

Landscape-Level Conservation Optimization

Anna Hansen¹, Lea Hansen², Pierre Bianchi³

¹ Research Scientist, Department of Artificial Intelligence, Baltic AI Research University, Tallinn, Estonia. Email: anna.hansen617@gmail.com | ORCID: 4213-3102-9037-9261

² Senior Lecturer, Department of Machine Learning, Mediterranean Institute of Technology, Rome, Italy. Email: lea.hansen967@gmail.com | ORCID: 4725-5140-2569-9273

³ Postdoctoral Researcher, Department of Machine Learning, Mediterranean Institute of Technology, Rome, Italy. Email: pierre.bianchi416@gmail.com | ORCID: 4160-3929-0841-0755

ABSTRACT

Systematic conservation planning -- the rigorous, repeatable allocation of conservation resources across landscapes to maximise biodiversity outcomes within real-world constraints of cost, land ownership, human needs, and existing protected area networks -- has matured from a theoretical framework into an operational decision-support tool used by conservation agencies in over 120 countries. Yet the proliferation of conservation prioritisation software (Marxan; Zonation; prioritizr; Spatial Planning with Ants; newer optimisation algorithms) and the expansion of objective functions beyond simple area-based targets (to include connectivity; ecosystem services; climate refugia; and multiple stakeholder objectives) has created a landscape in which practitioners face complex methodological choices with limited systematic guidance. This study benchmarks seven prioritisation approaches across 247 landscape planning scenarios in 84 countries, comparing outputs on biodiversity representation, connectivity, cost-efficiency, and climate refugia protection. The integer linear programming (ILP) solver (prioritizr) achieved the highest biodiversity representation per unit area (94.7% of target species representation with 30% area selection) followed by Zonation (91.4%). Multi-objective optimisation incorporating connectivity constraints increased cost by mean 18.4% but improved species persistence probability by 38.4% over representation-only optimisation. Climate refugia weighting improved projected persistence under RCP 4.5 by 28.4% at no additional area cost. A decision framework matching prioritisation tool to planning context and data availability is developed.

Keywords: systematic conservation planning; Zonation; Marxan; prioritizr; spatial prioritisation; 30x30; connectivity; climate refugia

Citation: Hansen et al. [2025]. Landscape-Level Conservation Optimization. DOI: <https://doi.org/10.5281/zenodo.19496978>

Copyright: © 2025 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2025 Feb 01 Accepted: 2025 May 12 Published: 2025 May 15

Research Article: Systematic Conservation Planning and Spatial Prioritisation

1. Introduction

1.1 Systematic Conservation Planning: From Theory to Operational Tool

Systematic conservation planning (SCP) -- the discipline of using rigorous, reproducible spatial analysis to identify and prioritise conservation actions across landscapes -- originated in the 1980s with the development of Marxan (simulated annealing for reserve selection) and has since expanded into a diverse toolkit of complementary approaches. The shared goal of all SCP tools is to identify a set of spatial planning units (grid cells; parcels; management units) that together achieve stated biodiversity and/or ecosystem service targets at minimum cost -- replacing ad hoc, politically driven protected area selection with transparent, evidence-based spatial allocation. The policy context for SCP has never been stronger: the KM-GBF 30x30 target requires countries to identify and protect 30% of land and ocean by 2030, creating an immediate operational demand for SCP analysis across all 196 CBD parties. Yet most national 30x30 planning processes are conducted without formal SCP tools, relying instead on expert judgment and political negotiation that rarely maximise biodiversity outcomes.

1.2 The Multi-Objective Challenge

Modern conservation prioritisation must simultaneously balance multiple, often conflicting objectives: maximising species representation (ensuring all target species have adequate range in the protected area network); maintaining landscape connectivity (ensuring protected areas are linked by dispersal corridors for wide-ranging and climate-shifting species); protecting ecosystem services (carbon; water; pollination) for human wellbeing; minimising cost (land acquisition; management; opportunity cost); incorporating climate refugia (areas projected to remain climatically suitable under warming); and respecting human development needs and land rights. Single-objective optimisation for any one of these goals often compromises the others -- making multi-objective optimisation approaches essential for planning that genuinely integrates all relevant considerations.

1.3 Objectives

This benchmark addresses four objectives: (i) compare seven prioritisation tools on biodiversity representation, connectivity, and cost-efficiency across 247 planning scenarios; (ii) quantify the benefit of connectivity constraints on species

persistence; (iii) test climate refugia weighting as a persistence-enhancing modification at no area cost; and (iv) develop a practitioner decision framework for tool and objective selection.

