

# SolaCam: A Deep Learning Model for Solar Radiation Estimation Using Consumer Cameras

Daisuke Sugiyama<sup>1,2</sup>, Ryo Onishi<sup>3</sup>, and Hironori Fudeyasu<sup>4</sup>

<sup>1</sup>Japan Agency for Marine-Earth Science and Technology, Yokohama, Japan

<sup>2</sup>Kagoshima University, Kagoshima, Japan

<sup>3</sup>Tokyo Institute of Technology, Tokyo, Japan

<sup>4</sup>Yokohama National University, Yokohama, Japan

(Manuscript received 17 May 2023, accepted 31 August 2023)

**Abstract** This study proposes a deep learning approach called SolaCam to accurately estimate solar radiation from the images captured by cameras. The proposed SolaCam performs deep learning by utilizing both image features and theoretical maximum solar radiation that vary with time and location. The trained model is capable of accurately estimating solar radiation on the ground surface from sky images captured by smartphones, fixed-point cameras, and other devices. The developed SolaCam can use a remote sensing function, which estimates solar radiation, on inexpensive camera-equipped devices.

**Citation:** Sugiyama, D., R. Onishi, and H. Fudeyasu, 2023: SolaCam: A deep learning model for solar radiation estimation using consumer cameras. *SOLA*, **19**, 246–252, doi:10.2151/sola.2023-032.

## 1. Introduction

To achieve a smart society, the use of high-density Internet of Things (IoT) sensors has become important in collecting information on cities at higher densities, which can be utilized for monitoring and predicting phenomena. Remote sensors, such as radiometers and solar radiation meters, are expensive. In contrast, digital cameras connected to the internet are inexpensive, and images of the sky and clouds taken by IoT sensors can be analyzed to extract useful information.

Recent advances in visual pattern recognition using deep learning (DL) methods have resolved many tasks that were previously challenging using machine learning (ML) methods (LeCun et al. 1989; Krizhevsky et al. 2012; Howard et al. 2019; Tan et al. 2019; Dosovitskiy et al. 2021). Image analysis using DL has also been applied in earth sciences; for example, estimating cloud optical thickness from images captured by sky cameras (Masuda et al. 2019).

This study targets camera images from inexpensive consumer webcams and mobile devices to build a DL model. Owing to the use of inexpensive camera devices rather than specialized IoT weather instruments, comprehensive observations would be obtained, which would contribute to urban or micro weather forecast simulations at an urban district level. Similar research on extracting weather observation data from consumer camera images includes the estimation of 0–10 cloud coverage (Onishi and Sugiyama 2017), and atmospheric temperature estimation (Chu 2018).

This study proposes a new DL-based method for accurately estimating solar radiation using information from camera equipment. The method involves a regression task using consumer camera images and shooting time information input to perform a real-value estimation of solar radiation directed at the ground surface (0–1.32 kW/m<sup>2</sup>). The image could be captured in any direction, roughly horizontal to the ground, where over 75% of the image should be of the sky. To achieve high estimation accuracy, the learning process included meteorological domain knowledge, such as the maximum amount of solar radiation outside the atmosphere, which can be calculated from the latitude and longitude of the camera location and the time of day.

## 2. Data and methods

### 2.1 Data augmentation

Sky images in five directions were captured over a six-year period using several fixed-point cameras installed on the roof of Yokohama National University (YNU), and data from the pyranometer on the same roof were obtained. Figure 1 shows an example of the observation environment and data collected. The images and solar radiation data were recorded every minute, taking a total of approximately five million images. However, images from fixed cameras only may be insufficient for learning diverse image features because of limited variations in shooting directions, lighting sources, and background locations. Therefore, a technique called data augmentation (Shorten et al. 2019), which increases the diversity of training data, was effectively used.



Fig. 1. Left: Camera and pyranometer installed on the roof of Yokohama National University. Images and solar radiation are measured every minute. Five cameras cover directions of Mt. Fuji (west), Machida (northwest), Tokyo (northeast), Minatomirai (southeast), and Kamakura (south). Right: Camera images (southeast direction) at 12:00 AM in July 2017, arranged for 30 days. The number in the upper left of each image shows the solar radiation in  $\text{kW/m}^2$ .

