

Boundary Aware U-Net for Glacier Segmentation

Bibek Aryal^{*1}, Katie E. Miles², Sergio A. Vargas Zesati¹, and Olac Fuentes¹

¹The University of Texas at El Paso, Texas, USA

²Aberystwyth University, Aberystwyth, UK

Abstract

Large-scale study of glaciers improves our understanding of global glacier change and is imperative for monitoring the ecological environment, preventing disasters, and studying the effects of global climate change. Glaciers in the Hindu Kush Himalaya (HKH) are particularly interesting as the HKH is one of the world's most sensitive regions for climate change. In this work, we: (1) propose a modified version of the U-Net for large-scale, spatially non-overlapping, clean glacial ice, and debris-covered glacial ice segmentation; (2) introduce a novel self-learning boundary-aware loss to improve debris-covered glacial ice segmentation performance; and (3) propose a feature-wise saliency score to understand the contribution of each feature in the multispectral Landsat 7 imagery for glacier mapping. Our results show that the debris-covered glacial ice segmentation model trained using self-learning boundary-aware loss outperformed the model trained using dice loss. Furthermore, we conclude that red, shortwave infrared, and near-infrared bands have the highest contribution toward debris-covered glacial ice segmentation from Landsat 7 images.

1 Introduction

Glacier delineation using remote sensing imagery has seen a growing use of deep learning in recent years [5, 14, 15, 34, 36]. This can be attributed to factors such as the availability of large-scale remote sensing data from multiple sources, the development of state-of-the-art deep learning architectures for image analysis, and the growing interest due to the impacts of climate change on

*Corresponding Author: baryal@miners.utep.edu

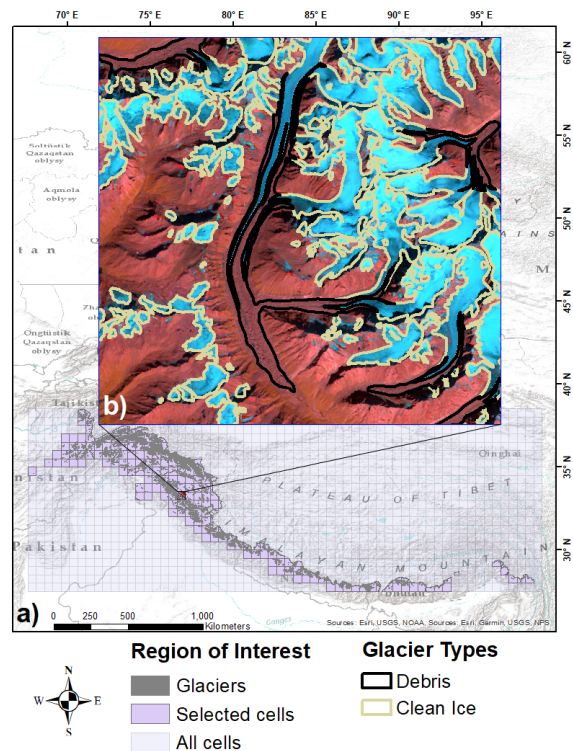


Figure 1: a) Spatially non-overlapping regions using fishnet grid. b) A zoomed image of one of the cells showing clean ice and debris glacier labels.

glaciers in recent decades. The Himalaya is one of the world's most sensitive regions to global climate change, with impacts manifesting at particularly rapid rates [17, 19]. Unsurprisingly, much research has been focused on mapping glaciers in the Hindu Kush Himalayas (HKH) [5, 6, 36]. Several studies have reported the performance of clean glacial ice and debris-covered glacial ice mapping in the HKH; however, most research has been focused on specific

glacier basins within the region and not across the region as a whole.

Glaciers form when snow compresses under its own weight and hardens over long timescales [12]. Near its formation, glacial ice has snow or ice surface cover and is known as clean glacial ice. As glacial ice slowly moves down valleys under gravity, avalanches can deposit debris (rocks and sediment) on top of the glacier. Glacial ice having a significant covering of dirt/rocks/boulders is known as debris-covered glacial ice. Clean glacial ice and debris-covered glacial ice appear differently in remote sensing imagery. The challenge lies in differentiating clean glacial ice from temporary snow/ice cover and debris-covered glacial ice from moraines and the surrounding valleys. The spectral uniqueness of clean glacial ice compared to surrounding terrain makes it relatively easy to identify and localize. However, the delineation of debris-covered glacial ice poses significant challenges because of its non-unique spectral signature.

