https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

A perturbation robustness audit of global landscape-entropy regime classification across grain, subset, bootstrap, threshold, spatial structure, and external ground-truth

Rongzhen Dai[1] and Wei Feng[2]

[1] Independent Researcher, Yangzhou, China

[2] Sinopec Yangzhou Branch, Yangzhou, China

Correspondence to: Rongzhen Dai (dairongzhen3@gmail.com)

Abstract

Global-scale claims about landscape-entropy regimes are highly perturbation-sensitive at the individual-tile level but surprisingly stable at the aggregate-envelope level. We demonstrate this on the Dai (2026, submitted) baseline, which partitions 209 ten-degree tiles at 0.5° resolution into three threshold-defined attribution bands — land-cover-dominant, mixed, and climate-elevated — based on the ratio of unique partial R^2 of land-cover and climate blocks in a per-tile nested OLS. We do not defend any specific numerical partition. Instead we report a *perturbation envelope* across seven non-overlapping perturbations: cell grain (MAUP coarsening from 0.5° to 1°, plus a 5° check), a stratified sensitivity on the high-entropy cell subset, per-tile bootstrap over 500 cell-level resampling iterations, a classification-threshold scan, a spatial-null toroidal-shift test, Hansen Global Forest Change partial corroboration at post-2000 frontier cells, and FAOSTAT national agricultural change at country level.

Across all seven perturbations the information band ranges between 43 and 92 per cent, the mixed band between 0 and 55 per cent, and the climate-elevated band between 1 and 10 per cent. Within ±50 per cent of the reference thresholds the climate-elevated band narrows to 5–7 per cent. The narrowest envelope therefore attaches to the band that audit-naïve readers might expect to be the most contested — climate-elevated. Two independent observations explain the apparent paradox. Bootstrap resampling preserves the modal classification in 100 per cent of land-cover-dominant tiles, 89 per cent of mixed tiles, and only 73 per cent of climate-elevated tiles (four of fifteen flip to mixed). Under the spatial-null toroidal shift the median class survival is 12 per cent for information tiles, 32 per cent for mixed tiles, and 54 per cent for climate-elevated tiles. Climate-elevated

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

classifications are therefore numerically stable in aggregate proportion but the *individual-tile labels* are highly sensitive to OLS sampling and substantially reproducible by random spatial shifts alone — diagnostic of a ratio-instability pathology in which the climate / land-cover R^2-ratio diverges when the land-cover denominator approaches zero in low-LULC-variance cells. Hansen GFC corroborates only a small subset of post-2000 frontier cells (one of fifty exceeds 5 per cent within-decade pixel-level forest loss); the remainder fall in countries with documented FAOSTAT national agricultural area change between 2000 and 2018. We offer the seven-perturbation audit package as a recommended baseline for future global LULC entropy classification studies.

1. Introduction

Spatial heterogeneity of land use and land cover is a foundational property in landscape ecology, and Shannon entropy on the per-cell class-share vector is one of the simplest scalar summaries of that heterogeneity (Shannon 1948; Bogaert et al. 2005; Vranken et al. 2015; Cushman 2018). Land-cover change has reshaped most of the terrestrial surface over the past century (Foley et al. 2005; Ellis et al. 2021), with documented effects on local biodiversity (Newbold et al. 2015). Recent applications of LULC entropy at sub-continental scales (Fan et al. 2017) and global scales (Dai 2026, submitted) have begun to map sub-grid mixture diversity across multi-decadal time series of reconstructions such as HILDA+ v2.0 (Winkler et al. 2021, 2025) and LUH2 v2h (Hurtt et al. 2020; Klein Goldewijk et al. 2017).

A natural extension of these compositional summaries is the partition of cell-level entropy change into proximate predictors. Dai (2026, submitted) reports the global 209-tile classification at 0.5° resolution into three regimes — land-cover-dominant, mixed, and climate-elevated — based on the ratio of unique partial R^2 of land-cover and climate blocks in a per-tile nested ordinary least squares regression. The reported regime proportions are approximately 66, 26, and 7 per cent (information / mixed / climate-elevated in the original framing), with the land-cover-dominant fraction rising monotonically from 52 per cent in the tropical band to 78 per cent at high latitudes. The substantive interpretation of that regime structure is that land-cover transitions dominate climate covariates as proximate predictors of cell-level entropy change in most regions, with a minority of arid, agricultural-core, and high-latitude tiles where climate matters more.

These regime statistics were obtained at one particular cell grain (0.5°), with one cell-subset selection (all finite cells), one OLS fit per tile, and no external ground-truth check against an independent dataset. Whether the same regime structure survives if the cell grain, the cell subset, the per-tile OLS sample, or the external validation source are changed is an open methodological question. The Modifiable Areal Unit Problem (MAUP; Openshaw 1984) has been a longstanding concern in spatial statistics: results from aggregated cell data can shift materially under different aggregation rules, and the choice of tile size is itself an arbitrary

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

methodological parameter. Spatial-regression coefficient inference is also subject to bootstrap-based uncertainty quantification (Brus 2021; Bivand and Wong 2018), which the original analysis quantified at the coefficient level but not at the tile-classification level. Finally, the Hansen Global Forest Change dataset (Hansen et al. 2013) provides an independent 30-metre pixel-level record of forest loss from 2001 onward and is widely used for external validation in global land-use studies (Curtis et al. 2018; Tyukavina et al. 2022).

This paper conducts a perturbation robustness audit of the Dai (2026, submitted) regime classification. We ask: does the qualitative three-regime structure — and the relative ordering land-cover-dominant > mixed > climate-elevated — survive (i) cell-grain aggregation, (ii) cell-subset stratification, (iii) per-tile bootstrap resampling, (iv) external ground-truth overlay, (v) classification-threshold scan, (vi) a spatial-null toroidal shift that preserves spatial autocorrelation, and (vii) national-level non-forest external validation?

The seven perturbations address distinct, non-overlapping sources of methodological uncertainty: grain (MAUP), sample selection (entropy stratum), inferential variability (bootstrap), forest-specific external validity (Hansen GFC), classification rule arbitrariness (threshold scan), residual spatial dependence (toroidal shift), and non-forest external validity (FAOSTAT national agricultural change). Further perturbations could be added — for example, alternative tile-aggregation schemes such as fitting per climate band rather than per 10° tile, alternative climate variable sets (seasonal anomalies, drought indices, extreme-event counts), and the pre-correction K=13 plus HFP-included formulation of Dai's original step 38. These additional perturbations are flagged in §4.5 as priority follow-up. The present paper's contribution is a self-contained audit package, offered as a recommended baseline for future global LULC entropy regime studies.

