

NepalDTM: A National 8-Metre Bare-Earth Digital Terrain Model of Nepal from Multi-Source Satellite Data Fusion

Achyut Paudel
achyut.pdl@gmail.com

October 2026

Abstract

NepalDTM is an open bare-earth elevation model of Nepal stored on an 8 m grid. Where the NASA High Mountain Asia 8 m stereo mosaic has a usable height, that mosaic is the surface. Elsewhere the surface is Copernicus GLO-30. A histogram gradient-boosting model, trained on 6,230,548 ICESat-2 points and 7,035,904 GEDI points, predicts how far the surface sits above the ground and subtracts that height. On 4,463,463 held-out ICESat-2 points inside Nepal, with whole 48 km blocks left out of each training fold, the root-mean-square error is 6.10 m (NMAD 2.61 m, bias -0.15 m). The same points give 9.07 m for Copernicus GLO-30, 8.47 m for GEDTM30, and 7.78 m for FABDEM. The error is 1.96 m below 300 m and 8.27 m between 2.5 and 4 km. Independent differential GPS on the Khumbu surface agrees to 2.49 m RMSE at glacier margins. Inside the delivered grid, 37% of the area uses the 8 m mosaic and 63% is Copernicus fill, so the posted spacing is not 8 m terrain detail on every cell. The tiles, the figures, and the code that rebuilds the model are public.

Keywords: NepalDTM; digital terrain model; bare earth; ICESat-2; GEDI; data fusion; Nepal

1 Introduction

Flood maps, landslide assessments, and engineering studies need the elevation of the ground. A digital surface model follows the top of trees, roofs, and snow. A digital terrain model is the bare ground underneath. Nepal has no public high-resolution bare-earth model, and the country runs from the Terai floodplain to the high peaks, so one national file has to serve both flat land and steep slopes.

Copernicus GLO-30 is the free global surface most often used instead (European Space Agency, 2021). It is a 30 m radar model. It remains several metres high where forest and buildings are present, which is the bias measured below. Thirty-metre pixels also omit river channels, terrace walls, and small scars that sit between the posts. The High Mountain Asia 8 m mosaics are a stereo surface of glacier and snow regions, built from commercial optical imagery (Shean, 2017). They are finer, and they are still a surface. FABDEM removes forests and buildings from a global 30 m radar surface (Hawker et al., 2022). GEDTM30 is an ensemble 30 m terrain model fit with ICESat-2 and GEDI (Ho et al., 2025). Because GEDTM30 uses ICESat-2, a score against those points is not an independent test of that model. We still report it, and we mark it as optimistic.

NepalDTM keeps the 8 m mosaic where the mosaic has a finite height, fills the gaps from Copernicus, and subtracts a predicted above-ground residual. The residual is learned from canopy

height (Tolan et al., 2024), land cover (Zanaga et al., 2022), and the shape of the surface, using ICESat-2 ATL08 (Markus et al., 2017; Neuenschwander and Pitts, 2019) and GEDI L2A (Dubayah et al., 2020) as the ground. Heights in the delivered tiles are EGM2008 orthometric metres (Pavlis et al., 2012). The comparison below uses held-out ICESat-2, and a separate differential-GPS survey that none of the models was trained on (Quincey et al., 2016).

2 Study area and product

The delivered grid covers 148,380 km² inside the national outline used to place the tiles. The product is 324 tiles of 24 km on a single UTM zone 45N grid (EPSG:32645), including a buffer around the border so that west Nepal, which lies in zone 44N, does not have a seam at the zone boundary. Each delivered core is 3,000 × 3,000 pixels at 8 m. A 1 km buffer is kept only while predictors are built, and is dropped from the file that is written.

Three rasters are stored for every tile. The surface file is the digital surface model. The bare-earth file is NepalDTM, the surface minus the predicted residual. The third file records the source of the surface, 1 where the HMA mosaic remains and 2 where Copernicus was used to fill, and a predicted absolute error. The code also writes three national VRT mosaics whose sources are the local tile folder.

3 Data

Table 1 lists the inputs. ICESat-2 and GEDI are the only height observations used to fit the correction. FABDEM and GEDTM30 are used only in the comparison. The Khumbu differential GPS is used only as an independent check.

Table 1: Inputs. Citations for each row are given in the text.