To verify the need for data augmentation in fixed-point cameras, a simple experiment was conducted using a convolutional neural network (CNN), which is effective for learning image features.

In this base experiment, only one camera oriented towards Mt. Fuji (west) was used, and the CNN estimation accuracy was measured by cross-validation in the time direction. Specifically, the first five years of the six-year period were used for training, and the resulting trained model was evaluated for the remaining year. The results showed that the root mean square error (RMSE) for the untrained year was  $0.064 \text{ kW/m}^2$ . However, when data from a camera in a different direction that was not used for learning was used for estimation with CNN, the accuracy was extremely low, with an RMSE of  $0.4 \text{ kW/m}^2$  or higher. Consequently, it was assumed that the system overfitted the features that depended on the fixed-point camera images of the learning target, such as brightness, contrast, ground position, and ground shadow. Based on this assumption, the generalization of fixed-camera learning to general cameras is necessary in addressing this problem.

Therefore, data augmentation (Advanced Setting S1) was performed on the images to be learned. The confusion matrix (Fig. S2) obtained from the CNN trained by the augmented dataset showed that the RMSE improved with data augmentation, becoming as low as  $0.12 \text{ kW/m}^2$  for target solar radiations less than  $0.7 \text{ kW/m}^2$  (7-th 8-quantile of dataset), and as low as  $0.25 \text{ kW/m}^2$  for targets higher than  $0.7 \text{ kW/m}^2$ .

The lower accuracy for larger radiation values may be attributed to seasonal variations in solar radiation. Maximum monthly solar radiation observed at the camera location (contained reflections by clouds and ground surfaces) was  $0.7 \text{ kW/m}^2$  in December and  $1.3 \text{ kW/m}^2$  in July. Alternatively, the number of images with solar radiation higher than  $0.7 \text{ kW/m}^2$  was small due to seasonal variations.

## 2.2 SolaCam model

The developed SolaCam model learned seasonal variations in solar radiation by calculating the theoretical maximum solar radiation from the latitude, longitude, and time, and conducting learning with images (Fig. 2). A neural network architecture was developed using PyTorch (Adam et al. 2019).

The architecture consisted of two parts: one for inputting images and the other for values. Sky images were processed by clipping and resizing to  $128 \times 128$  pixels for input into the CNN. The CNN used in this study was similar to those used in previous studies (LeCun et al. 1999; Onishi and Sugiyama 2017). The detailed architecture and formulae of the CNN are presented in Advanced Setting S3.

Input values crucial for calculating the solar radiation were latitude, longitude, and time. When using SolaCam estimation, location and time can be obtained from camera devices. A general geometric formula for meteorology was used