The earliest approaches for debris-covered glacial ice segmentation involving deep learning used multilayer perceptrons to estimate the supraglacial debris loads of Himalayan glaciers using pre-defined glacier outlines [7, 8]. More recent methods for glacier segmentation use Convolutional Neural Networks (CNNs) due to their success in image-based applications [5, 25, 36, 39]. Most recent approaches to learning-based image segmentation use variants of the U-Net architecture [30]. Originally introduced for biomedical image segmentation, the U-Net has seen successes in numerous applications involving satellite image segmentation [23, 29, 38]. Moreover, different architectures based on the U-Net have also been used for glacier segmentation in recent years [5, 36]. However, unlike segmentation for general images, the results when it comes to glacier segmentation are not very good, particularly for debris-covered glacial ice.

Here we present a variant of the U-Net and train it using multispectral images from Landsat 7 as inputs. Researchers have shown that the performance of deep learning models can be improved by learning multiple objectives from a shared representation [11]. Early approaches to learn multiple tasks use weighted sum of losses, where the loss weights are either constant or manually tuned [13, 21, 31]. We propose a method to combine two different loss functions - masked dice loss [2] and boundary

loss [9] - to simultaneously learn multiple objectives automatically during the training process for improved performance.

While deep learning models have been shown to perform well on various tasks involving computer vision, the interpretability of these models is limited. Deep neural networks are often considered black boxes, since their decision rules can not be described easily. Unlike coefficients and decision boundaries of simpler machine learning methods such as linear regression and decision trees, weights of neurons in deep neural networks can not be understood as knowledge directly. The development of transparent, understandable, and explainable models is imperative for the wide-scale adoption of deep learning models. Over the years, many have proposed different approaches to describe deep learning models [22, 37, 39]. One of the most widely used methods to envision which pixels in the input image affect the outputs the most is by visualizing saliency maps [33]. A saliency map is obtained by calculating the gradient of the given output class with respect to the input image by letting gradients backpropagate to the input. In the case of multispectral or hyperspectral images, spectral saliency [20] is used to visualize salient pixels of an image. Image saliency maps, computed independently for all channels on a multispectral image, can be used to visualize the contribution of each pixel in each channel toward the final output. We propose a method to quantify each channel's contributions towards the final label in the context of glacier segmentation using Landsat 7 imagery.

2 Dataset and Methodology

The HKH region covers an area of about 4.2 million km^2 from about 15° to 39° N latitude and about 60° to 105° E longitude extending across eight countries consisting of Afghanistan, Bangladesh, Bhutan, China, India, Myanmar, Nepal, and Pakistan [4]. The geographic extent of the glaciers within the HKH, however, ranges from about 27° to 38° N and about 67° to 98° E (Figure 1).

We downloaded the Landsat 7 images used for label creation using Google Earth Engine. Landsat 7 contains the Enhanced Thematic Mapper Plus (ETM+) sensor which captures multiple spectral bands as shown in Table 1. The thermal infrared

bands were upsampled from 60 meters to 30 meters resulting in all bands having a spatial resolution of 30 meters. The glacier outlines (labels) [3] were downloaded from International Centre for Integrated Mountain Development (ICIMOD) Regional Database System. (<http://rds.icimod.org/Home/DataDetail?metadataId=31029>) The glacier labels contain information on clean-ice and debris-covered glaciers in the HKH for regions within Afghanistan, Bhutan, India, Nepal, and Pakistan. The ICIMOD glacier outline labels used in this research were derived using the object-based image classification methods separately for clean-ice and debris-covered glaciers and fine-tuned with manual intervention [4].

Table 1: Landsat 7 bands description

Name	Description
B1	Blue
B2	Green
B3	Red
B4	Near Infrared
B5	Shortwave Infrared 1
B6_VCID_1	Low-gain Thermal Infrared
B6_VCID_2	High-gain Thermal Infrared
B7	Shortwave Infrared 2

The Landsat 7 images that were used for delineating glacier labels [4] overlap spatially. To avoid spatial overlap between train and test regions, we created polygon features representing a fishnet of rectangular cells for the entire geographical region. We then created a mosaic of all Landsat 7 images used for labeling into a single raster and clipped the raster mosaic to country boundaries for glacier labels (Figure 1) to avoid false negative glacier labels in the dataset. Finally, we discarded the rasters within the polygon cells that do not contain any glacier labels and downloaded clipped regions within selected cells. The Google Earth Engine code to replicate this process can be found in repository https://code.earthengine.google.com/?accept_repo=users/bibekarya17/get_hkh_tiff. The selected cells were then randomly split into train, validation, and test sets with no geospatial overlap. 1163 out of 1364 cells were filtered out to leave us with 141, 20, and 40 cells in the training, validation, and test sets respectively. Each cell

was then cropped into multiple sub-images of 512×512 pixels and the sub-images with less than 10% of pixels as glacier labels were discarded to reduce class imbalance. These sub-images are then normalized and provided as input to the model. There are 333, 68, and 98 sub-images in the training, validation, and test sets respectively. Every pixel within each sub-image can have one of four different classes as can be seen in Figure 2.

