https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

An ensemble machine learning approach for filling voids in surface elevation change maps over glacier surfaces

Cameron Markovsky[1], Summer Rupper[1], Simon Brewer[1], and Richard R. Forster[1]
[1]School of Environment, Society, and Sustainability, University of Utah, Salt Lake City, UT
Correspondence: Cameron Markovsky (cameron.markovsky@utah.edu)

Abstract. Glacier mass balance assessments in mountainous regions often rely on digital elevation models (DEMs) to estimate surface elevation change. However, these DEMs are prone to spatial data voids, particularly during historical reconstructions using older imagery. These voids, which are most common in glacier accumulation zones, introduce uncertainty into estimates of glacier mass balance and surface elevation change. Traditional void-filling methods, such as constant and hypsometric interpolation, have limitations in capturing spatial variability in elevation change. This study introduces a machine-learning-based approach using gradient-boosted tree regression (XGBoost) to estimate glacier surface-elevation change across voids. High Mountain Asia (HMA) is an ideal study area for assessing the accuracy of different void-filling approaches across glaciers with varying morphology and climatic settings. We compare XGBoost predictions to traditional void-filling methods across the Western and Eastern Himalayas using a dataset of DEM-derived elevation changes. Results indicate that XGBoost consistently outperforms simpler methods, reducing root mean square error (RMSE) and mean absolute error (MAE) while improving alignment with observed elevation changes. The study highlights the advantages of integrating multiple glaciological and topographic predictors, demonstrating the potential of machine learning to improve assessments of glacier mass balance and elevation change. Future research should explore additional predictors, such as climate data, to further enhance predictive accuracy.

1 Introduction

Estimating changes and temporal trends in regional ice mass is critical to understanding the impact of climate change on mountain glaciers. Mass balance assessments illustrate how glaciers respond to past climate variations and inform future changes. The remote and technical setting of many mountain glaciers has historically limited the collection of in situ and airborne data, particularly in regions such as High Mountain Asia and the Andes (World Glacier Monitoring Service (WGMS), 2025). This limited temporal scope in the historical data often necessitates the use of satellite-based methods to estimate regional glacier mass balance. Geodetic glacier mass balance methods use maps of surface elevation change generated by differencing digital elevation models (DEMs) from different timestamps to calculate ice volume change over time, which can be converted to ice mass change over time by adjusting for density (Bamber and Rivera, 2007).

While the complex, steep terrain surrounding mountain glaciers frequently poses challenges for collecting field measurements, it also complicates the generation of remotely sensed products. Spatial data gaps (hereafter referred to as “voids”) are

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

common in many DEMs derived from remotely sensed data. These voids, which are particularly likely in the accumulation zones of glaciers, due to a lack of contrast and increased cloud coverage (McNabb et al., 2019; Kääb, 2008; Berthier et al., 2018), need to be filled before estimating ice mass change. While advancements in optical and radar sensing technologies enable the generation of spatially complete, high-resolution DEMs (Shean et al., 2020), historical reconstructions of ice mass change often rely on older data products. Voids are common in these historical products, and the approach used to fill them can significantly alter estimates of ice mass change (McNabb et al., 2019; Kääb, 2008; Berthier et al., 2018).

Few previous studies have examined the impact of different void-filling methodologies on glacier mass change estimates, with some notable exceptions (e.g., McNabb et al., 2019; Kääb, 2008; Seehaus et al., 2020). However, these studies do not address the impact of different void-filling methods on large, continuous voids on glacier surfaces. McNabb et al. (2019), one of the only studies comparing void-filling techniques to observed elevation changes, artificially generates randomly distributed voids across a glacier's surface. Although these analyses provide insights into best practices for filling voids, it remains unclear how well the findings of this study can be applied to large, potentially unbounded swaths of a glacier's elevation data when filling voids, it remains unclear how well the findings of this study can be applied when large, potentially non-bounded swaths of a glacier's elevation data are missing.

Many traditional void-filling methods fall into two major categories: constant and hypsometric (McNabb et al., 2019). Constant methods often use the mean or median elevation change value for a surrounding radius of filled pixels (Maurer et al., 2019; McNabb et al., 2019) or all non-void pixels on the glacier (McNabb et al., 2019). Hypsometric methods aim to quantify the relationship between elevation and surface elevation change and use this relationship as the basis for void interpolation. Moreover, hypsometric methods can be subdivided into local (glacier-wide) and global (multi-glacier or regional scope) or global (multi-glacier or regional scope) categories (McNabb et al., 2019). These methods often use a binned elevation statistic or a polynomial interpolation to estimate elevation change (McNabb et al., 2019; Seehaus et al., 2020). These methods can either interpolate elevation values in DEMs before differencing or interpolate surface elevation change values after differencing. The latter techniques are considered in this study based on the findings of McNabb et al. (2019).

Recent studies have demonstrated the effectiveness of various machine learning algorithms for estimating parameters related to glacier mass balance (Maffezzoli et al., 2024; van der Meer et al., 2025; Bolibar et al., 2020). Machine learning (ML) offers a compelling alternative to traditional void-filling techniques, as it can integrate multiple variables and capture complex, nonlinear relationships between elevation change and other glaciological parameters. The problem of estimating glacier surface elevation change over voids is similar to those in ice thickness estimation (Maffezzoli et al., 2024), where incomplete data is supplemented using empirically or physically defined models. Both problems require integrating diverse attributes, such as elevation, slope, aspect, and glacier surface area, to infer missing values.

This study estimates surface elevation change values across glacier surfaces with large, continuous spatial data voids using an ML approach, specifically gradient-boosted tree regression. Gradient-boosted tree models have demonstrated strong performance in modeling geospatial and environmental parameters due to their ability to handle missing data, capture complex relationships, provide constraints on model uncertainty, and yield robust predictions with relatively low computational cost.

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

High Mountain Asia (HMA) contains one of the largest concentrations of non-polar glaciers in the world (Shean et al., 2020), and understanding ice mass trends in the region directly relates to variables such as the regional contribution to global sea-level rise and variations in water resource availability. Due to the limited field measurements in the area, estimating long-term trends in glacier mass change across HMA requires historical imagery. Studies have demonstrated the effectiveness of historical satellite imagery in reconstructing past glacial extents and surface elevations across HMA (Maurer et al., 2019; Yang et al., 2023). Specifically, Maurer et al. (2019) use a DEM extracted from Cold War-era (1971–1986) HEXAGON satellites (Maurer and Rupper, 2015). Similar to other historical data, these products are prone to voids. On average, voids account for 37% of glacial area across the Himalaya in historical DEMs generated by Maurer et al. (2019).

We use two regions of HMA as study sites to compare existing void-filling methods with our machine learning method when predicting surface elevation change: the Western and Eastern Himalaya (Figure 1a). These regions serve as ideal study sites since they meet three main criteria: high-quality surface elevation change data exists for recent decades (Shean, 2020), few historical in situ or airborne measurements are available, and both regions contain steep mountain glaciers in complex terrain prone to data voids during historical reconstructions (RGI Consortium, 2017; Maurer et al., 2019). This study aims to examine the accuracy of traditional void-filling methods for mountain glaciers, present a promising new methodology, and provide valuable insights for future glacier mass balance studies.

2 Data

The Western and Eastern Himalaya (Figure 1a) are selected to examine how each method's performance varies in areas with different glacierized areas, varying glacier-wide characteristics, and significant differences in regional glacier mass balance profiles (Figure 1b, c). Shean et al. (2020) reports near 100% coverage of elevation change data for both the Western and Eastern Himalaya, indicating few, if any, gaps over glacier surfaces. This dataset is constructed from void-filled SRTM-GL1 products (Shean et al., 2020) and thus provides a spatially complete dataset for comparing different void-filling methodologies.

