

APPLICATION PAPER

Discovering Effective Policies for Land-Use Planning with Neuroevolution

Daniel Young¹, Olivier Francon¹, Elliot Meyerson¹, Clemens Schwingshackl², Jacob Bieker³, Hugo Cunha⁴, Babak Hodjat¹ and Risto Miikkulainen^{1,5}

¹Cognizant AI Labs, San Francisco, CA, USA.

²Ludwig-Maximilians-Universität, Munich, Germany.

³Open Climate Fix, London, United Kingdom.

⁴Cognizant Technology Solutions, Brussels, Belgium.

⁵The University of Texas at Austin, Austin, TX, USA.

Keywords: Evolutionary Optimization, Neural Networks, Surrogate Modeling, Climate Change, Land-use Change, Carbon Budget

Abstract

How areas of land are allocated for different uses, such as forests, urban areas, and agriculture, has a large effect on the terrestrial carbon balance, and therefore climate change. Based on available historical data on land-use changes and a simulation of the associated carbon emissions and removals, a surrogate model can be learned that makes it possible to evaluate the different options available to decision-makers efficiently. An evolutionary search process can then be used to discover effective land-use policies for specific locations. Such a system was built on the Project Resilience platform [17] and evaluated with the Land-Use Harmonization dataset LUH2 [16] and the bookkeeping model BLUE [14]. It generates Pareto fronts that trade off carbon impact and amount of land-use change customized to different locations, thus providing a proof-of-concept tool that is potentially useful for land-use planning.

Impact Statement

Land use is one of the factors contributing to climate change. In some cases, there is a choice on how the land is used, e.g. for crops, range, pasture, or forest. If such decisions are made well, carbon emissions from land-use change (ELUC) can be mitigated. This paper shows how a surrogate model can be trained to predict ELUC, and a prescriptive model then evolved with the surrogate can optimize the changes globally. These models are combined into a proof-of-concept system that can be used to inform land-use decision-makers about the best choices available to them, thus helping them mitigate climate change.

1. Introduction

One of the main factors contributing to climate change is CO₂ emissions due to land-use change [8]. Land-use CO₂ emissions depend crucially on how much land area is allocated for different uses. Certain land-use changes such as afforestation can absorb carbon, and are considered a negative emission technology (NET) [24]. However not all land can be simply converted to forest; for example, land used for crops and pasture is essential for other purposes such as food supply. Land-use patterns must therefore be planned to minimize carbon emissions while maintaining a satisfactory level of benefits received from the land, or ecosystem services. It is also important to note that land-use planning on its own is not enough to reach net-zero emissions. However, it can help offset emissions from sectors that are harder to decarbonize.

Recently, several different approaches were taken in order to optimize land-use planning to minimize carbon emissions. One such approach did so for the city of Nanjing in China using multi-objective linear programming [3]. While effective, this optimization method cannot take advantage of nonlinearities. In another approach, multi-criteria decision-making (MCDM) was used to optimize urban development in the Rajshahi Metropolitan Area of Bangladesh [29]. This approach uses manually ordered weights on outcomes to create solutions that trade off multiple objectives [37]. Six pre-defined scenarios with fixed levels of trade-off are required in order to construct such weights. The land-use planning task was also represented as a causal machine learning task in the potential outcomes framework [31] with crop type as the treatment, environmental factors as the features, and net primary productivity (NPP, i.e. the uptake flux in which carbon from the atmosphere is sequestered by plants) as the outcome [12]. Causality between treatment and outcome was found to inform policy, but the approach does not create a comprehensive strategy. Thus, all three methods were effective in their approach but each had their own unique drawbacks.

In this paper, a new approach for land-use optimization was developed with these concerns in mind. It takes advantage of nonlinearities in the data and automatically generates a Pareto front of land-use planning policies that trade off multiple outcomes. This work was done as part of Project Resilience, a non-profit project hosted by the ITU agency of the United Nations [17]. The goal is to provide decision-makers with a tool to know how their land-use choices affect CO₂ fluxes in the long term, and make suggestions for optimizing these choices.

More specifically, the tool was designed to answer three questions: (1) For a geographical grid cell identified by its latitude and longitude, what changes to the land use can be made to reduce overall CO₂ emissions (or in some cases, maximize carbon absorption)? (2) What will be the long-term CO₂ impact of changing land use in a particular way? (3) What are the optimal land-use choices that can be made with minimal cost and maximal effect?

The approach is based on the Evolutionary Surrogate-assisted Prescription method [7]. The idea is to first utilize historical data to learn a surrogate model on how land-use decisions in different contexts affect carbon emissions from land-use change, and then use this model to evaluate candidates in an evolutionary search process for good land-use change policies.

As a result, a Pareto front is generated of solutions that trade off reduction in carbon emissions and the amount of change in land use. Each point in the Pareto front represents an optimal policy for that trade-off. To make the results trustworthy, the tool allows the decision-maker to explore modifications to these policies, and see the expected effect. In the future, it should also be possible to evaluate confidence of the predictions, evolve rule sets to make the policies explainable, and utilize ensembling and further objectives and preferences to make them more accurate. Thus, the tool harnesses several techniques in machine learning to provide a proof-of-concept for decision-makers in optimizing land-use decisions.

The project page can be found at <https://project-resilience.github.io/platform/projects/landuse.html>, and an interactive demo at <https://landuse.evolution.ml>.

2. Background

The surrogate-assisted evolutionary optimization method used on the project is first outlined below, followed by the project background as the first application of Project Resilience.

2.1. Evolutionary Surrogate-assisted Prescription

Evolutionary Surrogate-assisted Prescription [ESP; 7] is an approach for optimizing decision-making in a variety of domains (Figure 1). The main idea is that a decision policy can be represented as a neural network (or in some cases, as a set of rules), and a good policy can be discovered through population-based search, i.e. using evolutionary computation techniques, in order to prescribe actions in a given context. However, each candidate must be evaluated, which is difficult to do in many real-world

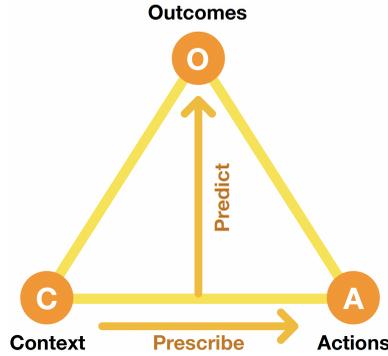


Figure 1. The Evolutionary Surrogate-assisted Prescription (ESP) method for decision optimization. A predictor is trained with historical data on how given actions in given contexts led to specific outcomes. It is then used as a surrogate in order to evolve prescriptors, i.e. neural networks that implement decision policies by prescribing actions for a given context, resulting in the best possible outcomes.

applications. Therefore, a surrogate model of the world is learned from historical data, predicting how good the resulting outcomes are for each decision in each context.

More formally, given a set of possible contexts \mathbb{C} and possible actions \mathbb{A} , a decision policy D returns a set of actions A to be performed in each context C :

$$D(C) = A, \quad (2.1)$$

where $C \in \mathbb{C}$ and $A \in \mathbb{A}$. For each such (C, A) pair there is a set of outcomes $O(C, A)$, and the Predictor P_d is defined as

$$P_d(C, A) = O, \quad (2.2)$$

and the Prescriptor P_s implements the decision policy as

$$P_s(C) = A, \quad (2.3)$$

such that $\sum_{i,j} O_j(C_i, A_i)$ over all possible contexts i and outcome dimensions j is minimized (assuming decreased outcomes are better). It thus approximates the optimal decision policy for the problem. The predictor can be learned from historical data on CAO triples. In contrast, the optimal actions A for each context C are not known, and must therefore be found through search.

More specifically in land-use optimization, context refers to the details of a given area such as current land use, a decision policy generates changes to the land use for the given area based on these details, and the outcomes are (1) the emissions from land-use change (ELUC) predicted by the surrogate as a result of this change, and (2) land-cost change, which is calculated as the percentage of the area changed.

ESP was first evaluated in reinforcement learning tasks such as function approximation, cart-pole control, and the Flappy Bird game, and found to discover significantly better solutions, find them faster, and with lower regret (deviation from the optimal reward) than standard approaches such as direct evolution, Deep Q-Networks [DQN; 23], and Proximal Policy Optimization [PPO; 33]. It thus forms a promising foundation for optimizing land use as well.

2.2. Project Resilience

A major application of ESP to decision-making was developed for optimizing strategies for non-pharmaceutical interventions (NPIs) in the COVID-19 pandemic [22]. Using case and NPI data from around the world, collected in 2020-2022 by Oxford University's COVID-19 project [13], a neural network model was developed to predict the course of the pandemic; this model was then used as a

surrogate in evolving another neural network to prescribe NPIs. It discovered different strategies at different phases of the pandemic, such as focusing on schools and workplaces early on, alternating policies over time, and focusing on public information and masking later. This application demonstrated the power of ESP in discovering good tradeoffs for complex real-world decision-making tasks.

Encouraged by this application, the XPRIZE Pandemic Response Challenge [36] was developed in 2020-2021 for NPI optimization using the same predictor+prescriptor approach. The competitors used several different machine learning or other techniques, and drew significant attention to AI-supported decision-making. The results were effective and in some cases used to inform decision-makers (e.g. on the 2021 surge in Valencia, Spain, and on the fall 2021 school openings in Iceland). An important conclusion was that it was possible to convene a group of experts—data scientists, epidemiologists, and public health officials—to build a useful set of tools to advise decision-makers on how to cope with and plan around a health disruption to society.

This experience with the XPRIZE Pandemic Response Challenge then led to Project Resilience [17]. The XPRIZE competition serves as a blueprint for a collaborative effort to tackle global problems. The goal of Project Resilience is to build a public AI utility where a global community of innovators and thought leaders can enhance and utilize a collection of data and AI approaches to help with better preparedness, intervention, and response to environmental, health, information, social equity, and similarly scoped problems. The project is hosted by the ITU agency of the United Nations, under its Global Initiative for AI and Data Commons, and contributes to its general efforts toward meeting the Sustainable Development Goals (SDGs).