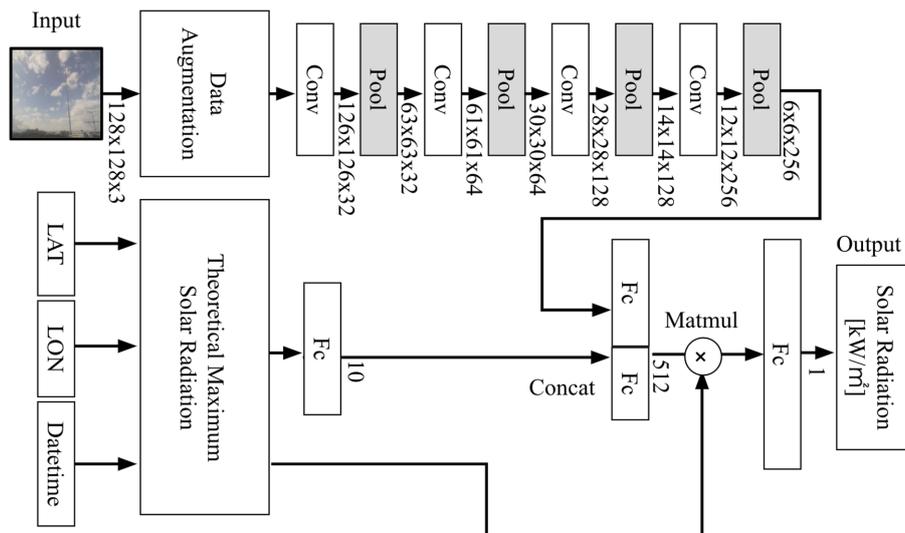


Fig. 2. Arrows refer to the data flow. The number next to the arrow indicates the size of each dimension of the output: Conv: convolutional layer; Pool: max pooling layer; and Fc: fully connected layer. The output of the convolutional neural network (CNN) is 502 nodes, and the output of Fc, the theoretical maximum solar radiation, is 10 nodes; the two outputs are concatenated. The Red Green Blue (RGB) values of the input image and the theoretical maximum solar radiation are normalized to 0.0–1.0.

to calculate the theoretical maximum solar radiation using the input values. The calculated value was normalized to a range of 0.0–1.0 and inputted into a fully connected (FC) layer. This FC layer was concatenated with the CNN FC layer and learned simultaneously. The output parameters from the output layer do not directly use solar radiation but rather use solar radiation divided by the maximum solar radiation. In other words, the product of the model's output value and maximum solar radiation is the solar radiation. This implies that the model learned the attenuation rate of the maximum solar radiation based on the sky conditions in the image. Generally, the observed solar radiation data were imbalanced, with exponentially fewer samples as the solar radiation approached high levels (the histogram of solar radiation used in this experiment is shown in Fig. S4). When optimizing a neural network using this type of imbalanced data, low solar radiation data are more heavily learned, resulting in substantially reduced accuracy at high solar radiation levels. Therefore, SolaCam applied a sampling method called weighted random sampling (Efraimidis et al. 2008). First, the training dataset were divided into groups of solar radiation at 0.1 kW/m<sup>2</sup> units and then weighted random sampling was performed, so that the data selected for the mini batch had uniform probability between groups. This enabled uniform probability learning to be achieved for each solar radiation group from 0.0 to 1.3 kW/m<sup>2</sup> in the data selected for mini batch stochastic gradient descent.

### 2.3 Dataset for training

To evaluate generalized performance of SolaCam, the YNU trained dataset was used, which consisted of sky images captured at one-minute intervals by four of the five cameras installed on the university roof (35.473°N, 139.591°E) from 2011 to 2016. The remaining one year of data from the unused camera facing south was used as the YNU test dataset. This allowed the generalized performance of SolaCam to be measured using unseen cameras, unexplored directions, unobserved landscapes, and future data (see Advanced Setting S5 for details on hyperparameter tuning).

### 2.4 Implementation of training

SolaCam was trained using approximately 3.5 million images from the YNU training dataset with Graphics Processing Units (GPU). Training was performed on supercomputers in Japan Agency for Marine-Earth Science and Technology (JAMSTEC) using the data-parallel method (Krizhevsky 2014) on eight GPUs (refer to Advanced Setting S6 for details on distributed learning and training hours). Although the training required significant amount of time, the inference time for SolaCam was less than one second, even on a Central Processing Unit (CPU). In addition, estimation can occur in less than 10 seconds even on a weak CPU installed onto an inexpensive device such as a Raspberry Pi.

### 2.5 Baseline analysis

The solar radiation data was analyzed to determine the optimal coefficient ( $k$ ), for the theoretical maximum solar radiation ( $M$ ) multiplied by  $k$  using a simple linear regression. The optimal coefficient ( $k$ ) was found to be 0.46 for all the YNU training dataset. This  $M \times k$  model was used as a baseline for comparison. Furthermore, the functionalities of multiple models were compared to reveal the effects of learning physical information.

Table 1. Comparisons of mean absolute errors (MAE), root mean square errors (RMSE), coefficient of determination called  $R^2$  score and balanced root mean square errors (Balanced-RMSE) of estimated solar radiation parameters per 1-minute for test dataset at Yokohama National University.