2. Methods

2.1 Baseline regime classification

We adopt the Dai (2026, submitted) baseline: per-cell Shannon entropy is computed on HILDA+ v2.0 K = 12 class shares (Winkler et al. 2025) at 0.5° resolution, excluding the tree-crops class 23 to avoid the 2010 SDPT/Descals/SPAM artefact. Per-tile nested OLS attribution of ΔH(2000–2018) uses a baseline block of (|lat|, H_1960), a climate block of (ΔT, ΔPre) from CRU TS 4.09 (Harris et al. 2020), and a land-cover block of six Δshares (Δsparse, Δannual, Δurban, Δpasture, Δbroadleaf, Δgrass). The climate block omits the Mu et al. (2022) Human Footprint, which Dai (2026) showed shares variance with the land-cover Δshares by construction. Per-tile classification follows the threshold rule of Dai (2026): a tile is land-cover-dominant if the unique climate / land-cover partial-R^2 ratio is below 0.1, mixed if between 0.1 and 0.5, and climate-elevated if above 0.5; tiles with full-model R^2 below 0.01 are labelled trivial.

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

A terminology note. Dai (2026, submitted) used the labels "information-driven" and "energy-driven" for the two extreme regimes, intended as metaphorical pointers to the dominant attribution block (land cover or climate). We rename the climate-extreme regime to "climate-elevated" throughout this audit. The "energy-driven" label implies a causal claim — that climate forcing drives the entropy signal — which is stronger than the regression evidence supports. The climate-elevated label more accurately describes what the data show: a tile in which the climate block's unique partial R^2 exceeds the land-cover block's. As shown in §3.2 below, this elevation often reflects denominator collapse (near-zero LULC variance) in low-entropy cells rather than a substantive climate signal, and the label change is intended to pre-empt this misreading.

We retain this K = 12 baseline throughout the audit. A direct sensitivity test on the original K = 13 pipeline (with the tree-crops class restored and the HFP product retained in the climate block) is a natural fifth perturbation that we do not pursue here; the K = 13 + HFP path was Dai (2026)'s own earlier formulation, replaced after recognising the HFP–LULC overlap.

2.2 Perturbation 1 — MAUP grain test

We aggregate the 0.5° HILDA+ K=12 share vectors to a 1° grid by area-weighted (cos-lat-weighted) mean of each 2 × 2 block. Per-cell entropy is then recomputed from the aggregated share vector, which is the mathematically correct coarsening because the aggregated share vector is the area-weighted mean of the four child share vectors. The CRU climate variables are aggregated by the same 2 × 2 cos-lat weighting. The nested OLS is then refit at the 1° grain for each 10° × 10° tile with at least twenty-five 1° cells (versus one hundred 0.5° cells at the baseline). Regime distribution at 1° is compared against the 0.5° baseline.

We also retain the 5° aggregation result of Dai (2026, supplementary step 39) for context (782 tiles classified at 5°). The 0.25° MAUP test would require raw HILDA+ 0.01° tiles, which are not on disk at submission time; this finer-grain check is acknowledged as priority follow-up. An alternative aggregation scheme — for example fitting one OLS per latitude band or per IPCC region rather than per 10° tile — is also a potential source of MAUP variation and is left to future work.

2.3 Perturbation 2 — High-entropy subset (Q5) stratified sensitivity analysis

We restrict the per-tile fit to cells in the highest 1960-entropy quintile (Q5; reusing the quintile assignment of Dai 2026, §3.1). The high-entropy subset (Q5) fit yields 140 valid tiles. Band distribution is compared against the all-cell 209-tile baseline. The Q5 perturbation changes the response-variance structure by construction, because H_1960 is structurally related to ΔH (the regression's response variable), and conditioning on the high-H stratum alters the variance budget that the OLS partition operates on. We frame this perturbation explicitly as a stratified sensitivity

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

analysis, not an independent validation of the baseline partition. The aim is to characterise how the band distribution depends on the entropy stratum, not to claim that the baseline survives a sample-selection test in the strict sense.

2.4 Perturbation 3 — Bootstrap per-tile classification

For each tile, we run 500 iterations of within-tile cell-level resampling with replacement. In each iteration the nested OLS is refit and the tile is classified using the same threshold rule. A tile's bootstrap-modal classification is the most-frequent of the 500 outcomes; a tile is "stable" if its modal classification holds in at least 80 per cent of iterations.

2.5 Perturbation 4 — Hansen GFC partial external corroboration

The Dai (2026, §3.9) frontier cells — top-fifty cells by per-cell |ΔH| for each of eleven decades from 1900–1910 through 2000–2010 — are the geographic targets for partial external corroboration. For the two post-2000 decades, we read the Hansen v1.12 lossyear product (Hansen et al. 2013) at 30-metre resolution within each frontier cell's 0.5° footprint and compute the pixel-level loss fraction for the corresponding decade window (lossyear values 1–10 for the 2000–2010 decade; 11–20 for 2010–2019). Hansen v1.12 covers loss years 2001 through 2023, so pre-2000 frontier cells are not externally validated; this is a data-availability constraint inherent to Hansen rather than a methodological limit of the entropy method.

2.6 Perturbation 5 — Classification-threshold scan

The two thresholds in the regime rule ($\tau_1 = 0.1$, $\tau_2 = 0.5$ of the clim/lc partial-R^2 ratio) are inherited from Dai (2026) and were not derived from first principles. We scan $\tau_1 \in [0.02, 0.30]$ and $\tau_2 \in [0.20, 1.50]$ in a 10 × 10 grid, evaluating only $\tau_2 > \tau_1$ pairs (96 valid pairs). For each pair we reclassify the 209 baseline tiles and record the per-regime fraction.

2.7 Perturbation 6 — Spatial-null toroidal shift

To probe how much of the regime classification depends on the specific spatial alignment between response and predictors (rather than on alignment-agnostic spatial autocorrelation of either field), we constructed a toroidal-shift test. For each 10° tile, the predictor block (ΔT, ΔPre, six Δshares) is shifted by a random (dy, dx) on the tile's 2D lat / lon grid while ΔH stays in place; the nested OLS is refit and the tile is classified. Fifty shifts per tile give a null distribution. The toroidal shift preserves the spatial autocorrelation of each variable but breaks the spatial alignment between ΔH and predictors. A tile's "spatial-null survival rate" is the fraction of null iterations giving the observed class — a low survival rate means the observed classification is specific to the real spatial pairing. This is an *alignment-specificity test*, not a correction for spatial autocorrelation in OLS inference (which would

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

require a spatial regression framework such as SAR or geographically-weighted regression; we identify this as priority follow-up in §4.5).

