

# All-Sky Image Segmentation of Clouds Using a Convolutional Neural Network

By

Maxwell Levin and Zoe Harrington

Advisors: Peter Drake, Jessica Kleiss, Erin Riley

An honors thesis written in partial fulfillment of the requirements  
for the degree of Bachelors of Arts

Lewis & Clark College  
Portland, Oregon

May 2019

## Abstract

The fraction of the sky covered by clouds has been a key meteorological observation since the 1800s. It is particularly useful for weather forecasting, climate model development, and solar energy forecasting. The cloud cover can be estimated from manual observation of the sky, active and passive ground-based and satellite remote sensing, and sky image analysis. Our work presents a novel derivation of fractional sky cover from all-sky images using a convolutional neural network. Our network achieved a final validation accuracy of 89.085%. Compared to the fractional sky cover obtained by the Total Sky Imager our network's measurement of the fractional sky cover had a root mean squared error that was 8.721% lower for *typical* data and 22.785% lower for *dubious* data. Our results show improved consistency with active remote sensing cloud fraction for problematic sky images.

## Introduction

Clouds play an important role in Earth's climate by modulating both incoming solar radiation and outgoing terrestrial radiation. Clouds can contribute to both a warming and a cooling influence on local climates, depending upon their altitude, optical thickness, aerosol concentration, and water phase. Furthermore, cloud cover and thickness can be altered by changed atmospheric aerosol concentrations and play a role in both positive and negative feedbacks with changing climate. The positive and negative feedbacks and the cloud changes due to aerosols currently provide the greatest source of uncertainty in climate model projections (Boucher et al., 2013).

Climate models typically cannot resolve individual clouds within a climate grid cell, so the cloud coverage is described statistically, typically by liquid water path or fractional cloud coverage. Climate models are assessed for accuracy by comparing hindcast models with observations. However, model-observation disagreement could be due to bias in either the model or the observation. Improved model predictions of future climate critically depends upon consistent and high-quality observations of cloud cover for model assessment.

The earliest quantitative observations of clouds were obtained by meteorologists recording the fractional sky cover (FSC) and cloud type from ground-based visual observations. FSC is the fraction of the observable sky covered by clouds at a given time. Manual observations of FSC were recorded at meteorological stations worldwide until the 1990s, when many meteorological stations began using automated ceilometers to determine cloud fraction (Dai et al., 2006). Ceilometers are pencil-beam lidars that record the presence of cloud at a point in the sky at high frequency (about 20 kHz), and return up to ten cloud base heights with temporal accuracy of 10 s. Temporal averages over longer time scales (typically 1 hour) result in the cloud fraction (CF), the fraction of time a cloud is directly overhead during a given averaging period. In recent decades, all-sky cameras such as the Total Sky Imager (TSI) (Long et al., 2006), Whole Sky Imager (WSI) (Feister et al., 2013), UCSD Sky Imager (USI) (Yang et al., 2014), and many other sky imaging rigs developed by research groups worldwide have presented the opportunity of recovering FSC observations comparable to manual observations. Sky imagers offer a strictly quantitative and reproducible estimate of FSC, and are not subject to differences between individual manual observers (Huo and Lu, 2012; Liu et al., 2013). However, digital images alone do not discern the

difference between low-, high-, thin-, and thick-clouds and clear sky. Towards this end, much work has been done to develop methods to segment sky images. Here *segmentation* refers to the process of dividing the sky image into multiple meaningful and distinct regions (Zhu et al., 2016). See Tapakis and Charalambides (2013) for an excellent review of image processing approaches to sky image segmentation.

## **Machine Learning Approaches to Cloud Image Segmentation**

Machine learning models offer an alternative way to tackle challenging or ill-defined problems. Whereas traditional programming techniques rely on an extensive prior knowledge to a problem, machine learning techniques are able to learn from the ground up and improve their performance on problems with practice. Machine learning models such as artificial neural networks are inspired by actual neural networks and loosely mimic their functionality through the implementation of layers of interconnected neurons. Within a layer, each neuron computes a function of some of the output values from the previous layer. For example, a neuron might multiply each previous output value by a weight, add them together, and then output the result if it is positive, zero otherwise. Neurons that are connected have a weight that is associated with the strength of the connection between them. For a supervised machine learning model, these weights are updated during the training process through a combination of gradient descent and backpropagation techniques. In its essence, a machine learning model is a large mathematical function that can be fine-tuned to adapt to the problem at hand.

