Nowcasting Solar Radiance Over Oahu

Kyle Hart
Information and Computer Sciences
University of Hawai'i at Mānoa

Giuseppe Torri
Atmospheric Sciences
University of Hawai'i at Mānoa

Peter Sadowski
Information and Computer Sciences
University of Hawai'i at Mānoa
peter.sadowski@hawaii.edu

Abstract

We use satellite data from GOES-17 and deep learning to predict solar radiance with a 10-60 minute forecast horizon. Neural networks were trained on data covering the the Hawaiian islands from 2019, and tested on 2020 data. Our 10-minute forecasts of solar radiance achieve an RMSE of 31 $\mathrm{Wm^{-2}sr^{-1}\mu m^{-1}}$, a significant improvement over a simple persistence model benchmark (45 $\mathrm{Wm^{-2}sr^{-1}\mu m^{-1}}$ on the same data). These results suggest that the approach could potentially be used by energy companies to more efficiently manage power-generators.

1 Introduction

The city of Honolulu generates more solar power per capita than any other major US city by far [1], but this creates unique challenges. As the Hawaiian islands are independent from the mainland power-grid, they are especially vulnerable to the volatility of solar energy production. Sudden, drastic changes in atmospheric conditions can lead to spikes or shortfalls in net solar production on the grid, especially in Honolulu and the island of Oahu (population 1 million). The ability to forecast changes in *solar irradiance* — the amount of sunlight reaching the Earth's surface — due to clouds would assist grid operators in managing power generators more efficiently. These operators must decide when to start additional generators and *which* generators to start; more efficient generators require 30 minutes to start up, while less efficient generators require only 10 minutes.

Intriguingly, the Hawaiian islands have unique properties that could make it possible to forecast atmospheric conditions better than in other locations. As the islands are surrounded by ocean with predictable trade wind patterns, approaching clouds can be tracked in satellite images. We hypothesize that methods from computer vision can automate this process, and investigate the use of deep learning to forecast cloud coverage on the short-term 10-60 minute *nowcasting* horizon. While similar approaches have been used for predicting radar-observed precipitation patterns [2, 3], this work models direct observations from the latest generation of space-based Geostationary Operational Environmental Satellites (GOES).

2 Methods

2.1 Data Source

The GOES-17 satellite began operations in February 2019, and was designed to provide high-resolution atmospheric measurements over the Pacific Ocean in near real-time with the explicit goal of improving forecasting capabilities. The Advanced Baseline Imager (ABI) conducts full-disk

AI for Earth Sciences Workshop at NeurIPS 2020.

observations once every 5 to 15 minutes, measuring 16 spectral bands from visible to long-wave infrared. These are measurements of solar *radiance*, the amount of light observed by the satellite and measured in Watts per area per steradian per wavelength: $\mathrm{Wm^{-2}sr^{-1}\mu m^{-1}}$. Here we focus on the problem of forecasting solar radiance, and leave for future work the problem of using satellite-observed radiance to predict ground-based irradiance.

This work focuses on data from Band 2 ($0.60 - 0.68\mu m$) because it has the highest spatial resolution–0.5 km–and it matches the peak of absorption range of crystalline silicon photovoltaic (PV) cells. The National Oceanographic and Atmospheric Administration (NOAA) has made GOES-17 imaging data available on the Google Cloud Platform, and we constructed our dataset by pulling all full-disk images for Band 2 of the API taken in 2019 and the first two months of 2020.

2.2 Training Dataset

Data from 2019 was used for model training and development, while data from 2020 was reserved as a clean test set for final evaluation. A random selection of days in 2019 were used as a validation set for hyper-parameter tuning and model comparison during development.

A single data sample consisted of a time *t* for which we had a fixed number of input and output timesteps, e.g. 5 input timesteps (corresponding to observations at *t*-0, *t*-10, *t*-20, *t*-30, and *t*-40 minutes) and an output timestep corresponding to predictions at 10 or 30 minutes into the future. Thus, predictions were not made for the very beginning or very end of each day. For each timestep, solar radiance is represented by a 192-by-192 pixel image, with each pixel covering a $0.5 \times 0.5\ \mathrm{km}^2$ area. During training, we included samples from four different islands: Oahu, Maui, Kauai, and Hawaii.

Solar radiance values were normalized by scaling by a constant factor of 800, so that all the values were non-negative and less than 1. The target values were normalized in the same way as the input, and were predicted using a *softplus* layer with mean squared error (MSE).

2.3 Model

The forecasting models explored were convolutional neural networks (CNN) with 3D convolutions over the two spatial dimensions and one temporal dimension. Models were trained using TENSORFLOW[4] and KERAS[5]. Over a hundred different models were trained to optimize hyper-parameters using the SHERPA hyper-parameter optimization framework [6]. The best model was selected based on the MSE on the 2019 validation set. This model was then fine-tuned on 2019 training data from Oahu, using the validation set again for early stopping. We present the performance results on the clean test set from 2020.