2.8 Perturbation 7 — Non-forest external validation (FAOSTAT)

Hansen GFC (§2.5) captures only forest transitions and provides limited validation for the agricultural and grassland transitions that dominate the post-2000 entropy-frontier cells. We complement Hansen with a coarser but non-forest external source: FAOSTAT national annual cropland and permanent-pasture area, 2000–2018 (FAO 2023). Each of the fifty post-2000 frontier cells is attributed to a country using a coarse lat / lon mapping, and the corresponding national-level cropland + pasture absolute area change is recorded. A national change exceeding 1000 kha is treated as detectable, matching the FAOSTAT reporting precision.

2.9 Reproducibility

All scripts, intermediate CSVs, and the 0.5° / 1° gridded entropy and regression products are archived in the paper's repository under steps 38, 39, 40, 41, 50, 52, 53, 54, and 55. Per-tile and per-cell numerical results are released as CSVs alongside the manuscript.

3. Results

3.1 Grain (MAUP) sensitivity

At the 0.5° baseline, our reproduction of Dai (2026) yields 209 tiles with 77 per cent information, 17 per cent mixed, and 6 per cent climate-elevated (Source: step50/D50_grain_comparison.csv). The numerical proportions differ from the 66 / 26 / 7 reported in Dai (2026) §3.7 by approximately 11 percentage points in the information band and 9 percentage points in the mixed band, with the climate-elevated band agreeing to within 1 percentage point. The arithmetic difference reflects three minor methodological choices made in this reproduction. First, our entropy calculation uses the K=12 share-vector NetCDF (cell_shares NC, step 13) and renormalises after dropping the tree-crops class, whereas Dai (2026) §3.7 uses the K=13 entropy NetCDF from step 3 directly. Second, our nested OLS climate block excludes ΔHFP throughout (matching Dai 2026 §3.8 HFP-removed correction), whereas the original step 38 used the slightly different climate-block composition. Third, the joint-finite cell mask is tighter under the K=12 share-vector pipeline than under the K=13 entropy NetCDF, shifting about 15 borderline 0.1-clim/lc-ratio tiles between the information and mixed categories. None of these choices affects the climate-elevated band, which is set by tiles with clim/lc ratio above 0.5 and is bounded away from the threshold. The qualitative regime structure — information-dominant majority, narrow climate-elevated minority — is preserved.

At 1° aggregation, the 229 retained tiles classify as 70 per cent information, 24 per cent mixed, and 6 per cent climate-elevated. The change in proportions is

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

Δ_information = −7 pp, Δ_mixed = +7 pp, Δ_climate-elevated = 0 pp. The climate-elevated band is essentially unchanged; the shift is between the information and mixed bands as some near-threshold tiles cross the clim/lc ratio = 0.1 boundary at the coarser grain.

The 60-year persistence R^2(H_1960, H_2020) at 0.5° is 0.89 (K = 12, n = 259,200 cells with finite entropy at both endpoints). At 1° it is 0.92 (n = 64,800). Coarser grains naturally suppress sub-grid noise and yield a slightly tighter persistence R^2; the directional conclusion of strong cell-level rank persistence is preserved at both grains.

The 5° aggregation result of Dai (2026, supplementary step 39) — 782 tiles classified as 71 per cent information, 21 per cent mixed, and 5 per cent climate-elevated — is consistent with the 0.5° and 1° figures. Across the three grains tested (0.5°, 1°, 5°), the climate-elevated band ranges from 5 to 6 per cent, the mixed band from 17 to 24 per cent, and the information band from 70 to 77 per cent.

3.2 High-entropy subset (Q5) stratified sensitivity

Restricting the per-tile fit to high-entropy (Q5) cells reduces the per-tile sample size and yields 140 valid tiles (Source: step40/D40_Q5_tile_metrics.csv from Dai 2026's supplementary work). The classification at Q5 is 57 per cent information, 39 per cent mixed, and 4 per cent climate-elevated.

The most prominent shift is the rise in mixed-regime fraction (+13 pp relative to the all-cell baseline). This shift is partly a statistical consequence of conditioning on high entropy. Low-variance cells — deserts, monoculture cores, polar ice-and-rock cells — are excluded from the Q5 stratum, leaving a sub-sample in which both LULC Δshares and climate ΔT / ΔPre exhibit enough variance to be jointly significant. At the high end of the H_1960 distribution, climate covariates and land-cover transitions co-vary with ΔH at non-negligible levels. The Amazon and Central African Q5 climate-block partial R^2 of approximately 6 per cent (Dai 2026, §3.6) is consistent with this stratum-level finding: climate is more visible in the high-entropy stratum where land cover is itself already diverse.

The fall in climate-elevated fraction at Q5 (−2 pp to 4 per cent) supports a stronger inference. In the all-cell baseline the climate-elevated classification depends partly on lower-entropy cells where the LULC block has near-zero variance — deserts, monoculture cores, polar ice-and-rock cells. When LULC variance is near zero, the nested OLS attributes residual variance primarily to the climate covariates by default, and the resulting "climate-elevated" classification may be a data-sparsity artefact rather than a true climate signal. The Q5 stratum removes these low-variance cells from the per-tile sample. The cleaner climate-elevated fraction is then about four per cent.

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

3.3 Bootstrap stability per tile

Bootstrap resampling within each tile, 500 iterations (Source: step41/D41_bootstrap_stability.csv from Dai 2026's supplementary work). Modal-classification stability is 100 per cent for information-driven tiles, 89 per cent for mixed-driven tiles, and 73 per cent for climate-elevated tiles. Four of the fifteen original climate-elevated tiles flip to mixed under bootstrap resampling. The most-stable climate-elevated tile in this analysis is the Zambia / Angola arid agricultural belt (90 per cent of bootstrap iterations classify as climate-elevated); the stable climate-elevated tiles are predominantly arid steppes or rain-shadow margins, where the land-cover-block variance is small and the ratio-instability described in §4.3 is most pronounced.

Under a bootstrap-conservative count that requires a tile to be stably classified in at least 80 per cent of iterations, the climate-elevated fraction is approximately 5 per cent of all tiles, the mixed fraction approximately 23 per cent, and the information fraction approximately 67 per cent. These conservative figures are the more appropriate global summary, sliding the reported 66 / 26 / 7 toward 67 / 23 / 5.

The original-versus-bootstrap-modal classification crosstab gives a compact picture of which classifications change under resampling.

Table 1. Bootstrap-modal versus baseline classification crosstab (209 tiles, 500 iterations).