Machine learning models have been widely successful in areas of computer vision such as image classification, object detection, and semantic segmentation (Zhu et al. 2016, Liu et al. 2018). One method that has had enormous success in these areas is convolution. Fully convolutional networks (FCNs) are machine learning models that are made up entirely of convolutional and pooling layers (Long et al., 2014). A convolutional layer is a collection of “filters” which slide over the layer’s inputs and apply a dot-product at each location. Pooling layers are similar to convolutional layers in that they apply sliding “filters” to the layer’s inputs, but instead of a dot product, these layers typically take the maximum value in the filtered region at each location in the layer’s input. Shelhamer et al. (2016) showed that FCNs can successfully be applied to the task of image segmentation by improving on the previous best results in image segmentation datasets.

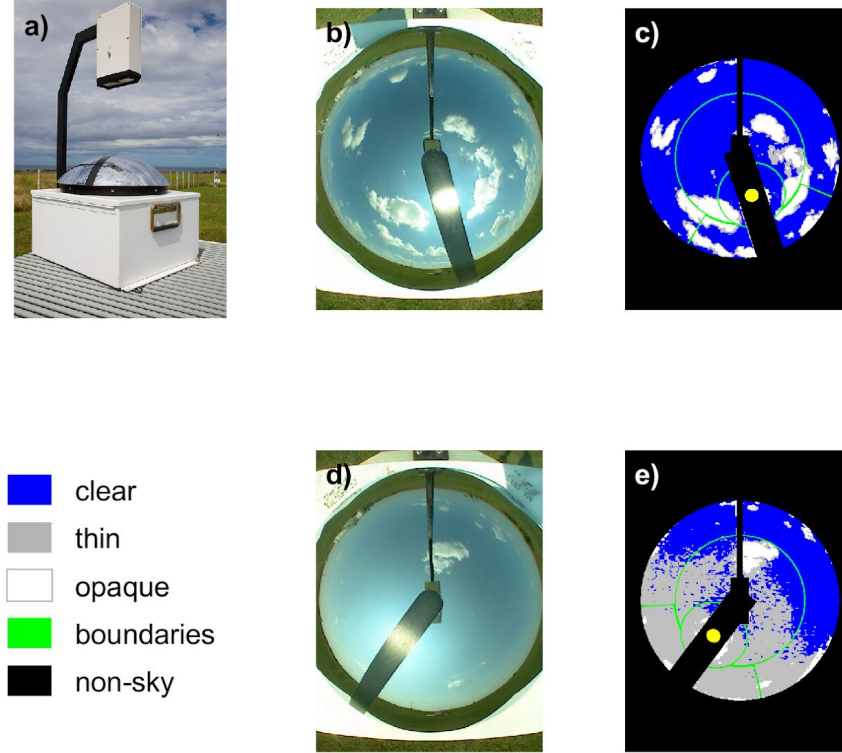
Other recent machine learning approaches to the task of sky image segmentation include the use of feature extraction (Guo et al., 2017), ensembles of binary classifiers (Dev et al., 2016), a graph-cut algorithm (Liu et al., 2015), superpixels (Liu et al., 2014, Dev et al., 2017), multilayer perceptrons and support vector machines (Taravat et al., 2015), and random forests and Bayesian classifiers (Cheng & Lin, 2017). See Dev et al. (2016) for a more detailed summary of recent machine learning applications in the field of sky image segmentation.

## **Methods**

### **Sky Image Dataset**

The sky image dataset we used was obtained from a Total Sky Imager (TSI) operating in the Southern Great Plains (SGP) site operated by the Atmospheric Radiation Measurement (ARM) program (Ackerman & Stokes 2003). The TSI consists of a digital camera suspended above a domed mirror that reflects a 360° image of the sky (Figure 1a). A hemispheric sky image (e.g. Figure 1b) is captured every

30 seconds during daylight hours. Images are 24 bit RGB images with a dimension of 480 x 640 pixels. A shadowband is affixed on the domed mirror to block the reflection of the sun. A servo-motor in the instrument housing causes the mirror rotate to align the shadowband with the sun throughout the day.