3 Results

3.1 10-Minute Forecast Horizon

First we predicted solar radiance on 10-minute forecast horizon, since this is the cadence of the data and shorter horizons are easier to predict. The model achieves a root mean squared error (RMSE) of 25 $\mathrm{Wm^{-2}sr^{-1}\mu m^{-1}}$ on a validation set of 2019 data, and 31 $\mathrm{Wm^{-2}sr^{-1}\mu m^{-1}}$ on 2020 test data. This was compared against a benchmark *persistence* model: a model in which the solar radiance does not change from the time of prediction. Our model performs significantly better than the persistence model, which only achieves 45 RMSE on the same test set. Figure 1 shows typical forecasts over a one hour interval selected from the test set; the model is able to predict the movement of clouds with different speeds, but unable to predict many of the finer details.

3.2 30- to 60-Minute Forecast Horizon

One approach to forecasting further into the future is by recursively feeding the model's predictions back into itself. This works for predicting the trajectory of moving clouds, but it also tends to predict fainter and fainter clouds with each timestep as the model predictions regress towards the mean (see Figure 2). Thus, we expect a model trained to directly predict solar radiance at the desired timestep to be more accurate.

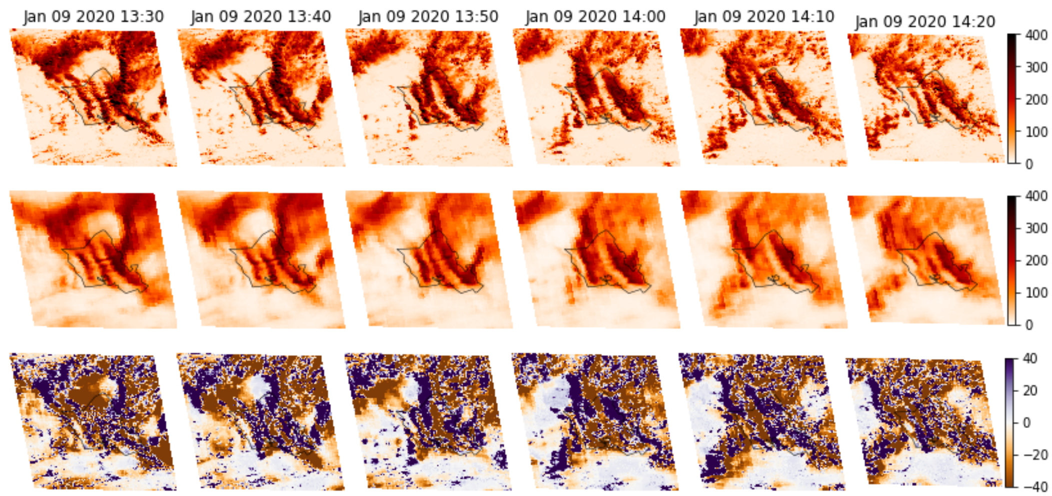

Figure 1: Observed solar radiance (top), predictions with 10-minute forecast horizon (middle), and prediction error, i.e. observed minus predicted (bottom), over Oahu at five sequential timesteps at 10-minute intervals. This typical example shows tradewinds pushing clouds West-ward at different rates. The predictions anticipate these cloud movements. The blurriness of the predictions reflects the inability of the model to predict fine-grained details. The patterns in the error plots show that the clouds in the top left and bottom left move faster than model expects.

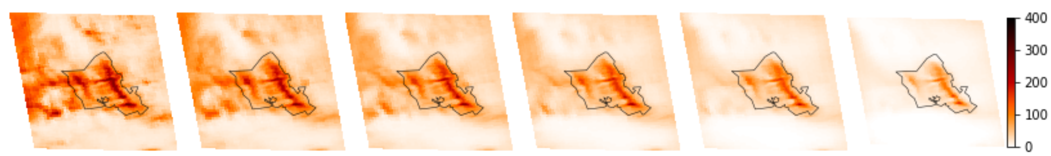

Figure 2: Example predictions at 10, 20, 30, 40, 50, and 60 minutes using the recursive approach to longer forecasts. The model predicts that the clouds continue moving to the left, but they get fainter with each step. A model that makes 30 minute predictions directly does not have this problem.

CNNs with the same model architecture were trained for 20, 30, 40, 50, and 60-minute forecast horizons. Figure 3 shows typical forecasts over a one-hour interval for the model with a 30-minute forecast horizon. This model predicts far less detail than the 10-minute horizon model, but the model is still able to predict cloud trajectories. It achieves an RMSE of 41 $\mathrm{Wm^{-2}sr^{-1}\mu m^{-1}}$ on the 2020 test set, compared to 56 $\mathrm{Wm^{-2}sr^{-1}\mu m^{-1}}$ for the persistence model. A video of one full day of predictions from the 30-minute forecast model can be seen at the following link: `https://drive.google.com/file/d/19NjHmSzQaMW2y6i_dhsd5peoTvOBAV5y/view?usp=sharing`

Finally, a single model was trained to predict the radiance at multiple forecast horizons from 10-60 minutes. This combined *multi-task* model performed approximately the same as, or slightly worse than, the individual models (Figure 4), but the deep learning approaches always perform better than the persistence model. As expected, RMSE increases as the forecast horizon increases for all models.

