Using artificial intelligence to optimize ecological restoration for climate and biodiversity

DANIELE SILVESTRO[1,2,3,4], STEFANO GORIA[3,5], E-PING RAU[2], RENATO AUGUSTO FERREIRA DE LIMA[6], BEN GROOM[7,8], PIOTR JACOBSSON[9] THOMAS STERNER[3,10], ALEXANDRE ANTONELLI[2,3,11,12,13]

[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] *Department of Biological Sciences, ESALQ, Universidade de São Paulo, São Paulo, Brazil*
[7] *LEEP Institute, Department of Economics, University of Exeter Business School, Exeter University, Exeter, UK*
[8] *Grantham Research Institute on Climate Change and the Environment, London School of Economics, London, UK*
[9] *Department of Biology, University of Fribourg, Fribourg, Switzerland*
[10] *Department of Economics, University of Gothenburg, Gothenburg, Sweden*
[11] *Royal Botanic Gardens, Kew, Richmond, UK*
[12] *Antonelli Foundation for Biodiversity Research and Conservation, Nova Friburgo, Brazil*
[13] *Department of Biology, University of Oxford, Oxford, UK*

Correspondence: daniele[at]captaintech.ai, a.antonelli[at]kew.org

Abstract

The restoration of degraded ecosystems is critical for mitigating climate change and reversing biodiversity loss. Depending on the primary objective – such as maximizing carbon sequestration or protecting threatened species – and within the boundaries of budget constraints, different spatial priorities have been identified at global and regional scales. 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 among 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 to allow a robust evaluation of biodiversity and climate outcomes. We find through a series of simulations that even low-to-moderate consideration of biodiversity in restoration projects leads to the selection of restored areas that disproportionately improve the conservation of threatened species, while resulting in a relatively smaller total amount of carbon captured. We propose a data-driven valuation of biodiversity credits in relation to carbon credits, enabling the design of a bundled financial model that could support restoration efforts even in areas previously excluded for economic reasons. Applying our methodology to plant diversity in the Atlantic Forest of eastern South America, one of the most biodiverse and threatened ecosystems globally, we demonstrate its practical utility in guiding real-world restoration and quantifying the essential trade-offs between climate and nature outcomes. Our study provides a robust, scalable methodological pathway to optimize the outcomes of restoration efforts for climate and nature.

Introduction

Ecological restoration is one of the most powerful nature-based solutions to tackle the twin planetary crises of biodiversity loss and climate change[1]. In recognition of its importance, 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 make many other contributions to people and the natural environment[2,3].

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[4], there are clear trade-offs when designing restoration projects aiming to maximize carbon capture or biodiversity recovery[5]. For example, the planting of exotic trees across Africa, although they grow quickly and capture large amounts of carbon (in principle helping to mitigate

climate change, if the biomass produced is retained long-term), has led to detrimental effects on native biodiversity[6]. Trade-off effects have been widely discussed in the literature e.g.[7,8], 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 to be in the range of USD 598-824 billion annually[9] – 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[10,11]. 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; iii) link our predictions to the emerging concept of biodiversity credits alongside carbon credits to design financially viable restoration programs; and iv) demonstrate our optimization framework on a large empirical dataset from the Atlantic Forest in South America. Our results show that identifying trade-offs between carbon and biodiversity is crucial for appropriately valuing biodiversity credits and for incentivizing 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)[12]. The implementation is based on 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 and 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 other element of the CAPTAIN framework is an *agent* (i.e. a virtual representation of a manager or social planner) who monitors the environment and implements a restoration program pursuing one or multiple pre-defined restoration targets. The action performed by the agent consist in the selection of areas within which the disturbance is reduced and natural restoration can occur. the agent uses a deep neural network to map the current state of the environment to a selection of areas for restoration. The environment evolves over time through a number of steps (e.g. years, here set to 30) during 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. The intervention of the agent is assigned a reward based on its outcome, for instance its contribution to preventing species extinctions, and models are optimized to maximize the reward based on the outcomes of the agent actions.

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;[13] 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*[14], using differently weighted rewards that ranged from one focusing exclusively on carbon (Policy 1) to one targeting only biodiversity outcomes (Policy 5), with three models exploring intermediate solutions. All policies shared the same weight on costs, thus implying equal budget constraints available for project implementation. 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. S2-S3), all datasets differed in the distribution of simulated species, threats, and costs, in order to cover a wide range of possible real-world scenarios. Consequently, the placement of the selected restored areas varied among them. However, at the level of each simulation, the policies clearly prioritized different areas (Fig. S1).

Policy 1 favored the restoration of areas that resulted in the highest net carbon gain (24% on average after 30 times steps; Fig. 1a). 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. S1). The result was a more modest carbon gain (just above 11% on average; Fig. 1a) but a substantially stronger reduction of the extinction risk compared with Policy 1. For instance, the average number of Critically Endangered species at the end of the simulation was 36% lower under Policy 5 compared to Policy 1 (Fig. 1b). Similarly, the number of species classified as Endangered decreased on average by 73% when giving priority to biodiversity over carbon. These experiments also showed that while carbon accumulated through time following near-linear trajectories, extinction risks followed more complex paths (Fig. 1b-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, their sensitivity to disturbance, and connectivity among suitable areas.

The distributions of outcomes at the end of the 30 time-step restoration show a substantial level of consistency across 100 simulations (Fig. S4). 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. 2a). 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. 1b, 2a). 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 2b), while Policy 1 resulted in significantly worse impacts on extinction risks compared to most other policies (Figs. 2, S4b). 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 rose to 71–76% (354–378 species) under all other policies (Fig. S4b). 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. S4d). 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. S4i), 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. 2c). This reflected the fact that all policies were run under the same budget constraints.

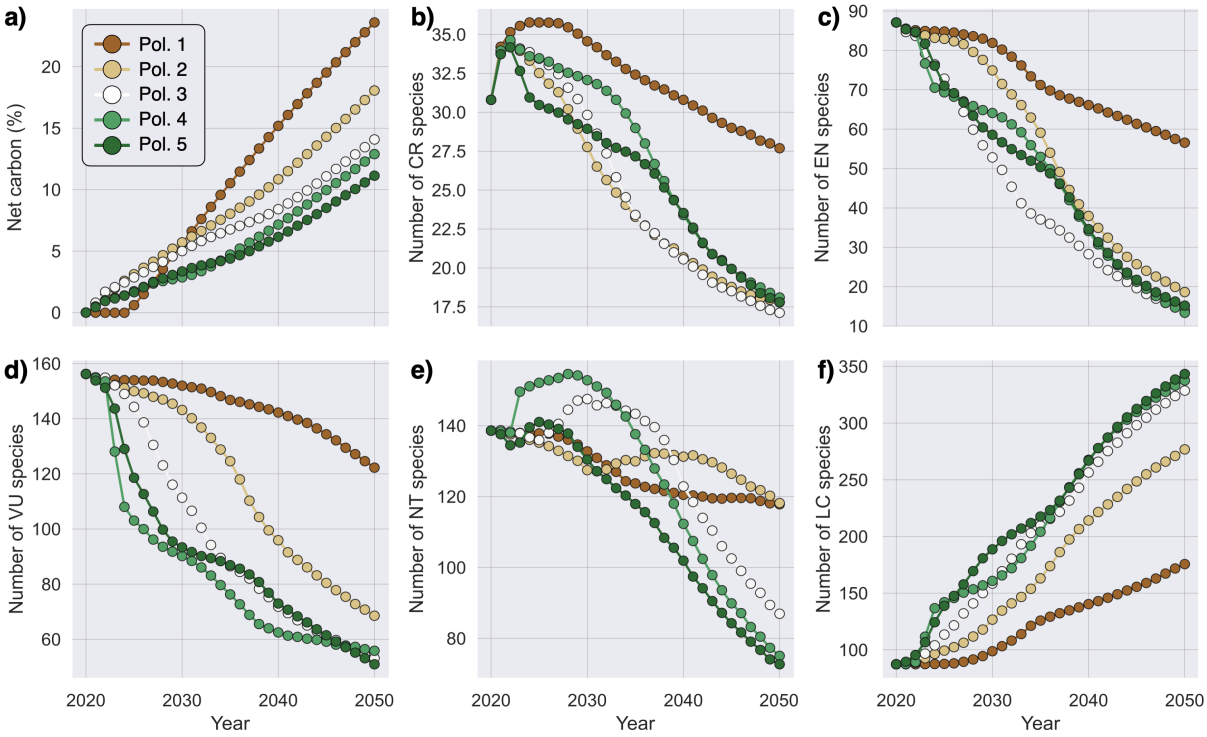

Figure 1: Accumulation of carbon and evolution of extinction risks across 30 time-steps averaged over the 100 simulations. The results are shown for Policies 1–5, which differ in the relative weight assigned to carbon sequestration (highest in Policy 1) and reduction of extinction risks (highest in Policy 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 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

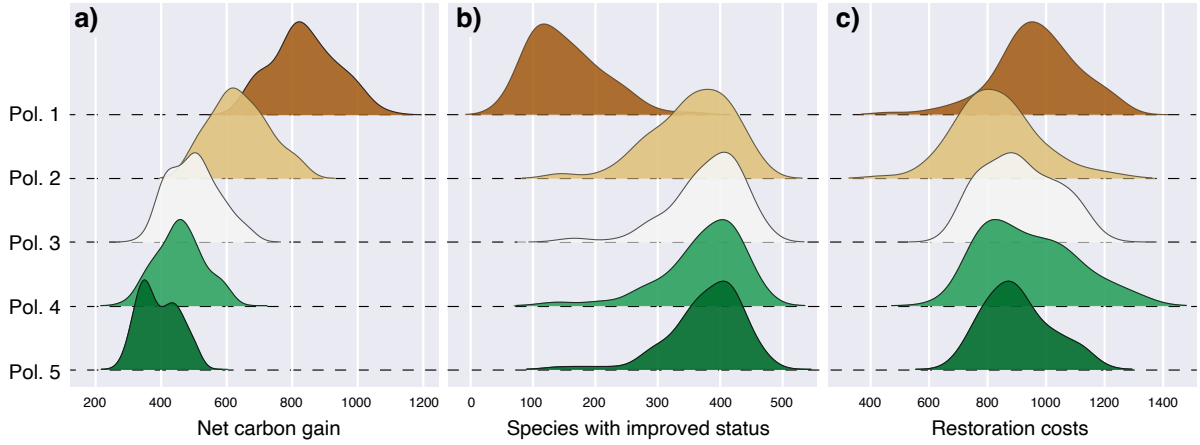

Figure 2: Summary statistics of key outcomes from 100 simulations run under five trained restoration models, Policies 1–5. The net carbon (a) represents the net gain in amount of carbon (here expressed in arbitrary units – see Methods) in the system after 30 time-steps during which 3,000 cells were restored. 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.

credits such that the cumulative value of carbon credits accumulated with Policy 1 matched exactly its restoration costs (Eq. 14; Fig. 3a), Under this definition of carbon credits, all other policies resulted in a net financial loss that ranged from -10% on average in Policy 2 to about -50% in Policy 5 (Fig. 3a).

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. While certain credit schemes are based on area units, such as one credit equaling one hectare, here we adopted a more qualitative approach focused on actual biodiversity outcomes for species. In this simple definition, for demonstration, one credit corresponds to one species moving from a threatened category before restoration to a non-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 financial value of a biodiversity credit. We therefore explored different valuations relative to a carbon credit. When setting the value of one biodiversity credit (β) equal to one carbon credit (χ), 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. 3b). Increasing the relative value of biodiversity credits

($\beta = 2\chi$ and $\beta = 4\chi$) led to improved financial outcomes of all policies (Fig. 3c, d). Yet, we found that Policy 2 consistently led to the best outcomes, reflecting the non-linear improvement in biodiversity outcomes observed in Fig. S4b. 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 3a).

