



MindEarth

VALIDATION REPORT

ME-D10 v4

Validation of a global 10 m digital terrain model

SUMMARY

ME-D10 is a global 10 m bare-earth digital terrain model (DTM) covering all land except Antarctica. This report characterizes the product against airborne LiDAR and photogrammetric reference DTMs from twenty-one countries on six continents, on spatially separated validation and test sets, six external regions, and 252 hydrological catchments.

On a large validation set, ME-D10 reduces the shift-invariant RMSE, which measures terrain-shape error, from 3.22 m to 1.56 m, 1.61 m below the next-best public product (FABDEM). Slope RMSE falls from 4.1° to 2.1° and aspect error from 33.0° to 16.6°, with comparable margins at native 30 m support. Its effective resolution is 46 m, against 101 to 113 m for the public 30 m products. The improvement holds across six external regions, including a continent that contributed no training data (Australia) and two national surveys that played no part in development (the Netherlands and Slovenia). On Slovenia's alpine karst, the hardest region tested, ME-D10 removes 55% of the input's shape error. The improved terrain closes 31% of the drainage-network gap to 1 m LiDAR references in Australian catchments, rising to 48% in the most rugged, and 23% to 40% in four external regions.

1. The product

This report validates ME-D10 v4, a global 10 m bare-earth DTM, against airborne LiDAR and photogrammetric reference DTMs from twenty-one countries, documenting its accuracy relative to public products, its level and datum behavior, its transfer to unsurveyed terrain, and its fitness for hydrological use.

ME-D10 is a 10 m refinement of GEDTM30 [1], an open, globally consistent 30 m DTM. Unlike products that inherit the 2000-era SRTM snapshot or subtract modeled canopy heights from

a radar surface (FABDEM [2]), GEDTM30's bare-earth estimate is anchored to more than 30 billion spaceborne-LiDAR ground detections [1]. It nevertheless retains residual biases near 1 m and limited detail relative to its 30 m grid, and its published accuracy against reference LiDAR DTMs is 2.1 to 7.1 m RMSE, depending on site [1].

ME-D10 sharpens GEDTM30's terrain shape at wavelengths below about 1.3 km and leaves its local mean level unchanged. The output therefore inherits GEDTM30's vertical reference (EGM2008), its absolute level, and its coverage.

2. Validation design

2.1 Reference data

We validate against twenty-two national and regional high-resolution DTMs in twenty-one countries across six continents, all LiDAR-derived except two. The datasets include Denmark at 0.4 m resolution; Spain, Finland, and Portugal at 2 m; Kruger, Mexico, Cape Town, Colombia, and Argentina at 5 m; and the ten others (the United States, Germany, Japan, New Zealand, Norway, Sweden, Austria, Estonia, France, and Chile) at 1 m. Australia spans all three of 1, 2, and 5 m across its state agencies. Only reference data from 2017 onward are used, and datasets whose DTMs visibly retain bridges or trees were excluded. For each survey, acquisition dates were attached at the finest granularity the providing agency supports, from per-tile to per-collection-project. Each reference is mean-aggregated from its native resolution to the 10 m grid. Australia, the Netherlands (0.5 m), and Slovenia (1 m) contributed no data to model development; the latter two are scored as external providers (Section 3.4) rather than drawn into the validation set.

The providers' published vertical accuracies set the measurement floor for every comparison. The LiDAR sources publish 5 to 30 cm figures; Austria publishes a deliberately conservative ± 0.5 m envelope over its merged state campaigns; and the two photogrammetric surveys (Kruger and Colombia) are meter-class.

The surveys also use about twenty distinct vertical references, essentially one national height system per survey. Absolute comparisons against these references therefore mix product error with datum differences, one reason the primary metric scores shape (Section 2.3).

