate species with a range of narrow to wide climatic tolerances. We also introduce a distance factor, to ensure that species ranges are also spatially constrained e.g. through limited dispersal ability and not only determined by habitat suitability. The effect of distance was simulated through an exponential distribution with rate parameter d randomly drawn for each species from a uniform distribution $\mathcal{U}(0.05, 1)$. Thus, for any set of coordinates (x, y) the relative population density of a species is $\text{MVN}(V_{x,y}, V_{m,n}, \Sigma)\exp(-d\Delta(x, y, m, n))$, where $\Delta(\cdot)$ is the cartesian distance between (m, n) and (x, y). We use this formula to define the habitat suitability $H_{c,s}$ of a species s in each cell c, as used in Equation 1.

We used the empirical mean annual precipitation layer (rescaled between 0 and 1) to split cells into three categories with high (> 0.4), intermediate (< 0.4) and low precipitation (< 0.25). We use these categories to simulate different vegetation types, e.g. representing one dominated by trees, a mixed open habitat, and a fully open habitat (Fig. S1a). These vegetation types are then used to characterize properties of the species that inhabit them. Specifically, we simulated three categories of species to represent different growth forms, e.g. herbaceous plants, shrubs, trees. We assigned each species to one of the categories randomly with probability set equal to the proportion of the overlap of their simulated range with low, intermediate and high precipitation, respectively.

We drew the species demographic growth rates from a Weibull distribution:

$$G_s \sim 1 + W(0.75, s)$$

where $s = 10$, 2, and 0.1 for the first, second, and third growth forms, respectively. This results in faster average demographic growth for simulated 'herbaceous' species. Similarly, we drew the species-specific carbon value from a gamma distribution:

$$C_s \sim \Gamma\left(2, \frac{m}{2}\right)$$

where $m = 0.5$, 1, or 2 is the mean of the distribution applied to the first, second, and third growth forms, respectively. We draw the species-specific carrying capacity K_s (i.e. the maximum number of individuals that can live in a cell; Equation 1) from a beta distribution that generates

many species with low abundance and few species with high abundance:

$$K_s \sim 2 + 100 \times B(0.6, 1.2)$$

This leads to a mean carrying capacity a species per cell of ≈ 35 ranging from ≈ 2 to ≈ 102. We assigned these values by setting the largest ones to species of the first growth forms, the intermediate values to those of the second, and the smallest values to species of the third. Thus, species with fast growth and small carbon tended to have higher carrying capacity than slow growing species with high carbon value (reflecting the realistic differences in wood densities and carbon storage among woody species). Examples of the simulated species' natural abundances resulting from both habitat suitability and species-specific carrying capacity are shown in Fig. S1d.

Disturbance and restoration costs – We simulated spatially heterogeneous disturbance through random multivariate normal distributions rescaled to range between 0 and 1, which we multiplied by exp(-rescaled elevation) to simulate a scenario in which anthropogenic disturbance decreases on average with elevation (given the increased difficulties in accessing natural resources or performing agriculture; see example of simulated disturbance in Fig. S1c). We assigned a sensitivity to disturbance for each species based on a symmetric U-shaped beta distribution:

$$S_s \sim B(0.5, 0.5)$$

thus generating several species with either very low or very high sensitivity. Prior to applying any restoration policy, the environment was set to evolve for a number of steps until reaching an equilibrium under the simulated disturbance.

The restoration costs were linked to disturbance, so that on average the cost of restoration was higher for highly degraded cells (reflecting lesser opportunities for natural regeneration and higher distance to seed sources). We added some stochasticity in this pattern by computing costs as the product between disturbance and another rescaled random multivariate normal distribution (see example of simulated disturbance in Fig. S1c). We set the top 5% cells with the highest cost as unrestorable, further raising their cost to exceed the budget. These cells are meant to represent areas where restoration is not possible e.g. because that would imply displacing inhabited areas or activities and infrastructure that cannot be moved.

Evolution of the species extinction risk – We initialized the species extinction risk based on their range size, on their sensitivity, and on the effect of disturbance:

$$r = \left(\frac{A_s}{\max(A_s)}\right)^l \times \frac{A_s(1 - S_s\delta)}{A_s} \tag{13}$$

where A_s is the natural range of the species (number of occupied cells) in the absence of disturbance. Thus, the value r decreases with small natural range sizes and with increasing sensitivity and/or disturbance (Strassburg et al., 2020). The exponent l (here set to $l = 0.3$) has the effect of flattening the difference in initial extinction risk based on natural range size. We then mapped the value to an initial extinction risk $X_t = 0$ (Table 1). We note that this approach to simulate species extinction risk, while based on common assumptions about species range size and decline due to anthropogenic disturbance, is arbitrary and should be replaced with actual risk assessments (e.g. from the IUCN Red List) when running the optimization with real world data.

