Graphical Abstract

Using Political-Ecological Models to Sustain Biodiversity

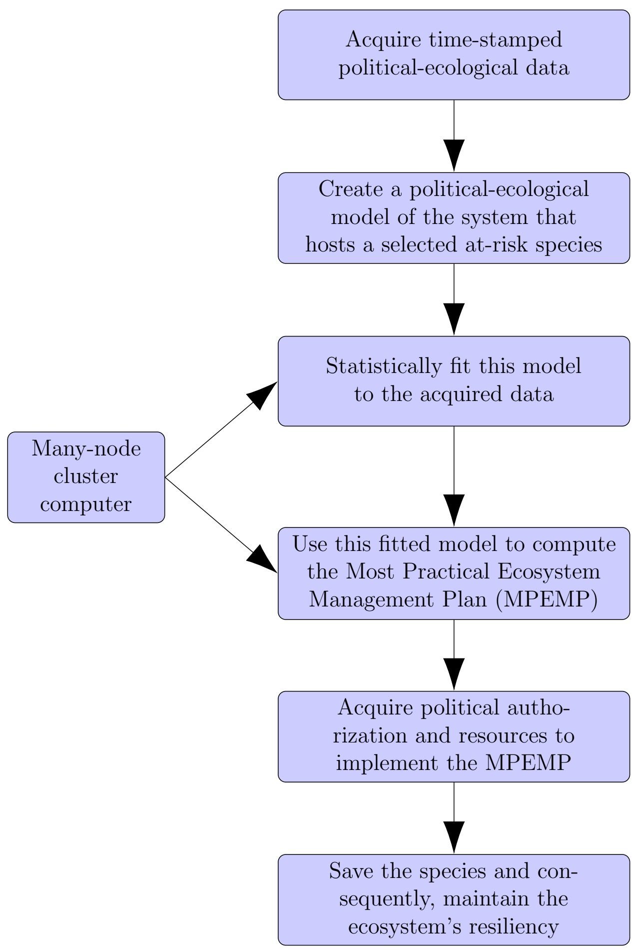

Highlights

Using Political-Ecological Models to Sustain Biodiversity

- A new toolkit finds politically feasible and effective ecosystem management plans.
- A new cluster computing optimization algorithm is described.
- This algorithm permits the computation of these ecosystem management plans.
- Statistical estimation of political-ecological model parameters also become possible.
- A new model captures the East African cheetah-hosting political-ecological system.

Using Political-Ecological Models to Sustain Biodiversity

Abstract

Because of the finality of a species' global extinction, there is a need to focus on stopping such extinction events from happening. The way forward is to find and implement politically feasible and ecologically effective projects that head off extinction events. This article delivers a software toolkit that implements one way to do this. This toolkit provides an organization the means to (1) build a political-ecological model; (2) fit this model to a political-ecological data set; and finally, (3) use this model to compute the *most practical ecosystem management plan* (MPEMP). This model-based approach to first understanding the political issues surrounding the conservation of a particular species and then second, finding a conservation plan that works with these political realities – is hamstrung by the expensive computations needed to first, fit a political-ecological model to data and then second, compute the MPEMP from this fitted model. Therefore, a new optimization algorithm is presented that overcomes this challenge when run on either a commercial or home-grown cluster computer. This new algorithm finds the global solution to an optimization problem characterized by constraints and a black-box, stochastic objective function. This toolkit is illustrated by finding the MPEMP for conserving the cheetah (*Acinonyx jubatus*) population across Kenya and Tanzania.

Keywords: extinction crisis, biodiversity conservation, political-ecological models, model credibility, ecosystem management, robust statistical estimators, optimization algorithms, high performance computing

Abbreviations:

CA: Consistency Analysis
IBM: Individual Based Model
MDAS: Multiple Dimensions Ahead Search
MPEMP: Most Practical Ecosystem Management Plan

NGO: Nongovernmental Organization
PACSA: Parallel Asynchronous Coupled Simulated Annealing
SA: Simulated Annealing
SA-MDAS: Simulated Annealing – Multiple Dimensions Ahead Search
SDE: Stochastic Differential Equation
TSCC: Triton Shared Computing Cluster

1. Introduction

Earth is in the midst of an extinction crisis (Torres-Romero et al., 2024), (Garber et al., 2024) with many species becoming globally extinct every year (Verones et al., 2022). Global extinction of a species is irreversible. Assuming that the end goal of environmental and ecological conservation efforts is to sustain ecosystems in states that are close to those circa 2025, it can be argued that conservation efforts that directly contribute to stemming irreversible ecosystem state changes should be given the highest priority both politically and financially. To support these efforts, this article argues that the highest priority should be given to (a) developing *political-ecological* models that can credibly gauge likelihoods of global extinction events; and (b) developing ecosystem management plans based on these models that can reduce these likelihoods.

This article argues that *biodiversity conservation* efforts in particular, should focus on understanding the political-ecological processes that lead to global extinction events so that politically feasible and ecologically effective ecosystem management plans can be identified that, when implemented, have the greatest chance of stopping these events from happening.

This article describes a free software toolkit for developing such plans. Armed with this toolkit, an organization first statistically estimates the parameter values of a political-ecological model. They then use this fitted model to compute the *most practical ecosystem management plan* (MPEMP) (Haas, 2011, Chapter 4). Accordingly, this toolkit is referred to here as the *MPEMP constructor*.

This article takes a broad view of what constitutes biodiversity to include species that may not have any commercial value but are listed as Vulnerable or Endangered on the IUCN Red List (Panwar et al., 2023), (Testa et al., 2025), (Tobias et al., 2025). Such species include many terrestrial mammals, marine mammals, fish, and many plants.

The private sector, although paying lip service to the conservation of such species, actually has a strong economic interest in biodiversity conservation that is focused mainly on the preservation of potentially profitable genetic resources. The commercial value of these resources was approximately USD44 trillion in 2022 (Medlong et al., 2022). The total budget of NGOs engaged in preserving endangered species by comparison, is in the neighborhood of USD12 billion (Wan, 2023). The efforts of the many professionals engaged in work to curb biodiversity loss, although laudable, is dwarfed by the efforts in the private sector to profit from biodiversity resources. In other words, the big money is being spent on preserving species that have high commercial potential rather than on species that, in-part, define the word, "wild."

Some efforts are being made, however, to encourage more private sector spending on species who have no commercial value but do have (1) *existence value*, (2) *bequest value*, or (3) *ecosystem function value*, i.e., the species performs important ecosystem-support functions. Existence value is the perceived value of knowing that a species exists, and bequest value is the perceived value of conserving a species for future generations (Ressurreição et al., 2012).

Beverdam et al. (2025) for example, call for a "blending" of private and public sector funds to conserve species that are not commercially valuable. An example of this type of financing mentioned by these authors is the so-called "Rhino Bond" that targets the conservation of the commercially valueless black rhinocerous (*Diceros bicornis*) in Africa.

1.1. The way forward

The concept that this article is offering a realization of, is a toolkit that can find conservation plans that are ecologically effective and politically feasible to implement. After describing how this concept is realized in software, this article gives a proof of this concept by finding a management plan for the East African cheetah that requires minimal changes to the political beliefs of those groups surrounding this biodiversity conservation challenge while predicting that the cheetah population will remain viable through the planning horizon.

As an example of this article's approach to biodiversity conservation, Haas and Ferreira (2017) build an agent/individual based model of the political-ecological system that surrounds rhinoceros (*Ceratotherium simum*) poaching in South Africa. These authors simulate several management options and predict that under current management policies, rhinos in South Africa will become extinct around the year 2036.

Such models, however, are in their infancy but need to quickly mature because at present extinction rates, many species including most large mammals will be globally extinct by about 2055 (Ceballos et al. , 2015). In other words, if present trends continue, a significant proportion of the earth's species will soon be gone due mainly to the activities of the most recent (circa 2025) several generations of humans.

This article gives a tested toolkit, namely, the MPEMP constructor that can be used to (a) model such political-ecological systems; and (b) based on such models, compute ecosystem management plans that, when implemented, fend off species extinction events. The MPEMP constructor streamlines the construction of an integrated model composed of agent-based submodels of political processes that interact with an individual-based population dynamics submodel of an at-risk species.

What mechanisms of biodiversity loss would such models represent? Habitat loss is often pointed to as the principal driver of global biodiversity loss (Hanski, 2011). But recently, one study could find no statistical difference between loss of habitat and *direct exploitation*, i.e., intentional harvesting of wildlife either legally or illegally (Jaureguiberry et al., 2022). Illegal harvesting and illegal trading of wildlife is referred to as *wildlife trafficking*. The scale of exploitation-curbing biodiversity conservation projects that is needed to slow the globe's ongoing biodiversity losses, is much larger than the sum total of currently active projects. To help address this disparity, the MPEMP constructor is engineered to focus on the modeling of direct exploitation drivers of biodiversity loss.

Recently, the literature has called for a more holistic view of biodiversity conservation rather than a focus on single species preservation (Tobias et al., 2025). This view, called "process-based" by Tobias et al. (2025) argues that maintaining ecological processes that drive critical ecosystem functioning, should be the primary goal of biodiversity conservation initiatives. These processes include adaptation, gene flow, dispersal, and trophic interactions.

Acknowledging that the goal of avoiding the extinction of species that have evolved over millions of years is not to be abandoned, Tobias et al. (2025) call for a synthesis of species-centric and process-based approaches to biodiversity conservation:

> "We advocate integration and communication across the two primary cultures of conservation—species-centric and process-based—as the most effective progress will occur when these two missions operate in tandem and synergistically ((Tobias et al., 2025))."

Process-based approaches to biodiversity conservation, however, have yet to wrestle with direct exploitation effects on an ecosystem's functioning.

Such integration is straightforward within the methods detailed in this article. For instance, *functional diversity* (FD) is one way to describe the healthy functioning of an ecosystem. FD can be quantified from remotely-sensed NDVI data (Li et al., 2025). And, *trophic transfer* between predators and their prey can be quantified with the *AB ratio* reviewed in Carroll et al. (2019). These two metrics can easily be added to the species abundance objective function employed in this article to arrive at an analysis and planning workflow that integrates species-centric goals and process-based goals.

The MPEMP maximally increases the probability of a species' survival while requiring the least change in the beliefs held by identifiable groups of humans in those countries that host a selected at-risk species. Indeed, according to Haas (2024c), any project that is intended to sustain biodiversity needs to have its MPEMP computed so that the conservation project can either be modified to enhance its conservation effectiveness – or abandoned altogether and replaced by a project that implements the MPEMP. The MPEMP constructor guides the user through this MPEMP computation.

An organization would use the MPEMP constructor to complete the following five steps.

1. Acquire a data set that pertains to the political-ecological system surrounding a selected at-risk species.
2. Create a political-ecological model of this system.

3. Compute statistical estimates of this model's parameters.
4. Display and assess the political-ecological actions generated by this fitted model.
5. Compute the MPEMP from this fitted model and implement it.

This five-step procedure has application in a business-oriented approach to conserving a selected at-risk species (Haas, 2022), (Haas, 2024c). Briefly, revenue from a commercial *offering* funds a *biodiversity project*. This project operationalizes the MPEMP that has been computed for that species and the political-ecological system that hosts it.

This article proceeds as follows. A description is given in Section 2 of this *interacting influence diagrams* modelling architecture (Haas, 2011, Chapter 2). Then, as an example, this architecture is used to build a model of the political-ecological system that hosts the East African cheetah (*Acinonyx jubatus*) population. A new cluster computer-based optimization algorithm is presented in Section 3 that has developed out of proposals advanced by Haas (2024). The above model is fitted via this new algorithm to a *political-ecological actions history* data set in Section 4. Statistical estimates of this model's parameter values found via Consistency Analysis (CA) (Haas, 2011, Chapter 3) are computed by running the new optimization algorithm on the Triton Shared Computing Cluster (TSCC) at the San Diego Supercomputer Center (San Diego Supercomputer Center, 2025). In Section 5, this statistically estimated model is used to compute the associated MPEMP. Section 6 contains a list of challenges that need to be overcome before such models can be widely used to identify extinction-avoiding management plans. Some conclusions are reached in Section 7.