Dataset	Spacing	What it is	Role
HMA DEM 8 m	8 m	Stereo surface	Base surface
Copernicus GLO-30	30 m	Radar surface	Fill, and a predictor
ICESat-2 ATL08	~20 m along track	Ground height	Training and the headline test
GEDI L2A	25 m footprint	Ground under canopy	Training, half weight
Canopy height	1 m	Vegetation height	Predictor
WorldCover 2021	10 m	Land-cover class	Predictor
FABDEM V1-2	30 m	Bare-earth model	Comparison only
GEDTM30 v1.2	30 m	Bare-earth model	Comparison only
Khumbu dGPS	points	Ground survey, 2015–2016	Independent check

ICESat-2 points are kept when the height is valid, the segment is snow-free, the beam is a strong beam, and the multiple-scattering warning is at most 1. A threshold on the reported height uncertainty is not applied. In steep terrain that uncertainty grows with slope and would remove the mountain points the test is meant to include. GEDI points are lowest-mode elevations that pass the SlideRule quality and degrade filters, come from a power beam, and have sensitivity above 0.95. Each GEDI point is given half the weight of an ICESat-2 point, because the footprint is larger and the heights are noisier. After those filters, 13,266,452 points lie on the tiles: 6,230,548 from ICESat-2 and 7,035,904 from GEDI, in 2,259 pass groups and 94 blocks of 48 km.

4 Method

4.1 Surface on each tile

The HMA mosaic is stored on the WGS84 ellipsoid. An EGM2008 offset is added before it is compared with anything else. Copernicus GLO-30 is already an orthometric surface. On a tile, mosaic pixels whose difference from Copernicus is more than 40 m from the median difference are removed. Those pixels are blunders from cloud, snow, or a failed stereo match. If more than a thousand mosaic pixels remain, Copernicus fills the holes. Otherwise the whole tile surface is Copernicus. The source flag is 1 where the mosaic remains and 2 on the fill.

4.2 Predictors

Every predictor is built at 8 m. Canopy height is reduced to the cell as the mean, the 75th percentile, and the maximum. Land cover is the majority WorldCover class. From the surface itself the model sees slope, the sine and cosine of aspect, terrain ruggedness, and topographic position over neighbourhoods of about 100 m and 500 m. It also sees the difference between the surface and Copernicus, the source flag, and the surface height. The target is that surface minus the altimeter ground height. Residuals larger than 50 m in absolute value are dropped before the fit.

4.3 Correction model

The learner is scikit-learn’s histogram gradient-boosting regressor (Pedregosa et al., 2011), with absolute-error loss, 500 iterations, a learning rate of 0.05, at most 63 leaves, at least 40 samples in a leaf, and L2 regularisation of 1. The land-cover class is treated as categorical. Each fit uses at most 3,000,000 of the 13,266,452 points. The bare-earth value written in a tile is the surface minus the predicted residual. A second model, fit to ICESat-2 only, predicts the absolute error and is stored on the error band.

4.4 How the scores are computed

Two held-out tests are run before the final fit. One holds out whole satellite passes, so a pass used to test is absent from that fold’s training. The other holds out whole 48 km blocks, in five folds. The block test is the headline, because a pass can still cross a region the model has otherwise seen. The single model written into the tiles is fit after those tests. The national ICESat-2 numbers below are the block out-of-fold residuals, not a rescore of that final model on points it was fit to.

A point enters the comparison with Copernicus, GEDTM30, and FABDEM only when every model has a value there. The national row and the province rows keep only points inside the national outline. That set has 4,463,463 points. Elevation-zone, forest, slope, and mosaic-versus-fill rows use the 6,227,829 points where all four models have a value, including points in the tile buffer outside the outline. The spatial out-of-fold RMSE on all 6,230,548 ICESat-2 points, before that common-support cut, is 5.49 m. That figure is not comparable with GEDTM30 or FABDEM, and it is not the number in Table 2.

RMSE is the root-mean-square error. NMAD is 1.4826 times the median absolute deviation from the median. Bias is the mean error. A positive bias means the model is too high.

5 Results

5.1 National accuracy

Table 2 is the national comparison. NepalDTM’s block-holdout RMSE is 6.10 m, against 9.07 m for Copernicus, 8.47 m for GEDTM30, and 7.78 m for FABDEM. The bias is -0.15 m. Copernicus remains 3.87 m high, which is the trees and buildings still in that surface. Holding out satellite passes instead of blocks changes the NepalDTM RMSE from 6.10 m to 6.01 m. The two tests agree. GEDTM30’s 8.47 m should be read with the caveat above: it was fit using ICESat-2.

Table 2: Held-out ICESat-2 error inside Nepal, in metres, on 4,463,463 points.

