https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

Projecting global changes in land use and ecosystem services using SEALS (Spatial Economic Allocation Landscape Simulator) version 2.0

Justin A. Johnson[1,2], Chiara Colesanti Senni[2,3], Sumil K. Thakrar[2,4], Patrick José von Jeetze[5]

[1]Department of Applied Economics, University of Minnesota; St. Paul, MN, USA.
[2]NatCap TEEMs, University of Minnesota; St. Paul, MN, USA
[3]Department of Finance, University of Zurich; Zurich, Switzerland.
[4]School of GeoSciences, University of Edinburgh; Edinburgh, UK.
[5]Potsdam Institute for Climate Impact Research (PIK), Potsdam, Germany

Correspondence to: Justin A. Johnson (jajohns@umn.edu)

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

Abstract. Land-use change alters ecosystems and the services they provide, but calculating these effects requires high-resolution land-use maps. Ecosystem service models rely on satellite-derived maps for present and historical conditions, using spatial resolutions of 10 to 300 meters. Scenarios of future land-use dynamics, however, are typically provided at resolutions of 0.25 arc-degrees (~30 km at the equator) or coarser, far too coarse for ecosystem-service models. Here, we describe and evaluate the Spatial Economic Allocation Landscape Simulator (SEALS) model, a global, computationally tractable model that downscales coarse land-use projections into high-resolution maps suitable for ecosystem service analysis. SEALS estimates 3.12 million parameters applied to spatial layers for land constraints, transition suitability and multi-scale adjacency convolutions, identified separately for each 1-degree tile globally. The parameters are trained on observed land-use change from 2000 to 2015 and validated on withheld data from 2020. SEALS can generate global 10 arc-second land-use maps in minutes on a standard laptop. We show that SEALS places change 1.9 times more accurately than a random allocation of the same change and that the advantage persists when near-misses are credited. Passing the downscaled maps to the InVEST habitat-quality model reproduces habitat quality calculated from observed land cover more closely than coarse accounting in nine of ten deforestation regions, showing that resolving where change falls improves a downstream ecosystem-service estimate. As a demonstration we downscale eight Shared Socioeconomic Pathway scenarios to 2100 and release the resulting global 10 arc-second land-use maps.

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

1 Introduction

Economic activity is a major driver of global change in land use and land cover (hereafter "land use"), including deforestation and agricultural production (IPBES, 2019). The impact these changes have on nature depends on their fine-scale pattern (Kim et al., 2018; Bukovsky et al., 2021; Bagstad et al., 2018). Understanding how these changes in nature affect human wellbeing requires ecosystem-service models, which depend on land-use information at a matching resolution, typically 10 to 300 meters (Chaplin-Kramer et al., 2019).

High-resolution land-use maps derived from satellite data, such as those from the European Space Agency's Climate Change Initiative (ESA-CCI) (ESA, 2017; Lamarche et al., 2017), can be used to estimate current or historic provision of ecosystem services (Chaplin-Kramer et al., 2019). However, these maps capture only observed conditions and cannot be used to explore future projections or alternative policy scenarios. Integrated assessment models (IAMs), such as those underlying the Shared Socioeconomic Pathway (SSP) scenarios (Riahi et al., 2017) and their land-use representation (Popp et al., 2017; Hurtt et al., 2020), link economic activity to land-use change across space, but they operate at resolutions too coarse to resolve many biophysical processes or ecosystem services (Britz et al., 2011).

In this paper we present version 2.0 of the Spatial Economic Allocation Landscape Simulator (SEALS), a model that downscales land-use data from regional or coarse inputs to fine resolution. SEALS has been pivotal in developing Earth-economy models (GTAP-InVEST; Johnson et al., 2023, 2025) and IAMs (MAgPIE; von Jeetze et al., 2023, 2025). Its methods have been partially documented in these papers, but there is not yet a comprehensive, peer-reviewed article describing the full model. This paper fills that gap: it describes the model in full, evaluates where it places land-use change against a withheld period, tests whether that placement improves a downstream ecosystem-service estimate, and compares SEALS with alternative downscaling approaches.

Several existing models estimate high-resolution land-use change from coarser inputs, using spatial suitability layers (for example soil, climate, topography, and accessibility) together with allocation rules that determine which land-use conversions occur and where. Existing models, however, are each constrained in at least one of three ways: the output is too small to be global at fine resolution, the allocation coefficients are asserted rather than estimated from observed change, or the placement of change is never validated against a period the model did not see. The seminal models in this space, such as CLUMondo (van Asselen & Verburg, 2013; Wolff et al., 2018) and its variants, are computationally limited to scales much smaller than the 129,600 by 64,800 pixels in ESA-CCI maps, typically bound by memory constraints. This means the generated outputs are either not global or are much coarser than 10 arc-seconds. The GLOBIO 4 downscaling model (Schipper et al., 2020) can produce global maps at the 10 arc-second

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

resolution that ecosystem-service models need, but the coefficients applied to suitability layers are defined as assumptions rather than estimated from observed land-use change.

Other downscalers do estimate their parameters from data, including generalized additive models fitted per land-use class (Hoskins et al., 2016) and prior-corrected multinomial logit models of the kind implemented in downscalr (Krisztin et al., 2023), but these fit a suitability or transition response and validate that fit, rather than calibrating how change is placed within a cell and testing that placement on a withheld period, and they have been demonstrated at regional to continental extent or at coarser resolution (30 arc-seconds). Section 5.3 and Supplement Section S4 compare these models with SEALS in detail.

What distinguishes SEALS is the combination of four properties: allocation parameters estimated against observed land-use change and allowed to vary in space; adjacency effects represented as calibrated convolutions acting at several distances; out-of-sample validation of where change is placed, on a withheld period; and, highly optimized, parallel computation able to downscale global, high-resolution maps quickly. To our knowledge no existing model combines these at global extent at resolutions necessary for ecosystem service models. SEALS estimates its parameters from observed historical change and uses them to allocate projected change through empirically grounded relationships between land-use change and its spatial drivers. SEALS is open source and produces global land-use projections at 10 arc-second resolution in minutes on a standard laptop and can handle 1/3 arc-second resolution on a high-performance workstation. Users can supply their own land-use classes (recalibrating the parameters for them), spatial predictors, and output resolution, or draw on a curated set of global, publicly available base data distributed with the model.

Earlier studies applied simplified versions of SEALS to a single land-use type, such as maize expansion (Suh et al., 2020), and to specific conservation goals (Johnson et al., 2020, 2021b). Here we present the underlying model, which downscales all land-use types simultaneously, is designed to apply to any coarse land-use projection, and uses empirically calibrated parameters. Version 2.0 (the version documented in this paper), is able to downscale both coarse raster inputs (such as those derived from IAMs, such as the LUH2 data), and regional vector inputs (such as country-level estimates of land-use change). Section 5.2 describes the larger modelling pipeline that carries both regional and gridded projections through to fine-scale scenario outputs; this paper documents and evaluates the fine-downscaling step.

This paper describes and evaluates the SEALS model. Section 2 defines the model, presenting the downscaling problem and its conservation constraint, the three components of transition suitability, and the allocation procedure. Section 3 describes how the parameters are estimated from observed land-use change, how the resulting placement is validated on a withheld period, and the settings of the application reported here. Section 4 reports the results: the held-out validation of placement skill, a test of whether resolving placement improves an ecosystem-

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

service estimate, a demonstration in which SEALS downscales eight SSP-RCP scenarios to 2100 (whose maps we release as a dataset), and the model's computational performance. Section 5 discusses what the resolved placement provides, where SEALS sits in the modelling pipeline, how it relates to other downscaling models, and its limitations; Section 6 concludes. The scenario runs are included as a demonstration of the calibrated model and as a resource for users.

We train SEALS on observed land-use change from 2000 to 2015 and evaluate it against observed change in 2020, then use it to downscale global land-use change for each SSP-RCP scenario to 10 arc-second resolution. Resolving where land-use change occurs matters because many ecosystem services depend on the spatial arrangement of that conversion at scales below 1 km, as well as on how much land converts. We demonstrate this with habitat quality, a service whose degradation depends on the proximity of conversion to the habitat that remains. Habitat quality is also a common measure of biodiversity change, so resolving where conversion falls supports the assessment of biodiversity loss as well as ecosystem services. We show out of sample that the fine-resolution placement reproduces observed habitat quality more closely than a coarse accounting that knows only how much each region converts, and than a fine-resolution map that places the same change at random.

2 Model description

SEALS downscales a coarse projection of land-use change to fine resolution by learning spatial patterns from historical change and applying them to allocate the projected future change (Figure 1). It estimates, for each fine-scale target region (a fine cell within a coarse source cell), a suitability for transition to each land-use class from three components: land constraints, land suitability, and adjacency effects. SEALS ranks the target regions by their suitabilities and allocates any proposed coarse change to those ranked most suitable for it. This section gives the formulation of each step; Section 3 describes how the parameters are estimated and how the resulting placement is validated.

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

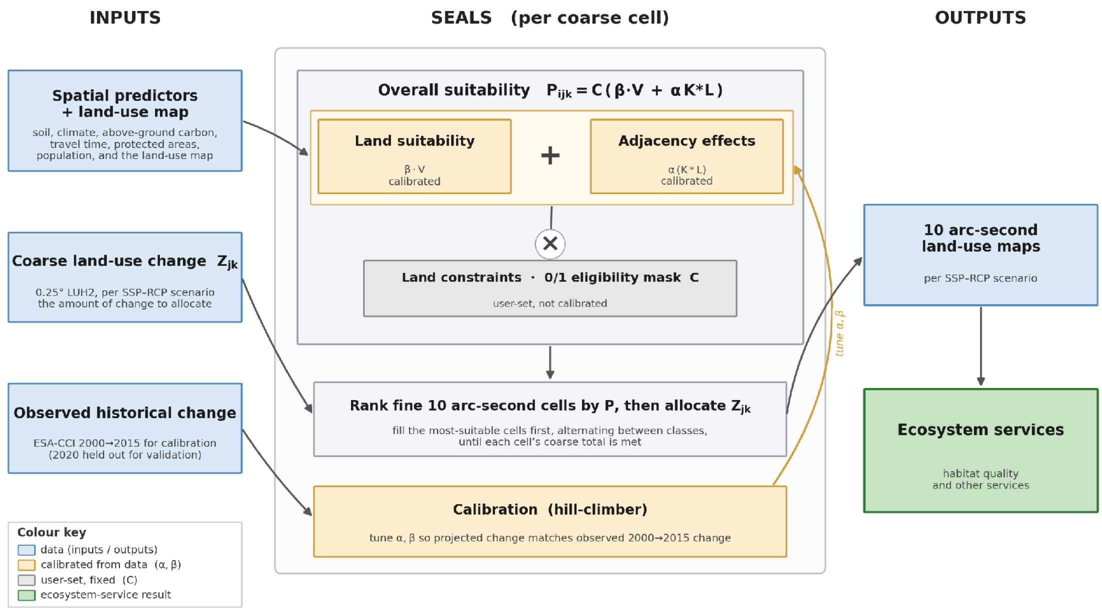

Figure 1. Overview of the SEALS workflow. For each coarse cell, land suitability (from the spatial predictors) and adjacency effects (convolutions of neighboring land-use classes) combine under a 0/1 land-constraint eligibility mask into an overall suitability. SEALS then ranks the fine 10 arc-second cells by suitability and allocates the coarse land-use change by taking the classes in turn, one cell at a time, until each class's coarse total is met. The coefficients are calibrated by hill-climbing so that projected change matches the observed ESA-CCI 2000-2015 change. The 10 arc-second land-use outputs feed ecosystem-service models such as habitat quality.