2.1 Dataset preprocessing

This study uses data from Shean (2020), derived from satellite stereo photogrammetry, to analyze glacier surface elevation change in HMA. To ensure consistency with these data, we define glacier extents using the Randolph Glacier Inventory (RGI) v6.0 and apply filtering criteria to refine the dataset. Glaciers exhibiting surging activity, as identified by Guo et al. (2023), are excluded to prevent the inclusion of anomalous elevation fluctuations that do not reflect long-term mass balance trends. Similarly, we remove any glaciers with evidence of calving (RGI Consortium, 2017).

Given the significance of debris cover in modulating glacier melt, we further limit our selection to glaciers with available debris cover data from Scherler et al. (2018). The percentage of a glacier's surface covered by debris is calculated and incorporated as a glacier-wide characteristic. Although this study uses the rate of debris-covered area, more detailed debris thickness maps, such as those from McCarthy et al. (2022), could be integrated into future analyses.

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

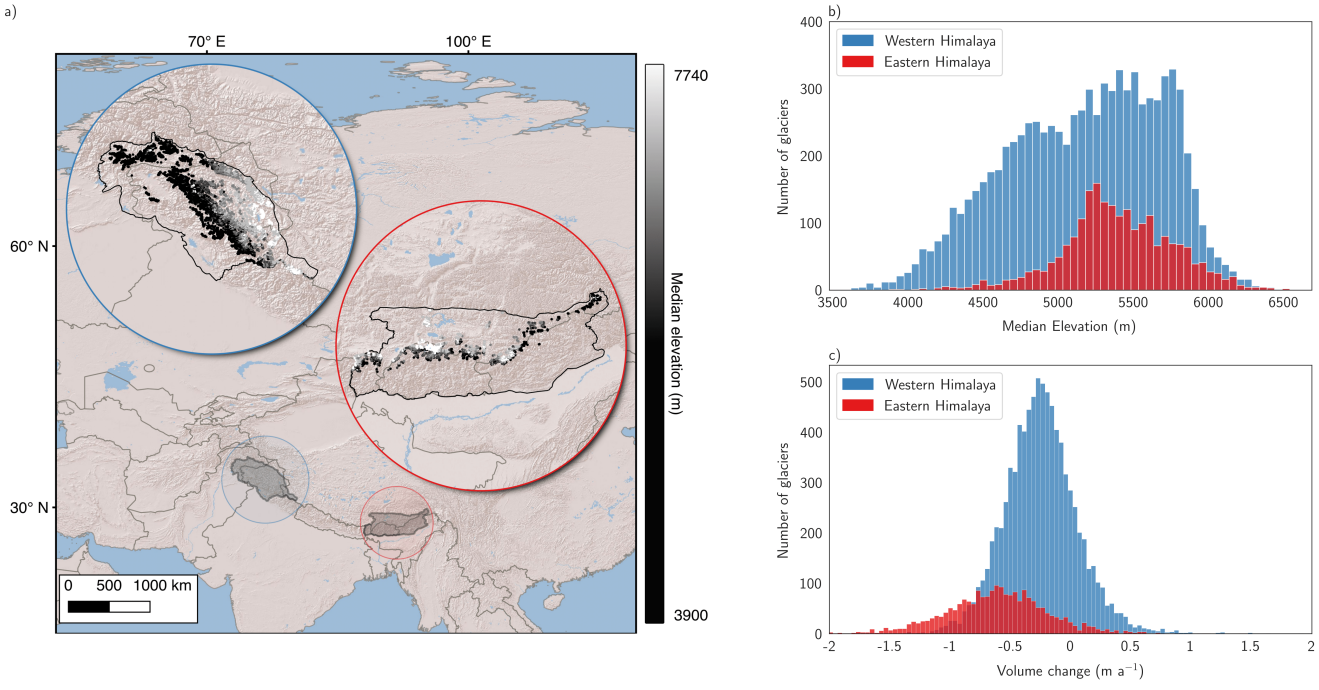

Figure 1. Study area in High Mountain Asia. (a) ESRI global shaded relief map with insets highlighting the Western Himalaya (blue) and the Eastern Himalaya (red). Glacier centroids in each region are plotted and gray-scaled shaded by median glacier elevation. (b) The distribution of glacier median elevation reported in RGI v6.0 for each region. (c) The distribution of glacier mean elevation change reported in Shean (2020) for each region.

To construct a spatially resolved dataset, we clip the Copernicus Global 30m Digital Elevation Model (OpenTopography, 2021) to the outlines of selected glaciers. This DEM is converted into a structured grid of points, each representing a 30 m × 30 m pixel, with surface elevation change values from Shean (2020) sampled at each point (point and pixel are used interchangeably henceforth). Additional glacier-wide attributes, including area, slope, and aspect, are joined to each point (RGI Consortium, 2017). By incorporating these attributes, the dataset accounts for large-scale glaciological trends and localized topographic influences on elevation change.

To improve data reliability, we apply an outlier-filtering process, removing any points where surface-elevation change values exceed three standard deviations from the glacier-wide mean, as reported in Shean et al. (2020). This step helps eliminate potential artifacts introduced during DEM generation and minimize the impact of localized anomalies, such as crevasses or icefall collapses, that do not reflect broader glacier dynamics. This filtering approach follows standard practices in DEM-based glacier change assessments (Maurer et al., 2019), ensuring that every void-filling method uses representative, high-confidence elevation change data.

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

Table 1. Summary statistics for Western and Eastern Himalaya study regions. Number of glaciers and pixels reflect the final, filtered dataset. Glacier area is reported from RGI v6.0. Mean and standard deviation (σ) of surface elevation change values ($\frac{dh}{dt}$) are reported for the period of 2000–2018 from Shean et al. (2020).

Region	Number of glaciers	Number of points	Mean glacier area (km^2)	Mean $\frac{dh}{dt}$ (m a^{-1})	σ $\frac{dh}{dt}$ (m a^{-1})
Western Himalaya	8,659	7,753,768	0.914	-0.249	0.308
Eastern Himalaya	2,172	2,817,783	1.290	-0.609	0.455
Combined	10,831	10,571,551	0.989	-0.321	0.372

2.2 Dataset description

The dataset includes glaciers from the Western and Eastern Himalaya, with the regional extents delineated according to the HIMAP project (Bolch et al., 2019). The Western Himalaya covers approximately 190,000 km², whereas the Eastern Himalaya spans around 165,000 km² (Figure 1a). Despite their comparable spatial extent, the number of usable glaciers and the total number of available pixels differ considerably (Table 1). The Western Himalaya contains roughly four times as many glaciers and over three times as many pixels as the Eastern Himalaya. After data processing and filtering, 12.90% of Western Himalaya glaciers and 19.26% of Eastern Himalaya glaciers were removed from the dataset. The discrepancy in the number of glaciers deemed unsuitable between the two regions is due to the greater prevalence of heavily debris-covered and calving glaciers in the Eastern Himalaya.

Along with differences in glaciated area, the Western and Eastern Himalaya exhibit different glacier elevation distributions (Figure 1b) and mass balance trends (Figure 1c). The Western Himalaya exhibits a regional mass balance of -0.32 ± 0.08 m w.e. a^{-1} for the period 2000–2018, whereas the Eastern Himalaya exhibits a more negative mass balance of -0.52 ± 0.15 m w.e. a^{-1} for the same period (Shean et al., 2020). These values highlight stronger glacier thinning in the Eastern Himalayas as seen in Fig. 1c.

When examining each region at the pixel level, we observed differences in the elevation distributions and surface elevation change trends (Figures 2a, b). The Western Himalaya exhibits a greater concentration of pixels at mid-range elevations, with a peak around 5,000–5,500 m, whereas the Eastern Himalaya has a broader distribution extending to higher elevations. The mean elevation change trends also differ between the two regions. In the Western Himalaya, elevation change remains relatively stable at a slightly negative value at lower elevations, then increases at higher elevations. In contrast, the Eastern Himalaya shows greater variability, with steeper fluctuations in elevation, particularly at the tails of the elevation distribution.