Model	RMSE	NMAD	Bias
NepalDTM, 48 km blocks	6.10	2.61	-0.15
NepalDTM, satellite passes	6.01	2.54	-0.13
Copernicus GLO-30	9.07	3.51	$+3.87$
GEDTM30 v1.2	8.47	3.96	$+1.80$
FABDEM V1-2	7.78	3.14	$+1.90$

5.2 Elevation zones

Figure 1 and Table 3 split the 6,227,829-point common set by the elevation of the surface. The bars labelled “Ours (8 m)” are NepalDTM. The lowest error is below 300 m, 1.96 m RMSE, against 4.53 m for Copernicus. The largest error, and the largest gap relative to Copernicus, is between 2.5 and 4 km, 8.27 m against 12.88 m. Above 4 km the NepalDTM RMSE falls to 5.33 m. GEDTM30 is worse than Copernicus in that highest band (8.98 m against 7.55 m). These zone rows include buffer points outside the national outline, so their counts sum to 6,227,829 rather than to the 4,463,463 points in Table 2.

Table 3: Spatial-holdout RMSE in metres by surface elevation, on the 6,227,829-point common set.

Zone	n	NepalDTM	Copernicus	GEDTM30	FABDEM
Below 300 m	1,251,581	1.96	4.53	2.12	2.46
300–1,000 m	487,932	5.18	8.53	5.87	6.16
1–2.5 km	1,076,095	6.24	8.69	7.57	7.33
2.5–4 km	776,214	8.27	12.88	10.45	10.27
Above 4 km	2,636,007	5.33	7.55	8.98	7.50

5.3 Provinces

Figure 2 compares NepalDTM with Copernicus in the seven provinces. These rows do partition the national set: the counts in Table 4 sum to 4,463,463. Every province is lower for NepalDTM than for Copernicus. Madhesh, which is almost entirely lowland, is 1.78 m against 3.88 m. Karnali is the highest provincial RMSE, 6.86 m against 9.88 m. GEDTM30 and FABDEM are in the table and not in the figure. In the high provinces they remain above NepalDTM. In Madhesh, GEDTM30 is 1.98 m and FABDEM is 2.04 m, close to NepalDTM’s 1.78 m and still above it.




Figure 1: Spatial-holdout RMSE by surface elevation. Bar labels are rounded to 0.1 m. Table 3 has the values to 0.01 m. “Ours (8 m)” is NepalDTM.

5.4 Forest

Figure 3 is the canopy correction on the common set. Under forest the RMSE falls from 11.20 m for Copernicus to 6.82 m ($n = 1,616,752$). Off forest it falls from 6.89 m to 4.94 m ($n = 4,611,077$). The correction does more where there is canopy, and it still moves the non-forest error. Under forest, GEDTM30 is 7.84 m and FABDEM is 7.95 m, between NepalDTM and Copernicus. Off forest, FABDEM is 6.78 m, close to Copernicus at 6.89 m, and GEDTM30 is 7.82 m.

5.5 Shape of the error

Figure 4 is the distribution of the spatial residual against the Copernicus residual, drawn from a sample of the point file. The NepalDTM errors sit closer to zero. The Copernicus distribution is shifted positive, the same high bias as in Table 2. The RMSE printed on that figure is computed on the sample, so it is not the 6.10 m national figure.

Figure 5 plots the same residual against elevation, and Figure 6 plots it against slope. The white line is an RMSE trend through the sample. Error grows through the middle elevations, eases in the highest band, and widens once slopes become steep. That is the same pattern as Table 3 and Table 5.

Table 5 and Figure 7 give the slope classes on the full common set, not on the sample. Below 20 degrees the RMSE is 3.26 m. From 20 to 40 degrees it is 6.07 m. At 40 degrees and above it is 10.81 m ($n = 516,657$), against 15.91 m for Copernicus. Part of the steep-slope error is the reference itself: a horizontal miss of a few metres becomes a large height error on a steep face.




Figure 2: Spatial-holdout RMSE by province, NepalDTM (“Ours”) against Copernicus, on points inside Nepal.

Table 4: Spatial-holdout RMSE in metres by province, on the 4,463,463 points inside Nepal.

Province	n	NepalDTM	Copernicus	GEDTM30	FABDEM
Bagmati	423,665	5.92	9.64	8.86	8.23
Gandaki	745,860	5.99	8.63	9.14	7.99
Karnali	1,325,510	6.86	9.88	9.39	8.55
Koshi	640,238	6.30	9.42	9.39	8.55
Lumbini	547,250	5.13	7.57	5.76	5.74
Madhesh	234,163	1.78	3.88	1.98	2.04
Sudurpashchim	546,777	6.28	9.67	7.68	7.54

5.6 Where the 8 m surface exists

Figure 8 splits the common set by the source flag. Where the base surface is the HMA mosaic, the NepalDTM RMSE is 4.00 m ($n = 3,199,580$). Where the base surface is Copernicus fill, it is 6.71 m ($n = 3,028,249$). Copernicus itself scores 6.82 m on the mosaic cells and 9.49 m on the filled cells. The correction cannot add valley and ridge detail that the base surface never contained. Of the area inside the national outline, 37% of the cells use the mosaic and 63% are Copernicus fill. The posted 8 m spacing is the storage grid everywhere. It is the spacing of the terrain shape only on the mosaic share.