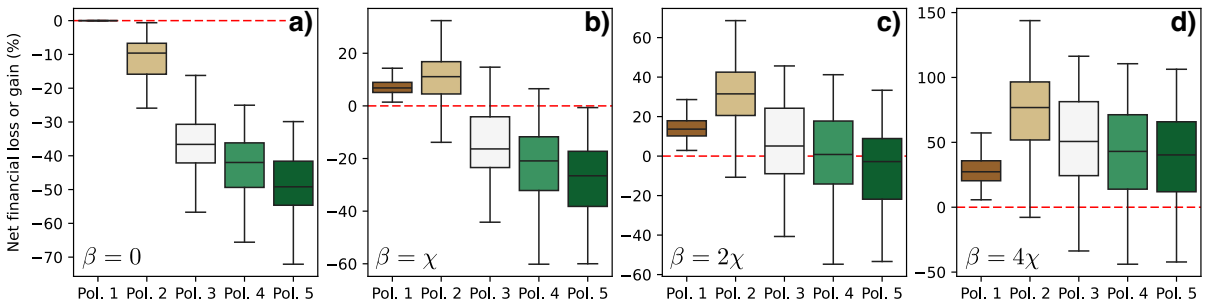

Figure 3: Financial outcomes of restoration policies (net percentage loss or gain, i.e. Φ in Equation 15) under different valuations of biodiversity credits. Negative values indicate a net economic loss where restoration costs exceed the cumulative value of carbon and biodiversity credits, while positive 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 ($\Phi > 0$).

Restoration priorities and trade-offs in the Atlantic Forest

We demonstrated our multi-objective reinforcement learning model on an empirical dataset from the Atlantic Forest – one of the world's most biodiverse and threatened ecosystems[15]. Our analysis was based on a recently published dataset[15], which included $> 3,700$ species of trees and shrubs. The data included species distribution models from 900,000 curated occurrences, an

approximation of per-species above-ground carbon and extinction risk based on the IUCN Red List classification scheme. We also incorporated a proxy for current disturbance levels based on land use, restoration costs based on accessibility, elevation, and value of agricultural production (see *Methods and Supplementary Information*; Fig. S9).

We ran three models based on Policies 1, 3, and 5 described above. Simulation experiments carried out outside of the region used for model training showed that the optimized policies can be applied to different contexts with reasonable performance (Fig. S1, S5). However, our results also suggest that models achieve higher and more consistent performance when applied within the region used in their optimization. For this reason, we fine-tuned the simulation-trained models with the Atlantic Forest data, to identify the 30% with highest restoration priority (i.e. 4,800 of the 16,263 spatial units considered here.

The selected areas showed substantial differences among policies highlighting partial mismatches between predicted restoration potential in terms of carbon and biodiversity outcomes (Fig. 4a). Specifically, 63% of the areas selected under a carbon-focused policy overlapped with those selected under a policy that accounted for both carbon and biodiversity, while the overlap dropped to 34.5% when comparing against a policy focused only on biodiversity outcomes. Meanwhile, policies 3 and 5 overlapped by 53.3%. The predicted outcomes reflected the objectives the models had been optimized for, showing clear tradeoffs between the potential for carbon sequestration through restoration and the ability to protect species and reduce extinction risks. Using species-specific proxies for above-ground carbon and calibrating the average carbon per spatial unit to 98.5 tonnes per hectare (Mg C / ha) in the Atlantic Forest[16], we found that Policy 1 led to an average predicted net carbon sequestration ranging from 69.3 to 70 Mg C / ha, across 30 runs (Fig. 4b). For comparison, Policies 3 and 5 led to a predicted carbon sequestration of 54.3–58.8 and 30.1–36.9 Mg C / ha, respectively, thus leading to a decrease in the order of 20 to 50%.

In contrast, areas selected under Policy 1 left 42–55 species under no protection, a number that dropped to 3–16 under Policy 3 and to 2–8 under Policy 5 (Fig. 4c). Additionally, Under Policy 1, between 15 and 38 of the 420 Critically Endangered and Endangered species included in the dataset (391–453 across imputations of missing data), remained in these high risk categories at the end of the simulation. This indicates that the selected areas failed to restore a sufficient fraction of the species suitable habitats (set to 25% for CR species and 10% for EN species) required to move to a lower risk category under out thresholds (Table 1). The median number of high risk species dropped to 0 under Policy 3, albeit with some degree of variability among runs (0–4), which reduced further under Policy 5, where at most 1 species was left in that group.

The total cost of Policy 3 was about 1% lower than that of Policy 1 and 16% lower than Policy 5. However, all policies showed that that the cost of selected cells cover most of the range of affordable cells (that is, excluding urban areas; Fig. 4d). This indicates that the restoration

areas optimized by CAPTAIN still included expensive cells despite a negative reward associated with implementation costs.

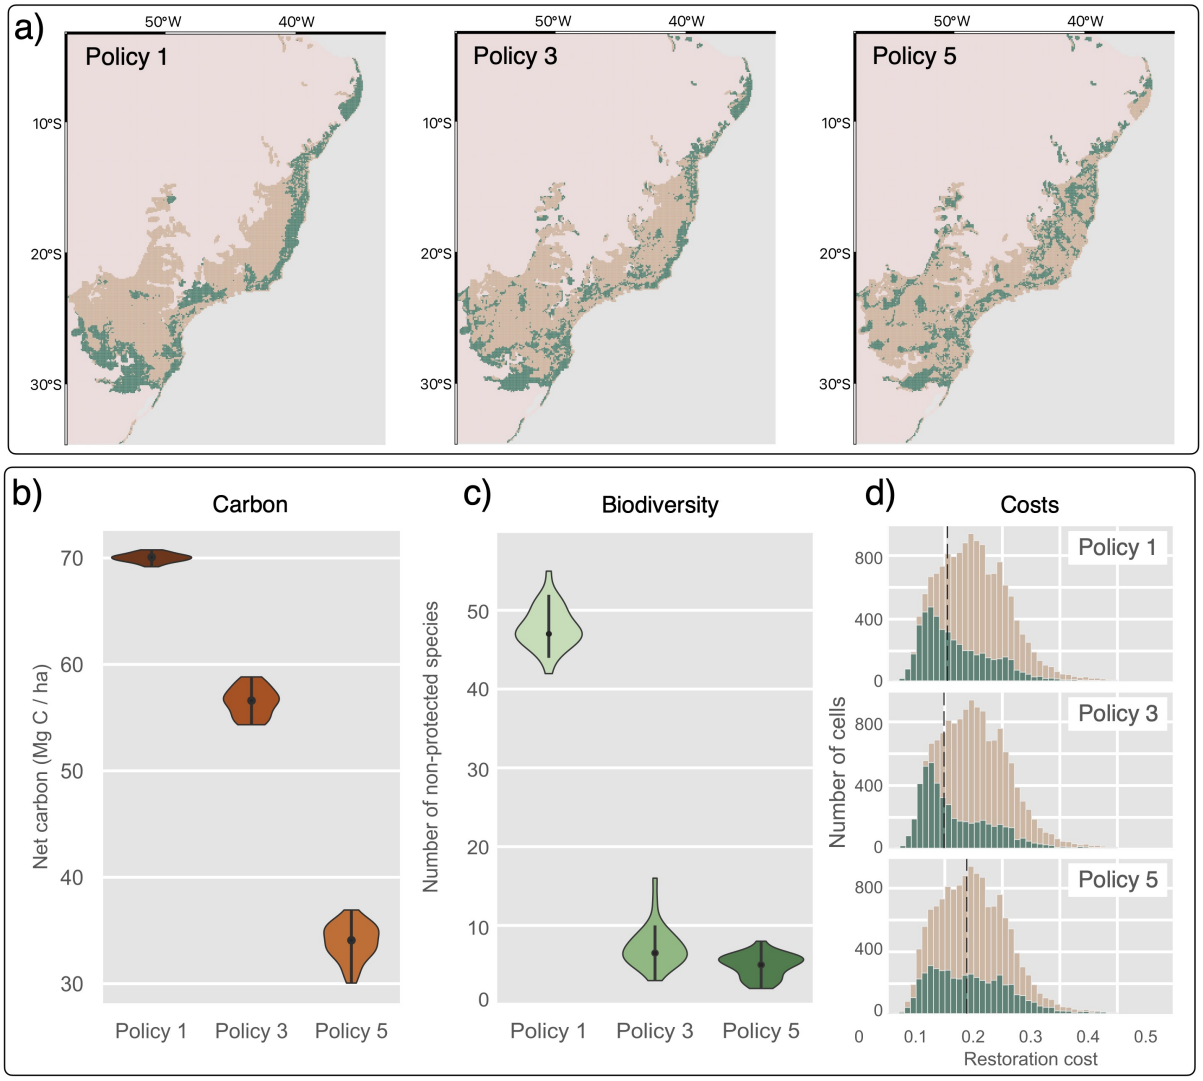

Figure 4: Results of CAPTAIN optimizations in the Atlantic Forest based on a dataset including habitat suitability maps for 3,789 species of trees and shrubs, and proxies for costs and land use. The data are analyzed under three policies (as defined in Table 2) that where focused on carbon (Policy 1), reduction of extinction risks (Policy 5) or both (Policy 3). All models also used a negative reward for costs and were constrained to protect the same number of grid cells (i.e. 30% of the total). The selected areas, after fine-tuning the simulation-trained models with the empirical data, are only partly overlapping, indicating substantial tradeoffs among priorities (a). The predicted outcomes also vary substantially among policies in terms of expected amount of net carbon sequestration (b) and number of species occurring only outside of protected areas (c). All policies selected cells with a wide range of costs (d)

Discussion

Reinforcement learning as a tool for guiding restoration

We demonstrated the power 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[17] between net carbon gain and reduction in extinction risks obtained through restoration, while keeping the extent of restored area constant. This simulation and optimization framework allows us to evaluate the synergies and trade-offs among restoration policies with different objectives.

As our models are trained within a biologically and environmentally dynamic environment, the solutions incorporate the effects of connectivity among restored areas, which favor migration processes and range expansion underlying the recovery of our simulated species. They also account for species-specific features such as climatic niche, potential for carbon sequestration, and sensitivity to anthropogenic disturbance. The simulations show that the spatial distribution of biodiversity and carbon priorities differ, as observed in real-world scenarios[4,18]. 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.

Our analysis of empirical data from the Brazilian Atlantic Forest confirms the trends observed with simulated data. A restoration program optimized exclusively for carbon sequestration fails to adequately protect biodiversity, even when based on natural restoration as opposed to monoculture plantations. We have not considered monocultures in our study, given their well-known harmful impact on the environment (e.g.,[6]) and since ecological restoration projects in Brazil no longer use a single species. On the other hand, including the reduction of extinction risks into the optimization will typically also lead to a reduction in net carbon gain. This is because threatened species are not necessarily the fastest-growing ones (such as the giant timber trees in the genus *Cariniana*, which take decades to reach maturity) or those with largest biomass (such as the many treelets in the genus *Myrcia*).