**Figure 1:** a) The Total Sky Imager (TSI). (b,c) example *typical* TSI sky image (30-Aug-2016 17:55:30) and corresponding TSI decision image, (d,e) example *dubious* TSI sky image (31-Aug-2015 19:48:00) and corresponding TSI decision image.

Onboard processing creates a decision image corresponding to each sky image, following the thresholding image segmentation technique in Long et al. (2006). The ratio of the red to blue channel of the image is computed for each image pixel, which is then thresholded to classify the pixel as clear sky, thin cloud, or opaque cloud. The fractional sky cover is computed as the ratio of cloudy pixels to total pixels. The results of the classification are generally in good agreement with other instrumental observations of cloud coverage (Boers et al. 2010, Wu et al. 2014, Wacker et al. 2015). However, there are known problems with accurate cloud detection near the sun and the horizon due to enhanced image brightness. Long et al. (2010) presented a correction for the fractional sky coverage to account for these errors, and these problem regions are commonly excluded from the TSI FSC product.

Shallow cumulus clouds, also known as fair weather cumulus clouds, are convective clouds with limited vertical development (Stull, 1985). Shallow Cumulus form over marine (Zhao and Di Girolamo, 2007) and continental environments (Berg and Kassianov, 2008), and play an important role in earth's radiative budget. Shallow cumulus clouds are generally white with well-defined boundaries and blue sky in between clouds. By definition, shallow cumulus clouds occupy a fraction of the sky, such that it is neither clear nor overcast conditions. We leverage these important properties of shallow cumulus to create our experimental design.

The possible times of shallow cumulus between May and September of the years 2012-2017 were determined from the Shallow Cumulus Value Added Product (VAP) (ARM 2000, Lim et al., 2019), which indicates hour-long periods that may have shallow cumulus. These time periods were divided into 15-minute segments, for which the average TSI opaque-cloud-only fractional sky cover within the inner 100° field-of-view was computed (Kassianov et al., 2005). Radar / lidar data from the Ka-band Zenith Radar Active Remote Sensing of Clouds (KAZRARSCL) VAP was averaged over a 30-minute period centered on the TSI interval to provide cloud base height (CBH), cloud top height, low-cloud fraction (with CBH < 3 km), and total cloud fraction. These time periods were further discriminated based on the following criteria 1) All valid radar/lidar measurements for the entirety of the 30-minute period. This reduces the data to 13,385 time periods. 2) The radar / lidar cloud fraction from low altitudes (CBH < 3 km) from must be within 0.1 of the total cloud fraction from all altitudes in order to remove cases with upper-level clouds. This reduces the data to 9406 time periods. The remaining data were then divided into two classes based on the 15-minute average TSI thin-cloud fractional sky coverage ( $FSC_{thin}$ ): *Typical* images have 15-min  $FSC_{thin} < 0.3$ , *Dubious* images have 15-min  $FSC_{thin} \geq 0.3$ .

This results in 102,005 *typical* images (with 30 images in each 15-minute period), and 4744 *dubious* images. Cumulus clouds have very high contrast (bright white) so their uncertainty (thin cloud) is expected to be low. Cases with high thin cloud amount correspond to the slow degradation of the TSI camera that increases the red channel. The agreement of TSI FSC with CF from the radar / lidar product is better for *typical* images (Figure 5a) than for *dubious* images (Figure 5b).

We allocate 60% of the available *typical* images for training our network, 20% for validating our network during the training process, and reserve the remaining 20% for a final test set which we save for a later date. This results in roughly 60,216 sky and decision image pairs in our training set and 20,071 sky and decision image pairs in our validation set for *typical* data.

Additionally, we allocate 60% of the available *dubious* images for validating our network's performance after training, and save the remaining 40% for a test set to be used on a later date. This results in about 2846 *dubious* sky images and decision pairs being used in our *dubious* validation set. No training is done on *dubious* data.


All images were cropped from 640x480 to 480x480 pixels to eliminate many non-sky pixels and speed up training. The training outputs were further preprocessed by removing (setting to black) any pixels that indicated the sun location within the shadowband or that were green to indicate sky regions (Figure 1).

