Detailed Glacier Area Change Analysis in the European Alps with Deep Learning

Codruț-Andrei Diaconu^{1,2}Jonathan L. Bamber^{2,3}codrut-andrei.diaconu@dlr.dej.bamber@tum.de¹German Aerospace Center (DLR)²Technical University of Munich³University of Bristol

Abstract

Glacier retreat is a key indicator of climate change and requires regular updates of the glacier area. Recently, the release of a new inventory for the European Alps showed that glaciers continued to retreat at about $1.3\% a^{-1}$ from 2003 to 2015. The outlines were produced by manually correcting the results of a semi-automatic method applied to Sentinel-2 imagery. In this work we develop a fully-automatic pipeline based on Deep Learning to investigate the evolution of the glaciers in the Alps from 2015 to present (2023). After outlier filtering, we provide individual estimates for around 1300 glaciers, representing 87% of the glacierized area. Regionally we estimate an area loss of $-1.8\% a^{-1}$, with large variations between glaciers. Code and data are available at https://github.com/dcodrut/glacier_mapping_alps_tccml.

1 Introduction

Glaciers are critical components in the Earth system, and an essential water resource for both industrial and domestic use[1]. They act as a low-pass filter for high-frequency, inter-annual variability in weather, which makes them a valuable and sensitive climate change indicator on multi-annual timescales [2]. One of the observations for assessing glacier health is their area[3], which has been classified as an Essential Climate Variable (ECV) [4]. Glaciers in the European Alps experienced a large volume loss over the last two decades[5] and, as a consequence, they also retreated significantly (about -1.3% change in area per year), as shown by Paul *et al.* [6] in the latest inventory available for this region.

There are more than 200,000 glaciers worldwide and methods requiring manual intervention are, therefore, extremely labour intensive[3]. Fully automated methods are, however, challenging for many reasons, *e.g.* cast shadow, local clouds, difficulties in distinguishing debris-covered segments from surrounding rocks or seasonal/perennial snow from glacier ice *etc.* [7]. As a consequence, significant variations can occur even within outlines manually delineated by experts[6].

Recently, there has been significant progress in developing fully-automated methods for glacier mapping, mainly based on Deep Learning approaches. Xie *et al.* [8] provides a comparison of multiple architectures, with DeepLabV3+[9] and GlacierNet[10] (an extended U-Net[11] architecture) achieving the best (IOU) scores. Baraka *et al.* [12] use also U-Net for the Hindu Kush Himalayan region and inicluded a web tool that can be used for visualizing and correcting the predictions.

Here, we utilize a Deep Learning method with additional data processing steps in order to investigate the evolution of the glaciers in the Alps between 2015 and the present (2023). For this we train a segmentation model using the existing (manually corrected) inventory outlines. When evaluating the model on unseen data we find good agreement between the estimated areas and the inventory ones. Next, we develop a pipeline which automatically downloads Sentinel-2 data by minimizing the cloud coverage for each glacier separately. Finally, we apply the model on data from 2023 and then

Tackling Climate Change with Machine Learning: workshop at NeurIPS 2023.

estimate the area change both at individual glacier and regional level. The processed dataset ready for training and evaluating glacier mapping models, together with the code to (re)train the models are freely available at https://github.com/dcodrut/glacier_mapping_alps_tccml.

2 Data

Inventory data. The first step is to collect the images which were employed to build the inventory [6], using exactly the same dates, to ensure that the outlines and the images correspond perfectly. The inventory contains 4395 glaciers with a total area of ca. 1806km^2 and it was built using Sentinel-2 images mainly from 2015 (with some from 2016 and 2017). The previous inventory (*i.e.* RGI v6 [13], from 2003) contains 3927 glaciers with a total area of ca. 2092km^2 . When comparing them, a regional-level shrinkage can be estimated, *i.e.* -14% (-1.2% a⁻¹). However, a more realistic estimate is -1.3% a⁻¹ when taking into account that the new inventory contains glaciers which were previously missing. Moreover, this is still considered a lower bound due to the inclusion of seasonal and perennial snow in a few sub-regions [6].

Given that the resolution of Sentinel-2 data is 10m, we decided to use only the glaciers with an area larger than 0.1 km^2 . Although this reduces the number of glaciers sampled (1646, *i.e.* about 37%), the percentage of glacierized area covered is close to 95%. To facilitate further analyses, we additionally drop 13 small glaciers (average area = 0.16 km^2) which are not represented in the previous inventory from 2003. This results in 1633 glaciers, covering 94.36% of the current inventory area, which will be used to train and validate our model.