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

3 Methods

3.1 Artificial void generation

Generating an artificial void is necessary to evaluate and compare void-filling methods under controlled yet realistic conditions. Since natural voids in historical digital elevation models (DEMs) are highly variable and often concentrated in glacier accumulation zones, simulating voids that mimic their spatial characteristics ensures that each void-filling method is tested on a consistent and representative basis. Our study's framework allows for the direct assessment of each method's performance and enables validation against observed surface elevation change values.

As previously discussed, the prevalence of voids over glacier accumulation zones informed the generation of artificial voids exclusively in the upper 50% of a glacier's elevation distribution. To generate an artificial void for each glacier, a seed pixel is randomly selected within the upper half of the elevation values. From this seed pixel, a continuous void is radially "grown" to cover approximately 37% of the glacier's total area. This methodology aims to replicate the spatial characteristics of voids observed in historical DEM reconstructions over the Himalaya, as described in Maurer et al. (2019). Depending on the shape of a glacier's outline and its elevation distribution, multiple void areas may be generated to reach 37% of the glacier's area if a single suitable void is not possible. The generated voids mimic the observed voids caused by cloud cover, sensor limitations, and steep terrain (Maurer et al., 2019; Berthier et al., 2018), ensuring that the void-filling methods are tested under fairly realistic conditions. Figure 2 illustrates the distribution of elevation and surface elevation change values for both non-void and artificial void pixels.

3.2 Constant void filling: local mean

The most straightforward void-filling approach assigns a single constant value to all void pixels. Specifically, we use the mean elevation change from all available non-void pixels within each glacier to fill void values. This method assumes uniform glacier-wide thinning or thickening and serves as a baseline for evaluating more sophisticated techniques. The mean elevation change consistently outperformed the median value for nearly all glaciers. This study considers only the mean elevation change, given its superior performance and the findings of McNabb et al. (2019).

3.3 Hypsometric void filling: local binned elevation

We implement a hypsometric approach to account for elevation-dependent variations in glacier surface change. Following the methodology detailed in McNabb et al. (2019), non-void pixels within each glacier are grouped into elevation bins with a default bin size of 50 meters if the glacier's total elevation range exceeds 500 meters. Ten equally spaced bins are used instead for glaciers with smaller elevation ranges. The mean elevation change within each bin is applied to all void pixels in that elevation range. If a void pixel falls outside existing bins, its value is estimated using linear interpolation based on surrounding bins. This method captures the expected elevation-dependent gradients observed in mountain glaciers (Jakob et al., 2021;

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

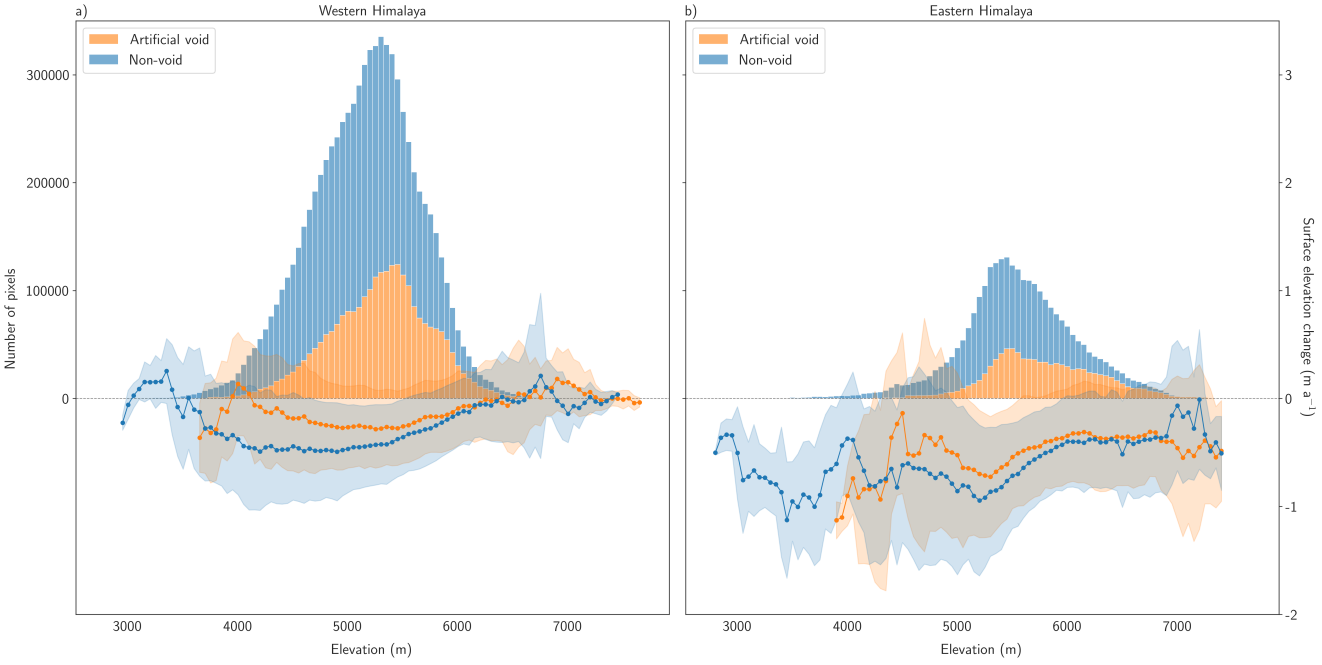

Figure 2. Plots of glacier surface elevation change values relative to elevation for the Western (a) and Eastern Himalaya (b) over the period 2000–2018 (Shean, 2020) (line graph). Elevation is binned using widths of 50 m. Surface elevation change shows the mean value for each bin and the shaded area reflects the standard deviation of surface elevation change for each bin. Number of artificial void (orange) and non-void (blue) pixels within each 50 m elevation bin is shown in the bar graphs.

McGrath et al., 2017). Like the constant method, the mean elevation change was more effective than the median in reducing reconstruction errors and, therefore, is the only one considered in this study.

3.4 Machine learning void filling: boosted tree regression

We employ a novel ML-based void-filling method using gradient-boosted tree regression to estimate surface-elevation change across glacier surfaces. This method is well-suited for modeling complex, nonlinear interactions between glaciological and topographic variables. Machine learning has proven effective in previous studies for estimating glaciological properties such as ice thickness (Maffezzoli et al., 2024), point mass balance (van der Meer et al., 2025), and annual mass balance (Bolibar et al., 2020). We implement this technique using the XGBoost Python library (Chen and Guestrin, 2016), leveraging its efficiency in handling large datasets and its ability to integrate multiple predictors. We train the model using non-void pixels across all glaciers, learning spatial patterns in elevation change to make informed predictions for void pixels.

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

3.4.1 Feature selection

We include the following features to capture spatial, topographic, and glaciological controls on glacier change. Table 2 summarizes and contextualizes these features.