## The Convolutional Neural Network

We constructed an artificial neural network to segment sky images. The input to the network consists of the red, green, and blue channels of the sky image. The output of the network is a similar image, but with a channel for each category of pixel: blue (clear sky), gray (thin cloud), white (thick cloud), or black (non-sky, such as the reflection of the TSI shadowband). Thus the function of the network is similar to that of the thresholding algorithm of the TSI.

Our network consists primarily of a series of five convolutional layers (Figure 2). The input to the network is the preprocessed sky image with 3 channels (red, green, blue). The first three layers have 32 channels each; the final output layer has 4 for the different pixel classifications (clear sky, thin cloud, thick cloud, and non-sky). All convolutional kernels are 3x3, following the suggestion of Simonyan & Zisserman (2015). The final convolutional layer takes input from the previous layer and also the original

input image, echoing He et al., (2015). Two large max pooling layers help detect large features such as the TSI's shadowband. 100x100 pooling would be too computationally costly, but we were able to get acceptable results with one layer of 1x100 pools and one of 100x1 pools. All layers are 480x480; this is padded with zeros where kernels in the following layer would spill over the edge of an image.



**Figure 2.** Schematic of the design of our convolutional neural network. The network is composed of five convolutional layers and two max pooling layers. The network utilizes several skip connections, notably a skip connection from the input layer to the final convolutional layer (output).

Whenever our network is run as a classifier, either in training or validation, pixels that are *always* (across the entire data set) marked black by the TSI are automatically marked black. This spares the network the burden of detecting the camera reflection or the corner regions that show the TSI housing rather than the sky.

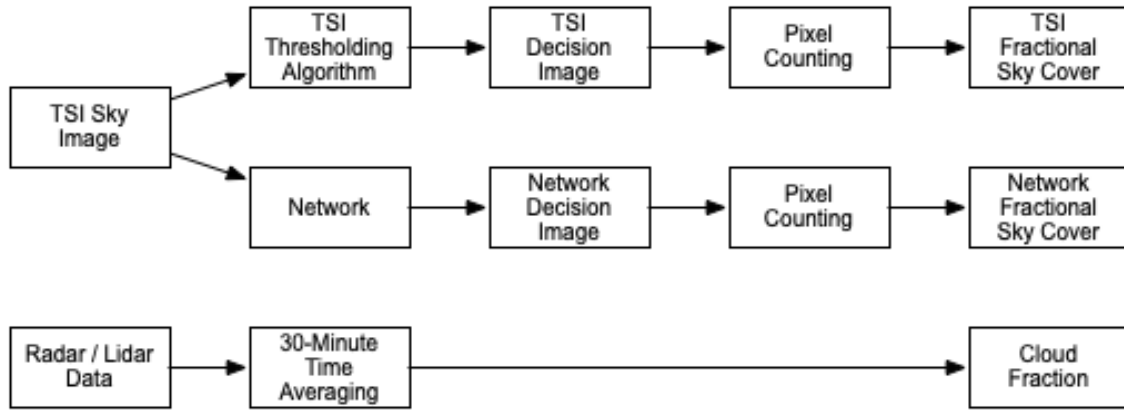
The network was constructed using Tensorflow for GPU and run on a GTX 1080 ti with 11GB of graphics memory. We found that the maximum batch size we could use with this setup is 23 sky images, so training was conducted in batches of 23 images. At each training step, a random batch of 23 images is selected from the training set and every sky image in the batch is processed by our network. The network's decision image output for each sky image input is compared to the corresponding TSI decision image in order to calculate error. Error here is defined as the number of pixels in the network output that differ from the TSI decision image. We take the average of the errors across the outputs in this batch of images and use this to perform gradient descent, an optimization strategy designed to reduce the model's error by modifying the strength of connections between neurons in the network.

The network's error on the training and validation set is monitored and reported every 10 training steps. The network is set to train for 10,000 batches. The validation accuracy represents the model's ability to generalize what it has learned to data that it has not seen before. To prevent the model from overfitting the training data (essentially memorizing the data) and help it generalize, early stopping is implemented using the model's performance on the validation data. That is, if no accuracy gains in the

validation data have been seen for 300 batches then training is cut short and the version of the model which obtained the highest validation accuracy during training is saved.