2.1 Definition of the downscaling problem

Our goal is to assign each fine target region i, inside coarse source region j, a land-use class k, written $\hat{Y}_{ijk}$, consistent with the coarse totals. Each source region j contains many target regions i, and the assignment is subject to the constraint that:

$$\sum_i (\hat{Y}_{ijk} - \hat{Y}^0_{ijk}) w_{ij} = Z_{jk}$$

where $\hat{Y}^0_{ijk}$ is the baseline (initial) state of each target region, Z_{jk} is the areal change of class k projected for source region j, and w_{ij} is the ground area of target region i. The weight w_{ij} turns the 0/1 occupancy into a physical area (hectares), so the left side is the change in the area of class k within the cell. Allocation proceeds by placing expansions in the highest-ranked locations in the calibrated suitability score (P_{ijk}, discussed below) to meet its Z_{jk}. Where a class is contracting, these locations are derived from where it loses area to another class's expansion. Expansions are allocated in whole fine cells, the smallest unit that can be placed. Class totals can differ from their targets when eligible cells are exhausted or when competing expansions displace contracting classes. Supplement Section S2.2 quantifies these deviations for a sample region.

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

Supplement Section S2.2 describes this allocation routine in more detail and illustrates it for a sample region in Mato Grosso, Brazil.

The source and target grids form a nested pyramid: each source region j contains an integer number of target regions i with no fractional overlap, which allows area changes to be accounted for within each source region, since the near-equal target areas w_{ij} sum within each source region to that region's area. In our application the source (coarse) resolution is 900 arc-seconds (about 0.25 degrees, the native LUH2 resolution) and the target (fine) resolution is 10 arc-seconds (about 300 m), giving a 90 by 90 block of target cells per source cell. A coarse projection must be resampled onto this pyramid before downscaling, and only alignment to the pyramid matters, so SEALS is agnostic to a projection's native resolution or source (Supplement Section S2.1).

SEALS represents each land-use class as a binary layer, $\hat{Y}_{ijk} \in \{0,1\}$, recording whether target region i holds class k. A discrete classification (a region is either forest or cropland) requires exactly one class per region:

$$\forall\, i,j \ \exists\, k' \ \ s.t. \ \ \hat{Y}_{ijk'} = 1 \ \ and \ \ \hat{Y}_{ijk} = 0 \ \ \forall\, k \neq k'$$

so exactly one class k' has $\hat{Y}_{ijk'} = 1$. A continuous classification (a region is 65% forest and 35% cropland) relaxes this to $\sum_k \hat{Y}_{ijk} = 1$, the class fractions summing to one. The formulation accommodates both, and here we use the discrete case, which produces maps like the ESA-CCI. Sections 2.2 to 2.4 describe each of the three components that define transition suitability and Section 2.5 combines them.

2.2 Land constraints

Land constraints act as a mask. The operator C takes the ordered pair $\langle k, k_{ij0} \rangle$ of a candidate class k and the region's initial class k_{ij0}, and returns 1 for an allowed transition and 0 for a disallowed one, forcing the suitability of any disallowed transition to zero. For example,

$$P_{ijk'} = 0 \ \ if \ \ (k' = urban \ \& \ k_{ij0} = water)$$

Land constraints are not hard-coded. The user specifies them as input parameters, and can relax a constraint or set weights between 0 and 1 rather than a hard 0 or 1. Unlike the other parameters, land constraints are not empirically calibrated. Because the suitabilities only rank target regions for allocation, rather than form a probability distribution, zeroing some of them does not require rescaling the rest.

The user may also define specific land constraints to explore policy. For example, inputting a layer that rules out expansion in a new protected area would simulate how land-use change would be affected elsewhere. The constraint values used in this application are given in Table S1; all are binary.

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

2.3 Land suitability

Land suitability is the part of P_{ijk} that depends on spatial variables. We collect these in a vector V^n, each element of which may or may not contribute to suitability. We omit the indices i and j for brevity, although each element of V^n in general varies across source and target regions.

SEALS weights each variable in V^n by a parameter $\hat{\beta}^n \in (-\infty,\ +\infty)$, so the suitability contribution is their weighted sum, the dot product

$$\hat{\beta} \cdot V = \sum_n \hat{\beta}^n V^n$$

The suitability response can vary across space: forest may convert to cropland more readily in South America than in Africa at equal covariates. SEALS therefore shares each parameter within regions q (for example continents or agro-ecological zones) larger than the source regions, so that each q holds enough observed change to calibrate against: $\hat{\beta}^n_{ijk} = \hat{\beta}^n_{i'j'k} = \hat{\beta}^n_{qk}$ for all i, j, i', j' in q. In this application each q is one 1-degree calibration tile, and the released coefficient table is indexed by these tiles. The specific spatial variables used in this application are listed in Section 3.4 and Table S1.

2.4 Adjacency effects

Adjacency effects capture how the suitability of any region for transition to any land-use class depends on the land-use classes of surrounding areas. Land-use types follow strong spatial patterns because neighboring types affect each other's probability of transition: an agricultural field surrounded by other fields is unlikely to become urban land, whereas a field at the edge of a city is likely to. SEALS represents these effects as convolutions of the land-use map with distance kernels, which can act at several spatial scales, with a calibrated coefficient for the effect of each neighboring class on each target class.

For all i, j we can assign a distance metric $D(i, j, i', j')$ between the target zone i in source zone j and any target zone i' in any source zone j'. We choose D to be the Euclidean distance between the centroids of i in j and i' in j'.

Then we define the adjacency kernel, evaluated on the grid, as a logistic function of distance:

$$K_{i,j}(\sigma, i', j') = \frac{1}{1 + e^{\,(4/\sigma)(D(i,j,i',j') - \sigma)}}$$

where σ sets the distance at which a neighbor's influence falls to one half (described below). We convolve this kernel with the current land use,

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

$$L'_{ijkt_0} = \left(K(\sigma) * L_{kt_0}\right)_{ij} = \sum_{i'j'} K_{ij}(\sigma, i', j') L_{i'j'kt_0}$$

After this convolution, every target zone carries some information about the adjacent land, weighted by distance through the kernel. In the unconvolved map, by contrast, each target zone records only its own land-use type. The convolution extends across source-region boundaries, which are not treated as edges.

The adjacency convolutions use this unnormalized logistic kernel. A separate Gaussian appears only in the calibration loss of Section 3.1.

Adjacency effects can act at multiple spatial scales, set by the parameter σ. SEALS allows several scales by letting the user choose a set of values, $\sigma_m = \sigma_1, \sigma_2, \ldots$

The adjacency effect is then given by

$$\sum_{l} \sum_{m} \hat{\alpha}_{qlk}^{(m)} \sum_{i'j'} K_{ij}(\sigma_m, i', j') L_{i'j'lt_0}$$

Here l indexes the neighboring land-use class and m the scale. $L_{i'j'lt_0}$ is the baseline indicator of class l at neighboring region (i', j'), and the convolution spreads that class's presence over distance. The coefficient $\hat{\alpha}_{qlk}^{(m)} \in (-\infty, +\infty)$ is the effect of class l, at scale m, on the suitability for class k, and is calibrated within each region q like $\hat{\beta}_{qk}$ above.

Figure 2 shows the kernel itself (Panel a): it shows how a single cell of cropland, or of urban land, raises the cropland suitability of its neighbors, drawn both as a curve against distance and as the corresponding two-dimensional kernel. Panel b shows the result of convolving these kernels across Minnesota, United States of America, which turns the cropland and urban maps into smooth suitability surfaces. Their sum, the third map, is the combined effect of nearby cropland and urban land on cropland suitability.

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

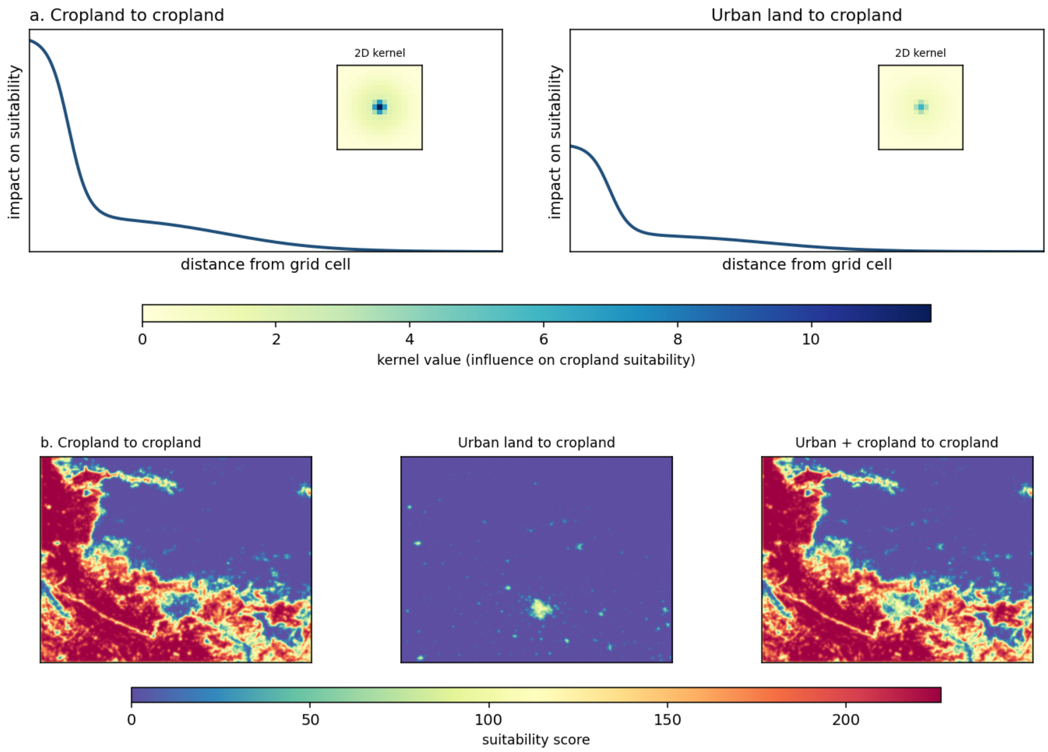

Figure 2. Adjacency effects, computed with the logistic kernel and the calibrated coefficients. (a) The effect of a single grid cell of cropland (*left*) or urban land (*right*) on the cropland suitability of its neighbors, plotted against distance, with the corresponding isotropic two-dimensional kernel inset. The inset shows the kernel as a small image, the source cell at the center and each pixel colored by the kernel weight at that offset (color scale below panel a). The weight fades within a few cells as the logistic kernel decays. Both effects are positive and decay monotonically. Cropland's is the stronger. (b) The same kernels convolved across Minnesota, USA: the cropland-to-cropland surface (*left*), the urban-to-cropland surface (*middle*), and their sum (*right*). Higher values mark locations more suitable for cropland under the adjacency component.

