Snow Depth Extraction From Time-Lapse Imagery Using a Keypoint Deep Learning Model

Catherine Breen[1], William Ryan Currier[2], Carrie Vuyovich[3], Zhongqi Miao[4], Laura Prugh[1]

1. Department of Environmental and Forest Sciences, University of Washington, Seattle, WA, USA. 2. Physical Sciences Laboratory, National Oceanic and Atmospheric Administration, Boulder, CO, USA. 3. Hydrological Sciences Laboratory, NASA Goddard Space Flight Center, Greenbelt, MD, USA, 4. AI For Good Research Lab, Microsoft, Seattle, WA, USA

Paper Link:

Summary

Current approaches for extracting snow depth from snow poles typically rely on time intensive manual photo processing. We present a method that uses a keypoint detection model to automatically observe snow height across a network of sites trained on 2021 SnowEx images. To assess model generalizability, we tested the model on sites in Washington state. We demonstrate that, especially when trained using a subset of site-specific data, a keypoint detection model can accelerate snow pole automation. This algorithm brings the hydrology community one step closer to a generalized snow pole detection model, and we call for a future model that integrates across time-lapse images from additional locations.

Objectives

1. Assess keypoint model ability for snow depth extraction on 2021 SnowEx data in Colorado (Fig. 1).
2. Assess for factors that influence model accuracy, establishing best practices for model success.
3. Fine tune model using images from timelapse camera network in Washington, identify number of images required for +/- 5 cm accuracy.

Fig. 1. Example image from Colorado (left) and Washington (right) show differences in lighting conditions and ecosystems.

Methods

Train model on 90% of 2021 SnowEx Colorado Images → w/ GPU → Test model on Colorado **and Washington images** → Fine-tune model using subset of Washington images → w/ CPU → Evaluate for improvement in accuracy with additional training data

Results

The initial model, the CO-model, performed well on Colorado images but poorly on the Washington images (Table 1; Fig. 2), but when fine-tuned on images from Washington, accuracy improves (Fig. 3; Figure 4). Shadows and patchy snow increase error.