Present data (2023). In order to maximize the observation period and thus increasing the signal-to-noise ratio, we utilise, where possible, the most recent data, *i.e.* end of summer 2023. According to the most recent measurements from the Glacier Monitoring in Switzerland (GLAMOS), 2023 was a strong melt year [14]. This, in turn, reduces the occurrence of spurious late-season snow cover, thus offering better conditions for glacier mapping.

A common criterion used to download optical data is by choosing the tile (*i.e.* 100km x 100km for Sentinel-2) with the smallest percentage of cloud coverage for each region of interest. If we follow this strategy and then compute the average cloud coverage per glacier (using the inventory outlines), we obtain an average of 4%. However, if rather than restricting to one single tile we instead use the least cloudy five tiles (centered on 01.09.2023 \pm 15 days) and then choose the best for each glacier individually, we significantly reduce the cloud coverage to 0.1%. This is lower than the average in the inventory images *i.e.* 2.7%.

3 Methodology

Model architecture. To segment the glaciers, we use the U-Net architecture [11], with a ResNet34 backbone[15], and the implementation from [16]. This architecture was found to be one of the best performing in [8], with a relatively smaller model size compared to the other methods evaluated. We extend the input from three to six channels, to accommodate the following inputs:

- five Sentinel-2 bands: blue (B2), green (B3), red (B4), NIR (B8) and SWIR (B12), which we found the most informative;
- surface elevation, obtained from NASADEM[17] (30m resolution) and processed using the Open Global Glacier Model[18]. The surface elevation should help for debris-cover on glaciers [10], Central Europe being one of the regions with the highest percentage of debris cover[19].

Model training. We train the model to predict the probability that each pixel is glacier or not, using patches of 256x256 pixels (*i.e.* 2.56x2.56 km). Given that we apply the model only on glacierized regions, we sample patches only if the center is on the glacier, which also helps in balancing the two classes. This implies that the model sees only the glaciers and a maximum buffer of 1.28km around, which also helps in reducing the data volume. We train the model with PyTorch [20] and PyTorch Lighting[21], with a binary cross entropy loss. The pixels which are either missing or covered by clouds/shadows are excluded when computing the loss. The bands are individually normalized using the mean and standard deviation of the training samples.



Figure 1: Overview of our approach, exemplified for a glacier in the Obersulzbach valley, Austria (47.1° N, 12.3° E), with an inventory area of 4.03km^2 in 2016. We apply the model in 2016, where some miss-classified sections can be observed, and then in 2023, where the model again misses those segments (they are debris-covered and seem to become completely disconnected from the main glacier body). The error being systematic, it will not affect significantly the area change estimate. For this glacier we obtain an area loss of $-1.66\% \text{ a}^{-1}$.

Table 1: Performance metrics for each of the five testing CV folds.

subregion	#patches	#glaciers	Accuracy	IOU	Precision	Recall	F1
r_1	1855	349	0.953	0.794	0.875	0.896	0.878
r_2	1321	234	0.955	0.862	0.924	0.926	0.923
r_3	1084	184	0.960	0.879	0.931	0.937	0.933
r_4	2146	406	0.964	0.836	0.916	0.903	0.905
r_5	2301	437	0.951	0.769	0.951	0.796	0.857
		$\mu \pm \sigma$:	0.96 ± 0.01	0.83 ± 0.05	0.92 ± 0.03	0.89 ± 0.06	0.90 ± 0.03

Model evaluation. We use a five-fold cross-validation (CV) scheme to evaluate the model. The split is done geographically, from West to East, to ensure that the testing scores capture the generalization ability of the model. Another important reason is that, ultimately, we want one prediction for each glacier. With this scheme, depicted in Figure A1, we train five models and then collect only the inferences from the testing areas, thus covering the entire region.

Area (change) estimation. Given the significant volume loss observed over the 2000-2019 period[5], with a mean elevation change of -1.02 ± 0.21 m a⁻¹, we can assume that glaciers in this region do not grow over the 2015-2023 period. This allows us to extract the changes in the areas by applying the model for each glacier but only for the pixels within the inventory outlines, thus excluding the predictions outside these. However, we do not use the areas from the inventory as the reference value but the predicted ones such that, if the model makes systematic errors, they will cancel out, as in the case illustrated in Figure 1. Therefore, for each glacier, we calculate the area change per year as $(A_{2023}^p - A_y^p)/(2023 - y)$ where A^p denotes the predicted area and y denotes the inventory year, so $y \in \{2015, 2016, 2017\}$. If we then divide this annual rate by A_y^p , we can express the area change rate in percentages per year. In order to decrease the noise in our estimates, we proceed with some filtering steps, which are detailed in the Appendix A. This leaves 1285 glaciers, which is only 29% from the inventory total, but covering 1572 km², a significant fraction of the glacierized area (87%).

