DEEP LEARNING FOR REAL-TIME AVALANCHE DETECTION IN WEBCAM IMAGES

James Fox[1,*], Anna Siebenbrunner[2,3], Sandra Reitinger[1], David Peer[1,4], Antonio Rodríguez-Sánchez[1]

[1]Department of Computer Science, University of Innsbruck, Austria
[2]Lo.La Peak Solutions GmbH, Innsbruck, Austria
[3]Faculty of Geo- and Atmospheric Sciences, University of Innsbruck, Austria
[4]DeepOpinion, Innsbruck, Austria

ABSTRACT: Continuous snow avalanche monitoring is essential to enable rapid responses to avalanche incidents. However, existing satellite-based methods of remote avalanche detection are typically unsuitable for ongoing monitoring due to long satellite revisit intervals and insufficient spatial resolution. This paper proposes a novel approach to automating avalanche detection via analysis of webcam streams with deep learning models. To assess the viability of this approach, we trained convolutional neural networks on a publicly-released dataset of 4090 mountain photographs and achieved avalanche detection F1 scores of 92.9 % per image and 64.0 % per avalanche. Notably, our models do not require a digital elevation model, enabling straightforward integration with existing webcams in new geographic regions. The paper concludes with findings from an initial case study conducted in the Austrian Alps and our vision for operational applications of trained models. The code and dataset are available at github.com/j-f-ox/avalanche-detection.

Keywords: avalanche detection, avalanche monitoring, avalanche, snow avalanche, remote sensing, deep learning

1. INTRODUCTION

Timely detection of snow avalanches is essential for urgent rescue operations and the swift enactment of safety measures around critical infrastructure in mountainous regions. Ongoing data on avalanche occurrences also play an important role in avalanche forecasting and informing risk management strategies (Schweizer and Herwijnen, 2013). However, avalanche monitoring currently relies heavily on field work (Hafner et al., 2022), which not only exposes observers to avalanche-prone terrain but also yields incomplete data with biases towards easily accessible locations in fair weather conditions (Eckerstorfer et al., 2016; Schweizer et al., 2015). Remote sensing can address these limitations by enabling the systematic collection of data at scale.

Although prior research into automating avalanche detection has focused on satellite data, the lengthy intervals between satellite revisits, which can span days, typically render them unsuitable for continuous monitoring. Satellite imagery may also lack the requisite spatial resolution to identify smaller avalanches (Hafner et al., 2021; Eckerstorfer et al., 2016). To overcome these disadvantages, we propose utilizing ground-based web cameras (webcams) as a data source for ongoing remote avalanche monitoring.

This paper presents findings from an initial study exploring the viability of automated avalanche detection in webcam images. After reviewing related work (section 2), we describe the methodology employed to train convolutional neural networks (CNNs) on a publicly released dataset of avalanche photographs (section 3) and analyze model performance on unseen images (section 4). We then outline a proposed workflow for operational applications of trained models and discuss the outcomes of a preliminary case study conducted in the Austrian Alps (section 5). Our conclusion considers study limitations and directions for future work (section 6).

2. BACKGROUND

This section summarizes existing approaches to automated avalanche detection in visual data and discusses the potential that webcams offer as an underutilized data source for avalanche monitoring.

2.1. Related Work

Early work into automated avalanche detection focused on images captured by airborne sensors (Bühler et al., 2009; Lato et al., 2012). More recent studies have identified avalanches in satellite data using feature-based clustering (Vickers et al., 2016; Eckerstorfer et al., 2019) and CNNs (Kummervold et al., 2018; Sinha et al., 2019; Bianchi et al., 2021; Hafner et al., 2022) with a particular

*Corresponding author address:
James Fox, University of Innsbruck,
Innrain 52, 6020 Innsbruck, Austria;
tel: +43 512 507-0;
email: james.fox@student.uibk.ac.at

Data Source	Revisit Time (days)	Study
Leica ADS40-SH52	–	Bühler et al. (2009) Lato et al. (2012)
QuickBird	1 - 3.5	Lato et al. (2012)
Sentinel-1	6 - 12	Vickers et al. (2016) Kummervold et al. (2018) Sinha et al. (2019) Eckerstorfer et al. (2019) Bianchi et al. (2021) Kapper et al. (2023)
SPOT 6/7	1 [1]	Hafner et al. (2022)

Table 1: A comparison of existing methods for automated avalanche detection in visual data. Aside from the work of Lato et al. (2012) and Sinha et al. (2019), each model requires a DEM or a manual mask for every input image which greatly complicates application to new areas.

focus on backscatter values. A comparison of these methods and their associated data sources is presented in Table 1.

