

Snow Depth Measurement via Time Lapse Photography and Automated Image Recognition

Kevin S. J. Brown
Steven R. Fassnacht

Department of Ecosystem Science and Sustainability
Colorado State University

January 2019

CWI Completion Report No.233



Colorado
State
University

Acknowledgements

We thank the Natural Resources Conservation Service Colorado Snow Survey Program for their logistical support of the study sites collocated at Snow Telemetry stations; the partners working on the Upper Gunnison Basin Watershed Function Science Focus Area project for their support of the Crested Butte Pumphouse site; and the Center for Snow and Avalanche Studies for their support of the Swamp Angel Study Plot.

This research was undertaken by Colorado State University. The Results and their Implications do not necessarily reflect the opinions of the Colorado Water Institute or the Colorado Water Conservation Board.

This report was financed in part by the U.S. Department of the Interior, Geological Survey, through the Colorado Water Institute. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Government.

Additional copies of this report can be obtained from the Colorado Water Institute, E102 Engineering Building, Colorado State University, Fort Collins, CO 80523-1033 970-491-6308 or email: cwi@colostate.edu, or downloaded as a PDF file from <http://www.cwi.colostate.edu>.

Colorado State University is an equal opportunity/affirmative action employer and complies with all federal and Colorado laws, regulations, and executive orders regarding affirmative action requirements in all programs. The Office of Equal Opportunity and Diversity is located in 101 Student Services. To assist Colorado State University in meeting its affirmative action responsibilities, ethnic minorities, women and other protected class members are encouraged to apply and to so identify themselves.

ABSTRACT

Seasonal snow is a crucial component of water supply in Colorado and the western United States. Measurement of snow accumulation through the winter and spring allows water managers to forecast water supply for the growing season and take actions to ease flooding and drought. The Natural Resources Conservation Service's (NRCS) snow telemetry (SNOTEL) network provides real-time data at a high cost per station and at single points. An evaluation of existing field measurements of snow depth taken in 2009 and 2010 was undertaken to determine if fine resolution depth measurements are justified. Fassnacht et al. (*in press*) showed that the snow depth variability can be substantial even at fine resolution. However, these data required extensive labor to collect and only represented one measurement in time. A low-cost method to measure snow variability around these stations or in underrepresented areas could improve snow forecasts by quantifying the representativeness of data from the current network. To this end, we trialed a method combining time lapse photography and computer vision techniques to find snow depth at five sites in Colorado during water year 2018. Different site configurations were trialed, and a best operating procedure was determined. The data gathered were not more accurate than current ultrasonic or laser snow depth measurement technologies. However, the low cost and versatility of this method may make it more applicable in certain situations.

KEYWORDS

Image Recognition, Snow Depth, Measurement, Photogrammetry, Spatial Variability

This page was intentionally left blank

Table of Contents

JUSTIFICATION OF WORK PERFORMED	1
REVIEW OF METHODS USED	4
DISCUSSION OF RESULTS AND THEIR SIGNIFICANCE	10
PRINICIPAL FINDINGS, CONCLUSIONS, AND RECOMMENDATIONS.....	12
SUMMARY	13
REFERENCES	14
APPENDIX A: Site Layout Detail	16
APPENDIX B: Further Detail on Image Recognition Algorithm	21

JUSTIFICATION OF WORK PERFORMED

Snow Depth as an Indicator of Snow Water Equivalent

When quantifying water storage in the snowpack, snow water equivalent (SWE) is the most straightforward measure and varies substantially at the watershed scale (López-Moreno et al., 2013). However, snow depth is easier to measure and can be converted to SWE if snow density is known (Elder et al., 1991). Density usually varies less than snow depth or SWE and tends to increase through the season (López-Moreno et al., 2013). The spatial pattern of snow depth can be modeled across individual watersheds at a low temporal resolution using landscape characteristics (Meiman, 1968), such as elevation, slope, aspect, canopy density, etc. The combination of spatial snow depth and density will yield SWE distributed across a watershed or area of interest. The ability to measure snow depth at remote sites across a watershed at a low cost could build toward a basin-wide estimation of SWE from snow depths (Rice and Bales, 2010).