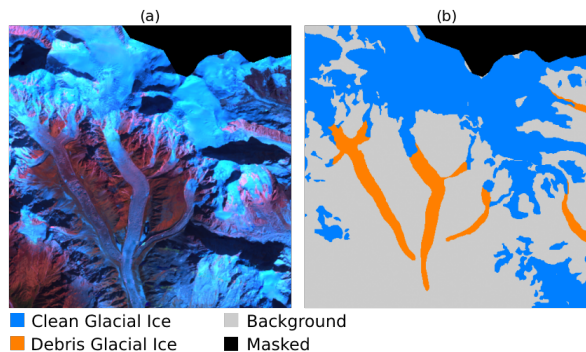


Figure 2: (a) Sample sub-image, (b) Corresponding Clean Glacial Ice, Debris Glacial Ice, Background, and Masked labels

The step-by-step processing we followed to prepare input features for the model is shown in Figure 3. The distribution of pixels for train, validation, and test set across different classes is shown in Table 2 and highlights that the distribution of pixels across different sets is similar and labels are heavily imbalanced across classes.

Table 2: Labels Distribution - Random Sampling

split	background	clean	debris	masked
train	72.44%	21.77%	2.44%	3.35%
val	68.69%	23.22%	3.24%	4.85%
test	70.16%	22.97%	2.65%	4.21%

clean = clean glacial ice

debris = debris-covered glacial ice

We used a modified version of the U-Net architecture [30] as shown in Figure 4. Each input sub-image is 512×512 pixels in size. Zero padding was added during each convolution operation to make the output labels the same size as input sub-images. We replaced the Rectified Linear Unit (ReLU) in the original U-Net architecture with Gaussian Er-

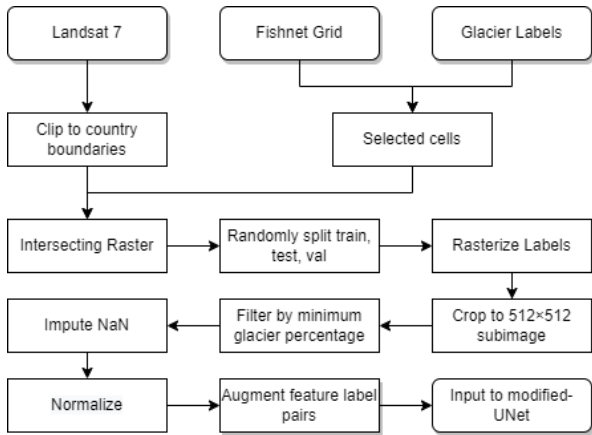


Figure 3: Input preprocessing

ror Linear Units (GELU) [16]. We applied batch normalization after each convolution operation and spatial dropout [35] of 0.1 after each down-sampling and up-sampling block to reduce overfitting. We also randomly modified 15% of the training samples by either rotating (90°, 180°, 270°) or flipping (horizontal/vertical) the input sub-images to the model. We trained the modified U-Net architecture for 250 epochs using the Adam optimizer and evaluated the performance based on precision, recall, and Intersection over Union (IoU).

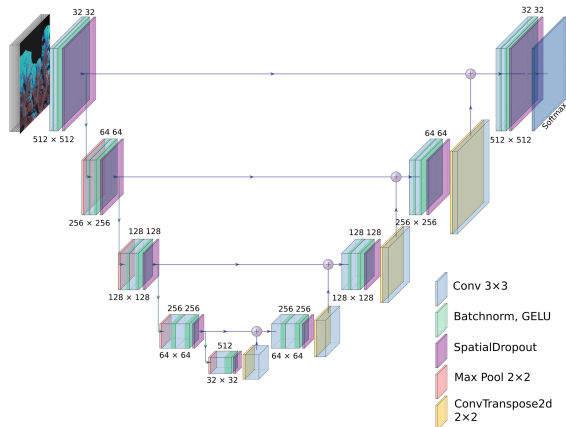


Figure 4: Our modified U-Net architecture has 32 feature maps in the first convolution layer. We also introduce batch normalization and spatial dropout in the modified architecture.

We trained two separate models, one for segmenting clean glacial ice and one for debris-covered glacial ice, and combined the outputs to produce the final segmentation map. Definitions of what constitutes debris-covered glacial ice vary widely, however, as a glacier does not have to be fully covered by debris to be classified as debris-covered glacial ice [24]. Therefore, for the pixels where debris-covered glacial ice labels overlapped with clean glacial ice labels on the final segmentation map, the output label was set as debris-covered glacial ice. The code to replicate our process can be found in the GitHub repository (https://github.com/Aryal007/glacier_mapping.git).