While no solution can thus achieve maximum levels of biodiversity outcomes and carbon sequestration simultaneously, our empirical analysis for the Atlantic Forest shows that the loss of biodiversity protection inflicted by focusing only on carbon would be substantially larger (ca. 7 times more species under no protection) than the reduction of carbon sequestration (ca. 25% decrease) due to the incorporation of biodiversity outcomes (see Policies 1 vs 3 in Fig. 4B-C). The results also indicate that the incorporation of carbon outcomes leads to only a minor loss of biodiversity benefits (Policies 3 vs 5 in Fig. 4B-C).

Our multi-objective reinforcement learning framework provides a robust tool to explore tradeoffs between the arguably most important objectives of ecological restoration: carbon

sequestration and biodiversity uplift, allowing for the optimization of these potentially conflicting metrics. However, we acknowledge that final decisions of where restoration efforts will occur depends on additional considerations other than the ones explored here - such as the local interests and support of Indigenous and local communities, political commitments at national and regional levels, availability of seedling supply, accessibility, human capacity, and other resources.

Toward a data-driven valuation of biodiversity credits

With our simulation experiments, we propose that a 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 sequestration 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. In principle, this could mean that high-integrity carbon credits, sometimes marketed as carbon credits with co-benefits and sold at a premium cost, could 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 global 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. We chose this definition since we considered a species-based metric to be more appropriate than an area-based one for the calculation of biodiversity uplifts core to the concept of biodiversity credits: "a measured and evidence-based unit of positive biodiversity outcome"[19]. 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 use population recovery, geographic range expansion, or improved ecosystem services as metrics[20,21]. 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[22]. In practice, the reference metric may have to be chosen so that it can be monitored accurately and regularly in the real world, for instance, through the re-assessment of extinction risks of selected species[23,24].

Biodiversity credits can be valuated 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[20]. That is because a ton of sequestered carbon provides the same contribution to climate change mitigation wherever it is sequestered (except for potentially losing value through time if the sequestration is not permanent[25]), while the value of the restoration and recovery of one species is linked to where the species occurs and its environmental context. Estimating trade-offs can help us define actionable and context-dependent valuation for *in situ* biodiversity credits. In our simulated datasets, 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.

Application to a real-world scenario

We demonstrated how our framework can be applied to real-world scenarios by analyzing a large empirical dataset from one of the most highly threatened and diverse ecosystems in the world, considered a Global Biodiversity Hotspot – the Atlantic Forest. Our analysis shows that the reinforcement learning approach implemented in CAPTAIN can utilize diverse and multi-dimensional empirical data to identify priority areas for restoration, while balancing potentially conflicting targets and constraints for carbon sequestration, biodiversity outcomes, and implementation costs. Within a 30% restoration scenario in the Atlantic Forest, we identified substantial discrepancies between carbon and biodiversity priorities, pointing to tradeoffs that need to be accounted for to achieve positive change for climate and nature.

Policies focused primarily on carbon sequestration tended prioritize long and cohesive stretches of coastal areas (Fig. 4A). This encompassed most of the Brazilian coast within the tropical belt (between 23 degrees of latitude), probably due to the fast growing rates and larger carbon sequestering potential in those areas. In contrast, a biodiversity-only prioritization (Fig. 4C) identified a more scattered set of areas, probably due to the heterogeneous distribution of threatened species across the Atlantic Forest biome.

While our analysis was performed on empirical data, it is based on simplified assumptions and approximations, which provide opportunity for future improvements. In particular, our use of land use as a proxy for anthropogenic disturbance and its modeled effect of on plant diversity and abundance could be refined by characterizing disturbance patterns (severity, frequency, and projected future temporal trends) within each land use type using high-resolution data, and modeling their effects on plant diversity and abundance based on principles of plant community dynamics. Similarly, restoration costs –here approximated based on accessibility and land use– could be calibrated by empirical estimates of implementation and opportunity costs[26]. Additionally, the simulated restoration process through natural recolonization based on dispersal and growth could be improved based on the findings of the ongoing restoration experiments, that can also inform on how successfully different species can be re-introduced in restored habitats[27–30]. Finally, the objectives and needs of Indigenous and local communities, which are essential for fair and effective conservation of nature[31], should 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.

Conclusion

The controlled simulations and empirical analyses performed in this study demonstrate the power and practical utility of a reinforcement learning framework to select priority areas for ecological restoration that substantially benefit climate and biodiversity. Focusing on a single restoration objective would result in missed opportunities[3], whereas integrating both carbon capture and biodiversity conservation in a multi-objective prioritization allows us to predict the outcomes of different restoration policies.

Our data-driven approach can also be used to model the financial viability of restoration efforts, by exploring how costs may be distributed between key funding mechanisms such as carbon credits and biodiversity credits[10]. This in turn can serve as a tool to estimate how a biodiversity crediting system could help achieving both climate and biodiversity targets. Future software implementations could include additional sources, such as payments for ecosystem services (clean water, pollination) that increase over time as well-functioning ecosystems are recovered.

Ecosystem restoration is increasingly recognized as a crucial opportunity in the conservation of nature and stabilization of Earth's climate[1,5]. 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 critical role in guiding society's efforts to tackle the twin crises of biodiversity loss and climate change.

Methods

Reinforcement learning optimization of restoration

We optimized spatial restoration planning developing new functionalities within the CAPTAIN software[12]. 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 sequestered 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)[13]. When working on empirical data, this can be effectively replaced by actual IUCN Red List assessments or model-based predictions[24,32]. 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[12], the features of each cell feed into a fully connected deep neural network with `tanh` activation function[33] 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[34] 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[22], 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[23,35,36]. 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[35,37,38].

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.

The CAPTAIN optimization algorithm trains a neural network policy to maximize the expected reward from protective actions by implementing a hybrid approach that combines parallelized evolution strategies (a genetic strategy algorithm) with concepts from classical policy gradient methods, i.e. a advantage function[12,39,40]. The algorithm operates through a repeated two-phase process (epochs): first, the policy parameters are randomly perturbed and evaluated via parallel environment simulation episodes; second, the results from these parallel runs are aggregated, and the parameters are updated using a stochastic gradient estimate. The parameter update is weighted by the advantage function, which measures the improvement of the current average reward relative to the running reward (weighted average of rewards across different epochs)[12]. This approach allows the policy to be optimized without the computational burden of calculating explicit gradients and enabling parallelization of the optimization task.

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. S2a). The properties include its coastline, elevation and cell-specific climate average values obtained from CHELSA[41], here including the mean annual temperature, mean annual precipitation and the minimum temperature of the coldest month. We used the `gdgtm` python library to generate the initial rasters (see Supplementary Data and Codes). 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. S2a). 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 $\approx$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. S2d.

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. S2c). 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. S2c). 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[5]. 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 calculated 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[42], 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 the initial species extinction risks ($X_{s,t=0}$, only applied in simulated data and replaced with Red List assessment when using empirical data) and their reassessment at time t based on their level of protection (Criterion 1) and population trends (Criterion 2). The re-assessed extinction risk is $X_{s,t} = X_{s,t=0} + c'_{s,t} + c''_{s,t}$.

Initial status		Reassessment at time t			
$X_{s,t=0}$	r	Criterion 1	$(O_{s,t} \cap M_t)/A_s$	Criterion 2	$O_{s,t}/O_{s,t=0}$
LC	>0.95	$c'_{s,t} = 0$	< 0.10	$c''_{s,t} = 0$	< 0.10
NT	0.9-0.95	$c'_{s,t} = -1$	0.10 - 0.25	$c''_{s,t} = 1$	0.10 - 0.25
VU	0.8-0.9	$c'_{s,t} = -2$	0.25 - 0.50	$c''_{s,t} = 2$	0.25 - 0.50
EN	0.6-0.8	$c'_{s,t} = -3$	0.50 - 0.80	$c''_{s,t} = 3$	0.50 - 0.80
CR	<0.6	$c'_{s,t} = -4$	> 0.80	$c''_{s,t} = 4$	> 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[17] 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).

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

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 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.

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[43].

Empirical analysis

We analyzed an extensive dataset of tree biodiversity in the Brazilian Atlantic Forest (AF)[15]. The dataset included occurrence data for thousands species of woody plants in the AF, with per-species information on their maximum height, wood density, and conservation status from IUCN Red List assessments[13].

We compiled spatial data gridded into 16,376 cells at a 10-km resolution within the Atlantic Forest ecoregion (around 1.2×10^6 km^2 in area) as defined by the WWF Global 200 project[44]. We generated species distribution models (SDM) based on cell-specific WorldClim bioclimatic variables[45] and on species occurrence data[15], and generated raster layers representing disturbance and restoration costs based on publicly available geospatial data products. All data pro-

cessing and analyses were run in R 4.3.3, using the packages `concaveman`, `ENMTools`, `flexsdm`, `future`, `fuzzyjoin`, `geodata`, `magrittr`, `sf`, `terra`, `tidyterra`, and `tidyverse`[46–57].

Species distribution models – We obtained the 19 WorldClim historical bioclimatic variables at 5-arc minutes resolution[45], cropped to the Atlantic Forest ecoregion and reprojected to the 10-km resolution. We examined the correlation between the bioclimatic variables to exclude highly collinear variable pairs, retaining the following 7 variables of temperature (°C) and precipitation (mm): 1) annual mean temperature, 2) mean diurnal temperature range (mean of monthly maximum temperature - minimum temperature), 3) annual temperature range, 4) annual precipitation, 5) seasonality of precipitation (coefficient of variation), 6) precipitation of the warmest quarter, and 7) precipitation of the coldest quarter.

Among the retained variables, only variable pairs 1-5 ($r = 0.65$) and 2-7 ($r = 0.61$) exhibited coefficients of correlation larger than 0.6 (Fig. S7): as these are both pairs of a temperature variable and a precipitation variable and they may contain important environmental information despite the correlation, we retained both variable pairs.

We obtained geolocated occurrence data for 4,950 species occurring in the Atlantic Forest[15], and eliminated data points outside of the rectangular bounding box of the Atlantic Forest ecoregion to exclude geographically distant data points. The remaining data set contained in total 722,448 records and 4,948 species. Of these, 3105 species had abundant data points ($n \geq 30$) and 1,631 species had sparse data points ($3 \leq n < 30$); 212 species with insufficient data points ($n < 3$) were excluded from the analysis. To mitigate potential spatial biases, we performed spatial thinning for species with abundant data points, using the `ENMTools::trimdupes.by.raster` function to retain one data point at maximum per pixel. To prevent excessive loss of species information, we only accepted the thinned results if they also contained ≥ 30 data points; otherwise, we retained the original unthinned original data. In total, 1,843 species were thinned, and for the thinned species 80% data points were retained per species on average.

We defined the spatial extent of the training area (from which background points are sampled) for each species as an alpha shape (concavity = 3) surrounding all occurrence data points of the species, buffered by twice the median inter-point distance among all data points and cropped to the Atlantic Forest ecoregion. We sampled 10,000 background points within the training extent of each species[58], assembled the occurrence data points and the background points, and extracted the environmental values at all points to construct the model data. We partitioned model data into training and testing sets for cross-validation, using complete random k-fold partition using the `flexsdm::part_random` function: we used 4-fold partition for the 3,105 species with abundant data points, and 2-fold partition for the 1631 species with sparse data points.