Throughout the simulations, the extinction risks of all species were re-evaluated at each step, and we used two criteria to re-assess the conservation status of a species. First, we calculated the number of restored cells within which a species s occurred at time t, indicated as $O_{s,t} \cap M_t$ relative to the species natural range size A_s. Thus, increasing values of $(O_{s,t} \cap M_t)/A_s$ indicate that a proportionally higher number of restored cells is occupied (potentially after re-colonization) by species s, leading to a decreased extinction risk (Table 1, Criterion 1). Second, we calculate the ratio between the occupied geographic range at time t and the occupied range at the beginning of the simulation $O_{s,t} = 0$, i.e. prior to enforcing restoration programs. This ratio is used to verify if the species is in decline, e.g., if the restoration is ineffective or if other factors, such as climate change, negatively affect a species' range. We used the thresholds reported in Table 1 (Criterion 2) to determine if the species was in decline and to what extent that affected the species extinction risk. Finally, any species among the initial simulated pool with occupied range size of 0 was moved to the highest risk category and considered as possibly extinct in the study system. These criteria are inspired by some of the Red List guidelines (IUCN Standards and Petitions Committee, 2025), specifically looking at species range sizes, and recovery or declining trends through time. While we used the occupied range as the main metric, the criteria can be defined based on alternative metrics, e.g., based on census population size and its temporal variation or fragmentation of the species range.

Table 1: Thresholds used to set species extinction risks and their reassessment.

$X_{s,t=0}$	r	Criterion 1	$(O_{s,t} \cap M_t)/A_s$	Criterion 2	$O_{s,t}/O_{s,t=0}$
LC	>0.95	$X_{s,t} = X_{s,t=0}$	< 0.10	$X_{s,t} = X_{s,t=0}$	< 0.10
NT	0.9-0.95	$X_{s,t} = X_{s,t=0} - 1$	0.10 - 0.25	$X_{s,t} = X_{s,t=0} + 1$	0.10 - 0.25
VU	0.8-0.9	$X_{s,t} = X_{s,t=0} - 2$	0.25 - 0.50	$X_{s,t} = X_{s,t=0} + 2$	0.25 - 0.50
EN	0.6-0.8	$X_{s,t} = X_{s,t=0} - 3$	0.50 - 0.80	$X_{s,t} = X_{s,t=0} + 3$	0.50 - 0.80
CR	<0.6	$X_{s,t} = X_{s,t=0} - 4$	> 0.80	$X_{s,t} = X_{s,t=0} + 4$	> 0.80

Observed features – The agent was set to observe the entire environment at each step, extracting several features for each cell, which form the input of the neural network and determine the next action. We used the same 16 features for all policies, independently of the

reward. The first features were the disturbance in the cell and the mean disturbance in the 5×5 cells around it (features 1-2). The latter feature provides a measure of the spatial distribution of disturbance, for instance, differentiating between a disturbed cell neighboring a pristine area and one placed within a broader degraded area. We then included the potential carbon net change in the cell calculated as the difference between the total carbon in the cell under natural conditions and the current amount of carbon, i.e. $C_{c,\text{natural}} - C_{c,t}$. The same measure is also provided as a mean value in the 5×5 cells around the cell (features 3-4). The features additionally include the number of CR, EN, VU and NT species in the cell based on the current populations occupying the cells and the mean number of CR and EN in the surrounding 5×5 cells (features 5-10). We also included the number of species in each of the five extinction risk categories, but counting only species that do not occur in any of the previously restored cells, if any (features 11-15). These features serve as a measure of complementarity among cells so that the model can choose to restore cells with a different species composition compared with areas already restored. Finally, we included the cell's cost (feature 16). All features were rescaled between 0 and 1 before being fed to the neural network.

Training – We trained models with 250 species running simulations for 30 time steps (such as years). The agent was set to restore 100 cells at each step, i.e., reaching 3,000 restored cells at the end of the simulation, i.e., corresponding to about 30% of the land cells in our grid to mimic the 30% restoration target of the Kunming-Montreal Global Biodiversity Framework. We used a parallel genetic strategies algorithm with 8 environments evolving in parallel and ran 5,000 epochs to train the models. We used small neural networks for each cell, connecting the 16 features to a first hidden layer of 3 nodes, followed by a second hidden layer of 2 nodes. The parameters of the network were shared among all cells. We trained different models exploring a gradient of optimality (Lin et al., 2019) going from favoring the carbon outcome to favoring the biodiversity outcome. We achieved this by assigning different weights to the reward (see Equation 12). Specifically, we trained five models, keeping the weight of the budget reward constant at 0.5 and assigning different weights to carbon and biodiversity such that the sum of the weights was kept constant and equal to 1 (Table 2).

Experiments on the reference region– After training the models, we used them on 100 simulations each including 500 species and initialized following the same approach utilized to generate the training data. These simulations were run on the same reference region used for the optimization of the policies. Although these simulations were based on the same simulation settings, they differed among each other being based on different random draws for all parameters. Thus, the distribution of species, disturbance, and costs and the traits associated with each species (carbon, sensitivity, growth rates) varied randomly across simulations. This allowed us to evaluate the results of different models across a range of different The simulated environments in their initial state, i.e. prior to restoration, included on average 30 CR species (standard deviation, sd: 18), 87 EN (sd: 39), 156 VU (sd: 31), 139 NT (sd: 30), and 87 LC (sd: 56) species. As during

Table 2: Weights used in the different trained models and assigned to the carbon reward (w_C), biodiversity reward (w_X), and budget reward (w_ε). We considered a range of models maintaining constant the weight on the budget reward, while progressively moving from a carbon-focused optimization (Policy 1) to a biodiversity-focused one (Policy 5). The intermediate policies allowed us to explore the trade-offs between these potentially competing objectives.