Method	MAE [kW/m <sup>2</sup> ]	RMSE [kW/m <sup>2</sup> ]	$R^2$ score	Balanced-RMSE [kW/m <sup>2</sup> ]
a) Baseline	0.095	0.16	0.6	0.28
b) <b>SolaCam (Proposed Model)</b>	<b>0.045</b>	<b>0.11</b>	<b>0.83</b>	<b>0.13</b>
c) SolaCam without physics knowledge	0.055	0.11	0.81	0.3
d) SolaCam without weighted random sampler	0.047	0.11	0.83	0.32
e) EfficientNet-B3	0.079	0.16	0.64	0.4

### 3. Results and discussion

#### 3.1 Estimation results using SolaCam

Table 1 lists the accuracy of SolaCam, and with model comparison results are shown in Table 1 (b)–(d). The results in Table 1 were analyzed to compare the accuracy and performance of the models. Generally, lower values of mean absolute error (MAE), RMSE and  $R^2$  score indicate lower overall errors and higher accuracy. Notably, the solar radiation data in this study were imbalanced with a few high solar radiation values; therefore, even if the RMSE values were good, poor accuracy in the high solar radiation range could possibly be observed with a small sample size. To evaluate performance, the average accuracy for high and low solar radiations was used as a balanced-RMSE metric, which can be defined as follows:

$$RMSEa = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}, \quad Y < 0.7$$

$$RMSEb = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}, \quad Y \geq 0.7$$

$$Balanced \ RMSE = \frac{RMSEa + RMSEb}{2}$$

The threshold of  $Y = 0.7 \text{ kW/m}^2$  was chosen based on the discussion in Subsection 2.1.

Based on Table 1, SolaCam is more accurate than the baseline model in all the four metrics, i.e., MAE, RMSE,  $R^2$  score and Balanced RMSE. Specifically, in the Balanced RMSE, SolaCam is more accurate than the baseline model by 54%. It should be noted that, from the MAE, RMSE and  $R^2$  score values, the balanced-RMSE increased for all the five methods, indicating that the model does not perform well in estimating high solar radiation. A comparison of (b) and (c) shows that incorporating physics knowledge into SolaCam can improve the accuracy of the MAE, RMSE and  $R^2$  score. The value of the balanced RMSE shows that the estimation accuracy for high solar radiation is also less likely to deteriorate. Comparison of (b) and (d) shows that weighted random sampling considerably improves the accuracy of estimating high solar radiation. Additionally, a comparison of (b) and (e) shows that even using a model with many layers, such as EfficientNet, there is an overfitting tendency, resulting in lower accuracy. A confusion matrix for analyzing the estimation accuracy in each solar radiation range is presented in more detail in Fig. S7.

Figure 3 shows examples of images and SolaCam estimations for the YNU test dataset. Figure S8 shows visual explanations using Score-CAM of the estimation in Fig. 3 (Wang et al. 2020). When examining the estimation results for each image in Fig. 3 and Fig. S8, solar radiation can be estimated reliably even if changes exist in image features due to weather conditions, such as when raindrops are visible (b), or when the sun is captured (c). Moreover, as shown in (c), estimation is possible even when the sun is visible, and expensive equipment to conceal the sun is unnecessary. Additionally, images with higher brightness values are commonly thought to correspond to higher solar radiation values; however, images with rain or thin clouds, such as in (b) and (e), have higher brightness values than those in clear weather. Therefore, image features such as clouds, thickness and type, and weather-related changes are considered more critical than brightness for estimating solar radiation. Figure 3f shows an example where the error becomes large in a pattern with thick clouds, but the sun is visible. This error is probably caused because solar radiation changes over a short period of time when the sun is locally hidden by thick clouds.

These results indicate that incorporating physical knowledge into DL substantially improves accuracy. Additionally, the inclusion of weighted random sampling, which is effective for learning imbalanced data, resulted in the highest accuracy of SolaCam in estimating high levels of solar radiation.