2.2 Separation of validation from development

Each survey is cut into patches on a uniform 1,280 m grid, and 128 km blocks of patches are assigned in a checkerboard to either development or validation, then eroded so that no validation patch lies within 10 km of a development patch (median 26 km). Where an area was surveyed more than once, all vintages fall on the same side of the split. Input-reference error decorrelates within about 1 km, so that separation is an order of magnitude larger. Validation patches are also required to be mostly valid reference, not mostly water, and to carry at least 0.5 m of relief.

The validation set is 50,000 patches, half from Australia and half from held-out blocks of the other surveys. Two further test sets were drawn only after development was complete, and each is scored exactly once: a primary 20,000-patch set at least 5 km from every patch used

in development, validation, or external scoring, and a 14,970-patch far-field set at least 10 km from the same.

Six external regions were also scored with the release configuration, none having informed any development decision: Australia, the Netherlands, and Slovenia contributed no development data at all, while Kruger, Colombia, and Hawaii are held-out areas at least 10 km from any development patch.

2.3 Metrics

Our primary metric is shift-invariant RMSE (SI), the RMSE after each patch's mean error (its bias) is subtracted. SI measures terrain shape, the property ME-D10 changes, and is insensitive to datum differences among the references. We also report standard absolute-accuracy metrics (RMSE, bias, MAE, NMAD), SSIM, LE90 of the shape error (the 90th percentile of per-pixel absolute shape error), and four terrain derivatives: slope, aspect, gradient (the quadrature mean of the two component-gradient RMSEs, in m per m), and Laplacian (4-neighbor stencil). Aspect error is reported only where the reference slope exceeds 5 degrees.

We compare ME-D10 against GEDTM30 and two widely used public elevation products, FABDEM [2] and Copernicus GLO-30 [3], twice: once on the common 10 m grid, with the 30 m products bilinearly upsampled, and again at native 30 m support, with ME-D10 and the reference mean-aggregated onto each product's own grid, so that no product is disadvantaged by resampling. Within each patch the same valid mask applies to every product. These are product benchmarks, not independent-algorithm comparisons.

For confidence intervals we bootstrap 128 km spatial blocks (about 600) rather than individual patches, so spatially correlated errors do not narrow the intervals.

3. Results

3.1 Accuracy against public products

GEDTM30 has a shift-invariant RMSE of 3.22 m; ME-D10 improves it to 1.56 m, a reduction of 1.66 m (95% CI 1.47 to 1.86 m). The public products score 3.16 m (FABDEM) and 3.71 m (GLO-30), leaving ME-D10 1.61 m (1.41 to 1.82 m) below the better of them. Only FABDEM improves on GEDTM30's shape at all, by about 50 mm (95% CI 20 to 80 mm), a thirtieth of ME-D10's margin. GLO-30, a radar surface that retains canopy and buildings, shows the +3.1 m bias expected when it is scored against bare-earth LiDAR, so its robust columns are the meaningful ones.

Table 3. Cross-product comparison on the 50,000-patch validation set.

(a) common 10 m grid	ME-D10	GEDTM30	FABDEM	Copernicus GLO-30
SI-RMSE (m)	1.56	3.22	3.16	3.71
RMSE (m)	2.54	3.79	3.73	7.05
bias (m)	0.17	0.17	0.46	3.13
MAE (m)	1.36	2.00	1.87	3.53
NMAD (m)	1.00	1.30	1.12	1.42
SSIM (higher better)	0.911	0.758	0.856	0.777
slope RMSE (°)	2.08	4.07	4.00	4.05
aspect RMSE (°)	16.6	33.0	32.7	31.9
gradient RMSE (m/m)	0.048	0.089	0.098	0.099
Laplacian RMSE (m)	0.012	0.017	0.023	0.024
LE90, shape (m)	1.71	4.37	4.19	5.46