2.5 Overall transition suitability

The three components combine into the overall suitability of a land-use transition:

$$P_{ijk} = C(\langle k, k_{ij0} \rangle) \left(\sum_{l} \sum_{m} \hat{\alpha}_{qlk}^{(m)} \sum_{i'j'} K_{ij}(\sigma_m, i', j') L_{i'j'lt_0} + \sum_{n} \hat{\beta}_{qk}^{n} V_{ij}^{n} \right)$$

The adjacency parameters $\hat{\alpha}_{qlk}$ and the suitability parameters $\hat{\beta}_{qk}$ are empirically calibrated, as described in Section 3. The constraint mask C is set by the user and is not calibrated. This

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

application uses no other multiplicative factors, so C is the only term multiplying the additive sum.

2.6 Allocating changes in land use

Given the suitabilities P_{ijk}, SEALS allocates a coarse projected change Z_{jk} to the fine target regions i within each source region j. In each source region it ranks the target regions by their suitability P_{ijk} for each land-use class k and fills the most suitable regions first, allocating across classes as described below until each class's coarse total Z_{jk} is met.

SEALS alternates between the changing classes. It assigns a single fine cell to each class in turn, always that class's most suitable still-available cell, and repeats this cycle until every class has met its source total Z_{jk} (Johnson et al., 2021b). Allocating one class in full before the next would let whichever class is taken first claim the cells a later class would also occupy. Within a cycle the classes are taken in a fixed index order, but because only one cell is assigned to each class per cycle, no class systematically pre-empts another and the result does not depend on an arbitrary class ordering. Other land-use models fix an allocation hierarchy that resolves competition by fiat, typically urban, then cropland, then pasture (Section 5.3).

Because the allocation ranks target regions within one source region at a time, and because the coarse totals are met within that source region, each source region can be processed independently of every other. Section 3.6 describes how the implementation exploits this.

3 Parameter estimation, validation design, and application

The model of Section 2 is parameterized by the variables C, $\hat{\alpha}_{qlk}$, and $\hat{\beta}_{qk}{}^{n}$, whose values are to be determined. Although sensible parameters may be chosen based on expert opinion (Suh et al., 2020), the SEALS model follows an empirical approach. Here, we describe an approach to estimate parameters in SEALS that is achieved by (a) defining a similarity metric (Section 3.1) for how closely two land-use maps agree, and (b) estimating the parameters using a greedy, coordinate-wise search, a hill-climber that optimizes a change-weighted form of this metric (Section 3.2). Section 3.3 sets out how the calibrated placement is validated on a withheld period, Section 3.4 the settings of the application, Section 3.5 the ecosystem-service test, and Section 3.6 the computational implementation.

3.1 Loss function

We score a parameter set by how closely its predicted change matches the observed change. For each changing class we split the change into a baseline-relative expansion (cells that gain the

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

class) and a contraction (cells that lose it), convolve each of the resulting maps with a Gaussian kernel, and sum the absolute differences of the expansion and contraction parts separately:

$$S = \sum_{i,j,k} \left(\left| \left(G(\sigma_S) * \Delta\hat{Y}_k^+\right)_{ij} - (G(\sigma_S) * \Delta Y_k^+)_{ij} \right| + \left| \left(G(\sigma_S) * \Delta\hat{Y}_k^-\right)_{ij} - (G(\sigma_S) * \Delta Y_k^-)_{ij} \right| \right)$$

where $G(\sigma_S) *$ is a Gaussian blur with standard deviation σ_S (a separate parameter from the adjacency scales σ_m), a Gaussian convolution distinct from the logistic adjacency kernel K of Section 2.4. $\Delta\hat{Y}_k^+$ and $\Delta\hat{Y}_k^-$ are the predicted expansion and contraction maps of class k, and $\Delta Y_k^+, \Delta Y_k^-$ the observed ones. The sum runs over all target cells i, j and the changing classes k (the five classes that change in this application). The blur standard deviation is $\sigma_S = 7$ fine cells. Blurring before differencing rewards agreement between nearby locations even where pixels do not exactly coincide, and keeping expansion and contraction separate stops a gain and a loss in the same neighborhood from cancelling.

Calibration maximizes a change-weighted ratio, the total projected change (the number of fine cells that change class) divided by $S + 1$, which rewards a parameter set for placing a large amount of change in close agreement with the observations. The amount of change is, however, set by the coarse scenario, subject to the allocation limitations in Section 2.1, so the calibrated parameters control only where change is placed, not how much. Maximizing the ratio is therefore equivalent to minimizing the placement error, except in the rare coarse cell whose required change exceeds its eligible land, where slightly less can be placed. Supplement Figures S1 and S2 make this concrete on a sample region in Mato Grosso, Brazil. Figure S1 runs SEALS forward from the 2000 baseline under the observed coarse change and sets the projected 2015 map beside the observed one. The blurred difference between them is the similarity S that calibration drives down, and it is largest where the projection placed change in the wrong location. Figure S2 breaks the same comparison out by land-use class, showing that the cropland and forest transitions, which dominate change in this region, account for most of the similarity, while the classes that barely change contribute little.

3.2 Calibration

To calibrate the model, we use a greedy coordinate search, a step-wise, coordinate-wise hill-climber, that does the following:

(a) sets the parameters at an initial starting value (for example, those in Table S1);

(b) runs the SEALS model on an initial land-use map, L_{ijkt_0}, and a coarse-grained change in land-use (*i.e.*, Z_{jk}) calculated from an observed land-use map from a later year;

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

(c) uses the loss function to find how similar the predictions are to the observed land-use map, and

(d) iteratively updates the parameters based on the value of the loss function, until no further improvements can be found.

Once we have found the correctly calibrated parameter set, we can use these parameters in the equation for P_{ijk} given in Section 2.5. Then, for any change in land-use Z_{jk}, we can arrive at our desired high-resolution land-use map Y_{ijk} by again allocating the changes in land-use using the method described in Section 2.6.

We calibrate the parameters separately for each 1-degree tile on Earth (the regions q of Section 2.3), retaining only the 17,360 tiles in which the LUH2 scenarios project some future change, since coefficients are never used in tiles where no change is allocated (for instance, oceanic tiles). The coordinate-wise search converges to a local optimum of the loss and reports a single parameter vector per tile with no interval around it; Supplement Section S5.3 discusses the consequences.

3.3 Validation design and metrics

The calibration of Section 3.2 fits the parameters to the observed 2000-to-2015 change, so a close in-sample similarity (Supplement Figures S1 and S2) shows only that the model can reproduce the change it was trained on. To test whether the learned placement transfers to a period it never saw, we hold the fitted coefficients fixed and apply them to a later, held-out interval.

We take the coarse 2015-to-2020 change, aggregated from the observed ESA maps to the 0.25 degree cells, and downscale it from the 2015 map using the frozen 2000-to-2015 coefficients, giving a projected 2020 map, which we compare to the observed 2020 map. Because the coarse cell totals are supplied to the model, the test measures only where it places a given amount of change, not whether it predicts the amount.

We score placement with the figure of merit (FoM; Pontius et al., 2008): over the pixels that change class in either the observed or the projected map, the fraction that change in both and to the same class. The FoM is 0 if no projected change coincides with an observed change and 1 if they coincide exactly.

We interpret the FoM against a null that places the same change the model placed, but at random: we take the model's own projected change in each coarse cell and scatter it to random pixels within that cell, so the only difference from SEALS is suitability-based versus random location. We benchmark against this random placement rather than against a no-change (persistence) prediction, whose figure of merit is zero by construction over the union of changed

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

pixels and so cannot serve as a skill floor (Pontius et al., 2008). The fair random placement, handed the same change and differing from SEALS only in where it falls, is the appropriate null.

The figure of merit credits only exact-pixel hits, penalising a change predicted one cell from where it occurred as heavily as one predicted far away. As a complementary, proximity-aware test we relax the hit criterion: a predicted change counts as a hit when an observed change of the same class lies within a radius of r pixels of it, scored over the union of changed pixels, which reduces to the exact figure of merit at $r = 0$. We report the ratio of the SEALS score to the random-placement score as a function of r; a ratio that stays above one as r grows shows that the placement skill is not an artifact of requiring exact-pixel hits.

Both metrics are computed globally over all land pixels. The placement test is complemented by the ecosystem-service test of Section 3.5, which asks whether the placement skill measured here carries through to a quantity that depends on the spatial configuration of land use.

3.4 Application to the SSP-RCP scenarios

As a demonstration of the calibrated model, we apply SEALS to downscale the coarse, global land-use change projections of Hurtt et al. (2020) to a 10 arc-second global grid, for the following combinations of Shared Socioeconomic Pathways (SSP) and Representative Concentration Pathways (RCP) (hereafter, "SSP-RCP scenarios"): SSP 1, RCP 1.9; SSP 1, RCP 2.6; SSP 2, RCP 4.5; SSP 3, RCP 7.0; SSP 4, RCP 3.4; SSP 4, RCP 6.0; SSP 5, RCP 3.4; and SSP 5, RCP 8.5. The SSP-RCP scenarios are standardized pathways exploring how socioeconomic, environmental, and technological change might occur in the coming decades.

We consider 7 discrete land-use classes, aggregated from the more detailed land-cover classification of the European Space Agency Climate Change Initiative (ESA-CCI): urban, cropland, grassland, forest, other natural, water, and barren and other. The coarse source regions j are the 1.036 million 0.25 degree (~30 kilometer) grid cells of the global land-use projection for each SSP-RCP scenario, and the fine target regions i are the 10 arc-second pixels within them. Using the calibration algorithm of Section 3.2, a unique set of parameters is estimated for each 1-degree (~110 km) tile on Earth. The adjacency effects act at two scales, σm = 1 and 5 fine cells, a short-range and a long-range effect.

Beyond the land-use maps, the land-suitability predictors for this application are sand content (%), silt content (%), soil bulk density, soil cation exchange capacity, soil organic content, strict protected area, air temperature (°C), travel time to market (min), presence of wetlands, altitude (meters), above-ground carbon (Mg ha^{-1}), clay content (%), soil pH, population, and precipitation (mm) (see Table S1 for all data sources and citations).

For calibration, the baseline is the 2000 ESA map and the parameters are fit to the observed change to 2015, while the 2020 map is held out to assess performance (Section 3.3). Once the optimized parameters are found, the SSP-RCP projections, expressed as changes from 2015, are applied to the 2015 ESA map for 2030, 2050, 2070 and 2100.

With 43 spatial layers, 5 changing land-use classes, and 17,360 valid tiles, the calibration estimates 3.12 million coefficients across the 36 non-constraint layers. With the 7 fixed land constraints the released file holds 3.73 million in total, listed by type in Table S1 and released in full in the accompanying data repository.