- Spatial coordinates (Easting x, Northing y): Latitude and longitude provide a georeferenced position for each pixel, allowing the model to capture spatial trends in elevation change. At 30 m resolution, these coordinates help contextualize a pixel within its broader environment, including its regional climate and glacier-specific characteristics. Local and regional dependencies arise from variations in topography and climate, making spatial coordinates essential for modeling trends in elevation change.
- Elevation (z): As a primary driver of surface elevation change, elevation directly influences surface elevation change. Incorporating elevation enables the model to quantify hypsometric relationships while interacting with other predictors such as slope, aspect, and debris cover.
- Glacier area ($area$): Area reflects overall glacier size and is roughly proportional to ice mass, influencing its response time to climatic forcing. Including glacier area helps account for variations in surface elevation change trends between small and large glaciers.
- Mean glacier slope ($slope$): Slope is critical for estimating ice flow dynamics and meltwater drainage patterns. Including the mean glacier slope provides a key control on the variability of elevation change.
- Mean glacier aspect (Sine component $\sin_{aspect}$, Cosine component $\cos_{aspect}$): The orientation of a glacier affects the amount of solar radiation it receives, influencing melt rates. Moreover, a glacier's orientation relative to storm and moisture flux directions influences melt and accumulation rates. We decompose the mean glacier aspect into sine and cosine components to treat it as a continuous, rather than circular, variable.
- Maximum, median, minimum glacier elevation (z_{max}, z_{med}, z_{min}): These metrics characterize a glacier's vertical structure. The elevation range contextualizes differences in glacier thinning rates, particularly between lower-elevation zones prone to higher melt and higher-elevation zones prone to lower melt.
- Hypsometric index (HI): This metric, derived from a glacier's maximum, median, and minimum elevations, describes the distribution of ice across the glacier. The HI is calculated using the methodology detailed in McGrath et al. (2017) and provides a base-level assessment of how climatic forcings affect ice mass change. Although the variables used to calculate HI are also model features, we include HI to provide further context about how the relationship between these variables relates to surface elevation change. Given the ML model's ability to capture non-linear relationships among variables, it may be possible to omit either the glacier elevation variables or the HI from the input features in future work.
- Percentage of debris-covered area (dc_{area}): Supraglacial debris influences glacier melt by modifying energy fluxes between the atmosphere and glacier surface. Including the percentage of debris-covered area as a predictor helps account for variations in elevation change linked to supraglacial factors.

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

Table 2. Summary of features used in XGBoost model. Features in bold represent per-pixel parameters, and surface elevation change represents the target variable.

Feature	Variable name	Units	Method	Source
Easting	$\boldsymbol{x}$	**m**	**Sampled**	**COPDEM30**
Northing	$\boldsymbol{y}$	**m**	**Sampled**	**COPDEM30**
Elevation	$\boldsymbol{z}$	**m**	**Sampled**	**COPDEM30**
Glacier surface area	$area$	km^2	Imported	RGI v6.0
Glacier mean slope	$slope$	degrees	Imported	RGI v6.0
Maximum glacier elevation	z_{max}	m	Imported	RGI v6.0
Median glacier elevation	z_{med}	m	Imported	RGI v6.0
Minimum glacier elevation	z_{min}	m	Imported	RGI v6.0
Sine component of mean glacier aspect	$\sin_{aspect}$	rad	Calculated	RGI v6.0
Cosine component of mean glacier aspect	$\cos_{aspect}$	rad	Calculated	RGI v6.0
Hypsometric index	HI	none	Calculated	RGI v6.0
Percentage of debris-covered area	dc_{area}	none	Calculated	Scherler et al., 2018; RGI v6.0
Surface elevation change	$\frac{dh}{dt}$	m a^{-1}	Sampled	Shean, 2020

3.4.2 Model tuning and training

We train our XGBoost model on a single dataset that combines pixels from the Western and Eastern Himalayas. Rather than training separate models for each region, we use a unified model to leverage spatial predictors and capture regional differences in glacier behavior. Initial testing indicated that a unified model performs comparably to individual regional models, supporting the use of a single-model framework.

The training data consists of all non-void pixels ($n = 7,104,981$), which represent approximately 67% of the total dataset. The remaining artificial void pixels are randomly split into equal-sized validation and test sets. We tune the model by minimizing the root mean square error (RMSE) on the validation set. Although we use RMSE for model evaluation, alternative metrics, such as mean absolute error (MAE) or a custom loss function, can be easily incorporated into the framework if desired.

We employ the Bayesian optimization framework Optuna (Akiba et al., 2019) to tune hyperparameters. This method iteratively refines the hyperparameter search space based on prior iterations, reducing computational cost compared to traditional grid search while allowing a broader initial search space. We optimize the model's hyperparameters by training the model over $n = 200$ trials. To prevent overfitting, we implement early stopping rounds. Early stopping rounds halt training if validation performance does not improve after $n = 50$ rounds. This strategy enhances model generalization, aiding its performance on unseen data.

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

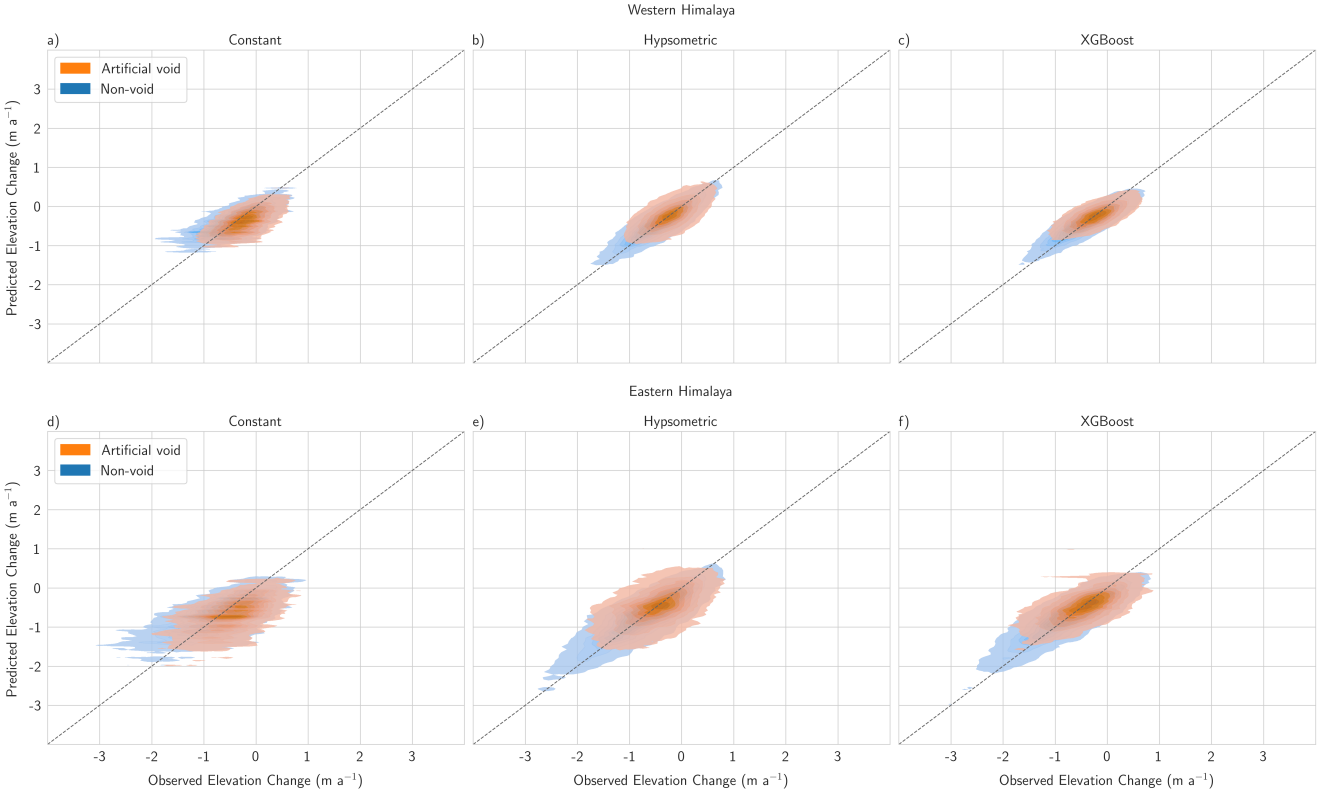

Figure 3. Kernel density estimate (KDE) plots comparing observed and predicted surface elevation change values for different void-filling methods in each region. The top row (a, b, c) displays the results of the Western Himalaya while the bottom row (d, e, f) displays the results of the Eastern Himalaya. Each sub-plot shows the distribution of non-void pixels (blue) and artificial void pixels (orange) with the gray dashed line representing a 1:1 relationship.