## Experimental Design

After training is complete, the network is used to process all available *typical* training and validation sky images and all available *dubious* validation sky images into decision images. From these network-processed decision images the FSC is calculated in the zenith region as the percentage of thick and thin cloud pixels divided by the total number of not non-sky pixels. To evaluate the performance of the network the computed FSC is compared with the corresponding TSI FSC values and the radar / lidar *cloud fraction* (CF) values. The cloud fraction is used as an independently obtained measurement of the same observable sky to verify the accuracy of the decision image's FSC. One of our advisors found that for TSI decision images with a low percentage (<30%) of thin clouds the FSC as measured by the TSI is frequently in agreement with the CF recorded by the KAZRARSCL product. An example of this comparison is shown in Figure 5.




**Figure 3.** Diagram of the data flow. The TSI Sky Image is used by both the TSI Thresholding Algorithm and the Network to create decision images. Analysis of these decision images yields an estimate for FSC. The Radar / Lidar data is averaged over a period of time to yield the cloud fraction (CF) statistic. The CF statistic is later used as a measure of ground truth for comparing the FSC values.

The data flow of the experiment is best visualized by Figure 3.

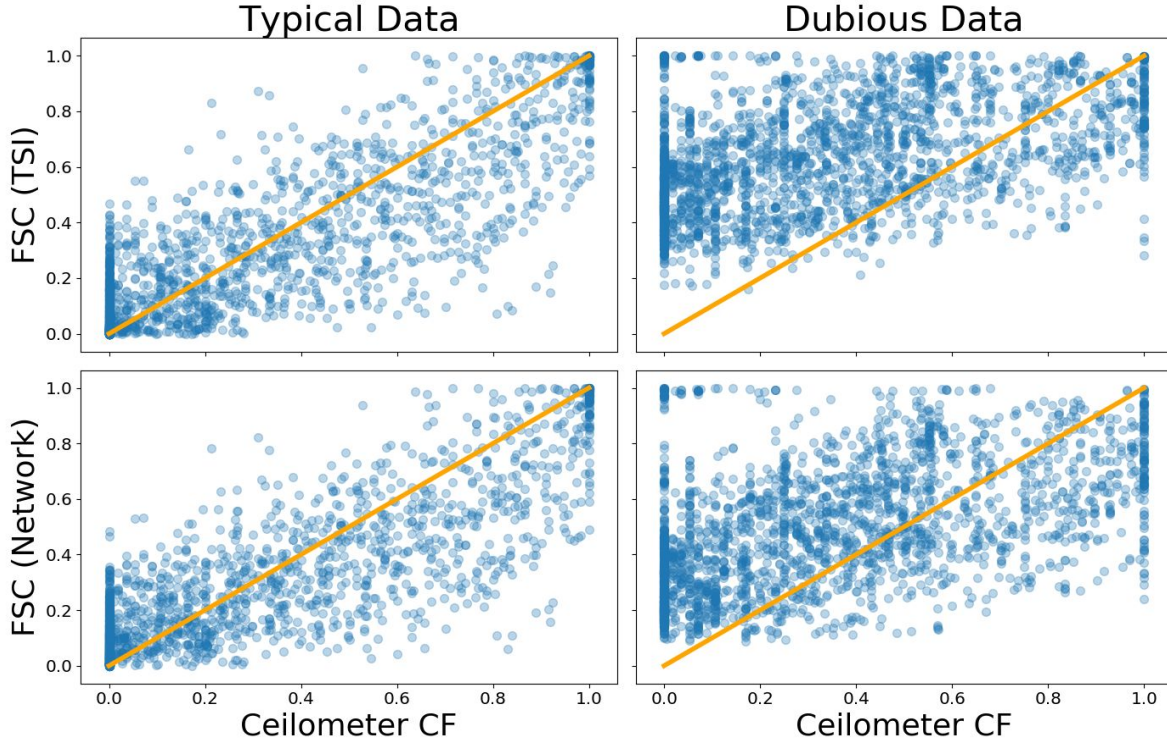
## Results

Figure 4 displays the network's performance on *typical* training and validation data during the training stage of the experiment. Because early stopping was used the network stopped training after 300 batches without improvement in the validation accuracy, training for a total of 1200 batches. The network achieved a final training accuracy of 88.913% and a final validation accuracy of 89.085%.



