Exploring GeoAI Methods for Supraglacial Lake Mapping on Greenland Ice Sheet

Xuanshu Luo
Technical University of Munich
Munich, Germany
xuanshu.luo@tum.de

Paul Walther
Technical University of Munich
Munich, Germany
paul.walther@tum.de

Wejdene Mansour
Technical University of Munich
Munich, Germany
wejdene.mansour@tum.de

Balthasar Teuscher
Technical University of Munich
Munich, Germany
balthasar.teuscher@tum.de

Johann Maximilian Zollner
Technical University of Munich
Munich, Germany
maximilian.zollner@tum.de

Hao Li
Technical University of Munich
Munich, Germany
hao_bgd.li@tum.de

Martin Werner
Technical University of Munich
Munich, Germany
martin.werner@tum.de

ABSTRACT

The ACM SIGSPATIAL Cup 2023 proposed the challenge to identify and map supraglacial lakes in Greenland in satellite imagery. The peculiarities of supraglacial lakes pose a hard problem for semantic segmentation and object detection tasks because the definition of a lake is ill-fitted to the inner workings of such approaches. For example, lakes are often covered by ice and snow and narrow streams can connect distinct lakes, which is not directly translatable to the semantic segmentation of water. It is also not well-posed for object detection, especially the identity relation - what is a lake, what is not (yet) a lake, and what are two lakes is challenging. In this context, we worked on adapting semantic segmentation using the Segment Anything Model and instance segmentation using Mask R-CNN to the setting. The latter ended up superior in our own evaluation and even got ranked second among all participants. We are proud that our approach has led to competitive performance. The source code is available from https://github.com/tum-bgd/GISCup23.

CCS CONCEPTS

• **Computing methodologies** → **Image segmentation**; • **Information systems** → Geographic information systems.

KEYWORDS

supraglacial lakes, image segmentation, satellite imagery, computer vision, mask R-CNN, segment anything model

ACM Reference Format:
Xuanshu Luo, Paul Walther, Wejdene Mansour, Balthasar Teuscher, Johann Maximilian Zollner, Hao Li, and Martin Werner. 2023. Exploring GeoAI Methods for Supraglacial Lake Mapping on Greenland Ice Sheet. In *The 31st ACM International Conference on Advances in Geographic Information Systems (SIGSPATIAL '23), November 13–16, 2023, Hamburg, Germany.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3589132.3629971

SIGSPATIAL '23, November 13–16, 2023, Hamburg, Germany

ACM ISBN 979-8-4007-0168-9/23/11…$15.00
https://doi.org/10.1145/3589132.3629971

1 INTRODUCTION

Climate change in recent years has led to increasing ice sheet melting rates all over the earth [14, 15]. Accumulated snow on the Greenland ice sheet melts into supraglacial lakes every summer. These lakes form earlier and at higher elevations than ever before, thus are an essential aspect of climate research [1, 3]. However, a quantitative study of all these effects requires manual annotation and image analysis, which does not scale to large spatial regions. Considering that such lakes are very small compared with the size of the Greenland ice sheet, it might be unable to uncover large-scale dynamics.

The ACM SIGSPATIAL Cup 2023[1] asked for methods to automatically map such supraglacial lakes on satellite imagery. Challenge participants were given access to preprocessed imagery of four dates containing only limited spectral information covering the visible colors red, green, and blue and to corresponding hand-drawn annotations on parts of these images. It remains open whether the interest of the domain scientist is more towards object detection (e.g., identifying lakes as objects) or lake segmentation (e.g., finding out on a selected spatial scale which pixels belong to a supraglacial lake).

Generating accurate annotations of supraglacial lakes in this dataset is hardly as trivial as solving the semantic segmentation problem of "water or not" [9]. This is because lakes can be partly covered by ice and snow. The challenge further integrates this vagueness by a set of rules on the lakes: lakes should be contiguous, simple polygons of sufficient size; lakes can contain floating ice; and finally, lakes connected by narrow streams should remain two different lakes. In addition, the evaluation measure is a domain-specific construction in which smaller and larger polygons are accepted in quite a wide range. In this context, we investigate