We then fitted an SDM for each species: for the 3,781 species with ≥ 15 data points, we fitted a general additive model (GAM) using the `flexsdm::fit_gam` function, and for the 955 species

with < 15 data points, we adopted an ensemble of small models (ESM) approach using the `flexsdm::esm_gam` function, which consists of fitting bivariate GAM models with all pair-wise combinations of predictors, and calculating the ensemble average of model outputs weighted by their performance metrics (Somers' D: $D = 2 \times (\text{AUC} - 0.5)$)[59]. For the model output of each species, we calculated the following five performance metrics: sensitivity, specificity, omission rate, true skill statistic (TSS), and area under curve (AUC). We excluded species with low performance (sensitivity < 0.7, specificity < 0.7, TSS < 0.5, and AUC < 0.8), retaining results for 3885 species and excluding 856 species in the end (Figure S8). Finally, for the retained species, we generated geographical predictions of continuous suitability scores maps over the Atlantic Forest ecoregion using the `flexsdm::sdm_predict` function.

Disturbance and restoration costs – We defined the disturbance level within each 10-km pixel of the Atlantic Forest ecoregion based on its land cover. We obtained the MapBiomas Collection 10 - Land Cover and Use dataset[60], which classifies each 30-m pixel within Brazil into one of 38 land cover classes. In addition, we also obtained the MapBiomas Collection 9 - Age of Secondary Vegetation maps in 2023[61], and created a mask of secondary forests representing pixels where the age of secondary vegetation > 0. We cropped the datasets to the Atlantic Forest ecoregion, reprojected to the 10-km resolution, and calculated the number of 30-m pixels belonging to each land cover class within each 10-km pixel, additionally classifying pixels within the secondary forest mask into a novel class of "secondary forest" (ID: 76). We reclassified the MapBiomas classes into the following 9 broad classes:

- Forest (IDs: 1, 3, 5, 6, 49),
- Grassland (4, 12, 29, 50),
- Secondary forest (76),
- Plantation (9),
- Cultivation (14, 15, 18, 19, 20, 21, 31, 35, 36, 39, 40, 41, 46, 47, 48, 62),
- Coastal area (23),
- Water bodies (11, 32, 26, 33),
- Urban area (24, 25, 30, 75), and
- Others.

We calculated the proportion of each broad class within each 10-km pixel. We then calculated the average disturbance score of each 10-km pixel by assigning a disturbance score to each broad class, namely, forest and grassland: 0, secondary forest: 0.3, plantation: 0.5, cultivation: 0.7, coastal area: 0.5, water bodies: 1, urban area: 1, others: 0.5. We then calculated the weighted average of all disturbance scores weighted by the proportion of the broad classes.

We defined the restoration costs using three data sources: (1) accessibility and (2) elevation, which are proxies of implementation cost, and (3) value of agriculture production, which is a proxy of opportunity cost.

- For accessibility, we retrieved data of travel time to cities with more than 5,000 inhabitants (`geodata::travel_time`)[62].
- For elevation, we retrieved global elevation data (`geodata::elevation_global`) that are based on the Shuttle Radar Topography Mission (SRTM) 90-m resolution data[63] and supplemented by the GTOP30 dataset for high latitudes[64].
- For the value of agricultural production, we retrieved SPAM data of aggregated crop production value (USD) per area (`geodata::crop_spam`).

SPAM pixels with empty values were filled in by interpolation from edges using the GDAL script `gdal_fillnodata`, with the `md` (maximum number of pixels to be searched for values to interpolate) set to 10. All three data layers were cropped and reprojected using the GDAL command `gdalwarp`, resampled to the 10-km resolution with bilinear interpolation using the `terra::resample` function, and then scaled with min-max scaling to the range of $[0, 1]$. We then calculated restoration costs as the arithmetic average of the three scaled variables.

Species traits – We included species-specific features in our analysis, allowing us to incorporate extinction risks, proxies for carbon sequestration potential, and life-history traits. From the original data[15] we obtained growth habit (categorized as shrub, small tree, and large tree), which we used to approximate relative values of growth rates. We used the same Weibull distribution used in the simulated data to draw species-specific growth rates, with scale parameters set to 10, 2, and 0.1 for shrub, small trees, and large trees, respectively. This led to faster average demographic growth rates for shrubs compared to large trees. While we acknowledge that these are approximations, empirical estimates of growth rates, where available, could be used to replace this parameterization.

We randomly drew the values describing species-specific sensitivity to disturbance from the same U-shaped beta distribution used in our previous simulations. To simulate the fact that trees might be more sensitive to land use changes we ranked the randomly drawn sensitivity values and assigned the highest values to large trees, followed by small trees, leaving the smallest values to shrubs. Again, these approximations could be replaced by empirical estimates of sensitivity to disturbance, where available. We expressed the carrying capacity of a species in a cell as a relative number between 0 and 100, scaled by the habitat suitability in the cell, as predicted by the species distribution models. The simulated abundance of a species in a cell was further affected by a mortality rate proportional to the product between disturbance and species sensitivity (Eq. 3).

We used the available Red List assessments as species extinction risk estimation[13]. Here we imputed missing extinction risk classification (i.e. unevaluated or data-deficient species) through random resampling proportionally to the empirical frequencies of each class among assessed species. While this is a crude approximation, more sophisticated approaches could be used to predict extinction risks using machine learning imputation methods[32,65]. We let the extinction risk of species evolve over time as a function of species recovery after restoration using the re-assessment criteria used in our simulations (Table 1).

We used per-species biomass as a proxy for a species' ability to sequester carbon, based on wood density (WD) and maximum height (H)[15] We estimated the maximum diameter at breast height (DBH) for each species from maximum height (H), using the generic model 3 from[66]: $\log(\mathrm{H}) = 1.029 + 0.567 \times \log(\mathrm{DBH})$ from which we obtain $\mathrm{DBH} = \exp((\log(\mathrm{H}) - 1.193)/0.529)$. Next we estimated maximum above-ground biomass (AGB, kg) per tree for each species from wood density (WD), DBH and H, using the pantropical model from[67]: $\mathrm{AGB} = \exp(-29.77 + \log(\mathrm{WD} \times \mathrm{DBH}^2 \times \mathrm{H}))\ 0.0509 \times \mathrm{WD} \times \mathrm{DBH}^2 \times \mathrm{H}$. As in the analysis of simulated data, we estimated the relative above-ground carbon in a cell based on the species-specific AGB estimated as shown above and the simulated abundance of the species in the cells. We then calibrated this relative quantification of carbon per cell, by rescaling them such their average equalled 98.5 tonnes per hectare, based on previous empirical estimates for the Atlantic Forest[16]. We then used the difference between the initial carbon and the total carbon after restoration as a measure of the sequestration generated by the restoration policy.

Model fine-tuning and predictions – We fine-tuned the models by running another training procedure for 1100–1700 epochs starting from the parameter values from the models trained on simulated data (Policies 1, 3, and 5). The additional training was based on a batch of 6 parallel environments, each containing a random subset of 1,000 species to speed up the run time.

We then used the fine-tuned model on the full set of species to identify restoration priorities. As in the analyses of simulated data, we set the models to select 4,800 grid cells (i.e. approximately 30%) within 30 time steps. Thus, at each time step 160 cells were selected for conservation, and the disturbance in protected cells was set to 0, allowing existing species in the cell to reach their carrying capacity and species from neighboring cells to recolonize them based on habitat suitability and a dispersal process modeled as in our previous simulations.

As some of the parameters in our simulations were random drawn or imputed (namely, the extinction risks for species un-assessed species, the growth rates, and sensitivity to disturbance), we repeated the restoration predictions 30 times, resampling each time these parameters. We then summarized the results by taking ranking cells based on how frequently they were selected for restoration and took the top 4,800 cells as the final outcome.

Data and code availability

All the codes and data used to reproduce the results shown in this study are included in the Supplementary Materials and will be uploaded to a public and permanent repository with a DOI in zenodo.org upon acceptance of the paper.

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. and E-P.R. are 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.

Generative AI statement

Generative AI was not used in producing the contents of this article.

*

References

[1] Nathalie Seddon, Alexandre Chausson, Pam Berry, Cécile AJ Girardin, Alison Smith, and Beth Turner. Understanding the value and limits of nature-based solutions to climate change and other global challenges. *Philosophical Transactions of the Royal Society B*, 375(1794):20190120, 2020.

[2] Alice Di Sacco, Kate A Hardwick, David Blakesley, Pedro HS Brancalion, Elinor Breman, Loic Cecilio Rebola, Susan Chomba, Kingsley Dixon, Stephen Elliott, Godfrey Ruyonga, et al. Ten golden rules for reforestation to optimize carbon sequestration, biodiversity recovery and livelihood benefits. *Global Change Biology*, 27(7):1328–1348, 2021.

[3] Pedro HS Brancalion, Fangyuan Hua, Francis H Joyce, Alexandre Antonelli, and Karen D Holl. Moving biodiversity from an afterthought to a key outcome of forest restoration. *Nature Reviews Biodiversity*, pages 1–14, 2025.

[4] Martin Jung, Andy Arnell, Xavier De Lamo, Shaenandhoa García-Rangel, Matthew Lewis, Jennifer Mark, Cory Merow, Lera Miles, Ian Ondo, Samuel Pironon, et al. Areas of global importance for conserving terrestrial biodiversity, carbon and water. *Nature Ecology & Evolution*, 5(11):1499–1509, 2021.

[5] Bernardo BN Strassburg, Alvaro Iribarrem, Hawthorne L Beyer, Carlos Leandro Cordeiro, Renato Crouzeilles, Catarina C Jakovac, André Braga Junqueira, Eduardo Lacerda, Agnieszka E Latawiec, Andrew Balmford, et al. Global priority areas for ecosystem restoration. *Nature*, 586(7831):724–729, 2020.

[6] Catherine L Parr, Mariska Te Beest, and Nicola Stevens. Conflation of reforestation with restoration is widespread. *Science*, 383(6684):698–701, 2024.

[7] Karen D Holl and Pedro HS Brancalion. Tree planting is not a simple solution. *Science*, 368(6491):580–581, 2020.

[8] Fangyuan Hua, L Adrian Bruijnzeel, Paula Meli, Philip A Martin, Jun Zhang, Shinichi Nakagawa, Xinran Miao, Weiyi Wang, Christopher McEvoy, Jorge Luis Peña-Arancibia, et al. The biodiversity and ecosystem service contributions and trade-offs of forest restoration approaches. *Science*, 376(6595):839–844, 2022.

[9] Andrew Deutz, Geoffrey M Heal, Rose Niu, Eric Swanson, Terry Townshend, Li Zhu, Alejandro Delmar, Alqayam Meghji, Suresh A Sethi, and John Tobin-de la Puente. Financing nature: Closing the global biodiversity financing gap. *The Paulson Institute, The Nature Conservancy, and the Cornell Atkinson Center for Sustainability*, 256, 2020.

[10] Alexandre Antonelli, Ximena Rueda, Robert Calcagno, and Pauline Nantongo Kalunda. How biodiversity credits could help to conserve and restore nature. *Nature*, 634(8036): 1045–1049, 2024.

[11] Martine Maron, Amrei von Hase, Fabien Quétier, Laura J Sonter, Sebastian Theis, and Sophus OSE zu Ermgassen. Biodiversity offsets, their effectiveness and their role in a nature positive future. *Nature Reviews Biodiversity*, pages 1–14, 2025.

[12] Daniele Silvestro, Stefano Goria, Thomas Sterner, and Alexandre Antonelli. Improving biodiversity protection through artificial intelligence. *Nature sustainability*, 5(5):415–424, 2022.

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

[14] Richard S Sutton, Andrew G Barto, et al. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.

[15] Renato AF de Lima, Gilles Dauby, André L de Gasper, Eduardo P Fernandez, Alexander C Vibrans, Alexandre A de Oliveira, Paulo I Prado, Vinícius C Souza, Marinez F. de Siqueira, and Hans Ter Steege. Comprehensive conservation assessments reveal high extinction risks across atlantic forest trees. *Science*, 383(6679):219–225, 2024.