(b) native 30 m support	ME-D10	GEDTM30	FABDEM	Copernicus GLO-30
SI-RMSE (m)	1.45	3.04	3.01	3.57
RMSE (m)	2.48	3.65	3.60	6.99
bias (m)	0.17	0.17	0.46	3.14
MAE (m)	1.32	1.93	1.80	3.50
NMAD (m)	0.98	1.25	1.09	1.40
SSIM (higher better)	0.926	0.787	0.847	0.799
slope RMSE (°)	0.98	2.32	2.37	2.50
aspect RMSE (°)	7.5	19.0	20.1	20.6
gradient RMSE (m/m)	0.020	0.048	0.049	0.051
Laplacian RMSE (m)	0.002	0.005	0.005	0.005
LE90, shape (m)	1.52	4.11	3.97	5.23

A shape-only comparison would favor ME-D10, which refines shape while the public products do not, so the comparison is broader. ME-D10 posts the lowest error on every metric in Table 3 except bias, which it inherited from GEDTM30 unchanged; this includes the absolute metrics (RMSE, MAE, NMAD) and SSIM. Since ME-D10 leaves the input's level untouched, its absolute-error improvement comes entirely from the sharper shape (Section 3.4). The lead holds across the whole error distribution: the 90th-percentile shape error is 1.71 m against FABDEM's 4.19 m, and ME-D10 removes more than half of GEDTM30's shape error at every percentile up to the 99th. The terrain derivatives show the same margins: slope, aspect, and gradient errors are roughly half the baseline's, and Laplacian error falls by about a quarter at 10 m and by more than half at 30 m, while no public product improves any of GEDTM30's derivatives by more than a few percent.

Re-scoring at native 30 m support changes neither conclusion: the FABDEM shape gap is 1.56 m at 30 m against 1.61 m at 10 m, and derivative errors remain roughly halved.

The test sets confirm the validation result. On the primary test set (5km spacing) ME-D10 scores 1.49 m against an input of 3.66 m; on the 10 km far-field set, 1.31 m against 3.17 m. Error shows no growth with separation at any distance tested.

3.2 Where the improvement is largest

ME-D10's shape error is lower than GEDTM30's on 99.8% of validation patches (median per-patch improvement 0.73 m; even the 90th-percentile patch improves by 0.18 m). Split into terciles of input difficulty, ME-D10 removes 54% of the input's shape error on easy terrain (0.45 to 0.21 m), 59% on medium (1.43 to 0.58 m), and 51% on hard (5.38 to 2.63 m). The removed fraction is nearly flat, so the absolute improvement grows with difficulty, from 0.25 m on the easiest tercile to 2.75 m on the hardest.

The margin over FABDEM grows the same way, from 0.19 m on the easiest tercile to 2.66 m on the hardest, where ME-D10 also roughly halves FABDEM's slope and aspect errors. Stratifying instead by land cover shows the same pattern: ME-D10 leads FABDEM in every GLAD class [4], with the biggest gains over tree cover and built-up areas, and the least over stable cropland. Both stratifications point at one mechanism: the correction is largest where the 30 m products depart most from bare earth and the 30 m grid under-resolves the relief.

3.3 Effective resolution and long-wavelength fidelity

A 10 m posting does not by itself establish 10 m detail. We define effective resolution as the wavelength at which the magnitude-squared coherence between a product and the reference falls to 0.5, measured from Welch cross-spectra over the 10,013 fully valid patches of the primary test set (radially averaged, patch mean removed).

Table 4. Effective resolution.

Product	Effective resolution	Coherence at 90 m	Coherence at 40 m
ME-D10	46 m	0.88	0.33
GLO-30	101 m	0.43	0.00
FABDEM	105 m	0.38	0.00
GEDTM30	113 m	0.28	0.00

ME-D10 resolves structure to about 46 m; every public 30 m product resolves no finer than about 100 m. At 90 m, where GEDTM30's coherence has fallen to 0.28, ME-D10 holds 0.88.