Staffed by volunteers from around the world, Project Resilience was started in 2022 with the design of a minimum viable product (MVP) platform. The goal is to

- Develop an architecture to pull input and output data hosted by third parties;
- Build an application programming interface (API) for third parties to submit models.
- Develop code to compare both predictors and prescriptors in third-party models and produce a set of performance metrics;
- Build a portal to visualize the assessment of predictors and prescriptors to include generations of key performance indicators (KPIs) and comparison across models;
- Develop an ensemble model for predictors and prescriptors; and
- Provide a user interface to decision-makers to help them in their decision-making process.

Land-use optimization is the first application of this MVP, as described below.

3. Land-Use Optimization Task

The data sources are described first, followed by how they are used as the context, action, and outcome variables. The elements of the approach include the predictor and prescriptor models, a method for modeling uncertainty, and ensembling the results.

3.1. Data

The data for carbon emissions (Emissions resulting from Land-Use Change, ELUC) originate from the Bookkeeping of Land-Use Emissions model [BLUE; 14], which attributes carbon fluxes to land-use activities. A bookkeeping model is a semi-empirical statistical model that calculates CO₂ fluxes based on changes in the carbon content of biomass and soil caused by land-use changes. Bookkeeping models employ different carbon densities for different vegetation types (plant functional types in BLUE) and use response functions to determine the temporal dynamics of the fluxes. Advantages of bookkeeping models include their ability to separate anthropogenic from natural CO₂ fluxes and to attribute fluxes to specific land-use change events [6]. Bookkeeping models thus differ from more general climate models and land surface models that are based on parameterizations of physical, chemical, and biological processes. BLUE is one of four bookkeeping models that are used to provide ELUC estimates for the

annually published Global Carbon Budget [GCB; 8–10]. BLUE is particularly useful because it provides spatially explicit ELUC estimates. In this paper, BLUE is used to estimate committed emissions due to land-use changes, which means that all the emissions or removals caused by a land-use change event are attributed to the year of the event. The emission (or removal) estimates thus relate to the long-term impact of a land-use change event, i.e. up to multiple decades [14, 15]. While in principle BLUE can be used as the surrogate model for ESP, in practice its calculations are too expensive to carry out on demand during the search for good policies. Therefore, a number of BLUE simulations were performed covering a comprehensive set of land-use changes for 1850-2022, resulting in a dataset that could be used to train an efficient surrogate model.

The Land-Use Change (LUC) data is provided by the Land-Use Harmonization project [LUH2; 16]. LUH2 provides data on fractional land-use patterns, underlying land-use transitions, and key agricultural management information, annually for the time period 850-2100 at 0.25×0.25 degree resolution. LUH2 provides separate layers for states, transitions, and management. Version LUH2-GCB2022 from the Global Carbon Budget 2022 [9] was the basis for this paper. It is based on the HYDE3.3 historical database [18, 19] and uses the 2021 wood harvest data provided by the FAO agency of the United Nations. It consists of the following land-use types:

Primary: Vegetation that has never been impacted by human activities (e.g. agriculture or wood harvesting) since the beginning of the simulation

- primf: Primary forest
- primn: Primary nonforest vegetation

Secondary: Vegetation that is recovering from previous human disturbance (either wood harvesting or agricultural abandonment)

- secdf: Secondary forest
- secdn: Secondary nonforest vegetation

Urban

- urban: Urban areas

Crop

- crop: All crop types pooled into one category

Pasture

- pastr: Managed pasture land
- range: Grazed natural land (i.e., grazing on grassland, savannah, etc.).

BLUE calculates yearly ELUC estimates based on the land-use changes provided by LUH2. Urban areas are pooled into a single type; their carbon contents are too low and the total urban area on the planet is too small for differences among them to have a significant effect on the CO₂ flux. All crop types are also pooled into a single crop category. BLUE considers that the degradation of primary to secondary land often lowers the carbon content in vegetation and soils, and it considers pastures and the degradation of natural grasslands and savannahs by grazing. A complete overview of the model can be found in [14].

It is important to note two caveats in this data. First, the 0.25-degree spatial resolution of LUH2 is rather coarse for land-use recommendations. Second, different crop types do have noticeable differences when making decisions at a local level. To address these concerns, in the future more fine-grained land-use estimates can be used and crop type can be added as a feature variable, and such extensions would allow for more accurate prescriptions at the local level. In the meantime, the current versions of LUH2 and BLUE are sufficient to build a proof-of-concept system that demonstrates the power of land-use planning through neuroevolution.

Together the ELUC and LUC datasets form the basis for constructing the context, action, and outcome variables, as will be described next.

3.2. Decision-making Problem

The modeling approach aims to understand the domain in two ways: (1) In a particular situation, what are the outcomes of the decision maker's actions? (2) What are the actions that result in the best outcomes, i.e. the lowest ELUC with the smallest change in land use? The data is thus organized into context, action, and outcome variables.

Context describes the problem the decision maker is facing, i.e. a particular grid cell, a point in time when the decision has to be made, and the current land use at that time. More specifically it consists of:

- Latitude and Longitude, representing the cell on the grid. Climate information loosely corresponding to plant functional types is implicitly represented in the location of the cell, as shown in Appendix A.
- Area, representing the surface of the cell. Cells close to the equator have a bigger area than cells close to the poles.
- Year, useful to capture historical decisions: The same cell may have been through many land-use changes over the years.
- Land use, representing the percentage of land in the cell used by each type above. For simplicity, all crop types were pooled into a single type. The percentages do not necessarily reach 100% in each grid cell; a 'nonland' type (e.g. typically sea, lake, etc.) is assumed to fill the remainder.

Actions represent the choices the decision-maker faces. How can they change the land? In the study of this paper, these decisions are limited in two ways:

First, decision-makers cannot affect primary land. The idea is that it is always better to preserve primary vegetation, such as pristine rainforests; destroying it is not an option given to the system. Technically, it is not possible to re-plant primary vegetation. Once destroyed, it is destroyed forever. If re-planted, it would become secondary vegetation, for which BLUE assumes lower carbon content than for primary vegetation [14].

Second, decision-makers cannot affect urban areas. The needs of urban areas are dictated by other imperatives, and optimized by other decision makers. Therefore, the ESP system cannot recommend that a city should be destroyed, or expanded. Further constraints on crop usage for food production are discussed in Section 5.4

Outcomes consist of two conflicting variables. The primary variable is ELUC, i.e. emissions from land-use change. It consists of all CO₂ emissions attributed to the land-use change, in metric tons of carbon per hectare (tC/ha), obtained from the BLUE simulation. A positive ELUC means carbon is emitted to the atmosphere and a negative number means carbon is taken up by vegetation. For simplicity, all future CO₂ emissions or removals are attributed to the initial year of the land-use change; temporal dynamics in CO₂ fluxes are not taken into account. Similarly, the full land-use change is assumed to take place; incremental changes are not taken into account. The secondary variable is the cost of the land-use change, represented by the percentage of land that was changed. This variable is calculated directly from the actions.

There is a tradeoff between these two objectives: It is easy to reduce emissions by changing most of the land, but that would come at a huge cost. Therefore, decision-makers have to minimize ELUC while minimizing land-use change at the same time. Consequently, the result is not a single recommendation, but a Pareto front where each point represents the best implementation of each tradeoff given a balance between the two outcomes.

4. Models

The system consists of the predictor, trained with supervised learning on the historical data, and the prescriptor, discovered through a evolutionary search in which each candidate is evaluated with the predictor.

Prediction: Given the context and actions that were performed, the predictive model estimates the outcomes. In this case, since the cost outcome (i.e. the percentage of changed land) can be calculated directly, only the ELUC is predicted by the model. That is, given the land use of a specific location, and the land-use changes that were made during a specific year, the model predicts the long-term CO₂ emissions directly caused by these land-use changes. The model does not predict non-ELUC emissions such as forest fires that may lead to a change into secondary forest.

Any predictive model can be used in this task, including linear regression, random forest, or a neural network; they are each evaluated below. As usual in machine learning, each model is fit to the existing historical data and evaluated with left-out data.

Prescription: Given context, the prescriptive model suggests actions that optimize the outcomes. The model has to do this for all possible contexts, and therefore it represents an entire strategy for optimal land use. The strategy can be implemented in various ways, like sets of rules or neural networks. The approach in this paper is based on neural networks.

The optimal actions are not known, but the performance of each candidate strategy can be measured (using the predictive model), therefore the prescriptive model needs to be learned using search techniques. Standard reinforcement learning methods such as PPO and DQN are possible; the experiments in this paper use evolutionary optimization, i.e. conventional neuroevolution [35]. As in prior applications of ESP [7, 22], the network has a fixed architecture of two fully connected layers; its weights are concatenated into a vector and evolved through crossover and mutation.

5. Experiments

This section describes the main results of the project so far, including different ways to create good predictors and the evolution of effective prescriptors based on them. The behavior of the discovered prescriptors is analyzed quantitatively and through various ablation and extension studies. As a concrete implementation of the mechanisms, an interactive demo of the entire land-use recommendation system is described.

5.1. Prediction

In preliminary experiments, prediction performance was found to differ between major geographical regions. To make these differences explicit, separate models were trained on different subsets of countries: Western Europe (EU), South America (SA), and the United States (US), making up 1.12, 9.99, and 7.26 percent of the global dataset, respectively. The EU subset contains the UK, France, Germany, the Netherlands, Belgium, Switzerland, and Ireland. The SA subset includes Brazil, Bolivia, Paraguay, Peru, Ecuador, Colombia, Venezuela, Guyana, Suriname, Uruguay, Argentina, and Chile. These countries were chosen as a representative subset to make the experiments computationally less expensive.