[16] Kauane Maiara Bordin, Adriane Esquivel-Muelbert, Joice Klipel, Rayana Caroline Picolotto, Rodrigo Scarton Bergamin, Ana Carolina da Silva, Pedro Higuchi, Elivane Salete Capellesso, Márcia Cristina Mendes Marques, Alexandre F Souza, et al. No relationship between biodiversity and forest carbon sink across the subtropical brazilian atlantic forest. *Perspectives in Ecology and Conservation*, 21(2):112–120, 2023.

[17] Xi Lin, Hui-Ling Zhen, Zhenhua Li, Qing-Fu Zhang, and Sam Kwong. Pareto multi-task learning. *Advances in neural information processing systems*, 32, 2019.

[18] Camilo Alejo, Amy Luers, Andrea Ventimiglia, and H Damon Matthews. Maximizing nature-based solutions using artificial intelligence to align global biodiversity, climate, and water targets. *bioRxiv*, pages 2025–01, 2025.

[19] Biodiversity Credit Alliance. Definition of a biodiversity credit. *Issue Paper*, (3): Available at: www.biodiversitycreditalliance.org/wp–content/uploads/2024/05/ Definition–of–a–Biodiversity–Credit–Rev–220524.pdf, 2024.

[20] Keith Douglas Farnsworth, Adewale Henry Adenuga, and RS De Groot. The complexity of biodiversity: A biological perspective on economic valuation. *Ecological Economics*, 120:350–354, 2015.

[21] Aafke M Schipper, Jelle P Hilbers, Johan R Meijer, Laura H Antão, Ana Benítez-López, Melinda MJ de Jonge, Luuk H Leemans, Eddy Scheper, Rob Alkemade, Jonathan C Doelman, et al. Projecting terrestrial biodiversity intactness with globio 4. *Global Change Biology*, 26(2):760–771, 2020.

[22]

[23] América P Durán, Jonathan MH Green, Christopher D West, Piero Visconti, Neil D Burgess, Malika Virah-Sawmy, and Andrew Balmford. A practical approach to measuring the biodiversity impacts of land conversion. *Methods in Ecology and Evolution*, 11(8): 910–921, 2020.

[24] Alexandre Antonelli, Kiran L Dhanjal-Adams, and Daniele Silvestro. Integrating machine learning, remote sensing and citizen science to create an early warning system for biodiversity. *Plants, people, planet*, 5(3):307–316, 2023.

[25] Ben Groom and Frank Venmans. The social value of offsets. *Nature*, 619(7971):768–773, 2023.

[26] Pedro H.S. Brancalion, Paula Meli, Julio R.C. Tymus, Felipe E.B. Lenti, Rubens M. Benini, Ana Paula M. Silva, Ingo Isernhagen, and Karen D. Holl. What makes ecosystem restoration expensive? a systematic cost assessment of projects in brazil. *Biological Conservation*, 240:108274, 2019. ISSN 0006-3207. doi: https://doi.org/10.1016/j.biocon.2019.108274. URL `https://www.sciencedirect.com/science/article/pii/S0006320719301934`.

[27] Lívia D Audino, Stephen J Murphy, Ludimila Zambaldi, Julio Louzada, and Liza S Comita. Drivers of community assembly in tropical forest restoration sites: role of local environment, landscape, and space. *Ecological Applications*, 27(6):1731–1745, 2017.

[28] Tao Wang, Lingbo Dong, and Zhaogang Liu. Factors driving native tree species restoration in plantations and tree structure conversion in chinese temperate forests. *Forest Ecology and Management*, 507:119989, 2022.

[29] Miriam Selwyn, Joan Pino, and Josep Maria Espelta. Disentangling the importance of intrinsic and extrinsic seed dispersal factors in forest restoration success: a global review. *Restoration Ecology*, 31(4):e13868, 2023.

[30] Laura HP Simões, Marielos Peña-Claros, Joannès Guillemot, Rens Brouwer, Frans Bongers, Lourens Poorter, Catherine T de Almeida, Danilo RA Almeida, Miguel Cooper, Mathieu Decuyper, et al. Recovery of tree species functional composition in eucalypt plantations with natural regeneration differs among canopy strata. *Forest Ecology and Management*, 594:122952, 2025.

[31] James T Erbaugh, Nabin Pradhan, James Adams, Johan A Oldekop, Arun Agrawal, Dan Brockington, Rose Pritchard, and Ashwini Chhatre. Global forest restoration and the importance of prioritizing local communities. *Nature Ecology & Evolution*, 4(11):1472–1476, 2020.

[32] Alexander Zizka, Tobias Andermann, and Daniele Silvestro. Iucnn–deep learning approaches to approximate species' extinction risk. *Diversity and Distributions*, 28(2): 227–241, 2022.

[33] Andrea Apicella, Francesco Donnarumma, Francesco Isgrò, and Roberto Prevete. A survey on modern trainable activation functions. *Neural Networks*, 138:14–32, 2021.

[34] Neil D Burgess, Natasha Ali, Jacob Bedford, Nina Bhola, Sharon Brooks, Alena Cierna, Roberto Correa, Matthew Harris, Ayesha Hargey, Jonathan Hughes, et al. Global metrics for terrestrial biodiversity. *Annual Review of Environment and Resources*, 49(1):673–709, 2024.

[35] Louise Mair, Leon A Bennun, Thomas M Brooks, Stuart HM Butchart, Friederike C Bolam, Neil D Burgess, Jonathan MM Ekstrom, EJ Milner-Gulland, Michael Hoffmann, Keping Ma, et al. A metric for spatially explicit contributions to science-based species targets. *Nature Ecology & Evolution*, 5(6):836–844, 2021.

[36] Alison Eyres, Thomas Ball, Michael Dales, Tom Swinfield, Andy Arnell, Daniele Baisero, América Paz Durán, Jonathan Green, Rhys E Green, Anil Madhavapeddy, et al. Life: A metric for quantitively mapping the impact of land-cover change on global extinctions. *Earth and Environmental Sciences*, 2024.

[37] C Pimiento, F Leprieur, D Silvestro, JS Lefcheck, C Albouy, DB Rasher, M Davis, JC Svenning, and JN Griffin. Functional diversity of marine megafauna in the anthropocene. sci adv 6: eaay7650, 2020.

[38] Rikki Gumbs, Claudia L Gray, Monika Böhm, Ian J Burfield, Olivia R Couchman, Daniel P Faith, Félix Forest, Michael Hoffmann, Nick JB Isaac, Walter Jetz, et al. The edge2 protocol: Advancing the prioritisation of evolutionarily distinct and globally endangered species for practical conservation action. *PLoS Biology*, 21(2):e3001991, 2023.

[39] Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems*, 12, 1999.

[40] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*, 2017.

[41] Dirk Nikolaus Karger, Michael P Nobis, Signe Normand, Catherine H Graham, and Niklaus E Zimmermann. Chelsa-trace21k–high-resolution (1 km) downscaled transient temperature and precipitation data since the last glacial maximum. *Climate of the Past*, 19 (2):439–456, 2023.

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

[43] MA Drupp, MC Hänsel, EP Fenichel, M Freeman, C Gollier, Ben Groom, GM Heal, PH Howard, A Millner, FC Moore, et al. Accounting for the increasing benefits from scarce ecosystems. *Science*, 383(6687):1062–1064, 2024.

[44] David M Olson and Eric Dinerstein. The global 200: Priority ecoregions for global conservation. *Ann. Mo. Bot. Gard.*, 89(2):199, 2002.

[45] Stephen E Fick and Robert J Hijmans. WorldClim 2: new 1-km spatial resolution climate surfaces for global land areas: NEW CLIMATE SURFACES FOR GLOBAL LAND AREAS. *Int. J. Climatol.*, 37(12):4302–4315, October 2017.

[46] Edzer Pebesma. Simple features for R: Standardized support for spatial vector data, 2018.

[47] Edzer Pebesma and Roger Bivand. Spatial data science: With applications in R, 2023.

[48] Joël Gombin, Ramnath Vaidyanathan, and Vladimir Agafonkin. concaveman: A very fast 2D concave hull algorithm, 2020.

[49] Dan Warren and Russell Dinnage. ENMTools: Analysis of niche evolution using niche and distribution models, 2025.

[50] Santiago J E Velazco, Brooke Rose, André F A Andrade, Ignacio Minoli, and Janet Franklin. flexsdm: An R package for supporting a comprehensive and flexible species distribution modelling workflow, 2022.

[51] Henrik Bengtsson. A unifying framework for parallel and distributed processing in R using futures. *The R Journal*, 13(2):208–227, 2021.

[52] Robert J Hijmans, Márcia Barbosa, Aniruddha Ghosh, and Alex Mandel. geodata: Download geographic data, 2024.

[53] Stefan Milton Bache and Hadley Wickham. magrittr: A forward-pipe operator for R, 2022.

[54] Robert J Hijmans. terra: Spatial data analysis, 2025.

[55] Hadley Wickham, Mara Averick, Jennifer Bryan, Winston Chang, Lucy D'agostino McGowan, Romain François, Garrett Grolemund, Alex Hayes, Lionel Henry, Jim Hester, Max Kuhn, Thomas Lin Pedersen, Evan Miller, Stephan Milton Bache, Kirill Müller, Jeroen Ooms, David Robinson, Dana Paige Seidel, Vitalie Spinu, Kohske Takahashi, Davis Vaughan, Claus Wilke, Kara Woo, and Hiroaki Yutani. Welcome to the tidyverse, 2019.

[56] Diego Hernangómez. Using the tidyverse with terra objects: the tidyterra package. *Journal of Open Source Software*, 8(91):5751, 2023.

[57] R Core Team. R: A language and environment for statistical computing, 2024.

[58] Roozbeh Valavi, Gurutzeta Guillera-Arroita, José J Lahoz-Monfort, and Jane Elith. Predictive performance of presence-only species distribution models: a benchmark study with reproducible code. *Ecol. Monogr.*, 92(1):e01486.

[59] Frank T Breiner, Antoine Guisan, Ariel Bergamini, and Michael P Nobis. Overcoming limitations of modelling rare species by using ensembles of small models. *Methods Ecol. Evol.*, 6(10):1210–1218, October 2015.

[60] MapBiomas. MapBiomas collection 10: Land cover and use - brazil coverage 2024. `https://storage.googleapis.com/mapbiomas-public/initiatives/brasil/collection_10/lulc/coverage/brazil_coverage_2024.tif`, 2025. Accessed: 2025-10-9.

[61] MapBiomas. MapBiomas brasil collection 9. `https://brasil.mapbiomas.org/en/map/colecao-9/`, 2024. Accessed: 2025-10-20.

[62] Andy Nelson, Daniel J Weiss, Jacob van Etten, Andrea Cattaneo, Theresa S McMenomy, and Jawoo Koo. A suite of global accessibility indicators. *Sci. Data*, 6(1):266, November 2019.

[63] T G Farr and M Kobrick. Shuttle radar topography mission (SRTM), 2000.

[64] U.S. Geological Survey. GTOPO30: Global 30-arc-second elevation data, 1996.

[65] Joaquim Estopinan, Pierre Bonnet, Maximilien Servajean, François Munoz, and Alexis Joly. Modelling species distributions with deep learning to predict plant extinction risk and assess climate change impacts. *arXiv preprint arXiv:2401.05470*, 2024.

