

Metadata-Augmented Neural Networks for Cross-Location Solar Irradiation Prediction from Satellite Images

Kuan-Ying Lee
r03922165@ntu.edu.tw
National Taiwan University
thingnario Inc.

Hsin-Fu Huang
h@thingnario.com
thingnario Inc.

Hung-Yueh Chiang
r05922005@ntu.edu.tw
National Taiwan University

Hu-Cheng Lee
r05922174@ntu.edu.tw
National Taiwan University

Winston H. Hsu
whsu@ntu.edu.tw
National Taiwan University

Wen-Chin Chen
wcchen@csie.ntu.edu.tw
National Taiwan University

ABSTRACT

Photovoltaics (PV) energy has witnessed a rapid growth in the passing years [13], raising the importance of monitoring the efficiency of PV generation. Monitoring can be based on measuring the ratio of real to optimal production in a period of time. Optimal production is derived from solar irradiation¹ that conventionally referred from preset pyranometers. In this work, we proposed to estimate solar irradiation in a time period through processing a series of satellite images (See Fig. 1). With a novel fusion method of important single-valued metadata (Sun position) and an autoencoder to leverage the spatiotemporal recurrence of satellite images, our method could reliably provide irradiation measurements in scenarios where pyranometers are unavailable (e.g. arbitrary location) or malfunction (e.g. dust, miscalibration). Our method performed on par with the state-of-the-art meteorological model and predictions closely aligned with ground truths from meticulously maintained pyranometers. Moreover, the Mean Absolute Percent Error (MAPE) is 25% lower than that of Solargis, a commercialized solar data provider designated by The World Bank Group.

KEYWORDS

solar irradiation, remote sensing, neural networks, computer vision

ACM Reference Format:

Kuan-Ying Lee, Hsin-Fu Huang, Hung-Yueh Chiang, Hu-Cheng Lee, Winston H. Hsu, and Wen-Chin Chen. 2018. Metadata-Augmented Neural Networks for Cross-Location Solar Irradiation Prediction from Satellite Images. In *MileTS '19: 5th KDD Workshop on Mining and Learning from Time Series, August 5th, 2019, Anchorage, Alaska, USA*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/1122445.1122456>

¹Solar irradiance integrated over time is solar irradiation. We used these two terms interchangeably in this paper.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MileTS '19, August 5th, 2019, Anchorage, Alaska, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9999-9/18/06...\$15.00

<https://doi.org/10.1145/1122445.1122456>




Figure 1: We attempt to leverage the global coverage of satellite images to achieve cross-location irradiation prediction, which is widely used in efficiency monitoring and PV site selection.

We observed the generality of Sun position and proposed a novel fusion method between images and metadata. Hence, our model could be applied on a much broader range while only trained on a small area. [Best viewed in color.]

1 INTRODUCTION

PV grew the fastest among renewable energy, with a growth rate over 50% in 2017 and contributed for almost half of the net addition in renewable energy [4]. PV installation is projected to grow exponentially to 400 GW by 2030, over four times the number in 2016 [10]. This uprising trend makes monitoring of PV generation efficiency more and more critical [20].

Operators often use Performance Ratio (PR) as the primary metric for efficiency evaluation. PR is defined as the ratio between the actual and the theoretical energy outputs, which is directly affected by solar irradiance. Practitioners often resort to pyranometers for solar irradiance. Yet, several issues exist concerning pyranometers.

To begin with, Pyranometer installment could be costly. The price starts from \$800 for a modest-quality all the way up to \$5000 for a highest-quality pyranometer². In addition, pyranometers are affected by solar dome effect and dusts, which necessitates regular cleaning [11, 14]. On top of that, pyranometers require recalibration from time to time [9]. Also, during site selection for new solar plants, there are no preset pyranometers. Hence, one could often only refer to observatories nearby for annual irradiation. Such measurements often seriously deviate from the ground

²Reference from Omni Instrument. <https://www.omniinstruments.co.uk>

truth irradiation (cf. Fig. 5 in supplementary material), a determinant impact factor for estimating the potential of a solar site.