All of these studies utilize data sources that are not continuously available, making them unsuitable for ongoing avalanche monitoring. Moreover, most of these methods require a DEM or manual mask for each region of interest to exclude areas where avalanches are unlikely to occur. Although this filtering reduces the search space, it significantly complicates application, especially in new locations, as a suitable DEM or mask must be obtained and aligned with each image.

We have identified two distinct approaches to avalanche detection in previous work. In *classification*-based approaches, models are trained to classify whether an entire image patch contains any avalanches, whereas *segmentation* models seek to outline each individual avalanche in a pixel-wise manner. Consequently, image classification can be seen as a simpler subtask of avalanche segmentation. Although segmentation models require more detailed training data, they are then able to localize multiple avalanches in a single image.

2.2. The Potential of Webcams

Webcam imagery has been used for automated snow depth retrieval (Fromm and Adams, 2016), tracking snow cover evolution (Valt et al., 2013), and identifying snow and fog (Baumer et al., 2023). In the domain of avalanche monitoring, Hafner et al. (2023a) recently proposed a human-assisted approach to avalanche segmentation in webcam images, whereby computer-generated suggestions are iteratively corrected by a human annotator. Moreover, some avalanche warning services manually review webcam data to obtain an overview of the current avalanche situation (Stucki, 2006).

However, to the best of our knowledge, no existing studies fully automate avalanche detection in photographs taken from the ground. Our work, therefore, represents a novel approach in this domain. Provided that webcam images are available within a few minutes of an avalanche event, they hold the potential to expedite response times for avalanche events that pose risks to human life or critical infrastructure.

3. METHODOLOGY

We trained deep learning models to assess the feasibility of automated avalanche detection in ground-based visible-spectrum photography. A central objective was to ensure that the models could make predictions without a DEM or knowledge of the camera location, as this greatly simplifies their application in new areas of interest. The models were trained on a dataset of 4090 photographs taken in the winters of 2000/2001 - 2021/2022 at 1276 locations throughout the Alps.[2] Details about the dataset annotations and model training are provided below.

3.1. Data Annotation

Domain experts can identify the avalanche release mechanism in photographs due to distinctive visual features. We chose to label training images by avalanche release mechanism into the categories of glide, loose-snow, and slab avalanche (denoted by labels `GLIDE`, `LOOSE`, and `SLAB`). This multiclass annotation represents an extension of previous stud-

[1] Source: docs.sentinel-hub.com/api/latest/data/airbus/spot/ accessed on August 31, 2023.

[2] The dataset is available at researchdata.uibk.ac.at/records/h07f4-qzd17

ies which only consider the presence or absence of avalanches.

Labeling data by avalanche release mechanism has two advantages: firstly, this provides more contextual information during model training. Secondly, the resulting models are then able to predict the release mechanism in previously unseen images, offering more detailed insights into the avalanche situation and providing value for avalanche forecasting.

Each image in the dataset was manually annotated by avalanche experts,[3] as is common practice in the domain of avalanche detection (Hafner et al., 2023b; Bianchi et al., 2021). The annotators outlined and classified each visible avalanche and each image was subsequently assigned an *overall* label corresponding to the release mechanism with the most visible pixels (see Figure 1). A fourth overall label, `NONE`, was introduced to denote the 1071 images devoid of visible avalanches.

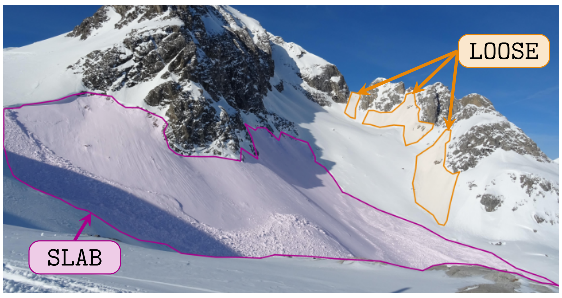

Figure 1: An enlarged dataset image containing a slab avalanche and three smaller loose-snow avalanches. Annotators were instructed to outline the entire avalanche from the release zone to the accumulation zone to encompass all relevant visual features. The image has overall label `SLAB` as this release mechanism has the most visible pixels.