[66] Vinicius Costa Cysneiros, Allan Libanio Pelissari, Tatiana Dias Gaui, Luan Demarco Fiorentin, Daniel Costa de Carvalho, Telmo Borges Silveira Filho, and Sebastião do Amaral Machado. Modeling of tree height–diameter relationships in the atlantic forest: effect of forest type on tree allometry. *Canadian Journal of Forest Research*, 50(12):1289–1298, 2020. doi: 10.1139/cjfr-2020-0060.

[67] Jérôme Chave, Christophe Andalo, Sandra Brown, Michael A Cairns, Jeffrey Q Chambers, Derek Eamus, Horst Fölster, François Fromard, Niro Higuchi, Tatuo Kira, et al. Tree allometry and improved estimation of carbon stocks and balance in tropical forests. *Oecologia*, 145(1):87–99, 2005.

SUPPLEMENTARY INFORMATION

Supplementary Methods and Results

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 climatic contexts, by running restoration experiments on a set of 100 additional datasets of varying geographic location, extent, and resolution (Fig. S3). 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. S3). 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 were optimized.

We found that, despite such wide geographic variation and no fine-tuning to regional conditions, the policies consistently 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. S5a–c). However, the patterns were less consistent in regards to changes in species number across risk category (Fig. S5d–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. S1b), 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 of restoration policies that target both carbon and extinction risk reduction (Fig. S6a–d; Table S2).

Comparison with greedy search optimization

We carried out analyses of the simulated datasets using simple metrics to select the restoration areas instead of CAPTAIN's neural network to evaluate the performance of our trained models against simpler non-parametric approaches. We ran three models targeting carbon and cost, extinction risk and cost, or all three factors. The models used a subset of the 16 features monitored by the agent in the other experiments.

Non-parametric model (NPM) 1 focused on carbon and costs and used ratio between the potential carbon in a cell (feature 3 in our CAPTAIN policies) and its restoration cost (feature

16) as priority metric to select restored areas. NPM 2 targeted carbon, extinction risk, and costs. We summarized the importance of a cell for biodiversity as the weighted average of the number of CR, EN, VU, NT, and LC species (all rescaled between 0 and 1). To maximize the coverage of protected species we only counted LC species that are not already found within previously restored cells (feature 15 in our CAPTAIN policies). the weights were proportional to those used in our CAPTAIN reward function (Eq. 10). Restoration priorities were set proportional to the ratio between the sum of carbon and biodiversity metrics and costs. Finally, NPM 3 focused on extinction risk and costs and selected the cells with highest ratio between the biodiversity metric and costs.

Through the analysis of 100 simulated datasets we found that NPM1 achieved a performance similar to that of the equivalent neural network model (Policy 1) with comparable net carbon gains and slightly lower costs (Table). The biodiversity outcomes of NPM2-3 (i.e. models that included extinction risk reduction in their objectives) were however generally worse than those obtained from trained policies. For instance NPM3 (which did not include a carbon objective) included on average 32 fewer species in the selected areas and its restoration policies compared to the equivalent trained model (Policy 5). It also led to the recovery of 133 previously threatened species as opposed to the more than 200 species achieved by trained models. The restoration costs of NPM2-3 were also slightly higher than those of equivalent trained policies (Table).

Testing a larger neural network

We compared the performance of the small neural network utilized in our CAPTAIN experiments against that of a larger model. To this aim we trained additional models with two hidden layers of 16 and 8 nodes, using the same training procedure and simulated datasets. These models therefore included about 7 times more parameters than the small ones.

After running the trained models on our set of 100 simulations, we found that the larger models did not yield substantially improved outcomes compared to the the small models (Table). Costs, carbon and biodiversity metrics suggest that the richer parameterization does not lead to a significant change in performance under our simulated scenarios. These results are likely a function of the relatively small number of features used in our analyses and larger models might be beneficial in scenarios with more aspects of the environment being observed by the agent. Although the construction of the model with parameter-sharing across grid cells and rescaled features makes the number of parameters independent of the size of the area and the number of species, further exploration is needed to assess whether larger models might lead to a better performance with datasets spanning a wider extent and more complex biodiversity patterns.

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. S3).

	Model	Cost median (q2.5, q97.5)	Sp. in restored areas median (q2.5, q97.5)	Net carbon median (q2.5, q97.5)
non-parametric	1	941 (591, 1260)	383 (227, 449)	846 (688, 1021)
	2	1128 (846, 1723)	431 (339, 484)	690 (528, 929)
	3	1096 (789, 1700)	452 (364, 488)	629 (469, 879)
small NN	Pol. 1	968 (665, 1225)	323 (218, 437)	838 (676, 1026)
	Pol. 2	820 (599, 1150)	479 (462, 489)	636 (495, 815)
	Pol. 3	899 (718, 1128)	483 (475, 491)	502 (391, 660)
	Pol. 4	916 (720, 1278)	482 (470, 492)	456 (337, 597)
	Pol. 5	883 (723, 1131)	484 (476, 492)	387 (307, 503)
larger NN	Pol. 1	1015 (693, 1316)	406 (235, 467)	809 (646, 1033)
	Pol. 2	817 (654, 1065)	476 (457, 488)	653 (513, 918)
	Pol. 3	833 (602, 1064)	483 (469, 491)	532 (404, 798)
	Pol. 4	885 (658, 1205)	481 (468, 491)	500 (382, 718)
	Pol. 5	924 (705, 1178)	483 (474, 492)	448 (340, 628)

	Model	Effectively protected sp. median (q2.5, q97.5)	Sp. with improved status median (q2.5, q97.5)
non-parametric	1	73 (32, 159)	172 (85, 317)
	2	119 (48, 231)	263 (119, 415)
	3	133 (52, 233)	295 (131, 435)
small NN	Pol. 1	62 (23, 125)	140 (62, 272)
	Pol. 2	178 (63, 247)	361 (218, 453)
	Pol. 3	204 (71, 268)	387 (250, 463)
	Pol. 4	201 (56, 274)	385 (206, 458)
	Pol. 5	201 (65, 268)	386 (239, 463)
larger NN	Pol. 1	84 (51, 145)	207 (136, 286)
	Pol. 2	165 (55, 237)	352 (197, 441)
	Pol. 3	201 (61, 259)	383 (237, 454)
	Pol. 4	205 (65, 282)	363 (203, 444)
	Pol. 5	208 (61, 279)	385 (235, 455)

Table S3: Summary of the outcomes of 100 simulations performed under different restoration models. Non-parametric models 1–3 (see *Supplementary Methods and Results*) were based on simple metrics targeting carbon and costs(1); carbon, biodiversity and costs (2); and biodiversity and costs (3). The other models (Policies 1–5) were based on CAPTAIN-trained neural networks of 3 and 2 nodes in two hidden layers (small NN) and 16 and 8 nodes (larger NN). The effectively protected species quantify the number of species that started from a threatened category (CR, EN, VU) at the beginning of the simulations but moved to NT or LC after restoration, while the species with improved status are species that moved to a lower extinction risk. All values are given as median among simulations and 2.5 and 97.5% quantiles rounded to integer.

Supplementary Figures

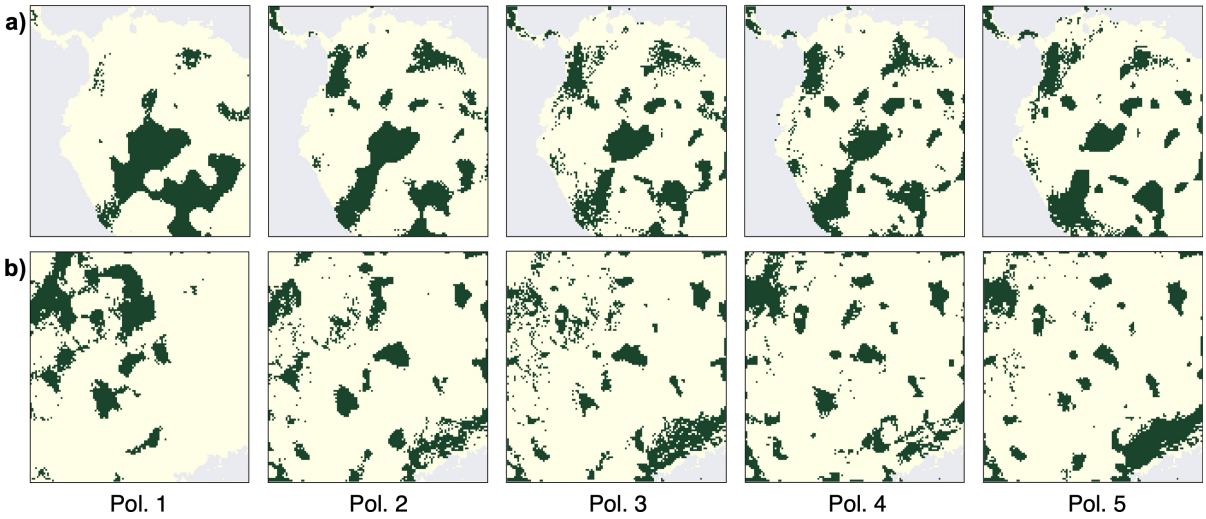

Figure S1: Changes in spatial restoration priorities depending on the relative weights of carbon (highest in Policy 1) and biodiversity (highest in Policy 5). a) Example of the restored areas in one of the 100 simulations performed on the reference region used to optimize the models (Fig. S2) selected under 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. 1). 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. S3).

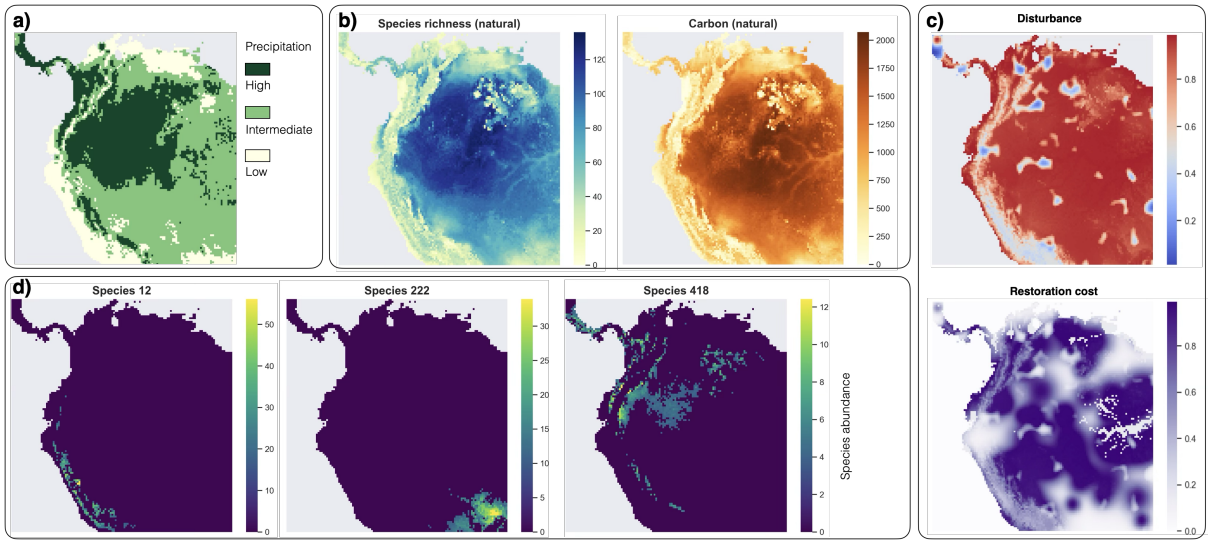

Figure S2: Map of the arbitrary area used for model training and subsequent simulations with examples of the empirical and simulated data used in our CAPTAIN runs. The physical properties of the area are based on the real topography and climate of the selected portion of Central and South America (a). 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. S3.