[1] https://sigspatial2023.sigspatial.org/giscup/, last accessed Oct. 18, 2023

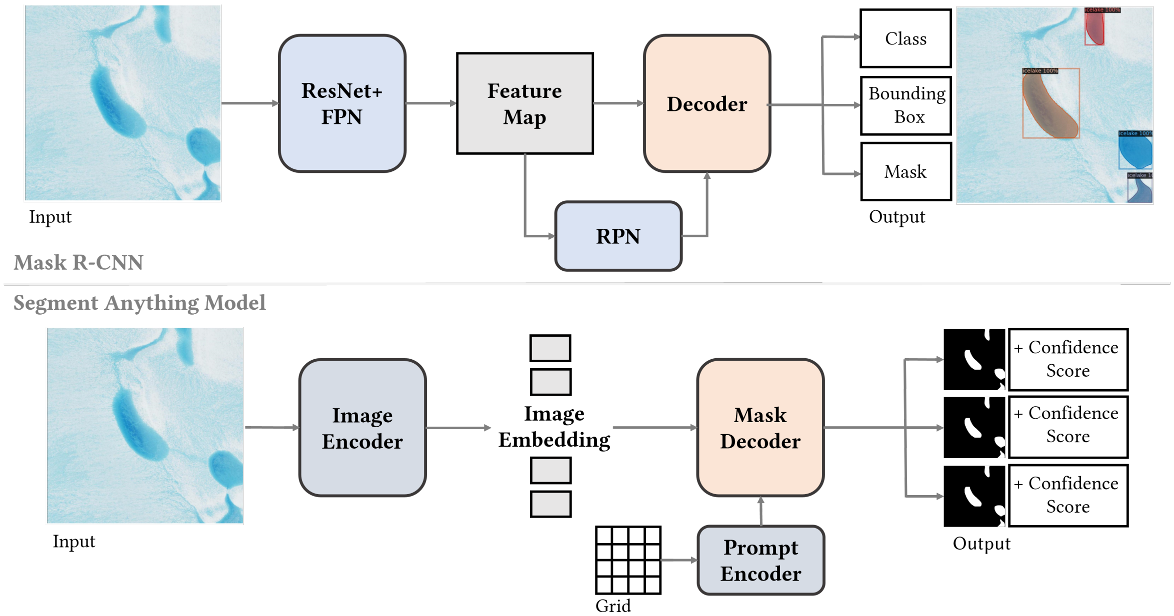

Figure 1: The two main methods explored - SAM (top) and MaskRCNN (bottom).

two approaches, out of which we only submit a single one to the challenge: we apply an instance segmentation model (e.g., object detection followed by binary segmentation) based on Mask R-CNN [5] and a semantic segmentation model based on the Segment Anything Model (SAM) [8] followed by a topological integration strategy to solve the challenge.

The remainder of the paper is structured as follows: Section 2 shortly reviews the baseline methods selected for this challenge. Section 3, shortly describes the dataset and some observations. Section 4 explains the experimental procedure and results for two methods and gives some evaluation results on the publicly available fragment of the dataset. Section 6 concludes the paper.

2 SELECTED METHODS

The continent-wide mapping of meltwater systems is non-economic with land-based approaches due to the region's high variability, large area, and difficult reachability. Hence, multiple attempts to use satellite data have been reported in the literature [2, 4, 11, 16]. Concerning deep learning, one approach is image segmentation based on U-Net [3]. However, recent advances in semantic segmentation, most notably foundation models such as SAM [8] or advanced region proposal based networks like Mask R-CNN [5] have seemingly not been adapted to this scenario yet. This paper explores the latter two architectures and finds quite interesting baseline behavior before tuning the models.

2.1 Mask R-CNN: A Region-Based Segmentation Model

The Mask R-CNN model can be considered an extension to the Faster R-CNN model [5]. In such models, an image is encoded into feature maps, which are then fed through a region proposal network (RPN) [13] to find areas that likely contain an object, as shown in 1. The result of the region proposal and the image features are combined to predict bounding boxes of objects in the image, including their class labels with confidence scores and a binary segmentation of their footprint.

The model is trained on a dataset containing object outlines as bounding boxes and segmentation masks inside. Furthermore, the visual feature extraction can be pre-trained on large datasets such as ImageNet. In this challenge, we utilized the individual lake polygons, computed their axis-aligned bounding boxes, and rasterized a mask from the polygons to fine-tune a Mask R-CNN.