3.2. Model Training

The dataset was split into 3612 training images and 478 images reserved exclusively for model evaluation. In order to prevent information from the test set being leaked into model training, images taken at the same approximate location were grouped into the same split. In each training run, ∼10 % of the training images were allocated for model validation and used to update the model parameters. Every training run was then repeated three times with different validation sets on a GeForce RTX 3090 GPU.[4]

Data augmentation techniques were applied during each training epoch to improve model generalization, including small variations in color and scaling and horizontal flips. Color and downslope orientation are important visual markers in avalanche images, so care was taken not to modify these excessively.

Previous work has explored both classification and segmentation-based approaches to automating avalanche detection, as discussed in subsection 2.1. For this reason, we decided to train classification and segmentation models on the same set of images to facilitate a numerical and qualitative comparison of model results. We utilized the ResNet architecture with 152 layers for image classification (He et al., 2016) and version three of the You Only Look Once (YOLO) network (Redmon and Farhadi, 2018) with spatial pyramid pooling (Huang et al., 2020) for avalanche segmentation.

3.2.1. Image Classification

For the image classification task, we trained ResNet152 on the overall image labels to assess the viability of avalanche detection in entire photographs. The network was pretrained on the ImageNet dataset and then finetuned for avalanche detection via transfer learning. We used the Adam optimizer (Kingma and Ba, 2015) with a learning rate of 2.25×10^{-3} to minimize the cross-entropy loss on the training set and cropped each image to a different random square each epoch before passing it to the network to improve model generalization. Test images were center-cropped for reproducibility.

3.2.2. Avalanche Segmentation

For the avalanche segmentation task, we trained YOLOv3-SPP on the avalanche bounding boxes. Training was carried out using an open-source framework from Ultralytics (Jocher, 2020) with an Intersection Over Union (IOU) threshold of 0.2 and a confidence threshold of 0.25.

4. RESULTS

This section contains numerical results and a qualitative analysis of predictions made by the trained classification and segmentation models.

4.1. Metric Definitions

We consider the standard definition of the metrics

$$\text{precision} = \frac{\mathtt{TP}}{\mathtt{TP} + \mathtt{FP}},$$

$$\text{recall} = \frac{\mathtt{TP}}{\mathtt{TP} + \mathtt{FN}},$$

$$\text{F1 score} = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}},$$

[3] "Experts" refers to people working in the snow and avalanche industry or conducting research in this field.

[4] The code used to train models is available at github.com/j-f-ox/avalanche-detection

Model	Task	Detection F1	Detection Recall	Multiclass F1	Multiclass Recall
ResNet152	Classification	92.9 ± 1.1	91.5 ± 2.7	81.5 ± 1.1	82.1 ± 0.9
YOLOv3-SPP	Segmentation	64.0 ± 0.6	58.1 ± 3.5	56.8 ± 0.4	53.6 ± 3.5

Table 2: Model results on the test set (mean ± standard deviation) for input size 896 px. The YOLOv3-SPP results are reported at IOU 0.05 and at the confidence threshold which maximizes the model F1 score.

where `TP` refers to the number of true positive classifications, `TN` to true negatives, `FP` to false positives, and `FN` to false negatives.

The concept of a "negative" avalanche prediction is represented by the image label `NONE` for the image classification task and by the lack of a predicted bounding box for the avalanche segmentation task. Scores are then calculated per image for classification models and per bounding box for the segmentation models.

A further consideration is that models were trained on data labeled by avalanche release mechanism into categories `GLIDE`, `LOOSE`, and `SLAB` (and `NONE` in the case of image classification). However, in scenarios such as rescue operations, the primary concern is the models' capability to detect avalanches regardless of whether the predicted release mechanism is correct. For this reason, we consider *multiclass* scores on the original dataset labels and *detection* scores for which the avalanche release mechanisms are combined. Note that the detection scores are, by definition, always greater or equal to the corresponding multiclass scores.

4.2. Test Results

Our models achieved mean avalanche detection F1 scores of 92.9 % per image and 64.0 % per bounding box on 478 previously unseen images (see Table 2). The lower numerical scores for segmentation are expected as this is a significantly harder subtask of the classification task. Although the multilabel scores were lower than those for avalanche detection alone, both models demonstrated the ability to distinguish between the avalanche classes, supporting our hypothesis that release mechanisms encode relevant information for avalanche detection.