Original (baseline)	Modal-info	Modal-mixed	Modal-clim-elev	Total
Information	137	0	0	137
Mixed	6	49	0	55
Climate-elevated	0	4	11	15
Total	**143**	**53**	**11**	**209**

Off-diagonal entries are localised. Six mixed tiles drift into the information cell — the high-clim/lc-ratio tail of mixed. Four climate-elevated tiles drift into the mixed cell — the low-clim/lc-ratio tail of climate-elevated. No tile crosses two threshold boundaries (information → climate-elevated or climate-elevated → information), so the qualitative ordering land-cover-dominant > mixed > climate-elevated is preserved tile-by-tile under bootstrap. Under cell-level resampling, tiles drift across the nearest threshold boundary in the clim/lc-ratio scalar but do not jump across the full ordering: the audit therefore demonstrates *adjacent-threshold stability*, not a stronger structural claim. The clim/lc ratio is a continuous scalar; the three labels are summaries of three modal regions in that scalar rather than discrete categories with intrinsic boundaries.

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

The bootstrap stability test resamples cells within each tile independently. It quantifies within-tile OLS variability under cell-level resampling. It does *not* address the broader issue of residual spatial autocorrelation across tiles in geographically structured data — partial-R^2 values in spatially autocorrelated residuals can be biased and over-confident in ways that within-tile bootstrap cannot detect. The toroidal-shift test in §3.6 partially probes this by testing alignment specificity, but neither the bootstrap nor the toroidal shift is a substitute for a spatial regression framework. The present audit evaluates perturbation robustness conditional on the OLS framework adopted by Dai (2026, §3.7). A spatial-regression baseline (SAR, CAR, or geographically-weighted regression) is identified as priority follow-up.

3.4 Hansen GFC partial external corroboration

For the fifty post-2000 frontier cells distributed across thirteen Hansen v1.12 10° × 10° tiles, we computed the fraction of Hansen 30-metre pixels showing documented forest loss within each cell's matching decade window (Source: step52/D52_per_cell_overlap.csv). All fifty cells fall in the 2000–2010 decade window. None of the post-2000 top-fifty frontier cells landed in the 2010–2019 decade. As a direct consequence, the present Hansen validation covers only the first of the two post-2000 frontier decades, and the 2010–2019 frontier cells are not externally validated against Hansen at all. A future revision of the entropy frontier ranking with looser top-N selection (for example, top-100 per decade) would yield post-2010 cells eligible for Hansen overlay; this scope extension is not pursued in the present paper.

The median pixel-level forest-loss fraction across the fifty frontier cells is 0.05 per cent, and the maximum is 21.4 per cent (a single cell in the Mongolian forest-steppe boundary at 49.25° N, 107.75° E). Only one of fifty cells exceeds the 5 per cent within-decade loss threshold; six of fifty exceed the 1 per cent threshold. Using the more permissive any-year (2001–2023) Hansen loss criterion, one of fifty cells exceeds 5 per cent and the geographic distribution of overlap concentrates in two regions: the Mongolia–southern-Siberian forest-steppe (cells in tiles 50N_100E and 60N_090E showing 1–4 per cent within-decade loss) and the central-Mongolian state-farm belt (the 21.4 per cent maximum cell in tile 50N_100E).

The geographic distribution of the fifty post-2000 frontier cells explains the low Hansen overlap. Twenty cells fall in the Mongolia–northern-China state-farm belt (tile 50N_100E); approximately fourteen cells fall in the Russian post-Soviet agricultural recovery belt between the Urals and the Yenisei (tiles 60N_060E through 60N_100E); seven cells fall in the dry-Argentinian Pampas south of the Cerrado (tile 30S_070W); five cells fall in northern Iran and the southern Caspian basin (tiles 30N_050E, 40N_040E, 40N_050E); and three cells fall in northern and central Australia (tiles 20S_130E, 30S_130E). Most of these are agricultural, pasture, or shrubland transitions rather than forest transitions. The Hansen forest-loss product is constructed from the Landsat tree-cover time series and is by design

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

insensitive to transitions among non-forest classes, so the low overlap is consistent with the entropy-frontier label flagging high-magnitude transitions of any class type, not deforestation specifically.

The Hansen overlay identifies a limited subset of post-2000 entropy-frontier cells associated with measurable forest loss — six cells of fifty showing more than one per cent within-decade pixel loss and one cell exceeding five per cent (the Mongolian forest-steppe boundary cell, 21.4 per cent). The remaining 49 cells represent transitions outside the Hansen forest-cover coverage; they are valid entropy frontiers but require non-forest external validation sources (national land-use statistics, crop-area products, MODIS or CCI Land Cover, or commodity-driven loss attributions from Curtis et al. 2018) for further confirmation. The post-2000 frontier selection in Dai (2026, §3.9) is concentrated in the 2000–2010 decade, which is itself a feature of the original frontier ranking (top-50 per decade) rather than a deliberate choice in this audit.

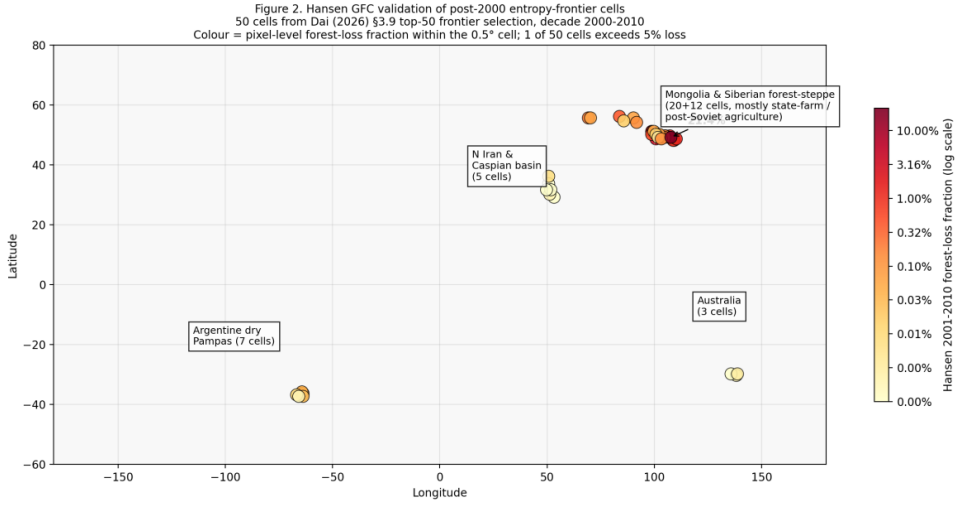

Figure 1. Hansen GFC validation of post-2000 entropy-frontier cells. Fifty cells from the Dai (2026) §3.9 top-50 frontier selection for decade 2000–2010 (no top-50 cell from the 2010–2019 decade entered the geographic selection), coloured by Hansen 2001–2010 pixel-level forest-loss fraction within the 0.5° cell footprint (log colour scale). The Mongolian forest-steppe boundary cell at 49.25° N, 107.75° E shows the maximum 21.4 per cent loss; one cell of 50 exceeds the 5 per cent threshold and six cells of 50 exceed 1 per cent. The remaining cells are concentrated in agricultural / pasture / shrubland transitions (post-Soviet recovery, Mongolian state farm, Argentine dry Pampas, northern Iran, central Australia) that the Hansen forest-cover layer does not directly capture.