2.2 Segment Anything: A Foundation Model for Segmentation

SAM, introduced in 2023, poses a foundation model pre-trained on the one million images SA-1B dataset. It claims to perform promising in few and single shot learning of instance segmentation tasks [8]. As shown in Figure 1, the model features an autoencoder architecture with two parts. One part is a two-fold encoder, which a) transforms the input visuals to a one-dimensional embedding and b) offers the possibility for prompt encoding. That means the user can mark special points of interest with additionally supplied points, bounding boxes, or descriptive texts. The other part is a mask decoder that translates the image and prompt embeddings to binary masks with a confidence score.

3 DATASET AND OBSERVATIONS

The ACM SIGSPATIAL Cup dataset contains four multi-part satellite images covering six regions of the Greenland ice sheet together with a geo-package of manually labeled polygons identifying

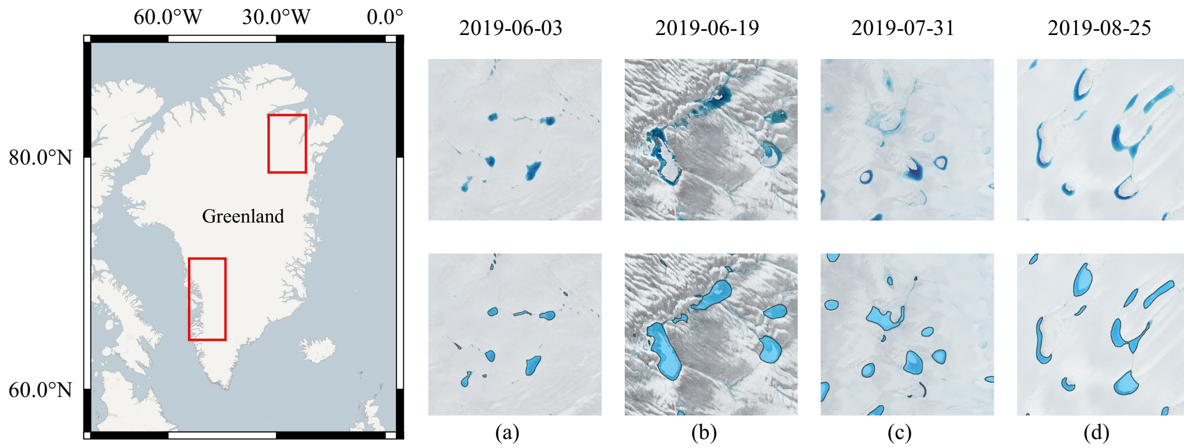

Figure 2: Examples of Supraglacial Lake Mapping Results. (a)-(d) are mapping results based on multi-temporal Sentinel-2 Imagery from 2019-06-03, 2019-06-19, 2019-07-31, 2019-08-25, respectively.

supraglacial lakes[2]. It has to be noted that labels were only provided for half of the image region, while the organizers of ACM SIGSPATIAL Cup kept the other half for the challenge evaluation.

Initial visual inspection of the labels indicated some inconsistencies, probably resulting from human bias. Concretely, this may originate from the question to which extent snow-covered lakes should be included, which lakes to consider separately, and whether lakes can be treated as a stream or too small and thus should not be included. This poses an additional challenge for the algorithms as they must fit the domain and labeling behavior to achieve correct results.

4 METHOD SELECTION AND TUNING

Since the ACM SIGSPATIAL Cup challenge mimics an instance segmentation challenge, two fundamental approaches, namely SAM and Mask R-CNN, were chosen for further evaluation. These two algorithms pose different state-of-the-art approaches to image segmentation. SAM features a large training basis and a lot of trainable weights, and Mask R-CNN includes a region attention-based architecture that helps to not only predict labels for single pixels but also make sense of larger areas.

As a first step, we apply the following preprocessing before model fine-tuning. The provided satellite images are cut into 1024×1024 pixel tiles with a step size of 512. Tiles containing more than 50% blank pixels are excluded, as are tiles that are unlikely to contain lakes (e.g., coast, snowy inland, etc.). The latter is achieved by filtering on the ratio of *blue* colors. Different thresholds were used for the train and test set. Like this, we only keep tiles with at least 105 *blue* pixels for inference. Due to the posing of the challenge of the ACM SIGSPATIAL Cup, no additional split is applied on the train and test data, as the trained algorithm should leverage the whole labeled dataset's knowledge.