Both models had the lowest recall for loose-snow avalanches (see Figure 2), possibly because this label had the lowest support in the dataset (369 training images). The addition of more loose-snow avalanche images to the dataset could address this disparity.

4.3. Qualitative Analysis

We conducted a comprehensive review of predictions made by the segmentation and classification models with the highest validation scores to investigate model capabilities and misclassifications. For the ResNet152 model, we employed Gradient-weighted Class Activation Mapping (Grad-CAM) heatmaps (Selvaraju et al., 2020) to visualize which image regions received the most attention. In the case of the YOLO model, we utilized IOU values to quantify the degree of overlap between the predicted bounding boxes and the ground truth.

The classification and segmentation models were both able to identify avalanches at a range of scales and in various lighting conditions, as is evident in Figure 3a and Figure 3b. However, the models were sometimes unable to locate smaller or less visible loose-snow avalanches, as in Figure 3c. Moreover, we observed that features such as buildings, ski tracks, fences, and large rock faces sometimes resulted in misclassifications. This problem could potentially be mitigated by deliberately augmenting the dataset with photographs containing these features.

Interestingly, the models sometimes detected very small avalanches or avalanches in shadow which had not been identified by the annotators. This highlights the potential that AI-assisted annotation offers to improve the quality of data annotation; an idea which is expanded on in section 5.

In summary, both the classification and segmentation models demonstrated promising results. Although further exploration is required to assess the practical utility of each approach, these encouraging findings underscore the potential of ground-based photography for automated avalanche detection.

5. OPERATIONAL APPLICATIONS

We believe that ground-based imagery will play a significant role in progressing towards real-time automated avalanche monitoring in the coming years. Although experts can identify avalanches in webcam images (Hafner et al., 2023a), manual monitoring becomes impractical at scale due to the vast number of possible avalanche locations and a scarcity of trained observers. Automating the analysis of webcam streams could therefore provide experts with a continuous overview of avalanche activity spanning entire regions in near real time. The Alps contain an extensive network of freely accessible mountain webcams, which could be integrated into such a monitoring system without incur-

(a) Results for ResNet152.

True \ Predicted	GLIDE	LOOSE	SLAB	NONE
GLIDE	90.1	1.6	4.6	3.7
LOOSE	5.7	68.8	9.9	15.6
SLAB	1.5	4.1	85.6	8.8
NONE	4.1	3.3	8.6	84.0

(b) Results for YOLOv3-SPP.

True \ Predicted	GLIDE	LOOSE	SLAB	Bg
GLIDE	50.2	0.6	2.6	46.6
LOOSE	0.7	34.2	3.5	61.6
SLAB	1.2	3.5	68.3	27.0
Bg	54.6	29.6	15.8	0.0

Figure 2: Normalized confusion matrices for the ResNet152 and YOLOv3-SPP models on the test set. For the YOLO model, “Bg” refers to the image background.

(a) A large slab avalanche viewed from a distance, including a long upper- and shorter lower crown fracture. The ResNet152 model often focused on the crown fracture of slab avalanches.

(b) A photograph containing four glide cracks which were identified by both models.

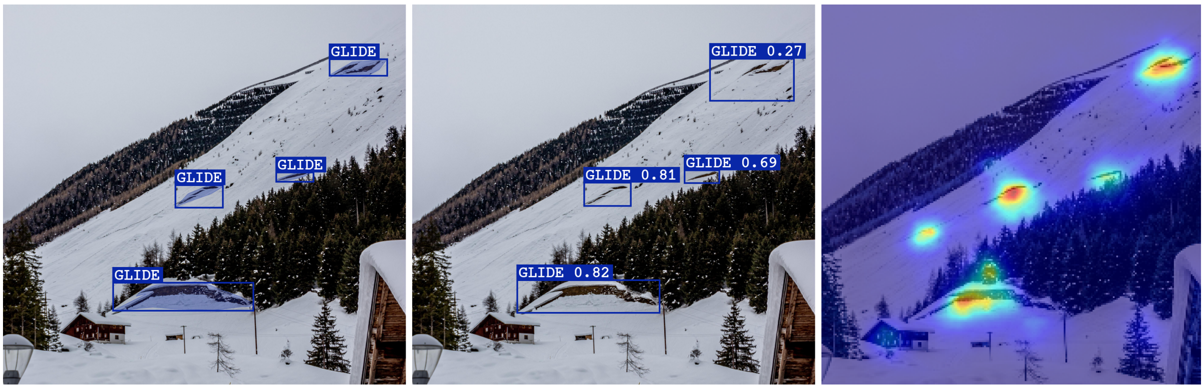

(c) An image with overall label `SLAB` containing two slab avalanches and four loose-snow avalanches. Both models identified the slab avalanches but struggled to detect the smaller loose-snow avalanches in the foreground.