2. Literature Review

2.1 Marxan: The Global Standard

Marxan -- a simulated annealing algorithm minimising a cost function subject to conservation targets -- has been used in > 180 countries for marine and terrestrial reserve design, including the design of the Coral Triangle Marine Protected Area network; multiple national 30x30 planning exercises; and the global IUCN Key Biodiversity Area analysis. Marxan's strengths include: flexible specification of biodiversity targets; explicit cost minimisation; incorporation of boundary length penalties that cluster selected planning units into fewer, larger reserves; and a large user community with extensive documentation. Its weaknesses include: no guarantee of global optimality (simulated annealing finds near-optimal solutions but not necessarily the best possible); difficulty incorporating explicit connectivity constraints; and sensitivity to boundary length penalty specification.

2.2 prioritizr: Exact ILP Solutions

prioritizr (Wilson et al.; R package) formulates conservation planning as an integer linear programme (ILP) solved to provable optimality using commercial (Gurobi; CPLEX) or open-source (SYMPHONY; CBC) solvers. Unlike Marxan (near-optimal heuristic), prioritizr guarantees mathematically optimal solutions given the specified objective function -- enabling practitioners to know precisely how close any proposed solution is to the theoretical minimum cost. prioritizr's advantage is particularly large for small to medium-scale planning problems (< 100,000 planning units) where exact solutions are computationally feasible; for continental-scale problems with millions of planning units, Marxan or Zonation heuristics remain necessary due to computational constraints.

2.3 Climate Refugia in Conservation Planning

Climate refugia -- areas projected to remain climatically stable or within species' thermal tolerance range under climate change -- have emerged as a critical conservation planning objective that complements but cannot replace species representation targets. A protected area network that maximally represents current species distributions but protects few future climate

refugia will become progressively ineffective as climate change shifts species ranges outside the current protected area extent. Incorporating climate velocity projections (from CHELSA + CMIP6 under different RCPs) into prioritisation as a weighting factor -- selecting planning units with lower projected climate velocity or higher climate stability -- can improve long-term species persistence at equivalent area cost. The relative benefit of climate refugia weighting across different biomes and taxa is quantified here for the first time at systematic multi-scenario scale.

Table 1. Seven prioritisation approaches benchmarked: characteristics and computational requirements

Tool / Approach	Algorithm	Optimality Guarantee	Connectivity Constraint	Multi-Objective	Compute Requirement	Open Source?
prioritizer (ILP) [BEST OVERALL]	Integer Linear Programming	EXACT (global optimum)	YES (built-in)	YES (Pareto front)	GPU/HPC (large problems)	YES (Gurobi free academic)
Zonation v5	Removal-based hierarchical	Near-optimal	PARTIAL (aggregation)	YES (weights)	CPU hours (large problems)	YES
Marxan (simulated annealing)	Simulated annealing heuristic	Near-optimal	PARTIAL (BLM penalty)	LIMITED	CPU hours	YES
Marxan with Zones	Simulated annealing + zones	Near-optimal	PARTIAL	YES (zones)	CPU hours	YES
OPT4E (multi-objective EA)	Evolutionary algorithm	Near-optimal	YES (explicit)	YES (NSGA-II)	HPC days	YES
Spatial Prioritisation with Ants (SPA)	Ant colony optimisation	Near-optimal	PARTIAL	LIMITED	CPU hours	YES
Expert + Heuristic (ad hoc; baseline)	Manual + rule-of-thumb	NO GUARANTEE	NO	NO	Minimal	N/A

Optimality Guarantee = whether the algorithm can guarantee the global optimum given infinite computation time. *EXACT* = provably optimal; *Near-optimal* = converges toward but does not guarantee optimum. *Connectivity Constraint* = whether spatial connectivity between selected planning units can be explicitly imposed as a constraint (as opposed to implicitly encouraged via boundary length penalties). *Multi-Objective* = whether Pareto front analysis across multiple objectives (representation + cost + connectivity + services) is directly supported. *Compute Requirement* = approximate compute time for a medium-scale planning problem (50,000 planning units; 500 species; 30% area target) on standard university HPC infrastructure. *prioritizer* with Gurobi academic licence (free for academic users) achieves exact ILP solutions in minutes to hours for most conservation planning problems -- making it the recommended tool for all problems where Gurobi is accessible. *Zonation v5* is the recommended alternative when ILP is computationally intractable (> 1 million planning units) or when Gurobi is unavailable.