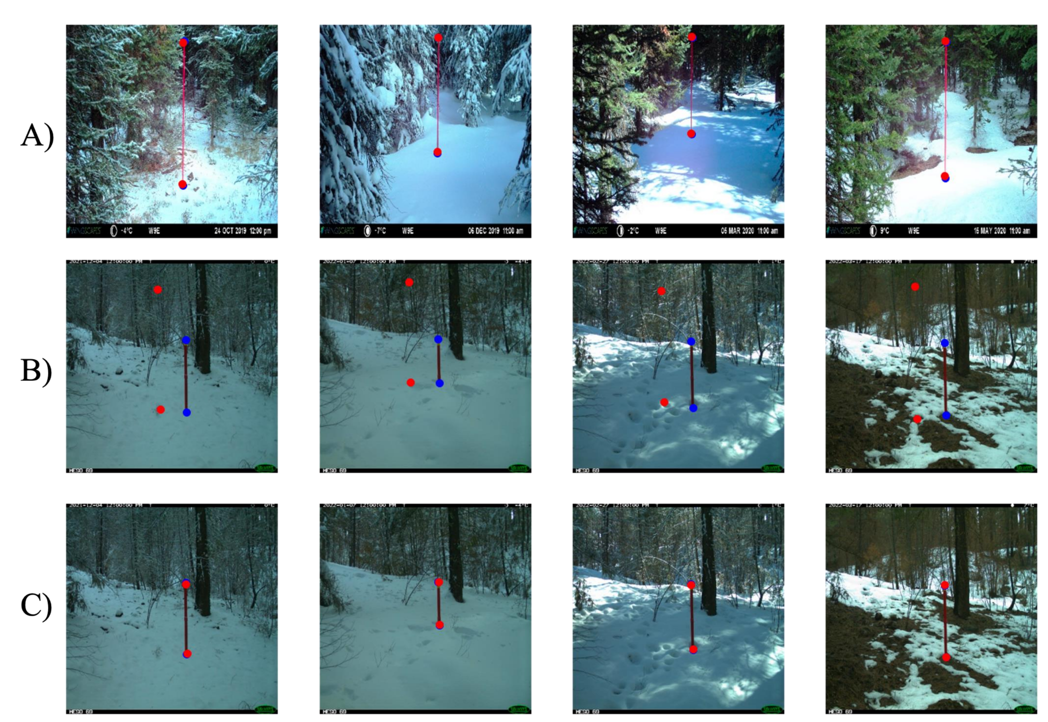

Fig. 2. Blue dots represent actual top and bottom locations, the red dot represents the predicted top and bottom point. (a) Shows Colorado site and (b) Washington site when the model is trained only on a 90% random subset of Colorado data. Row (c) shows that when a subset of WA images is included in training (i.e., the fine-tuned model), the model improves prediction of top and bottom.

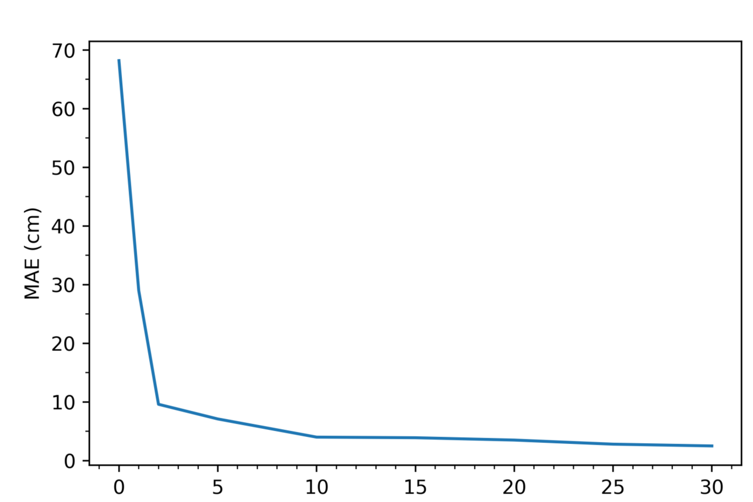

Fig. 3. Increasing images in fine-tuning step increases accuracy of model on novel data.

Table 1. Model, epoch, training size, region, and metrics for four different models.

Model	Epochs	Train size	Data region	Test size	RE (cm)	MAE (cm)	R^2
CO-only	35	8749	Colorado	972	-2.21 ± 2.96	1.30	0.99
			Washington	1770	68.21 ± 28.85	68.21	0.42
WA-only	100	120	Washington	1410	-2.07 ± 6.08	6.42	0.92
			Colorado	1749	-60.43 ± 41.18	60.85	0.44
Fine-tuned	42	120	Washington	1410	-1.35 ± 5.08	3.98	0.96
			Colorado	972	-78.0 ± 39.76	78.14	0.52
CO + WA	72	8917	Washington	170	-0.70 ± 1.30	1.14	0.99
			Colorado	972	-1.55 ± 2.09	2.05	0.99

Note. The CO-only model refers to the model trained on a random 90% subset of the Colorado images. The WA-only model refers to a model trained on a subset of Washington images. The "Fine-tuned model" refers to the Colorado model, fine-tuned with Washington images. The CO + WA model was trained on a combined 90% subset of the Colorado and Washington images.