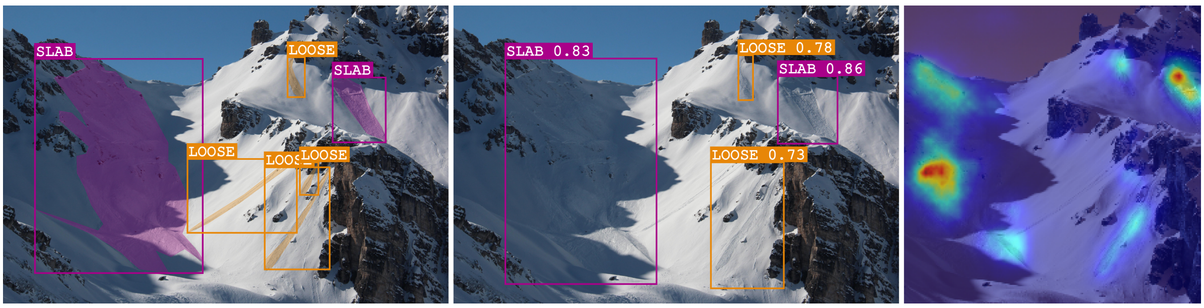

Figure 3: Example model predictions for previously unseen images. The ground truth is highlighted on the left, YOLOv3-SPP predictions with IOU values are shown in the center, and Grad-CAM heatmaps showing the ResNet152 attention are displayed on the right.

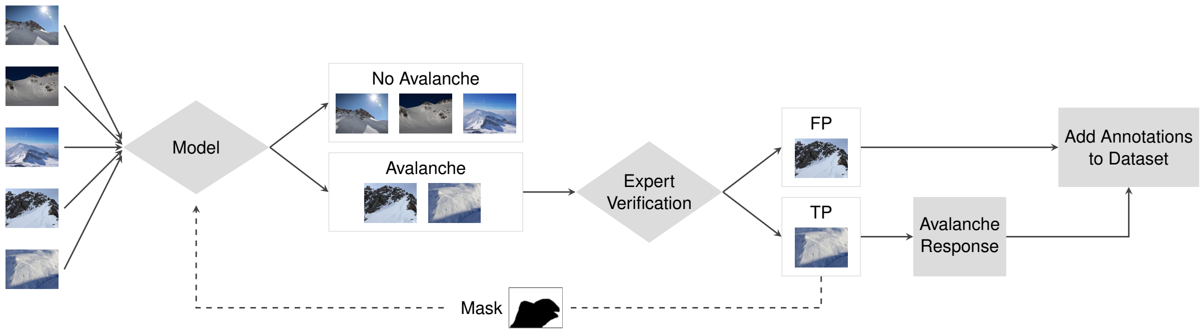

Figure 4: A possible workflow for trained models. A single model could monitor input webcam images and pass potential avalanche images to a human expert for verification. The expert could then feed confirmed avalanche outlines back into the model to prevent duplicate avalanche warnings. Retraining the model on corrected true positive (TP) and false positive (FP) data could iteratively improve model performance.

ring additional installation or maintenance costs.[5] Avalanche warning services could then install additional webcams in strategic locations (such as around critical infrastructure), focusing on areas with elevated avalanche risk.

In addition to expediting avalanche responses in remote areas, trained models could also enable the systematic collection of avalanche data over extended periods at higher temporal resolutions than are available from field observations or satellites. This data could then be cross-referenced with environmental conditions at the time of avalanche releases to improve avalanche forecasting models. The remainder of this section outlines a possible workflow for the operational application of trained models and findings from an initial case study in the Tyrolean Alps.

5.1. Proposed Workflow

Our vision for operational deployment of models centers around collaboration between humans and artificial intelligence (AI). At present, expert assessment represents the gold standard for avalanche detection in visual data (Hafner et al., 2023b). It would therefore be advisable to have a human expert verify AI-generated avalanche predictions and determine the most suitable course of action, especially considering the substantial resources required for avalanche response efforts.