3 Experiments

3.1 Self-learning Boundary-aware Loss

The subject of this section of our work lies at the intersection of two branches of research, which are penalizing misalignment of label boundaries by using a boundary-aware loss and learning multi-task weights during the training process. We propose a combined loss ($\mathcal{L}_{Combined}$) that is a weighted sum of masked dice loss (\mathcal{L}_{MDice}) and boundary loss ($\mathcal{L}_{Boundary}$), as described in Equation 1. We also compare the performance of our methods using the modified U-Net to the standard U-Net trained on cross entropy loss (\mathcal{L}_{CE}).

$$\mathcal{L}_{Combined} = \alpha \times \mathcal{L}_{MDice} + (1 - \alpha) \times \mathcal{L}_{Boundary} \quad (1)$$

The value of hyperparameter α can be set manually between 0 and 1. Having an α of 1 is equivalent to training the model exclusively using masked dice loss and an α of 0 is the same as training the model exclusively using boundary loss. However, tuning the value of α manually for the best results is an expensive process. In order to learn the weights for $\mathcal{L}_{Boundary}$ and \mathcal{L}_{MDice} through backpropagation, we initially set α to 0.5 and let the model find the best value of α . However, we observe that without any constraints on the value of α , the network updates α such that $\mathcal{L}_{Combined}$ is minimized without necessarily having to minimize \mathcal{L}_{MDice} or $\mathcal{L}_{Boundary}$. This results in poor performance. Inspired by [18] for weighing two different loss func-

Table 3: Performance comparisons between standard U-Net trained using cross entropy loss and modified U-Net trained using combined loss and self-learning boundary-aware loss

Loss (\mathcal{L})	$\mathcal{L}_{weight(s)}$	clean			debris		
		Precision	Recall	IoU	Precision	Recall	IoU
\mathcal{L}_{CE}	—	89.39%	84.65%	70.16%	0.00%	0.00%	0.00%
$\mathcal{L}_{Combined}$	0	68.40%	0.25%	0.25%	3.00%	99.75%	3.00%
$\mathcal{L}_{Combined}$	0.1	79.82%	79.66%	66.31%	50.93%	49.30%	33.43%
$\mathcal{L}_{Combined}$	0.5	81.60%	80.77%	68.33%	51.16%	46.92%	32.41%
$\mathcal{L}_{Combined}$	0.9	81.60%	80.77%	68.33%	51.16%	46.92%	32.41%
$\mathcal{L}_{Combined}$	1	80.31%	80.65%	67.34%	46.00%	44.10%	29.05%
\mathcal{L}_{SLBA}	<i>Dynamic</i>	81.59%	80.55%	68.17%	51.97%	53.81%	35.94%

*clean = clean glacial ice; debris = debris-covered glacial ice

tions, we propose Self-Learning Boundary-Aware loss (\mathcal{L}_{SLBA}) that is a combination of \mathcal{L}_{MDice} and $\mathcal{L}_{Boundary}$.

$$\mathcal{L}_{SLBA} = \frac{1}{2\alpha_1^2} \times \mathcal{L}_{MDice} + \frac{1}{2\alpha_2^2} \times \mathcal{L}_{Boundary} + |\ln(\alpha_1 \times \alpha_2)| \quad (2)$$

In the case of \mathcal{L}_{SLBA} , α_1 and α_2 both are initially set to 1 and we let the model find the best value for α_1 and α_2 through backpropagation. In Table 3 we show performance for different values of α in the case of $\mathcal{L}_{Combined}$ and performance of \mathcal{L}_{SLBA} . One advantage of using \mathcal{L}_{SLBA} over $\mathcal{L}_{Combined}$ is that there is no extra hyperparameter that requires fine-tuning. All experiments in Table 3 use eight features from Landsat 7 imagery as inputs.

best for debris-covered glacial ice segmentation and eliminates the need to fine-tune loss weights. We can also see that the model fails to converge when training solely on boundary loss ($\alpha = 0$) and training on glacier boundaries by incorporating boundary loss along with masked dice loss results in an overall improvement in performance for debris-covered glacial ice regardless of the weighting factor. Figure 5 shows the weights for masked dice loss ($\frac{1}{2\alpha_1^2}$) and the weights for boundary loss ($\frac{1}{2\alpha_2^2}$) vs. epoch during training for debris-covered glacial ice. The optimal values for α_1 and α_2 are calculated to be 0.9569 and 1.045 for clean glacial ice segmentation and 0.952 and 1.05 for debris-covered glacial ice segmentation for \mathcal{L}_{SLBA} .