4 Results

4.1 Variability of surface elevation change estimates

We compare the root mean square error (RMSE) and mean absolute error (MAE) to assess each void-filling method's performance at the pixel level (Table 3). The results in Table 3 demonstrate that the XGBoost model consistently achieves the lowest RMSE and MAE values across all regions. In the Western Himalaya, XGBoost achieves an RMSE of 0.280 m a^{-1} and an MAE of 0.207 m a^{-1}. The ML model outperforms the constant method by a wide margin and narrowly outperforms the hypsometric method. Although all three methods exhibit higher errors in the Eastern Himalaya, the pattern of performance remains consistent. The XGBoost method performs best with an RMSE of 0.431 m a^{-1} and an MAE of 0.310 m a^{-1}. XGBoost maintains its advantage with the lowest overall error when considering the combined dataset.

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

Table 3. Performance comparison of void-filling methods (Constant, Hypsometric, and XGBoost) based on root mean squared error (RMSE) and mean absolute error (MAE) for non-void and void pixels in the Western and Eastern Himalayas. All reported values are in m a^{-1}.

Region		Constant		Hypsometric		XGBoost	
		RMSE	MAE	RMSE	MAE	RMSE	MAE
Western Himalaya	Non-void	0.304	0.228	0.207	0.153	0.214	0.161
	Void	0.357	0.271	0.317	0.229	0.280	0.207
Eastern Himalaya	Non-void	0.453	0.336	0.320	0.231	0.315	0.229
	Void	0.510	0.392	0.508	0.360	0.431	0.310
Combined	Non-void	0.349	0.256	0.242	0.173	0.245	0.179
	Void	0.405	0.304	0.379	0.265	0.328	0.235

The performance of each method can be further contextualized by examining the distribution of elevation change predictions compared to observed values (Figure 3; Table 3). The constant method exhibits a relatively wide spread of predictions, particularly for artificial void regions, where it fails to capture local variations in elevation change. This method tends to produce lower estimates of surface elevation change. The presence of this bias is consistent with our expectations, given the data's skew towards lower elevations. The hypsometric approach improves predictions slightly, especially in areas where voids align with regions of consistent elevation change trends. The XGBoost model demonstrates the closest alignment with observed values, with a tighter distribution around the 1:1 line and fewer outliers. These results indicate that the ML approach effectively captures local variations and relationships between elevation change and topographic features. The wider spread observed in void regions across all methods suggests greater uncertainty in areas where data were initially missing. However, the reduced error in XGBoost predictions indicates that incorporating multiple predictors improves the accuracy of estimates in these challenging regions.

The comparison of model performance across the Western and Eastern Himalaya (Figure 3) reveals notable spatial differences. In the Western Himalaya, the errors for all methods are relatively symmetrically distributed, suggesting a more uniform glacier response in this region. The XGBoost model performs best, with a lower RMSE and MAE, followed by the hypsometric method, while the constant method remains the least reliable. In the Eastern Himalaya, the error distributions are wider, indicating greater variability in elevation change. The constant method performs particularly poorly in this region, likely due to the more complex topography and variations in glacier dynamics. The XGBoost model still outperforms the other approaches in this region, but exhibits slightly larger errors than in the Western Himalaya, suggesting that additional predictors may be needed to capture the variability in this region fully.

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

4.2 Impact of void-interpolation on glacier-wide estimates

To further evaluate model performance, we analyze the impact of each void-filling method on glacier-wide volume change estimates (Figure 4; Table 4). The constant method exhibits the broadest distribution of errors when compared to observed glacier-wide elevation change values for both the Western (Figure 4a) and Eastern (Figure 4b) Himalaya and expresses a noticeable bias toward both underestimations. The constant method's bias toward negative surface elevation change values is expected due to the removal of the accumulation zone. It is consistent with the per-pixel results (Figures 3a, d).

Table 4. Performance comparison of void-filling methods (Constant, Hypsometric, and XGBoost) on glacier-wide elevation change estimates compared to values reported in Shean (2020). All values are in m a^{-1} except for *Pct. in* σ, which represents the percentage of glaciers with estimates within one standard deviation of the reported value.

Region	Method	Mean ± std. dev.	Median	Minimum	Maximum	RMSE	Pct. in σ
Western Himalaya	Constant	-0.039 ± 0.088	-0.035	-0.788	0.706	0.096	99.54
	Hypsometric	0.004 ± 0.860	0.003	-0.600	0.762	0.086	99.61
	XGBoost	-0.004 ± 0.094	-0.003	-0.793	0.752	0.094	99.42
Eastern Himalaya	Constant	-0.042 ± 0.126	-0.045	-1.249	0.736	0.132	99.59
	Hypsometric	0.012 ± 0.126	0.007	-0.720	0.854	0.126	99.49
	XGBoost	0.012 ± 0.117	0.006	-1.172	0.973	0.117	99.49

The hypsometric method reduces the spread of errors for both regions. Moreover, it is not biased toward negative values of surface elevation change. This adjustment suggests that the hypsometric method better captures surface elevation change trends relative to each glacier's structure. Overall, XGBoost performs similarly to the hypsometric method. However, we observe some differences in performance between regions. For the Western Himalaya, the hypsometric method exhibits slightly smaller error quantities (Table 4), which is reflected in a somewhat narrower error distribution (Figure 4a). In comparison, XGBoost outperforms the hypsometric method in both error quantity (Table 4) and error distribution (Figure 4b).

Although the error distributions of the two methods are tailed towards both positive and negative residuals, we observed better overall performance in the Western Himalaya. This result is consistent with the per-pixel results (Table 3). The increased performance for the Western Himalaya may reflect the larger number of glaciers in the region and the simpler distribution of surface elevation change values (Figure 2a) when compared to the Eastern Himalaya (Figure 2b). The superior performance of XGBoost in the Eastern Himalaya may suggest that the ML model provides more consistent predictions across varying glacier conditions.