Snow Depth Measurement

Snow depth is the most straightforward snowpack property to measure. Snow monitoring in the Western US started with James Church at Mount Rose, Nevada in 1905 (NRCS, 2014). This led to the manual snow survey program by the Soil Conservation Service in the 1930s. These monthly snow course measurements have been supplemented by the automated snow telemetry (SNOTEL) network that started in the late 1970s. The SNOTEL stations initially measured SWE and precipitation, and subsequently temperature and snow depth measurements were added (NRCS, 2014).

Snow depth can vary significantly around SNOTEL stations (Fassnacht et al., *in press*) and snow courses (Fassnacht and Hultstrand, 2015). Additional automated depth measurements in the area can capture variations in due to terrain variables that are not well represented by point measurements (Rice and Bales, 2010). Better characterizing this variability, even on a relatively small scale, can improve estimates of SWE estimates across varied terrain.

In the winters of 2009 and 2010, Fassnacht et al. (*in press*) completed multiple snow depth surveys over 1km² areas centered on two SNOTEL stations. The surveys reiterated that snow depths vary greatly over small scales, but also collected snow depth points in a variety of configurations. By subsampling the points collected, the best point configuration for closely approximating the local variability was found (Figure 1; Fassnacht et al., *in press*). For an acceptable error of 5%, three of 11 points were on average sufficient. However, this did vary from one to 11 points across each sampling domain. Similar results were found across a different domain (Figure 2; Fassnacht et al., 2017). Use of multiple measurements can provide an estimate of variability (Fassnacht and Hultstrand, 2015; Fassnacht et al., *in press*). Additionally, the terrain and land cover characteristics associated with snow depth variation could be determined to assist in planning future sampling strategies.




Figure 1: Mean absolute difference from the pixel mean snow depth for all samples per pixel as a function of the number of points per pixel for the three sampling dates using 11 points in a row within one digital elevation model pixel. The 5% difference threshold is shown as the dotted line [after Fassnacht et al., *in press*].




Figure 2: Representativeness of snow depth measurements versus the “true” mean (from all 11 points) shown as a function of number of percent of the total number of measurement location (42) versus the number of samples (one to 11) per measurement location, with 1% and 5 % acceptable differences/error for the two land cover types (evergreen and non-evergreen) [from Fassnacht et al., 2017]).

Automated Image Recognition in the Context of Snow

Recently there have been a number of studies interested in using automated image recognition to determine snowpack properties. Most of these efforts have been focused on albedo determination starting with Corripo (2004), but a few have dabbled in using relative pixel color of exposed depth staffs determine snow depth. Garvelmann et al. (2012) worked to find snow depth via this technique but did not achieve satisfactory results. The application of recent advances in image recognition and computer vision could improve the accuracy of snow depth obtained from images via automated processes.

REVIEW OF METHODS USED

Field Site Setup

The premise of site setup was to place a weatherproofed time-lapse camera and a number of depth staffs in its field of view. This setup has been used by other researchers before but with the images examined manually to extract snow depth. This is quite time consuming. The images produced from our sites were to be processed automatically. This did not necessitate any special considerations on the site configuration aside from ensuring the staffs not occlude each other from the camera's point of view.

We used consumer-grade game cameras due to their availability and low price. The Stealthcam G34 12 MP camera (stealthcam.com) was selected as it had good weatherproofing, long battery life, and high resolution at a low cost. Additionally, these cameras maintained an internal clock and wrote metadata into the image files, simplifying subsequent sorting and categorization. The cameras were mostly collocated at SNOTEL sites. They were attached to the SNOTEL meteorological tower or T-post and pointed with unobstructed sightlines towards a large open area. (Figure 3)



Figure 3: Game camera mounted T-post at Center for Snow and Avalanche Studies' Swamp Angel study plot, a non-SNOTEL site.

Previous studies had used a variety of materials to construct the snow depth staffs, including wood boards and PVC pipe. We selected regular steel T-posts as they offered stability, durability, and simple installation with the major drawback of weight. At several sites, 10-foot T-posts were required due to the historical record of deep snow. Before deployment, the posts were painted red to increase contrast. Initial trials of image capture with different colors at varying distances suggested that red provided the best contrast against a vegetated background.