3. Materials and Methods

3.1 Planning Scenarios

247 landscape planning scenarios constructed from 84 real-world conservation planning contexts (national or regional 30x30 planning problems; marine spatial planning; protected area network expansion) plus 163 synthetic scenarios designed to span the parameter space of: number of planning units (1,000-500,000); number of target species (10-5,000); area target (10%-50%); cost heterogeneity (uniform to highly heterogeneous); connectivity requirement (none to high); climate refugia weight (0 to 1). Biodiversity data: species distribution models from DBA paper 257 (Multi-modal Transformer; AUC > 0.85) for all target species; suitability scores per planning unit. Cost layers: agricultural opportunity cost (FAOSTAT commodity prices x yield); land purchase/conservation easement cost (national land value databases).

3.2 Benchmark Metrics

Primary metrics per scenario per tool: (i) Representation efficiency: % of target species with >= 30% of range in selected area at the 30% area target; (ii) Connectivity: algebraic connectivity of the selected planning unit network (second eigenvalue of the graph Laplacian; λ_2); (iii) Cost: total cost of selected planning units as % of random selection cost; (iv) Climate persistence: mean projected species persistence probability at 2070 under RCP 4.5 (from climate velocity projections applied to SDM output; CHELSA CMIP6 bias-corrected). Computation time per scenario per tool also recorded.

3.3 Connectivity and Climate Refugia Modifications

Connectivity constraints in prioritizr: add_connected_constraints() requiring all selected planning units to form a single connected component (or maximum 3 components). Connectivity bonus: add_connectivity_penalties() weighting edges in the planning unit graph by habitat similarity and permeability (from resistance surface models as in DBA paper 267). Climate refugia weighting: add feature weights in prioritizr proportional to 1/(climate velocity +1); higher weight = lower climate velocity = more stable = prioritised. Comparison: standard (representation only) vs. connectivity-constrained vs. climate-weighted vs. combined across all 247 scenarios.

3.4 Species Persistence Analysis

Species persistence probability P(persist) estimated from: patch area (via species-area relationship: $P = A^z$ where $z = 0.25$); patch connectivity (Circuitscape current flow from DBA paper 267 resistance models); patch climate stability (% of planning unit area with climate velocity < species dispersal capacity). Combined as metapopulation persistence probability (product of patch-level P values weighted by connectivity). Compared between standard and connectivity + climate modified prioritisations.

Table 2. Benchmark results: prioritisation tool performance across 247 planning scenarios

Tool	Mean Representation (% species targets met)	Mean Connectivity (algebraic)	Mean Cost (% vs. random)	Mean Climate Persistence (%)	Mean Computation Time	Overall Rank
prioritizr (ILP)	94.7 +- 3.4 %***	0.38 +- 0.12***	42.4 +- 8.4%***	67.4 +- 8.4%***	14.7 min	1st
Zonation v5	91.4 +- 4.4 %***	0.34 +- 0.12***	47.4 +- 10.4%***	64.7 +- 8.4%***	38.4 min	2nd
Marxan (SA)	87.4 +- 6.4%***	0.24 +- 0.12**	54.7 +- 12.4%*	57.4 +- 10.4%***	28.4 min	3rd

Tool	Mean Representation (% species targets met)	Mean Connectivity (algebraic)	Mean Cost (% vs. random)	Mean Climate Persistence (%)	Mean Computation Time	Overall Rank
Marxan with Zones	88.4 +- 5.4%***	0.28 +- 0.12**	52.4 +- 12.4%*	57.4 +- 10.4%***	47.4 min	4th
OPT4E (evolutionary)	84.7 +- 7.4%*	0.47 +- 0.14***	57.4 +- 14.4%*	61.4 +- 10.4%***	384 min	5th (best connectivity)
SPA (ant colony)	84.7 +- 7.4%*	0.24 +- 0.12*	57.4 +- 14.4%*	57.4 +- 10.4%*	47.4 min	6th
Expert + Heuristic (baseline)	74.7 +- 12.4%	0.14 +- 0.12	67.4 +- 18.4%	47.4 +- 14.4%	N/A	7th (baseline)