In this collaborative approach, a model would continuously monitor numerous webcams and present potential avalanche images to an expert for assessment (see Figure 4). The confirmed or corrected annotations could then be saved and the model could be retrained on this expanded dataset, iteratively improving model performance over time. By automating the initial filtering process, such a workflow could enable a single observer to oversee avalanche responses on a regional scale.

When analyzing images from webcam streams, it is likely that the same avalanche will appear in subsequent frames. However, generating repeated warnings for the same avalanche during live monitoring may introduce undesirable redundancy. To prevent duplicated alerts, human experts could mask confirmed avalanches during prediction until the next precipitation. In this case, experts could review predicted bounding boxes (if using a segmentation-based approach) or attention heatmaps (if using classification models) to assist in localizing the avalanches within the image.

Training models for avalanche detection entails a natural trade-off between maximizing detections and minimizing false positive predictions. In the context of avalanche monitoring, false positives can be refuted easily by experts, whereas false negatives delaying the detection of an avalanche could have severe consequences. Future research could therefore explore deliberately biasing models to prioritize the reduction of false negatives.

5.2. Case Study

A case study was conducted in collaboration with the Tyrolean Avalanche Warning Service during March and April 2023 as a first step towards operational implementation. Webcam data from 30 locations across the Austrian Alps were analyzed to assess the robustness of trained models in operational settings. It is intended to continue developing this system during the next winter season.

During the case study, we observed that mountain huts were often misclassified as glide avalanches. Edges in the snow due to ski tracks or snow groom-

[5] For examples of existing webcam networks, refer to foto-webcam.eu (www.foto-webcam.eu/), feratel (www.feratel.com/en/webcams.html), and bergfex (www.bergfex.at/oesterreich/webcams/).

ing were also sometimes misclassified as slab avalanches. Augmenting the dataset with images containing these features mostly resolved this problem in preliminary experiments and should be considered in future dataset iterations.

A further observation from the case study was the pronounced impact of image scale on model performance. Results were significantly improved for higher-resolution images and for images where avalanches constituted a larger proportion of the image. In practice, segmenting webcam images into smaller tiles and analyzing each tile separately is likely to be necessary. However, due to the considerable variability in scale exhibited by avalanches and mountains, determining the optimal tile size remains open for further investigation.

6. CONCLUSIONS

This paper has described a novel approach to remote avalanche detection using webcams. Analyzing webcam streams with deep learning models could enable near real-time avalanche monitoring at higher temporal resolutions than existing satellite-based methods. Our models achieved F1 scores of 92.9 % for image classification and 64.0 % for avalanche segmentation, demonstrating the potential that ground-based photography offers to advance operational avalanche risk management.

A key advantage of our models is their ability to make predictions without a DEM or location-specific filter. This enables integration with pre-existing webcams with minimal overhead, making remote avalanche monitoring feasible at scale. In contrast to previous studies, we labeled data by avalanche release mechanism, thus enabling models to differentiate between glide, loose-snow, and slab avalanches. The remainder of the paper discusses study limitations and directions for future work.

6.1. Limitations

An intrinsic limitation of visible-spectrum photography is its requirement for relatively clear visibility and adequate ambient lighting. For truly continuous monitoring, webcam-based monitoring could be combined with either ground-based or satellite synthetic aperture radar (SAR) sensors. A further consideration is that most of the dataset images were captured in daylight and fair weather conditions. Evaluating the models on images captured in heavy precipitation, mist, partial cloud occlusion, or darkness would enable a more realistic assessment of model performance in real-world conditions.

Although utilizing expert annotations as the ground truth is standard practice in avalanche monitoring, this can introduce subjectivity in delineating avalanche boundaries (Hafner et al., 2023b) which may be seen as a limitation. To address this, annotators were provided with objective annotation guidelines. Future research could explore using multiple annotators for each image or integrating model predictions into the annotation pipeline (see Figure 4) to increase the consistency of annotations.

6.2. Future Work

In the upcoming winter season, we intend to carry out a more extensive case study on live webcam data which will contribute both to dataset expansion and informing future work. Further research could also investigate weighting models to prioritize recall over precision or the practical implications of classification versus segmentation-based approaches.

Models could additionally be extended with object detection to identify humans in webcam frames prior to an avalanche, allowing for improved response times in rescue efforts. Preprocessing techniques such as automated skyline detection or masking rocks could also be explored to improve model robustness.