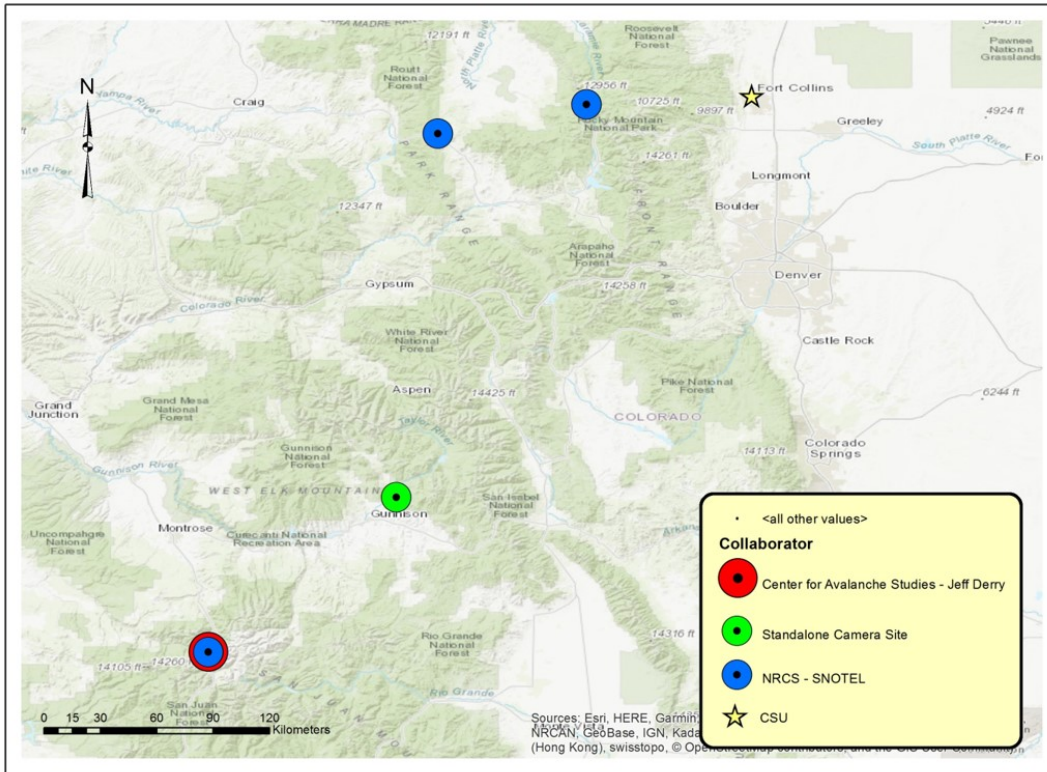
A different configuration of staffs was used at each site (Appendix A), with the key criteria being that all of the staffs were in view of the camera and did not obstruct each other. The number of staffs placed would increase image processing time. The configurations at each site were partly informed by space and sightline constraints, but were also modelled on previous snow depth surveys (Kashipazha, 2012; Meromy et al., 2013; Fassnacht et al., 2017; Fassnacht et al., *in press*). The site layouts are presented in the appendix, and elaboration on the performance of each configuration is discussed in the recommendations section.

Site Locations

Six sites were placed around Colorado for the winter and spring of 2017-2018 (Figure 4). Four of them were placed at SNOTEL sites, which had safe winter access and a tower to place the camera on. The other two sites were collocated with other ongoing experiments; one was at the Colorado Snow and Avalanche Center (snowstudies.org) Senator Beck Basin near Silverton, and the other was on the East River near Crested Butte. In addition, this work was an opportunity to collaborate and share data with other researchers. Installing sites around the state helped overcome the local impact of a low snowpack, in particular the dry snow year in the southern part of Colorado.

Snow Depth Camera Sites

Prepared by Kevin Brown (kbrown137@colostate.edu)



Coordinate System: GCS North American 1983
Datum: North American 1983
Units: Degree

Figure 4: Site Locations.

Site Maintenance and Data Retrieval

An important limitation at these sites is that they are non-telemetered. The images must be retrieved from the cameras manually. Images were taken every 30 minutes and were around 4 MB each. The batteries lasted 6-8 weeks, and the SD cards were large enough to accommodate the data. Site visits were undertaken every 5-6 weeks, which provided an opportunity to check in on the condition of the equipment. If data acquisition is required more often, more frequent site visits could be necessary.