	w_C	w_X	w_ε
Policy 1	0.5	0	0.5
Policy 2	0.4	0.1	0.5
Policy 3	0.25	0.25	0.5
Policy 4	0.1	0.4	0.5
Policy 5	0	0.5	0.5

training, the simulated datasets were then let evolve for 30 time-steps with the agent interacting with the environment based on each of the five trained policies and restoring 3,000 cells. We evaluated the outcomes of the different policies by calculating the amount of carbon gained through restoration, the changes achieved in terms of number of species falling in each of the five extinction risk categories, the number of species occurring within restored areas at the end of the simulation, and the total cost of restoration. This allowed us to quantify the trade-off between carbon and biodiversity outcomes along a gradient of optimal solutions.

Experiments on different environments– We also evaluated the outcomes of our trained models on a set of 100 environments that varied in their spatial extent, resolution, and geographic location. Specifically, we generated 100 additional rasters of 128×128 cells with resolution at the equator of 2.5, 10, or 25 Km. Their placement globally was random but conditioned on having at least 70% of cells on land (Fig. S2). For each raster we extracted elevation and climate data and simulated 500 species following the procedure described above. We then ran the five optimized policies on these datasets to evaluate the ability of our trained models to generalize outside of the region on which they where optimized.

Carbon and biodiversity credits

We implemented a simplified concept of carbon credits by assigning a monetary value to the amount of net carbon gain (or potentially loss) incurred during the simulations. As carbon (like restoration costs and budget) is set in arbitrary units in our simulations, we defined the value of a carbon credit unit (χ) as the value of the net carbon gain obtained through Policy 1, namely the policy placing the highest weight on carbon, leading to the full compensation of the restoration costs:

$$\chi \frac{\sum_c C_{c,t} - \sum_c C_{c,t=0}}{\sum_c \sum_t \varepsilon_{c,t}} := 1 \tag{14}$$

where χ is the value of one carbon credit unit based on the simulated systems used for carbon and budget, the fractions numerator is the net carbon gain at time t compared to the initial state ($t = 0$) and the denominator is the total cost of restoration. Thus, under this definition, the carbon credits accumulated through the implementation of Policy 1 results, by construction, in full compensation of the cumulative restoration costs, with no economic loss at the end of the simulation. We note that this definition is arbitrary and could be replaced by other definitions, including empirical valuations of carbon in the cap-and-trade system of the European Union, or others.

The same carbon credit value χ was then applied to the other four policies where carbon was given lower weights (Policies 2–5) in the training phase with the inclusion of a positive weight assigned to the biodiversity reward. This allowed us to evaluate to what extent the carbon credits accumulated through restoration policies that optimized for both carbon and biodiversity (Policies 2–4) or biodiversity only (Policy 5) were able to compensate for the restoration costs.

Next, we selected a simple metric as a reference to define a biodiversity credit and developed an approach to valuate the credit relative to the value of a carbon credit. Specifically, we quantified the biodiversity credits as the number of species that moved from a threatened class (CR, EN, VU) at the beginning of the simulation to a non-threatened class (NT, LC) after restoration. This differs to current practice in some of the emerging biodiversity credit markets, such as in Colombia where one biodiversity credit unit equals one hectare regardless of how many species were benefited by the restoration and in which ways. However, the same approach utilized here can be used for alternative definitions of biodiversity credits.

We then evaluated the net financial loss or gain of each policy after stacking the value of carbon and biodiversity credits under different valuations of biodiversity credits. We indicate with β the value of one biodiversity credit obtained from moving a species from a threatened to a non-threatened status after restoration. Given that we are using arbitrary units for costs and carbon credits, we defined the value of a biodiversity credit relative to that of a carbon credit. This allowed us to assess which policy would become financially viable or preferable if e.g. the value of one biodiversity credit equals that of a carbon credit (i.e. $\beta = \chi$). We tested valuations of $\beta = 0, \chi, 2\chi,$ and 4χ and calculated the net financial loss or gain of a policy (Φ) as a percentage of the restoration cost:

$$\Phi_{\text{Pol. } i} = \left(\frac{\chi C_{\text{Pol. } i} + \beta S_{\text{Pol. } i}}{\varepsilon_{\text{Pol. } i}} - 1 \right) \times 100 \tag{15}$$

where $C_{\text{Pol.}i}$ is the net carbon gain of the ith policy, $S_{\text{Pol.}i}$ is the policy's net biodiversity outcome (as defined above), and $\varepsilon_{\text{Pol}i}$ is the cost of restoration.