**Figure 4:** Accuracy of network versus batches of images. In this case the network trained on batches of 23 images for 1200 batches. The network achieved a final training accuracy of 88.913% and a final validation accuracy of 89.085%.






**Figure 5.** Top row from left to right: a) TSI compared to radar / lidar cloud fraction for *typical* data. b) TSI compared to radar / lidar for *dubious* data. Bottom row from left to right: c) Neural network output compared to radar / lidar for *typical* data, and d) Neural network output compared to radar / lidar for *dubious* data.

Figure 5 displays a collection of scatter plots comparing FSC to CF for both the TSI and the network output on a sample *typical* and *dubious* validation data. The orange one-to-one line marks perfect agreement between CF and FSC. The blue points represent the FSC value on the y-axis and the corresponding radar / lidar cloud fraction on the x-axis for each randomly selected observation.

Figure 6 reports the Root Mean Squared Error (RMSE) for each scatter plot in Figure 5. Error is calculated as the difference between CF and FSC for each point. Thus each reported RMSE in Figure 6 is the square root of the average of the squared errors across each collection of data points. By nature of RMSE, the RMSE is on the same scale as the original data values used in its calculation. Since FSC and CF both range from 0 to 1, so does the RMSE. We found that on typical data the TSI has a RMSE of 0.172 and the Network has a RMSE of 0.157. On dubious data the TSI has a RMSE of 0.395 and the Network has a RMSE of 0.305.



**Figure 6:** The root mean squared error (RMSE) between FSC and CF observations for a sample of *typical* and *dubious* validation data. On *typical* data the TSI has a RMSE of 0.172 and the Network has a RMSE of 0.157. On *dubious* data the TSI has a RMSE of 0.395 and the Network has a RMSE of 0.305.

## Discussion

Figures 5 and 6 indicate that the network we developed produces outputs similar in FSC to those of the TSI for *typical* validation data. In Figures 5a) and 5c) we see that the distributions are similar in shape, and loosely follow the one-to-one line with the CF data. In Figure 6 we see that the RMSE values for *typical* data are very close, with the network producing a slightly lower RMSE, indicating slightly better agreement with the CF than the TSI with the CF. Further investigation is required to determine if this difference in RMSE values is significant.

For *dubious* validation data the network produces outputs which are less similar to those of the TSI. In Figures 5b) and 5d) we see that the shape of the distribution of the network's FSC vs. CF data has changed slightly so that it is closer to the orange one-to-one line. This is confirmed by the result in Figure 6 for *dubious* data. Here we find that the network produces a much lower RMSE value than the TSI, indicating better agreement with the CF for *dubious* validation data. The difference in RMSE between the TSI and the network in these cases is likely due to algorithmic differences in how the network and the TSI compute decision images.

The network's FSC values generally seem to be lower than the TSI's FSC values and more in agreement with the CF, potentially indicating that our network handles problematic conditions such as sun glare or condensation better than the TSI. Further investigation is required in this matter.

Future work involving our neural network could explore how applicable our neural network is to cloud types other than shallow cumulus. For instance, it could be trained on decision images of additional clouds types including clear and overcast skies. However, the neural network must not be trained on low quality decision images to avoid replication of errors in the TSI's algorithm. Our work only investigated shallow cumulus in order to differentiate between correct and problematic automated segmentation.

Correctly identifying the shadowband and removing it from analysis was a major hurdle and achievement in this work.

Our analysis demonstrates improved FSC retrievals for cases with camera reddening, resulting in elevated thin cloud FSC as measured by the TSI algorithm. Future work could also address other errors that occur — for example sun glare. With a labelled collection of correct and questionable segmented images the neural network could be trained and tested to improve upon known segmentation errors. The TSI segmentation algorithm uses only the red-to-blue ratio, and the neural network also uses the three color channels as input. Improving the neural networks performance in future work could be achieved through including other sources of information, such as the collocated lidar / radar CF, during the training phase. This could possibly help the neural network identify and improve the segmentation of problematic images.

Future work could explore a broader range of neural network designs, to investigate in how far out neural network design shown in Figure 2 affects the results presented in this paper. A number of network architectures were explored in development of this work. Dramatic improvement was obtained after inclusion of the long max pooling layers, which enabled the correct identification and removal of the shadowband.