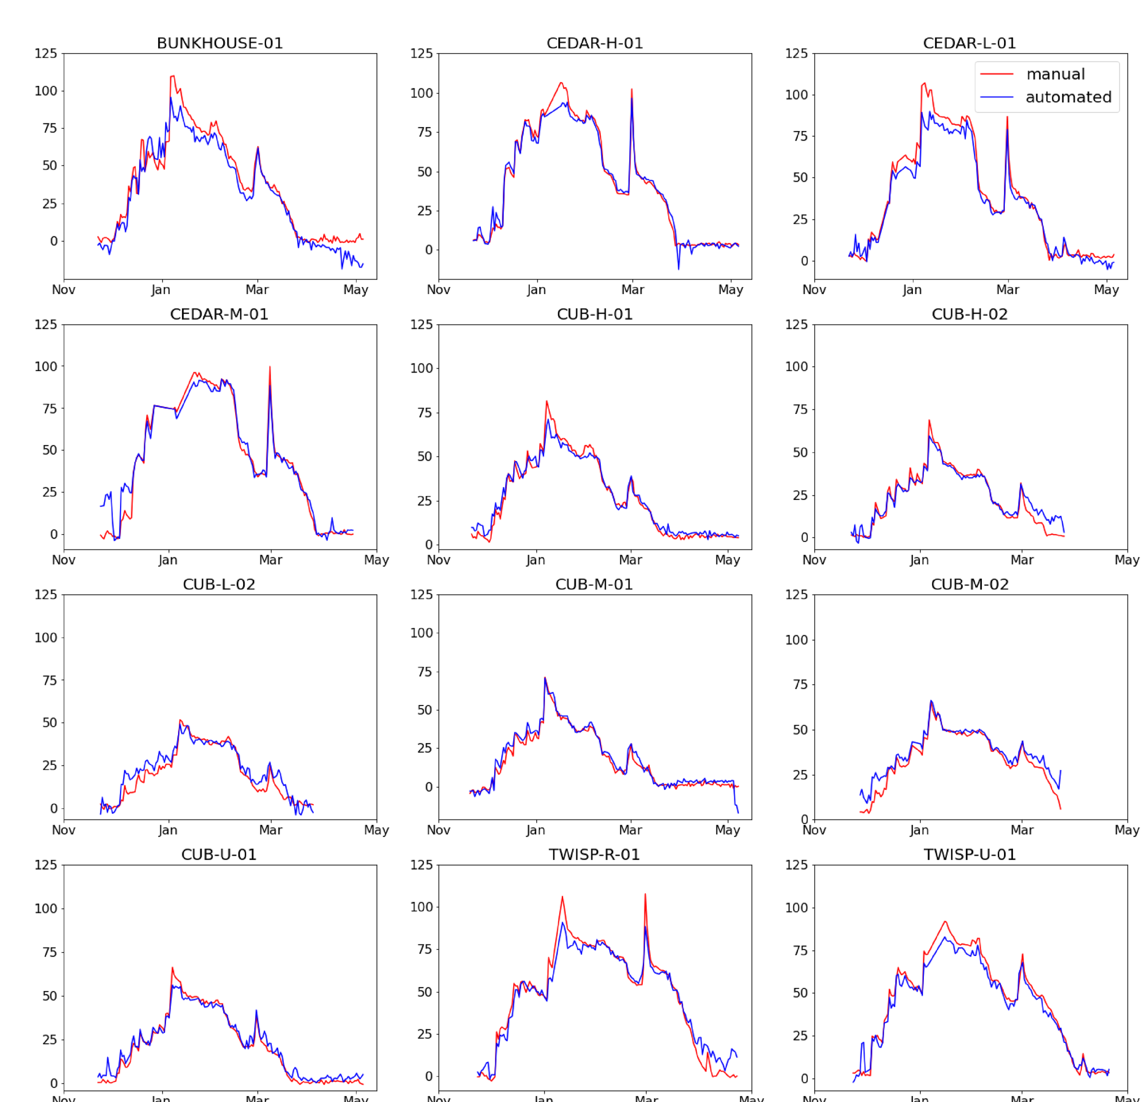

Fig. 4. Fine-tuned model can achieve accuracy of 3.98 cm ($R^2 = 0.96$) on novel data.

Codebase

Model demo and re-training available at:
https://github.com/catherine-m-breen/snowpoles