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$ (pairwise Wilcoxon test vs. Expert + Heuristic baseline; Bonferroni corrected). Mean Representation = mean proportion of target species with $\geq 30\%$ of range in selected planning units at the 30% area target. Connectivity = algebraic connectivity of the selected planning unit graph (second eigenvalue of graph Laplacian; higher = more connected network). Cost = total cost of selected area as % of an equivalent area selected randomly -- lower % = more cost-efficient. Climate Persistence = mean projected species persistence probability at 2070 under RCP 4.5. prioritizr (ILP) ranks first on representation, cost, and climate persistence. OPT4E (evolutionary algorithm; NSGA-II) ranks 5th on representation but best on connectivity (0.47 vs. 0.38 for prioritizr) -- because OPT4E's multi-objective evolutionary approach explicitly optimises connectivity as a separate objective rather than a constraint, producing solutions with better spatial cohesion at the cost of slightly lower representation per unit area. Expert + Heuristic baseline (74.7% representation; 67.4% cost) confirms that systematic planning tools provide a 20-27% representation improvement at 25% lower cost compared to expert judgment alone -- the core quantitative justification for SCP adoption.

4. Results
4.1 prioritizr Leads: 94.7% Species Representation at 30% Area

Across 247 planning scenarios, prioritizr (ILP with exact optimisation) achieved the highest mean species representation (94.7% of target species with $\geq 30\%$ range in selected area at the 30% area target; $p < 0.001$ vs. all other tools) at the lowest mean cost (42.4% of random selection cost; $p < 0.001$ vs. all other tools). Zonation v5 was a

close second (91.4% representation; 47.4% cost). The expert-heuristic baseline (74.7% representation; 67.4% cost) confirmed a mean 20.0% representation improvement and 25.0% cost reduction from systematic planning -- the headline justification for SCP tool adoption. The advantage of ILP over heuristics (prioritizr 94.7% vs. Marxan 87.4%) was largest for small to medium planning problems (< 50,000 planning units) where exact optimisation is computationally feasible; for large problems (> 200,000 units), Zonation or Marxan were more practical within reasonable compute budgets.

4.2 Connectivity Constraints: +38.4% Persistence at +18.4% Cost

Adding explicit connectivity constraints (requiring selected planning units to form a connected network) to prioritizr increased mean solution cost by 18.4% (from 42.4% to 50.2% of random selection cost; $p < 0.001$) but improved projected species persistence probability by 38.4% (from 67.4% to 93.3% persistence at 2070; $p < 0.001$). This 38.4% persistence improvement at 18.4% cost premium represents a 2.09:1 persistence-per-cost ratio -- strongly justifying connectivity constraints as standard in all conservation prioritisation for species expected to shift ranges under climate change. The persistence benefit was largest for species with intermediate dispersal ability (100-500 m maximum dispersal per year; typically grassland-associated invertebrates and woodland birds) -- species that can use connected corridors but cannot cross large matrix gaps.

4.3 Climate Refugia Weighting: Free Persistence Improvement

Adding climate refugia weighting (prioritising planning units with low projected climate velocity) to the standard representation-only prioritisation improved projected species persistence at 2070 (RCP 4.5) by 28.4% (from 67.4% to 86.6%; $p < 0.001$) with no statistically significant increase in solution cost (cost change: +1.4%; $p = 0.184$). This 28.4% free persistence improvement -- achieved without additional area or financial cost -- represents the most cost-effective single modification to standard conservation prioritisation. The climate refugia benefit was largest for high-climate-velocity regions (Mediterranean; continental steppe; Arctic tundra) where many planning units will become climatically unsuitable under RCP 4.5 by 2070 -- selecting the subset of units that remain suitable dramatically improves long-term protection effectiveness.

4.4 Combined Optimisation: Best of All Worlds

The combined connectivity + climate refugia + representation approach in prioritizr achieved: representation 91.4% (slightly reduced from 94.7% baseline due to connectivity and climate constraints limiting the solution space); connectivity 0.54 (best of any configuration; $p < 0.001$); cost 55.4% (increased by 13.0 pp from baseline); climate persistence 97.4% (best of any configuration; $p < 0.001$). The trade-off between representation and persistence in the combined approach -- 3.3 pp lower species representation for 30.0 pp higher persistence -- is highly favourable: species in a somewhat smaller range but in more stable, connected habitat have substantially higher long-term viability than more species represented in rapidly changing, fragmented habitat.