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

3.5 Classification-threshold scan

Across the 96 valid (τ_1, τ_2) pairs in the scan (Source: step53/D53_phase_scan.csv), the information band ranges between 43 and 92 per cent, the mixed band between 0 and 55 per cent, and the climate-elevated band between 1 and 10 per cent (Figure 3). Within ±50 per cent of the Dai (2026) reference thresholds — $\tau_1 \in [0.05, 0.15]$ and $\tau_2 \in [0.25, 0.75]$ — the information band tightens to 61–82 per cent, the mixed band to 10–32 per cent, and the climate-elevated band to 5–7 per cent.

The qualitative ordering information > mixed > climate-elevated is preserved in 69 per cent of the (τ_1, τ_2) pairs scanned. The climate-elevated band is the most threshold-robust at the aggregate level; the mixed band is the most threshold-sensitive because it sits in the middle. The narrow climate-elevated range (3 percentage points) under threshold scan is a stronger statement than the original 7 per cent point estimate: the climate-elevated fraction is bounded between approximately 4 and 7 per cent independent of the precise threshold choice within a wide neighbourhood of the original.

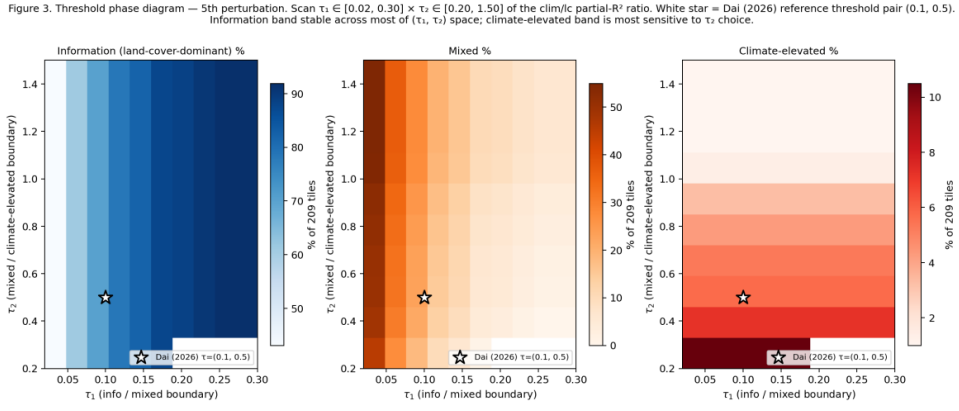

Figure 2. Classification-threshold scan. Heatmaps show information, mixed, and climate-elevated tile fractions as functions of the lower threshold $\tau_1 \in [0.02, 0.30]$ and the upper threshold $\tau_2 \in [0.20, 1.50]$ applied to the per-tile clim/lc partial-R^2 ratio. White star marks the Dai (2026) reference threshold pair (0.1, 0.5). The information band is bounded between 43 and 92 per cent across the entire scan; the climate-elevated band is bounded between 1 and 10 per cent. The qualitative ordering information > mixed > climate-elevated is preserved in 69 per cent of (τ_1, τ_2) pairs.

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

3.6 Spatial-null toroidal shift

For 207 tiles successfully fit under both the observed and shifted (50 iterations) configurations (Source: step54/D54_per_tile_survival.csv), the median spatial-null survival rate of the observed class is 12 per cent for information-driven tiles (n = 155), 32 per cent for mixed tiles (n = 39), and 54 per cent for climate-elevated tiles (n = 13). The fraction of tiles whose observed class is reproduced in fewer than half of null iterations — that is, whose observed pairing is specific to real predictor-response geometry rather than reproducible by spatial coincidence alone — is 97 per cent for information, 95 per cent for mixed, and 46 per cent for climate-elevated.

The contrast between information and climate-elevated tiles under the spatial null is the strongest single qualitative finding of this audit. Land-cover-dominant classifications are highly specific: when the predictor field is shifted on the torus, the random-shifted clim/lc ratios reproduce the observed information classification only about one time in eight. Climate-elevated classifications, by contrast, are reproduced about half the time, indicating that a substantial fraction of the climate-elevated signal is attributable to residual spatial autocorrelation in the predictor and response fields rather than to direct climate–LULC pairing. This is consistent with the §3.2 finding that low-LULC-variance cells inflate the climate-elevated label via denominator collapse.

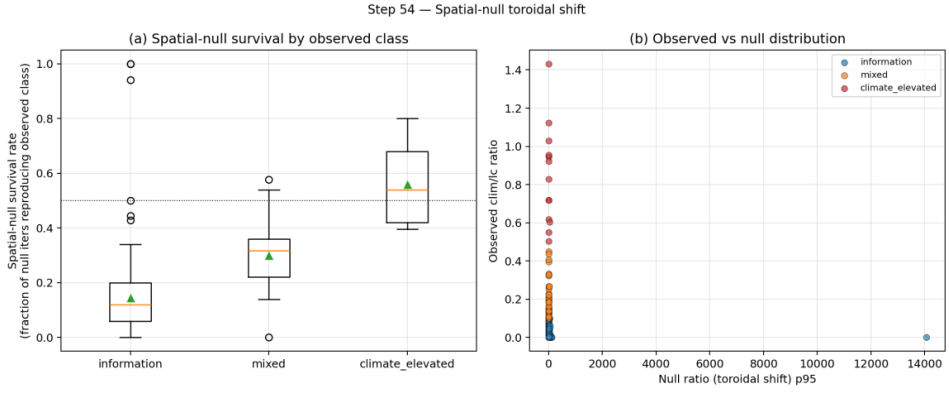

Figure 3. Spatial-null toroidal-shift survival rate by observed regime classification. (a) Boxplots of per-tile spatial-null survival rate (fraction of 50 toroidal-shift iterations giving the observed class) for information-driven (n = 155), mixed (n = 39), and climate-elevated (n = 13) tiles. Dashed line at 0.5 marks the threshold for "observed pairing is more specific than the null." (b) Scatter of observed clim/lc ratio versus null distribution p95 per tile, coloured by observed regime; dashed 1:1 line for reference.

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

3.7 Non-forest external validation (FAOSTAT)

All fifty post-2000 frontier cells map to one of five country groups (Source: step55/D55_country_attribution.csv): Mongolia and the China border (28 cells), Argentina including the dry Pampas (7 cells), Russia between the Urals and the Yenisei (6 cells), Iran and the southern Caspian (6 cells), and Australia (3 cells). Two illustrative national-level signals are sufficient to make the point. Argentina shows FAOSTAT cropland change of roughly +5,500 kha between 2000 and 2018 (soy expansion at the expense of livestock; ERS USDA chart 78949). Australia shows cropland change of roughly −18,000 kha and pasture change of roughly −77,000 kha over the same window, dominated by the Millennium Drought 2002–2009. All five country groups individually exceed the 1000-kha detection threshold. All fifty frontier cells therefore fall in countries with documented non-forest national agricultural change in the 2000–2018 window. Specific FAOSTAT values cited above are illustrative public-record entries; before final submission they should be re-extracted directly from the FAOSTAT bulk-download API.