2. An Integrated Political-Ecological Model

2.1. Why do submodels need to be integrated?

Given a model composed of political submodels and an ecosystem submodel, if actions that are generated by political submodels do not impact the ecosystem submodel at model-generated time points – and ecosystem submodels actions do not in-turn, affect these political submodels, then feedback loops between political processes and ecological processes cannot emerge.

For instance, if a climate model is exogenously set to a schedule of reductions in carbon emissions forcing, there is no mechanism for modelling the vicissitudes of political support for climate policy. Such vicissitude might take the form of a sequence of governmental administrations alternately supporting and then not supporting climate change mitigation policies. The jagged time series of carbon emissions that results from this support-nonsupport policy record cannot be realistically simulated unless this sequence of different administrations and the effects of their actions on earth's climate is represented in the simulation model. Hence, there needs to be an *integrated* model of political processes interacting with ecological processes. In this article, the main ecological process that needs to be modeled is the population dynamics of an at-risk species.

2.2. Advantages of an individual-based model of an at-risk species

As Netz et al. (2022) note,

> "To allow for mathematical analysis, models of predator–prey coevolution are often coarse-grained, focussing on population-level processes and largely neglecting individual-level behaviour. As selection is acting on individual-level properties, we here present a more mechanistic approach: an individual-based simulation model for the coevolution of predators and prey on a fine-grained resource landscape, where features relevant for ecology (like changes in local densities) and evolution (like differences in survival and reproduction) emerge naturally from interactions between individuals" (Netz et al., 2022).

2.3. Example: The East African cheetah population

The cheetah is listed as Vulnerable on the IUCN Red list and as Endangered by the Namibian government (Milloway, 2025). World Population (2025) reports 938 cheetah in Tanzania and 715 in Kenya. These two countries share a border and hence have the potential of interacting with each other politically. Therefore, as an example, a political-ecological model is built of the cheetah-hosting political-ecological system enclosed by these two countries. This model consists of submodels of several groups that affect the cheetah population, and an ecosystem submodel of East African cheetah population dynamics. All of these submodels interact with each other through time.

This agent/individual-based political-ecological model of the cheetah-hosting ecosystem contained within Kenya and Tanzania, is new.

2.3.1. Submodels interact through a bulletin board

At each time step, each group submodel and the ecosystem submodel read actions directed against themselves from a bulletin board. Conditional on a read-in input action, a group submodel computes the expected value of *overall goal attainment* that they believe they will receive if they implement a particular action-target combination. After making this computation for all action-target combinations in their repertoire, they post to the bulletin board the combination that has the highest expected overall goal attainment. The Ecosystem Management Actions Taxonomy (EMAT) (Haas, 2024b) dictates what actions a submodel recognizes and what it holds in its repertoire of output actions.

The ecosystem submodel reacts to actions directed against it by adjusting its output of cheetah and prey abundance across time.

2.3.2. Group submodels

Submodels represent Kenya's presidential office, the Kenya Wildlife Service, the rural residents of Kenya, and the pastoralists of Kenya. Similar submodels are constructed for Tanzania. A ninth group submodel is constructed of a conservation-focused nongovernmental organization (NGO) that runs conservation projects in both of these countries. Haas (2011, Chapter 2), Haas and Ferreira (2017), and Haas (2025) detail the cognitive theory and causal flow that these submodels use to decide what action to implement based on their beliefs and those actions that have been directed against them (called *input actions*). These group decision making submodels make decisions that they believe will further their own set of goals without regard to the goals of other groups and, except for the wildlife protection agency groups, without regard to what effects their actions might have on the ecosystem or the abundance of any particular species.

2.3.3. Cheetah-hosting ecosystem submodel

The ecosystem submodel tracks the dynamics of the trans-Kenya-Tanzania cheetah population (abundance through time). This submodel interacts with the above rural resident and pastoralist group submodels. Cheetah are individually modeled following the individual-based model (IBM) paradigm. For simplicity, however, this IBM interacts with a stochastic differential equation (SDE) submodel of cheetah prey such as Thomson's gazelle (*Gazella thomsoni*) (Fitzgibbon, 1990). Both the cheetah IBM and the herbivore SDE represent population dynamics stochastically.

This synthetic predator-prey submodel is new and extends a model described in Kimbrell and Holt (2005).

Parameterization of this submodel follows a set of data-based values reported in Kelly et al. (1998):

> "Data are presented on the demography and reproductive success of cheetahs living on the Serengeti Plains, Tanzania over a 25-year period. Average age at independence was 17.1 months, females gave birth to their first litter at approximately 2.4 years old, interbirth interval was 20.1 months, and average litter size at independence was 2.1 cubs. Females who survived to independence lived on average 6.2 years while minimum male average longevity was 2.8 years for those born in the study area and 5.3 years for immigrants" (Kelly et al., 1998).

Working from this information, Table 1 contains the parameter values used in the cheetah IBM and the prey SDE.

Parameter	Value
Cheetah	
Life expectancy	6.2 (females); 2.8 (males)
Age at maturity	1.42
Sexual maturity	2.4 (females)
Interbirth interval	1.675
Herbivores	
Initial abundance (minor poaching)	2250
(moderate poaching)	4500
(severe poaching)	8750
Birth rate − death rate (minor poaching)	-0.003
(moderate poaching)	-0.014
(severe poaching)	-0.017
White noise multiplier (minor poaching)	0.001
(moderate poaching)	0.001
(severe poaching)	0.001

Table 1: Ecosystem submodel parameter hypothesis values. Cheetah IBM parameter values are derived from Kelly et al. (1998). The temporal unit is *years*. The submodel uses only the averaged cheetah life expectancy (4.5 years). Submodel output of herbivore abundance depicts a negative trend through time that is influenced by the amount of herbivore poaching.

A data file is created based on estimated cheetah abundance reported in World Population (2025).

3. A New Optimization Algorithm

A new optimization algorithm is introduced herein that combines Simulated Annealing (SA) (Corana et al., 1987) and Multiple Dimensions Ahead Search (MDAS) (Haas, 2020). This new algorithm is called Constrained Stochastic Synchronous Coupled Simulated Annealing – MDAS (SA-MDAS).

3.1. Background

Consider an optimization problem wherein an objective function in n dimensions (hereafter, *variables*) is to be *minimized*. A *black-box optimization algorithm* can handle all types of deterministic objective functions without assuming anything about their smoothness or presence of discontinuities. To-date, however, there has been little work on the development of algorithms that can optimize a noisy (hereafter, *stochastic*) black-box objective function.

As reviewed in Haas (2024), one algorithm that is capable of finding the *global* solution to a deterministic black-box objective function is Parallel Asynchronous Coupled Simulated Annealing (PACSA) of Gonçalves-e-Silva et al. (2018). This algorithm however, does not recognize constraints and is not designed for stochastic objective functions. Haas (2024) proposes for future work, the development of a general purpose black-box optimization algorithm that would handle *bound constraints* and, by incorporating the algorithm of Branke et al. (2008) into PACSA, stochastic objective functions. This envisioned program would compute statistical estimates of the parameters of a political-ecological model. As explained next, however, a combination of theory and computational experience has led to the development of an algorithm that is different than the one envisioned in Haas (2024).

One final background note: In the eighteenth century, the person who organized an opera production was referred to as an *impresario*. The individuals who performed the opera be they singers, dancers, instrumentalists or conductors, were (and are) referred to as *performers* (Holmes, 1994). These terms are used here to refer to different roles given by an optimization algorithm to different compute nodes who collectively, make up a cluster computer (Werstein et al., 2006).

3.2. Rationale for a new algorithm

In SA-MDAS, an impresario posts tasks to a *JavaSpace* (Haas, 2020) for performers to take and complete. Once completed, a performer posts a task's results back to the JavaSpace. The impresario then takes these results from the JavaSpace and uses them to decide where next to search. Here, a JavaSpace is implemented via the *GigaSpaces* XAP (Ciatto et al., 2020). SA-MDAS is implemented in the JAVA language because JAVA is (a) computationally efficient; (b) easily parallelized; and (c) easy to read and hence, easy to maintain (Haas, 2020).

When the cooling schedule of Aarts and Korst (1989) is employed within PACSA, computational experience has shown that the use of equation (1) in

Gonçalves-e-Silva et al. (2018) causes the maximum accepted function value term to become large as iterations increase. This happens because some of the performers begin to return poor (large) objective function values when their solution chains enter suboptimal subspaces. This increased value of the maximum-function term causes the acceptance probability to become small – causing progress towards an optimal solution to essentially stop. For this reason, the asynchronous inter-performer coupling scheme developed in Gonçalves-e-Silva et al. (2018) is not used in SA-MDAS.

Regarding stochastic objective functions, Bouttier and Gavra (2019) provide a convergence proof for their approach to handling such functions within an SA optimization algorithm. For this reason, the Bouttier and Gavra (2019) approach to handling stochastic objective functions is employed in SA-MDAS rather than the approach taken by Branke et al. (2008).

3.3. Algorithm summary

SA-MDAS uses both SA and MDAS to solve an optimization problem that has a (possibly) stochastic black box objective function and continuous variables that are bound-constrained. Implicit constraints are also accommodated. To ensure finite-time convergence, SA-MDAS employs the cooling schedule of Aarts and Korst (1989, chapter 4).

SA-MDAS employs many performers in order to both increase the chance that a global solution is found, and to reduce the time it takes to find it, called the *wall clock time* (Jiang and Singh, 2010). This is accomplished in-part through a new method developed herein called Multiple Periodic Exchange (MPE) – a scheme similar to the class of parallel SA algorithms dubbed *periodic exchange schemes* by Lee and Lee (1996).

A well-known characteristic of SA is its slow rate of convergence (Guilmeau et al., 2021). This drawback is addressed in SA-MDAS by switching the search algorithm from a parallel SA-like search algorithm to a parallel Hooke-Jeeves-like search algorithm namely, MDAS.

The result of this *two-phase* architecture (Ferreiro et al., 2019) is an algorithm that takes advantage of SA's global search capability enhanced through cluster computing – but that avoids SA's proclivity for slow convergence by switching to an efficient local search algorithm (MDAS) that also leverages a cluster computing environment. The switch happens only after significant effort has been directed towards finding a subspace that has a high likelihood of containing the global minimum. Specifically, this switch is delayed until the SA step length along any variable becomes so small that the likelihood of a move into a radically different subspace also becomes small.

The Appendix contains details of how SA-MDAS handles continuous variables, stochastic objective functions, and messaging between compute nodes. This Appendix also contains comparisons of SA-MDAS with other optimization algorithms.

3.4. Algorithm

1. Global search phase:
   (a) Set values for n_s (see Appendix), and the chain length. Also, find initial values for each variable's step length, and the control parameter t ("temperature"). Initialize this latter parameter so that the percentage of moves accepted is between 70 and 90. Run this task on the impresario compute node alone.
   (b) Within SA's inner loop, start each of m performer nodes at the same initial solution.
   (c) Run these performers simultaneously but independently over one Markov chain.
   (d) Always accept solutions that deliver a score value that is smaller than *current_value*. Accept other solutions with probability

$$\exp\left\{-(trial\_value - current\_value)/t\right\}. \tag{1}$$

   (e) Block until all performers have finished their respective chains.
   (f) Update the sample size used to compute the average value of a stochastic objective function. Also update each variable's step length.
   (g) Update t via the rule developed in Aarts and Korst (1989, Chapter 4).
   (h) Return if a solution has been found such that the score function has been reduced to 90% of its initial value. Otherwise, continue to Step *i*.