Above the patch scale, the question is whether ME-D10 alters wavelengths it does not refine. We cut 20.48 km squares from co-registered mosaics of ME-D10, GEDTM30, and the reference. Slovenia is the informative region: its reference has 3,000 to 25,000 times more long-wavelength signal power than the input's error in the band. There ME-D10's coherence with the reference is 1.00 at every wavelength from 1.3 to 20 km, identical to the input's, while below the patch scale ME-D10 pulls ahead (0.75 against 0.11 at 90 m), and its error power over 1.3 to 20 km is 61% lower ($n = 20$ squares). ME-D10 therefore changes nothing above the patch scale except to reduce error, and adds no mid-wavelength artifacts.

3.4 Absolute level and vertical datum

By design, ME-D10's per-patch mean offset against the reference is essentially 0: correlation 0.9998, per-patch difference RMS 0.035 m, median absolute difference 5 mm. Both products have a global signed bias of 0.17 m, and both have patch-scale mean offsets with an RMS of

2.01 m. Shape error and offset add as squares to give each product's absolute RMSE (ME-D10: $1.56^2 + 2.01^2 = 2.54^2$; GEDTM30: $3.22^2 + 2.01^2 = 3.79^2$). The standard-RMSE improvement (1.25 m) is smaller than the shape improvement (1.66 m) entirely because of this inherited level error, not because of any shape deficiency, and not a datum artifact. The offset distribution is heavy-tailed and tied to difficulty (median absolute offset 0.62 m, 90th percentile 2.92 m, 99th 7.78 m; correlation with input shape error 0.59). Where GEDTM30's shape is poor its level is poor too, and ME-D10 corrects only the shape.

The same holds in every external region. For each, we measured ME-D10's median signed offset against the local reference and compared it with GEDTM30's (Table 5). ME-D10 matches the input's offset within 6 mm in every region except Colombia, where the difference is 5 cm.

Table 5. External regions: shape error and median signed offset against the local reference.

Region	Patches	GEDTM30 SI (m)	ME-D10 SI (m)	Shape error removed	Offset: GEDTM30 (m)	Offset: ME- D10 (m)
Australia	43,276	2.84	1.55	45%	+0.266	+0.265
Netherlands	24,540	1.16	0.44	62%	-0.303	-0.308
Slovenia	12,509	6.10	2.72	55%	+0.167	+0.173
Kruger	11,352	0.88	0.42	52%	+0.518	+0.513
Colombia	1,745	3.34	2.63	21%	+0.388	+0.435
Hawaii	139	2.38	1.08	55%	+0.027	+0.023

External regions are scored with a single-model variant of the release configuration (area-averaged difference about 15 mm). The Kruger and Colombia references are photogrammetric, with meter-class published accuracy (Section 2.1).

3.5 Transfer to terrain without local reference data

Australia

Withholding Australia from training makes evaluation there a measure of transfer to a continent that contributed no LiDAR supervision. Nearly half (46%) of the Australian validation patches lie in two terrain regimes that scarcely exist outside Australia, a subtropical arid floodplain grassland and a semi-arid valley-floor grassland, 97% and 95% Australian by global extent; the development surveys contain no patches of either and only poor proxies. ME-D10 nonetheless removes 43% of the input's shape error on the Australian validation half and 45% over the wider 43,276-patch Australian region.

To quantify what that missing coverage costs, we retrained the model with 8.5% of its training patches replaced by Australian patches from four regimes the other surveys could not supply. The gain was 25 mm on Australian shape (95% CI 12 to 38 mm), about 1.5% of the 1.66 m total improvement, concentrated on the hardest Australian terrain as expected. Transfer without Australian supervision therefore delivers about 98% of what this targeted coverage achieves, bounding the cost of missing regimes near the surveyed terrain (not of biomes with no surveyed analogue, Section 5); within that scope, the held-out result reads as an estimate of performance on unsurveyed regions.

Netherlands

The Dutch survey tests the low-relief coastal and polder terrain the other regions under-represent. Over 24,540 patches and 92 million valid ground pixels, ME-D10 reduces shape error from 1.16 to 0.44 m (−0.716 m, 95% CI −0.798 to −0.567 m). The Netherlands is one compact region, so the 128 km checkerboard yields only 11 blocks and the interval is correspondingly wide. Mask erosion around the building and tree holes that pervade Dutch bare earth leaves a lower valid fraction than elsewhere, so the product is characterized here on unobstructed ground.