The gridded downscaling described and evaluated in this paper is one of two input modes SEALS supports. The model also accepts regional-polygon inputs (for example country- or agro-ecological-zone totals), the mode used in the GTAP-InVEST applications (Johnson et al., 2023, 2025); because a polygon does not contain an integer block of fine cells, its conservation and validation differ from the gridded case, and we leave a formal treatment to future work.

3.5 Ecosystem-service estimation: habitat quality

To show how SEALS supports ecosystem-service estimation at fine resolution, and whether resolving the placement improves the estimate, we compute habitat quality with the InVEST model (Sharp et al., 2020). InVEST assigns each pixel a quality between 0 and 1 from its habitat suitability and the degradation it suffers from nearby threats, whose influence decays with distance. We treat forest and other natural land as habitat and cropland and urban land as threats, with distance-decay parameters following the InVEST User Guide's forest-cropland example where available and set to plausible values otherwise, with robustness to that choice tested in Supplement Section S3. The reference is InVEST habitat quality calculated from observed land cover. For the held-out 2015-to-2020 period we compute habitat quality on four maps, the observed 2020 map, the SEALS-downscaled map, a coarse map that spreads each cell's conversion uniformly within it, and a fine-resolution map that places the same per-cell change as SEALS at random, and compare the SEALS, coarse, and random-fine estimates to the observed.

3.6 Implementation and computational design

Rank-and-allocate downscaling requires ordering the candidate cells within each conservation unit by suitability. The size of that unit, not the fine resolution, determines whether the problem partitions. Because SEALS allocates within each coarse source region and never moves change between source regions (Section 2.1), the ranking is local by construction: in the application reported here a source region contains a 90 by 90 block of 8,100 target regions, and the working set for the allocation is one such block together with the convolved adjacency layers that overlap it. Each block can therefore be processed and written independently, and discarded once written, so memory use is governed by the block size and the number of blocks processed concurrently rather than by the size of the global grid.

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

SEALS is implemented in Python with the allocation loop compiled (Cython), and it processes source-region blocks in parallel across the available CPU cores. The adjacency convolutions are evaluated across source-region and tile boundaries rather than treating them as edges. The same partitioning applies to the calibration, which runs independently for each 1-degree tile. Section 4.4 reports the resulting runtime, memory use, and scaling with the number of worker processes, and Supplement Section S4.5 contrasts this design with the scaling strategies of related models.

4 Results

4.1 Held-out validation of placement skill

We first ask whether the placement learned from 2000-to-2015 change transfers to a period the calibration never saw. Figure 3 illustrates the test for a sample region of Kalimantan, an active forest-to-cropland frontier, setting the in-sample projection of 2015 beside the held-out projection of 2020. The projected map resembles the observed one in the held-out year much as in the training year; the quantitative result is the global figure of merit below, not the visual match in this one window.

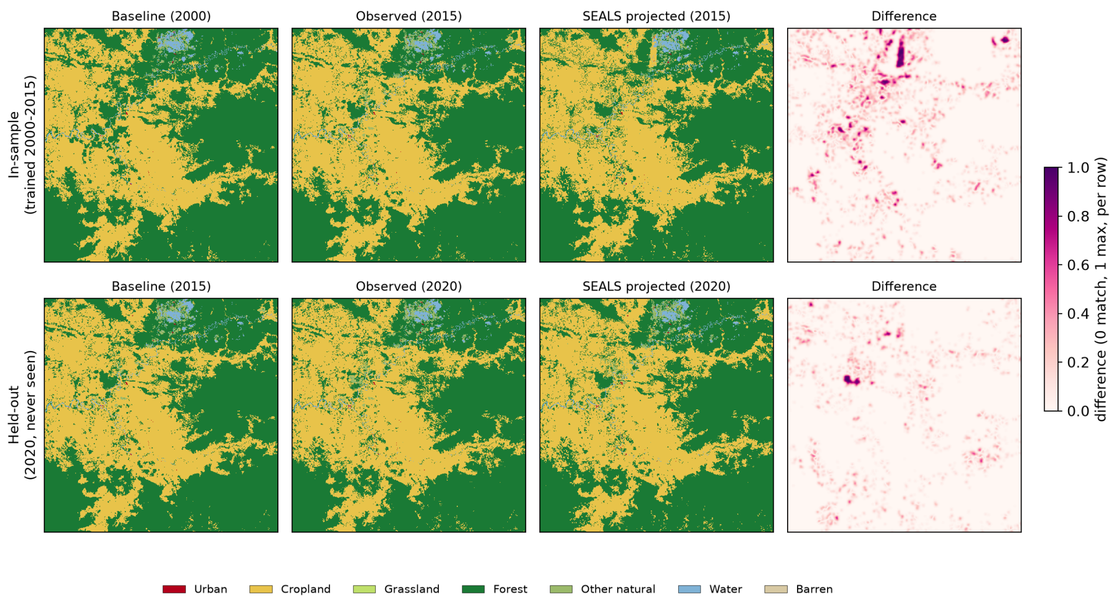

Figure 3. In-sample versus held-out placement, for a sample region of Kalimantan (Borneo; 111 to 113°E, 1°S to 1°N), an active forest-to-cropland frontier. *Top row, in-sample:* the 2000 baseline, the observed 2015 map, and the SEALS-projected 2015 map (run forward from 2000 under the observed coarse change), the period the parameters were calibrated on. *Bottom row, held-out:* the 2015 baseline, the observed 2020 map, and the SEALS-projected 2020 map (run forward from

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

2015 under the observed 2015-to-2020 coarse change), a year the calibration never saw. The region loses forest to cropland in both periods (about 5% of pixels change in-sample and 2% held-out), and the projected map resembles the observed one in the held-out year much as in the training year. This panel is illustrative. The quantitative out-of-sample skill is the global figure of merit in the text, not the visual match in this one window. The difference panels (*right*) show the blurred per-location disagreement between projected and observed change, summed over the five changing classes. Each is normalized to its own row, so they mark where the two differ within a period and are not comparable in magnitude between the two periods.

Globally, SEALS scores a figure of merit of 0.0405 on the held-out 2015-to-2020 change against 0.0215 for the random placement of the same change within each coarse cell, a ratio of 1.88, and the ratio is unchanged across random seeds (Section 3.3). SEALS therefore places held-out change about 1.9 times more accurately than a random allocation of the same change, on a metric that credits only exact-pixel agreement.

Two limits hold the absolute FoM low. First, over a five-year interval much of the ESA change is small or reflects classification noise, so exact-pixel agreement is hard for any method. Second, the downscaling reproduces only the net change each coarse cell records. Offsetting gains and losses within a cell cancel, so SEALS places only about half the gross change present in the observed map and cannot match change it never receives, however well it is positioned. The result is therefore best read as a relative one: the calibrated placement is about 1.9 times more accurate than chance out of sample, on a metric whose absolute ceiling is well below one for any downscaler that receives only net coarse change.

Because the figure of merit penalises a change predicted one cell from where it occurred as heavily as one predicted far away, we also score the placement with the proximity-aware metric of Section 3.3 (Figure 4). SEALS is more accurate than the random placement at every tested tolerance radius: by a ratio of 1.88 at radius zero, where the measure equals the figure of merit, falling from 1.71 at a one-pixel tolerance to 1.56 at a three-pixel tolerance and 1.39 at a ten-pixel tolerance. The advantage persists when near-misses are credited, so the placement skill is not an artifact of requiring exact-pixel hits.

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

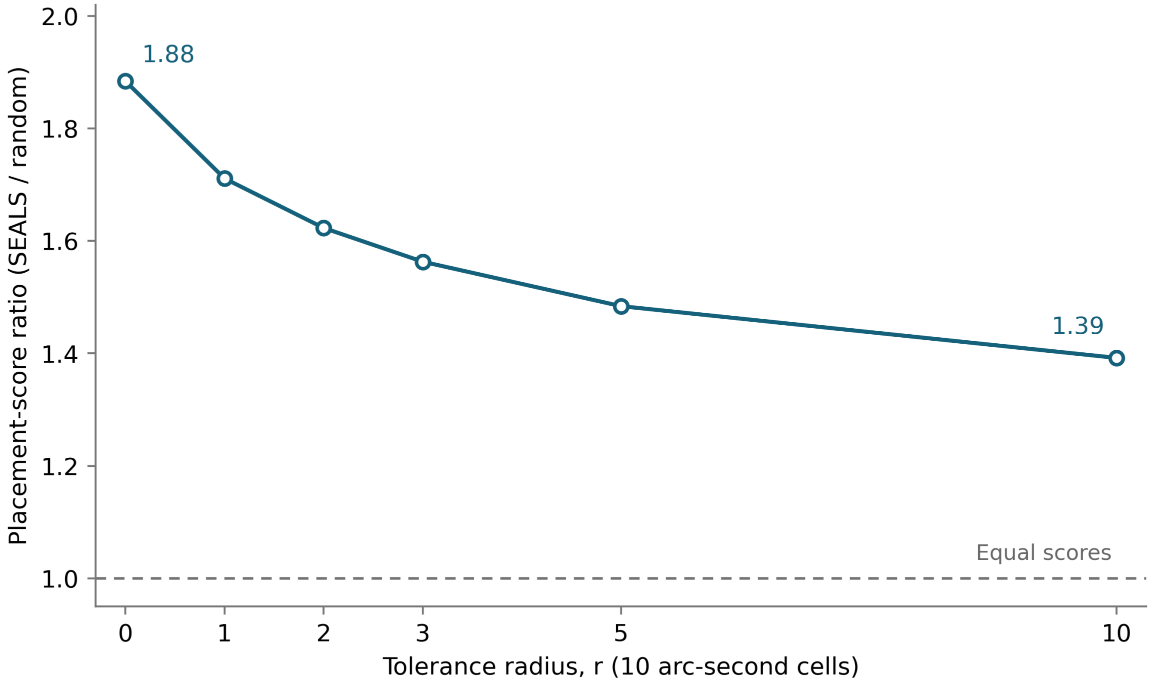

Figure 4. Proximity-aware placement skill on the held-out 2015-to-2020 land-cover change. The ratio of the SEALS placement score to the score of a random placement of the same change within each coarse cell, evaluated at tolerance radii of 0, 1, 2, 3, 5 and 10 fine cells (10 arc-seconds per cell). A predicted change is counted as a hit when an observed change to the same class lies within r horizontal or vertical cell steps. Each score is divided by its changed-pixel union (Section 3.3), and at radius zero equals the figure of merit. Points mark evaluated radii; connecting lines are guides. The dashed line indicates equal scores. SEALS scores higher than random at every tested radius.

4.2 Habitat quality and the value of resolution

We test whether resolving the location of conversion improves a habitat-quality estimate during the held-out period. We illustrate this using the Habitat Quality model in the InVEST suite of ecosystem service models (Sharp et al., 2020). This model estimates habitat degradation based on the proximity of threats such as cropland and urban land, and is widely used as an indicator of biodiversity condition (von Jeetze et al., 2023). Resolving where conversion falls therefore changes the habitat-quality estimate, and because real deforestation concentrates along frontiers, the fine-resolution placement should match observations more closely than a coarse one.