5.7 Differential GPS

The comparisons above all use ICESat-2. Figure 9 is a separate survey: differential GPS collected on and beside Khumbu Glacier between October 2015 and October 2016 (Quincey et al., 2016). These points were not used to train the model. The figure scores the delivered *surface* tile, not the




Figure 3: Spatial-holdout RMSE on forest and non-forest cells of the common set. “Ours” is NepalDTM.

Table 5: Spatial-holdout RMSE in metres by slope, on the 6,227,829-point common set.

Slope	<i>n</i>	NepalDTM	Copernicus	GEDTM30	FABDEM
Below 20°	3,163,548	3.26	5.31	4.50	4.22
20–40°	2,547,624	6.07	8.89	8.26	7.56
40° and above	516,657	10.81	15.91	16.67	14.71

bare-earth tile. RMSE is 2.49 m at glacier margins ($n = 83$, bias -1.08 m), 2.96 m on the glacier surface ($n = 388$, bias $+0.86$ m), and 1.02 m on the wind-tower grid ($n = 85$, bias $+0.46$ m).

The bare-earth tile on the same points is a different result. RMSE is 3.89 m at the margins, 6.01 m on the glacier surface, and 6.59 m on the wind-tower grid. On the glacier surface the bare-earth model is 4.85 m low. The correction is treating ice texture as something to remove. Glacier values in the bare-earth file should be read as a corrected surface, not as bedrock.

5.8 Tile edges

Figure 10 compares the step across a tile boundary with the step between neighbouring pixels inside a tile. There are 591 edges. The median absolute step across an edge is 2.08 m. The median step inside a tile is 2.06 m. The 99th percentiles are 13.1 m across and 12.4 m inside. The national mosaic does not add a step at the tile boundary beyond the step the terrain already has. A search for a horizontal shift against ICESat-2, run separately in each province, returns 0 m east and 0 m north. The median of the predicted absolute error stored on the error tile is 4.3 m, and the 90th percentile is 8.8 m.




Figure 4: Error distributions for NepalDTM (“Ours”) and Copernicus, from a sample of the spatial out-of-fold residuals. The RMSE in the legend is the sample RMSE, not Table 2.

6 Limitations

The 8 m spacing is the stored grid. Terrain shape at that spacing exists where the HMA mosaic supplies the surface, 37% of the delivered area. Elsewhere the hills and valleys are still the Copernicus 30 m shape, with a canopy adjustment on top. Above 40 degrees the RMSE exceeds 10 m. Glacier ice is the other weak place: the bare-earth model sits several metres below the 2015–2016 glacier-surface GPS. The predicted-error band is a model of the absolute residual, not a second survey.

7 Data and code

Cloud Optimized GeoTIFF tiles, with the surface, the bare earth, and the predicted absolute error, are at <https://pub-7a51429102ce456f85c0aee664338f8f.r2.dev/cog/>. The web map reads <https://pub-7a51429102ce456f85c0aee664338f8f.r2.dev/NepalDTM.pmtiles>. The code that rebuilds the model, and the validation file this paper quotes, are at <https://github.com/achyut/NepalDTM>. From the code directory, `python run.py` all downloads the inputs, fits the correction, writes the tiles, and runs the validation. A NASA Earthdata token in `~/.earthdata_token` is required and is not in the repository. The code writes three VRT mosaics, `nepal_dem_8m.vrt`, `nepal_dtm_8m.vrt`, and `nepal_aux_8m.vrt`, whose sources are the local `out/` tiles. Figures 1–10 were drawn from that validation file and from the spatial out-of-fold point file by `make_figures.py`. Code is MIT. The NepalDTM tiles are CC BY 4.0. Input datasets keep their own licences.




Figure 5: Spatial out-of-fold residual of NepalDTM against surface elevation. The white line is RMSE in 500 m bins of a sample of the point file.