#### 3.2 Estimation accuracy at different locations from YNU training dataset

To further confirm the generalized accuracy, the accuracy of solar radiation estimation was measured for camera images obtained from locations different from the training dataset. A fixed-point camera (facing southeast) and a pyra-

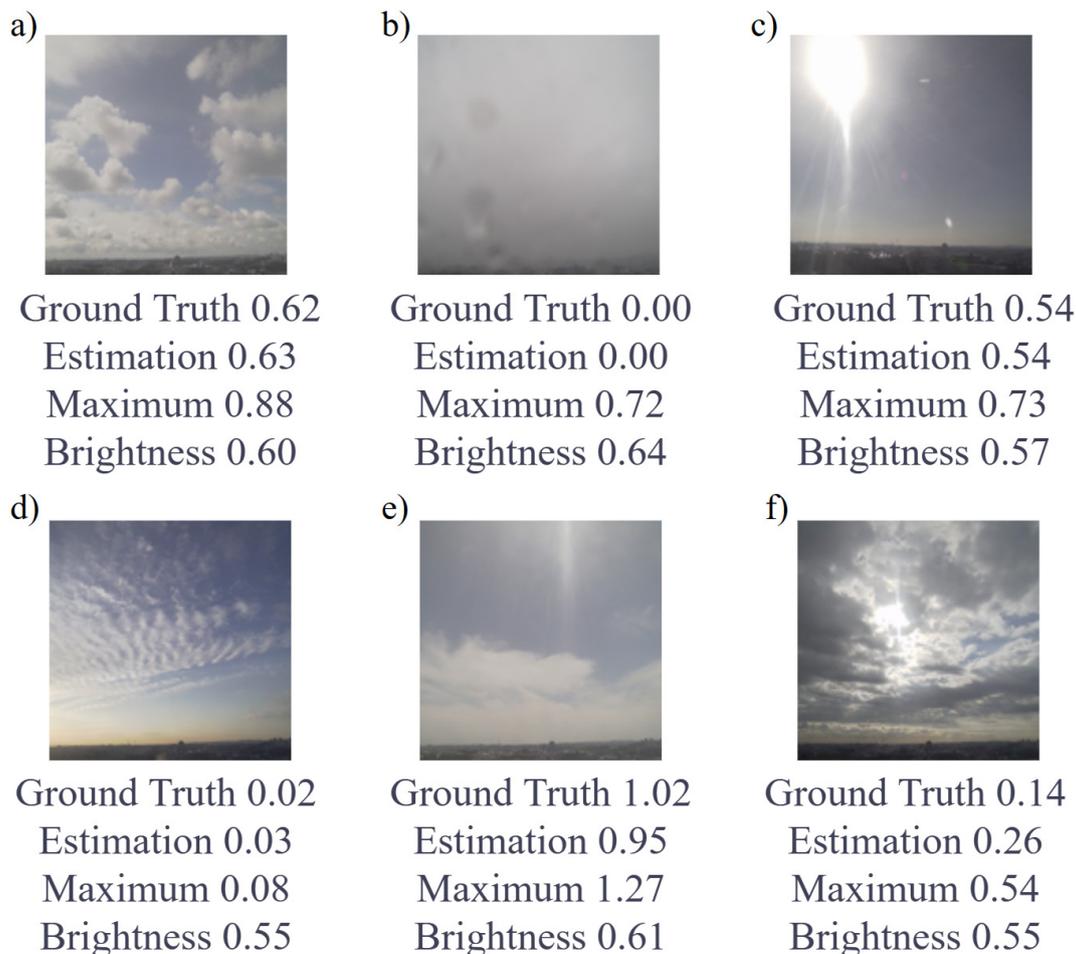


Fig. 3. Examples of SolaCam estimation model for the YNU test dataset: a) clear sky, b) rain, c) sun light, d) early morning cirrocumulus, e) high solar radiation near noon, and f) a pattern of large error. Maximum value is the theoretical maximum solar radiation. Brightness is the pixel-wise average of the brightness values (0.0 to 1.0) of the image.