We test this prediction on the held-out 2020 maps, computing InVEST habitat quality for the SEALS, coarse-uniform, and random-fine maps and comparing each to the observed 2020 map (Section 3.5). We focus on ten deforestation regions, because that is where conversion is large and concentrated enough for its placement to affect the habitat-quality estimate. Across these regions the SEALS estimate is closer to the observed than the coarse estimate in nine, and closer than the random-fine estimate in the same nine, by large margins (Figure 5; Supplement Table S2). Spreading conversion uniformly places a threat near almost every remaining habitat pixel

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

and over-degrades the landscape. The random-fine placement does the same despite its fine resolution, with a habitat-quality error that tracks the coarse map's in every region, whereas SEALS concentrates conversion where deforestation actually occurred and reproduces habitat quality calculated from observed land cover. The advantage holds in-sample as well as out of sample, across threat distances from 2 to 15 km, and across a range of InVEST habitat-sensitivity and threat-weight parameters.

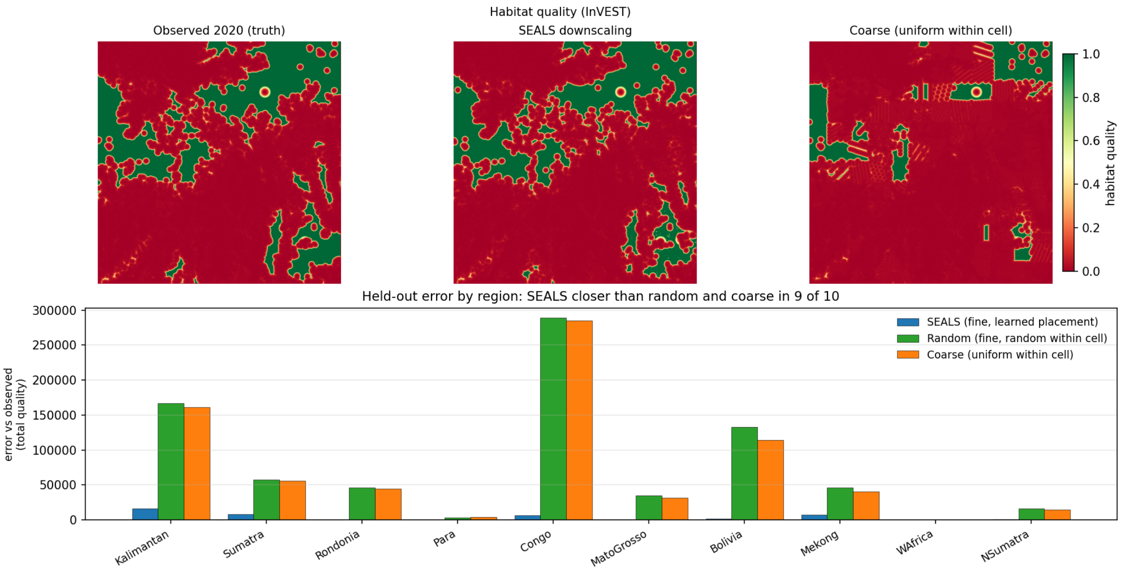

Figure 5. Habitat quality (InVEST) on the held-out 2020 maps. Top: the habitat-quality map over a Kalimantan window, computed on the observed map, the SEALS-downscaled map, and a coarse map that spreads each cell's conversion uniformly within it. The SEALS map tracks the observed one, while the coarse map over-degrades habitat by placing threats throughout each cell. Bottom: the error of each estimate relative to the observed, across ten deforestation regions, for the coarse map, a fine-resolution map carrying the same per-cell change placed at random within each cell, and the SEALS map. The SEALS estimate is closer than both the coarse and the random-fine estimate in nine of the ten regions, and the random-fine estimate tracks the coarse one.

For this habitat-quality model and these regions, calibrated placement improves agreement with estimates derived from observed land cover relative to both uniform and random placement.

4.3 Demonstration: a 10 arc-second land-use dataset for eight SSP-RCP scenarios

Having evaluated the calibrated model, we apply it to the eight SSP-RCP scenarios of Section 3.4 to produce high-resolution land-use maps for 2030, 2050, 2070 and 2100 (Figure 6 for global maps, Figure 7 for a regional illustration). The maps resolve the position of the change within each coarse cell. SEALS concentrates cropland expansion, and the matching loss of forest and other natural land, along the frontiers implied by the spatial patterns it learned from observed change, rather than spreading it evenly across the cell (Figure 6). The regional example shows that as cropland advances across Central Africa under SSP3 RCP7.0, forest cover in the window

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

falls from 52% in 2015 to 7% by 2100, concentrated where the model places conversion (Figure 7). We do not analyse the scenarios further here; the runs demonstrate the calibrated model on the coarse inputs it is designed for, and the full dataset is publicly available (see “Data availability”) as the first comprehensive set of land-use maps for the SSP-RCP scenarios at a resolution suitable for ecosystem-service models such as InVEST.

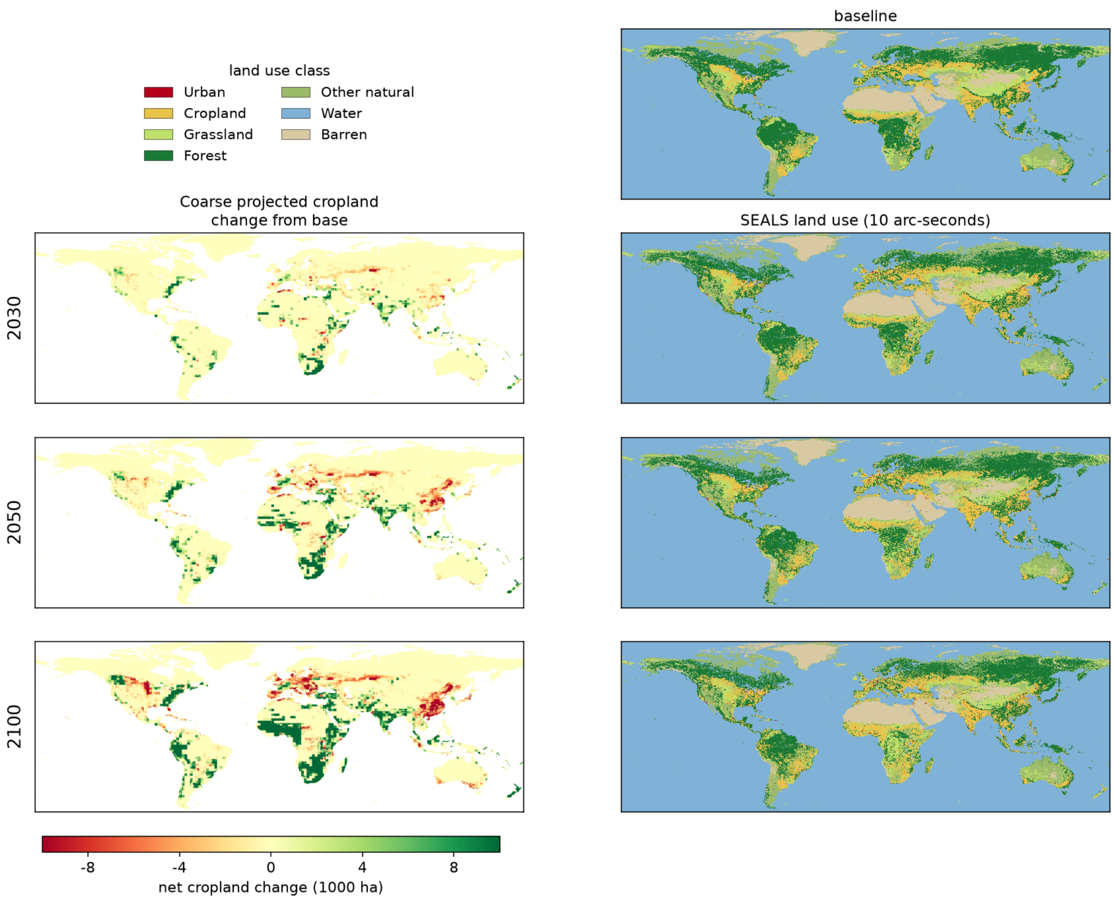

Figure 6. Global coarse-versus-fine land-use change under SSP3 RCP7.0. Top right: the 2015 ESA land-use baseline. For 2030, 2050, and 2100 (rows): the coarse projected cropland net change from 2015 (0.25°, left) and the fine SEALS cropland fraction change (10 arc-seconds, aggregated for display, right). SEALS concentrates the same coarse change into specific fine-scale locations.