Upon examining the global dataset, a linear relationship was found between |ELUC| and land-use change with an R-squared value of 0.92. Therefore, a linear regression (LinReg) was fit to the data as the first predictive model. In preliminary experiments, the LinReg model was found to perform the best when current land use, latitude, longitude, cell area, and time were removed from the inputs. Thus, the difference in land use was used as the only input feature in the experiments. Data from the years 1851–2011 (inclusive), totaling 38,967,635 samples globally, was used to train the LinReg models; the 2,420,350 remaining samples from 2012–2021 (inclusive) were used for testing.

Second, a random forest model (RF) was trained on the same data but now includes the full set of input features, in order to take advantage of RF's ability to decide which features are most useful. The forest consisted of 100 decision trees with unrestricted maximum depth; at each split, a random subset of n features was considered (where n is the square root of the total number of features). To make the training time of the Global RF model feasible, it was trained with data from only 1982–2011. Preliminary experiments showed that constraining the training set in this manner did not decrease

Model	Time (s)	EU	SA	US	Global
LinReg (EU)	0.047	0.033	0.172	0.169	0.206
LinReg (SA)	0.457	0.137	0.153	0.061	0.110
LinReg (US)	0.331	0.139	0.146	0.035	0.073
LinReg (Global)	4.644	0.139	0.150	0.035	0.074
RF (EU)	17.697	0.064	0.211 [†]	0.161 [†]	0.218 [†]
RF (SA)	209.688	0.133 ^{*†}	0.071 ^{*†}	0.074 [†]	0.126 [†]
RF (US)	111.701	0.163	0.185 [†]	0.032 [*]	0.094 [†]
RF (Global)	417.647	0.041 ^{*†}	0.076 ^{*†}	0.028 [*]	0.045 ^{*†}
NeuralNet (EU)	10.711	0.025 ^{*†}	0.277	0.286	0.334
NeuralNet (SA)	103.696	0.248	0.100 [*]	0.562	0.399
NeuralNet (US)	73.141	0.136 [†]	0.225	0.024 ^{*†}	0.150
NeuralNet (Global)	1649.193	0.046 [*]	0.110 [*]	0.025 ^{*†}	0.050 [*]

Bold: the best model in each region; ^{*}: 99% confidence the model outperforms LinReg;

[†]: 99% confidence that RF outperforms NeuralNet, or vice-versa

Table 1. Average training time and mean absolute error (MAE) in tC/ha of the trained models of each type trained on each region and evaluated on that region as well as all other regions. Statistical significance was estimated in each case with the unpaired *t*-test. The NeuralNet models are more accurate RF on EU and US, RF more accurate than NeuralNets on SA and globally; both are more accurate than LinReg in all regions (except RF on EU). Overall, the specialized models are more accurate than the Global models. However, RF models do not extrapolate well (as seen in Figure 2), which is necessary for evaluating prescriptor candidates. In contrast, NeuralNets do, and they are also otherwise reasonably accurate. Therefore, the Global NeuralNet model was used to evolve the prescriptors.

performance significantly. Both the linear regression and random forest models were trained using the scikit-learn library [26].

Third, a neural network (NeuralNet) was trained for the same task with the full set of features and data from 1851–2011 using the PyTorch library [25]. The network was fully connected, composed of an input layer, a single hidden layer of size 4096, and an output layer. The output layer took as input the concatenation of the input layer and the hidden layer in order to take advantage of the observed linear relationships more easily [4]. It was trained for three epochs with batch size 2048 using the AdamW optimization algorithm [20]. After each epoch, the learning rate was scaled by a factor of 0.1.

LinReg, RF, and NeuralNet models were trained separately on each region and globally. While LinReg is deterministic and needed to be trained only once, RF and NeuralNet are stochastic in their initialization and training and were therefore trained 30 times to evaluate statistical significance. All models were tested on all regions with data from the most recent years, i.e. 2012–2021, and evaluated on mean absolute error (MAE) of predicted ELUC. As shown in Table 1, the LinReg models were fast to train, but they performed the worst (i.e., they have the largest mean absolute error in almost all cases), suggesting that the linear relationship is not enough to account for the data fully. RF performed significantly better than LinReg in all regions but EU, significantly worse than NeuralNet in EU and US, and outperformed NeuralNets in SA and globally. However, RF models do not generalize/extrapolate well to large changes in land use that occur in the dataset infrequently (or not at all for some land types; Figure 2), which turns out to be important when evolving prescriptors. Many prescriptor candidates suggest large land-use changes, and the predictor model needs to be able to evaluate them. In contrast, NeuralNets capture nonlinearities, which explains why they outperform LinReg in all regions. Most importantly, they extrapolate well, which suggests that they are a better choice for the surrogate model in prescriptor evolution. Additionally, LinReg and NeuralNet models have many fewer parameters than RF (taking 4KB and 363KB on disk respectively vs. 14GB for RF).

In order to visualize the behavior of the predictor models, they were queried with a synthetic dataset created with every possible conversion from 100% land type A to 100% land type B (the types primf, primn, and urban were excluded from B as was explained in Section 3.2). Because the RF and NeuralNet

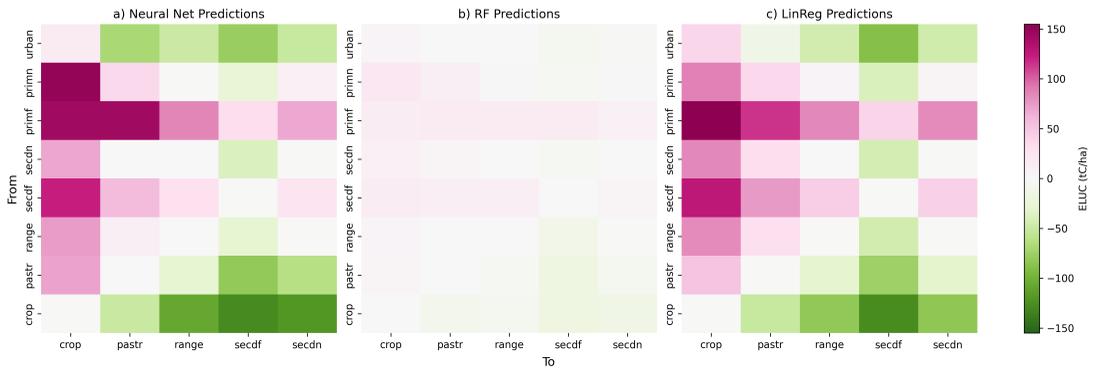


Figure 2. Visualizing the differences in model behavior. Predictions for ELUC (tC/ha) are created for the Global models using synthetic data created by changing 100% of land-use type A (row) to 100% of type B (column) in a 1% sample across the range of latitude, longitude, cell area, and year occurring in the test data. The results are averaged for each conversion from A to B. The models generally agree on the sign of ELUC, which in turn aligns with the sign generated by the BLUE model, suggesting that the results are reliable. The RF model is not able to extrapolate to large values, resulting in low predictions; LinReg and NeuralNet are similar but differ numerically, presumably due to the differences between linear and nonlinear predictions.

take latitude, longitude, cell area, and year as input, 24,204 (1%) latitude/longitude/cell-area/year combinations from the test dataset were sampled randomly and evaluated with each conversion. The results were averaged across the samples to get the overall result for the conversion. Figure 2 shows these results for the Global models. The first finding is that all models generally rank conversions to crops < pasture < forest in terms of reducing ELUC, and generally agree on the sign of ELUC. This ranking also generally aligns with those generated by the BLUE model. This consistency suggests that the results are reliable. Second, while LinReg and NeuralNet make similar predictions overall, their numerical values differ in many places, presumably due to the nonlinearities. Third, despite RF's generally high accuracy in known scenarios, it is not able to generalize/extrapolate to large changes and instead predicts low-magnitude ELUC values for all imposed land-use transitions.

Thus, since NeuralNet achieves similar or better performance than the RF, extrapolates well to extremes, and can learn nonlinear relationships in the data, the Global NeuralNet was used as the surrogate model to evolve the prescriptors.

5.2. Prescriptions

The prescriptors were fully connected neural networks with two layers of weights, implemented with PyTorch. The topology was fixed and only the weights were evolved. The input consisted of 12 float values, i.e. the eight land-use percentages and the latitude, longitude, and area of the geographical cell, as well as the year. The hidden layer contained 16 units with tanh activation.

During evolution, prescriptor candidates were evaluated on a fixed random subset of 1% of the Global dataset from years 1851–2011 with the Global NeuralNet predictor. Prescriptor output was a single vector of size five (the eight land-use percentages minus primf, primn, and urban, which were considered fixed), produced by the ReLU activation function. These five outputs were normalized to sum to the previous amount of land taken by these five types. The result is the prescriptor's suggestion for the percentage of land use of each type. The five outputs were combined with the fixed land-use values for primf, primn, and urban to compute the difference from the initial land use. This difference was concatenated with the initial land use as well as latitude, longitude, cell area, and year, and passed to

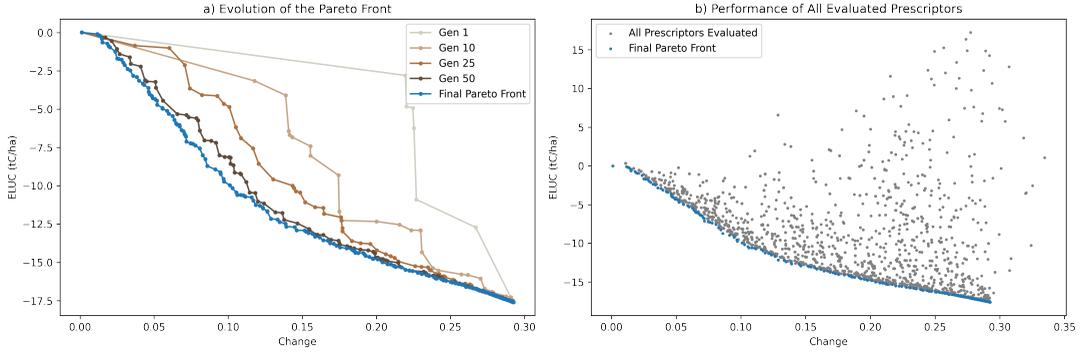


Figure 3. Evolution of prescriptors with the Global NeuralNet predictor. (a) The Pareto front moves towards the lower left corner over evolution, finding better implementations for the different tradeoffs of the ELUC and land-use change amount objectives. (b) Each prescriptor evaluated during evolution is shown as a dot, demonstrating a wide variety of solutions and tradeoffs. The final Pareto front (collected from all generations) is shown as blue dots. It constitutes a set of solutions from which the decision-maker can choose a preferred one.

the Predictor model to get the predicted ELUC metric. The cost outcome, i.e. the percentage of land-use change, was computed based on the difference between the five outputs and the corresponding inputs. All positive differences were summed for a given cell and divided by the total land-use percentage of the cell (excluding nonland). These metrics were aggregated across the evaluation data set to obtain the performance outcome of each prescriptor candidate. Thus, a candidate was evaluated based on its average performance in creating a strategy for each land area in the test set.