(i) Execute MPE by reinitializing every performer with the solution that gave the current smallest score value.
(j) Go to Step *c*.

2. Local search phase:
   (a) After the SA subalgorithm has returned, use the returned solution as the starting solution in an MDAS run. Set the MDAS algorithm to search forward three variables at a time.
   (b) Run MDAS to convergence and then stop.

3.5. SA-MDAS performance on analytic objective functions

Haas (2020) reports that MDAS correctly solves Bukin's F4 function (Mishra, 2006). Using $n_s = 60$ (Corana et al., 1987), a chain length of 10, and 10 performers, SA-MDAS finds the global minimum solution to Bukin's F4 function in 12,353 objective function evaluations. SA-MDAS fails to find the global minimum solution for Bukin's F6 function. But when the global search phase of SA-MDAS is allowed to run to convergence rather than switching to local search, the global minimum of this function is found after 4,716,030 function evaluations. Bukin's F4 function has a pathological number of *non-global minima* (Hasanzadeh et al., 2022) – as does Bukin's F6 function. Clearly, SA-MDAS is capable of finding the global minimum of a highly multi-modal function when nearly unlimited computing power is available.

Real-world functions are not necessarily as pathological as Bukin's F6. Hence, the performance of SA-MDAS on Bukin's F4 gives gives some credence to the idea that it can find nearly optimal solutions to real-world conservation optimization problems. To support this supposition, SA-MDAS is assessed in the next Section by seeing how well it statistically estimates the parameters of a political-ecological model built to represent a real world conservation challenge.

4. Statistical Estimation of a Political-Ecological Model

4.1. Overview of the CA estimator

The CA estimator of Haas (2011, Chapter 3) produces a set of *consistent* parameter values such that model output balances agreement with data versus agreement with cognitive theories of decision making. To do this, an objective function is defined to be the priority-weighted sum of standardized measures of the agreement with a data set (g_S), and the agreement with the model's joint probability distribution across its stochastic nodes (hereafter, *distribution*) under *hypothesis* parameter values (g_H). The value of $ch \in (0,\ 1)$ is the relative priority given to having the *consistent* (estimated) model's distribution agree with the one specified by the hypothesis parameter values – versus agreeing with the data (here, an observed actions history): $g_{CA} = (ch)g_H + (1 - ch)g_S$.

When ch is 0, CA becomes a frequentist statistical estimator. When ch is 1, CA finds parameter values that result in the model maximally agreeing with both decision making theory and with ecological theory that collectively, dictate how the political-ecological system ought to behave.

CA fits a model in two stages. Stage I consists of finding parameter values so that the model matches as many observed actions as possible. Then, this Stage I percentage of action-agreement (match fraction) is computed. Next, Stage II adjusts these Stage I parameter values until the model's distribution is as close as possible to the hypothesis distribution while maintaining Stage I's match fraction. The idea behind this two-stage approach is to minimize the occurrence of *jump discontinuities* (Schober and Prestin, 2023) caused by one or more group submodels switching to different action-target combinations due to a small change in a parameter's value as the algorithm evaluates the objective function at different points in the solution space.

4.2. Objective function

Because of the above-mentioned potential for jump discontinuities, it is important to not accept any trial moves that reduce the match fraction from that achieved in Stage I. This agreement is enforced by adding a large penalty to the objective function if a move causes a reduction in the agreement with the observed actions history. Use of such a penalty function to, in-effect, represent an implicit constraint, works in SA-MDAS because numerical instabilities that might be caused by explosive numerical derivative computations cannot happen because SA-MDAS does not compute such derivatives.

Stage II uses a stochastic agreement function for g_H that is the negative of the Hellinger distance between the consistent and hypothesis distributions.

Letting $P(P = i) = p_i$ (and $P(Q = i) = q_i$), for two discrete distributions, this distance is

$$H(P, Q) = \sqrt{\frac{1}{2}\sum_{i=1}^{d}(\sqrt{p_i} - \sqrt{q_i})^2} \tag{2}$$

where d is the number of discrete values that the random variable P (and Q) can take on (Suresh, 2021).

4.3. Stage I's action-matching algorithm

Stage I finds starting parameter values by executing two subalgorithms. The first subalgorithm entails the sequential matching of action-target combinations with those observed. This subalgorithm is as follows.

1. At each time point, check the match between the observed action-target combination and the one generated by the group submodel. If they do not match, replace the submodel's action-target combination with that observed.
2. If this replacement causes the overall fraction of matches to become smaller, reject this replacement.
3. If the end-time has not been reached, go to the next group submodel or next time point. Otherwise, write this *desired actions history* to a file and exit.

The second subalgorithm proceeds by adjusting submodel parameters until model output matches as many action-target combinations in the desired actions history as possible – regardless of its agreement with the hypothesis distribution.

4.4. Results

SA-MDAS is used to statistically estimate parameter values of the political-ecological model of East African cheetah trafficking as follows.

4.4.1. Actions history

The STAR protocol of Haas (2024b) is used to collect actions reported in online press sources concerning cheetah management in Kenya and Tanzania. Doing so yields 105 actions over the time period 2019 through May 2025 (Table 2).

Date	Actor	Action	Target(s)
2019.55	kenpas	starve_due_to_drought	kenpas
2019.71	kenrr	poach_for_food	three,kenpres,kenepa,ngo
2019.99	tanrr	poach_for_food	one,cheetaheco
2021.20	ngo	fund_rural_development_project	one,chetaheco
2022.17	kenepa	translocate_animals	one,chetaheco
2022.34	tanrr	plant_trees	one,cheetaheco
2023.67	kenrr	poach_for_cash	one,chetaheco
2023.73	kenrr	poach_for_food	three,kenpres,kenepa,ngo
2024.49	kenrr	poach_for_cash	one,chetaheco
2025.18	kenrr	poach_for_cash	one,chetaheco

Table 2: Selected political-ecological actions extracted from online sources with the STAR protocol of Haas (2024b)

.

4.4.2. Model fitting

The parameters that determine how likely Kenya rural residents and Tanzania rural residents think they will be arrested after they poach cheetah for cash are adjusted so that the model's output agrees with the observed actions history while exhibiting a joint probability distribution across its nodes that maximally agrees with the distribution defined to represent theoretical propositions about these beliefs. Adding-in parameters to represent Kenya pastoralist and Tanzania pastoralist beliefs is planned for future work.

Here, the hypothesis distribution values of these parameters represent the theory that a typical rural resident believes there is little chance they will have any interaction with police after they poach a cheetah for cash. Equal priority is given to agreement with the hypothesis distribution versus agreement with observations ($ch = 0.5$).

This CA is implemented as a 54-variable optimization problem. These variables are the parameters that determine the Scenario Imminent Interaction with Police (SIIWP) node in the Kenya rural resident submodel and the Tanzania rural resident submodel under the conditioning event that a rural resident decides to implement the output action, *poach cheetah for cash*.

SA-MDAS employed three performers with each performer accessing 10 *threads* when performing a parallel Monte Carlo simulation. Each such simulation required 1000 simulated realizations of a submodel's stochastic nodes. Such a simulation was run whenever a submodel received a new input action-target combination at some step in the time interval that the model was run over.

Because parameters in the ecosystem submodel were not being adjusted, in order to reduce wall clock time, ecosystem submodel calculations were turned off during this run of the SA-MDAS algorithm. Doing so reduced the objective function's evaluation time from 44 seconds to 15 seconds.

Due to a modest computing budget, the SA-MDAS algorithm was allowed to run for up to 3000 function evaluations during each phase. The run's wall clock time was 6.6 hours. Phase one converged after 2991 function evaluations, and phase two was terminated after completing 3043 function evaluations due to exceeding the maximum number of function evaluations. After the run had finished, the algorithm had increased the value of g_{CA} by

23.3% (Table 3).

At this solution, ecosystem computations were turned on and the objective function computed one final time (Table 4). Including the ecosystem submodel defined by its hypothesis distribution parameter values significantly affects the model's overall fit as quantified by g_{CA} (Table 4).

Agreement measure	Symbol	Initial value	Consistent value	Percent increase
Agreement with observed actions history (match fraction)	$g_S^{(Grp)}$	0.491	0.528	0.075
Agreement with observed cheetah abundance	$g_S^{(Eco)}$	-0.141	0.0	1.000
Average agreement between consistent group submodels and hypothesis submodels	$g_H^{(Grp)}$	-0.486	-0.419	0.138
Overall agreement	g_{CA}	-0.382	-0.294	0.230

Table 3: CA agreement measures before and after the SA-MDAS run wherein ecosystem computations were kept off during the computation of the final values. Agreement between the ecological submodel's consistent distribution and its hypothesis distribution is not computed because no parameters within this submodel are being estimated.

Agreement measure	Symbol	Initial value	Consistent value	Percent increase
Agreement with observed actions history (match fraction)	$g_S^{(Grp)}$	0.528	0.509	-0.360
Agreement with observed cheetah abundance	$g_S^{(Eco)}$	-0.141	-0.137	0.028
Average agreement between consistent group submodels and hypothesis submodels	$g_H^{(Grp)}$	-0.486	-0.419	0.138
Overall agreement	g_{CA}	-0.382	-0.338	0.114

Table 4: CA agreement measures before and after the SA-MDAS run with ecosystem computations turned on when computing the final values.

Figure 1 displays the actions history generated by the CA-estimated model along with those observed actions that are matched by the CA-estimated model. To exhibit typical interaction patterns that are obscured in Figure 1, Figure 2 exhibits a closeup of the period from 2024 to 2025.

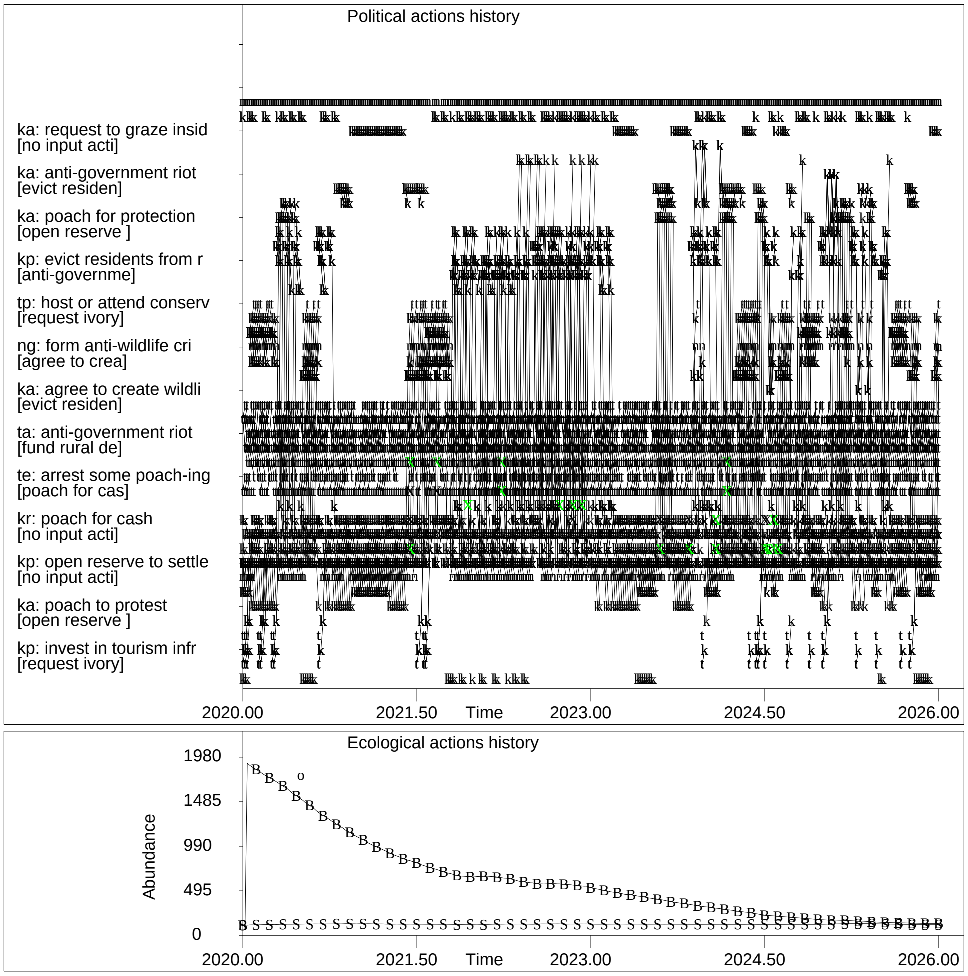

Figure 1: Model-generated actions history from 2020 through 2029. Green crosses are observed actions matched by the the CA-estimated model. The **o** symbol denotes an observation on cheetah abundance.