Finally calculated, for any given policy i, the value that should be attributed to a biodiversity credit so that the policy results in the same net financial outcome as a restoration program

focused on carbon (i.e. our Policy 1):

$$\frac{\chi C_{\text{Pol. 1}} + \beta S_{\text{Pol. 1}}}{\varepsilon_{\text{Pol. 1}}} := \frac{\chi C_{\text{Pol. } i} + \beta S_{\text{Pol. } i}}{\varepsilon_{\text{Pol. } i}} \tag{16}$$

The value of one biodiversity credit is therefore, after solving Equation 16 for β:

$$\beta = \chi \frac{C_{\text{Pol. } i}\, \varepsilon_{\text{Pol. 1}} - C_{\text{Pol. 1}}\, \varepsilon_{\text{Pol. } i}}{S_{\text{Pol. 1}}\, \varepsilon_{\text{Pol. } i} - S_{\text{Pol. } i}\, \varepsilon_{\text{Pol. 1}}} \tag{17}$$

Using this formula, we derive the value of a biodiversity credit (relative to the value of a carbon credit) required to make the financial outcome (defined in Eq. 15) of any optimal restoration policy equal to the profit of a carbon-oriented policy (our Policy 1), that is such that $\Phi_{\text{Pol. } i} = \Phi_{\text{Pol. 1}}$. We note that under this definition, the value of biodiversity credits is context specific, as it depends on the overall area considered and its restorable carbon and biodiversity. Within the area, the restoration of cells that contribute the most to reducing the extinction risk of species (e.g., cells that protect rare and threatened species) will lead to the highest biodiversity credits (Drupp et al., 2024).

Funding

A.A., B.G. T.S. and D.S. are members of the BIOPATH research programme funded by the Swedish Foundation for Strategic Environmental Research MISTRA (F 2022/1448). A.A. is also funded by the Swedish Research Council (VR: 2019-05191, 2024-04303) and the Royal Botanic Gardens, Kew; B.G. is funded by Dragon Capital and the NERC/UKRI projects BIOADD (ref: NE/X002292/1), BIOESG (NE/X016560/1) and RENEW (ref: NE/W004941/1); D.S. by ETH Zurich and the Swiss National Science Foundation (PCEFP3_187012); and T.S. by the Kamprad Foundation. The funders had no role in the contents of this study.

*

References

Alejo, C., A. Luers, A. Ventimiglia, and H. D. Matthews. 2025. Maximizing nature-based solutions using artificial intelligence to align global biodiversity, climate, and water targets. bioRxiv Pages 2025–01.

Antonelli, A., K. L. Dhanjal-Adams, and D. Silvestro. 2023. Integrating machine learning, remote sensing and citizen science to create an early warning system for biodiversity. Plants, people, planet 5:307–316.

Antonelli, A., X. Rueda, R. Calcagno, and P. Nantongo Kalunda. 2024. How biodiversity credits could help to conserve and restore nature. Nature 634:1045–1049.

Apicella, A., F. Donnarumma, F. Isgrò, and R. Prevete. 2021. A survey on modern trainable activation functions. Neural Networks 138:14–32.

Brancalion, P. H., F. Hua, F. H. Joyce, A. Antonelli, and K. D. Holl. 2025. Moving biodiversity from an afterthought to a key outcome of forest restoration. Nature Reviews Biodiversity Pages 1–14.

Burgess, N. D., N. Ali, J. Bedford, N. Bhola, S. Brooks, A. Cierna, R. Correa, M. Harris, A. Hargey, J. Hughes, et al. 2024. Global metrics for terrestrial biodiversity. Annual Review of Environment and Resources 49:673–709.

Deutz, A., G. M. Heal, R. Niu, E. Swanson, T. Townshend, L. Zhu, A. Delmar, A. Meghji, S. A. Sethi, and J. Tobin-de la Puente. 2020. Financing nature: Closing the global biodiversity financing gap. The Paulson Institute, The Nature Conservancy, and the Cornell Atkinson Center for Sustainability 256.

Di Sacco, A., K. A. Hardwick, D. Blakesley, P. H. Brancalion, E. Breman, L. Cecilio Rebola, S. Chomba, K. Dixon, S. Elliott, G. Ruyonga, et al. 2021. Ten golden rules for reforestation to optimize carbon sequestration, biodiversity recovery and livelihood benefits. Global Change Biology 27:1328–1348.

Drupp, M., M. Hänsel, E. Fenichel, M. Freeman, C. Gollier, B. Groom, G. Heal, P. Howard, A. Millner, F. Moore, et al. 2024. Accounting for the increasing benefits from scarce ecosystems. Science 383:1062–1064.

Durán, A. P., J. M. Green, C. D. West, P. Visconti, N. D. Burgess, M. Virah-Sawmy, and A. Balmford. 2020. A practical approach to measuring the biodiversity impacts of land conversion. Methods in Ecology and Evolution 11:910–921.

Erbaugh, J. T., N. Pradhan, J. Adams, J. A. Oldekop, A. Agrawal, D. Brockington, R. Pritchard, and A. Chhatre. 2020. Global forest restoration and the importance of prioritizing local communities. Nature Ecology & Evolution 4:1472–1476.

Estopinan, J., P. Bonnet, M. Servajean, F. Munoz, and A. Joly. 2024. Modelling species distributions with deep learning to predict plant extinction risk and assess climate change impacts. arXiv preprint arXiv:2401.05470 .

Eyres, A., T. Ball, M. Dales, T. Swinfield, A. Arnell, D. Baisero, A. P. Durán, J. Green, R. E. Green, A. Madhavapeddy, et al. 2024. Life: A metric for quantitively mapping the impact of land-cover change on global extinctions. Earth and Environmental Sciences .