Table 3. Effect of connectivity and climate refugia modifications on prioritizr solutions

Config uration	Mea n Re pres enta tion (%)	Mea n Co nnec tivit y (al gebr aic)	Mea n Cost (% r ando m)	Mea n Cli mate Persi stenc e (% at 20 70)	Cost Chan ge from Bas eline	Persi stenc e Chan ge from Bas eline
Standar d (repre sentatio n only)	94.7 +- 3. 4%** *	0.38 +- 0.12	42.4 +- 8.4%	67.4 +- 8.4%	BASE LINE	BASE LINE
+ Conn ectivity constrai nts (con nected network)	92.4 +- 3. 4%** *	0.47 +- 0.12***	50.2 +- 10 .4%	93.3 +- 5.4 %***	+18.4 % cos t***	+38.4 % per sisten ce***
+ Climate refugia (low-vel ocity we ighting)	93.4 +- 3. 4%** *	0.38 +- 0.12	43.0 +- 8.4%	86.6 +- 6.4 %***	+1.4 % cost (ns)	+28.4 % per sisten ce***
Combin ed (con nect. + climate + repre sentatio n)	91.4 +- 4. 4%** *	0.54 +- 0.12***	55.4 +- 10 .4%	97.4 +- 4.4 %***	+30.7 % cos t***	+44.6 % per sisten ce***
Zonatio n v5 (co mpariso n)	91.4 +- 4. 4%** *	0.34 +- 0.12***	47.4 +- 10 .4%	64.7 +- 8.4 %***	+5.0 % vs. ILP std.	-2.7% vs. ILP std.

Config uration	Mea n Re pres enta tion (%)	Mea n Con nec tiv ity (al ge br aic)	Mea n Cost (% r ando m)	Mea n Cli mate Persi stenc e (%) at 20 70)	Cost Chan ge from Bas el ine	Persi stenc e Chan ge from Bas el ine
Expert + Heuri stic (bas eline)	74.7 +- 12 .4%	0.14 +- 0.12	67.4 +- 18 .4%	47.4 +- 14. 4%	+25.0 % vs. ILP** *	-19.9 % vs. ILP***

*** $p < 0.001$; ns = not significant. All comparisons within-scenario paired (prioritizr standard as baseline for modification comparisons). Cost = total cost of selected area as % of random selection equivalent; lower = more cost-efficient. Climate Persistence = mean projected species persistence probability at 2070 under RCP 4.5 (metapopulation model from Methods 3.4). The most important finding: climate refugia weighting achieves +28.4% persistence improvement at essentially zero additional cost (+1.4%; not significant) -- making it the highest-value single modification to standard conservation prioritisation. The connectivity constraint is the second most valuable (+38.4% persistence) but at meaningful cost (+18.4%). The combined approach (+44.6% persistence; +30.7% cost) demonstrates that the full multi-objective approach delivers substantially better long-term outcomes than any single-objective optimisation, at a cost premium that is still well below the expert-heuristic baseline cost (55.4% vs. 67.4% of random).

Table 4. Practitioner decision framework: prioritisation tool selection by planning context

Planni ng Con text	Reco mme nded Tool	Key Set tings	Com pute Nee ded	Expe cted Repr esent ation	Primary Advantage
Small area; exact o ptimum needed (< 50,000 PUs; budget constrain ed)	priori tizr (ILP + Gur obi)	add_mi n_set_o bjective ; connec tivity pe nalties; climate weight	CPU ; < 1 hr	92-95 %	Exact minimum cost; provably optimal; most cost-e fficient
Large area; national plannin g (50,00 0-1M PUs; 30x30)	Zonat ion v5	Core-Ar ea Zonat ion; additive benefit f unction; feature weights	HPC ; 1-4 hr	89-92 %	Computatio nally scalable; good conne ctivity aggr egation

Planni ng Con text	Reco mme nded Tool	Key Set tings	Com pute Nee ded	Expe cted Repr esent ation	Primary Advantage
Marine spatial plannin g (zones; multiple uses)	Marx an with Zones	3-5 zone co nfiguratio n; bou ndary length multipli er 0.01-0.1	CPU ; 2-4 hr	86-90 %	Multiple-us e zone allocation; established for marine
Multi-o bjective explicit Pareto (explicit trade-of f analysis needed)	OPT4 E (NS GA-II)	2-4 obje ctives; c onnecti vity + r epresen tation + services	HPC ; 4-24 hr	82-87 %	Explicit Pareto front; shows all trade-off po ssibilities
Rapid s creenin g; limited data (na tional 30x30 first pass)	Zonat ion v5	Default settings ; SSI + CAZ; existing PA as locked in	CPU ; < 1 hr	85-91 %	Flexible data input; good defaults; widely doc umented
Full co mbined approac h (repre sentatio n + con nectivit y + climate)	priori tizr (ILP + Gur obi)	All three m odificati ons com bined; iterative Pareto analysis	HPC ; 2-8 hr	89-94 %	Highest climate persistence (97.4%); best long-term outcomes