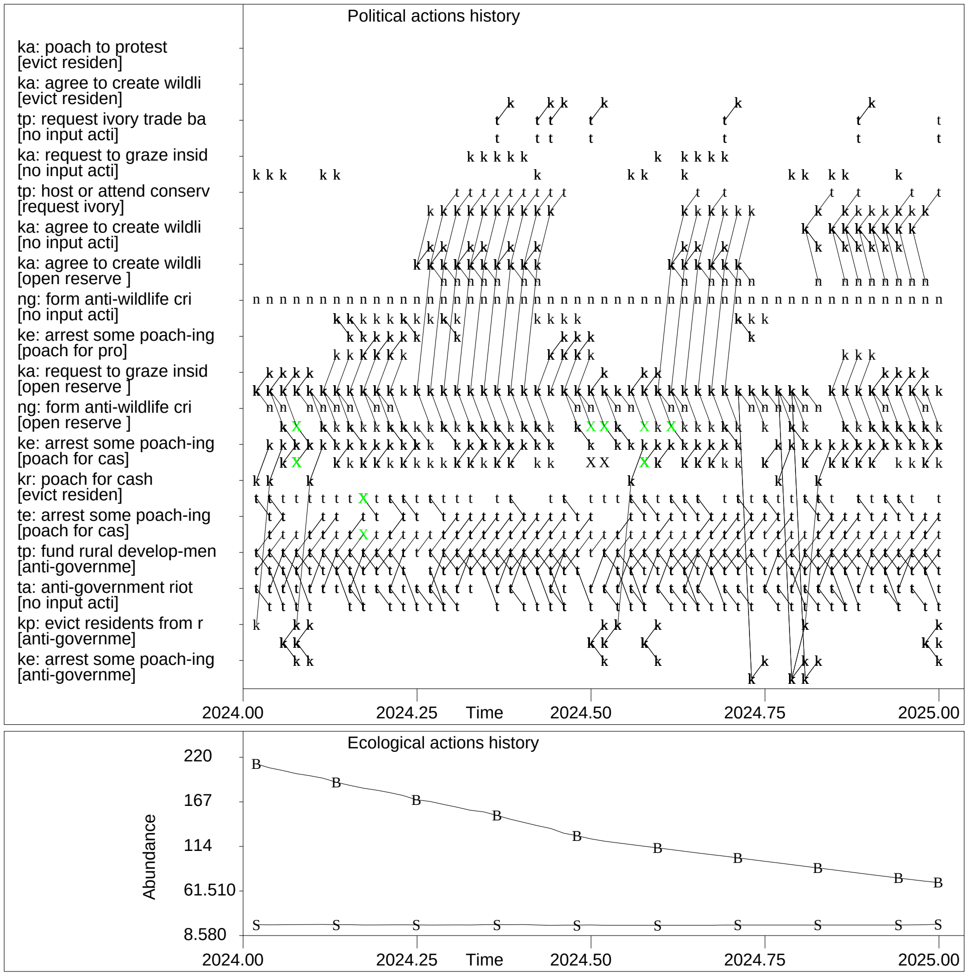

Figure 2: Model-generated actions history from 2024 to 2025.

This example shows that SA-MDAS can find CA estimates of a political-ecological model's parameters in a practical amount of wall clock time.

5. Finding the MPEMP

Haas and Ferreira (2017) state that

> "A more general method of developing management policies is the *most practical ecosystem management plan* (MPEMP) of Haas (2011, Chapter 4). The MPEMP emerges from the pattern of group behaviors that results from modifying one or more group belief systems. These modifications are such that the agreement

> between group belief systems that are estimated from data – and the belief systems that produce group actions that cause a desired ecosystem state, is maximized. In other words, the MPEMP is the sequence of group behaviors that occur from the least change in existing group beliefs systems that still achieves ecosystem state goals" (Haas and Ferreira, 2017).

Clearly, finding the MPEMP involves solving an optimization problem. This optimization problem is constrained by the requirement that any solution needs to produce values of the ecosystem's output variables that are close to the desired values at the desired point in time. This implicit constraint is incorporated into SA-MDAS as a penalty function in a manner similar to the CA constraint of maintaining a maximal fraction of model-to-observed action matches.

Only group submodel parameters can be variables in this constrained optimization problem. CA-estimated parameter values are used as starting values for all group submodels. During the optimization's search, however, all ecosystem submodel parameters are held at their CA-estimated values. Doing so represents the assumption made in the MPEMP computation that ecosystem dynamics are not under anthropogenic control but group belief systems are. Therefore, realistic ecosystem management plans should be restricted to making small modifications to human beliefs – and hence behavior rather than attempting to make changes to those ecological mechanisms that produce the modeled ecosystem's dynamics.

The preferred solution to the MPEMP optimization problem is a local one rather than a global one. This is because the MPEMP is the plan that requires the least change in existing beliefs (as represented by the model's consistent distribution) needed to redirect human behaviors enough to allow the ecosystem to reach a desired state. Hence, a solution that is close to the existing set of beliefs needs to be found.

The objective function in an MPEMP optimization problem needs to include the ecosystem submodel. This because the effect of group actions on the ecosystem needs to be detected every time the objective function is evaluated. Including the ecosystem submodel in the political-ecological model's simulation, however, can increase the objective function's evaluation time.

Further, when the ecosystem submodel is stochastic as in the East African cheetah example herein, the method of Bouttier and Gavra (2019) for solving a stochastic, black-box optimization problem needs to be employed. This method involves taking the average of repeated evaluations of the stochastic objective function at a trial solution point. This averaged value has a smaller variance than the original stochastic objective function.

The above discussion reveals that there are two computationaly expensive factors that complicate the evaluation of an MPEMP objective function. These are: The need to include the ecosystem submodel in the political-ecological model's simulation, and the need to compute an average over several function calls each time the optimization algorithm requires an objective function value. These two factors can cause the MPEMP objective function to have a long evaluation time. For instance, in the example below, the objective function's evaluation time is about 59 seconds. This long evaluation time in-turn, causes a long wall clock time before a solution to the MPEMP optimization problem is found. And long wall clock times can be expensive.

5.1. MPEMP algorithm

5.1.1. Needed definitions

1. The vector
$$\mathcal{B} = \left[\mathcal{B}^{(Grp)\prime}\ \mathcal{B}^{(Eco)\prime}\right]'$$
contains parameters of the group submodels, and the ecosystem submodel, respectively.
2. Let the vector, $\mathbf{Q}(\mathcal{B})$ contain the monitored ecosystem submodel variables whose values are generated by the ecosystem submodel using parameter values contained in $\mathcal{B}$.
3. Let $\mathbf{q}_d$ contain the desired ecosystem state in terms of $\mathbf{Q}(.)$ along with the time when these values are to be achieved.
4. Identify those actions that, if taken, would contribute the most towards the ecosystem submodel producing the values in $\mathbf{q}_d$. And, identify those actions that, if ceased, would raise the likelihood of the ecosystem submodel producing the values in $\mathbf{q}_d$. Collect all of these desirable and undesirable actions into a set called $\mathbf{c}_{MPEMP}$.

5.1.2. Algorithm

The following mathematical form of the MPEMP algorithm extends the one reported in Haas (2020).

1. Compute CA estimates of selected model parameters, i.e., update $\mathcal{B}_H$ to the most recent set of consistent parameter values: $\mathcal{B}_C$.
2. Compute $\mathbf{q}_H \equiv E\left[\mathbf{Q}(\mathcal{B}_H)\right]$.
3. Specify $\mathbf{q}_d$ and $\mathbf{c}_{MPEMP}$.
4. Compute initial values for $\mathcal{B}^{(Grp)}$ with CA's **Initialize** step.
5. Compute

$$\mathcal{B}_{\text{MPEMP}} = \arg\max_{\mathcal{B}^{(Grp)}} \left\{ g_H^{(Grp)}(\mathcal{B}) - \frac{||E[\mathbf{Q}(\mathcal{B})] - \mathbf{q}_d||}{||\mathbf{q}_H - \mathbf{q}_d||} \right\} \quad (3)$$

   under the set of constraints specified by $\mathbf{c}_{MPEMP}$.

Note that during the search in Step 5, $\mathcal{B}_H^{(Eco)}$ is unchanged.

5.1.3. Quantifying political feasibility

The MPEMP algorithm implements one way to quantify the concept of a politically feasible ecosystem management plan: Associate political feasibility with $g_H(\mathcal{B}_{\text{MPEMP}}^{(Grp)}) \in (-\infty,\ 0]$ where $\mathcal{B}_{\text{MPEMP}}^{(Grp)}$ contains the parameters of the decision making submodels whose values have been modified from those in $\mathcal{B}_H^{(Grp)}$ in such a way that now, the sequence of output actions taken by the different groups in the model causes a desired ecosystem state at a desired future time point ($\mathbf{q}_d$).

A measure of a plan's political feasibility can be defined as

$$\psi \equiv g_H^{(Grp)}(\mathcal{B}_{\text{MPEMP}}^{(Grp)})/(|g_H^{(Grp)}(\mathcal{B}_H)| + 0.000001). \quad (4)$$

The numerator is the agreement of the MPEMP distribution with the hypothesis distribution where large negative values indicate poor agreement. The first term in the denominator is the absolute value of the agreement of the model's distribution with the hypothesis distribution – using parameter values from the hypothesis distribution itself. When approximation error is zero, this term is zero. Hence, $g_H^{(Grp)}(\mathcal{B}_{\text{MPEMP}}^{(Grp)}) \leq |g_H^{(Grp)}(\mathcal{B}_H)|$.

A plan having a value of $\psi << -1.0$ will face stiff political resistance to its implementation because significant changes to the belief systems of one or more groups needs to happen – while a plan having a value close to -1.0 should not face such strong political headwinds.

5.2. Results

The desired ecosystem state is 800 cheetah across Kenya and Tanzania in the year 2029.

The parameters to be modified during the search for the MPEMP are those used in the above CA example with the addition of the parameters defining the node: Number of Cheetah Poached (NMPOACHED) under the proposed action of *poach cheetah for protection* for both Kenya rural residents and Tanzania rural residents.

These parameters are included in the MPEMP computation in order to study the feasibility of a two-pronged approach to cheetah conservation: Discouraging rural residents from poaching cheetah for cash while at the same time increasing antipoaching measures. To these ends, the SIIWP parameters are included to guide the amount of belief-change that would be needed to discourage rural residents from poaching cheetah for cash – and the NMPOACHED parameters to find the needed increase in antipoaching measures to stop rural residents from poaching cheetahs for protection. Optimizing these two sets of parameters simultaneously results in smaller changes to these parameters relative to the changes that would be required if only one of these sets was changed without changing the other. And smaller changes give the plan greater political feasibility.