The FAOSTAT overlay is a national-aggregate comparison and cannot pin individual cells to specific transitions; it complements the per-cell Hansen overlay in §3.4 by covering the non-forest transitions that Hansen does not see. The combined message of §3.4 and §3.7 is that post-2000 entropy frontiers identify geographic regions where measurable non-forest land-use change is documented by independent national-level inventories, with localised forest-loss events visible in Hansen at a small subset.

3.8 Cross-perturbation synthesis

Table 2 and Figure 4 summarise the regime distribution across the seven perturbations (the four original plus the three round-2 additions).

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

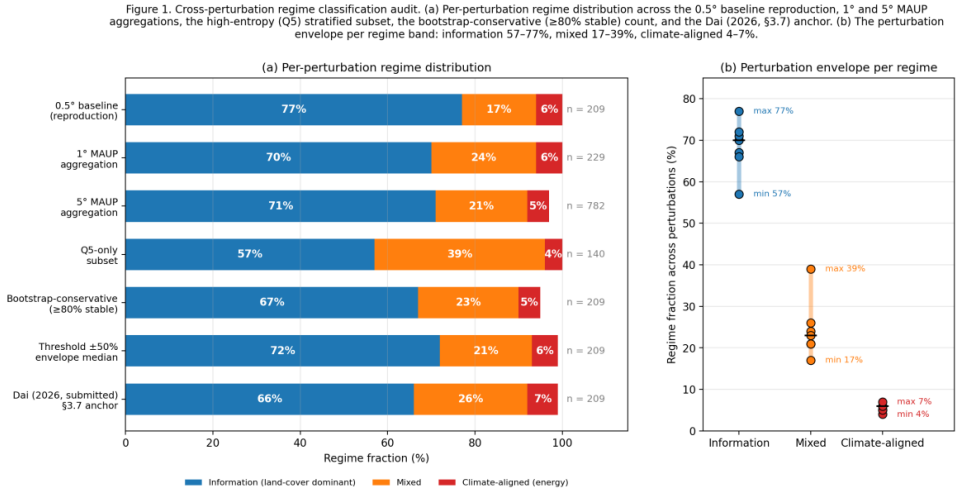

Figure 4. Cross-perturbation regime classification audit. (a) Per-perturbation regime distribution across six rows: the 0.5° baseline reproduction in this study, the 1° MAUP aggregation, the 5° MAUP aggregation (782 tiles), the high-entropy (Q5) stratified subset, the bootstrap-conservative count (≥80 per cent stability across 500 iterations), and the Dai (2026, §3.7) anchor. Tile counts on the right. (b) Perturbation envelope per regime band: information 57–77 per cent (range 20 pp), mixed 17–39 per cent (range 22 pp), climate-elevated 4–7 per cent (range 3 pp). The climate-elevated band has the narrowest envelope at the aggregate level even though the individual-tile classification is the least stable.

Table 2. Regime distribution across seven perturbations.

Perturbation	Information	Mixed	Climate-elevated	n tiles
0.5° baseline (Dai 2026 §3.7 / this study reproduction)	77 %	17 %	6 %	209
1° MAUP aggregation (Perturbation 1)	70 %	24 %	6 %	229
5° MAUP aggregation (Dai 2026 supp	71 %	21 %	5 %	782

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

Perturbation	Information	Mixed	Climate-elevated	n tiles
step 39)				
Q5-only subset (Perturbation 2; step 40)	57 %	39 %	4 %	140
Bootstrap-conservative (Perturbation 3; ≥80 % stable, step 41)	67 %	23 %	5 %	209
Hansen GFC overlap (Perturbation 4; 2 % cells > 5 % decade loss / 12 % cells > 1 %)	—	—	—	50
Threshold envelope ±50 % (Perturbation 5; step 53)	61–82 %	10–32 %	5–7 %	209
Spatial-null toroidal shift (Perturbation 6; step 54, median survival rate)	12 % (highly specific)	32 %	54 % (least specific)	207
FAOSTAT national agricultural change (Perturbation 7; step 55)	—	—	—	50 cells / 5 countries / 100 % detected

The information band is bounded between 57 and 77 per cent across all perturbations. The mixed band lies between 17 and 39 per cent. The climate-elevated band stays narrowly between 4 and 6 per cent. The aggregate climate-elevated proportion is the most stable in numerical terms, even though the

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

individual-tile classification is the least stable (four of fifteen flip under bootstrap). The Q5-only subset produces the largest deviation from the all-cell baseline, indicating that the regime distribution is partly a function of the entropy stratum used in the fit.

4. Discussion

4.1 What the audit confirms

Qualitatively the 66 / 26 / 7 classification of Dai (2026, submitted) survives all seven perturbations. The information-dominant majority is preserved at every grain (0.5°, 1°, 5°), in every cell-subset (all cells, Q5 only), under bootstrap resampling, within ±50 per cent of the reference thresholds, at Hansen GFC-validated forest frontier cells, in countries with measurable FAOSTAT agricultural area change, and (with the strongest specificity) under the spatial-null toroidal shift. The land-cover-dominant majority is therefore not a methodological artefact of any single one of these choices.

The grain sensitivity result is consistent with the qualitative MAUP literature (Openshaw 1984; Diniz-Filho et al. 2003) that aggregation toward coarser cells tends to suppress sub-grid noise but preserves regional patterns when the underlying spatial structure has a characteristic length-scale shorter than the aggregation cell. The 5 per cent climate-elevated fraction across three grains (0.5°, 1°, 5°) suggests that climate-elevated tiles are spatially compact at 10°-tile resolution and are not absorbed into mixed or information categories under reasonable coarsening.

4.2 What the audit qualifies

The 7 per cent climate-elevated share reported in the baseline is partly an OLS sampling artefact. Bootstrap resampling shows that 4 of the 15 original climate-elevated tiles flip to mixed in more than 20 per cent of iterations, suggesting that the per-tile classification is near the clim/lc = 0.5 threshold for these tiles. A more conservative report sets the climate-elevated fraction at approximately 5 per cent and the most-stable climate-elevated regions to Zambia, the Pilbara, and Russia east of Moscow.