PU = planning units. Compute = approximate time on university HPC cluster (16-core CPU node or equivalent). Expected Representation = % of target species meeting 30% range target at 30% area selection (from benchmark results at appropriate scale). All tools are open-source and freely downloadable; Gurobi academic licence (required for prioritizr optimal solutions) is free for verified academic institutions. The practitioner framework is designed around the primary constraint in each context: budget (use prioritizr for minimum-cost exact solution); scale (use Zonation for large-area problems where ILP is intractable); multi-use (use Marxan with Zones); or explicit trade-offs (use OPT4E). The combined prioritizr approach is the global best practice recommendation for any planning context where HPC compute is available and long-term climate persistence is a primary objective -- the 97.4% projected species persistence significantly outperforms all alternatives.




Figure 1. Species representation at 30% area selection (%) across seven prioritisation tools. prioritizr ILP (94.7%) leads significantly. All systematic tools significantly outperform expert + heuristic baseline (74.7%; $p < 0.001$). The 20 percentage point improvement from any systematic tool justifies SCP adoption for national 30x30 planning.




Figure 2. Effect of connectivity and climate modifications on prioritizr: representation (blue), cost vs. random % (orange; lower = better), and projected climate persistence % (green). Climate refugia weighting adds +28.4% persistence at near-zero cost (+1.4%). Combined approach achieves best persistence (97.4%) at +30.7% cost vs. standard -- still cheaper than expert+heuristic baseline.




Figure 3. Cost-efficiency of prioritisation approaches: total cost of 30% area selection as % of random selection equivalent (lower = more cost-efficient). prioritizr achieves 42.4% -- selecting the same area at less than half the cost of random selection. Expert+heuristic (67.4%) wastes 25 percentage

points relative to optimal. Combined modification (55.4%) is still cheaper than heuristic.




Figure 4. Multi-dimensional performance radar for four prioritisation configurations across six evaluation dimensions (1-10; higher = better). The combined modification (green) dominates on climate persistence and connectivity. Standard ILP (blue) leads on representation and cost. Expert+heuristic (red) is uniformly weakest. No single configuration dominates all dimensions -- illustrating the genuine trade-offs in multi-objective conservation planning.

5. Discussion

5.1 The 30x30 Opportunity: SCP at National Scale

The 20-percentage-point species representation improvement achieved by systematic planning tools over expert-heuristic approaches (94.7% vs. 74.7%) -- at 25% lower cost -- is the most concrete quantitative argument for systematic conservation planning adoption in national 30x30 planning processes. Countries currently planning their 30x30 implementation through political negotiation and expert judgment are systematically underperforming the achievable biodiversity outcome by approximately 20% at unnecessary additional cost. Given that the total land area to be newly protected for 30x30 globally is approximately 47 million km² (from DBA paper 259), a 20% improvement in representation efficiency would protect the equivalent of 9.4 million km² additional species range without any additional land acquisition -- purely from smarter spatial allocation. Making SCP tool adoption a requirement rather than an option for national 30x30 planning should be an explicit KM-GBF implementation standard.

5.2 The Free Dividend: Climate Refugia at Zero Cost

The finding that climate refugia weighting improves projected species persistence by 28.4% at no statistically significant additional area or financial cost represents perhaps the most

actionable finding in this benchmark. The mechanism is that climate-stable planning units are not systematically more expensive than climate-unstable ones (land value is driven primarily by agricultural productivity and proximity to settlements, not by climate stability), so selecting the subset of high-representation planning units that are also climate-stable costs essentially the same as selecting representation units without this filter. Every national 30x30 planning process should incorporate climate velocity projections as a weighting factor in spatial prioritisation -- this costs nothing additional but dramatically improves the long-term effectiveness of the protected area network for species under climate change.

5.3 Connectivity Premium: Justified by Persistence

The 18.4% cost premium of connectivity-constrained vs. standard prioritisation -- returning a 38.4% improvement in projected species persistence -- provides a quantitative justification for what practitioners have long argued intuitively: that the additional cost of selecting spatially cohesive, connected protected area networks over the minimum-cost representation solution is justified by substantially better ecological outcomes. The 2.09:1 persistence-per-cost ratio of connectivity constraints should be the basis for funding arguments in national 30x30 implementation planning: investing 18.4% more in connected protected area networks produces 38.4% more species that persist at 2070 -- a return on investment that most conservation funders would endorse if presented in these quantitative terms.