The aggregated ELUC and the land-use change outcomes were minimized with an evolutionary search on the neural network weights. Conventional neuroevolution with crossover and mutation operators on a vector of network weights was combined with NSGA-II [5] to make the search multiobjective. In order to seed evolution with good coverage of the search space, two prescriptor models were trained with backpropagation and injected into the initial population: One to prescribe no land-use changes at all and the other to prescribe as much secondary forest (secdf) as possible (which was observed to reduce ELUC the most). This process takes advantage of the Realizing Human Expertise through AI (RHEA) framework [21] for distilling human intuition into evolutionary search. Details and the impact of this seeding mechanism are discussed in Section 5.5. The rest of the population was created randomly through an orthogonal initialization of weights in each layer with a mean of zero and a standard deviation of one [32]. Evolution was run for 100 generations with the following parameters:

- mutation_type: gaussian_noise_percentage
- mutation_factor: 0.1
- mutation_probability: 0.2
- population_size: 100
- parent_selection: tournament
- initialization_distribution: orthogonal.

Figure 3 demonstrates the progress of evolution towards increasingly better prescriptors, i.e. those that represent better implementations of each tradeoff of the ELUC and land-use change objectives. These solutions are referred to as "Evolved Prescriptors" in the rest of the paper. They represent a wide variety of tradeoffs and a clear set of dominant solutions that constitute the final Pareto front. That set is returned to the decision-maker, who can then select the most preferred one to implement.

To demonstrate the value of evolutionary search, the Evolved Prescriptors were compared against two heuristic baselines (Figure 4). Each of them converts as much land as possible to secdf, which has the lowest ELUC, up to a set threshold. The first is a simple baseline called "Even Heuristic", where

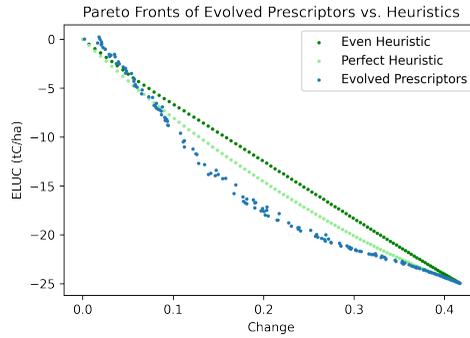


Figure 4. The Pareto fronts of Evolved Prescriptors vs. heuristic baselines, with ELUC and land-use change evaluated on the Global test set. The Evolved Prescriptors achieved better solutions than the baselines in the middle-change region where the land-use changes matter the most, demonstrating that they can take advantage of nonlinear relationships in land use to discover useful, non-obvious solutions.

an equal proportion of land up to the threshold is taken from all land-use types and converted to secdf. The second is an optimal linear baseline called "Perfect Heuristic", where as much land as possible is first taken from the land-use type that has the highest weight in the Global LinReg model, then from the land-use type that has the second highest weight, and so on until either the given land-use change threshold is reached or there is no land left to convert to secdf. These heuristic solutions were created with evenly distributed land-use change thresholds between 0% and 100% in order to generate heuristic Pareto fronts. All three methods were then evaluated on a random 1% sample of the Global test set from the years 2012–2021.

As Figure 4 shows, the Evolved Prescriptors and the heuristics all converged in the high-change region of the Pareto front, where 100% forest prescriptions dominate. The heuristics performed slightly better in the low-change region, where the predictors may not be accurate enough to distinguish between small changes. However, in the most useful middle-change region, the Evolved Prescriptors outperformed both heuristics, demonstrating that they are likely exploiting the nonlinear relationships well.

The performance of each prescription method can be quantified in terms of hypervolume, which measures how much of the solution space it dominates. It is calculated as the area above and to the right of the Pareto front up to maximum possible change and zero ELUC, divided by the area of the entire search space (which starts from zero change and lowest possible ELUC, which in turn is defined as the ELUC of the Perfect Heuristic at the maximum change). The Evolved Prescriptors outperform the Perfect Heuristic, which outperforms the Even Heuristic, with hypervolumes of 0.616, 0.566, and 0.512, respectively.

5.3. Prescriptor Behaviors

It is interesting to analyze how the Evolved Prescriptors were able to improve upon the linear heuristics, which were constructed to be quite good to begin with. One representative from each method was chosen from the middle-change region and compared (Figure 5a). The heuristic methods both prescribe 19.8% change and -12 and -14 ELUC on average, and the closest Evolved Prescriptor that dominates them averages 19.7% change and -18 ELUC. In Figure 5b these averages were expanded to show the results for actual samples in the test set (only 10% of those samples are plotted for readability). The samples for the Even Heuristic are hidden under those of the Perfect Heuristic: When these heuristics cannot reach their allocated thresholds, they change everything to secondary forest, resulting in overlapping points; their points similarly overlap at the vertical line indicating the change threshold. Note that the Evolved Prescriptor allocates more change to several samples than the Heuristic: This is because the

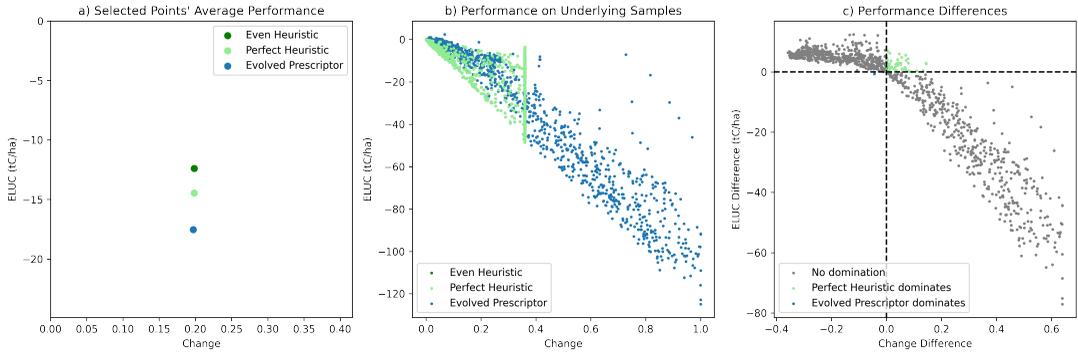


Figure 5. Comparing a selected Evolved Prescriptor with the Perfect Heuristic. (a) The average performance of the Evolved Prescriptor dominates that of the heuristic. (b) The averages are expanded into actual samples in the test set (subsamped for readability). The samples for the Even Heuristic are largely hidden under the samples for the Perfect Heuristic. The Evolved Prescriptor suggests many more large changes than the heuristic. (c) The differences in change percentage and ELUC between the Evolved Prescriptor and the Perfect Heuristic for each test sample, with color indicating which solution dominates. Surprisingly, this particular Evolved Prescriptor dominates the Perfect Heuristic only on a single individual sample. Thus, evolution discovered the insight that it is possible to do well globally by utilizing a few cases where large change is possible.

heuristic's change threshold is set to 0.35 (in order to get to ~20% change on average) whereas the Evolved Prescriptor has no such restriction.

The differences in change percentage and ELUC between the Perfect Heuristic and a typical Evolved Prescriptor are plotted for each sample in the test set in Figure 5c. The heuristic dominates in the samples in the upper right quadrant (where both differences are above 0, i.e. the prescriptor's change and ELUC are both larger than those of the heuristic) and the prescriptor in the bottom left quadrant. Surprisingly, the Evolved Prescriptor dominates the Perfect Heuristic in only 1 of the 24,204 test samples. How can this be?

The reason is that evolution discovered a way to take advantage of large changes. As shown in Figure 5b, the trained prescriptor is not constrained by a limit on every change it makes: As long as its changes remain below the limit on average, it can make a few large changes. As a result, evolution discovered a strategy where it makes only a few large changes, and makes them in cells where they have the largest impact. Thus, even though the Evolved Prescriptor almost never beats the Perfect Heuristic in any individual cell, when averaged across all cells, it dominates the heuristic. In other words, pick your battles right, and you'll win overall.

In order to see what the recommendations look like, the average policies of the three approaches are illustrated in Figure 6a. The greatest amount of negative change made is to crop and positive is to secondary forest, i.e. all approaches recommend converting cropland to forest, which makes sense. The greatest departure from the heuristics is that the Evolved Prescriptors recommend more crop change and almost no rangeland change.

Where the Evolved Prescriptors picked their battles compared to the Perfect Heuristic is shown in Figure 6b. Certain areas such as the Amazon rainforest were left untouched by both, because the system does not allow primary forests (that capture much carbon) to be changed. The Evolved Prescriptors avoided change in dry regions, such as the US Southwest, Sahara, Middle East, Central Asia, and Northwestern Australia, where presumably only limited change is effective. The prescriptors recommended more change in tropical, temperate, and continental regions like the US Midwest, Gulf of Guinea, South and Southeast Asia, Continental Europe, and the Pampas region of South America. These locations align with the change profile in Figure 6a: The areas where more crops can grow are changed more, whereas plains and desert regions corresponding to rangeland are barely touched.