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

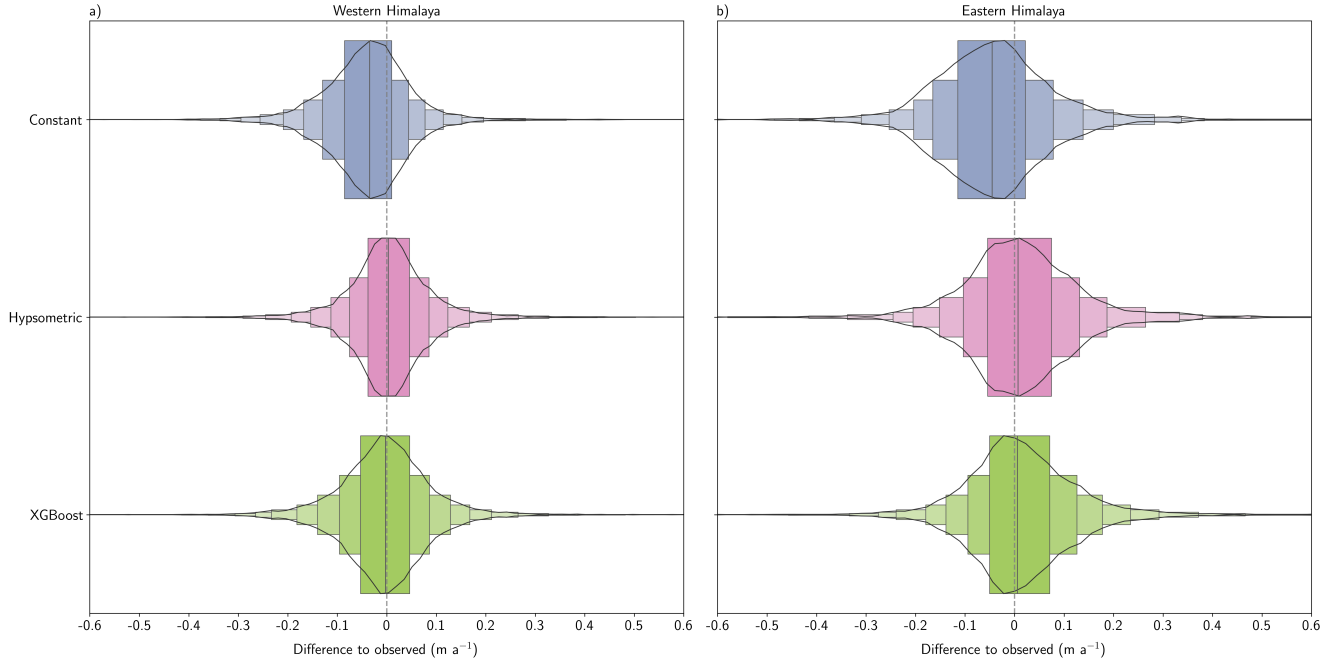

Figure 4. Distribution of residuals (difference between glacier-wide predicted and observed surface elevation change values) for different void-filling methods in the Western (a) and Eastern (b) Himalaya. Violin plots show the spread and density of residuals for each method.

4.3 Model interpretability

Machine learning models, such as XGBoost, are often described as "black boxes" due to the limited information provided on their internal decision-making processes. Although these models are effective at learning complex, non-linear functions, their results can be difficult to interpret because of the opaque decision-making structure of these algorithms. In this study, we use SHapley Additive exPlanations (SHAP) values (Lundberg et al., 2020), which quantify each feature's contribution to the predicted elevation change. These values allow us to directly explore the influence of each input feature on our model output and make the model more transparent.

We use the Toshain (Rupal) Glacier (RGI60-14.20157) in the Western Himalaya (Figure 5) and the Langmusang Glacier (RGI60-15.10755) in the Eastern Himalaya (Figure 6) as case studies to examine the spatial distribution of predictor importance. For both glaciers, the SHAP TreeExplainer (Lundberg et al., 2020) is applied to all artificial void pixels to determine the relative importance of each feature for the predicted surface elevation change. We examine the distribution of features with the two highest SHAP values in the void area and gain insight into regional influences on predictions.

Toshain Glacier, which covers approximately 44 km^2, exhibits negative elevation change values near the tongue of the glacier, and neutral to slightly positive values at upper elevations (Figure 5a). Easting (x) is the feature with the highest SHAP value for the majority of void pixels, followed by elevation (z) and northing (y) (Figure 5b). Spatial coordinates (x, y)

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

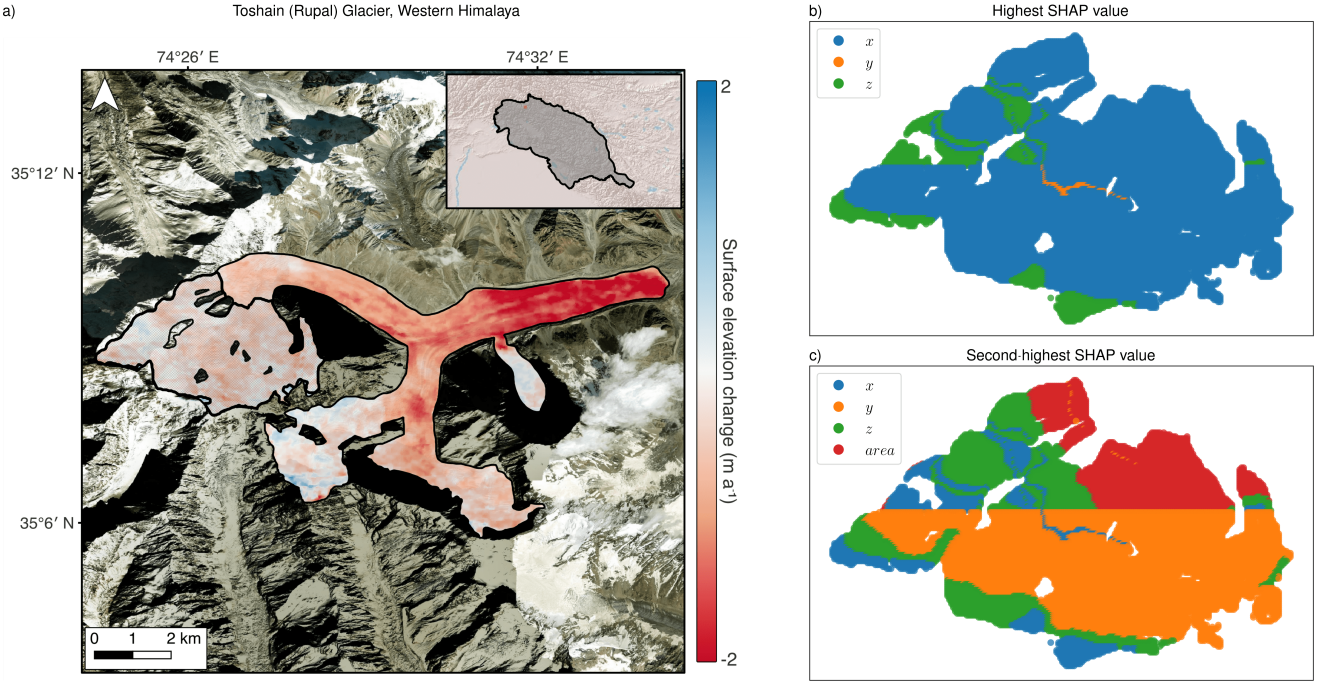

Figure 5. Glacier outline, artificial void area, and surface elevation change distribution for Toshain Glacier (RGI60-14.20157) in the Western Himalaya. (a) ESRI global satellite imagery with the glacier outline from RGI v6.0 and the observed surface elevation change rates across the glacier from Shean (2020), with red indicating thinning and blue indicating thickening. The artificial void area is represented by the bold line and gray shading while the inset map shows the location of the glacier within the Western Himalaya. (b) The predictive feature in the XGBoost model with the highest SHAP value. (c) The predictive feature in the XGBoost model with the second-highest SHAP value.

and elevation (z) comprise the majority of the second-highest SHAP values, along with glacier area ($area$) serving as a key predictor for the lowest-elevation void pixels (Figure 5c).

In comparison, the Langmusang glacier, which covers approximately 32 km^2, exhibits a less gradual elevation change profile. Considerable thinning is observed across the glacier, with areas of localized thickening at higher elevations (Figure 6a). Elevation (z) is the feature with the highest SHAP value for 100% of void pixels on this glacier (Figure 6b). Maximum elevation (z_{max}), median elevation (z_{med}), and easting (x) have the second-highest SHAP values for most void pixels (Figure 6c). Median elevation (z_{med}) is the second most important predictor for a small section of void pixels in the southeast section of the artificial void area.