Farnsworth, K. D., A. H. Adenuga, and R. De Groot. 2015. The complexity of biodiversity: A biological perspective on economic valuation. Ecological Economics 120:350–354.

Groom, B. and F. Venmans. 2023. The social value of offsets. Nature 619:768–773.

Gumbs, R., C. L. Gray, M. Böhm, I. J. Burfield, O. R. Couchman, D. P. Faith, F. Forest, M. Hoffmann, N. J. Isaac, W. Jetz, et al. 2023. The edge2 protocol: Advancing the prioritisation of evolutionarily distinct and globally endangered species for practical conservation action. PLoS Biology 21:e3001991.

Holl, K. D. and P. H. Brancalion. 2020. Tree planting is not a simple solution. Science 368:580–581.

Hua, F., L. A. Bruijnzeel, P. Meli, P. A. Martin, J. Zhang, S. Nakagawa, X. Miao, W. Wang, C. McEvoy, J. L. Peña-Arancibia, et al. 2022. The biodiversity and ecosystem service contributions and trade-offs of forest restoration approaches. Science 376:839–844.

IUCN Standards and Petitions Committee. 2025. Guidelines for Using the IUCN Red List Categories and Criteria. Version 16 (http://www.iucnredlist.org/documents/RedListGuidelines.pdf).

Jung, M., A. Arnell, X. De Lamo, S. García-Rangel, M. Lewis, J. Mark, C. Merow, L. Miles, I. Ondo, S. Pironon, et al. 2021. Areas of global importance for conserving terrestrial biodiversity, carbon and water. Nature Ecology & Evolution 5:1499–1509.

Karger, D. N., M. P. Nobis, S. Normand, C. H. Graham, and N. E. Zimmermann. 2023. Chelsa-trace21k–high-resolution (1 km) downscaled transient temperature and precipitation data since the last glacial maximum. Climate of the Past 19:439–456.

Lin, X., H.-L. Zhen, Z. Li, Q.-F. Zhang, and S. Kwong. 2019. Pareto multi-task learning. Advances in neural information processing systems 32.

Mair, L., L. A. Bennun, T. M. Brooks, S. H. Butchart, F. C. Bolam, N. D. Burgess, J. M. Ekstrom, E. Milner-Gulland, M. Hoffmann, K. Ma, et al. 2021. A metric for spatially explicit contributions to science-based species targets. Nature Ecology & Evolution 5:836–844.

Parr, C. L., M. Te Beest, and N. Stevens. 2024. Conflation of reforestation with restoration is widespread. Science 383:698–701.

Pimiento, C., F. Leprieur, D. Silvestro, J. Lefcheck, C. Albouy, D. Rasher, M. Davis, J. Svenning, and J. Griffin. 2020. Functional diversity of marine megafauna in the anthropocene. sci adv 6: eaay7650.

Red List of Threatened Species. 2023. Red List of the International Union for the Conservation of Nature (http://www.iucnredlist.org).

Schipper, A. M., J. P. Hilbers, J. R. Meijer, L. H. Antão, A. Benítez-López, M. M. de Jonge, L. H. Leemans, E. Scheper, R. Alkemade, J. C. Doelman, et al. 2020. Projecting terrestrial biodiversity intactness with globio 4. Global Change Biology 26:760–771.

Seddon, N., A. Chausson, P. Berry, C. A. Girardin, A. Smith, and B. Turner. 2020. Understanding the value and limits of nature-based solutions to climate change and other global challenges. Philosophical Transactions of the Royal Society B 375:20190120.

Silvestro, D., S. Goria, B. Groom, T. Sterner, and A. Antonelli. 2023. The 30 by 30 biodiversity commitment and financial disclosure: Metrics matter. Available at SSRN 4594937 Page 10.2139/ssrn.4594937.

Silvestro, D., S. Goria, T. Sterner, and A. Antonelli. 2022. Improving biodiversity protection through artificial intelligence. Nature sustainability 5:415–424.

Strassburg, B. B., A. Iribarrem, H. L. Beyer, C. L. Cordeiro, R. Crouzeilles, C. C. Jakovac, A. Braga Junqueira, E. Lacerda, A. E. Latawiec, A. Balmford, et al. 2020. Global priority areas for ecosystem restoration. Nature 586:724–729.

Sutton, R. S., A. G. Barto, et al. 1998. Reinforcement learning: An introduction vol. 1. MIT press Cambridge.

Zizka, A., T. Andermann, and D. Silvestro. 2022. Iucnn–deep learning approaches to approximate species' extinction risk. Diversity and Distributions 28:227–241.

SUPPLEMENTARY INFORMATION

Supplementary Tables

	β/χ			
	median	st. dev.	25%	75%
Policy 1	0	0	0	0
Policy 2	0.80	2.25	0.51	1.20
Policy 3	2.65	3.93	1.72	4.17
Policy 4	2.94	7.66	2.17	5.01
Policy 5	3.51	6.05	2.56	5.39

Table S1: Biodiversity credit valuation estimated to reach the same net financial gain as in Policy 1 across all policies, computed based on Equation 17. The value are calculated across 100 simulations based on the reference region (on which models were trained).