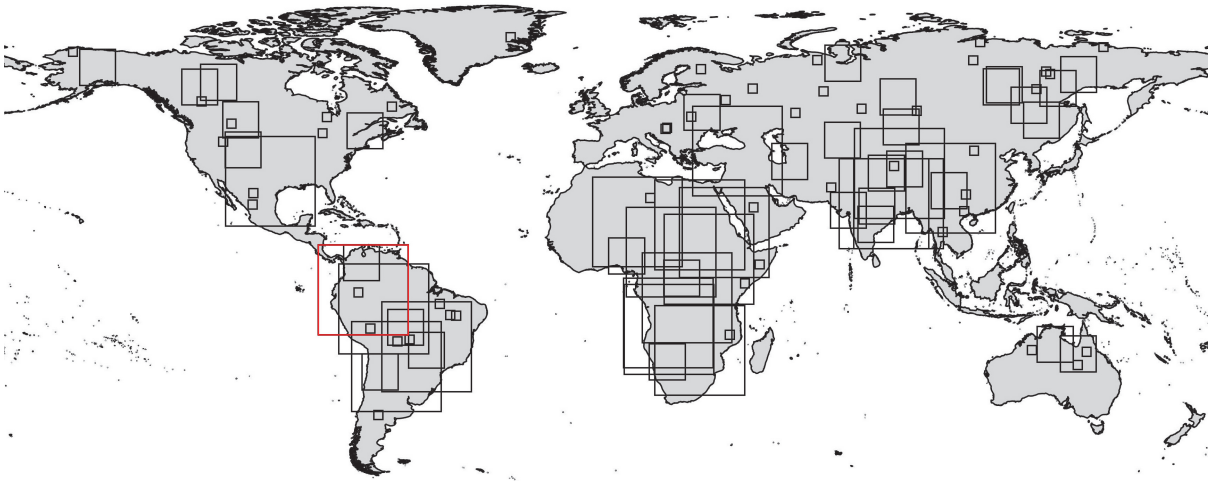

Figure S3: Maps showing the 100 randomly sampled rasters used as the basis for our simulations. The raster highlighted in red indicates the region used to train the models and run the simulations based on a fixed geographic extent (see Fig. S2 and *Methods*).

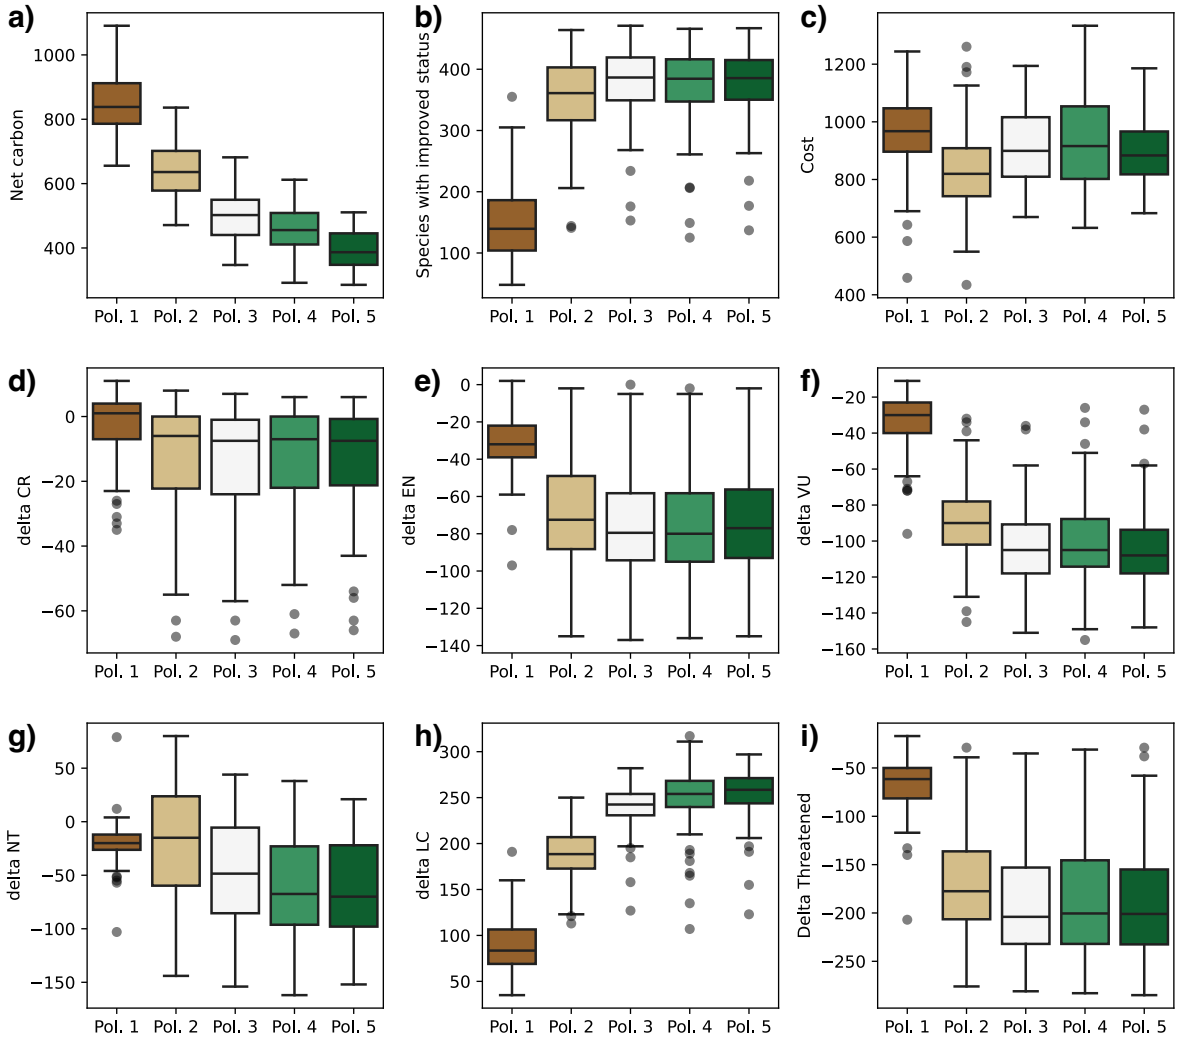

Figure S4: Summary statistics of key outcomes from 100 simulations run under five trained restoration models, Policies 1–5. The net carbon (a) represents the net gain in amount of carbon (here expressed in arbitrary units – see Methods) in the system after 30 time-steps during which 3,000 cells were 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).

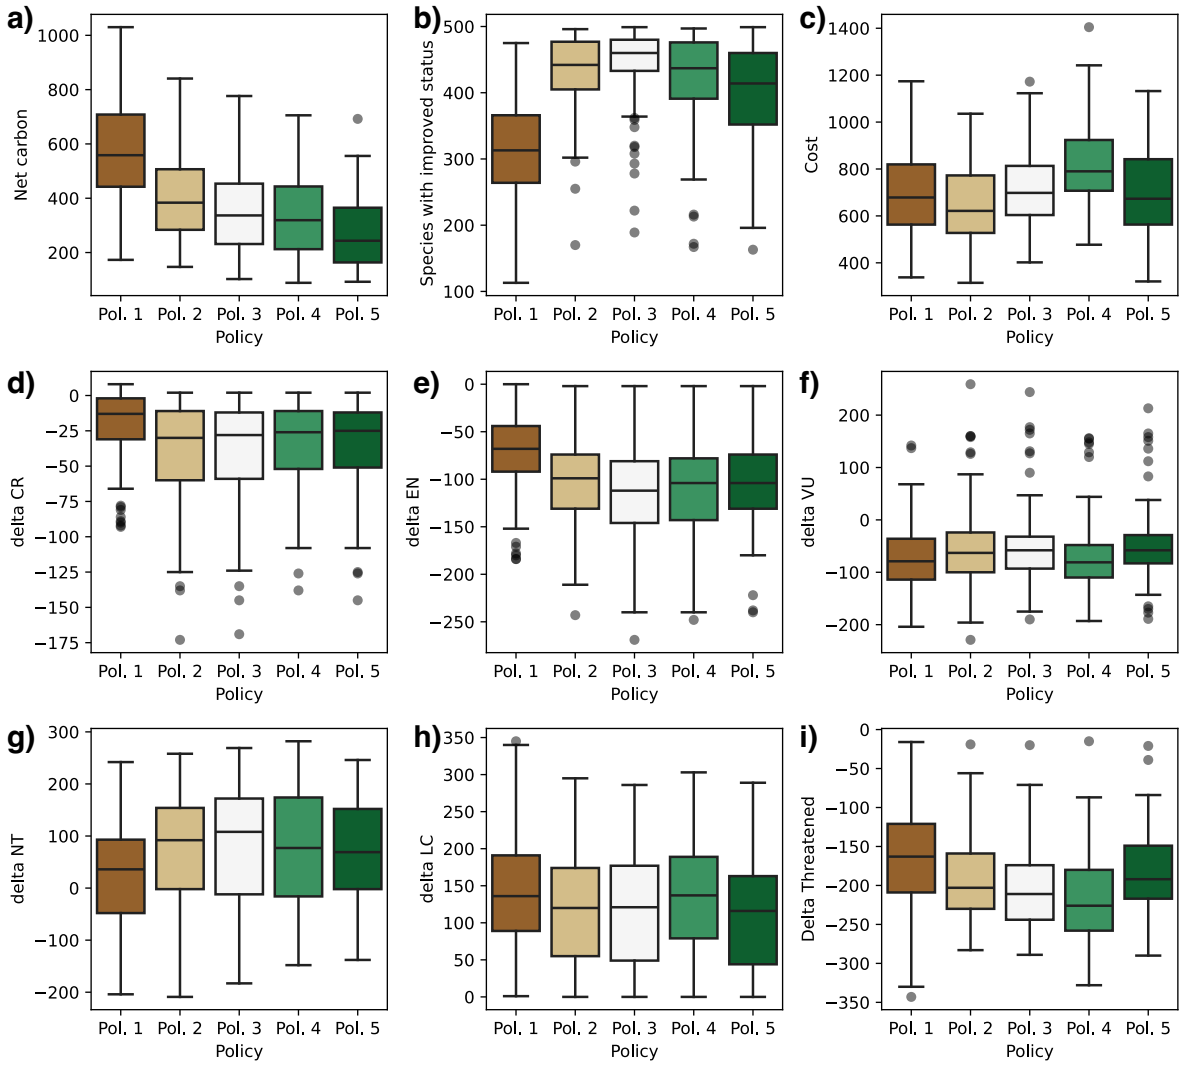

Figure S5: Summary of 100 simulations from random locations and variable resolution (Fig. S3) run under five trained restoration models, Policies 1–5. The general trends in terms of net carbon gain (decreasing from Policy 1 to 5; a), number of species with improved conservation status (higher from Policies 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.