## References

- Ackerman, T. P., & Stokes, G. M. (2003). The atmospheric radiation measurement program. *Physics Today*, 56(1), 38–44. <https://doi.org/10.1063/1.1554135>
- Atmospheric Radiation Measurement (ARM) user facility. (2000), updated hourly. Fair-Weather Shallow Cumulus Identification (Shallow Cumulus). 2012-05-01 to 2017-08-31, Southern Great Plains (SGP) Central Facility, Lamont, OK (C1). Compiled by K. Lim, L. Riihimaki and Y. Shi. ARM Data Center. <http://dx.doi.org/10.5439/1392569>.
- Berg, L. K., & Kassianov, E. I. (2008). Temporal Variability of Fair-Weather Cumulus Statistics at the ACRF SGP Site. *Journal of Climate*, 21(13), 3344–3358. <https://doi.org/10.1175/2007JCLI2266.1>
- Boers, R., de Haij, M. J., Wauben, W. M. F., Baltink, H. K., van Ulf, L. H., Savenije, M., & Long, C. N. (2010). Optimized fractional cloudiness determination from five ground-based remote sensing techniques. *Journal of Geophysical Research*, 115(D24). <https://doi.org/10.1029/2010JD014661>
- Boucher, O., D. Randall, P. Artaxo, C. Bretherton, G. Feingold, P. Forster, V.-M. Kerminen, Y. Kondo, H. Liao, U. Lohmann, P. Rasch, S.K. Satheesh, S. Sherwood, B. Stevens and X.Y. Zhang, (2013): Clouds and Aerosols. In: *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

- Cheng, H.-Y., & Lin, C.-L. (2017). Cloud detection in all-sky images via multi-scale neighborhood features and multiple supervised learning techniques. *Atmospheric Measurement Techniques*, 10(1), 199–208. <https://doi.org/10.5194/amt-10-199-2017>
- Dai, A., Karl, T. R., Sun, B., & Trenberth, K. E. (2006). Recent Trends in Cloudiness over the United States: A Tale of Monitoring Inadequacies. *Bulletin of the American Meteorological Society*, 87(5), 597–606. <https://doi.org/10.1175/BAMS-87-5-597>;
- Dev, S., Wen, B., Lee, Y. H., & Winkler, S. (2016). Machine Learning Techniques and Applications For Ground-based Image Analysis. ArXiv:1606.02811 [Cs]. Retrieved from <http://arxiv.org/abs/1606.02811>
- Feister, U., Shields, J., Karr, M., Johnson, R., Dehne, K., & Woldt, M. (2000). Ground-based cloud images and sky radiances in the visible and near infrared region from whole sky imager measurements. EUMP31, EUMETSAT Satellite Application Facility Workshop, Dresden, Germany, 20–22. Retrieved from [http://www.mpl.ucsd.edu/people/jshields/publications/dd\\_feist.pdf](http://www.mpl.ucsd.edu/people/jshields/publications/dd_feist.pdf).
- Guo, Yanming & Liu, Yu & Georgiou, Theodoros & S. Lew, Michael. (2017). A review of semantic segmentation using deep neural networks. *International Journal of Multimedia Information Retrieval*. DOI: <https://doi.org/10.1007/s13735-017-0141-z>.
- He, K., Zhang, X., Ren, S., & Sun, J. (2015). Deep Residual Learning for Image Recognition. ArXiv:1512.03385 [Cs]. Retrieved from <http://arxiv.org/abs/1512.03385>
- Huo, J., & Lu, D. (2012). Comparison of Cloud Cover from All-Sky Imager and Meteorological Observer. *Journal of Atmospheric and Oceanic Technology*, 29(8), 1093–1101. <https://doi.org/10.1175/JTECH-D-11-00006.1>
- Kassianov, E., Long, C. N., & Christy, J. (2005). Cloud-Base-Height Estimation from Paired Ground-Based Hemispherical Observations. *Journal of Applied Meteorology*, 44(8), 1221–1233. <https://doi.org/10.1175/JAM2277.1>
- Lim et al. (2019). “Long-term cloud type retrieval using a combination of active remote sensors and a total sky imager at the ARM SGP site,” submitted.
- Liu, L., Sun, X., Gao, T., & Zhao, S. (2013). Comparison of Cloud Properties from Ground-Based Infrared Cloud Measurement and Visual Observations. *Journal of Atmospheric and Oceanic Technology*, 30(6), 1171–1179. <https://doi.org/10.1175/JTECH-D-12-00157.1>.
- Liu, S., & Li, M. (2018). Deep multimodal fusion for ground-based cloud classification in weather station networks. *EURASIP Journal on Wireless Communications and Networking*, 2018 (1). <https://doi.org/10.1186/s13638-018-1062-0>.