In summary, the CA estimation modifies the SIIWP parameters concerning *poach cheetah for cash.* The MPEMP computation modifies the SIIWP parameters concerning *poach cheetah for cash* and the NMPOACHED parameters conditional on *poach cheetah for protection.*

5.2.1. MPEMP computation

To begin the optimization computation at a feasible solution, initial values of the SIIWP parameters were modified to produce the following output actions by both the Kenya rural resident group, and the Tanzania rural resident group: no *poach cheetah for cash* actions, no *poach cheetah for food* actions; but high probabilities for the actions, *poach cheetah for protection*, and *protest national park boundaries*.

Due to a modest computing budget, the sample size for averaging values of the stochastic objective function was fixed at three for the entire run. Not allowing this sample size to increase each pass through SA's outer loop means that the convergence proof of Bouttier and Gavra (2019) is only partially satisfied for this run. This proof states that if the algorithm of Bouttier and Gavra (2019) is followed, SA will converge to the global minimum of the stochastic objective function's expected value. Had this sample size been allowed to increase every pass through SA's outer loop, it would have equalled 39 when phase one finished.

Running on the TSCC, the SA-MDAS algorithm was restricted to 3000 objective function evaluations for each phase. The run employed three performers executing on 10 threads each. The run's wall clock time was 8.8 hours. The initial MPEMP objective function value was -0.779, and the final value was -0.711 for a 8.66% increase. Phase one converged after 1661 function evaluations and phase two terminated after 3037 function evaluations because it had exceeded its maximum number of function evaluations.

5.2.2. The computed MPEMP

This run produced an MPEMP that is projected to allow an expected cheetah population size of 278 by 2029. This is short of the desired goal of 800 cheetah by 2029 but avoids the forecast extinction event at the end of 2029 under the business-as-usual plan. The measure of the plan's political feasibility, ψ was computed to be -1.025.

Tables 5 and 6 show each parameter's definition, its CA value, and its MPEMP value. These parameter value changes indicate that the MPEMP is to (a) change the belief in being arrested for poaching cheetah for cash from being perceived as negligible to being perceived as likely; and (b) increase antipoaching measures to the point where Kenya and Tanzania rural residents each succeed in poaching less than six cheetah per week.

The number of cheetah poached per week by Kenya rural residents and the number of cheetah poached per week by Tanzania rural residents were not adjusted by the CA. Hence, these two parameters were held at their hypothesis values. These hypothesis values are slightly below the upper bound constraint shared by these parameters. This upper bound value is six. This value indicates a significant amount of poaching that aligns with the hypothesis-belief held by these rural residents that they will not be arrested for poaching cheetah.

Parameter	Node	CA value	MPEMP value	Change fraction	> 10% change?
1	Number_Poached	5.994	3.574	00.403	*
2	SIIWP	00.274	00.193	00.294	*
3	SIIWP	00.356	00.232	00.347	*
4	SIIWP	00.369	00.573	00.555	*
5	SIIWP	00.054	00.320	4.926	*
6	SIIWP	00.346	00.357	00.029	
7	SIIWP	00.599	00.322	00.461	*
8	SIIWP	00.166	00.553	2.337	*
9	SIIWP	00.402	00.119	00.702	*
10	SIIWP	00.431	00.326	00.244	*
11	SIIWP	00.410	00.036	00.911	*
12	SIIWP	00.133	00.739	4.558	*
13	SIIWP	00.455	00.224	00.507	*
14	SIIWP	00.346	00.329	00.047	
15	SIIWP	00.056	00.329	4.878	*
16	SIIWP	00.598	00.341	00.429	*
17	SIIWP	00.175	00.461	1.634	*
18	SIIWP	00.292	00.373	00.280	*
19	SIIWP	00.533	00.165	00.690	*
20	SIIWP	00.242	00.535	1.214	*
21	SIIWP	00.490	00.337	00.310	*
22	SIIWP	00.268	00.126	00.528	*
23	SIIWP	00.118	00.403	2.387	*
24	SIIWP	00.253	00.202	00.199	*
25	SIIWP	00.628	00.394	00.372	*
26	SIIWP	00.306	00.197	00.355	*
27	SIIWP	00.400	00.660	00.647	*
28	SIIWP	00.292	00.142	00.515	*

Table 5: Kenya rural resident parameter value changes needed to produce the MPEMP. The first parameter is the average number of cheetah poached per week by Kenya rural residents. The second set of parameters define the perceived outcome by Kenya rural residents under the decision option to *poach cheetah for cash*. Perceived outcomes are *will be arrested for poaching*, *will be evicted*, and *no interaction with police*.

Parameter	Node	CA value	MPEMP value	Change fraction	> 10% change?
29	Number_Poached	5.994	5.372	00.103	*
30	SIIWP	00.039	00.328	7.423	*
31	SIIWP	00.521	00.183	00.647	*
32	SIIWP	00.439	00.487	00.108	*
33	SIIWP	00.330	00.552	00.675	*
34	SIIWP	00.296	00.406	00.368	*
35	SIIWP	00.373	00.040	00.890	*
36	SIIWP	00.250	00.370	00.474	*
37	SIIWP	00.024	00.403	15.802	*
38	SIIWP	00.725	00.226	00.687	*
39	SIIWP	00.289	00.318	00.103	*
40	SIIWP	00.099	00.310	2.137	*
41	SIIWP	00.612	00.370	00.394	*
42	SIIWP	00.272	00.290	00.068	
43	SIIWP	00.400	00.352	00.118	*
44	SIIWP	00.328	00.356	00.087	
45	SIIWP	00.450	00.082	00.817	*
46	SIIWP	00.167	00.446	1.674	*
47	SIIWP	00.382	00.471	00.233	*
48	SIIWP	00.209	00.540	1.586	*
49	SIIWP	00.471	00.240	00.491	*
50	SIIWP	00.319	00.219	00.313	*
51	SIIWP	00.153	00.335	1.190	*
52	SIIWP	00.290	00.290	00.000	
53	SIIWP	00.557	00.374	00.327	*
54	SIIWP	00.132	00.320	1.426	*
55	SIIWP	00.048	00.252	4.259	*
56	SIIWP	00.820	00.427	00.478	*

Table 6: Tanzania rural resident parameter value changes needed to produce the MPEMP. The first parameter is the average number of cheetah poached per week by Tanzania rural residents. The second set of parameters define the perceived outcome by Tanzania rural residents under the decision option to *poach cheetah for cash*. Perceived outcomes are *will be arrested for poaching*, *will be evicted*, and *no interaction with police*.

5.2.3. The MPEMP's political consequences

Because of ψ's modestly negative value, this MPEMP is expected to face moderate political resistance. A practical policy for implementing this MPEMP is to simultaneously (a) visibly increase antipoaching enforcement in both countries, and (b) communicate such increases to the rural residents of Kenya and Tanzania (Kegamba et al., 2024), (Nachihangu et al., 2022).

This fitted model and the political insights derived from it, are new.

5.3. The MPEMP compared to other conservation planning tools

5.3.1. Marxan

Marxan helps users select a set of conservation reserves (hereafter, a set of *reserve patches*) that satisfy ecological, social, and financial goals. The Marine Spatial Planning group at the Nature Conservancy gives this overview of Marxan:

> "Marxan is decision support software for designing new reserve systems, reporting on performance of existing reserve systems, and developing multiple-use zoning plans. It is the most widely adopted site selection tool by conservation groups globally. Marxan is a stand-alone software program that provides decision support to teams of conservation planners and local experts identifying efficient areas that combine to satisfy a number of ecological, social and economic objectives. Given data on species, habitats, ecosystems and other biodiversity features, Marxan was designed to minimize the cost of selected sites while meeting all goals (Marine Spatial Planning (2025))."

Marxan finds these plans by solving the *set covering problem* (Liang et al., 2020), (Watts et al., 2021). Marxan is strictly spatial in that the tool does not recognize any temporal component in the spatial data it is given. Also, the user provides all parameter values and Marxan makes no effort to incorporate parameter uncertainties into its final solution. The only conservation plans that Marxan can deliver are maps identifying the location and geometry of each reserve patch.

Once computed, these reserve patches will presumably be forced onto landscape residents by authoritarian actions such as the purchase of land or the successful lobbying for the conversion of privately-held real estate into reserve patches.

The MPEMP by comparison, emerges from the political realities surrounding a biodiversity conservation challenge and allows plans to include non-command-and-control solutions such as community-led changes in landuse on private, non-reserve lands.

Marxan assumes that different sets of reserve patches can be mandated without requiring community acceptance of the plan. In other words, Marxan does not have readily-available inputs to represent data-derived political forces that determine the political feasibility of mandating a set of reserve patches. On the other hand, the central capability of the MPEMP constructor is its ability to model these forces and their effects on the management of an ecosystem.

The set covering problem is a classic, simple, linear, binary integer programming problem (Liang et al., 2020). Continuously-valued metrics are not allowed, nonlinear objective functions are not allowed, and integer variables with more than two values are not allowed. Marxan runs on a single-CPU computer. Hence, its ability to solve its set covering problem is, by design, computationally limited. Unfortunately, the set covering problem is known to be NP-hard (Liang et al., 2020) meaning that finding a solution to a complex planning problem involving many possible sets of reserve patches can take many months of wall clock time on a single-CPU computer to reach the optimal solution.

But real-world reserve-set planning problems usually involve many possible solutions. This is similar to the set of possible plans for effectively protecting an ecosystem within a larger political-ecological system – the problem that the present article solves. In contrast to Marxan, the present article has tackled this crippling computational-expense roadblock head-on with its cluster computing-based algorithm, SA-MDAS. SA-MDAS is designed to find the MPEMP in a practical amount of wall clock time as opposed to the several months of wall clock time that Marxan might spend in order to find an optimal set of reserve patches.

5.3.2. Zonation 5

The Finnish Environment Institute describes Zonation as follows:

> "The Zonation software enables prioritisation analyses of conservation values based on spatial data to support decision-making, conservation area planning, and the avoidance of negative ecological impacts. It can incorporate a wide range of data, including information on uncertainties and connectivity. The analyses can be highly detailed and extensive, depending on the available datasets (Finnish Environment Institute, 2025)."

In addition, Moilanen et al. (2022) describe the newest version of this spatial conservation planning tool: *Zonation Version 5*.

Similar to Marxan, this tool helps the user to identify a set of reserve patches within a given landscape. Differing from Marxan, however, this tool does not explicitly solve an optimization problem in order to find a set of patches that optimally meet a collection of target values on selected ecological metrics. Rather, Zonation 5 performs a nested, two-level sort on a large number of spatial units that collectively make up a landscape. This sorting operation produces a ranked list of reserve patch geometries. Many assumptions are made about what is ecologically preferred in order to place enough characteristics on each spatial unit to allow the sorting algorithm to compute rank-differences between spatial units. As with Marxan, Zonation 5 is strictly spatial in that the tool does not recognize any temporal component in the spatial data it is given.

Because Zonation 5 employs an efficient sorting algorithm, it can quickly rank-order a large number of spatial units on a single-CPU computer. Zonation 5 does allow some political forces to impact its solutions by allowing the user to input a set of "local preferences." As with Marxan, however, the user provides all parameter values with Zonation 5 making no effort to incorporate parameter uncertainties into its final rank-ordering of possible reserve patch geometries.

As an aside, in actuality, sorting can be cast as an optimization problem solved with a greedy search algorithm (Bauckhage and Welke, 2022). It is the greedy nature of the search algorithm that allows Zonation to perform efficient sorts on a single-CPU computer.

5.3.3. Assessment

These two tools only deliver a set of reserve patches across a given landscape. Neither tool allows, as the MPEMP computation does, a user to model the political-ecological system that the landscape is a part of – let alone provide a capability for statistically fitting such a model's parameters to a political-ecological data set.