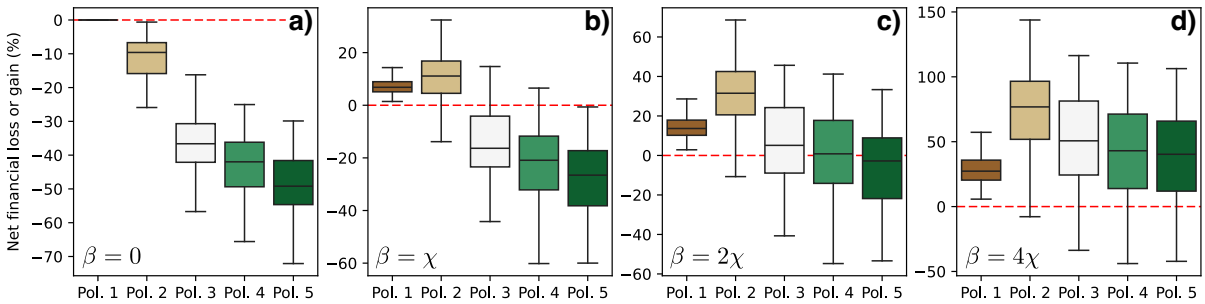

Figure S6: Financial outcomes of restoration policies (net percentage loss or gain, i.e. Φ in Equation 15) under different valuations of biodiversity credits, obtained from running the simulations across 100 random regions with different spatial location, extent and resolution. Negative values indicate 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 ($\Phi > 0$).

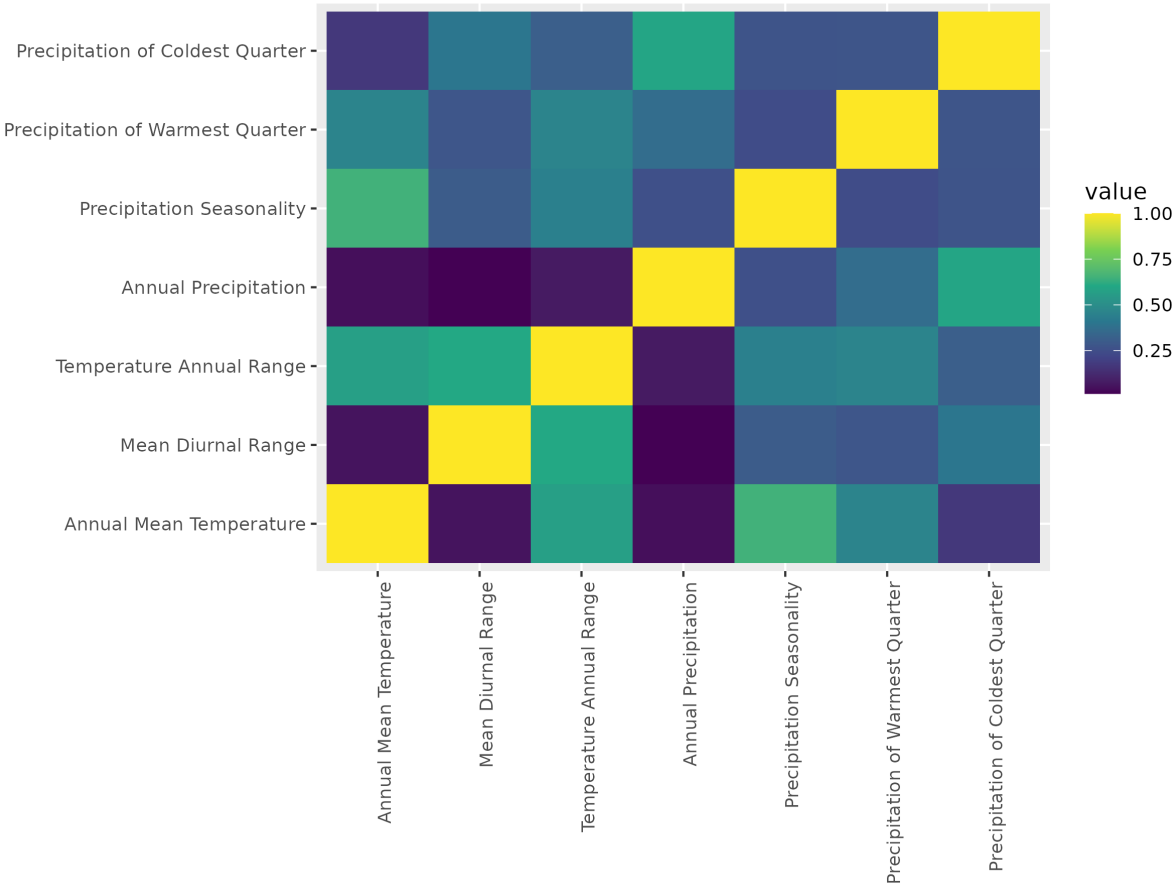

Figure S7: Correlation coefficients for the 7 bioclimatic variables retained for species distribution modeling.

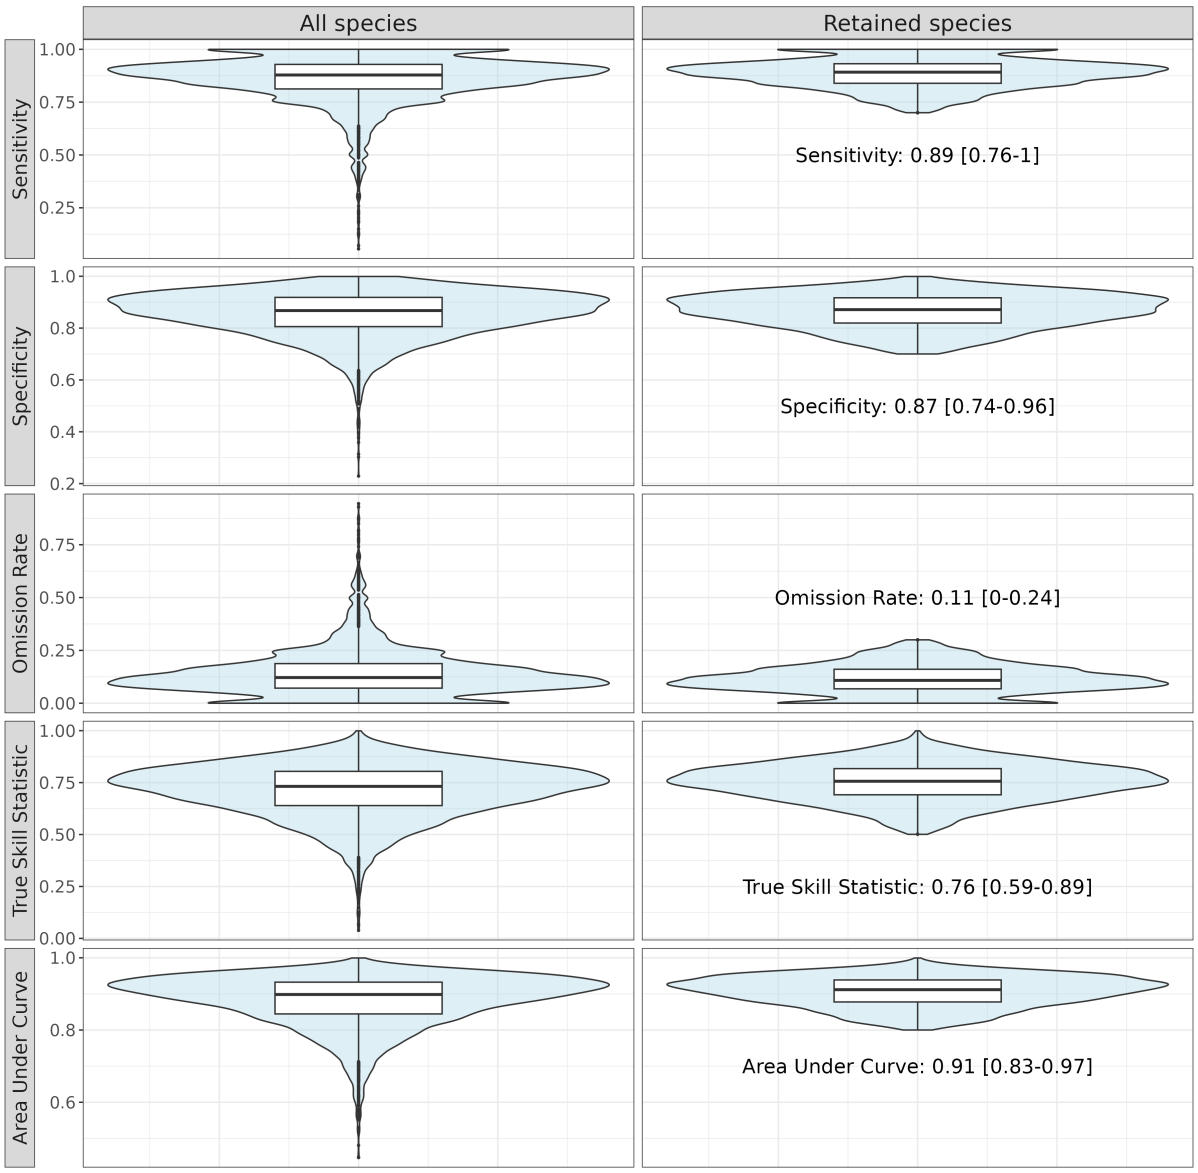

Figure S8: Violin plots and box plots of the five performance metrics of species distribution models across (a) all species and (b) retained species after applying exclusion criteria (sensitivity < 0.7, specificity < 0.7, True Skill Statistic < 0.5, and Area Under Curve < 0.8).

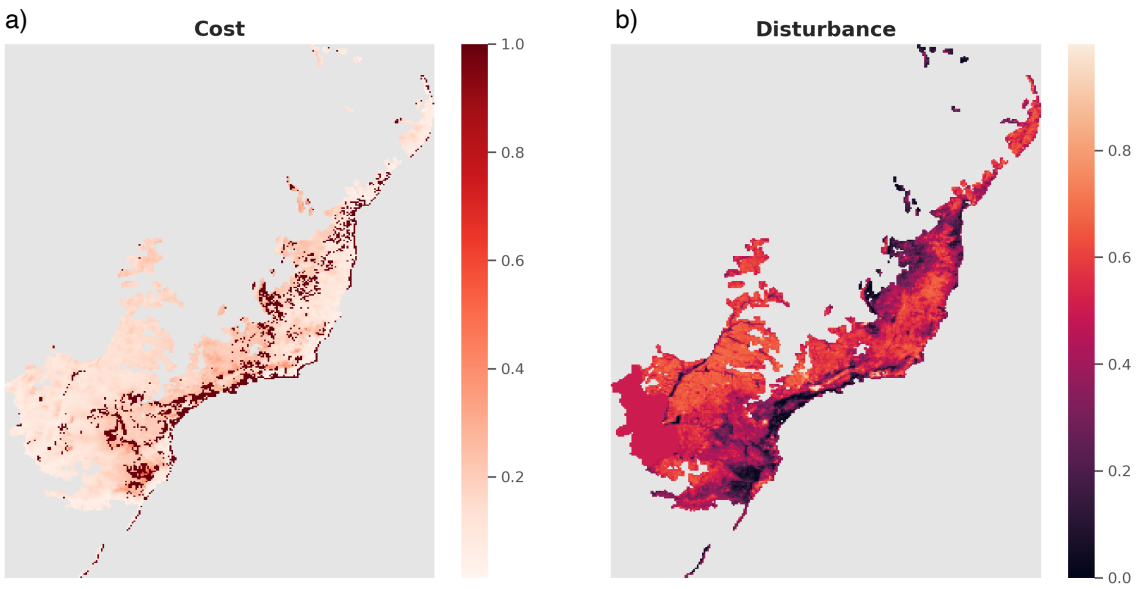

Figure S9: Distribution of restoration costs (a) and anthropogenic disturbance (b), approximated as described in the *Methods*.