	β/χ			
	median	st. dev.	25%	75%
Policy 1	0	0	0	0
Policy 2	2.64	18.09	0.26	6.50
Policy 3	4.03	45.29	1.31	10.19
Policy 4	6.06	13.58	2.09	11.67
Policy 5	4.49	30.96	0.00	11.20

Table S2: Biodiversity credit valuation estimated to reach the same net financial gain as in Policy 1 across all policies, computed based on Equation 17. The value are calculated across 100 simulations based on a set of random regions differing in size and location from the region used to train the policies (Fig. S2).

Supplementary Figures

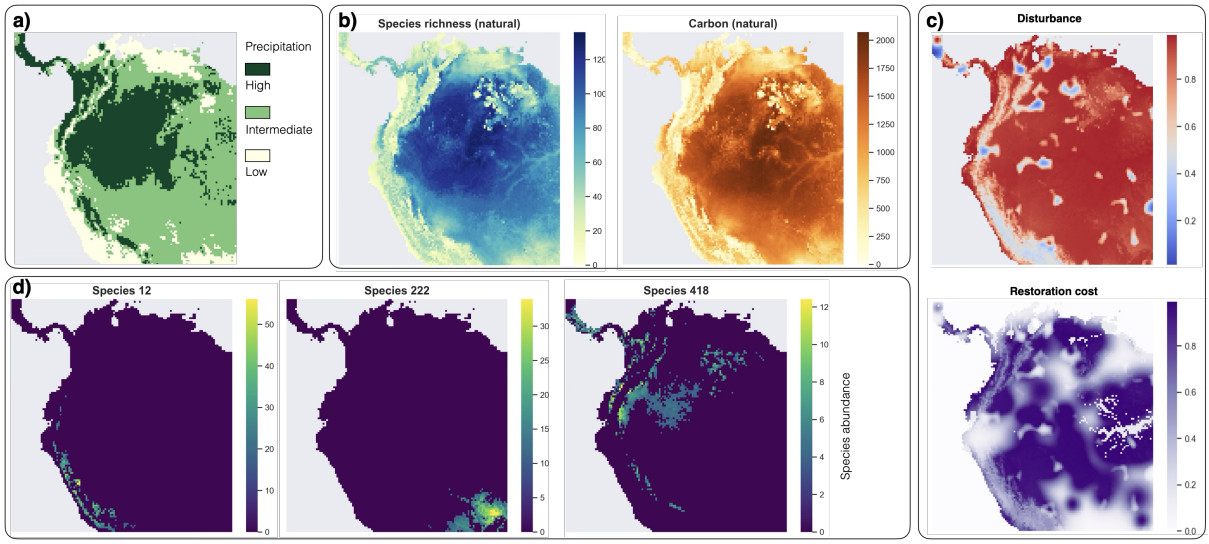

Figure S1: Maps of the physical properties of the arbitrary area used for model training, based on the topography and climate of a portion of Central and South America (a). The area was gridded into 128×128 cells and included empirical climate and elevation layers, while the 'natural' state of biodiversity and carbon (b) and the disturbance patterns and costs (c) were simulated and changed randomly across simulations. Species richness and carbon maps are obtained after stacking 500 simulated species (see examples in d). Note that only the map in panel a) is based on real data, while all other the panels show simulated data. Similar simulations were repeated on 100 additional random maps of different size and location (Fig. S2.

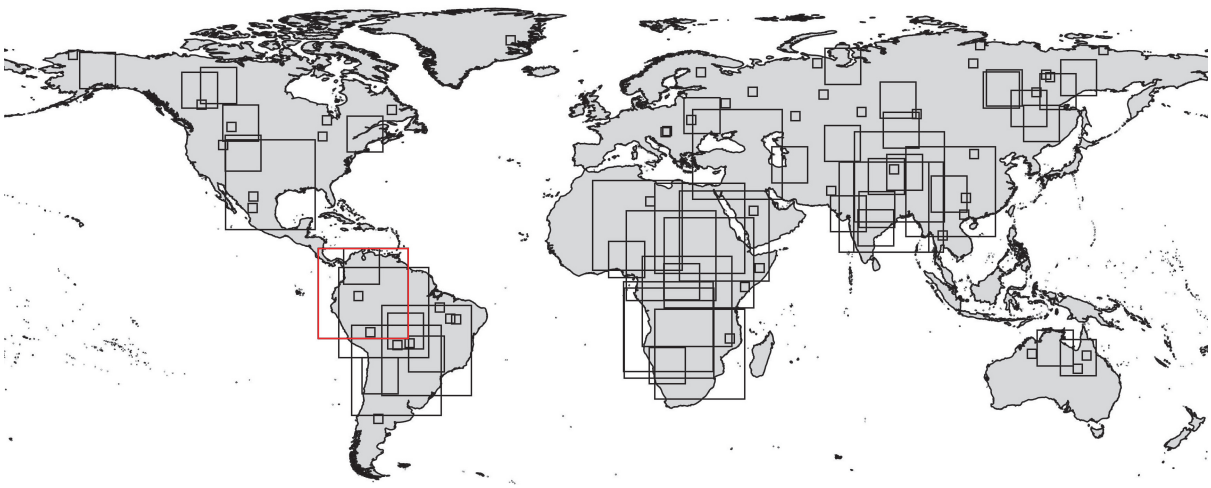

Figure S2: Maps showing the 100 randomly sampled rasters used as the basis for our simulations. The raster highlighted in red indicates the region used to train the models and run the simulations based on a fixed geographic extent (see Fig. S1 and *Methods*).