A major drawback that makes these tools insufficient in light of current political realities is their inability to allow a user to explore alternate conservation plans that might involve initiatives that are not a command-and-control gazzetting of land into reserve patches. Initiatives such as developing alternate livelihoods, installing wildlife control fencing, increasing antipoaching efforts, relocating animals, shooting marauding elephants, eradicating invasive plants, donating antipoaching equipment, planting trees, restoring mine tailings areas, compensating farmers for wildlife damages, and the downgrading, downsizing, and degazettement of protected areas (PADDD) Albrecht et al. (2021) – are not part of either tool's plan repertoire. See the file `emat.dfn` in the Supplementary Materials file, `MPEMP_constructor.zip` for a complete list of such alternate conservation actions as gleaned from political-ecological data acquired using the STAR protocol of Haas (2024b).

The critical deficiency of both of these tools, however, is that because all priority weights are user-specified, any plan delivered by either tool, unlike the MPEMP, lacks a data-derived, real-world measure of the plan's political feasibility.

6. Challenges

Before political-ecological modelling of global extinction events can guide conservation efforts, several challenges need to be overcome. Below is a priority-ordered list of some of these challenges.

1. Collaboration with the private sector is needed in order to tap their distributed autonomy; enormous reserves of talent; and enormous reserves of cash and credit (Haas, 2022), (Haas, 2024c), (Haas, 2025).
2. Better political-ecological data sets need to be acquired that, via cluster computer computations, can be used to statistically fit the param-

eters of these extinction event models. On the political side, observations need to be acquired on those private political agreements that often determine the fate of conservation legislation. On the ecological side, unrestricted and noninvasive monitoring of species abundance is needed. Currently, such data is often restricted because (a) owners refuse access to their properties, and/or (b) data managers fear that poachers will hack into the data's repository in order to locate wildlife (Lennox et al., 2020).

3. In order to make the statistical fitting of these models feasible to a wide spectrum of researchers, cluster computing resources need to be financially accessible to cash-strapped departments of ecology and departments of political science.

   One economical alternative to purchasing time on a commercial cluster computer is to organize some number of in-house computers into a cluster computer. With the JavaSpaces-based SA-MDAS algorithm, this is a straightforward four-step build:

   (a) Install on an impresario, a *dynamic domain name system* (DDNS) (mintdns, 2023) in order to give this computer a fixed domain name so that performers can locate its GigaSpace. A DDNS is acquired from an internet provider or from a DDNS provider. Fees can be either none; a fixed monthly charge; or a monthly, usage-based charge.
   (b) Verify that all performers can access the internet.
   (c) Install the JAVA Development Kit (free) and the GigaSpaces XAP (annual lease) on the impresario and all performers.
   (d) Compile and run on these computers, the JAVA code exercised in this article.

4. The credibility of political-ecological models in the eyes of their skeptics needs to be established by meeting Popperian credibility criteria that are based on frequentist statistical methods (Haas, 2024) – rather than relying on the lobbying efforts of *activist scientists* (Whipple, 2024). Such rigorous, statistically-based model validation results are needed to maximize the acceptance by the general public of extinction-event predictions. This would be in contrast to the acceptance of the declarations made by the authors of climate models concerning the causes

of climate change. For instance, fewer than 50% of Americans believe that climate change is driven by anthropogenic forces (EPIC, 2023) – the central mechanism that is programmed into almost all mainstream climate models. In order to marshal the necessary support for expensive biodiversity conservation projects, a much higher proportion of popular support will be needed than what is given to climate change mitigation policies.

Indeed, climate models do not currently meet rigorous statistical measures of model credibility. For instance, Michaels (2019) notes that recent research indicates that the averaging of predictions across many different climate models as is routinely reported by the Intergovernmental Panel on Climate Change (IPCC), increases prediction errors. This author suggests that work should focus on developing a single model that has low prediction errors. And Jain et al. (2023) find significant errors in model predictions of observed weather and suggest more computing resources may be needed to reduce this error.

A literature search failed to find a peer-reviewed article that reports on a climate model that produces predictions of out-of-sample observations to any specific, quantified level of accuracy. Nonetheless, such models are routinely used to justify climate-change mitigation policy decisions. This use of incompletely validated climate models has only fueled skepticism about the credibility of climate models. These skeptics instead, suspect that climate change policymaking is being driven by flawed models (Montford, 2022). Indeed, in a particularly worrisome passage, Montford mentions the current practice of "tuning" model parameters to increase a climate model's agreement with observations. Such tuning is outside any frequentist-based statistical method of parameter estimation.

5. Group decision making mechanisms need to be better understood. Such newly-derived mechanisms need to be programmed into each group decision making submodel that runs within a political-ecological model. Also, computational models need to be developed that, through improved architectures, realistically integrate political and ecological processes. One such architecture employed herein is that of agent-based

submodels communicating with each other and with an individual-based ecosystem submodel through a bulletin board.

6. Increased financial and career support is needed for interdisciplinary research on the integration of political and ecological processes. Within academia, this support will be aided when (a) journal editors encourage the publication of high-quality manuscripts that probe this interface (Huber-Sannwald, 2025), (O'Connor et al., 2021), (Van Bael, 2025); and (b) departmental committees encourage tenure-track faculty to publish at this interface (Washbourne et al., 2024). Outside academia, funding agencies need to seek out and fund high-quality proposals to build and test models of those political-ecological systems that contain at-risk species.

7. Conclusions

For purposes of biodiversity conservation, attention needs to focus on finding politically feasible projects aimed at conserving biodiversity. Earth is running out of time to save what remains of its biodiversity. MPEMPs need to be computed and implemented as soon as possible. This article shows that, thanks to the new optimization algorithm introduced herein, these computations are now feasible. Indeed, as explained in the Appendix, an optimization algorithm such as SA-MDAS that scales on the increasing availability of cluster computer nodes, is currently, the only feasible way to statistically fit a large political-ecological model to data.

The political-ecological models that will support these projects, can have their parameters statistically estimated via CA to political-ecological data sets acquired via the STAR protocol of Haas (2024b). Use of CA to perform this parameter estimation, allows cognitive theories of decision making to be represented in the final fitted models.

As per this article's title, the MPEMP is how a political-ecological model can help sustain biodiversity. The freely available MPEMP constructor can be immediately used by an organization to develop and implement politically feasible and ecologically effective plans for conserving those at-risk species that they have selected to help preserve.

Given the current negative trends in biodiversity, one might conclude that conservation research is futile and hence, should be abandoned. A recent meta-analysis, however, using data from many different conservation projects, finds that actually, conservation projects have been effective at stopping or at least slowing the decline of biodiversity across the globe (Oxford, 2024). Given this evidence, there is reason to believe that research products such as the MPEMP constructor offered in this article, can help build a conservation science that is capable of guiding successful efforts to curb the loss of species regardless of their commercial value.

Supplementary material

All source code and input files that constitute the MPEMP constructor described and exercised in this article are contained in the Supplementary Material files `idsrce.zip` (JAVA source code), and `MPEMP_constructor.zip` (linux shell scripts, input command files, and data files).

References

Aarts, E. and Korst, J. (1989), *Simulated Annealing and Boltzmann Machines,* New York: John Wiley and Sons.

Albrecht, R., Cook, C. N., Andrews, O., Roberts, K. E., Taylor, M. F. J., Mascia, M. B., and Kroner, R. E. G. (2021), "Protected Area Downgrading, Downsizing, and Degazettement (PADDD) in Marine Protected Areas," *Marine Policy,* 129: 104437. doi: 10.1016/j.marpol.2021.104437.

Bauckhage, C. and Welke, P. (2022), *ML2R Coding Nuggets: Sorting as a QUBO*, *Machine Learning Rhine-Ruhr*, St. Augustin, Germany. `https://pwelke.de/lecturenotes/ml2r-cn-sortingasqubo.pdf`

Beverdam, J., Hubacek, K., Scholtens, B., and Sijtsma, F. (2025), "Improving Biodiversity Resilience Requires Both Public and Private Finance: A life-Cycle Analysis of Biodiversity Finance," *Ecological Economics,* 234: 108607. doi: 10.1016/j.ecolecon.2025.108607.

Bouttier, C. and Gavra, I. (2019), "Convergence Rate of a Simulated Annealing Algorithm with Noisy Observations," *Journal of Machine Learning Research,* 20(4): 1-45. `http://jmlr.org/papers/v20/16-588.html`

Branke, J., Meisel, S. and Schmidt, C. (2008), "Simulated Annealing in the Presence of Noise," *Journal of Heuristics,* 14: 627-654. 10.1007/s10732-007-9058-7.

Carroll G., Holsman, K. K., Brodie, S., Thorson, J. T., Hazen, E. L., Bograd, S. J., Haltuch, M. A., Kotwicki , S. Samhouri, J. Spencer, P., Willis-Norton, E., and Selden, R. L. (2019), "A Review of Methods for Quantifying Spatial Predator–Prey Overlap," *Global Ecology and Biogeography,* 28: 1561–1577. doi: 10.1111/geb.12984.

Ceballos, G., Ehrlich, P. R., Barnosky, A. D., García A., Pringle, R. M., Palmer, T. M. (2015), "Accelerated Modern Human-Induced Species Losses: Entering the Sixth Mass Extinction," *Science Advances*, 1(5): e1400253. `https://doi.org/10.1126/sciadv.1400253`

Ciatto, G., Mariani, S., Serugendo, G. D., Louvel, M., Omicini, A., and Zambonelli, F. (2020), "Twenty Years of Coordination Technologies: COORDINATION Contribution to the State of Art," *Journal of Logical and Algebraic Methods in Programming,* 113(100531), ISSN 2352-2208, doi: 10.1016/j.jlamp.2020.100531.

Corana, A., Marchesi, M., Martini, C., and Ridella, S. (1987), "Minimizing Multimodal Functions of Continuous Variables with the "Simulated Annealing" Algorithm," *ACM Transactions on Mathematical Software,* 13(3): September, 262-280. `https://doi.org/10.1145/29380.29864`

EPIC (2023), *Who, Us? Fewer Than 50% of Americans Think Humans are the Main Cause of Climate Change,* Energy Policy Institute, The University of Chicago, April 18. `https://epic.uchicago.edu/news/who-us-fewer-than-50-of-americans-think-humans-are-the-main-cause-of-climate-change/`

Ferreiro, A. M., García-Rodríguez, J. A., Vázquez, C., Costa e Silva, E., and Correia, A. (2019), “Parallel Two-Phase Methods for Global Optimization on GPU,” *Mathematics and Computers in Simulation,* 156: 67-90. doi: 10.1016/j.matcom.2018.06.005.

Finnish Environment Institute (2025), *Zonation – A Practical Tool for Nature Conservation and Land Use Planning.* `https://www.syke.fi/en/environmental-data/calculators-and-tools/zonation`

Fitzgibbon, C. D. (1990), “Why Do Hunting Cheetahs Prefer Male Gazelles?,” *Animal Behaviour,* 40(5): 837-845. doi: 10.1016/S0003-3472(05)80984-4.

Garber, P. A., Estrada, A., Shanee, S., Svensson, M. S., Arregoitia, L. V., Nijman, V., Shanee, N., Gouveia, S. F., Nekaris, K. A. I., Chaudhary, A., Bicca-Marques, J. C., and Hansen, M. F. (2024), “Global Wildlife Trade and Trafficking Contribute to the World’s Nonhuman Primate Conservation Crisis,” *Frontiers in Conservation Science,* 5: 1400613, doi: 10.3389/fcosc.2024.1400613.

Gonçalves-e-Silva, K., Aloise, D., and Xavier-de-Souza, S. (2018), “Parallel Synchronous and Asynchronous Coupled Simulated Annealing,” *Journal of Supercomputing,* 74: 2841-2869. 10.1007/s11227-018-2327-4.