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

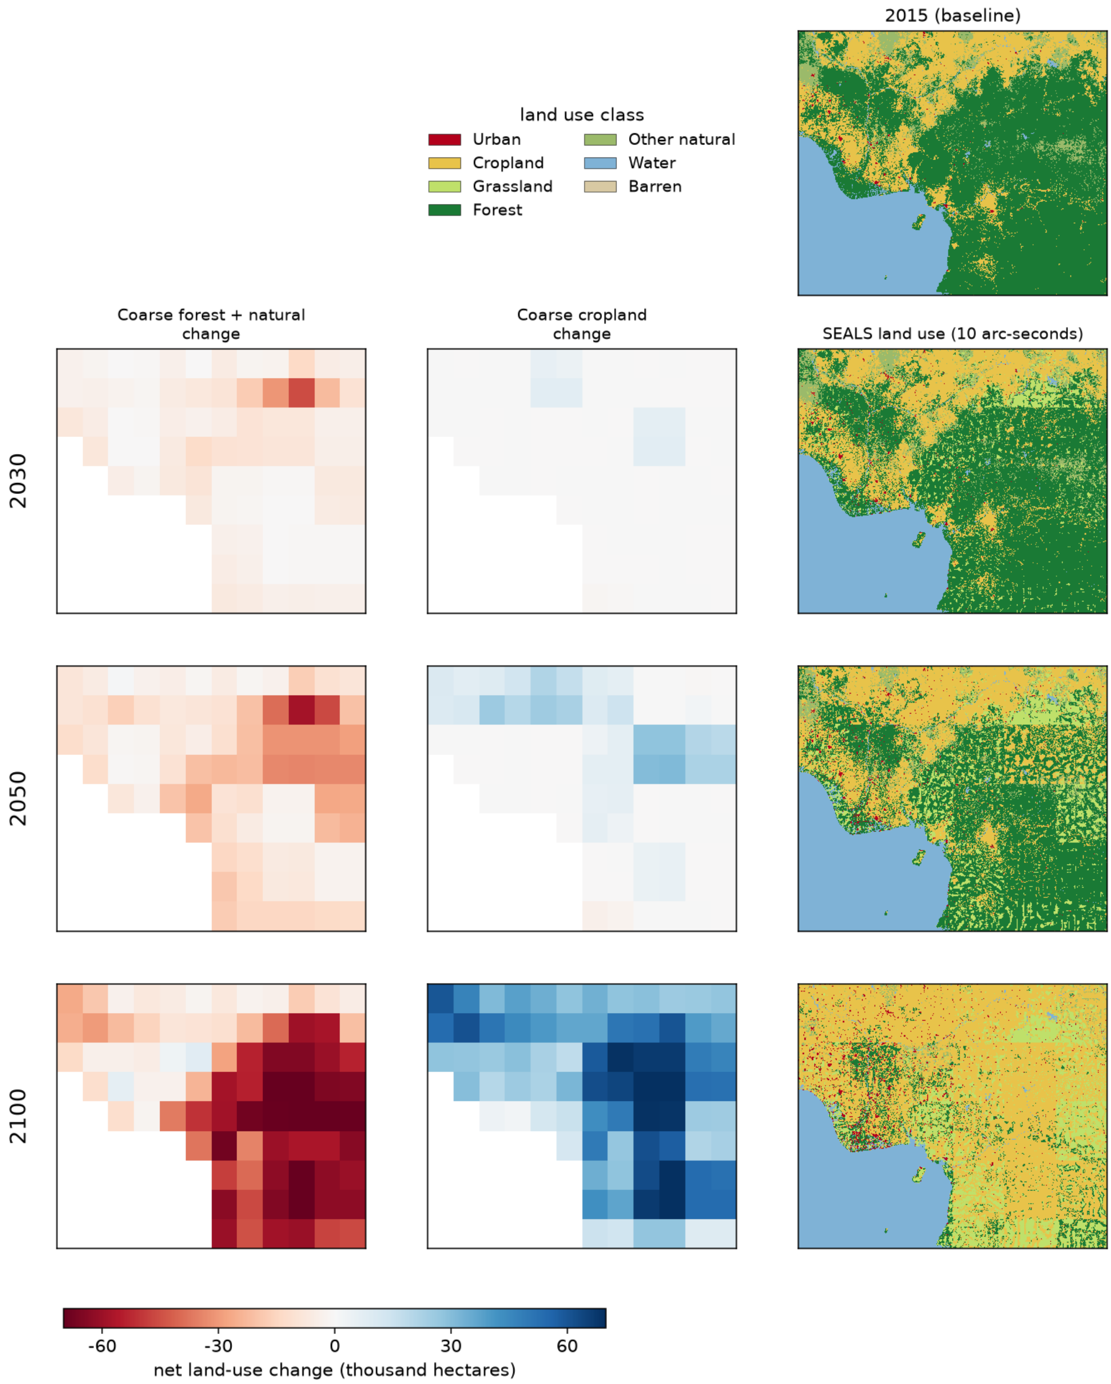

525 *Figure 7. Downscaling over time for a sample region of Central Africa (4-16°E, 1-10°N) under SSP3 RCP7.0. Top right: the observed 2015 ESA land-use baseline at 10 arc-seconds, above the SEALS column. For 2030, 2050, and 2100 (rows): the coarse LUH2 input change from 2015 in forest and other natural land and in cropland (net hectares, 0.25°, shared scale), and the SEALS 10 arc-second land use. SEALS allocates the coarse input change to fine-scale locations. Forest cover in the region falls from 52% in 2015 to 7% by 2100.*

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

4.4 Computational performance

Table 1 reports computational performance for global 10 arc-second downscaling. System A is an Apple M5 Max (Arm64), with an 18-core CPU. System B is MSI node cn1079, with two 64-core AMD EPYC 7002-series "Rome" processors, 128 physical cores, SMT disabled and approximately 503 GiB of node memory; 127 worker processes were used for exclusive-node execution. The time for one global scenario-year is approximately 8 minutes on System A and 1.5 minutes on System B; the speedup is approximately fivefold. The global calibration time is 120 hours for System A and 19 hours for System B. System A performed better per zone due to higher clock speeds but was slower overall due to lower core count.

Table 1. *Computational performance of the global 10 arc-second downscaling and calibration. 8.4 billion grid-cells.*

Quantity	System A	System B
CPU (architecture, model)	M5 Arm64 CPU	2 × 64-core AMD EPYC 7002-series "Rome" (cn1079)
Physical cores / hardware threads	18 cores (6 super + 12 performance) / 18 threads	128 physical cores / 128 threads (SMT disabled)
Memory	128 GiB	503 GiB
Worker processes used	18	127
Downscaling, one scenario-year, wall time (min)	8	1.5
Downscaling, peak resident memory (GB)	55 GiB	N/A
Processing time per zone, median	11.2 s per zone	14.2 s per zone
Global calibration, 17,360 tiles, wall time (h)	120	19
Calibration, peak resident memory (GB)	55 GiB	N/A

Because each source region is processed independently (Section 3.6), the runtime scales with the number of worker processes and the memory footprint scales with the number of blocks resident at once rather than with the size of the global grid. The same partitioning is what allows the model to run on low-memory machines and to scale to high-performance computing with thousands of CPUs, where 1/3 arc-second or finer allocation becomes feasible. Supplement Section S4.5 explains why this scaling follows from the allocation design rather than from the implementation language.

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

5 Discussion

5.1 What resolving placement provides

SEALS downscales coarse, global projections of land-use change to a 10 arc-second grid, using estimated relationships between the characteristics of land and its suitability for transition. The model is open source and replicable, and the user sets the land-use classification (recalibrating the parameters for it), the spatial predictors, the output resolution, and the spatial scales of the adjacency effects. Its parameters can be specified or estimated from observed change.

The held-out validation shows that the placement SEALS learns from observed change transfers to a period it never saw, and the habitat-quality test shows that this placement skill carries through to a downstream estimate. The ecosystem-service effects of land-use change depend on where within a coarse cell that change falls, which SEALS resolves and a coarse projection does not, with placement skill validated out of sample. For habitat quality, a service whose impact depends on that placement, the fine-resolution maps yield a more accurate estimate than coarse accounting, validated across ten deforestation regions (Section 4.2). Integrated assessment models report regional areas of cropland, pasture, and forest (Popp et al., 2017) but not their location. The same point extends to biodiversity. Habitat quality reflects the spatial configuration of land use, including the proximity and fragmentation of the remaining habitat, so resolving where change falls bears on biodiversity change as well as ecosystem services.

Even as it stands, the value of these maps for assessing land-use policy is their spatial resolution: the local consequences of land-use futures, for ecosystems and for the people exposed to them, depend on where change falls, not only on how much. Coarse projections report the amount but not the location. SEALS supplies the location, with that skill confirmed out of sample.

That location is immediately usable: the maps produced by SEALS are ready for spatially explicit impact models. Tools such as InVEST (Sharp et al., 2020) can read them directly, and an intercomparison of ecosystem service models found many of them constrained by the lack of high-resolution land-use input (Kim et al., 2018). Air quality models share that need: InMAP resolves concentrations to cells as small as 4 km (Thakrar et al., 2022), and health impacts depend on local exposure (Thakrar et al., 2024). Earlier high-resolution analyses of the interaction between people and nature relied on this kind of downscaling (Chaplin-Kramer et al., 2019; Johnson et al., 2021a; Suh et al., 2020). The local effects of the global SSP-RCP scenarios depend on the fine-scale location of land-use change, which SEALS resolves.

5.2 SEALS in the modelling pipeline

Although this paper focuses on describing and evaluating the coarse-to-fine downscaling step, SEALS sits within a larger modelling pipeline that carries regional or gridded projections through

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

to fine-scale scenario outputs (Figure 8). Automatic combination of the two downscaling steps enhances the applicability of SEALS to multiple modelling contexts.

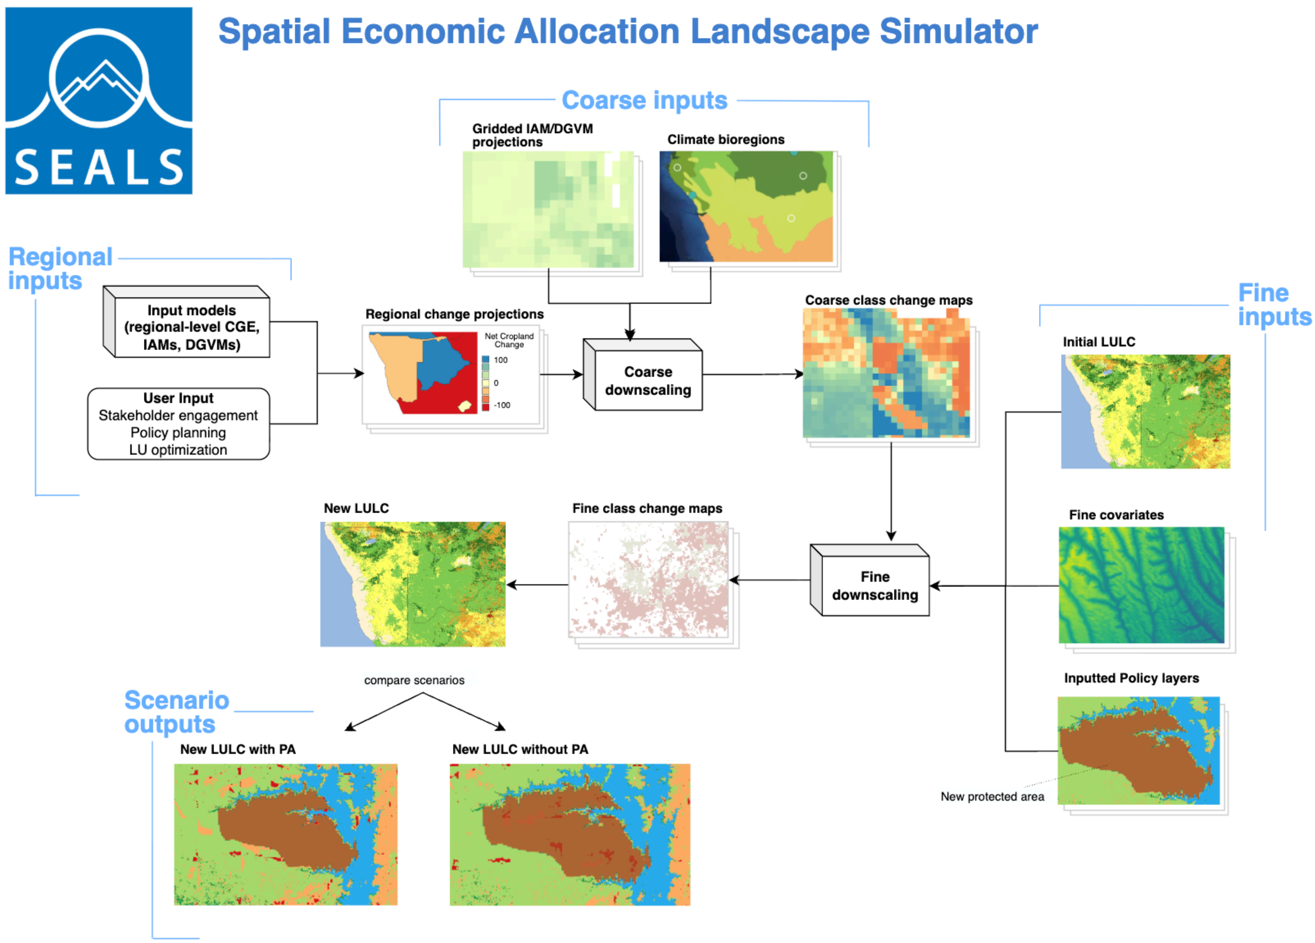

Figure 8. SEALS within the broader earth-economy modelling pipeline. Regional inputs (regional-level CGE, IAM, or DGVM projections, together with stakeholder input, policy planning, and land-use optimization) are first translated into regional change projections and downscaled to coarse gridded class-change maps; the fine-downscaling step then allocates that coarse change to a fine land-use map using the initial land use, fine covariates, and any imposed policy layers, and the resulting maps feed scenario comparisons (for example, with and against a new protected area). This paper documents and validates the fine-downscaling step [boxed]; the regional-to-coarse front end is used in the GTAP-InVEST applications (Johnson et al., 2023, 2025) and is not evaluated here. Section 5.3 compares SEALS with related models at this fine-downscaling step.