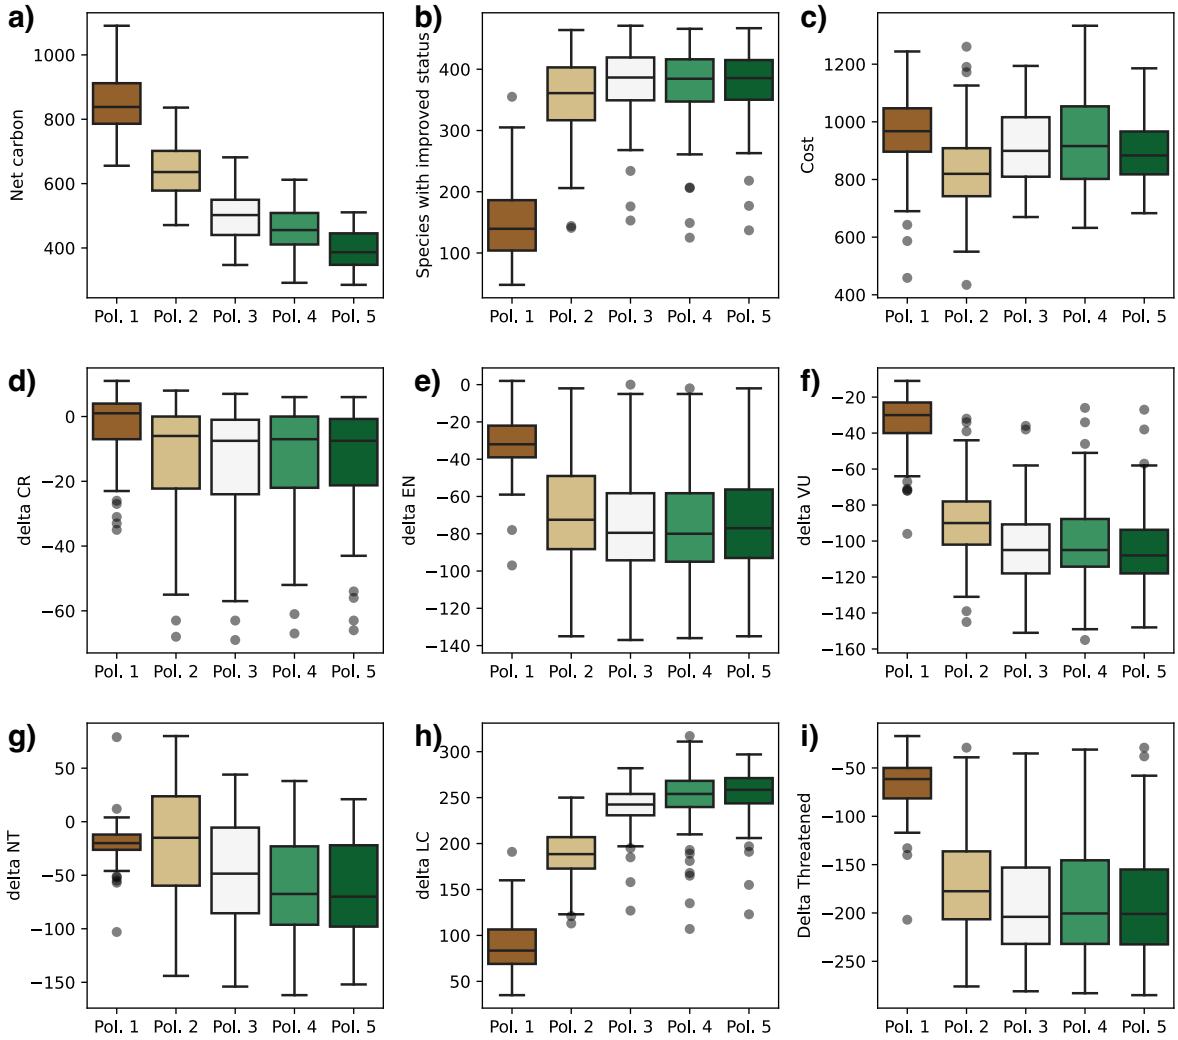

Figure S3: Summary statistics of key outcomes from 100 simulations run under five trained restoration models, Policies 1–5. The net carbon (a) represents the net gain in amount of carbon (here expressed in arbitrary units – see Methods) in the system after 30 time-steps during which 3,000 cells were restored. As the weight assigned to the carbon reward was lowered from 0.5 to 0 in favor of the biodiversity reward, the net carbon gain gradually decreased, although all restoration policies resulted in a net gain. In contrast, the number of species with improved conservation status, i.e. lower extinction risk, drastically increased as soon as a non-zero weight was assigned to the biodiversity reward (Policies 2–5; b). The implementation cost was similar across all policies (c), reflecting the fact that all models were trained with the same budget constraint. Plots d–h) show the difference between the number of species in each extinction risk category, going from highest risk (CR) to lowest (LC), at the beginning of the simulation and the corresponding number at the end of the simulation. Thus, negative numbers, here observed in d–g) mean that the number of species in the corresponding risk categories decreased. Positive numbers, here mostly observed in h), mean that the number of species in these categories increased. These plots show that all policies that considered biodiversity rewards in the optimization (Policies 2–5) resulted in substantial reductions of threatened species compared to a policy targeting carbon only. This is also evident when lumping together all threatened species (i).