3.2 Representation Analysis

To understand the contribution of each feature in the multispectral image toward the final label, we computed a Saliency Score (SS) for each feature by summing all pixels in the Saliency Map (SM) for that feature.

$$SS_{feature} = \sum_{i=0}^{c-1} \sum_{j=0}^{r-1} SM_{feature}(i, j) \forall feature \in Input \quad (3)$$

where:

r, c = number of rows, columns in saliency map

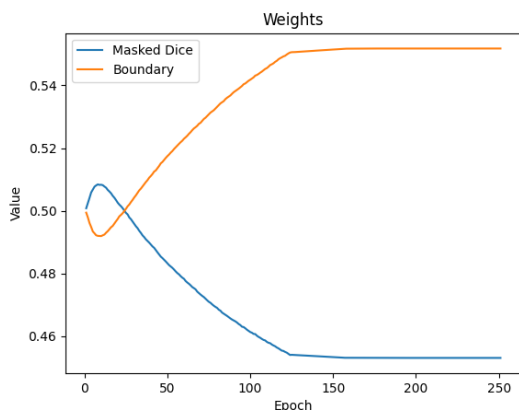


Figure 5: Masked Dice Loss weights and Boundary Loss weights vs. epoch for debris-covered glacial ice

From Table 3, we see that \mathcal{L}_{SLBA} performs the

Average feature saliency scores across all the images in the training samples are shown in Figure 6. The channel-wise contributions towards

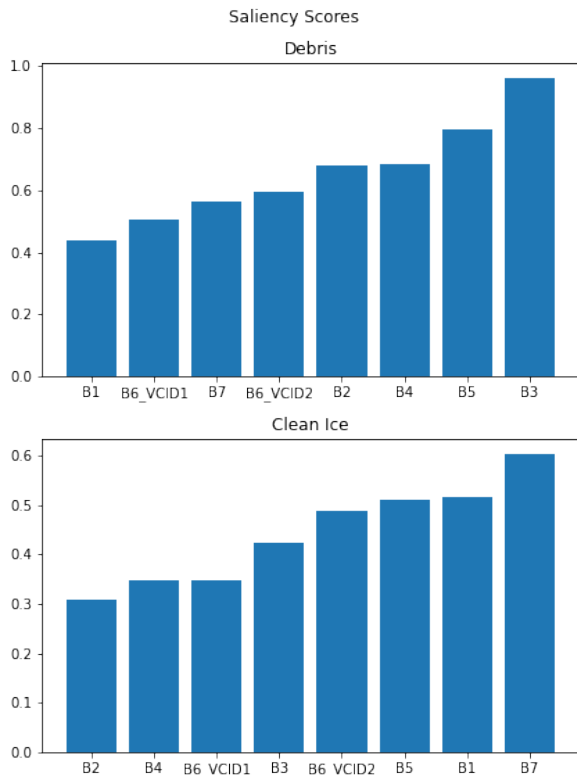


Figure 6: Average saliency scores for all sub-images in training set.

debris-covered glacial ice segmentation in decreasing order are: red, shortwave infrared 1, near infrared, green, high-gain thermal infrared, shortwave infrared 2, low-gain thermal infrared, and blue. Similarly, for clean glacial ice segmentation, the channel-wise contributions in decreasing order are: shortwave infrared 2, blue, shortwave infrared 1, high-gain thermal infrared, red, low-gain thermal infrared, near infrared, and green. As shown in Figure 6, the segmentation models have different high contributing channels for clean glacial ice and debris-covered glacial ice segmentation.

4 Discussion

Glaciers have been melting at an unprecedented rate in recent years due to global climate change [17, 19]. Glaciers are the largest freshwater reservoir on the planet [32], so it is necessary to understand the changes they undergo.

As a result, numerous approaches to automatically delineate glacier boundaries have been proposed [5, 6, 14, 15, 34, 36]. We frequently observe deep learning methods outperforming traditional machine learning methods for glacier segmentation in the literature [1, 5]. However, the results have not been very good, particularly in the case of debris-covered glacial ice.

In this work, we modify U-Net and train it using a novel loss function that allows the modified U-Net to focus on glacier boundaries. From Table 3, we see that standard U-Net is not able to detect debris-covered glacial ice in input sub-images. We can also see from Table 2 that only 2.44% of pixels in the training set correspond to debris-covered glacial ice. This shows that our proposed method is more robust than the original U-Net to imbalanced labels, which are common in remote sensing datasets.

From Figure 5, we can see how the weights change for \mathcal{L}_{SLBA} while training. A higher weight is assigned to masked dice loss at the beginning and the weights for boundary loss are gradually increased during training. The reason behind this could be that for an untrained model, it may be easier to learn glacier instances over trying to learn the boundaries. However once the network learns to label instances, it is easier to learn the glacier boundaries. This also explains why the model fails to converge when training solely on $\mathcal{L}_{Boundary}$ from scratch as can be seen from the results in Table 3.