[2]https://sigspatial2023.sigspatial.org/giscup/download.html, last accessed Oct. 16, 2023

4.1 Mask R-CNN

Our implementation utilizes a backbone based on `ResNet-101` [6] and Feature Pyramid Network [10] to extract latent feature maps of tiles as input for RPN and mask estimations. Pre-trained parameters provided by the detectron2 library [18] are adopted for fine-tuning using preprocessed supraglacial lake data. Figure 2 shows several estimation results on tiles. The fine-tuning procedure of Mask R-CNN is seemingly promising, as lakes primarily covered by snow and ice could still be segmented correctly. Finally, a series of post-processing steps w.r.t. *holes* in lakes, sizes, and narrow streams are applied to the Mask R-CNN estimation to generate the final result.

4.2 Segment Anything

The model weights of the SAM encoder were taken from the transformers python library [17]. At the same time, the decoder was fine-tuned using an Adam optimizer [7] and mean squared error loss, all implemented using PyTorch [12]. After fine-tuning SAM, the initial results without prompt encoding were promising, as shown in Figure 3. The concepts of the lakes were understood, and a large part of the lake regions was labeled correctly.

However, the main issue with this approach is the pixel-wise prediction resulting in fragmented segments, as highlighted in the middle image of Figure 3. Such over-segmented patches are considered independent lakes instead of being treated as a single big lake. While this is correct to the human eye, the evaluation method would render it a failure. Indeed, not a single patch comes close to covering about half of the ground truth area, thus resulting in multiple false positives and not a single true positive. A possible explanation for these results is the difference in the data pre-processing compared to that adopted for Mask R-CNN. Unlike for Mask R-CNN, every tile generated from the GeoTIFFs is used to train SAM. This means that the color distribution is heavily skewed towards white pixels, which could explain the overfitting of SAM to the white pixels. An idea to mitigate this issue is to use the probability maps for additional post-processing according to the idea depicted in the right picture of Figure 3. Instead of relying on a single global

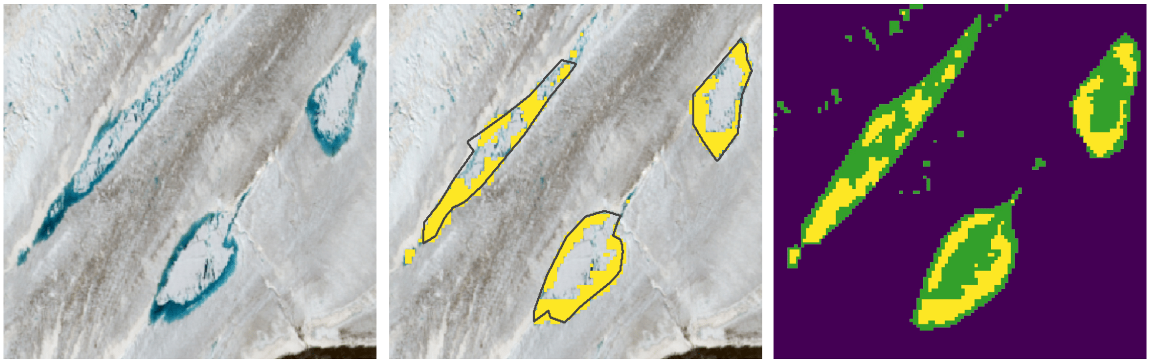

Figure 3: The input tile with three lakes (left), SAM-generated results (yellow), and the ground-truth polygons (blue)(middle), the idea of a probability map-based approach for improved contiguous areas (right).

probability threshold to generate the masks (yellow), a first lower threshold is used to obtain contiguous areas (green), and a higher second threshold (yellow) is applied to filter out segments that do not contain any high pixel probabilities. The resulting segments are anticipated to reflect superior overall confidence in a lake.

5 VALIDATION

A validation region is manually selected from training regions as a validation dataset. We calculate F1 scores on the validation set of two approaches using a similar method described by the ACM SIGSPATIAL Cup organizers. The Mask R-CNN solution has ultimately been selected for submission because it outperforms the SAM-based solution. Possible reasons for SAM, as a foundation model, not providing better performance in this task include a) Reasonable and effective fine-tuning is non-trivial as SAM is of large scale; b) pixel-wise estimation of SAM is hardly suitable for this task.