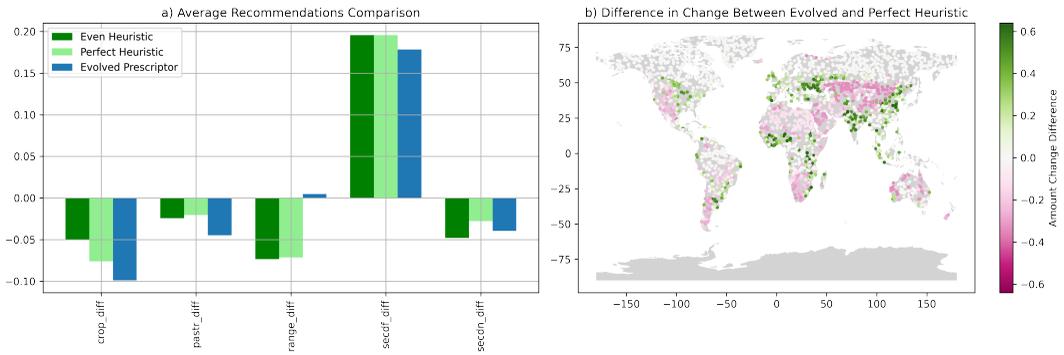


Figure 6. Characterizing the recommendations. (a) The strategies used by the Evolved Prescriptors and the heuristics are illustrated by plotting the average amount of change in each land-use type. Primarily, they all suggest converting cropland to secondary forest. The Evolved Prescriptor removes more crop on average, and generally takes advantage of more change. (b) The learned ability of the Evolved Prescriptor to allocate more or less change to certain regions. Primary forest is left untouched, dry regions are changed less, and tropical, temperate, and continental regions are changed more.

5.4. Constraining Crop Loss

While minimizing ELUC and change allows a decision-maker to find solutions along an important trade-off, it does not limit how much cropland the prescriptor is allowed to change, which could lead to food shortages. A step towards a more practical application is to optimize under such a constraint. To this end, the same evolutionary setup was run with a third objective of minimizing absolute change in crops. This approach assumes that the current level of crops is sufficient, however, the constraint could easily be set to an arbitrary value. Consequently, a seed model, trained through backpropagation, was included in the initial population to change as much land as possible to forest without removing any cropland. Such a seed makes it more likely for evolution to discover solutions that do not prescribe as much crop change. Evolution parameters were the same as in Section 5.2. These newly evolved prescriptor models were compared to the previous heuristics as well as to a new "No Crop" Heuristic, i.e. the Perfect Heuristic without modifying any cropland.

Three main findings resulted from this experiment. First, the No Crop heuristic performs quite poorly in the other two objectives, demonstrating that more sophisticated optimization is necessary to incorporate this constraint. Second, evolution was able to find solutions trading off all three objectives, allowing a decision-maker to fix the amount of crop to change before selecting their preferred solution. Third, evolution was able to discover solutions that outperformed the Perfect Heuristic in all three objectives. The trade-off can be seen in Figure 7a. The Pareto front is again plotted as in Figure 4, but this time, each Evolved Prescriptor is colored by the difference in how much it changed cropland on average compared to the Perfect Heuristic at the same level of change. There is a clear color gradient from more to less crop change perpendicular to the Pareto front such that generally the candidates that prescribed more crop change did the best and those that prescribed less did worse in terms of ELUC and change. Thus, a decision-maker can first choose how much crop they want to change, and then select the best prescriptor at that level along the Pareto front.

The solutions that outperform the Perfect Heuristic in all three objectives are shown in Figure 7a outlined in black. Each point represents a prescriptor that achieves a better ELUC with less change than the Perfect Heuristic while also prescribing less crop change. Figure 7b shows the average prescription of one such point in comparison to the heuristics in order to determine how it did so. The Evolved Prescriptor changed slightly less crop but compensated by removing more pasture. This strategy allowed it to prescribe less crop change, less change overall, and achieve a lower ELUC than the perfect heuristic, showing the power of evolution in the case of three objectives. Ultimately, the results of this experiment

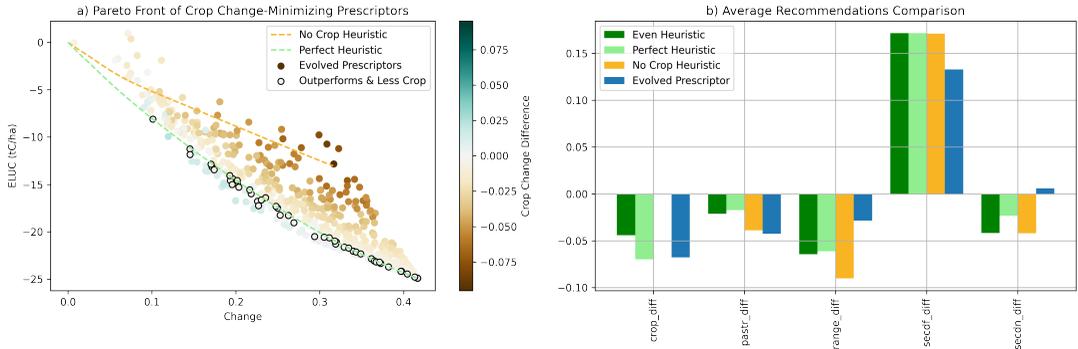


Figure 7. *Heuristics and Evolved Prescriptors with crop-change minimization as a third objective. (a) The final Pareto front of the Evolved Prescriptors is plotted against the Perfect and No Crop Heuristics. Each Evolved Prescriptor is colored according to how different its average prescribed crop change is from the Perfect Heuristic. The points that are outlined identify Evolved Prescriptors that outperform the Perfect Heuristic in all three objectives. (b) The average prescriptions of one such Evolved Prescriptor. It reduces pasture more in order to reduce crops less while still outperforming the heuristics along all objectives.*

show the ability of the proof-of-concept system to incorporate realistic constraints such as those on food production in order to make it practical in real-world decision-making.

5.5. Incorporating Human Expertise

When developing solutions to global societal problems with AI, incorporating expert human knowledge may improve performance while also adding accountability and explainability [21]. The first step was taken in Section 5.2, where two hand-designed solutions were included in the initial prescriptor population through RHEA.

In more detail, RHEA consists of (1) defining the problem so that experts can develop solutions to it, (2) gathering a set of solutions from the experts, (3) distilling these solutions into a population of evolvable models such as neural networks, and (4) discovering better models by evolving this population. The land-use optimization problem was already defined with context, action, and outcome variables and was therefore compatible with the framework. Next, the behaviors for prescribing no change and all change to secondary forest were identified as useful extremes. In order to distill these behaviors, two sets of labels for 10% of the training dataset were created in order to run backpropagation. One was constructed by not changing the land use at all and the other changed as much land as possible to secondary forest. Then the AdamW optimization algorithm was run to minimize mean squared error over 15 epochs to targets that prescribed either no change or full change for a sampling of land-use cases. The resulting two trained prescriptors were then injected into the initial population and evolution was run.

In order to determine how useful it is in general to incorporate human expertise in this manner, two further experiments were run. One was an ablation study where no RHEA was performed and the entire initial population was randomly orthogonally initialized. The second one was an extended seeding study where the Perfect Heuristic was queried with thresholds from 5% to 95% in increments of 5%, prescriptors were trained with backpropagation to imitate the results of each query, and the resulting 19 seed models were added to the initial population together with the original two extreme prescriptors. The original and the new experiments were each run 10 times. Pareto fronts for each generation were created by combining all 10 trials' candidates and then forming the Pareto front of the combined population. These Pareto fronts were also used to compute the hypervolume at each generation.

The results, shown in Figure 8a, demonstrate the power of incorporating human expertise. When the distilled seed models were removed, even with 10 trials, evolution was unable to discover the top half

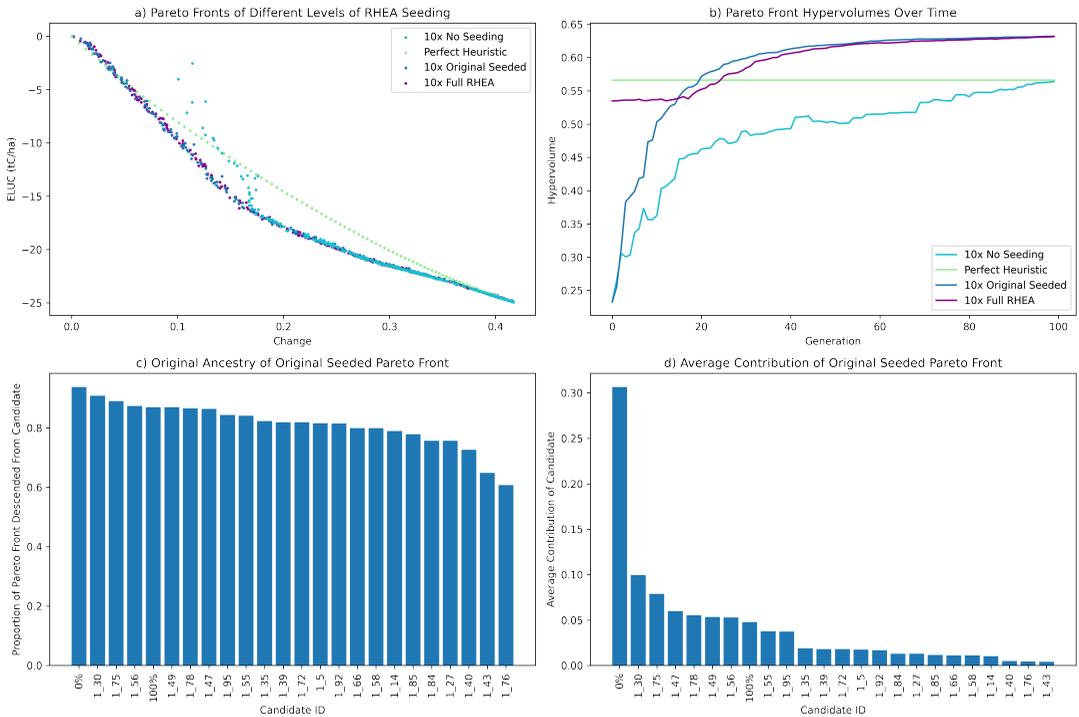


Figure 8. The value of incorporating human expertise. (a) Combined Pareto fronts over 10 trials for evolution run without initial seeding, the Perfect Heuristic, the original seeded Evolved Prescriptors (evolved with two extreme seeds), and a full RHEA evolution (with an extended seeding of the population) evaluated on the test dataset. The two extreme seeds achieved as strong results as extended seeding. However, without any seeds evolution fails to find the top half of the Pareto front. (b) The hypervolume comparison visualizes these results clearly as well. (c) The proportion of candidates in the Evolved Prescriptors Pareto front from Section 5.2 that have a given candidate from the first generation in its ancestry. The distilled 0% and 100% heuristics are ancestors of 94% and 87% of the Pareto front, respectively. (d) The average contribution of each first-generation candidate to the same Pareto front. The distilled 0% perfect heuristic makes up on average over 30% of total ancestry for the candidates in the final Pareto front. Thus, even limited seeding can have a significant effect on the quality of solutions, suggesting that RHEA is instrumental in constructing practical solutions to complex decision-making problems like land-use optimization..