Image Recognition Algorithm

Development of computer vision technology has been rapid in recent years, with applications including self-driving cars, kinesiology, traffic engineering, and psychology. While numerous image recognition techniques exist, they all use two stages: training and detection (Figure 5). The training stage uses images with known locations of target objects to ‘teach’ a detector. Then, that detector can be used with unlabeled images to find locations of target objects. More details are provided in Appendix B.




Figure 5: General image recognition workflow

This detector outputs bounding boxes. This bounding box is defined by pixel coordinates of its corners, which were converted to snow depth. This process is fairly straightforward and requires an image of the staff with no snow and the actual length of the staff exposed above ground. Once these are known, the pixel length of the bounding box in the image can be compared to the pixel length of the box with no snow, then scaled by the total height of the staff above ground with no snow (Equation 1).

$$d_s = L_0 - \frac{p_i}{p_0} L_0 \quad \text{or} \quad d_s = \left(1 - \frac{p_i}{p_0}\right) L_0$$

Equation 1: Equation to retrieve exposed staff length from pixel measurements. Where d_s is snow depth, p_0 is pixel length of staff with no snow, p_i is pixel length of exposed staff in image i , and L_0 is length of exposed staff with no snow.

Running an image through the detection step usually produces 50-100 bounding boxes, each with a confidence score assigned by the algorithm. These boxes are scattered all over the image, with most of the low-scoring returns being false positives (Figure 6). The question of how to sort through all these results and assign the final snow depth to each individual staff was one of the most difficult problems overcome during development. Each return could correspond to any of the staffs and most of the time they correspond to none.




Figure 6: Raw detection result.

In the end, the detected returns were split into lists based on which of the boxes in a baseline image were closest. Then, each of these lists was ranked by score, with only the top score of each list being accepted as the bounding box for that staff. For certain staffs in some images there are no positive returns nearby, or the best return nearby still has a very low confidence score. In these cases the bounding boxes were thrown out for that particular image, like the example in Figure 7.




Figure 7: Processed detection result.

The training and detection stages are more computationally intensive with larger images. The Stealthcam produced high resolution images at 3000x4000 pixels, which meant that training and testing new detectors was time-consuming. In the end, trainings with several hundred images took 20-40 minutes, while the detection step using the resulting detector took about 0.5 seconds for each image. With camera images taken every 30 minutes during daylight producing about 600 images per

month per site, this processing can also be time consuming but is still far more efficient than manually delineating those images.

DISCUSSION OF RESULTS AND THEIR SIGNIFICANCE

Image recognition accuracy

The accuracy of the image recognition algorithm was tested by using a validation set of images with supplied correct bounding boxes. The difference between the staff length found through a bounding box and the actual staff length was averaged for each site to produce a percent deviation (Figure 8).




Figure 8: Deviations in measured depth from actual depth in the verification dataset, by site

The difference in performance between sites was massive but was underlain by differences within sites. Some staffs at individual sites were consistently more recognizable to the algorithm than others. Those that were further away from the camera and more centered in the frame were more

frequently detected accurately. A higher fraction of the background being filled by snow versus trees or bare ground also greatly increased accuracy.

Swamp Angel had a 10 staff configuration with staffs placed close together, greatly decreasing algorithm efficiency. Red Mountain Pass had staffs placed on a sloping hill, with forest covering most of the background. Pumphouse was a compact site, with staffs taking up most of the field of view. The other two were similar in setup and the algorithm was much better at picking out the staffs consistently. Joe Wright's ~5% deviation is equivalent to about 6 inches on a 9-foot exposed staff. Averaging data from individual images from an entire day can cut down on this deviation somewhat.

Snow Depth Data




Figure 9: Early season snow depth data at Joe Wright for two staffs with no smoothing

The early processed data has a high amount of noise, but trends are still visible. The discrepancy between the two depth staffs in Figure 9 is interesting given that the two staffs were only five meters apart. The depths also varied widely from the SNOTEL depth measurement only 25 meters away. The evolution of these discrepancies during accumulation and melt are indicative of the high spatial variability of snow depth and SWE.