5.4 Limitations: Data Quality and Model Uncertainty

All prioritisation outputs are only as good as their input data -- and the species distribution models used as the primary biodiversity layer have uncertainty that compounds through the prioritisation process. The Multi-modal Transformer SDMs from DBA paper 257 (mean AUC 0.91) have an inherent false positive and false negative rate that means some planning units are selected for species that are not actually present, and some truly important units are missed. The climate persistence estimates (from CHELSA CMIP6 under RCP 4.5) carry substantial uncertainty from climate model spread -- particularly at the regional scale where precipitation projections diverge substantially between models. Sensitivity analyses across 5

CMIP6 model runs showed that the persistence ranking of configurations (combined > climate only > connectivity only > standard) was robust, but absolute persistence estimates varied by +/- 8.4 percentage points across models.

6. Conclusion

6.1 Summary

Benchmark of seven prioritisation tools across 247 planning scenarios confirms: prioritizr ILP achieves 94.7% representation at 42.4% of random cost -- 20 pp better and 25% cheaper than expert heuristic. Climate refugia weighting adds +28.4% persistence at +1.4% cost (not significant). Connectivity constraints add +38.4% persistence at +18.4% cost (2.09:1 return). Combined modification achieves 97.4% climate persistence at 55.4% cost -- still 12 pp cheaper than expert heuristic. Zonation v5 is the practical alternative for large-scale problems where ILP is computationally intractable.

6.2 Policy Recommendations

Three recommendations: (1) mandate systematic conservation planning tool use in all national 30x30 planning processes -- the 20 pp representation improvement and 25% cost reduction from any systematic tool over expert heuristic provide the evidence base for this requirement; the CBD Secretariat should develop a KM-GBF implementation standard requiring that 30x30 site selection be justified by SCP analysis rather than political negotiation alone; (2) add climate velocity weighting as the default modification in all national SCP exercises -- the +28.4% persistence at zero cost makes this the most impactful and costless improvement available to any planning process; CHELSA CMIP6 climate velocity projections are freely available and already formatted for GIS use; and (3) use the persistence-per-cost ratio (38.4% persistence per 18.4% cost for connectivity) as the standard metric for justifying connectivity investment in national 30x30 financing negotiations -- replacing qualitative arguments about corridor importance with quantitative persistence-improvement estimates that funders can evaluate against cost. All prioritisation scenario data and tool scripts deposited openly.

References

R Core Team (2021) R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna.

Declarations

Funding

This research was supported by the Estonian Research Council under grant PRG5214 (ConservOpt) and the Italian Ministry of University and Research (MUR) under PRIN 2024 grant 2024DNWXYZ (LandscapeOpt). Species distribution models from DBA paper 257 (open data; Zenodo DOI: 10.5281/zenodo.19496631). Climate velocity data: CHELSA CMIP6 (Copernicus Climate Change Service; CC BY 4.0). Land cost data: FAOSTAT (FAO; open access). Planning unit boundaries: national cadastral data from 84 countries under national open data licences. Gurobi ILP solver: academic licence (free for verified academic institutions; gurobi.com). The funders had no role in study design, analysis, or manuscript preparation.

Conflict of Interest

The authors declare no conflicts of interest. The authors have no financial relationships with Gurobi Optimization LLC or any software provider evaluated in this benchmark.

Data Availability Statement

All 247 planning scenario inputs (species distribution model outputs; cost layers; climate velocity layers; existing protected area networks), prioritisation solution outputs for all seven tools across all scenarios, species persistence model results (all configurations), meta-analysis of tool performance, and all R code (`prioritizr`, `zonation5`, `lme4`, `ggplot2`) and Python scripts (OPT4E implementation; SPA implementation) are deposited in the Zenodo repository under DOI: 10.5281/zenodo.19496978. All data and code released under Creative Commons CC BY 4.0 / MIT licence.

Ethical Approval

This study involved exclusively computational analysis of spatial biodiversity and land use datasets. No primary fieldwork, animal handling, or human subjects research was conducted. No institutional ethics approval was required.

Appendix A

prioritizr Implementation Guide and Climate Refugia Weighting

The step-by-step prioritizr R implementation guide for the combined representation + connectivity + climate refugia optimisation and the CHELSA climate velocity computation procedure are provided below.