When examining the second most important predictor, we observe some rough patterns when compared with the surface elevation change map. The area with maximum elevation (z_{max}) as the second-highest SHAP value somewhat mimics the area of surface thinning observed in the upper part of the glacier. Moreover, the small section that values median elevation approximately aligns with the thickening observed in the upper part of the glacier. Although these similarities do not allow us

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

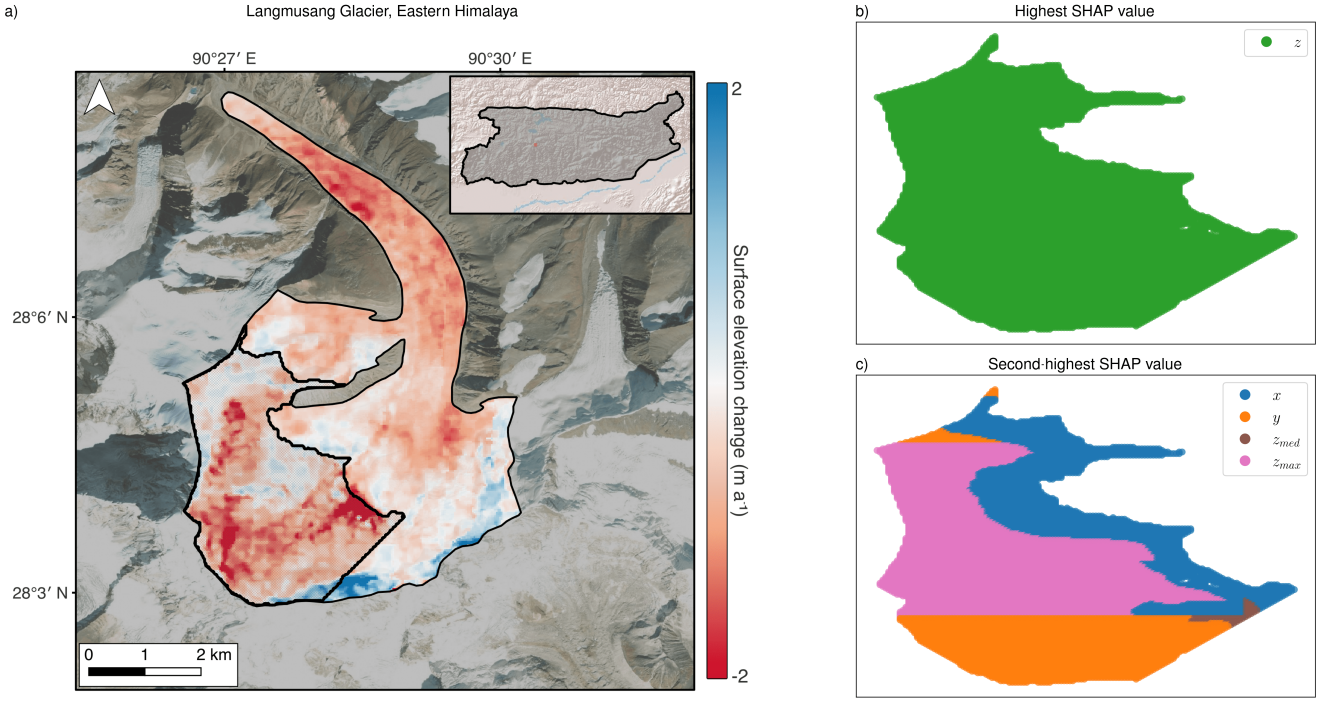

Figure 6. Glacier outline, artificial void area, and surface elevation change distribution for Langmusang Glacier (RGI60-15.10755) in the Eastern Himalaya. (a) ESRI global satellite imagery with the glacier outline from RGI v6.0 and the observed surface elevation change rates across the glacier from Shean (2020), with red indicating thinning and blue indicating thickening. The artificial void area is represented by the bold line and gray shading while the inset map shows the location of the glacier within the Eastern Himalaya. (b) The predictive feature in the XGBoost model with the highest SHAP value. (c) The predictive feature in the XGBoost model with the second-highest SHAP value.

to draw major conclusions, they may provide some insight into how the ML model uses glacier-wide elevation data to identify fine-resolution trends in surface elevation change.

The horizontal lines in SHAP plots represent a diagnostic signature of the underlying structure of the predictors and the model. Due to the piecewise manner in which gradient-boosted trees assign importance, we expect this pattern when using tree-based models like XGBoost. Since decision rules partition the feature space into discrete bins, the model produces repeated patterns of feature influence for all pixels that share the same categorical or binned attribute, resulting in visible horizontal banding.

5 Conclusions

This study evaluates the performance of different methods for estimating glacier surface elevation change in the Himalaya, focusing on traditional approaches (constant and hypsometric) and a machine learning approach (XGBoost). Our results high-

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

light the significant advantages of using data-driven approaches for void filling in glacier datasets, particularly in regions with complex topography and heterogeneous glacier dynamics.

The analysis demonstrates that the constant method, while straightforward, introduces substantial biases in surface elevation change predictions due to its inability to account for local variability in glacier elevation change. The hypsometric approach improves predictions by incorporating elevation-dependent trends but still struggles to capture finer-scale variations, particularly in the Eastern Himalaya, where glacier dynamics are more complex. In contrast, the XGBoost model consistently outperforms these traditional methods, showing a closer alignment with observed values, reduced error distributions, and improved predictive accuracy in both artificial void and non-void regions. Notably, the advantage of using XGBoost is more pronounced when examining surface elevation change predictions on a per-pixel level. The hypsometric method performs nearly as well as XGBoost when integrating predictions across a glacier's surface to estimate glacier-wide surface mass balance. This decrease in comparative performance indicates that our ML method would be the most impactful when examining ice dynamics or other parameters when the spatial distribution of surface elevation change is relevant.

Regional differences in model performance indicate that the Western Himalaya exhibits more uniform glacier responses, leading to lower overall errors across all methods. The Eastern Himalaya, characterized by greater topographic and glaciological complexity, presents a greater challenge for all interpolation-based approaches. XGBoost remains the most reliable predictor for the Eastern Himalaya, albeit with slightly higher errors than in the Western region. These findings suggest that incorporating additional predictors, such as climatic and high-resolution debris cover data, could further enhance model performance.

This study underscores the limitations of traditional void-filling methods and highlights the potential for machine learning models to improve surface elevation change estimates. While XGBoost provides a substantial improvement, residual errors indicate that further refinements, such as the integration of physical glacier process, may enhance predictive accuracy. Future work should explore the application of deep learning and hybrid modeling approaches to further optimize void-filling techniques.

In terms of runtime, the constant and hypsometric interpolation methods are extremely fast, requiring only simple arithmetic operations or bin-based averaging to generate predictions. The constant method is nearly instantaneous while the hypsometric method takes a couple of seconds depending on the number of pixels. In contrast, the XGBoost approach demands greater computational resources, both for hyperparameter tuning and model training. When optimizing and training the XGBoost model, each trial, which consists of 1,000 epochs, took about 20 seconds, resulting in a total training time of approximately an hour. This time is dependent on the available computing resources and could be shortened with further optimization or the choice of a different gradient-boosted tree model. While the improved accuracy and robustness of XGBoost's predictions justifies this added cost, it highlights an important trade-off: XGBoost provides the most reliable void filling performance and at the expense of longer runtimes and the need for a more powerful computing infrastructure. Depending on the number and distribution of voids, the hypsometric void-filling method may provide similar results quicker and without the need for increased computational resources.

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

Another benefit of using XGBoost is the identification of key predictors using SHAP values. Exploring the marginal contribution of each input feature provides valuable information for future glacier modeling work. Specifically, understanding which factors are more critical for filling surface elevation change maps can aid in the feature selection process when constructing emulator models for surface elevation change or surface mass balance. In addition, this knowledge helps to refine future efforts to utilize machine learning algorithms for filling voids in glacier DEMs and surface elevation change maps.

Overall, this research contributes to ongoing efforts to improve the accuracy of glacier mass balance assessments, which are imperative for understanding regional water resources and the impacts of climate change. By demonstrating the advantages of machine learning over traditional interpolation methods, this study provides valuable insights for future remote sensing and glaciological applications.