of the Pareto front, missing out on the most realistic decision-making use cases (i.e. those with small changes). In Figure 8b, the hypervolume of the No Seeding experiment barely reaches that of the Perfect Heuristic. However, just including one seed on each end of the Pareto front, as in the original seeded experiment, allowed evolution to achieve better results than the Perfect Heuristic. Additional seeds in the full RHEA case did not significantly improve the final solutions. Thus, simply injecting two extreme seeds based on human expertise was sufficient for this task.

The results of RHEA in a single run of the original two-seed experiment in Section 5.2 are further illustrated in Figures 8c and 8d (the results with other trials were qualitatively similar). Figure 8c shows the proportion of the final Pareto front that has a given candidate from the first generation as an ancestor. The no-change and full-change seed behaviors are ancestors of 94% and 87% of the final Pareto front, thus contributing to almost every prescriptor. The magnitude of that contribution is measured in Figure 8d by constructing a tree of parents for each prescriptor in the final Pareto Front and counting how many times each first-generation candidate was included in it. On average a model in the Pareto front is made

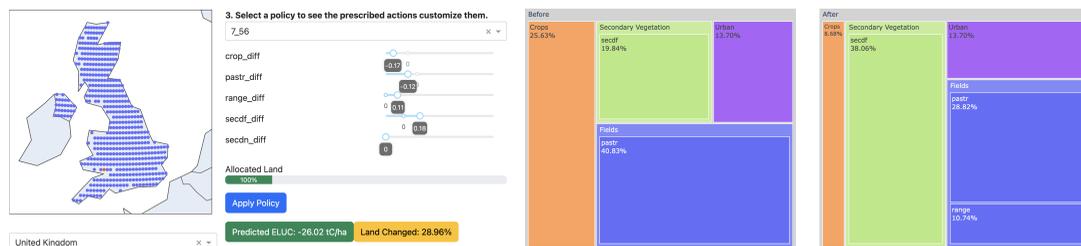


Figure 9. A suggested land-use change for a given location (screenshots from the demo at <https://landuse.evolution.ml>). The location is indicated by the red dot among the UK grid cells. One Evolved Prescriptor is chosen from the middle region of the Pareto front spanning minimal change and minimal ELUC. The current land use is shown on the left chart and the recommended one on the right chart, as well as the sliders on the left. This prescriptor recommends decreasing pasture and crops and increasing range and secondary forest, resulting in a 26.02 tC/ha decrease in carbon emissions with a 28.96% land-use change. The user can then select different solutions from the Pareto front and modify the sliders manually to explore alternatives.

up of over 30% of the distilled no-change behavior, mutated and crossed-over throughout the generations of evolution. This result aligns with the fact that the No Seeding experiment was unable to find the top left corner of the Pareto front: The seeded experiment needed to use the no-change solution frequently to fill in this top half. Very little behavior is taken from other randomly initialized models, showing the importance of the seeded model. Distilling this particular behavior was crucial in order to achieve the final Pareto front. The ultimate performance increase and explainability show how incorporating just a small amount of human feedback can improve the automatic decision-making process. In the future, as more real-world factors are taken into account, human expertise is likely to become even more valuable, and can be further applied with RHEA.

5.6. Interactive Evaluation

An interactive demo of the trained/evolved models is provided at <https://landuse.evolution.ml>. It allows the user to explore different locations, observe the prescribed actions for them, and see their outcomes. It is also possible to modify those actions and see their effect, thus evaluating the user’s skills as a decision-maker compared to that of machine learning.

As an example, Figure 9 shows the grid cells from which the user has selected one. The system recommends a change in land use that results in a 26.02 tC/ha reduction in carbon emissions with a 28.96% land-use change. Using the sliders, the user can then explore alternative tradeoffs and modify specific elements of each suggestion. For each variation, the user can then obtain the estimated carbon reduction and land-use change percentage, and in this manner gain an understanding of the possible attainable solutions and confidence in the preferred choice.

6. Future Work

The project so far has resulted in a promising proof-of-concept for utilizing ESP in land-use planning. In order to make it usable for real-world decision-makers, it needs to be extended with more fine-grained data. The models so far did not distinguish between the different types of crops and only partly considered regional variations in how land-use changes affect carbon stocks in soil and vegetation (through training different models for different regions). The assumed values for carbon stocks in soil and vegetation represent conditions of the recent past [14]; they do not consider environmental effects such as disturbances due to fires or droughts or better vegetation growth due to CO₂ effects or a warmer climate, which can substantially impact terrestrial CO₂ fluxes [e.g. 1, 27, 30]. Such additional features

could be added as context input to the predictor and prescriptor models, and the actions could be updated accordingly to prescribe more detailed policies. Further, estimates of CO₂ fluxes from land-use change also depend on the spatial resolution of the land-use change dataset [2, 11]. Extending the project with such details should lead to more precise and more fine-scale estimates of the impact of land-use change decisions, allowing decision-makers to utilize these dimensions in the real world.

Another immediate future work direction consists of improving the predictor accuracy through ensembling. Given that models trained on specific regions perform the best on those regions, there is a good chance that an ensemble of them will do even better. Moreover, there may be a benefit of including LinReg, RF, and NeuralNet versions of these models in the ensemble as well. There are several ensembling techniques in the Project Resilience codebase, such as averaging, sampling, and confidence-based ensembling, and contributors are encouraged to submit more in the future. The effect of such alternatives will be evaluated in future work, both on prediction accuracy and on driving the search for better prescriptors. The hypervolume of the Pareto front can be used to compare the different approaches quantitatively.

The predictor creates point estimates of ELUC. A technique called Residual Input-Output kernel [RIO; 28] can be used to estimate confidence in these predictions. The method builds a Gaussian Process model of the residual error, quantifying the uncertainty. RIO is part of the Project Resilience codebase but was not yet used in the study in this paper.

Project Resilience also provides an alternative approach to prescription: Representing the prescriptor as a set of rules [34]. This approach has the advantage that the behavior is transparent and explainable, as demonstrated in prior experiments in other tasks (including NPI optimization [22]). Rule-set evolution on land-use optimization will be evaluated in the future.

To make the system more immediately useful to decision-makers, the cost metric could be refined to take into account region-specific costs of changing each particular land-use type to another. The approach can also be extended with preferences and further objectives. For instance, the decision maker could choose which crops should be preferred, which may be more actionable than the large-scale changes in the current dataset. It may be possible to optimize additional objectives such as minimizing nitrate pollution and water consumption, and maintaining or increasing food production, similar to how minimizing crop change was optimized in Section 5.4. Further, instead of planning for a single year, it might be possible to develop prescriptors that recommend land-use changes for the next several years: For instance, what could be done in the next five years to reduce ELUC by a given target percentage. As data and simulations become more accurate in the future, such refined and extended decisions should become possible.

Even though land-use change is an important factor, it alone is not enough to reduce emissions to net zero. However, the framework designed in this project can be applied to other decision-making use cases as long as available context, action, and outcome data exist. A comprehensive such effort, covering energy policy and other factors, may eventually help reach climate-change goals.

7. Conclusion

Land-use policy is an area relevant to climate change where local decision-makers can have a large impact. In this paper, historical data and simulation technology are brought together to build an efficient machine-learning model that can predict the outcomes of these decisions efficiently for different contexts. An evolutionary optimization process is then used to identify effective policies, balancing the cost of land-use change and its effect in reducing carbon emissions. Machine learning methods can thus play an important role in empowering decision-makers to act on climate change issues.

A. Appendix: Using Location as a Proxy for Climate Information

Different locations have different emissions from land-use change due to the climate. They are affected by both the initial makeup of the land use and the conditions of the cell such as weather, soil, and type of

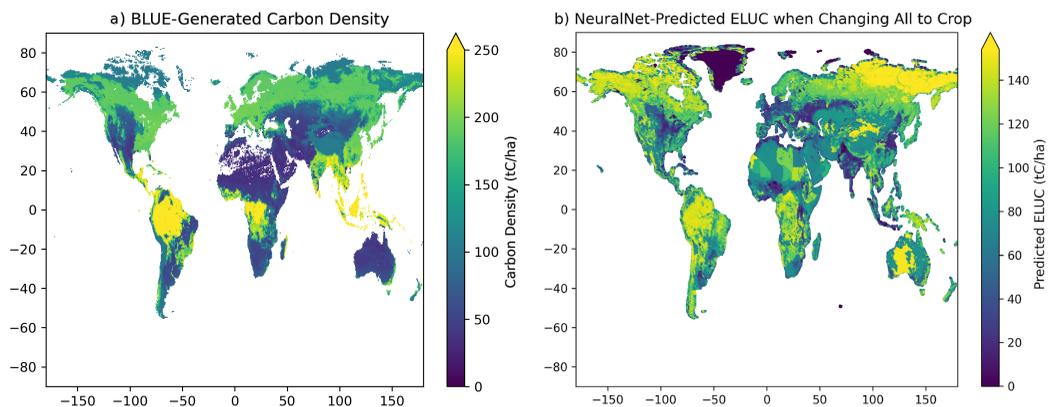


Figure 10. *Location as a Proxy for Climate. Different climates have different plant functional types (PFTs), which in turn have different carbon densities. (a) The carbon density, or potential carbon release, calculated with the BLUE model across the globe. (b) NeuralNet-predicted ELUC when each land cell is converted entirely to crop, acting as a proxy for carbon density. The same general qualitative trends appear across the globe with some outliers, showing the ability of the NeuralNet to model climate information implicitly through plant functional types.*

vegetation. The BLUE model takes this type of vegetation into account as plant functional types (PFTs) when calculating ELUC. In order to test if the predictor model encodes such climate information, an experiment was run to see if similar trends arose in both the predicted ELUC from the NeuralNet and the carbon densities (i.e. potential for carbon release) computed from the BLUE model using plant functional types. Similarities in the results would indicate that the predictor model is able to learn climate information from plant functional types even when they are not explicitly passed in as features.