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

5.3 Relation to other downscaling models

SEALS occupies the same position in the modelling chain as several other downscaling models, among them LandScaleR (Woodman et al., 2023), Demeter (Vernon et al., 2018; Chen et al., 2019), Dinamica EGO (Soares-Filho et al., 2002), and the econometric downscaler downscalr (Krisztin et al., 2023): each takes a coarse projection of land-use change as given and redistributes it within a coarse accounting unit, conserving that unit's totals. What separates them is where the allocation rules come from, what spatial information those rules use, how competition between expanding classes is resolved, and how far the implementation scales. Demeter's behaviour is set by a small number of user-specified rules, six of which Chen et al. (2019) later calibrated as global scalars; LandScaleR exposes essentially one stochasticity parameter, tuned so that the simulated landscape reproduces the observed degree of spatial aggregation, and deliberately uses no covariates; Dinamica EGO estimates transition probabilities by weights of evidence from observed change but assumes conditionally independent covariates and pre-selects candidate cells, both of which bias the realised transitions (Mazy and Longaretti, 2022); downscalr estimates a Bayesian multinomial logit and satisfies the coarse totals by adjusting probabilities rather than by selecting cells, so its output is a fractional allocation rather than a categorical map and it carries no adjacency term. All four resolve competition among expanding classes either by an asserted ordering or by proportional mixing, and LandScaleR redistributes change between neighbouring coarse cells where a cell cannot absorb its assigned change, whereas SEALS treats the coarse projection as authoritative.

Each of these designs has advantages SEALS lacks: Demeter's small parameter set is legible and easy to override deliberately, LandScaleR is the least data-hungry and the only one tuned directly against landscape configuration, Dinamica EGO is the most general modelling environment and the only one whose runtime was engineered for out-of-core execution, and downscalr carries a posterior on its parameters. What SEALS combines, and these do not, is spatially varying parameters estimated against observed change, multi-scale adjacency, allocation within the coarse source region, out-of-sample validation of the placed class-changes, and parallel execution at global extent and 10 arc-second resolution. Accounting for change within source regions also makes SEALS applicable when the coarse input is a set of regional and agro-ecological-zone totals rather than a grid, as when it downscales GTAP-AEZ, and it is the source of the tiling artefacts discussed below. Supplement Section S4 gives the model-by-model comparison, including a feature-by-feature summary (Table S3) and the computational-scaling argument, and Supplement Section S6.5 proposes the downscaler intercomparison that would put these models on a common footing.

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

5.4 Limitations and extensions

Several limitations bound the results in this paper. The habitat-quality estimate uses a single InVEST parameterization, so its absolute values are illustrative rather than definitive. The SEALS-versus-coarse comparison that underlies our conclusion, by contrast, holds across the parameter ranges we tested (Supplement Section S3). More substantially, SEALS assumes (by design) that the coarse projections are completely correct, which can result in "tiling artifacts" where the boundaries of the coarse data can be seen even in the high-resolution data. This is evident in Figure 7 in the year 2100, where the coarse projections of change convert nearly the entire coarse grid-cell to agriculture. SEALS uses the coarse totals as allocation targets without redistributing demand between source regions, subject to the allocation limitations in Section 2.1. An alternative approach would instead address this by updating the coarse inputs explicitly, for example through a Bayesian hierarchical framework that lets the model adjust the coarse inputs and carry uncertainty through to the maps. Supplement Section S5 sets out the limitations of the model itself, organised by the stage of the method they affect, including the greedy sequential allocation, the local calibration search whose uncertainty is not propagated, the loss function's emphasis on placement rather than configuration, and the evaluation of adjacency on the map at the start of each step.

The model invites several extensions, developed in Supplement Section S6. It reports a single projected map and could instead report uncertainty bounds or a distribution of outcomes. The classification could be widened beyond the seven classes to resolve transitions that matter for other services, and the habitat-quality demonstration carries over to other ecosystem services such as carbon storage, water yield, and air quality once the relevant response maps are attached. SEALS computes its adjacency effects as spatial convolutions with fixed, calibrated kernels. A natural extension is to replace these with learned convolutions, as in a convolutional neural network, which could let the model represent richer spatial relationships. SEALS also accepts gridded or regional-polygon inputs, which has let it downscale the GTAP-AEZ economic model across multiple regions and agro-ecological zones (Baldos & Corong, 2020; Lee et al., 2009; Johnson et al., 2021a).

6 Conclusions

This paper has described and evaluated version 2.0 of SEALS, a model that downscales coarse projections of land-use change to the resolution at which ecosystem-service models operate. Four results summarise the contribution.

First, SEALS combines, at global extent and 10 arc-second resolution, properties that existing downscalers offer only separately: allocation parameters estimated against observed land-use change and allowed to vary across 1-degree tiles, adjacency effects represented as calibrated convolutions acting at several distances, allocation against coarse totals within each source region, and a partitioned implementation that runs on a laptop and scales to high-performance computing. The full formulation, calibration procedure, and validation design are given in Sections 2 and 3.

Second, the placement the model learns transfers out of sample. On withheld 2015-to-2020 change, SEALS places change 1.9 times more accurately than a random allocation of the same change on an exact-pixel figure of merit, and the advantage persists, at ratios of 1.7 to 1.4, when near-misses within one to ten cells are credited.

Third, that placement skill matters for the quantities ecosystem-service models compute. For habitat quality, a service that depends on the proximity of conversion to remaining habitat, the SEALS map reproduces the estimate calculated from observed land cover more closely than coarse accounting, and than a fine-resolution map that places the same change at random, in nine of ten deforestation regions, and the result is robust to the InVEST parameterization.

Fourth, the calibrated model, its 3.73 million released parameters, and its 10 arc-second downscaling of eight SSP-RCP scenarios to 2100 are openly available, so that scenario-based ecosystem-service and biodiversity assessments can be carried out at the resolution their models require without each user recalibrating a downscaler.

The principal limitation is that SEALS takes the coarse projection as given and does not adjust it in response to fine-scale evidence, which is what makes it applicable to regional as well as gridded inputs but also inherits any error in the coarse totals. Relaxing that assumption, propagating parameter uncertainty into the maps, and an intercomparison of downscaling models on common inputs are the extensions we regard as most valuable.

Code availability

The current version of SEALS is available from the project website https://github.com/jandrewjohnson/seals under the Apache 2.0 License. The exact version of the model used to produce the results used in this paper is archived on Zenodo under DOI https://doi.org/10.5281/zenodo.21726278 (Johnson et al. 2026a). Additional documentation and the User Guide are available at https://justinandrewjohnson.com/seals.

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

Data availability

All input data, trained parameters and output data are available at: https://doi.org/10.5281/zenodo.21813238 (Johnson et al. 2026b) licensed under CC BY 4.0.

Author contributions

JAJ developed the original SEALS model. JAJ and CCS developed the current version (2.0.0) and performed the simulations. JAJ, CCS, SKT and PJvJ contributed to the methodology, analysis, and interpretation. JAJ and CCS prepared the manuscript with contributions from all co-authors.

Competing interests

The authors declare that they have no competing interests.

Acknowledgements

Artificial intelligence was used to assist in code creation of the SEALS Model, specifically Claude Opus 4.8 used through Claude Code.

Financial support

None.

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

References

Bagstad, K. J., Cohen, E., Ancona, Z. H., McNulty, S. G., and Sun, G.: The sensitivity of ecosystem service models to choices of input data and spatial resolution, Appl. Geogr., 93, 25–36, https://doi.org/10.1016/j.apgeog.2018.02.005, 2018.

Baldos, U. L. and Corong, E.: Development of GTAP 10 Land Use and Land Cover Data Base for years 2004, 2007, 2011 and 2014, 2020.

Britz, W., Verburg, P. H., and Leip, A.: Modelling of land cover and agricultural change in Europe: Combining the CLUE and CAPRI-Spat approaches, Agr. Ecosyst. Environ., 142, 40–50, https://doi.org/10.1016/j.agee.2010.03.008, 2011.

Bukovsky, M. S., Gao, J., Mearns, L. O., and O'Neill, B. C.: SSP-based land-use change scenarios: A critical uncertainty in future regional climate change projections, Earth's Future, 9, e2020EF001782, https://doi.org/10.1029/2020EF001782, 2021.

Chaplin-Kramer, R., Sharp, R. P., Weil, C., Bennett, E. M., Pascual, U., Arkema, K. K., Brauman, K. A., Bryant, B. P., Guerry, A. D., Haddad, N. M., Hamann, M., Hamel, P., Johnson, J. A., Mandle, L., Pereira, H. M., Polasky, S., Ruckelshaus, M., Shaw, M. R., Silver, J. M., ... Daily, G. C.: Global modelling of nature's contributions to people, Science, 366, 255–258, https://doi.org/10.1126/science.aaw3372, 2019.

Chen, M., Vernon, C. R., Huang, M., Calvin, K. V., and Kraucunas, I. P.: Calibration and analysis of the uncertainty in downscaling global land use and land cover projections from GCAM using Demeter (v1.0.0), Geosci. Model Dev., 12, 1753–1764, https://doi.org/10.5194/gmd-12-1753-2019, 2019.

ESA: Land Cover CCI Product User Guide Version 2.0, available at: https://maps.elie.ucl.ac.be/CCI/viewer/download/ESACCI-LC-Ph2-PUGv2_2.0.pdf, 2017.

Ferreira, B. M., Soares-Filho, B. S., and Pereira, F. M. Q.: The Dinamica EGO virtual machine, Sci. Comput. Program., 173, 3–20, https://doi.org/10.1016/j.scico.2018.02.002, 2019.

Hoskins, A. J., Bush, A., Gilmore, J., Harwood, T., Hudson, L. N., Ware, C., Williams, K. J., and Ferrier, S.: Downscaling land-use data to provide global 30" estimates of five land-use classes, Ecol. Evol., 6, 3040–3055, https://doi.org/10.1002/ece3.2104, 2016.