Slovenia

Slovenia's alpine karst is the hardest external terrain in this report, with an input shape error of 6.10 m at full coverage. ME-D10 reduces that to 2.72 m, an improvement of 3.38 m (95% CI of 2.64 to 4.91 m). By input-difficulty tercile, the improvement runs 1.52 m on easy terrain, 2.57 m on medium, and 4.94 m on hard.

Colombia

Colombia is the only tropical region and has the noisiest reference in the suite, with a shape error of 3.34 m for the input and 2.63 m for ME-D10, and ME-D10 posts the best score on every metric. ME-D10 beats the input on 89.5% of patches and reduces median per-patch shape error by 41% (0.75 against 1.19 m). A residual +4 to +8 m bias in three tiles from a single 2023 campaign, confined to ground under canopy the photogrammetric processing did not remove, sets Colombia's measurement floor and is disclosed rather than excluded.

Kruger and Hawaii

Both improve by more than half: Kruger from 0.88 to 0.42 m over 11,352 patches, and Hawaii from 2.38 to 1.08 m over 139 patches.

3.6 Hydrological fidelity

A flood modeler checks where water flows and where it pools, neither of which a shape metric establishes, so we score both directly: channel networks extracted from flow accumulation, and flood-extent footprints derived from height above nearest drainage (HAND). We evaluate these terrain-derived quantities rather than a calibrated flood model, so that differences are attributable to the DEM alone, with no forcing, roughness, or calibration choice between the surface and the score.

The evaluation units are HydroBASINS [5] level-12 catchments, 10 to 120 km² (median 62 km²), each lying entirely within contiguous high-resolution reference coverage (1 m LiDAR, except Kruger's 5 m photogrammetric DTM). Within each catchment we derive channel networks from GEDTM30, ME-D10, FABDEM, and the reference, processing all four identically with WhiteboxTools [6] (least-cost breaching, D8 flow accumulation, channels initiated above 10 ha of upstream area). Agreement with the reference network is scored as F1 within a 20 m buffer and reported as the fraction of the input's remaining gap a product closes, the standard skill score $(F1_{\text{product}} - F1_{\text{input}}) / (1 - F1_{\text{input}})$ [7]. The Australian catchments are the complete census of the 117 that qualify; neither the reference nor the model influenced catchment selection. The Dutch reference network is built from AHN4 with building and tree holes filled from the nearest ground.

Table 6. Drainage-network gap closure (F1, 20 m buffer), ME-D10 and FABDEM against GEDTM30. 95% CIs in brackets.

Region	Catchments	ME-D10	FABDEM
Australia	117	31% [28, 34]	-2% [-4, -1]
Slovenia	81	40% [37, 42]	-8% [-10, -6]
Kruger	31	39% [32, 46]	12% [10, 14]
Netherlands	23	27% [23, 31]	11% [9, 13]
Hawaii	single domain	30%	21%

On the 117 Australian catchments, ME-D10 closes 31% of the drainage-network gap to the reference (95% CI 28% to 34%; 30 to 32% across channel-initiation thresholds of 2 to 50 ha). The median reference-network F1 is 0.545 for ME-D10, 0.388 for GEDTM30, and 0.362 for FABDEM. The gain grows with terrain difficulty, from 19% on the flattest tercile of catchments to 48% on the most rugged, with non-overlapping CIs. FABDEM, processed identically, closes none of the Australian gap and degrades the input's network in Australia and Slovenia; where it does close some gap, ME-D10 closes 2.5 to 3.3 times more (Kruger, Netherlands) or 1.4 times more (Hawaii), so the gain reflects product skill rather than a weak input baseline. Flood extent also improves. First-order inundation footprints, cells whose HAND value falls below a given flood stage, close 23 to 32% of their gap to the reference across stages of 0.5 to 5 m. At the 1 m stage, ME-D10 detects 78% of the reference inundation against the input's 62%, while the false-alarm ratio falls from 33% to 24%.