Economics

Off-the-shelf hardware to measure snow depth include Ultrasonic Depth Sensors (e.g., Judd and Campbell SR50) and Laser Snow Depth Sensor (Lufft SMH30). These units cost between \$700 (Judd) and \$2800 (Lufft) and require an additional data logger and a meteorological tower. Yet, each of these sensors measures snow depth at a single point. The use of an automated camera system to estimate multiple snow depths offers major advantages in cost, with the total setup for a camera, SD cards, batteries, and five 10ft depth staffs totaling about \$300.

PRINCIPAL FINDINGS, CONCLUSIONS, AND RECOMMENDATIONS

The Need

As suggested in previous work (Neumann et al., 2006; Rice and Bales, 2010), multiple snow depth measurements provide a more robust estimate of snow depth than a single measurement. The assessment of existing field datasets (Fassnacht et al., 2017; Fassnacht et al., *in press*) illustrated that depending upon the acceptable error, at least three measurements were necessary to provide a reasonable snow depth estimate at a location (Figures 1 & 2). At all study sites, at least five stakes were installed for the estimation of snow depth.

Site Setup

The exact configuration of the sites is very adaptable. Any arrangement of staffs and camera where all staffs have a clear sightline to the camera produce usable data. That being said, placing staffs extremely close together did interfere with the image recognition and, counterintuitively, staffs that were placed further from the camera and took up less space in the image had more consistent detections than those placed closest. In addition, it is strongly recommended that future sites be set up so that the camera faces roughly north whenever possible. Glare washed out many of the images to the point of them being unusable, and in one case, consistent reflection of the low angled sun from the snow caused minor damage to the image sensor.

The Stealthcam cameras we used could run for 5-6 weeks with 8 AA batteries in the main case but have an input for external power. This raises the possibility of placing these sites in one location and letting them run for the entire winter and returning in the spring to collect the data.

Image Recognition

The ACF image recognition algorithm performed well in this application. Other methods will likely produce similar results, but the ACF method is optimized for color images. Its daily averaged accuracy for the most successful sites is generally around 2-4 inches. Sites prone to high amounts of glare or frequent fog have greatly decreased accuracies, as the returns are less consistent. Future work should explore other methods in order to increase the algorithm's accuracy for staffs that take up a large portion of the field of view, which even the final trained detector had difficulty recognizing. Manual measurements from the images are still the most accurate, but they are quite time consuming, so automated methods may be preferred when there are a large number of data points.

Future Work

In the near-term work will continue on improving the image recognition algorithm for sites where it is currently performing poorly. Further work will also examine new locations where these sites could be placed to provide useful data for other research efforts. They could be used to verify remotely sensed snow data or provide estimates of SWE in underrepresented areas. Deployment of several sites in close proximity could also give insight into local-scale snow depth variability at a small time step.

SUMMARY

At the onset of this work, an analysis of existing field data (Fassnacht et al., 2017; Fassnacht et al., *in press*) provided the impetus for the use of game cameras to provide multiple snow depth measurements at a location. Previous work has manually extracted snow depth, which can be quite

time consuming. Automating this process is possible using new computer vision techniques, raising the possibility of efficient, low-cost snow depth measurement via photogrammetry.