Future research could ultimately progress towards establishing a system for continuous year-round avalanche monitoring using webcams and human-AI collaboration. In addition to expediting responses to avalanche incidents, such a system could also collect information on avalanche releases over extended periods, contributing to the advancement of avalanche modeling and forecasting. Although avalanches will never be entirely predictable, this preliminary study demonstrates the potential that ground-based imagery offers for automating avalanche detection.

ACKNOWLEDGEMENT

This research project received financial support from the Austrian Research Promotion Agency (FFG grant no. 888986). We thank the Avalanche Warning Service Tyrol and the Avalanche Warning Service Bavaria for providing images.

REFERENCES

Baumer, J., Metzger, N., Hafner, E. D., Daudt, R. C., Wegner, J. D., and Schindler, K.: Automatic Image Compositing and Snow Segmentation for Alpine Snow Cover Monitoring, in: 2023 10th IEEE Swiss Conference on Data Science (SDS), pp. 77–84, doi:10.1109/SDS57534.2023.00018, 2023.

Bianchi, F. M., Grahn, J., Eckerstorfer, M., Malnes, E., and Vickers, H.: Snow Avalanche Segmentation in SAR Images With Fully Convolutional Neural Networks, IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 14, 75–82, doi:10.1109/JSTARS.2020.3036914, 2021.

Bühler, Y., Hüni, A., Christen, M., Meister, R., and Kellenberger, T.: Automated detection and mapping of avalanche deposits using airborne optical remote sensing data, Cold Regions Science and Technology, 57, 99–106, doi:10.1016/j.coldregions.2009.02.007, URL https://www.sciencedirect.com/science/article/pii/S0165232X09000251, 2009.

Eckerstorfer, M., Bühler, Y., Frauenfelder, R., and Malnes, E.: Remote sensing of snow avalanches: Recent advances, potential, and limitations, Cold Regions Science and Technology, 121, 126–140, doi:https://doi.org/10.1016/j.coldregions.2015.11.001, URL https://www.sciencedirect.com/science/article/pii/S0165232X15002591, 2016.

Eckerstorfer, M., Vickers, H., Malnes, E., and Grahn, J.: Near-Real Time Automatic Snow Avalanche Activity Monitoring System Using Sentinel-1 SAR Data in Norway, Remote Sensing, 11, doi:10.3390/rs11232863, URL https://www.mdpi.com/2072-4292/11/23/2863, 2019.

Fromm, R. and Adams, M. S.: Spatio-temporal Snow Depth Estimates with Time-Lapse Photography, in: International Snow Science Workshop 2016 Proceedings, pp. 628–630, Breckenridge, CO, USA, URL https://arc.lib.montana.edu/snow-science/item.php?id=2335, 2016.

Hafner, E., Oberson, L., Kontogianni, T., Daudt, R. C., Wegner, J. D., Schindler, K., and Bühler, Y.: Using interactive object segmentation to derive avalanche outlines from webcam imagery, doi:10.5194/egusphere-egu23-10867, URL https://meetingorganizer.copernicus.org/EGU23/EGU23-10867.html, conference Name: EGU23, 2023a.

Hafner, E. D., Techel, F., Leinss, S., and Bühler, Y.: Mapping avalanches with satellites – evaluation of performance and completeness, The Cryosphere, 15, 983–1004, doi:10.5194/tc-15-983-2021, URL https://tc.copernicus.org/articles/15/983/2021/, publisher: Copernicus GmbH, 2021.

Hafner, E. D., Barton, P., Daudt, R. C., Wegner, J. D., Schindler, K., and Bühler, Y.: Automated avalanche mapping from SPOT 6/7 satellite imagery with deep learning: results, evaluation, potential and limitations, The Cryosphere, 16, 3517–3530, doi:10.5194/tc-16-3517-2022, URL https://tc.copernicus.org/articles/16/3517/2022/, 2022.

Hafner, E. D., Techel, F., Daudt, R. C., Wegner, J. D., Schindler, K., and Bühler, Y.: Avalanche size estimation and avalanche outline determination by experts: reliability and implications for practice, preprint, Other Hazards (e.g., Glacial and Snow Hazards, Karst, Wildfires Hazards, and Medical Geo-Hazards), doi:10.5194/egusphere-2023-586, URL https://egusphere.copernicus.org/preprints/2023/egusphere-2023-586/, 2023b.

He, K., Zhang, X., Ren, S., and Sun, J.: Deep Residual Learning for Image Recognition, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.