Guilmeau, T., Chouzenoux, E., and Elvira, V. (2021), “Simulated Annealing: A Review and A New Scheme,” *IEEE Statistical Signal Processing Workshop (SSP),* Rio de Janeiro, Brazil, pp. 101-105, doi: 10.1109/SSP49050.2021.9513782.

Haas, T. C. (2025), “A Technology-Based Business Plan for Profitably Curbing Wildlife Trafficking,” *Sustainability Technology* (SusTech 2025), Santa Ana, California, April 20-23. See page 50 of: `https://ieee-sustech.org/wp-content/uploads/sites/261/2025/04/SusTech-2025-Program-Guide.pdf`

Haas, T. C. (2025b) *Profitable Biodiversity Website*, `https://profitablebiodiversity.com`.

Haas, T. C. (2024), "Models Vetted Against Prediction Error and Parameter Sensitivity Standards Can Credibly Evaluate Ecosystem Management Options," *Ecological Modelling*, 498, December, 11090 ("decreases" should be "increases" in the Graphical Abstract). https://doi.org/10.1016/j.ecolmodel.2024.110900

Haas, T. C. (2024), "Protocol to Discover Machine-Readable Entities of the Ecosystem Management Actions Taxonomy," *STAR Protocols*, Cell Press, Elsevier, 5(2), 103125: 1-12. https://doi.org/10.1016/j.xpro.2024.103125

Haas, T. C. (2024c), "A New Technology-Based Tool for Building Profitable Biodiversity-Conserving Offerings," *The European Journal of Sustainable Development,* 13(3): 1-13. https://doi.org/10.14207/ejsd.2024.v13n3p57

Haas, T. C. (2022), "Profitable Biodiversity," *Cogent Social Sciences*, 8(1): 1-24. https://doi.org/10.1080/23311886.2022.2116814 https://www.tandfonline.com/doi/full/10.1080/23311886.2022.2116814

Haas, T. (2020), "Developing Political-Ecological Theory: The Need for Many-Task Computing," *PLOS ONE*, November 24. 10.1371/journal.pone.0226861. https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0226861

Haas, T. C. (2011), *Improving Natural Resource Management: Ecological and Political Models*, a "Statistics in Practice" volume, cross-listed in the Environmental Management, Policy and Planning series, and the Environmental Economics and Politics series, Wiley-Blackwell, Oxford, U.K. ISBN: 978-0-470-66113-0.

Haas, T. C. and Ferreira, S. M. (2017), "Finding Politically Feasible Conservation Strategies: The Case of Wildlife Trafficking," *Ecological Applications*, 28(2): 473-494. https://doi.org/10.1002/eap.1662

Hanski, I. (2011), "Habitat Loss, the Dynamics of Biodiversity, and a Perspective on Conservation," *Ambio*, 40(3): 248-255. https://doi.org/10

.1007/s13280-011-0147-3 https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3357798/

Hasanzadeh, M. R., Keynia, F., and Hashemipour, M. (2022), "A New Index-Based Hyper-Heuristic Algorithm for Global Optimisation Problems," *IET Software*, 16(5): 493–515. doi: 10.1049/sfw2.12065.

Holmes, W. C. (1994), *Opera Observed: Views of a Florentine Impresario in the Early Eighteenth Century*, University of Chicago Press, ISBN: 9780226349701.

Huber-Sannwald, E. (2025), *Earth Stewardship*, Editor-in-Chief: Elisabeth Huber-Sannwald, Online, ISSN:2835-3617. https://esajournals.onlinelibrary.wiley.com/journal/28353617

Jain, S., Scaife, A. A., Shepherd, T. G., Deser, C., Dunstone, N., Schmidt, G. A., Trenberth, K. E. and Turkington, T. (2023), "Importance of Internal Variability for Climate Model Assessment," *npj Climate and Atmospheric Science,* 6: 68. https://www.nature.com/articles/s41612-023-00389-0

Jaureguiberry, P., Titeux, N., Wiemers, M., Bowler, D. E., Coscieme, L., Golden, A. S., Guerra, C. A., Jacob, U., Takahashi, Y., Settele, J., Daz, S., Molnr, Z., Purvis, A. (2022), "The Direct Drivers of Recent Global Anthropogenic Biodiversity Loss," *Science Advances,* 8(45): eabm9982. https://doi.org/10.1126/sciadv.abm9982 https://www.science.org/doi/pdf/10.1126/sciadv.abm9982

Jia, K., Duan, X., Wang, Z., Yi, T., Yan, L., and Chen, X. (2024), "A New Partition Method for DIRECT-type Algorithm Based on Minimax Design," *Journal of Global Optimization,* 88, 171–197. DOI: 10.1007/s10898-023-01297-6.

Jiang, P. and Singh, M. (2010), "SPICi: A Fast Clustering Algorithm for Large Biological Networks," *Bioinformatics,* April 15, 26(8): 1105-11. doi: 10.1093/bioinformatics/btq078.

Kegamba, J. J., Sangha, K. K., Wurm, P. A., Kideghesho, J. R., and Garnett, S. T. (2024), "The Influence of Conservation Policies and Legislations on Communities in Tanzania," *Biodiversity and Conservation,* 33: 3147–3170. 10.1007/s10531-024-02906-w.

Kelly, M. J., Laurenson, M. K., Fitzgibbon, C. D., Collins, D. A., Durant, S. M., Frame, G. W., Bertram, B. C., and Caro, T. M. (1998), "Demography of the Serengeti Cheetah (Acinonyx jubatus) Population: The First 25 years," *Journal of Zoology,* 244: 473-488. https://doi.org/10.1111/j.1469-7998.1998.tb00053.x

Kimbrell, T. and Holt, R. D. (2005), "Individual Behaviour, Space and Predator Evolution Promote Persistence in a Two-Patch System with Predator Switching," *Evolutionary Ecology Research,* 7: 53-71.

Lee, S-Y. and Lee, K. G. (1996), "Synchronous and Asynchronous Parallel Simulated Annealing with Multiple Markov Chains," *IIE Transactions on Parallel and Distributed Systems*, 7.

Lennox, R. J., Harcourt, R., Bennett, J. R., Davies, A., Ford, A. T., Frey, R. M., Hayward, M. W., Hussey, N. E., Iverson, S. J., Kays, R., Kessel, S. T., Mcmahon, C., Muelbert, M., Murray, T. S., Nguyen, V. M., Pye, J. D., DRoche, D. G., Whoriskey, F. G., Young, N., Cooke, S. J., (2020), "A Novel Framework to Protect Animal Data in a World of Ecosurveillance," *BioScience,* 70(6), June: 468–476. DOI: 10.1093/biosci/biaa035.

Li, W., Peng, Y., and Zhang, X. (2025), "Remote Sensing Enables Accurate Assessment of Functional Diversity Rather Than Species Diversity in Sandy Grasslands," *Biodiversity Informatics,* 19: 86-108. https://journals.ku.edu/jbi/article/view/24120/21751

Liang, S., Alanazi, K., and Hamoud, K. A. (2020), *Set Covering Problem, Cornell University Computational Optimization Open Textbook.* https://optimization.cbe.cornell.edu/index.php?title=Set_covering_problem

Marine Spatial Planning (2025), *MARXAN*, *The Nature Conservancy,* https://marineplanning.org/tools/software/marxan/

Medlong, J., Keyes, G., and McCaffrey, A. (2022), *COP15 | Biodiversity: US Legal and Business Perspectives*, *DLA Piper.* https://www.dlapiper.com/en/insights/publications/2022/12/cop15-biodiversity-us-legal-and-business-perspectives

Michaels, P. J. (2019), *Are Climate Models Overpredicting Global Warming? Cato Institute*, Commentary, January 31. https://www.cato.org/commentary/are-climate-models-overpredicting-global-warming

Milloway, O. (2025), *TWS 2024: Lead is 'Silently Poisoning' Captive Cheetahs, The Wildlife Society,* June 30. https://wildlife.org/tws-2024-lead-is-silently-poisoning-captive-cheetahs/

mintdns (2023), *The Domain Name System,* mintdns. https://dynamic.domains/dynamic-dns/providers-list/default.aspx

Mishra, S. (2006), "Some New Test Functions for Global Optimization and Performance of Repulsive Particle Swarm Method," *MPRA Paper*, N. 2718. https://mpra.ub.uni-muenchen.de/2718/1/MPRA_paper_2718.pdf

Moilanen, A., Lehtinen, P., Kohonen, I., Jalkanen, J., Virtanen, E. A., and Kujala, H. (2022), "Novel Methods for Spatial Prioritization with Applications in Conservation, Land Use Planning and Ecological Impact Avoidance," *Methods in Ecology and Evolution,* 13: 1062–1072. doi: 10.1111/2041-210X.13819.

Montford, A. (2022), "Ten Things Everyone Should Know About Climate Models," *The Global Warming Policy Foundation,* Note 32. https://www.thegwpf.org/content/uploads/2022/06/Montford-Climate-Models-1.pdf

Nachihangu, J., Massao, C. A., Nahonyo, C., Richard, U. (2022), "The Impact of Harvest-Based Initiatives on Wildlife Poaching around the Ugalla and Rungwa Game Reserves, *Global Ecology and Conservation,* 40, e02320, ISSN 2351-9894, 10.1016/j.gecco.2022.e02320.

Netz, C., Hildenbrandt, H., and Weissing, F. J.(2022), "Complex Eco-evolutionary Dynamics Induced by the Coevolution of Predator–Prey Movement Strategies," *Evolutionary Ecology,* 36: 1–17, `https://doi.org/10.1007/s10682-021-10140-x`

O'Connor, M. I., Mori, A. S., Gonzalez, A., Dee, L. E., Loreau, M., Avolio, M., Byrnes, J. E. K., Cheung, W., Cowles, J., Clark, A. T., Hautier, Y., Hector, A., Komatsu, K., Newbold, T., Outhwaite, C. L., Reich, P. B., Seabloom, E., Williams, L., Wright, A., and Isbell, F. (2021), "Grand Challenges in Biodiversity–Ecosystem Functioning Research in the Era of Science–Policy Platforms Require Explicit Consideration of Feedbacks," *Proceedings of the Royal Society B*, 28820210783, 10.1098/rspb.2021.0783.

University of Oxford News and Events (2024), *Landmark Study Definitively Shows That Conservation Actions are Effective at Halting and Reversing Biodiversity Loss, University of Oxford,* `https://www.ox.ac.uk/news/2024-04-26-landmark-study-definitively-shows-conservation-actions-are-effective-halting-and`.

Panwar, R., Ober, H., and Pinkse, J. (2023), "The Uncomfortable Relationship between Business and Biodiversity: Advancing Research on Business Strategies for Biodiversity Protection," *Business Strategy and the Environment,* 32(5): 2554-2566. doi: 10.1002/bse.3139.

Pichetti, L., De Sensi, D., Sivalingam, K., Nassyr, S., Cesarini, D., Turisini, M., Pleiter, D., Artigiani, A., and Vella, F. (2025), "Benchmarking Ethernet Interconnect for HPC/AI Workloads," (in) *Proceedings of the SC '24 Workshops of the International Conference on High Performance Computing, Network, Storage, and Analysis (SC-W '24),* IEEE Press, 869–875. `https://doi.org/10.1109/SCW63240.2024.00124`

Ressurreição, A., Zarzycki, T., Kaiser, M., Edwards-Jones, G., Dentinho, T. P., Santos, R. S., and Gibbons, J. (2012), "Towards an Ecosystem Approach for Understanding Public Values Concerning Marine Biodiversity Loss," *Marine Ecology Progress Series,* 467 (October): 15-28.

San Diego Supercomputer Center (2025), *Triton Shared Computing Cluster,* University of California, San Diego. Service. doi: 10.57873/T34W2R.