References

- Ralph Dubayah, James Bryan Blair, Scott Goetz, Lola Fatoyinbo, Matthew Hansen, Sean Healey, Michelle Hofton, George Hurtt, James Kellner, Scott Luthcke, John Armston, Hao Tang, Laura Duncanson, Steven Hancock, Patrick Jantz, Suzanne Marselis, Paul L. Patterson, Wenlu Qi, and Carlos Silva. The Global Ecosystem Dynamics Investigation: High-resolution laser ranging of the Earth’s forests and topography. *Science of Remote Sensing*, 1:100002, 2020. doi: 10.1016/j.srs.2020.100002.
- European Space Agency. Copernicus DEM — global and European digital elevation model, 2021.
- Laurence Hawker, Peter Uhe, Luntadila Paulo, Jeison Sosa, James Savage, Christopher Sampson, and Jeffrey Neal. A 30 m global map of elevation with forests and buildings removed. *Environmental Research Letters*, 17(2):024016, 2022. doi: 10.1088/1748-9326/ac4d4f.
- Yu-Feng Ho, Carlos H. Grohmann, John Lindsay, Hannes I. Reuter, Leandro Parente, Martijn Witjes, and Tomislav Hengl. GEDTM30: a global 30 m digital terrain model and derived multiscale terrain variables. *PeerJ*, 13:e19673, 2025. doi: 10.7717/peerj.19673.
- Thorsten Markus, Tom Neumann, Anthony Martino, Waleed Abdalati, Kelly Brunt, Beata Csatho, Sinead Farrell, Helen Fricker, Alex Gardner, David Harding, Michael Jasinski, Ron Kwok, Lori Magruder, Dan Lubin, Scott Luthcke, James Morison, Ross Nelson, Amy Neuenschwander, Stephen Palm, Sorin Popescu, C. K. Shum, Bob E. Schutz, Benjamin Smith, Yuekui Yang,




Figure 6: Spatial out-of-fold residual of NepalDTM against slope. The white line is RMSE in 5 degree bins.

and Jay Zwally. The Ice, Cloud, and land Elevation Satellite-2 (ICESat-2): Science requirements, concept, and implementation. *Remote Sensing of Environment*, 190:260–273, 2017. doi: 10.1016/j.rse.2016.12.029.

Amy Neuenschwander and Katherine Pitts. The ATL08 land and vegetation product for the ICESat-2 mission. *Remote Sensing of Environment*, 221:247–259, 2019. doi: 10.1016/j.rse.2018.11.005.

Nikolaos K. Pavlis, Simon A. Holmes, Steve C. Kenyon, and John K. Factor. The development and evaluation of the Earth Gravitational Model 2008 (EGM2008). *Journal of Geophysical Research: Solid Earth*, 117:B04406, 2012. doi: 10.1029/2011JB008916.

Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12: 2825–2830, 2011.

Duncan Quincey, Owen King, and C. Scott Watson. Ground control points for investigating mass loss processes on Khumbu Glacier, Nepal. British Geological Survey, National Geoscience Data Centre, accession 75529, 2016. URL <https://webapps.bgs.ac.uk/services/ngdc/accessions/index.html#item75529>. NERC grant NE/L002574/1.




Figure 7: Spatial-holdout RMSE by slope class on the common set. “Ours” is NepalDTM.

David Shean. High Mountain Asia 8-meter DEM mosaics derived from optical imagery, version 1. NASA National Snow and Ice Data Center Distributed Active Archive Center, 2017.

Jamie Tolan, Hung-I Yang, Benjamin Nosarzewski, Guillaume Couairon, Huy V. Vo, John Brandt, Justine Spore, Sayantan Majumdar, Daniel Haziza, Janaki Vamaraju, Theo Moutakanni, Piotr Bojanowski, Tracy Johns, Brian White, Tobias Tiecke, and Camille Couprie. Very high resolution canopy height maps from RGB imagery using self-supervised vision transformer and convolutional decoder trained on aerial lidar. *Remote Sensing of Environment*, 300:113888, 2024. doi: 10.1016/j.rse.2023.113888.

Daniele Zanaga, Ruben Van De Kerchove, Dirk Daems, Wanda De Keersmaecker, Carsten Brockmann, Grit Kirches, Jan Wevers, Oliver Cartus, Maurizio Santoro, Steffen Fritz, Myroslava Lesiv, Martin Herold, Nandin-Erdene Tsendbazar, Panpan Xu, Fabrizio Ramoino, and Olivier Arino. ESA WorldCover 10 m 2021 v200, 2022.




Figure 8: Spatial-holdout RMSE where the stacked surface is the HMA 8 m mosaic and where it is Copernicus fill. “Ours” is NepalDTM.




Figure 9: Khumbu differential GPS against the delivered surface tile. The orange bars are the absolute bias; the signed bias is printed on them.




Figure 10: Absolute elevation step across tile boundaries and inside tiles. Median and 99th percentile.