The benefit holds beyond Australia. Kruger closes 39% of its gap over 31 catchments, and Hawaii closes 23 to 36% across the match-buffer sweep (a single contiguous domain, so no catchment-level CI). The 23 Dutch catchments are drawn from the near-zero-relief polder and delta terrain where vertical fidelity matters most for flood modeling. There ME-D10 closes 27%, rising to 36% (95% CI 30% to 41%) on the engineered polder sub-basins, where sub-meter channel structure is exactly what the coarse input misses. Slovenia adds 81 alpine-karst catchments from a provider never seen in development; there ME-D10 closes 40%, rising from 35% on the flattest tercile to 45% on the most rugged, reproducing the Australian difficulty pattern, and its flood-extent footprints close 22 to 31% of their gap.

These scores measure terrain hydrological fidelity against high-resolution reference DTMs, a precondition for flood modeling rather than a demonstration of flood-model skill. The secondary regions rest on few catchments (23 Dutch, 31 Kruger, 81 Slovenian, no Hawaiian interval at all), so they are supporting evidence around the primary Australian results.

3.7 Buildings, acquisition dates, and observation year

Buildings

Over mapped building footprints [8], ME-D10 changes the GEDTM30 surface less than elsewhere (RMS change 1.75 m against 2.64 m on other land), shows no roof heights, and its error against the LiDAR ground there is below the all-area average (1.55 against 2.58 m).

Reference acquisition date

The reference surveys span 2017 to 2025, against GEDTM30's effective epoch of 2020. The gap does not predict error: per-patch absolute bias grows 0.04 m per year of gap with $r = 0.033$, and binned means are non-monotonic across gaps of 0 to 5 years. ME-D10 tracks the input's

relationship identically, so the temporal gap between reference and input drives neither product's error.

Observation year

The correction is guided by annual satellite observations, so any cell can be generated in any year of the observation record. A terrain model that can be rebuilt per year is a change instrument: differencing two years shows where the ground moved, over the same grid, from the same model, with no survey to commission. Because ME-D10 keeps GEDTM30's local mean level, what responds is within-patch shape, so this tracks erosion, slumping, and earthworks, but can't track low-wavelength regional subsidence.

Our measurements show the output does follow real ground change. Cyclone Gabrielle (New Zealand, 2023) supplies mapped landslide polygons and LiDAR flown both before and after the event, so the real elevation change is known pixel by pixel. Differencing two observation years that bracket the event gives a change field of our own, and the two change fields agree spatially: a correlation coefficient of 0.41 across the most-changed third of cells, and 0.81 over the pixels where the ground was measured by LiDAR to have moved more than 4 m. Differencing two years that both sit before the event, where no ground moved, correlates with the same LiDAR field at 0.00. Each era also fits its own era's ground better, with a shape error of 1.97 m against 2.05 m before the event and 2.00 m against 2.09 m after, so the two years are genuinely two surfaces. Hawke's Bay (New Zealand) and the 2018 Iburi landslides (Japan) reproduce the pattern.

Two numbers set the limits of that reading. Re-running 1,998 unchanged Australian cells with 2017 observations instead of 2025 moves the output by a median of 0.14 m RMS per tile, 16% of the typical correction, so a year-to-year difference smaller than that is likely noise rather than change. Vegetation change alone can also move the output: clear-cut harvest cells move 0.39 m against 0.12 m in matched controls, three times as much, though the movement stays at decimeter scale and inside the harvested stands. A difference is therefore worth reading when it is well clear of a decimeter, and it is a difference in shape only, since both years carry GEDTM30's absolute level and epoch.