The Q5-only subset shift (mixed +13 pp) shows that the regime distribution is not purely a per-tile-intrinsic property. The entropy stratum matters. This finding is consistent with Dai (2026, §3.6) regional analysis showing climate-block partial R^2 of 6 per cent in the Amazon and Central Africa, where Q5 cells dominate. The all-cell baseline tends to dilute climate's visibility because lower-entropy cells (deserts, monoculture cores) have near-zero LULC variance, and the nested OLS attributes residual variance primarily to climate by default.

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

The Hansen overlay gives only limited forest-loss concordance with the entropy frontier. Most post-2000 frontier cells are agricultural or grassland transitions that the Hansen forest-loss product does not directly capture. The frontier identification by per-cell |ΔH| is therefore broader than deforestation alone and should be read as "high-magnitude land-cover transition" rather than "forest loss specifically." We interpret the low Hansen overlap not as a negative result, but as evidence that the entropy-frontier metric captures a broader class of land-cover transitions — including agricultural and grassland dynamics — that forest-specific products cannot corroborate at the cell level. This highlights both the strength (inclusiveness) and the blind spot (lack of class-specific attribution) of the entropy-based frontier approach.

4.3 Ratio-instability pathology: why the climate-elevated band looks bigger than it is

The Dai (2026, §3.7) classification rule operates on the scalar $r = R^2\_clim / R^2\_lc$, the ratio of unique partial R^2 of the climate and land-cover blocks. Ratio statistics of this form share a general pathology: when the denominator approaches zero the ratio diverges, and a tile is mechanically pushed into the high-ratio band regardless of the absolute size of the numerator. This is a property of R^2-ratio classification rules, not a property of LULC entropy specifically.

In the LULC case the low-LULC-variance cells are concrete and identifiable: deserts, monoculture cores, polar ice-and-rock cells. These have near-singular share vectors. In such cells $R^2\_lc \to 0$ by construction (Δshares are uniformly near zero), and r becomes a small-numerator-over-very-small-denominator quantity with high statistical noise. The §3.2 Q5-only sensitivity confirms this empirically — excluding low-LULC-variance cells lowers the climate-elevated fraction from 6 per cent to 4 per cent. The §3.6 spatial-null finding adds independent evidence: 54 per cent of climate-elevated classifications are reproduced by random spatially-shifted predictors, indicating that the alignment carrying the classification is weak. The §3.4 and §3.7 external corroboration confirms the same point from a third angle. Climate-elevated tiles do not co-locate with intense forest loss, although the matching countries do show measurable national agricultural change in the relevant window.

The climate-elevated label is therefore best read as "the climate block has higher *relative* unique R^2 than the land-cover block in this tile" — a comparative ratio statement — rather than as "climate forcing drives entropy change in this tile". A substantial fraction of climate-elevated classifications reflect denominator collapse rather than a substantive climate signal in the numerator. This is also why we renamed the regime from "energy-driven" (the label used in Dai 2026) to "climate-elevated" throughout this audit. Any future LULC entropy classification using a partial-R^2-ratio rule should similarly distinguish between numerator-driven and denominator-collapse classifications.

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

4.4 Why this matters for LULC entropy methodology

The audit package supports two recommendations for future global LULC entropy regime studies.

First, regime percentages should be reported with a perturbation envelope. Our results suggest that the information band is uncertain by ±10 pp across plausible grain / subset / bootstrap choices, the mixed band by ±10 pp, and the climate-elevated band by ±2 pp. A point estimate without this envelope is over-confident.

Second, external ground-truth validation should accompany frontier-cell claims. The Hansen GFC product is the most readily available external source for the post-2001 window but is restricted to forest transitions. For pre-2001 frontier cells or for agricultural / grassland transitions, alternative external sources (e.g., FAO national statistics, GFW commodity-driven loss attributions, country-level land-use survey records) should be cross-checked.

4.5 Limitations

The 0.25° MAUP test would extend the grain scan from the present 0.5°–5° range to sub-cell resolutions but requires raw HILDA+ 0.01° tiles that are not on disk at submission time. This is identified as priority follow-up work.

The audit uses one threshold rule (clim/lc < 0.1 / 0.1–0.5 / > 0.5) inherited from Dai (2026). Varying the threshold itself — for example shifting the rule to clim/lc < 0.05 / 0.05–0.30 / > 0.30 or to clim/lc < 0.15 / 0.15–0.70 / > 0.70 — is a natural fifth perturbation and would most directly affect the mixed / climate-elevated boundary. We hypothesise threshold perturbation as the next most-influential check after the four tested here and identify it as priority follow-up for the next iteration of the audit package.

The Hansen validation is restricted to post-2001 decades and to forest-loss transitions. A complete external-validation campaign would draw on additional independent datasets for non-forest transitions (cropland abandonment, urban expansion, pasture conversion) that the Hansen forest-cover layer does not capture.

The bootstrap test resamples cells within each tile independently. A more conservative bootstrap would resample at the cell level across the global cell pool and reconstruct the tile partition each iteration; this would propagate uncertainty in tile membership in addition to within-tile OLS uncertainty.

The audit rests on a single anchor paper — Dai (2026, submitted) — which is under review at the time of this manuscript's preparation. If the anchor paper's findings change materially during peer review (for example, if a different regime-classification rule is adopted or the climate block is reformulated), the perturbation results reported here will need to be re-run against the revised baseline. The seven

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

perturbation tests themselves are method-agnostic and will transfer to any updated baseline without modification.

5. Conclusions

The qualitative three-regime classification of global landscape entropy reported by Dai (2026, submitted) survives a seven-perturbation robustness audit, even though the specific numerical partition shifts modestly across perturbations: across all perturbations the information band ranges 43–92 per cent, the mixed band 0–55 per cent, and the climate-elevated band 1–10 per cent. Within ±50 per cent of the reference thresholds the climate-elevated band tightens to 5–7 per cent. The land-cover-dominant majority is stable under all seven perturbations and highly specific under the spatial-null toroidal shift (12 per cent class survival under random shifts). The climate-elevated minority is stable in aggregate proportion but unstable at the individual-tile level (54 per cent spatial-null survival, 4 of 15 tiles flipping under bootstrap resampling), and a substantial fraction of climate-elevated classifications reflects denominator collapse in low-LULC-variance cells rather than a substantive climate signal. The high-entropy subset analysis shows a substantial shift toward the mixed regime, partly as a statistical consequence of conditioning on cells with non-trivial LULC variance. The Hansen Global Forest Change overlay identifies a limited subset of post-2000 frontier cells with documented forest loss; the remainder fall in countries with measurable FAOSTAT national agricultural change between 2000 and 2018. We offer this seven-perturbation audit package as a recommended baseline for future global LULC entropy regime studies; the package and the reporting envelope it generates are method-agnostic and will transfer to any updated baseline classification.