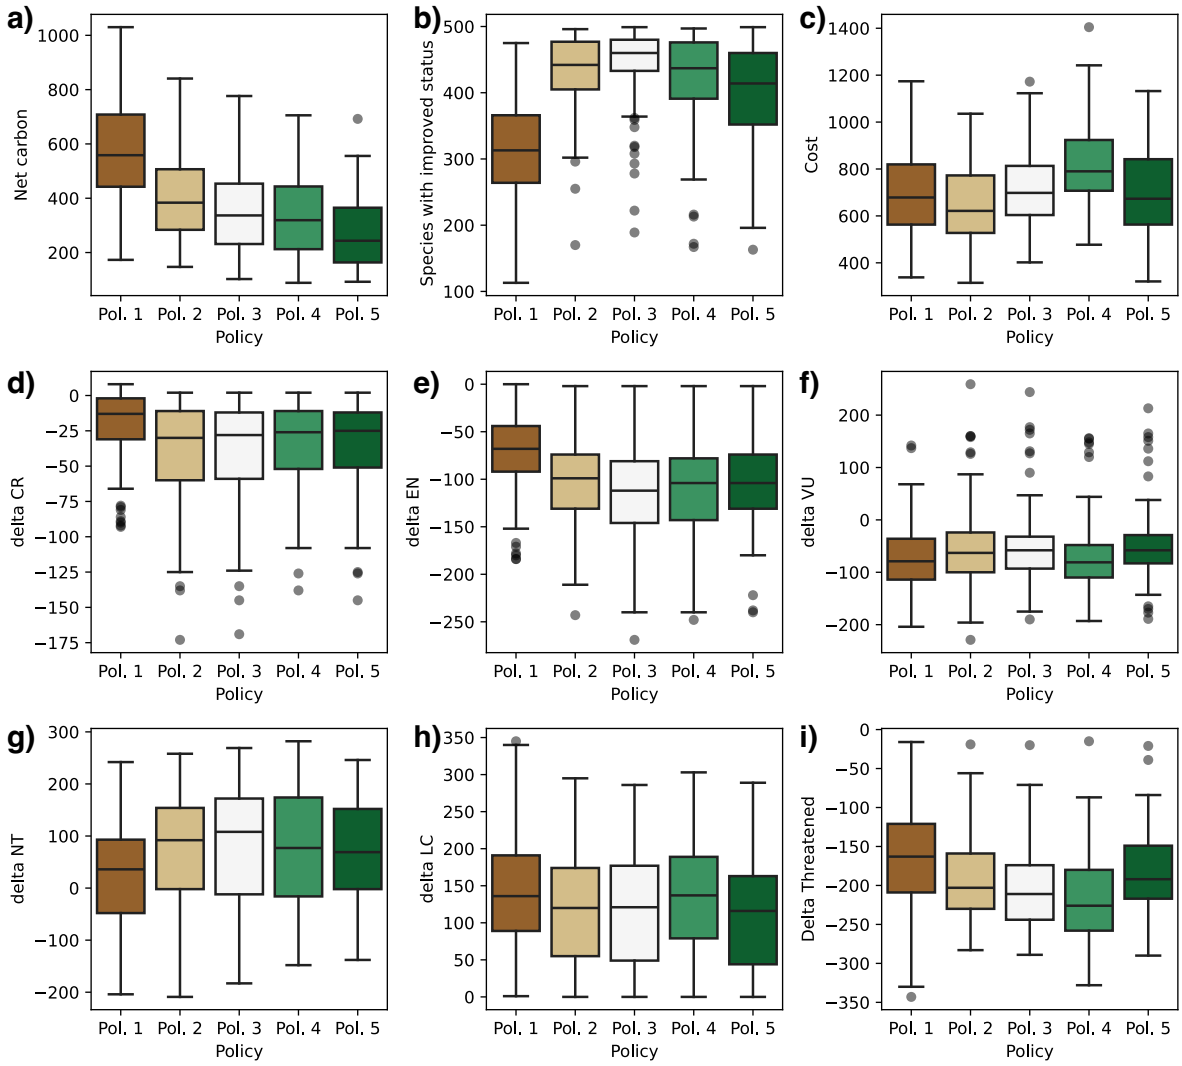

Figure S4: Summary of 100 simulations from random locations and variable resolution (Fig. S2) run under five trained restoration models, Policies 1–5. The general trends in terms of net carbon gain (decreasing from Policy 1 to 5; a), number of species with improved conservation status (higher from Policies 2–5; b), and costs (roughly constant across all policies) are similar to those observed across the simulations based on a fixed geographic extent, upon which the models were optimized. However the patterns are less consistent when looking at changes in species number across risk category (from CR to LC; d–h) or moving from a threatened to a non-threatened classification (i). These results suggest that the models (here trained on a single region) could be re-optimized or fine-tuned to best fit different regions.