Code and data availability. The code used to create the dataset, generate and fill artificial voids, as well as the trained models and resulting data for each point, can be found in a Git repository at https://github.com/cmarkovsky/ensemble_void_fill.

Author contributions. CM and SR designed the study; CM developed all code for artificial void generation, void interpolation, and figure creation; CM wrote the manuscript with all other authors contributing edits.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. This work was funded by NASA HiMAT 19-HMA19-0029 and NASA 80NSSC22K1866. Moreover, this work was also funded by NSF 2218664 and 1853881. We would like to thank the University of Utah Center for High Performance Computing (CHPC) for computational resources and providing input on model design.

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

References

Akiba, T., Sano, S., Yanase, T., Ohta, T., and Koyama, M.: Optuna: A next-Generation Hyperparameter Optimization Framework, in: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, Kdd '19, pp. 2623–2631, Association for Computing Machinery, New York, NY, USA, ISBN 978-1-4503-6201-6, https://doi.org/10.1145/3292500.3330701, 2019.

Bamber, JL. and Rivera, A.: A Review of Remote Sensing Methods for Glacier Mass Balance Determination, Global and Planetary Change, 59, 138–148, https://doi.org/10.1016/j.gloplacha.2006.11.031, 2007.

Berthier, E., Larsen, C., Durkin, W. J., Willis, M. J., and Pritchard, M. E.: Brief Communication: Unabated Wastage of the Juneau and Stikine Icefields (Southeast Alaska) in the Early 21st Century, The Cryosphere, 12, 1523–1530, https://doi.org/10.5194/tc-12-1523-2018, 2018.

Bolch, T., Shea, J. M., Liu, S., Azam, F. M., Gao, Y., Gruber, S., Immerzeel, W. W., Kulkarni, A., Li, H., Tahir, A. A., Zhang, G., and Zhang, Y.: Status and Change of the Cryosphere in the Extended Hindu Kush Himalaya Region, in: The Hindu Kush Himalaya Assessment, edited by Wester, P., Mishra, A., Mukherji, A., and Shrestha, A. B., pp. 209–255, Springer International Publishing, Cham, ISBN 978-3-319-92287-4 978-3-319-92288-1, https://doi.org/10.1007/978-3-319-92288-1_7, 2019.

Bolibar, J., Rabatel, A., Gouttevin, I., Galiez, C., Condom, T., and Sauquet, E.: Deep Learning Applied to Glacier Evolution Modelling, The Cryosphere, 14, 565–584, https://doi.org/10.5194/tc-14-565-2020, 2020.

Chen, T. and Guestrin, C.: XGBoost: A Scalable Tree Boosting System, in: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Kdd '16, pp. 785–794, Association for Computing Machinery, New York, NY, USA, ISBN 978-1-4503-4232-2, https://doi.org/10.1145/2939672.2939785, 2016.

Guo, L., Li, J., Dehecq, A., Zhiwei Li, Li, X., and Jianjun Zhu: A New Inventory of High Mountain Asia Surging Glaciers Derived from Multiple Elevation Datasets since the 1970s, https://doi.org/10.5281/ZENODO.7961207, 2023.

Jakob, L., Gourmelen, N., Ewart, M., and Plummer, S.: Spatially and Temporally Resolved Ice Loss in High Mountain Asia and the Gulf of Alaska Observed by CryoSat-2 Swath Altimetry between 2010 and 2019, The Cryosphere, 15, 1845–1862, https://doi.org/10.5194/tc-15-1845-2021, 2021.

Kääb, A.: Glacier Volume Changes Using ASTER Satellite Stereo and Icesat GLAS Laser Altimetry. A Test Study on EdgeØya, Eastern Svalbard, IEEE Transactions on Geoscience and Remote Sensing, 46, 2823–2830, https://doi.org/10.1109/TGRS.2008.2000627, 2008.

Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., and Lee, S.-I.: From Local Explanations to Global Understanding with Explainable AI for Trees, Nature Machine Intelligence, 2, 56–67, https://doi.org/10.1038/s42256-019-0138-9, 2020.

Maffezzoli, N., Rignot, E., Barbante, C., Petersen, T., and Vascon, S.: A Gradient-Boosted Tree Framework to Model the Ice Thickness of the World's Glaciers (IceBoost V1), https://doi.org/10.5194/egusphere-2024-2455, 2024.

Maurer, J. and Rupper, S.: Tapping into the Hexagon Spy Imagery Database: A New Automated Pipeline for Geomorphic Change Detection, ISPRS Journal of Photogrammetry and Remote Sensing, 108, 113–127, https://doi.org/10.1016/j.isprsjprs.2015.06.008, 2015.

Maurer, J. M., Schaefer, J. M., Rupper, S., and Corley, A.: Acceleration of Ice Loss across the Himalayas over the Past 40 Years, Science Advances, 5, eaav7266, https://doi.org/10.1126/sciadv.aav7266, 2019.

McCarthy, M., Miles, E., Kneib, M., Buri, P., Fugger, S., and Pellicciotti, F.: Supraglacial Debris Thickness and Supply Rate in High-Mountain Asia, Communications Earth & Environment, 3, 269, https://doi.org/10.1038/s43247-022-00588-2, 2022.

McGrath, D., Sass, L., O'Neel, S., Arendt, A., and Kienholz, C.: Hypsometric Control on Glacier Mass Balance Sensitivity in Alaska and Northwest Canada, Earth's Future, 5, 324–336, https://doi.org/10.1002/2016EF000479, 2017.

https://doi.org/10.5194/egusphere-2025-4404
Preprint. Discussion started: 18 November 2025

McNabb, R., Nuth, C., Kääb, A., and Girod, L.: Sensitivity of Glacier Volume Change Estimation to DEM Void Interpolation, The Cryosphere, 13, 895–910, https://doi.org/10.5194/tc-13-895-2019, 2019.

OpenTopography: Copernicus GLO-90 Digital Surface Model, https://doi.org/10.5069/G9028PQB, 2021.

RGI Consortium: Randolph Glacier Inventory - a Dataset of Global Glacier Outlines, Version 6, https://doi.org/10.7265/4M1F-GD79, 2017.

Scherler, D., Wulf, H., and Gorelick, N.: Supraglacial Debris Cover, https://doi.org/10.5880/GFZ.3.3.2018.005, 2018.

Seehaus, T., Morgenshtern, V. I., Hübner, F., Bänsch, E., and Braun, M. H.: Novel Techniques for Void Filling in Glacier Elevation Change Data Sets, Remote Sensing, 12, https://doi.org/10.3390/rs12233917, 2020.

Shean, D.: High-Mountain Asia Glacier Elevation Change Trend (Dh/Dt) Map for the Period Spanning 2000 to 2018, https://doi.org/10.5281/ZENODO.3872696, 2020.

Shean, D. E., Bhushan, S., Montesano, P., Rounce, D. R., Arendt, A., and Osmanoglu, B.: A Systematic, Regional Assessment of High Mountain Asia Glacier Mass Balance, Frontiers in Earth Science, 7, https://doi.org/10.3389/feart.2019.00363, 2020.

van der Meer, M., Zekollari, H., Huss, M., Bolibar, J., Sjursen, K. H., and Farinotti, D.: A Minimal Machine-Learning Glacier Mass Balance Model, The Cryosphere, 19, 805–826, https://doi.org/10.5194/tc-19-805-2025, 2025.

World Glacier Monitoring Service (WGMS): Fluctuations of Glaciers (FoG) Database, https://doi.org/10.5904/WGMS-FOG-2025-02B, 2025.

Yang, L., Zhao, G., Mu, X., Liu, Y., Tian, P., Puqiong, and Danzengbandian: Historical and Projected Evolutions of Glaciers in Response to Climate Change in High Mountain Asia, Environmental Research, 237, 117 037, https://doi.org/10.1016/j.envres.2023.117037, 2023.