4 Discussion and Conclusion

Our results demonstrate that deep learning can provide significantly better forecasts than the simple persistence model benchmark. Qualitatively, our model does more than just "blur" the persistence model (i.e. a "damped persistence" model) — rather, it predicts the trajectories of clouds by detecting movements over the five input timesteps and extrapolating cloud locations in future timesteps. Despite the underlying system being chaotic, the model is able to learn these patterns from data.

These initial results suggest that the model could already be accurate enough to inform energy-production decisions. However, more work is needed to assess the value of these predictions. We

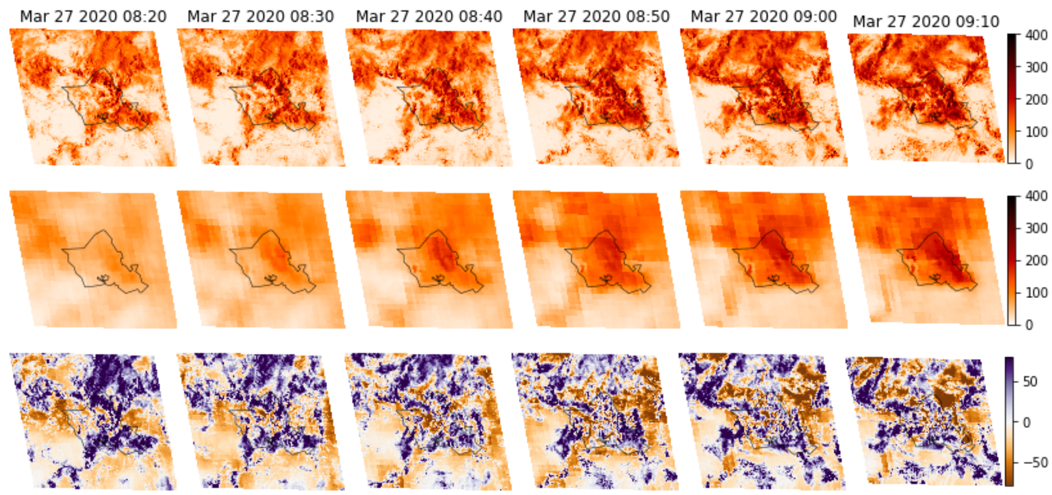

Figure 3: Observed solar radiance (top), predictions with 30-minute forecast horizon (middle), and prediction error, i.e. observed-predicted, (bottom) over Oahu at five sequential timesteps at 10-minute intervals. The predictions are significantly less sharp than those in Figure 1 because the model has a harder time forecasting this far, but we still observe the model predicting significant variation in solar radiance.

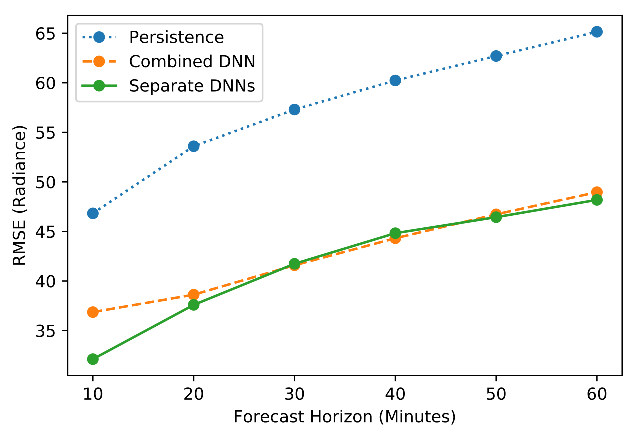

Figure 4: Prediction error (RMSE) vs. forecast horizon (minutes) on 2020 test data for the persistence model (blue), a combined deep neural network (DNN) model that forecasts for all six horizons simultaneously (orange), and separate DNNs trained to predict each horizon individually (green).

present average prediction error over the entire island of Oahu, but for energy-production we will need to focus on the solar assets concentrated in highly-populated areas. Furthermore, we expect forecasts to be more accurate in some areas than others due to geological features. Ultimately, future work will need to incorporate these predictions into a policy and perform back-testing comparisons to existing policies.

Acknowledgments

The authors would like to acknowledge support from Hawaiian Electric Company (HECO) and the UHM Office of Innovation and Commercialization.

References

[1] Adrian Pforzheimer, Elizabeth Ridlington, Ben Sonnega, and Emma Searson. Shining cities 2020: The top U.S. cities for solar energy, 2020. https://environmentamerica.org/feature/ame/shining-cities-2020.

[2] Xingjian SHI, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun WOO. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 802–810. Curran Associates, Inc., 2015.

[3] Shreya Agrawal, Luke Barrington, Carla Bromberg, John Burge, Cenk Gazen, and Jason Hickey. Machine learning for precipitation nowcasting from radar images, 2019.

[4] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.

[5] François Chollet. *Keras*. GitHub, 2015.

[6] Lars Hertel, Julian Collado, Peter Sadowski, Jordan Ott, and Pierre Baldi. Sherpa: Robust hyperparameter optimization for machine learning. *SoftwareX*, 12:100591, 2020. Software available at: https://github.com/sherpa-ai/sherpa.