REFERENCES

- Corripio, J. G., 2004. Snow surface albedo estimation using terrestrial photography, *Int. J. Remote Sens.*, 25, 5705-5729
- Elder, K., J. Dozier, and J. Michaelsen, 1991. Snow accumulation and distribution in an alpine watershed. *Water Resources Research*, 27(7), 1541-1552, [doi:10.1029/91WR00506].
- Fassnacht, S.R., and M. Hultstrand, 2015. Snowpack Variability and Trends at Long-term Stations in Northern Colorado, USA. *Proceedings of the International Association of Hydrological Sciences (Hydrologic Non-Stationarity and Extrapolating Models to Predict the Future)*, 371, 131-136, [doi:10.5194/piahs-371-131-2015].
- Fassnacht, S.R., D. Wyss, and S.M. Heering, 2017. Räumliches Denken – ein Beispiel im Schnee / A Spatial Thinking Research-Didactic Example in Snow. *Geoöko*, 37(3-4), 231-247.
- Fassnacht, S.R., K.S. Brown, E.J. Blumberg, J.I. López Moreno, T.P. Covino, M. Kappas, Y. Huang, V. Leone, and A.H. Kashipazha, in press. Distribution of Snow Depths Variability. *Frontiers of Earth Science*, 12(4), 10 pages [doi: 10.1007/s11707-018-0714-z].
- Garvelmann, J., Pohl, S., Weiler, M., 2013. From observation to the quantification of snow processes with a time lapse network. *Hydrol. Earth Syst. Sci.*, 17, 1415-1429, [doi: 10.5194/hess-17-1415-2013]
- Kashipazha, A.H., 2012. Practical snow depth sampling around six snow telemetry (SNOTEL) stations in Colorado and Wyoming, United States. Unpublished M.S. thesis, Watershed Science Program, Colorado State University, Fort Collins, Colorado USA.
- López-Moreno, J.I., S.R. Fassnacht, J.T. Heath, K. Musselman, J. Revuelto, J. Latron, E. Morán, and T. Jonas, 2013. Spatial variability of snow density over complex alpine terrain: implications for estimating snow water equivalent. *Advances in Water Resources*, 55, 40-52, [doi, 10.1016/j.advwatres.2012.08.010].
- Meiman, J.R., 1968. Snow accumulation related to elevation, aspect, and forest canopy. In *Proceedings of Snow Hydrology Workshop Seminar*, Canadian National Committee of the International Hydrological Decade, Fredericton NB, pp. 35–47
- Meromy, L., N.P. Molotch, T. Link, S.R. Fassnacht, and R. Rice, 2013. Subgrid variability of snow water equivalent at operational snow stations in the western USA. *Hydrological Processes*, 27(17), 2383-2400, [doi:10.1002/hyp.9355].
- Neumann, N., C. Derksen, C. Smith, and B.E. Goodison, 2006. Characterizing local scale snow cover using point measurements during the winter season. *Atmosphere-Oceans*, 44(3), 257-269, [doi:10.3137/ao.440304].



NRCS, 2014


Rice, R., and R.C. Bales, 2010. Embedded sensor network design for snow cover measurements around snow pillow and snow course sites in the Sierra Nevada of California. *Water Resources Research*, 46(3), W03537, [doi:10.1029/2008WR007318].

APPENDIX A: Site Layout Detail


Measurements in all images are in meters.

Key:

Camera	
Staff	



Joe Wright SNOTEL



Columbine SNOTEL



Red Mountain Pass SNOTEL



Pumphouse/East River



Swamp Angel/CSAS

APPENDIX B: Further Detail on Image Recognition Algorithm

The images supplied are usually sample images with bounding boxes drawn around the object or objects of interest. The algorithm will take this image set and identify features that differentiate the target objects from the background. After enough training images have been ingested, the algorithm will produce a detector, a set of weights and protocols that when plugged into a 'detect' function can pick out the target object from a scene. The training step is very computationally intensive but with high variability based on the type of algorithm. The detection step is where the detector object is applied to new images with no training data and outputs the corners of a shape circumscribed around positive detections (Figure 6).

There are two broad types of image recognition that should be discussed with relation to this project. There are numerous specific types of algorithms that fall into each category. The first category are systems designed to identify object types in a cluttered scene. In self-driving cars, this technology is used to differentiate different types of road signage, crosswalks, or other relevant details from a complex background. The other class of methods are specialized to pick out an individual type of object, such as a face, from a background while ignoring anything else. For picking out our uniformly colored and shaped snow staffs from a background, the latter was used.

The specific image recognition technique selected for the project is the Aggregate Channel Features (ACF) method. ACF recognition was first developed by Dollar et al. (2014) and has been implemented in the Matlab (www.mathworks.com) programming environment. Its basic mechanism takes the channels of an image (red, green, and blue), performs an array of transformations on them, and finds distinctive features in the results that are associated with the supplied positive training examples. The detector object generated from this process weights the transformations and features most frequently found around the object of interest. Importantly, the detector also looks at features that indicate the object of interest is not present in an area of the image. When the detector is applied to images with no training data, the algorithm picks out distinctive features throughout the image, and the ones most indicative of the target object's presence during the training stage will cause a positive detection.

In practice, the detector was trained from a selected subset of the game camera images, which had the bounding boxes drawn in manually. The training dataset eventually exceeded 400 images from all sites and produced a capable detector. There are a number of parameters that can be changed in both the training and detection steps, and detectors were generated with subsets of these images to determine the optimal values for this application.