- Liu, S., Zhang, Z., Xiao, B., & Cao, X. (2015). Ground-Based Cloud Detection Using Automatic Graph Cut. *IEEE Geoscience and Remote Sensing Letters*, 12(6), 1342–1346. <https://doi.org/10.1109/LGRS.2015.2399857>
- Long, C. N. (2010). Correcting for Circumsolar and Near-Horizon Errors in Sky Cover Retrievals from Sky Images. *Open Atmos. Sci. J.*, 4, 45–52.
- Long, C. N., Sabburg, J. M., & Calbó, J. (2006). Retrieving cloud characteristics from ground-based daytime color all-sky images. *Journal of Atmospheric & Oceanic Technology*, 23(5), 633–652.
- Shelhamer, E., Long, J., and Darrell, T. (2016). Fully Convolutional Networks for Semantic Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39, 4 (April 2017), 640-651. DOI: <https://doi.org/10.1109/TPAMI.2016.2572683>.
- Simonyan, K., & Zisserman, A. (2014). Very Deep Convolutional Networks for Large-Scale Image Recognition. *ArXiv:1409.1556 [Cs]*. Retrieved from <http://arxiv.org/abs/1409.1556>.
- Stull, R. B. (1985). A Fair-Weather Cumulus Cloud Classification Scheme for Mixed-Layer Studies. *Journal of Climate and Applied Meteorology*, 24(1), 49–56.
- Tapakis, R., & Charalambides, A. G. (2013). Equipment and methodologies for cloud detection and classification: A review. *Solar Energy*, 95, 392–430. <https://doi.org/10.1016/j.solener.2012.11.015>
- Taravat, A., Del Frate, F., Cornaro, C., & Vergari, S. (2015). Neural Networks and Support Vector Machine Algorithms for Automatic Cloud Classification of Whole-Sky Ground-Based Images. *IEEE Geoscience and Remote Sensing Letters*, 12(3), 666–670. <https://doi.org/10.1109/LGRS.2014.2356616>
- Wacker, S., Gröbner, J., Zysset, C., Diener, L., Tzoumanikas, P., Kazantzidis, A., ... Kämpfer, N. (2015). Cloud observations in Switzerland using hemispherical sky cameras. *Journal of Geophysical Research: Atmospheres*, 120(2), 695–707. <https://doi.org/10.1002/2014JD022643>
- Wu, W., Liu, Y., Jensen, M. P., Toto, T., Foster, M. J., & Long, C. N. (2014). A comparison of multiscale variations of decade-long cloud fractions from six different platforms over the Southern Great Plains in the United States. *Journal of Geophysical Research: Atmospheres*, 119(6), 3438–3459. <https://doi.org/10.1002/2013JD019813>
- Yang, H., Kurtz, B., Nguyen, D., Urquhart, B., Chow, C. W., Ghonima, M., & Kleissl, J. (2014). Solar irradiance forecasting using a ground-based sky imager developed at UC San Diego. *Solar Energy*, 103, 502–524. <https://doi.org/10.1016/j.solener.2014.02.044>.

- Zhao, G., & Di Girolamo, L. (2007). Statistics on the macrophysical properties of trade wind cumuli over the tropical western Atlantic. *Journal of Geophysical Research: Atmospheres*, 112(D10), D10204. <https://doi.org/10.1029/2006JD007371>
- Zhu, H., Meng, F., Cai, J., & Lu, S. (2016). Beyond pixels: A comprehensive survey from bottom-up to semantic image segmentation and cosegmentation. *Journal of Visual Communication and Image Representation*, 34, 12–27. <https://doi.org/10.1016/j.jvcir.2015.10.012>