Hurtt, G. C., Chini, L., Sahajpal, R., Frolking, S., Bodirsky, B. L., Calvin, K., Doelman, J. C., Fisk, J., Fujimori, S., Klein Goldewijk, K., Hasegawa, T., Havlik, P., Heinimann, A., Humpenöder, F., Jungclaus, J., Kaplan, J. O., Kennedy, J., Krisztin, T., Lawrence, D., ... Zhang, X.: Harmonization of global land use change and management for the period 850–2100 (LUH2) for CMIP6, Geosci. Model Dev., 13, 5425–5464, https://doi.org/10.5194/gmd-13-5425-2020, 2020.

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

IPBES: Global assessment report on biodiversity and ecosystem services of the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services, Zenodo, https://doi.org/10.5281/zenodo.5657041, 2019.

Johnson, J. A., Colesanti Senni, C., Thakrar, S. K., and von Jeetze, P. J.: jandrewjohnson/seals: v2.0.0, Zenodo [code], https://doi.org/10.5281/zenodo.21726278, 2026a.

Johnson, J. A., Colesanti Senni, C., Thakrar, S. K., and von Jeetze, P. J.: Data and trained coefficients for: "Projecting global changes in land use and ecosystem services using SEALS", Zenodo [data set], https://doi.org/10.5281/zenodo.21813238, 2026b.

Johnson, J. A., Ruta, G., Baldos, U., Cervigni, R., Chonabayashi, S., Corong, E., Gavryliuk, O., Gerber, J., Hertel, T., Nootenboom, C., and Polasky, S.: The Economic Case for Nature: A Global Earth-Economy Model to Assess Development Policy Pathways, World Bank, https://openknowledge.worldbank.org/handle/10986/35882, 2021a.

Johnson, J. A., Kennedy, C. M., Oakleaf, J. R., Baruch-Mordo, S., Polasky, S., and Kiesecker, J.: Energy matters: Mitigating the impacts of future land expansion will require managing energy and extractive footprints, Ecol. Econ., 187, 107106, https://doi.org/10.1016/j.ecolecon.2021.107106, 2021b.

Johnson, J., Baldos, U. L., Hertel, T., Nootenboom, C., Polasky, S., and Roxburgh, T.: Global Futures: Modelling the Global Economic Impacts of Environmental Change to Support Policy-making-Technical Report, 2020.

Johnson, J. A., Baldos, U. L., Corong, E., Hertel, T., Polasky, S., Cervigni, R., Roxburgh, T., Ruta, G., Salemi, C., and Thakrar, S.: Investing in nature can improve equity and economic returns, P. Natl. Acad. Sci. USA, 120, e2220401120, https://doi.org/10.1073/pnas.2220401120, 2023.

Johnson, J. A., Chaplin-Kramer, R., Chapman, M., Polasky, S., and Williams, B.: Earth-Economy Modeling: Advances in Linking Economic and Ecosystem Models, Annu. Rev. Resour. Econ., 17, 209–239, https://doi.org/10.1146/annurev-resource-013024-033043, 2025.

Kim, H., Rosa, I. M. D., Alkemade, R., Leadley, P., Hurtt, G., Popp, A., van Vuuren, D. P., Anthoni, P., Arneth, A., Baisero, D., Caton, E., Chaplin-Kramer, R., Chini, L., De Palma, A., Di Fulvio, F., Di Marco, M., Espinoza, F., Ferrier, S., Fujimori, S., ... Pereira, H. M.: A protocol for an intercomparison of biodiversity and ecosystem services models using harmonized land-use and climate scenarios, Geosci. Model Dev., 11, 4537–4562, https://doi.org/10.5194/gmd-11-4537-2018, 2018.

Krisztin, T., Wögerer, M., and Ringwald, L.: downscalr: Downscaling land-use change projections, R package version 0.1.0, International Institute for Applied Systems Analysis, https://github.com/tkrisztin/downscalr, 2023.

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

Lamarche, C., Santoro, M., Bontemps, S., D'Andrimont, R., Radoux, J., Giustarini, L., Brockmann, C., Wevers, J., Defourny, P., and Arino, O.: Compilation and Validation of SAR and Optical Data Products for a Complete and Global Map of Inland/Ocean Water Tailored to the Climate Modelling Community, Remote Sens., 9, 36, https://doi.org/10.3390/rs9010036, 2017.

Le Page, Y., West, T. O., Link, R., and Patel, P.: Downscaling land use and land cover from the Global Change Assessment Model for coupling with Earth system models, Geosci. Model Dev., 9, 3055–3069, https://doi.org/10.5194/gmd-9-3055-2016, 2016.

Lee, H.-L., Hertel, T. W., Rose, S., and Avetisyan, M.: An integrated global land use data base for CGE analysis of climate policy options, in: Economic Analysis of Land Use in Global Climate Change Policy, Routledge, 2009.

Mazy, F.-R. and Longaretti, P.-Y.: A formally correct and algorithmically efficient LULC change model-building environment, arXiv [preprint], https://doi.org/10.48550/arXiv.2203.17078, 2022.

Pontius, R. G., Boersma, W., Castella, J.-C., Clarke, K., de Nijs, T., Dietzel, C., Duan, Z., Fotsing, E., Goldstein, N., Kok, K., Koomen, E., Lippitt, C. D., McConnell, W., Mohd Sood, A., Pijanowski, B., Pithadia, S., Sweeney, S., Trung, T. N., Veldkamp, A. T., & Verburg, P. H.: Comparing the input, output, and validation maps for several models of land change. The Annals of Regional Science, 42(1), 11–37, https://doi.org/10.1007/s00168-007-0138-2

Popp, A., Calvin, K., Fujimori, S., Havlik, P., Humpenöder, F., Stehfest, E., Bodirsky, B. L., Dietrich, J. P., Doelmann, J. C., Gusti, M., Hasegawa, T., Kyle, P., Obersteiner, M., Tabeau, A., Takahashi, K., Valin, H., Waldhoff, S., Weindl, I., Wise, M., ... van Vuuren, D. P.: Land-use futures in the shared socio-economic pathways, Glob. Environ. Change, 42, 331–345, https://doi.org/10.1016/j.gloenvcha.2016.10.002, 2017.

Riahi, K., van Vuuren, D. P., Kriegler, E., Edmonds, J., O'Neill, B. C., Fujimori, S., Bauer, N., Calvin, K., Dellink, R., Fricko, O., Lutz, W., Popp, A., Cuaresma, J. C., Kc, S., Leimbach, M., Jiang, L., Kram, T., Rao, S., Emmerling, J., ... Tavoni, M.: The Shared Socioeconomic Pathways and their energy, land use, and greenhouse gas emissions implications: An overview, Glob. Environ. Change, 42, 153–168, https://doi.org/10.1016/j.gloenvcha.2016.05.009, 2017.

Schipper, A. M., Hilbers, J. P., Meijer, J. R., Antão, L. H., Benítez-López, A., de Jonge, M. M. J., Leemans, L. H., Scheper, E., Alkemade, R., Doelman, J. C., Mylius, S., Stehfest, E., van Vuuren, D. P., van Zeist, W.-J., and Huijbregts, M. A. J.: Projecting terrestrial biodiversity intactness with GLOBIO 4, Glob. Change Biol., 26, 760–771, https://doi.org/10.1111/gcb.14848, 2020.

Sharp, R. P., Douglas, J., Wolney, S., Johnson, J., and Wyatt, K.: InVEST User Guide, InVEST documentation, https://storage.googleapis.com/releases.naturalcapitalproject.org/invest-userguide/latest/index.html, 2020.

https://doi.org/10.5194/egusphere-2026-5027
Preprint. Discussion started: 16 September 2026

Soares-Filho, B. S., Coutinho Cerqueira, G., and Lopes Pennachin, C.: DINAMICA – a stochastic cellular automata model designed to simulate the landscape dynamics in an Amazonian colonization frontier, Ecol. Model., 154, 217–235, https://doi.org/10.1016/S0304-3800(02)00059-5, 2002.

Suh, S., Johnson, J. A., Tambjerg, L., Sim, S., Broeckx-Smith, S., Reyes, W., and Chaplin-Kramer, R.: Closing yield gap is crucial to avoid potential surge in global carbon emissions, Glob. Environ. Change, 63, 102100, https://doi.org/10.1016/j.gloenvcha.2020.102100, 2020.

Thakrar, S. K., Tessum, C. W., Apte, J. S., Balasubramanian, S., Millet, D. B., Pandis, S. N., Marshall, J. D., and Hill, J. D.: Global, high-resolution, reduced-complexity air quality modelling for PM2.5 using InMAP (Intervention Model for Air Pollution), PLOS ONE, 17, e0268714, https://doi.org/10.1371/journal.pone.0268714, 2022.

Thakrar, S. K., Johnson, J. A., and Polasky, S.: Land-Use Decisions Have Substantial Air Quality Health Effects, Environ. Sci. Technol., 58, 381–390, https://doi.org/10.1021/acs.est.3c02280, 2024.

van Asselen, S. and Verburg, P. H.: Land cover change or land-use intensification: Simulating land system change with a global-scale land change model, Glob. Change Biol., 19, 3648–3667, https://doi.org/10.1111/gcb.12331, 2013.

Vernon, C. R., Le Page, Y., Chen, M., Huang, M., Calvin, K. V., Kraucunas, I. P., and Braun, C. J.: Demeter – a land use and land cover change disaggregation model, J. Open Res. Softw., 6, 15, https://doi.org/10.5334/jors.208, 2018.

von Jeetze, P. J., Weindl, I., Johnson, J. A., Borrelli, P., Panagos, P., Molina Bacca, E. J., Karstens, K., Humpenöder, F., Dietrich, J. P., Minoli, S., Müller, C., Lotze-Campen, H., and Popp, A.: Projected landscape-scale repercussions of global action for climate and biodiversity protection, Nat. Commun., 14, 2515, https://doi.org/10.1038/s41467-023-38043-1, 2023.

von Jeetze, P. J., Weindl, I., Johnson, J. A., Borrelli, P., Panagos, P., Meyer, T., Humpenöder, F., Sauer, P., Dietrich, J. P., Lotze-Campen, H., and Popp, A.: Conservation outcomes of dietary transitions across different values of nature, Nat. Sustain., 8, 1130–1142, https://doi.org/10.1038/s41893-025-01595-9, 2025.

Wolff, S., Schrammeijer, E. A., Schulp, C. J. E., and Verburg, P. H.: Meeting global land restoration and protection targets: What would the world look like in 2050?, Glob. Environ. Change, 52, 259–272, https://doi.org/10.1016/j.gloenvcha.2018.08.002, 2018.

Woodman, T., Rueda-Uribe, C., Henry, R. C., Burslem, D. F. R. P., Travis, J. M. J., and Alexander, P.: Introducing LandScaleR: A novel method for spatial downscaling of land use projections, Environ. Modell. Softw., 169, 105826, https://doi.org/10.1016/j.envsoft.2023.105826, 2023.