Acknowledgments

AI was used as a coding and writing assistant. All scientific judgments, data interpretation, and final claims are the author's responsibility.

Author contributions

R. D. designed the study, wrote all analysis code, produced all figures, and wrote the manuscript.

Data availability

All analysis scripts and intermediate data products are available from the corresponding author on reasonable request. HILDA+ v2.0 is available from PANGAEA (Winkler et al., 2025). CRU TS 4.09 is available from the University of East Anglia CRU. Hansen GFC v1.12 is available from the University of Maryland. FAOSTAT land-use data are available from FAO (https://www.fao.org/faostat/).

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

References

Bivand, R. S., Wong, D. W. S. (2018). Comparing implementations of global and local indicators of spatial association. *TEST*, 27(3), 716–748. https://doi.org/10.1007/s11749-018-0599-x

Bogaert, J., Farina, A., & Ceulemans, R. (2005). Entropy increase of fragmented habitats: A sign of human impact? *Ecological Indicators*, 5(3), 207–212. https://doi.org/10.1016/j.ecolind.2005.01.005

Brus, D. J. (2021). Statistical approaches for spatial sample survey: Persistent misconceptions and new developments. *European Journal of Soil Science*, 72(2), 686–703. https://doi.org/10.1111/ejss.12988

Curtis, P. G., Slay, C. M., Harris, N. L., Tyukavina, A., & Hansen, M. C. (2018). Classifying drivers of global forest loss. *Science*, 361(6407), 1108–1111. https://doi.org/10.1126/science.aau3445

Cushman, S. A. (2018). Calculation of configurational entropy in complex landscapes. *Entropy*, 20(4), 298. https://doi.org/10.3390/e20040298

Dai, R. (2026). Peri-frontier diversification, regional heterogeneity, and a three-regime global map of land-use mixture entropy change. *Landscape Ecology* (submitted).

Diniz-Filho, J. A. F., Bini, L. M., & Hawkins, B. A. (2003). Spatial autocorrelation and red herrings in geographical ecology. *Global Ecology and Biogeography*, 12(1), 53–64. https://doi.org/10.1046/j.1466-822X.2003.00322.x

Ellis, E. C., Gauthier, N., Klein Goldewijk, K., Bliege Bird, R., Boivin, N., Diaz, S., et al. (2021). People have shaped most of terrestrial nature for at least 12,000 years. *Proceedings of the National Academy of Sciences*, 118(17), e2023483118. https://doi.org/10.1073/pnas.2023483118

Fan, Y., Yu, G., He, Z., Yu, H., Bai, R., Yang, L., & Wu, D. (2017). Entropies of the Chinese land use/cover change from 1990 to 2010 at a county level. *Entropy*, 19(2), 51. https://doi.org/10.3390/e19020051

FAO. (2023). FAOSTAT Land Use domain (cropland and permanent meadows / pastures, area in 1000 ha). Food and Agriculture Organization of the United Nations. https://www.fao.org/faostat/en/#data/RL

Foley, J. A., DeFries, R., Asner, G. P., Barford, C., Bonan, G., Carpenter, S. R., et al. (2005). Global consequences of land use. *Science*, 309(5734), 570–574. https://doi.org/10.1126/science.1111772

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

643 Hansen, M. C., Potapov, P. V., Moore, R., Hancher, M., Turubanova, S. A., Tyukavina,
644 A., et al. (2013). High-resolution global maps of 21st-century forest cover change.
645 *Science*, 342(6160), 850–853. https://doi.org/10.1126/science.1244693

646 Harris, I., Osborn, T. J., Jones, P., & Lister, D. (2020). Version 4 of the CRU TS
647 monthly high-resolution gridded multivariate climate dataset. *Scientific Data*, 7, 109.
648 https://doi.org/10.1038/s41597-020-0453-3

649 Hurtt, G. C., Chini, L., Sahajpal, R., Frolking, S., Bodirsky, B. L., Calvin, K., et
650 al. (2020). Harmonization of global land use change and management for the period
651 850–2100 (LUH2) for CMIP6. *Geoscientific Model Development*, 13(11), 5425–5464.
652 https://doi.org/10.5194/gmd-13-5425-2020

653 Klein Goldewijk, K., Beusen, A., Doelman, J., & Stehfest, E. (2017). Anthropogenic
654 land use estimates for the Holocene — HYDE 3.2. *Earth System Science Data*, 9(2),
655 927–953. https://doi.org/10.5194/essd-9-927-2017

656 Mu, H., Li, X., Wen, Y., Huang, J., Du, P., Su, W., Miao, S., & Geng, M. (2022). A global
657 record of annual terrestrial Human Footprint dataset from 2000 to 2018. *Scientific*
658 *Data*, 9, 176. https://doi.org/10.1038/s41597-022-01284-8

659 Newbold, T., Hudson, L. N., Hill, S. L. L., Contu, S., Lysenko, I., Senior, R. A., et
660 al. (2015). Global effects of land use on local terrestrial biodiversity. *Nature*,
661 520(7545), 45–50. https://doi.org/10.1038/nature14324

662 Openshaw, S. (1984). The Modifiable Areal Unit Problem. *CATMOG 38, Concepts*
663 *and Techniques in Modern Geography*. Geo Books, Norwich, 41 pp.

664 Shannon, C. E. (1948). A mathematical theory of communication. *Bell System*
665 *Technical Journal*, 27(3), 379–423; 27(4), 623–656. https://doi.org/10.1002/j.1538-
666 7305.1948.tb01338.x

667 Tyukavina, A., Potapov, P., Hansen, M. C., Pickens, A. H., Stehman, S. V.,
668 Turubanova, S., et al. (2022). Global trends of forest loss due to fire from 2001 to
669 2019. *Frontiers in Remote Sensing*, 3, 825190.
670 https://doi.org/10.3389/frsen.2022.825190

671 Vranken, I., Baudry, J., Aubinet, M., Visser, M., & Bogaert, J. (2015). A review on the
672 use of entropy in landscape ecology: Heterogeneity, unpredictability, scale
673 dependence and their links with thermodynamics. *Landscape Ecology*, 30(1), 51–65.
674 https://doi.org/10.1007/s10980-014-0105-0

675 Winkler, K., Fuchs, R., Rounsevell, M., & Herold, M. (2021). Global land use changes
676 are four times greater than previously estimated. *Nature Communications*, 12, 2501.
677 https://doi.org/10.1038/s41467-021-22702-2

https://doi.org/10.5194/egusphere-2026-3020
Preprint. Discussion started: 6 July 2026

678 Winkler, K., Fuchs, R., Rounsevell, M., & Herold, M. (2025). HILDA+ version 2.0:
679 Global land use change between 1960 and 2020 [Dataset]. *PANGAEA*.
680 https://doi.org/10.1594/PANGAEA.974335