4. The released product

The ME-D10 deliverable is a global 10 m map, sectioned by UTM zone. and produced per year (2017 - 2025). Open ocean sits at the EGM2008 zero: mean -0.0002 m, 5th percentile -0.052 m, with a flatness of 0.076 against GEDTM30's 0.221. Inland water bodies are flattened to GEDTM30's level, one level per water body.

Tiles predicted independently can disagree slightly at their edges, so the map is produced in large overlapping blocks with surrounding context and blended across the overlaps. A boundary audit puts the median seam step below 3 cm, with residual anisotropy confined to a single pixel row at 2% excess relief, below terrain noise.

5. Limitations

The reference data are concentrated in wealthy temperate countries. The six external regions, including the tropical Andes, indicate that the correction transfers, but hyper-arid deserts, closed-canopy tropical forest, high mountain Asia, and permafrost terrain have no evaluation of any kind, and the Australian experiment of Section 3.5 bounds the cost of missing regimes near the surveyed terrain, not of biomes with no surveyed analogue. A census of the planet's land surface, clustered into 1,500 surface types, places the development surveys in 462 of them (31%); the other 1,038 hold 75.9% of global land area. But this census flags only coverage, not error: a covered type is not guaranteed accurate, and an uncovered type is not known to be wrong. Additional altimetry checks over the unevaluated biomes is future work.

Spatial blocking during training and validation separates locations but not providers. A survey's processing pipeline and artifact profile can sit on both sides of the split. The Australian holdout controls this for one continent, with provider, datum, and terrain all unseen, and the Netherlands and Slovenia extend the control to two more national providers, on which ME-D10 keeps its shape advantage and adds no offset (Sections 3.4 and 3.5).

ME-D10's absolute accuracy is bounded by GEDTM30's level. Patch-scale mean offsets of 2.01 m RMS pass through unchanged (Section 3.4), and every score is bounded below by the published accuracy of its reference (Section 2.1).

References

- [1] Ho, Y.-F., Grohmann, C. H., Lindsay, J., Reuter, H. I., Parente, L., Witjes, M., & Hengl, T. (2025). GEDTM30: Global ensemble digital terrain model at 30 m and derived multiscale terrain variables. *PeerJ*, 13, e19673. <https://doi.org/10.7717/peerj.19673>
- [2] Hawker, L., Uhe, P., Paulo, L., Sosa, J., Savage, J., Sampson, C., & Neal, J. (2022). A 30 m global map of elevation with forests and buildings removed. *Environmental Research Letters*, 17(2), 024016. <https://doi.org/10.1088/1748-9326/ac4d4f>
- [3] European Space Agency. (2019). Copernicus global digital elevation model (GLO-30). Dataset, Copernicus Programme. <https://doi.org/10.5270/ESA-c5d3d65>
- [4] Potapov, P., Hansen, M. C., Pickens, A., Hernandez-Serna, A., Tyukavina, A., Turubanova, S., Zalles, V., Li, X., Khan, A., Stolle, F., et al. (2022). The global 2000-2020 land cover and land use change dataset derived from the Landsat archive: First results. *Frontiers in Remote Sensing*, 3, 856903. <https://doi.org/10.3389/frsen.2022.856903>
- [5] Lehner, B., & Grill, G. (2013). Global river hydrography and network routing: Baseline data and new approaches to study the world's large river systems. *Hydrological Processes*, 27(15), 2171-2186. <https://doi.org/10.1002/hyp.9740>
- [6] Lindsay, J. B. (2016). Whitebox GAT: A case study in geomorphometric analysis. *Computers & Geosciences*, 95, 75-84. <https://doi.org/10.1016/j.cageo.2016.07.003>
- [7] Wilks, D. S. (2011). *Statistical Methods in the Atmospheric Sciences* (3rd ed.). Academic Press.

- [8] Marconcini, M., Metz-Marconcini, A., Esch, T., & Gorelick, N. (2021). Understanding current trends in global urbanisation: The World Settlement Footprint suite. *GI_Forum*, 9(1), 33-38. https://doi.org/10.1553/giscience2021_01_s33