The carbon densities from the BLUE model are shown in Figure 10a. To compute the corresponding values in the model, first note that crops have a significantly lower carbon density than vegetation (this effect can be seen e.g. in Figure 2, where conversion to crop from anything else results in a high ELUC). Thus, converting all land to crop would be similar to releasing all the carbon from the land, and the resulting ELUC can thus be seen as an estimate of its carbon density. To obtain this estimate, all land types were converted to crops globally across all data points in the year 2020, and the resulting ELUC predicted for this change using the NeuralNet. The result is plotted on a map in Figure 10b.

Similar qualitative trends arise in both plots, with core primary forest regions like the Amazon and the Congo Basin having high carbon densities and predicted ELUCs. Secondary forest regions such as the eastern US, southern Canada, and southern China were picked up by the predictor as having a more moderate carbon density/ELUC. Additionally, drier regions like Central Asia tended towards lower carbon densities and ELUC values. The scales do not perfectly align, however it is more important to view overall trends rather than numeric similarities.

However, qualitatively higher ELUC values than carbon densities appeared in the Great Victoria Desert, northeastern Siberia, and northern Africa. This result was due to the fact that LUH2 classifies these regions as almost entirely primary nonforest vegetation, which the NeuralNet erroneously predicted to release a large amount of carbon. LUH2 uses country-level Food and Agriculture Organization (FAO) statistics when there is not enough data, and therefore defaults to a uniform distribution. While the predictor was able to pick up on trends in the other types of vegetation, it overestimated the ELUC for such regions.

To conclude, similar trends arise overall in carbon density estimates in BLUE and the model, indicating that the model captures climate information even when it is not explicitly given to it as

features. Further alignment could be performed in the future using more specific data, however for the purposes of this proof-of-concept study, the amount of information encoded in the current features is sufficient.

Acknowledgments. We would like to thank the BLUE, LUH2, and Project Resilience teams, in particular Amir Banifatemi, Prem Krishnamurthy, Gyu Myoung Lee, Gillian Makamara, Michael O’Sullivan, Julia Pongratz, and Jennifer Stave.

Funding Statement. This work received no specific grant from any funding agency, commercial or not-for-profit sectors.

Competing Interests. The authors report no competing interests.

Code and Data Availability Statement. Code and data for the experiments are available at https://github.com/Project-Resilience/mvp/tree/main/use_cases/eluc.

Author Contributions. Conceptualization: RM, OF, EM, JB, HC, BH; Data Curation: CS, OF, DY, JB, HC; Formal Analysis: EM, OF, DY; Investigation: RM, OF, DY, EM, BH; Methodology: RM, OF, DY, EM, BH; Project Administration: RM, OF, BH; Software: OF, EM, DY; Supervision: RM, OF, BH; Validation: RM, OF, EM, DY, BH; Visualization: RM, OF, EM, DY; Writing – Original Draft: RM, OF, DY, BH; Writing – Review & Editing: RM, OF, DY, EM, CS, BH.

References

- [1] Luiz E. O. C. Aragão, Liana O. Anderson, Marisa G. Fonseca, Thais M. Rosan, Laura B. Vedovato, Fabien H. Wagner, Camila V. J. Silva, Celso H. L. Silva Junior, Egidio Arai, Ana P. Aguiar, Jos Barlow, Erika Berenguer, Merritt N. Deeter, Lucas G. Domingues, Luciana Gatti, Manuel Gloor, Yadvinder Malhi, Jose A. Marengo, John B. Miller, Oliver L. Phillips, and Sassan Saatchi. 21st century drought-related fires counteract the decline of amazon deforestation carbon emissions. *Nature Communications*, 9:536, 2018.
- [2] Selma Bultan, Julia E. M. S. Nabel, Kerstin Hartung, Raphael Ganzenmüller, Liang Xu, Sassa Saatchi, and Julia Pongratz. Tracking 21st century anthropogenic and natural carbon fluxes through model-data integration. *Nature Communications*, 13:4350, 2022.
- [3] Yaxuan Cai and Kongqing Li. Investigating the impact of land use management on terrestrial carbon emissions in nanjing city: past trajectories and multi-scenario simulation prediction. *Clean Technologies and Environmental Policy*, 26(8):2721–2738, 2024.
- [4] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishvi Aradhya, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ispir, Rohan Anil, Zakaria Haque, Lichan Hong, Vihan Jain, Xiaobing Liu, and Hemal Shah. Wide & deep learning for recommender systems, 2016.
- [5] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II. In *Proceedings of Parallel Problem Solving from Nature: PPSN VI, 6th International Conference*, pages 849–858, 2000.
- [6] Lea Dorgeist, Clemens Schwingshackl, Selma Bultan, and Julia Pongratz. A consistent budgeting of terrestrial carbon fluxes. *Nature Communications*, 15(1):7426, 2024.
- [7] Olivier Francon, Santiago Gonzalez, Babak Hodjat, Elliot Meyerson, Risto Miikkulainen, Xin Qiu, and Hormoz Shahrzad. Effective reinforcement learning through evolutionary surrogate-assisted prescription. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2020)*, pages 814–822, 2020.
- [8] P. Friedlingstein, M. O’Sullivan, M. W. Jones, R. M. Andrew, D. C. E. Bakker, J. Hauck, P. Landschützer, C. Le Quéré, I. T. Lujikx, G. P. Peters, W. Peters, J. Pongratz, C. Schwingshackl, S. Sitch, J. G. Canadell, P. Ciais, R. B. Jackson, S. R. Alin, P. Anthoni, L. Barbero, N. R. Bates, M. Becker, N. Bellouin, B. Decharme, L. Bopp, I. B. M. Brasika, P. Cadule, M. A. Chamberlain, N. Chandra, T.-T.-T. Chau, F. Chevallier, L. P. Chini, M. Cronin, X. Dou, K. Enyo, W. Evans, S. Falk, R. A. Feely, L. Feng, D. J. Ford, T. Gasser, J. Ghattas, T. Gkritzalis, G. Grassi, L. Gregor, N. Gruber, Ö. Gürses, I. Harris, M. Hefner, J. Heinke, R. A. Houghton, G. C. Hurtt, Y. Iida, T. Ilyina, A. R. Jacobson, A. Jain, T. Jarníková, A. Jersild, F. Jiang, Z. Jin, F. Joos, E. Kato, R. F. Keeling, D. Kennedy, K. Klein Goldewijk, J. Knauer, J. I. Korsbakken, A. Körtzinger, X. Lan, N. Lefèvre, H. Li, J. Liu, Z. Liu, L. Ma, G. Marland, N. Mayot, P. C. McGuire, G. A. McKinley, G. Meyer, E. J. Morgan, D. R. Munro, S.-I. Nakaoka, Y. Niwa, K. M. O’Brien, A. Olsen, A. M. Omar, T. Ono, M. Paulsen, D. Pierrot, K. Pockock, B. Poulter, C. M. Powis, G. Rehder, L. Resplandy, E. Robertson, C. Rödenbeck, T. M. Rosan, J. Schwinger, R. Séférian, T. L. Smallman, S. M. Smith, R. Sospedra-Alfonso, Q. Sun, A. J. Sutton, C. Sweeney, S. Takao, P. P. Tans, H. Tian, B. Tilbrook, H. Tsujino, F. Tubiello, G. R. van der Werf, E. van Ooijen, R. Wanninkhof, M. Watanabe, C. Wilmart-Rousseau, D. Yang, X. Yang, W. Yuan, X. Yue, S. Zaehle, J. Zeng, and B. Zheng. Global carbon budget 2023. *Earth System Science Data*, 15(12):5301–5369, 2023.
- [9] P. Friedlingstein, M. O’Sullivan, M. W. Jones, R. M. Andrew, L. Gregor, J. Hauck, C. Le Quéré, I. T. Lujikx, Olsen A., G. P. Peters, W. Peters, J. Pongratz, C. Schwingshackl, S. Sitch, J. G. Canadell, P. Ciais, R. B. Jackson, S. R. Alin, R. Alkama, A. Arneth, V. K. Arora, N. R. Bates, M. Becker, N. Bellouin, H. C. Bittig, L. Bopp, F. Chevallier, L. P. Chini, M. Cronin, W. Evans, S. Falk, R. A. Feely, T. Gasser, M. Gehlen, T. Gkritzalis, L. Gloege, G. Grassi, N. Gruber, Ö. Gürses, I. Harris,