4 Results

We first report the standard performance metrics based on patches that fall into the testing region of each split. The results (Table 1) are comparable to those from previous studies on the Himalayan region[8], which also experiences significant debris cover[19]. We also observe that the scores vary from one sub-region to other, the best being obtained for the regions with a smaller number of larger glaciers (given that the data split ensures a similar total glacierized area in each sub-region).



Figure 2: Distribution of per-glacier area change rates, separated by their size class (same ones as in [6]). The highest possible rate is -12.5% which would imply that the glacier completely melted over the last eight years. The text shows the number of glaciers in each class (n) and the sum of their areas relative to the total glacierized area covered (A_%).

The relatively high recall shows that the models manage to recover a significant part of the glacier, therefore next we will focus on the predictions at glacier level. We estimate the areas as previously described and compare them to the inventory values from ~ 2015 , with which we find a very good agreement (details provided in Figure A3). Given that we only limit the predictions to the given inventory, our predictions are always an underestimate but with a relatively small bias, only -0.07km². The filtering steps remove the outliers and thus reduce the bias to -0.06km². Moreover, if we assume that a significant part of the bias is due to systematic errors (*e.g.* consistently missed debris-covered segments, as the one exemplified in Figure 1, or segments always shaded) we expect the area changes not to be significantly affected by errors in the model.

We show the distribution of the area change rates of all the glaciers separated by their size class in Figure 2. First, we observe that most of the glaciers, especially the small ones, have a significant negative change rate, which is in line with the negative volume changes estimated over the last two decades [5]. Second, we notice many small glaciers are losing a significant fraction of their surface area over the last eight years, some being close to complete deglaciation (*i.e.* disappearing over that time period). Finally, we observe a higher variability among the small glaciers, as they can have diverse topographical and morphological settings which influences their sensitivity to climate change processes[22], but it may also be due to a lower signal-to-noise ratio. To obtain a regional level estimate, we integrate over the total estimated areas, obtaining 1504km² and 1289km² in 2015 and 2023, respectively, resulting in a -1.8% a⁻¹ area loss.

5 Conclusions

We develop a fully-automatic pipeline to train a glacier mapping model for the European Alps using the most recent inventory and apply it on data from 2023. The predicted glacier masks are then used to estimate the areas and, after post-processing, provide estimates, both at glacier and regional level, for area changes from 2015-2023. Our regional estimate is -1.8% a⁻¹, which illustrates the high sensitivity of the glaciers in this region to climate change, with significant inter-glacier variability. Our detailed glacier-scale estimates could provide valuable constraints on future water availability and discharge for hydro-power production and other water-critical activities, which heavily rely on glacier-driven melt water [23].

Acknowledgments

CAD is supported by the Helmholtz Association through the joint research school Munich School for Data Science - MuDS (grant number: HIDSS-0006). JLB was supported by German Federal Ministry of Education and Research (BMBF) in the framework of the international future lab AI4EO (grant number: 01DD20001). We gratefully acknowledge the computing time granted by the Helmholtz Association's Initiative and Networking Fund on the HAICORE@FZJ partition.