Schober, K. and Prestin, J. (2023), "Analysis of Directional Higher Order Jump Discontinuities with Trigonometric Shearlets," *Mathematical Foundations of Computing,* 6(1): 14-40. doi: 10.3934/mfc.2021038.

Schumer, M. A. and Steiglitz, K. (1968), "Adaptive Step Size Random Search," *IEEE Transactions on Automatic Control,* AC-13(3): 270-276.

Suresh, A. T. (2021), "Robust Hypothesis Testing and Distribution Estimation in Hellinger Distance," *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS),* San Diego, California, USA. Proceedings of Machine Learning Research (PMLR): Volume 130.

Testa, F., Tosi, D., Tessitore, S., Todaro N. M., and Di Iorio, V. (2025), "Untangling Companies' Engagement With Biodiversity: A Systematic Literature Review and Research Agenda," *Business Strategy and the Environment,* 34(6): 6773-6790. doi: 10.1002/bse.4324.

Tobias, J. A., Bullock, J. M., Dicks, L. V., Forester, B. R., and Razgour, O. (2025), "Biodiversity Conservation Requires Integration of Species-Centric and Process-Based Strategies, *Proceedings of the National Academy of Sciences*, 122(31): e2410936122. doi: 10.1073/pnas.2410936122.

Torres-Romero, E. J., Fisher, J. T., Nijman, V., He, F., Eppley, T. M. (2024), "Accelerated Human-Induced Extinction Crisis in the World's Freshwater Mammals," *Global Environmental Change Advances,* 2: 100006, ISSN 2950-1385. `https://doi.org/10.1016/j.gecadv.2024.100006`

Van Bael, S. (2025), *Ecosphere: Socio-Ecological Systems,* Editor-in-Chief: Sunshine Van Bael, Online ISSN:2150-8925, Print ISSN:2150-8925. `https://esajournals.onlinelibrary.wiley.com/hub/journal/21508925/socio-ecological-systems`

Verones, F., Kuipers, K., Núñez, M., Rosa, F., Scherer, L., Marques, A., Michelsen, O., Barbarossa, V., Jaffe, B., Pfister, S., and Dorber, M. (2022), "Global Extinction Probabilities of Terrestrial, Freshwater, and Marine Species Groups for Use in Life Cycle Assessment," *Ecological Indicators,* 142: 109204, ISSN 1470-160X. https://doi.org/10.1016/j.ecolind.2022.109204

Wan, J. (2023), "Revealed: Big Conservation NGOs are Majority Governed by Finance Figures," *African Arguments,* August 23. https://africanarguments.org/2023/08/revealed-big-conservation-ngos-majority-governed-finance-africa-carbon-markets/

Washbourne, C.-L., Murali, R., Saidi, N., Peter, S., Pisa, P. F., Sarzynski, T., Ryu, H., Filyushkina, A., Campagne, C. S., Kadykalo, A. N., Ávila-Flores, G., and Amiar, T. (2024), "Navigating the Science Policy Interface: A Co-Created Mind-Map to Support Early Career Research Contributions to Policy-Relevant Evidence," *Environmental Evidence,* 13(15). 10.1186/s13750-024-00334-5.

Watts M., Klein, C. J., Tulloch, V. J. D., Carvalho, S. B., Possingham, H. P. (2021), "Software for Prioritizing Conservation Actions Based on Probabilistic Information," *Conservation Biology,* July 16, 35(4): 1299-1308. doi: 10.1111/cobi.13681.

Werstein, P., Situ, H., and Huang, Z. (2006), "Load Balancing in a Cluster Computer," *Seventh International Conference on Parallel and Distributed Computing, Applications and Technologies (PDCAT'06),* Taipei, Taiwan, pp. 569-577. doi: 10.1109/PDCAT.2006.77.

Whipple, T. (2024), "Scientists Lose Our Credibility When We Become Climate Activists," *The Times,* Friday May 17. https://www.thetimes.com/uk/environment/article/scientists-lose-our-credibility-when-we-become-climate-activists-63bwxwhdk

World Population Review (2025), *Category: Environment.* https://worldpopulationreview.com/country-rankings/cheetah-population-by-country

Appendix: SA-MDAS Details and Comparisons

Continuous variables

All parameters in the political-ecological model studied in Section 4 are continuous. Hence, a method is needed for applying SA to continuously-valued variables. Here, the method developed in Corana et al. (1987) is used. One way to specify the length of SA's Markov chain, i.e., the number of passes through SA's inner loop, is to define it to be the size of the neighborhood surrounding the current solution point (Aarts and Korst, 1989, p. 65), i.e., the number of points reachable in one move. Given a small step length, the number of possible solution space points accessible in one move can become large.

Corana et al. (1987) take a different approach. They define their n_s constant to be how many times each variable is subjected to trial moves within one pass through SA's inner loop. Each variable is subjected to these n_s trials sequentially. Consequently, the value $n \times n_s$ becomes in-effect, their chain length. These authors do not attempt to equate their de-facto chain length definition with the number of possible points accessible in one move – and instead, simply recommend n_s be set to 20. In SA-MDAS, however, a move always consists of a possible step-change on each and every variable. Therefore, in SA-MDAS, n_s is the number of times that moves are made within one pass through SA's inner loop.

As in Corana et al. (1987), SA-MDAS recomputes each variable's step length every pass through the outer loop. Except for initialization, step length updates in SA-MDAS use equations given in (Corana et al., 1987, p. 267).

The step length's "varying criterion," Corana et al. (1987) and the step lengths themselves need to be initialized. This is accomplished by first assuming that the median is returned by the uniform random number generator on the unit interval, i.e., the value 0.5. Then, the initial median step length is initialized to be the width of the widest set of the n bound constraints divided by 50 – and then multiplied by this median value. The chain length is fixed at the value $10 \times n_s$ where $n_s = 3$.

The continuous-variable scheme of Corana et al. (1987) for SA maintains an approximately 50% acceptance rate at every temperature, t. But this is accomplished by progressively making the step length along each variable, smaller. Hence, as the temperature decreases, the chance that SA will break into a significantly more optimal subspace becomes small since, as the temperature decreases, a typical move will be close to the currrent point. Hence, at some temperature, the global search capability of SA will become practically nonexistent.

Stochastic objective functions

SA-MDAS can handle stochastic objective functions using a method developed by Bouttier and Gavra (2019). These authors define a "time" variable, t. To avoid confusion with this article's notation for "temperature," u is used here to refer to this variable. In the Bouttier and Gavra (2019) method, instead of evaluating the objective function once when a "score" evaluation is needed, the function is evaluated n_u times at the same point in the solution space. These n_u values are then averaged and returned as the score value at that point. The size of n_u increases as u (time) increases.

The value of u is updated every pass through the outer loop with $u = u + \zeta$ where ζ is a realization from the Exponential(1.0) distribution. Because the expected value of this distribution is 1.0, this update is, on average, adding the value 1.0 to u every pass through the outer loop. Addition of this stochastic term rather than the addition of the fixed value 1.0 to u every pass through the outer loop is necessary to allow the chain of SA-generated solutions to be modeled as a continuous-time Markov process. This is the model that is assumed in the proof of their method's convergence.

Further, their method requires a function mapping u to n_u. One such function that is used by the authors in their numerical experiments is

$$n_u = f(u) = \lfloor u \rfloor \tag{5}$$

where $\lfloor . \rfloor$ is the floor function.

In Step (g) of the algorithm in Section 3.4, Bouttier and Gavra (2019) have the temperature, t decreased by some function of u such as $t = 1/u$. But doing so ignores the standard deviation of the objective function estimated after each pass through SA's inner loop. Hence, SA-MDAS employs the Aarts and Korst (1989) algorithm to decrement t for both deterministic objective functions and for stochastic objective functions.

Inter-performer coupling

By executing its MPE Step, SA-MDAS couples the performers to each other by allowing them to always begin their respective chains at the current best solution that was found across all performers during the last outer loop iteration.

Message-passing overhead

Spreading tasks across a cluster of performers as is done in SA-MDAS is known to have some amount of *latency* (Pichetti et al., 2024). Here, such latency is viewed as a minor issue because the optimization problem's objective function, being composed of interacting agent/individual stochastic submodels, is computationally expensive to evaluate. Hence, most of the compute time needed to estimate the parameters of the political-ecological model will typically be taken by objective function evaluations rather than by inter-node messaging.

To verify this hunch, an experiment was conducted using Bukin's F4 function. This function takes very little compute time. SA-MDAS was run on a Windows$^{\mathrm{TM}}$ laptop computer with 32 GB of memory running at 1.2 GHz. Doing so resulted in SA-MDAS requiring 1.62 seconds to compute 192K evaluations of Bukin's F4 function. SA-MDAS was next run on the TSCC to perform 1,400 evaluations of this same function using two performers communicating with the impresario through a GigaSpace. Wall clock time for these evaluations was seven minutes. Almost all of this time was devoted to inter-node message-passing. These values suggest that a GigaSpaces implementation of SA-MDAS requires about 0.3 seconds to complete a *round-trip*, i.e., the posting of a task to the GigaSpace, a performer down-loading the task, uploading the result, and the impresario taking the result off of the GigaSpace. Note that a round-trip does not include the objective function's evaluation.

Let x be the number of minutes needed to evaluate the objective function. Assume a 100-node cluster computer is available. Then, the break-even function evaluation time can be estimated by solving for x in:

$$192000x = 192000x/100 + 192000(0.3/60) \tag{6}$$

The left side of 6 is the time needed to compute 192K function evaluations on a single-processor computer, while the right side is the time needed to compute the same number of evaluations on a cluster computer running 100 nodes in parallel. Equation 6 can be written as $0.99x = 0.3/60$, i.e., x is approximately equal to the roundtrip time. Hence, for expensive functions, cluster computing can speed-up black-box optimization problems even when round-trip times are as large as 0.3 seconds. And further, this speed-up increases as the objective function's evaluation time increases.

Evaluating the political-ecological model's CA objective function in the example of Section 4, requires about 15 seconds. Hence, an optimization algorithm that scales on the increasing availability of cluster computer nodes, is currently, the only feasible way to statistically fit a large political-ecological model to data.

Comparisons with other optimization algorithms

The optimization problem of least sum-of-squares possesses a very smooth, unimodal objective function. Using only phase one (SA), SA-MDAS needed 3,780 evaluations to solve a 27-variable least sum-of-squares problem. Whereas a Hooke and Jeeves algorithm (Haas, 2020) needed 1,725 evaluations, and a Random Search algorithm (Schumer and Steiglitz (1968)) needed 1,568 evaluations. These latter two algorithms are local, nonstochastic, and do not scale well as neither can make use of more than one compute node.

Similar to PACSA, SA-MDAS phase one mainly uses parallel processing to store the end result of many chains in order to increase the chance that it will find the global minimum point – not necessarily to speed up the solution time (Gonçalves-e-Silva et al., 2018). In other words, the focus of SA-MDAS phase one is to find the global minimum, not to speed up the solution. Even with this caveat, SA-MDAS phase one takes only about twice as long to solve a 27-variable least-squares problem as do two strictly local-search algorithms.

SA performs global search through decisions to probabilistically accept uphill moves (worse objective function values) rather than delineating a partitioning of the solution space first and then evaluating the function at some point in each partition as is done for example, in Jia et al. (2024). Such a partitioning-then-evaluation approach to global optimization results in many objective function evaluations regardless of whether they end up being uphill or downhill values. Further, there is an implicit assumption that all points within a partition have function values that are similar to the value at the single point actually evaluated in that partition.

MDAS, being more focused on local search than SA, employs many performers to

1. Perform limited global search (up to the m dimensions (variables) being searched simultaneously), and
2. Produce a speedup from the sequential Hooke and Jeeves algorithm by evaluating the objective function at all m-dimensions-ahead search points simultaneously rather than sequentially.