Huang, Z., Wang, J., Fu, X., Yu, T., Guo, Y., and Wang, R.: DC-SPP-YOLO: Dense connection and spatial pyramid pooling based YOLO for object detection, Information Sciences, 522, 241–258, doi:10.1016/j.ins.2020.02.067, URL https://linkinghub.elsevier.com/retrieve/pii/S0020025520301511, 2020.

Jocher, G.: YOLOv5 by Ultralytics, URL https://github.com/ultralytics/yolov5, 2020.

Kapper, K. L., Goelles, T., Muckenhuber, S., Trügler, A., Abermann, J., Schlager, B., Gaisberger, C., Eckerstorfer, M., Grahn, J., Malnes, E., Prokop, A., and Schöner, W.: Automated snow avalanche monitoring for Austria: State of the art and roadmap for future work, Frontiers in Remote Sensing, 4, 1156 519, doi:10.3389/frsen.2023.1156519, URL https://www.frontiersin.org/articles/10.3389/frsen.2023.1156519/full, 2023.

Kingma, D. P. and Ba, J.: Adam: A Method for Stochastic Optimization, in: 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings, edited by Bengio, Y. and LeCun, Y., URL http://arxiv.org/abs/1412.6980, 2015.

Kummervold, P., Malnes, E., Eckerstorfer, M., Arntzen, I., and Bianchi, F. M.: Avalanche detection in Sentinel-1 radar images using convolutional neural networks, in: International Snow Science Workshop Proceedings, Innsbruck, Austria, 2018.

Lato, M. J., Frauenfelder, R., and Bühler, Y.: Automated detection of snow avalanche deposits: segmentation and classification of optical remote sensing imagery, Natural Hazards and Earth System Sciences, 12, 2893–2906, doi:10.5194/nhess-12-2893-2012, URL https://nhess.copernicus.org/articles/12/2893/2012/, 2012.

Redmon, J. and Farhadi, A.: YOLOv3: An Incremental Improvement, CoRR, abs/1804.02767, URL http://arxiv.org/abs/1804.02767, arXiv: 1804.02767, 2018.

Schweizer, J. and Herwijnen, A. v.: Can Near Real-Time Avalanche Occurrence Data Improve Avalanche Forecasting?, in: Proceedings of the 2013 International Snow Science Workshop, pp. 195–198, Grenoble - Chamonix Mont-Blanc, URL https://arc.lib.montana.edu/snow-science/item.php?id=1789, 2013.

Schweizer, J., Bartelt, P., and van Herwijnen, A.: Snow avalanches, in: Snow and ice-related hazards, risks and disasters, pp. 395–436, Elsevier, Amsterdam, doi:10.1016/B978-0-12-394849-6.00012-3, URL https://www.dora.lib4ri.ch/wsl/islandora/object/wsl%3A12141/, 2015.

Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., and Batra, D.: Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization, International Journal of Computer Vision, 128, 336–359, doi:10.1007/s11263-019-01228-7, URL https://doi.org/10.1007/s11263-019-01228-7, 2020.

Sinha, S., Giffard-Roisin, S., Karbou, F., Deschâtres, M., Karas, A., Eckert, N., Coléou, C., and Monteleoni, C.: Can Avalanche Deposits be Effectively Detected by Deep Learning on Sentinel-1 Satellite SAR Images?, in: Climate Informatics, Paris, France, URL https://hal.archives-ouvertes.fr/hal-02278230, 2019.

Stucki, T.: Lawinenprognose. Wie entsteht ein Lawinenlagebericht - Möglichkeiten und Grenzen, in: Lawinen und Recht. Proceedings zum Internationalen Seminar vom 6.-9. November 2005 in Davos, pp. 21–34, Davos, Switzerland, URL https://www.dora.lib4ri.ch/wsl/islandora/object/wsl%3A16752/, 2006.

Valt, M., Salvatori, R., Plini, P., Salzano, R., Giusto, M., and Montagnoli, M.: Climate Change: A New Software to Study the Variations of Snow Images Shot by Web Cam, in: Proceedings of the 2013 International Snow Science Workshop, pp. 1004–1008, Grenoble - Chamonix Mont-Blanc, URL https://arc.lib.montana.edu/snow-science/item.php?id=1902, 2013.

Vickers, H., Eckerstorfer, M., Malnes, E., Larsen, Y., and Hindberg, H.: A method for automated snow avalanche debris detection through use of synthetic aperture radar (SAR) imaging, Earth and Space Science, 3, 446–462, publisher: Wiley Online Library, 2016.