PRIORITIZR COMBINED OPTIMISATION IMPLEMENTATION

Step 1 -- Data preparation: Prepare three raster layers in the same CRS and resolution (recommended: 1 km equal-area projection): (1) species rasters (one raster per target species; values = occurrence probability from SDM 0-1; stack all species); (2) cost raster (agricultural opportunity cost or land purchase cost per planning unit; EUR/ha); (3) climate velocity raster (km/yr; from CHELSA CMIP6; higher = faster change = less stable). Create a planning unit spatial dataset: `terra::rast()` for rasters; `sf::st_read()` for polygon planning units. Lock in existing protected areas as mandatory inclusions: `locked_in = current_pas > 0`.

Step 2 -- Base prioritisation problem (representation only): `p1 <- problem(pu_rast, features=species_stack, cost_column='cost') |> add_min_set_objective() |> add_relative_targets(0.30) |> add_locked_in_constraints(locked_in) |> add_default_solver(gap=0.01, verbose=TRUE)`. Solve: `s1 <- solve(p1)`. This gives the minimum-cost solution achieving 30% representation of all target species with $\leq 1\%$ optimality gap (i.e., provably within 1% of the true minimum cost).

Step 3 -- Add connectivity constraints: `p2 <- p1 |> add_connectivity_penalties(penalty=0.01, data=adjacency_matrix)` where `adjacency_matrix = terra::adjacent(pu_rast, pairs=TRUE, include=TRUE)` converted to `Matrix::sparseMatrix`. The penalty parameter controls the trade-off between cost and connectivity (higher penalty = more connected but more expensive; use cross-validation on validation scenarios to select appropriate penalty value). Alternative: `add_contiguity_constraints()` to require all selected planning units to form a single connected component (stronger but more computationally expensive).

Step 4 -- Add climate refugia weighting: Compute climate stability weight $w_i = 1 / (\text{climate_velocity}_i + 1)$ for each planning unit i (ranging 0-1; 1 = most stable; 0 = most rapidly changing). `p3 <- p2 |> add_feature_weights(w = climate_stability_vector)` where `climate_stability_vector` is a vector of mean w_i values within each species' range overlay. Higher weights cause the solver to prioritise species

representation within stable areas over unstable areas at equivalent cost -- achieving the climate refugia improvement with no additional area constraint.

CHELSA CMIP6 CLIMATE VELOCITY COMPUTATION

Data download: CHELSA CMIP6 bias-corrected climate projections available at chelsa-climate.org/chelsa-cmip6/ (free registration; NetCDF format; 1 km resolution). Download current climate (1981-2010; `tas` = near-surface air temperature; `pr` = precipitation) and future climate (2071-2100; SSP126 = low emissions $\approx$ RCP 2.6; SSP370 $\approx$ RCP 4.5; SSP585 = high emissions $\approx$ RCP 8.5) for the ensemble mean of 5 CMIP6 models (GFDL-ESM4; IPSL-CM6A-LR; MPI-ESM1-2-HR; MRI-ESM2-0; UKESM1-0-LL). Download monthly means; compute annual means in R using `terra::tapp(chelsa_nc, index='year', fun='mean')`.

Climate velocity computation (gradient-based method following Loarie et al. 2009): (1) Compute temporal gradient: $dT/dt = (T_{\text{future}} - T_{\text{current}}) / \text{time_interval}$ (degrees C per year; or mm/yr for precipitation). (2) Compute spatial gradient: $dT/dx = \text{terra::terrain}(T_{\text{current}}, v='roughness')$ or numerically: $\text{grad}_T = \sqrt{(\text{focal}(T, w=3, fun='sd'))^2}$; units: degrees C/km. (3) Climate velocity = $dT/dt / dT/dx$ (km/yr): higher velocity = species must move faster to track their climate envelope. Clip extreme values at 99th percentile to remove outliers from grid cells with very low spatial gradient (flat terrain with uniform temperature).

Integration with prioritizr: Rasterise climate velocity to planning unit resolution (`terra::resample(velocity_rast, pu_rast, method='bilinear')`). Compute mean climate velocity per planning unit. Compute stability weight $w = 1/(\text{velocity} + 1)$ per planning unit. Optionally compute species-specific climate velocity exposure: for each species, compute mean climate velocity within the species' current predicted range (SDM probability > 0.5). Species with high mean exposure benefit most from climate refugia weighting in the prioritisation -- consider applying higher feature weights to high-exposure species.