- M. Hefner, R. A. Houghton, G. C. Hurtt, Y. Iida, T. Ilyina, A. K. Jain, A. Jersild, K. Kadono, E. Kato, D. Kennedy, K. Klein Goldewijk, J. Knauer, J. I. Korsbakken, P. Landschützer, N. Lefèvre, K. Lindsay, J. Liu, Z. Liu, G. Marland, N. Mayot, M. J. McGrath, N. Metzler, N. M. Monacci, D. R. Munro, S.-I. Nakaoka, Y. Niwa, K. O'Brien, T. Ono, P. I. Palmer, N. Pan, D. Pierrot, K. Pockock, B. Poulter, L. Resplandy, E. Robertson, C. Rödenbeck, C. Rodriguez, T. M. Rosan, J. Schwinger, R. Séférian, J. D. Shutler, I. Skjelvan, T. Steinhoff, Q. Sun, A. J. Sutton, C. Sweeney, S. Takao, T. Tanhua, P. P. Tans, X. Tian, H. Tian, B. Tilbrook, H. Tsujino, F. Tubiello, G. R. van der Werf, A. P. Walker, R. Wanninkhof, C. Whitehead, A. Willstrand Wranne, R. Wright, W. Yuan, C. Yue, X. Yue, S. Zaehle, J. Zeng, and B. Zheng. Global carbon budget 2022. *Earth Syst. Sci. Data*, 14:4811—4900, 2022.
- [10] P. Friedlingstein, M. O'Sullivan, M. W. Jones, R. M. Andrew, J. Hauck, P. Landschützer, C. Le Quééré, H. Li, I. T. Lujikx, A. Olsen, G. P. Peters, W. Peters, J. Pongratz, C. Schwingshackl, S. Sitch, J. G. Canadell, P. Ciais, R. B. Jackson, S. R. Alin, A. Arneeth, V. Arora, N. R. Bates, M. Becker, N. Bellouin, C. F. Berghoff, H. C. Bittig, L. Bopp, P. Cadule, K. Campbell, M. A. Chamberlain, N. Chandra, F. Chevallier, L. P. Chini, T. Colligan, J. Decayeux, L. Djeutchouang, X. Dou, C. Duran Rojas, K. Enyo, W. Evans, A., R. A. Feely, D. J. Ford, A. Foster, T. Gasser, M. Gehlen, T. Gkritzalis, G. Grassi, L. Gregor, N. Gruber, Ö. Gürses, I. Harris, M. Hefner, J. Heinke, G. C. Hurtt, Y. Iida, T. Ilyina, A. R. Jacobson, A. Jain, T. Jarníková, A. Jersild, F. Jiang, Z. Jin, E. Kato, R. F. Keeling, K. Klein Goldewijk, J. Knauer, J. I. Korsbakken, S. K. Lauvset, N. Lefèvre, Z. Liu, J. Liu, L. Ma, S. Maksyutov, G. Marland, N. Mayot, P. McGuire, N. Metzler, N. M. Monacci, E. J. Morgan, S.-I. Nakaoka, C. Neill, Y. Niwa, T. Nützel, L. Olivier, T. Ono, P. I. Palmer, D. Pierrot, Z. Qin, L. Resplandy, A. Roobaert, T. M. Rosan, C. Rödenbeck, J. Schwinger, T. L. Smallman, S. Smith, R. Sospedra-Alfonso, T. Steinhoff, Q. Sun, A. J. Sutton, R. Séférian, S. Takao, H. Tatebe, H. Tian, B. Tilbrook, O. Torres, E. Tourigny, H. Tsujino, F. Tubiello, G. van der Werf, R. Wanninkhof, X. Wang, D. Yang, X. Yang, Z. Yu, W. Yuan, X. Yue, S. Zaehle, N. Zeng, and J. Zeng. Global carbon budget 2024. *Earth System Science Data*, 2024. Preprint under review.
- [11] R. Ganzenmüller, S. Bultan, K. Winkler, R. Fuchs, F. Zabel, and J. Pongratz. Land-use change emissions based on high-resolution activity data substantially lower than previously estimated. *Environmental Research Letters*, 17:064050, 2022.
- [12] Georgios Giannarakis, Vasileios Sitokostantinou, Roxanne Suzette Lorilla, and Charalampos Kontoes. Towards assessing agricultural land suitability with causal machine learning, 2022.
- [13] Thomas Hale, Sam Webster, Anna Petherick, Toby Phillips, and Beatriz Kira. Oxford COVID-19 government response tracker. Blavatnik School of Government, 2020. <https://www.bsg.ox.ac.uk/research/covid-19-government-response-tracker>, accessed 11/20/2020.
- [14] Eberhard Hansis, Steven J. Davis, and Julia Pongratz. Relevance of methodological choices for accounting of land use change carbon fluxes. *Global Biogeochemical Cycles*, 29(8):1230–1246, 2015.
- [15] R. A. Houghton, J. E. Hobbie, J. M. Melillo, B. Moore, B. J. Peterson, G. R. Shaver, and G. M. Woodwell. Changes in the carbon content of terrestrial biota and soils between 1860 and 1980: A net release of CO₂ to the atmosphere. *Ecological Monographs*, 53:235–262, 1983.
- [16] G. C. Hurtt, L. Chini, R. Sahajpal, S. Frolking, B. L. Bodirsky, K. Calvin, J. C. Doelman, J. Fisk, S. Fujimori, K. K. Goldewijk, T. Hasegawa, P. Havlik, A. Heinimann, F. Humpenöder, J. Jungclaus, J. Kaplan, J. Kennedy, T. Kristzin, D. Lawrence, P. Lawrence, L. Ma, O. Mertz, J. Pongratz, A. Popp, B. Poulter, K. Riahi, E. Shevliakova, E. Stehfest, P. Thornton, F. N. Tubiello, D. P. van Vuuren, and X. Zhang. Harmonization of global land-use change and management for the period 850-2100 (LUH2) for CMIP6. Geoscientific Model Development Discussions, 2020. <https://doi.org/10.5194/gmd-2019-360>.
- [17] ITU. Project resilience, 2023. <https://www.itu.int/en/ITU-T/extcoop/ai-data-commons/Pages/project-resilience.aspx>, accessed 11/20/2023.
- [18] Stefan C. Dekker, Kees Klein Goldewijk and Jan Luiten van Zanden. Per-capita estimations of long-term historical land use and the consequences for global change research. *Journal of Land Use Science*, 12:313–337, 2017.
- [19] K. Klein Goldewijk, A. Beusen, J. Doelman, and E. Stehfest. Anthropogenic land use estimates for the holocene – hyde 3.2. *Earth System Science Data*, 9:927–953, 2017.
- [20] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.
- [21] Elliot Meyerson, Olivier Francon, Darren Sargent, Babak Hodjat, and Risto Miikkulainen. Unlocking the potential of global human expertise. In *Proceedings of the 38th Conference on Neural Information Processing Systems (NeurIPS 2024)*, 2024.
- [22] Risto Miikkulainen, Olivier Francon, Elliot Meyerson, Xin Qiu, Darren Sargent, and Elisa Canzani and Babak Hodjat. From prediction to prescription: Evolutionary optimization of non-pharmaceutical interventions in the COVID-19 pandemic. *IEEE Transactions on Evolutionary Computation*, 25:386–401, 2021.
- [23] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fiedjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumanar, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518:529–533, 2015.
- [24] National Academies of Sciences, Engineering, and Medicine. *Negative Emissions Technologies and Reliable Sequestration: A Research Agenda*. The National Academies Press, Washington, DC, 2019.
- [25] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library, 2019.

- [26] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [27] Julia Pongratz, Clemens Schwingshackl, Selma Bultan, Wolfgang Obermeier, Felix Havermann, and Suqi Guo. Land use effects on climate: Current state, recent progress, and emerging topics. *Current Climate Change Reports*, 7:99–120, 2021.
- [28] Xin Qiu, Elliot Meyerson, and Risto Miikkulainen. Quantifying point-prediction uncertainty in neural networks via residual estimation with an I/O kernel. In *Proceedings of the International Conference on Learning Representations*, 2020.
- [29] Md. Mostafizur Rahman and György Szabó. Sustainable urban land-use optimization using gis-based multicriteria decision-making (gis-mcdm) approach. *ISPRS International Journal of Geo-Information*, 11(5), 2022.
- [30] Thais M. Rosan, Stephen Sitch, Michael O’ Sullivan, Luana S. Basso, Chris Wilson, Camila Silva, Emanuel Gloor, Dominic Fawcett, Viola Heinrich, Jefferson G. Souza, Francisco Gilney Silva Bezerra, Celso von Randow, Lina M. Mercado, Luciana Gatti, Andy Wiltshire, Pierre Friedlingstein, Julia Pongratz, Clemens Schwingshackl, Mathew Williams, Luke Smallman, Jürgen Knauer, Vivek Arora, Daniel Kennedy, Hanqin Tian, Wenping Yuan, Atul K. Jain, Stefanie Falk, Benjamin Poulter, Almut Arneth, Qing Sun, Sönke Zaehle, Anthony P. Walker, Etsushi Kato, Xu Yue, Ana Bastos, Philippe Ciais, Jean-Pierre Wigneron, Clement Albergel, and Luiz E. O. C. Aragão. Synthesis of the land carbon fluxes of the Amazon region between 2010 and 2020. *Communications Earth & Environment*, 5:46, 2024.
- [31] Donald B Rubin. Causal inference using potential outcomes. *Journal of the American Statistical Association*, 100(469):322–331, 2005.
- [32] Andrew M Saxe, James L McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural network. In *Proceedings of the Second International Conference on Learning Representations (ICLR)*, 2014.
- [33] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv:1707.06347*, 2017.
- [34] Hormoz Shahrzad, Babak Hodjat, and Risto Miikkulainen. Evoter: Evolution of transparent explainable rule-sets. *arXiv:2204.10438*, 2022.
- [35] Kenneth O. Stanley, Jeff Clune, Joel Lehman, and Risto Miikkulainen. Designing neural networks through evolutionary algorithms. *Nature Machine Intelligence*, 1:24–35, 2019.
- [36] XPRIZE. Pandemic response challenge, 2023. <https://www.xprize.org/challenge/pandemicresponse>, accessed 11/20/2023.
- [37] Ronald R Yager. On ordered weighted averaging aggregation operators in multicriteria decisionmaking. *IEEE Transactions on systems, Man, and Cybernetics*, 18(1):183–190, 1988.