We presented methods to improve debris-covered glacial ice segmentation from remote sensing imagery using deep learning. While we were able to show significant improvements over existing methods, the IoU for debris-covered glacial ice still leaves much to be desired. The existing body of literature on the topic has shown that the performance for debris-covered glacial ice segmentation can be improved by incorporating thermal signatures [28] and topographical information [10, 26, 27] from other satellites. Since debris-covered glacial ice is common in low-gradient areas due to how it forms and has cooler surface temperatures compared to the surrounding non-glaciated regions, we suspect that adding this information can further help improve the performance of debris-covered glacial ice segmentation. We may also be able to see an improvement in performance by using images from the recently-launched Landsat 9 satellite, instead

of the Landsat 7 images used in this work. The Operational Land Imager 2 (OLI-2) and the Thermal Infrared Sensor 2 (TIRS-2) sensors on Landsat 9 provide data that is radiometrically and geometrically superior to instruments on the previous generation Landsat satellites. With the higher radiometric resolution, Landsat 9 can differentiate 16,384 shades of a given wavelength compared to only 256 shades in Landsat 7. Meanwhile, the TIRS-2 in Landsat 9 enables improved atmospheric correction and more accurate surface temperature measurements. Future work includes using the images captured through these improved sensors and incorporating additional information such as a digital elevation model for improving debris-covered glacial ice segmentation performance.

5 Conclusion

In this research study, we proposed a modified version of the U-Net architecture for large-scale debris-covered glacial ice and clean glacial ice segmentation in the HKH from Landsat 7 multispectral imagery and concluded that debris-covered glacial ice (IoU: 35.94%) is significantly harder to delineate compared to clean glacial ice (IoU: 68.17%)(Table 3). We also introduced two different methods to combine commonly-used masked dice loss and boundary loss to incorporate label boundaries into the training process. We show that the performance of debris-covered glacial ice segmentation can be improved by encouraging the deep learning model to focus on label boundaries. The performance can be improved further by correctly weighing loss terms. Furthermore, the relative weights can be learned automatically from the data during the training process using our proposed loss (\mathcal{L}_{SLBA}). Figure 7 shows the performance of the models trained using \mathcal{L}_{SLBA} on a sample image from the test set. We also introduced the concept of feature saliency scores to quantify the contribution of each feature (channel) in the input image toward the final label and concluded that the red, short-wave infrared, and near infrared bands contribute the most towards the final label for debris-covered glacial ice segmentation, while shortwave infrared 2, blue, shortwave infrared 1 bands contributed the most towards the final label for clean glacial ice segmentation.

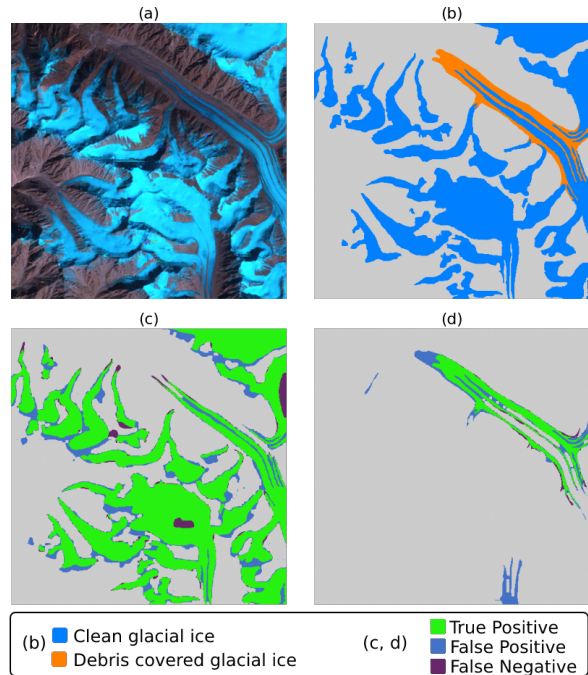


Figure 7: (a) Sample subimage from test set (b) Corresponding clean glacial ice and debris-covered glacial ice ground truth labels (c) True positive (TP), False positive (FP), False negative (FN) for clean glacial ice (IoU 79.17%) (d) TP, TP, FN for debris-covered glacial ice (IoU 59.19%)

6 Acknowledgements

We would like to thank Microsoft for providing us with the Microsoft Azure resources through their AI for Earth grant program (Grant ID: AI4E-1792-M6P7-20121005). We also acknowledge ICIMOD for providing a rich dataset which this work has been built on. This research was supported in part by the Department of Computer Science at The University of Texas at El Paso.