Table 2. Comparisons for the dataset in Japan Agency for Marine-Earth Science and Technology (JAMSTEC) at Yokohama Institute for Earth Sciences, and Yokohama National University (YNU) from November to December. For comparison, a) and b) are the same as in Table 1.

Method	Target	MAE [ $\text{kW}/\text{m}^2$ ]	RMSE [ $\text{kW}/\text{m}^2$ ]	Balanced-RMSE [ $\text{kW}/\text{m}^2$ ]
a) Baseline	YNU	0.059	0.11	0.18
b) SolaCam	YNU	0.026	0.06	0.08
f) Baseline	JAMSTEC	0.06	0.11	0.18
g) SolaCam	JAMSTEC	0.041	0.08	0.11

nometer were installed on the JAMSTEC roof ( $35.380^\circ\text{N}$ ,  $139.625^\circ\text{E}$ ). The distance between JAMSTEC and YNU is approximately 12 km. The input was a camera image captured roughly horizontal to the ground, such that over 75% of the image was of the sky.

The JAMSTEC solar radiation data were obtained for November to December 2022. SolaCam, trained with the YNU training dataset, was used to estimate solar radiation. Table 2 lists the estimation errors for the YNU test dataset and the JAMSTEC dataset for the same month. Figure S9 illustrates examples of these images.

The results showed similar estimation errors for SolaCam for both YNU and JAMSTEC data, and the errors were smaller than the baseline errors. This confirms the generality of SolaCam, i.e., it can estimate highly accurate solar radiation from camera images even at locations where the training dataset was not obtained.

### 3.3 Analysis of seasonal trends in results and detection of observation errors

To analyze the trend of results by season and identify observation errors, the daily cumulative transition of estimated solar radiation was shown by SolaCam, and the seasonal errors in the estimation are shown using a boxplot (Fig. 4).

Consequently, it was found that SolaCam had the largest errors in July, particularly during the summer months when

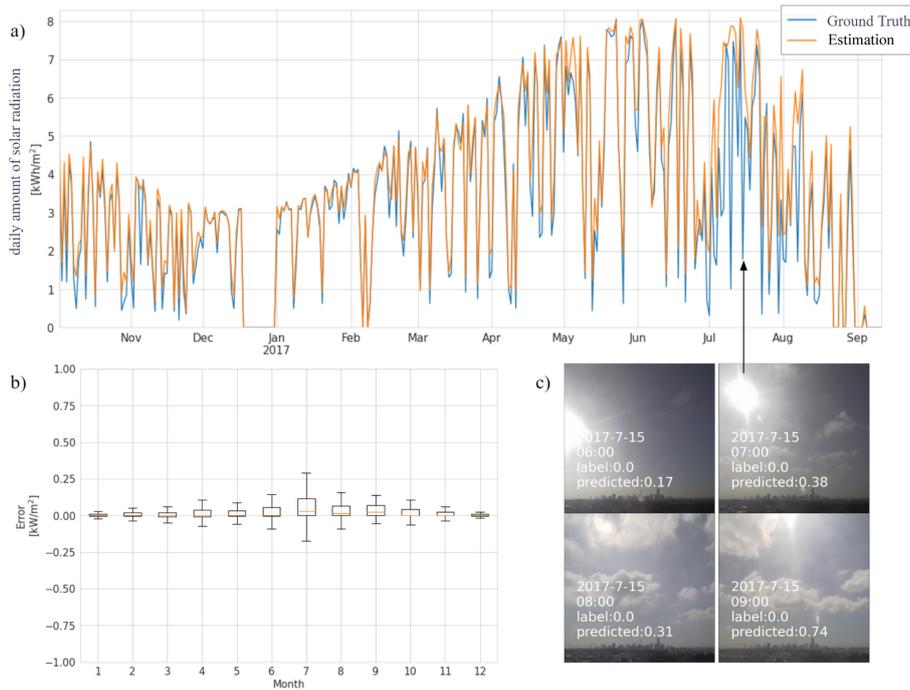


Fig. 4. a) Line plots of daily amount of solar radiation estimated by the SolaCam for YNU test dataset. The accuracy of SolaCam is mean absolute error:  $0.402 \text{ kW/m}^2$  and root mean square error:  $0.796 \text{ kW/m}^2$ . Note that the latter half of December is missing. b) Boxplot for errors of each month. Both ends of the box are quartiles; the outer lines are 1.5 times quartiles, and the orange line is median. c) Examples of camera images suggesting observation error on July 15, 2017, as measured solar radiations were  $0 \text{ kW/m}^2$ .