6 CONCLUSION

This paper investigates two approaches to the ACM SIGSPATIAL Cup 2023 challenge of supraglacial lake mapping on the Greenland ice sheet: Mask R-CNN for instance segmentation and SAM for semantic segmentation. These two models are fine-tuned with the given dataset. The Mask R-CNN solution is finally selected for submission as it obtains better results than SAM and excels in predicting continuous patches compared to individual pixel predictions. This resulted in an F1-score of 0.699 on the test set.

REFERENCES

[1] Robin E. Bell, Alison F. Banwell, Luke D. Trusel, and Jonathan Kingslake. 2018. Antarctic surface hydrology and impacts on ice-sheet mass balance. *Nature Climate Change* 8, 12 (2018), 1044–1052. https://doi.org/10.1038/s41558-018-0326-3

[2] Rebecca Dell, Neil Arnold, Ian Willis, Alison Banwell, Andrew Williamson, Hamish Pritchard, and Andrew Orr. 2020. Lateral meltwater transfer across an Antarctic ice shelf. *The Cryosphere* 14, 7 (2020), 2313–2330. https://doi.org/10.5194/tc-14-2313-2020

[3] Mariel Dirscherl, Andreas J. Dietz, Christof Kneisel, and Claudia Kuenzer. 2021. A Novel Method for Automated Supraglacial Lake Mapping in Antarctica Using Sentinel-1 SAR Imagery and Deep Learning. *Remote Sensing* 13, 2 (2021). https://doi.org/10.3390/rs13020197

[4] Anna Ruth W. Halberstadt, Colin J. Gleason, Mahsa S. Moussavi, Allen Pope, Luke D. Trusel, and Robert M. DeConto. 2020. Antarctic Supraglacial Lake Identification Using Landsat-8 Image Classification. *Remote Sensing* 12, 8 (2020), 1327. https://doi.org/10.3390/rs12081327

[5] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. 2017. Mask R-CNN. In *Proceedings of the IEEE international conference on computer vision*. 2961–2969.

[6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.

[7] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. https://doi.org/10.48550/arXiv.1412.6980

[8] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. 2023. Segment Anything. arXiv:2304.02643 [cs.CV]

[9] Hao Li, Johannes Zech, Christina Ludwig, Sascha Fendrich, Aurelie Shapiro, Michael Schultz, and Alexander Zipf. 2021. Automatic mapping of national surface water with OpenStreetMap and Sentinel-2 MSI data using deep learning. *International Journal of Applied Earth Observation and Geoinformation* 104 (2021), 102571.

[10] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2117–2125.

[11] Mahsa Moussavi, Allen Pope, Anna Halberstadt, Luke Trusel, Leanne Cioffi, and Waleed Abdalati. 2020. Antarctic Supraglacial Lake Detection Using Landsat 8 and Sentinel-2 Imagery: Towards Continental Generation of Lake Volumes. *Remote Sensing* 12, 1 (2020), 134. https://doi.org/10.3390/rs12010134

[12] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. arXiv:1912.01703 [cs.LG]

[13] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* 28 (2015).

[14] The Imbie Team. 2018. Mass balance of the Antarctic Ice Sheet from 1992 to 2017. *Nature* 558, 7709 (2018), 219–222. https://doi.org/10.1038/s41586-018-0179-y

[15] The Imbie Team. 2020. Mass balance of the Greenland Ice Sheet from 1992 to 2018. *Nature* 579, 7798 (2020), 233–239. https://doi.org/10.1038/s41586-019-1855-2

[16] Andrew G. Williamson, Neil S. Arnold, Alison F. Banwell, and Ian C. Willis. 2017. A Fully Automated Supraglacial lake area and volume Tracking ("FAST") algorithm: Development and application using MODIS imagery of West Greenland. *Remote Sensing of Environment* 196 (2017), 113–133. https://doi.org/10.1016/j.rse.2017.04.032

[17] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Online, 38–45. https://www.aclweb.org/anthology/2020.emnlp-demos.6

[18] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. 2019. Detectron2. https://github.com/facebookresearch/detectron2.