References

- [1] B. Aryal. *Glacier Segmentation in Satellite Images for Hindu Kush Himalaya Region*. PhD thesis, 2020.
- [2] B. Aryal, S. M. Escarzaga, S. A. Vargas Zesati, M. Velez-Reyes, O. Fuentes, and C. Tweedie.

- Semi-automated semantic segmentation of arctic shorelines using very high-resolution airborne imagery, spectral indices and weakly supervised machine learning approaches. *Remote Sensing*, 13(22):4572, 2021. doi: 10.3390/rs13224572.
- [3] S. Bajracharya, S. Maharjan, and F. Shrestha. Clean ice and debris covered glaciers of HKH region. 2011. doi: 10.26066/RDS.31029.
- [4] S. R. Bajracharya and B. R. Shrestha. The status of glaciers in the Hindu Kush-Himalayan region. 2011. doi: 10.53055/ICIMOD.551.
- [5] S. Baraka, B. Akera, B. Aryal, T. Sherpa, F. Shrestha, A. Ortiz, K. Sankaran, J. M. Lavista Ferres, M. A. Matin, and Y. Bengio. Machine learning for glacier monitoring in the Hindu Kush Himalaya. In *NeurIPS 2020 Workshop on Tackling Climate Change with Machine Learning*, 2020. URL <https://www.climatechange.ai/papers/neurips2020/57>.
- [6] R. Bhambri and T. Bolch. Glacier mapping: a review with special reference to the Indian Himalayas. *Progress in Physical Geography*, 33(5):672–704, 2009. doi: 10.1177/0309133309348112.
- [7] M. P. Bishop, J. F. Shroder Jr, and J. L. Ward. SPOT multispectral analysis for producing supraglacial debris-load estimates for Batura glacier, Pakistan. *Geocarto International*, 10(4):81–90, 1995. doi: 10.1080/10106049509354515.
- [8] M. P. Bishop, J. F. Shroder Jr, and B. L. Hickman. SPOT panchromatic imagery and neural networks for information extraction in a complex mountain environment. *Geocarto International*, 14(2):19–28, 1999. doi: 10.1080/10106049908542100.
- [9] A. Bokhovkin and E. Burnaev. Boundary loss for remote sensing imagery semantic segmentation. In *International Symposium on Neural Networks*, pages 388–401. Springer, 2019. doi: 10.1007/978-3-030-22808-8_38.
- [10] T. Bolch and U. Kamp. Glacier mapping in high mountains using DEMs, Landsat and ASTER data. In *Proceedings of the 8th International Symposium on High Mountain Remote Sensing Cartography*. Karl-Franzens-Universität Graz, 2005.
- [11] R. Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997. doi: 10.1023/A:1007379606734.
- [12] K. M. Cuffey and W. S. B. Paterson. *The physics of glaciers*. Academic Press, 2010. ISBN 9780080919126.
- [13] D. Eigen and R. Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2650–2658, 2015.
- [14] J. Florath, S. Keller, R. Abarca-del Rio, S. Hinz, G. Staub, and M. Weinmann. Glacier monitoring based on multi-spectral and multi-temporal satellite data: A case study for classification with respect to different snow and ice types. *Remote Sensing*, 14(4):845, 2022. doi: 10.3390/rs14040845.
- [15] Q. He, Z. Zhang, G. Ma, and J. Wu. Glacier identification from Landsat8 OLI imagery using deep U-NET. *ISPRS Annals of Photogrammetry, Remote Sensing & Spatial Information Sciences*, 5(3), 2020. doi: 10.5194/isprs-annals-V-3-2020-381-2020.
- [16] D. Hendrycks and K. Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. doi: 10.48550/arXiv.1606.08415.
- [17] W. W. Immerzeel, A. Lutz, M. Andrade, A. Bahl, H. Biemans, T. Bolch, S. Hyde, S. Brumby, B. Davies, A. Elmore, et al. Importance and vulnerability of the world’s water towers. *Nature*, 577(7790):364–369, 2020. doi: 10.1038/s41586-019-1822-y.
- [18] A. Kendall, Y. Gal, and R. Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.