solar radiation is the highest. To investigate the cause of these errors, the estimated results for July were examined in detail, and the estimated cumulative solar radiation on several days were found to be substantially different from the actual value. Despite the presence of the sun in the images for these days, the observed data were  $0 \text{ kW/m}^2$ , indicating observational errors in the solar radiation sensors of the YNU test dataset or local changes in the observation conditions that were not captured in the images.

Therefore, it can be concluded that errors in estimating solar radiation during the summer months when there was a high amount of solar radiation were large, and some of these errors were due to missing observational data on certain days. This result highlights the potential for detecting observation errors that were previously unknown before using the SolaCam. For example, by comparing the estimated solar radiation from camera images with the actual data measured by a solar radiation sensor, observation errors could be realized that would otherwise be unnoticed.

### 3.4 Influence of change in time window

The amount of solar radiation depends on whether the sun is covered by clouds. Even on a clear day, a small, isolated cloud covering the sun can significantly decrease solar radiation, meaning that the observed solar radiation could change markedly locally within a few minutes. However, the estimated solar radiation from a camera image that captures a wide sky area from a distance is considered the average solar radiation over space. Since space and time averaging are closely related, SolaCam captures time-averaged solar radiation. Table 3 shows the estimation errors for solar radiation averaged over different time windows, including 1 min, 10 min, 30 min, and 1 h moving averages. The results demonstrate that the estimation error for the 1 min average was large, while it was notably lower (by over 30%) for the 30 min and 1 h averages. It should be noted that SolaCam estimate is more sensitive to the time windows than the baseline. If the space and time scales are related with a representative velocity of  $10 \text{ m/s}$ , a 30 min or 1 h time window corresponds to a spatial window of several tens of kilometers. This indicates that the present SolaCam captures solar radiation averaged over an area of several tens of kilometers.

## 4. Conclusion

A DL-based SolaCam that accurately estimates solar radiation based on camera images, latitude, longitude, and time is proposed. The SolaCam is more accurate in estimating 1 min average solar radiations than the baseline model by

Table 3. Estimation errors of SolaCam for the YNU test dataset with moving averaging for different time windows.

Method	Balanced-RMSE [kW/m <sup>2</sup> ]	MAE Improvement Rate From 1-min.
a) Baseline 1-min.	0.25	—
a-1) Baseline 10-min.	0.24	4.00%
a-2) Baseline 30-min.	0.23	8.00%
a-3) Baseline 1-hour	0.23	8.00%
b) SolaCam 1-min.	0.13	—
b-1) SolaCam 10-min.	0.094	20.00%
b-2) SolaCam 30-min.	0.074	31.10%
b-3) SolaCam 1-hour	0.064	40.00%

54% in terms of the balanced RMSE, which evaluates the accuracy for high and low solar radiations. The learned physical information enabled the system to estimate solar radiation with high accuracy even for locations that were different from those where the training data were acquired.

SolaCam has high accuracy for the estimation of 1 h and 30 min average radiations rather than 1 min average ones. This means that SolaCam captures solar radiation averaged over an area of several tens of kilometers.

Camera images used for solar radiation estimation should be preprocessed (e.g., trimmed and aligned) to match the specific format. Relaxation of this restriction can be a topic for future investigations. Other future investigations that can improve SolaCam are described in Future works S10.

## Acknowledgements

This research was partially supported by Grant-in-Aid for Transformative Research Areas, Grant Number JP20B207 and KAKENHI Grant 21K03658. SolaCam were trained using supercomputers in JAMSTEC.