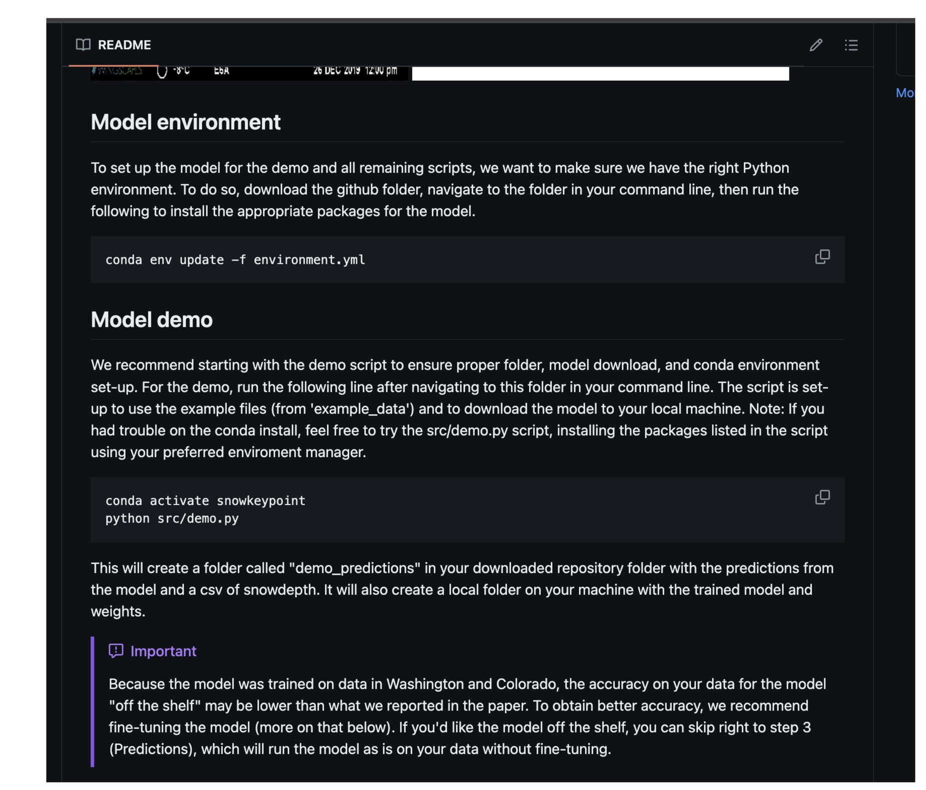

README

Model environment

To set up the model for the demo and all remaining scripts, we want to make sure we have the right Python environment. To do so, download the github folder, navigate to the folder in your command line, then run the following to install the appropriate packages for the model.

```
conda env update -f environment.yml
```

Model demo

We recommend starting with the demo script to ensure proper folder, model download, and conda environment set-up. For the demo, run the following line after navigating to this folder in your command line. The script is set-up to use the example files (from 'example_data') and to download the model to your local machine. Note: If you had trouble on the conda install, feel free to try the src/demo.py script, installing the packages listed in the script using your preferred enviroment manager.

```
conda activate snowkeypoint
python src/demo.py
```

This will create a folder called "demo_predictions" in your downloaded repository folder with the predictions from the model and a csv of snowdepth. It will also create a local folder on your machine with the trained model and weights.

> Important
>
> Because the model was trained on data in Washington and Colorado, the accuracy on your data for the model "off the shelf" may be lower than what we reported in the paper. To obtain better accuracy, we recommend fine-tuning the model (more on that below). If you'd like the model off the shelf, you can skip right to step 3 (Predictions), which will run the model as is on your data without fine-tuning.

Conclusions

- A keypoint detection model to automate snow depth measurements can achieve a mean absolute error (MAE) equal to 1.14 cm when trained on sites of interest.
- Model shows accuracy for snow depth both in and out of canopy locations and throughout the winter season.
- When "fine-tuned" using 10 images per camera, the model achieves accuracy within 4 cm on a novel network comprising 12 cameras.
- Fine-tuning can occur with CPU, and codebase is available on GitHub.

Acknowledgments

We thank Chris Hiemstra, Yasmine Hentati, Ben Sullender, and Daniel Hogan, for help installing poles and time-lapse cameras. We also thank Alicia Joseph for SnowEx data access. Jessica Christley and Carl Sack provided image processing support for SnowEx data labeling. Dan Morris provided advice on method development and access to cloud computing resources. Members of the Prugh lab provided valuable feedback on manuscript drafts. This work was enabled in part by participation in CV4Ecology Summer Workshop, supported by the California Resnick Sustainability Institute. This research was supported by NASA Graduate Fellowship Program Grant 80NSSC19K1673 as well as Microsoft's AI for Earth grant program (ID: 00010003316).