- [19] P. D. Kraaijenbrink, M. Bierkens, A. Lutz, and W. Immerzeel. Impact of a global temperature rise of 1.5 degrees Celsius on Asia’s glaciers. *Nature*, 549(7671):257–260, 2017. doi: 10.1038/nature23878.
- [20] S. Le Moan, A. Mansouri, J. Y. Hardeberg, and Y. Voisin. Saliency for spectral image analysis. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 6(6):2472–2479, 2013. doi: 10.1109/JSTARS.2013.2257989.
- [21] B. Leibe, A. Leonardis, and B. Schiele. Robust object detection with interleaved categorization and segmentation. *International Journal of Computer Vision*, 77(1):259–289, 2008. doi: 10.1007/s11263-007-0095-3.
- [22] W. Luo, Y. Li, R. Urtasun, and R. Zemel. Understanding the effective receptive field in deep convolutional neural networks. *Advances in neural information processing systems*, 29, 2016. URL <https://proceedings.neurips.cc/paper/2016/file/c8067ad1937f728f51288b3eb986afaa-Paper.pdf>.
- [23] J. McGlinchy, B. Johnson, B. Muller, M. Joseph, and J. Diaz. Application of UNet fully convolutional neural network to impervious surface segmentation in urban environment from high resolution satellite imagery. In *IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium*, pages 3915–3918. IEEE, 2019. doi: 10.1109/IGARSS.2019.8900453.
- [24] K. E. Miles, B. Hubbard, T. D. Irvine-Fynn, E. S. Miles, D. J. Quincey, and A. V. Rowan. Hydrology of debris-covered glaciers in high mountain Asia. *Earth-Science Reviews*, 207:103212, 2020. ISSN 0012-8252. doi: 10.1016/j.earscirev.2020.103212. URL <https://www.sciencedirect.com/science/article/pii/S0012825220302580>.
- [25] Y. Mohajerani, M. Wood, I. Velicogna, and E. Rignot. Detection of glacier calving margins with convolutional neural networks: A case study. *Remote Sensing*, 11(1):74, 2019. doi: 10.3390/rs11010074.
- [26] N. Mölg, T. Bolch, P. Rastner, T. Strozzi, and F. Paul. A consistent glacier inventory for Karakoram and Pamir derived from Landsat data: distribution of debris cover and mapping challenges. *Earth System Science Data*, 10(4):1807–1827, 2018.
- [27] F. Paul, C. Huggel, and A. Kääb. Combining satellite multispectral image data and a digital elevation model for mapping debris-covered glaciers. *Remote sensing of Environment*, 89(4):510–518, 2004.
- [28] R. Ranzi, G. Grossi, L. Iacovelli, and S. Taschner. Use of multispectral aster images for mapping debris-covered glaciers within the glims project. In *IGARSS 2004. 2004 IEEE International Geoscience and Remote Sensing Symposium*, volume 2, pages 1144–1147. IEEE, 2004.
- [29] C. Robinson, A. Ortiz, K. Malkin, B. Elias, A. Peng, D. Morris, B. Dilkina, and N. Jojic. Human-machine collaboration for fast land cover mapping. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2509–2517, 2020. doi: 10.1609/aaai.v34i03.5633.
- [30] O. Ronneberger, P. Fischer, and T. Brox. U-Net: convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. ISBN 978-3-319-24574-4.
- [31] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun. Overfeat: integrated recognition, localization and detection using convolutional networks. *arXiv preprint arXiv:1312.6229*, 2014. doi: 10.48550/arXiv.1312.6229.
- [32] I. A. Shiklomanov et al. Global water resources. *Nature and resources*, 26(3):34–43, 1990.
- [33] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013. doi: 10.48550/arXiv.1312.6034.

- [34] S. Tian, Y. Dong, R. Feng, D. Liang, and L. Wang. Mapping mountain glaciers using an improved U-Net model with cSE. *International Journal of Digital Earth*, 15(1):463–477, 2022. doi: 10.1080/17538947.2022.2036834.
- [35] J. Tompson, R. Goroshin, A. Jain, Y. LeCun, and C. Bregler. Efficient object localization using convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 648–656, 2015.
- [36] Z. Xie, U. K. Haritashya, V. K. Asari, B. W. Young, M. P. Bishop, and J. S. Kargel. GlacierNet: a deep-learning approach for debris-covered glacier mapping. *IEEE Access*, 8:83495–83510, 2020. doi: 10.1109/ACCESS.2020.2991187.
- [37] J. Yosinski, J. Clune, A. M. Nguyen, T. J. Fuchs, and H. Lipson. Understanding neural networks through deep visualization. *CoRR*, abs/1506.06579, 2015. URL <http://arxiv.org/abs/1506.06579>.
- [38] P. Zhang, Y. Ke, Z. Zhang, M. Wang, P. Li, and S. Zhang. Urban land use and land cover classification using novel deep learning models based on high spatial resolution satellite imagery. *Sensors*, 18(11), 2018. ISSN 1424-8220. doi: 10.3390/s18113717. URL <https://www.mdpi.com/1424-8220/18/11/3717>.
- [39] M. Zheng, X. Miao, and K. Sankaran. Interactive visualization and representation analysis applied to glacier segmentation. *ISPRS International Journal of Geo-Information*, 11(8), 2022. ISSN 2220-9964. doi: 10.3390/ijgi11080415. URL <https://www.mdpi.com/2220-9964/11/8/415>.