Edited by: K. Ito

## Supplements

In the auxiliary material, settings for training SolaCam; data augmentation, neural network architecture, validation dataset and multi-GPUs training, are shown in Advanced Settings S1, S3, S5 and S6. Confusion matrix for results of Subsection 2.1, Tables 1 and 3 are shown in Fig. S2, S7. Histogram of solar radiation in the YNU dataset is shown in Fig. S4. Visualizations of basis of estimation are shown in Fig. S8. Examples of sky images at JAMSTEC are shown in Fig. S9. Future works of SolaCam are described in S10.

## References

- Chu, W. T., K. C. Ho, and A. Borji, 2018: Visual weather temperature prediction. *2018 IEEE Winter Conf. on Applications of Computer Vision (WACV)*, Lake Tahoe, NV, USA, 234–241, doi:10.1109/WACV.2018.00032.
- Dosovitskiy, A., L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, T. Unterthiner, and X. Zhai, 2021: An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale, *ICLR2021 Oral*, <https://doi.org/10.48550/arXiv.2010.11929>.
- Efrimidis, P., and P. Spirakis, 2008. Weighted random sampling. *Encyclopedia of Algorithms*, M.-Y. Kao, Ed., Springer US, Boston, MA, 1024–1027, [https://doi.org/10.1007/978-0-387-30162-4\\_478](https://doi.org/10.1007/978-0-387-30162-4_478).
- Howard, A., M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, Q. V. Le, and H. Adam, 2019: Searching for MobileNetV3. *Proc. the IEEE/CVF International Conference on Computer Vision*, 1314–1324.
- Kingma, D. P., and Ba J. Adam, 2015: A method for stochastic optimization. *Proc. the 3rd international conference on learning representations (ICLR2015)*, Hilton San Diego Resort & Spa, San Diego.
- LeCun, Y., P. Haffner, L. Bottou, and Y. Bengio, 1999: Object recognition with gradient-based learning. *Shape, Contour and Grouping in Computer Vision. Lecture Notes in Computer Science*. Springer, Berlin, Heidelberg, 1681.
- Masuda, R., H. Iwabuchi, K. S. Schmidt, A. Damiani, and R. Kudo, 2019: Retrieval of cloud optical thickness from sky-view camera images using a deep convolutional neural network based on three-dimensional radiative transfer. *Remote Sens.*, **11**, 1962, doi:10.3390/rs11171962.
- Onishi, R., and D. Sugiyama, 2017: Deep convolutional neural network for cloud coverage estimation from snapshot camera images. *SOLA*, **13**, 235–239, doi:10.2151/sola.2017-043.
- Paszke, A., S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, 2019: PyTorch: An Imperative Style, High-Performance Deep Learning Library. *Neural Information Processing Systems*.
- Shorten, C., and T. M. Khoshgoftaar, 2019: *A survey on Image Data Augmentation for Deep Learning*. *J Big Data* **6**, 60 2019, <https://doi.org/10.1186/s40537-019-0197-0>.
- Sujan, G., C. Deo Ravinesh, D. Casillas-Pérez, S. Salcedo-Sanz, E. Sharma, and M. Ali, 2022: Deep learning CNN-LSTM-MLP hybrid fusion model for feature optimizations and daily solar radiation prediction. *Measurement*, **202**, 111759, ISSN 0263-2241, doi:10.1016/j.measurement.2022.111759.
- Tan, M., and Q. Le, 2019: EfficientNet: Rethinking model scaling for convolutional neural networks. *Proc. the 36th International Conf. on Machine Learning, PMLR*, **97**, 6105–6114.
- Wang, H., Z. Wang, M. Du, F. Yang, Z. Zhang, S. Ding, P. Mardziel, and X. Hu, 2020: Score-CAM: Score-weighted visual explanations for convolutional neural networks. *Proc. the IEEE/CVF Conf. on Computer Vision and Pattern Recognition Workshops*, 24–25.