References

- W. W. Immerzeel, A. F. Lutz, M. Andrade, A. Bahl, H. Biemans, T. Bolch, S. Hyde, S. Brumby, B. J. Davies, and A. C. Elmore, "Importance and vulnerability of the world's water towers," *Nature*, 2020.
- [2] S. Legg, "IPCC, 2021: Climate change 2021-the physical science basis," Interaction, 2021.
- [3] E. Berthier, D. Floricioiu, A. S. Gardner, N. Gourmelen, L. Jakob, F. Paul, D. Treichler, B. Wouters, J. M. Belart, and A. Dehecq, "Measuring glacier mass changes from space-a review," *Reports on Progress in Physics*, 2023.
- [4] S. Bojinski, M. Verstraete, T. C. Peterson, C. Richter, A. Simmons, and M. Zemp, "The concept of essential climate variables in support of climate research, applications, and policy," *Bulletin of the American Meteorological Society*, 2014.
- [5] R. Hugonnet, R. McNabb, E. Berthier, B. Menounos, C. Nuth, L. Girod, D. Farinotti, M. Huss, I. Dussaillant, and F. Brun, "Accelerated global glacier mass loss in the early twenty-first century," *Nature*, 2021.
- [6] F. Paul, P. Rastner, R. S. Azzoni, G. Diolaiuti, D. Fugazza, R. Le Bris, J. Nemec, A. Rabatel, M. Ramusovic, G. Schwaizer, and C. Smiraglia, "Glacier shrinkage in the Alps continues unabated as revealed by a new glacier inventory from Sentinel-2," *Earth System Science Data*, 2020.
- [7] F. Paul, N. E. Barrand, S. Baumann, E. Berthier, T. Bolch, K. Casey, H. Frey, S. P. Joshi, V. Konovalov, and R. Le Bris, "On the accuracy of glacier outlines derived from remote-sensing data," *Annals of Glaciology*, 2013.
- [8] Z. Xie, V. K. Asari, and U. K. Haritashya, "Evaluating deep-learning models for debris-covered glacier mapping," *Applied Computing and Geosciences*, 2021.
- [9] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-decoder with atrous separable convolution for semantic image segmentation," in *ECCV*, 2018.
- [10] 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*, 2020.
- [11] 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*, 2015.
- [12] S. Baraka, B. Akera, B. Aryal, T. Sherpa, F. Shresta, A. Ortiz, K. Sankaran, J. L. Ferres, M. Matin, and Y. Bengio, "Machine learning for glacier monitoring in the Hindu Kush Himalaya," *Tackling Climate Change with Machine Learning workshop at NeurIPS*, 2020.
- [13] W. T. Pfeffer, A. A. Arendt, A. Bliss, T. Bolch, J. G. Cogley, A. S. Gardner, J.-O. Hagen, R. Hock, G. Kaser, C. Kienholz, *et al.*, "The randolph glacier inventory: A globally complete inventory of glaciers," *Journal of glaciology*, 2014.
- [14] GLAMOS. "The Swiss Glaciers 1880-2018/19, Glaciological Reports No 1-140, Yearbooks of the Cryospheric Commission of the Swiss Academy of Sciences (SCNAT), published since 1964 by VAW / ETH Zurich." (2023), [Online]. Available: https://doi.glamos.ch/figures/index.html (visited on 2023-09-23).
- [15] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in CVPR, 2016.
- [16] P. Iakubovskii, Segmentation models pytorch, https://github.com/qubvel/segmentation_ models.pytorch, 2019.
- [17] NASA JPL, Nasadem merged dem global 1 arc second v001, 2020. DOI: 10.5067/MEASURES/ NASADEM/NASADEM_HGT.001. [Online]. Available: https://lpdaac.usgs.gov/products/ nasadem_hgtv001.
- [18] F. Maussion, A. Butenko, N. Champollion, M. Dusch, J. Eis, K. Fourteau, P. Gregor, A. H. Jarosch, J. Landmann, and F. Oesterle, "The open global glacier model (OGGM) v1. 1," *Geoscientific Model Development*, 2019.
- [19] S. Herreid and F. Pellicciotti, "The state of rock debris covering Earth's glaciers," *Nature Geoscience*, 2020.
- [20] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, and L. Antiga, "Pytorch: An imperative style, high-performance deep learning library," *NeurIPS*, 2019.
- [21] W. Falcon and The PyTorch Lightning team, PyTorch Lightning, 2019. DOI: 10.5281/zenodo.3828935.
 [Online]. Available: https://github.com/Lightning-AI/lightning.
- [22] M. Huss and M. Fischer, "Sensitivity of very small glaciers in the Swiss Alps to future climate change," *Frontiers in earth science*, 2016.
- [23] B. Schaefli, P. Manso, M. Fischer, M. Huss, and D. Farinotti, "The role of glacier retreat for Swiss hydropower production," *Renewable energy*, 2019.

A Appendix



Figure A1: Cross-validation scheme with a geographic split

Filtering steps

First, the aim is to drop the results for which the model does not perform well according to the testing results. This can be due to *e.g.* very high percentage of debris coverage, a significant part covered by topographical shadows or presence of clouds which were not masked correctly. Thus we keep only the glaciers which have a recall larger than 75%, as shown in Figure A2. Second, we drop the glaciers for which either of the start or end years has more than 20% masked-out pixels. For the glaciers where there are still some masked-out pixels, we apply a correction to the estimated area: we assume a recall equal to the one obtained on the non-masked pixels, such that the values from the two years are comparable. Finally, once we compute the rates as described in Section 3, we drop the most extreme 5% of the values.



Figure A2: Dropping the glaciers with a small recall. The text shows the number of glaciers we loose, their total area and the relative percentages.



Figure A3: Comparing the predicted glacier-wide areas to the inventory ones, before and after the automatic filtering steps.