These reasons motivated us to develop a framework that could report solar irradiation cross-locationally and historically. The framework could aid in both efficiency monitoring and site selection. With further investigation, we discovered among several influencing factors such as cloud coverage, aerosol concentration, terrain and etc., cloud is of the most decisive factor [19]. Based on such insight, we made a novel proposal to apply ConvNet on off-the-shelf satellite images. Our contributions summarize as follows:

- We collected SATI Dataset for future research on cross-location irradiation prediction.
- We identified the crucial role of meteorological domain knowledge and investigated fusion of two diverse modalities. e.g. image and metadata.
- We uncovered the spatiotemporal recurrence inside satellite images and designed an autoencoder to leverage such property to unsupervisedly enhance model training.
- We achieved promising performance to the state-of-the-art meteorological framework in cross-location irradiation prediction with simply satellite images.
- The MAPE of our method is 25% lower than that of Solrgis, a commercialized solar data provider.

2 RELATED WORK

Many work have been focusing on solar irradiance forecasting with different approaches and media. Researchers have tried Numeric Weather Prediction schemes [17], all-sky camera [7] and wireless sensor network of cheap light sensors [1] for solar irradiance forecasting.

Different to the above approaches, which required additional equipment for irradiance forecasting, researchers [3, 18] utilized the weather forecasting from nearby locations as the foundation for irradiance forecasting at the target site.

These methods have mitigated the needs of instrumentation to a certain extent. Yet, the supposition of existence of weather stations nearby is still too strong (the average distance to the nearest station in Taiwan is roughly 17 km. 9 km for the USA before removal of stations without irradiation.) Such resolution renders the transferability of weather quantities from nearby stations questionable considering the fickle nature of weather (cf. Fig. 5 in supplementary material.)

Also, the aforementioned work all dedicated to the temporal aspect (changing time only) of solar irradiance prediction, that is, forecasting the solar irradiance of a fixed location. Few addressed the spatial aspect, namely, the prediction of solar irradiance at any place.

Such service is beneficial not only to monitoring of PV power generation, which is commonly based on expensive and meticulously-tuned pyranometers, but also to site selection where the estimation of the potential accumulated irradiation at an arbitrary geolocation is needed. We hence proposed a framework capable of reporting cross-location solar irradiation both currently and historically.

The most related work to ours is Kosmopoulos et.al (2018) [15], which provided cross-location irradiance through the use of neural networks on results obtained from radiative transfer model (RTM)

LibRadtran, which achieved state-of-the-art estimation of solar irradiance.

While their method did provide a cross-location assessment to solar irradiance, it required numerous types of data input inclusive of aerosol optical density, water vapor and ozone column, which are rather inapplicable in practical scenarios. On the contrary, we relied simply on easily accessible satellite images and achieved a comparable and likely complementary performance.

3 SATELLITE IRRADIATION DATASET (SATI)

SATI contains two sub datasets, SATI-Taiwan and SATI-BSRN. SATI-Taiwan contains 30 stations located in Taiwan, of which we held out 3 stations for validation. SATI-BSRN contains 2 stations – Fukuoka, Japan (33.5822, 130.3764) and Kwajalein (8.7200, 167.7310) and we used it solely to validate cross-location prediction capability of the proposed method. SATI-Taiwan and SATI-BSRN contain pairs of six satellite image patches (each separated by 10 minutes) cropped from the Japanese satellite Himawari-8 and a hourly irradiation ground truth from Central Weather Bureau and Baseline Surface Radiation Network [16] respectively.

There are 550*550 pixels for the RGB visible channel and the resolution is 1 km/pixel, while for infrared channel, there are 550*550*1 pixels with resolution being 2 km/pixel. Note that the ground truth irradiation we referred to in both the two sub datasets were from stations where pyranometers are under regular cleaning and calibration of radiation experts and meteorologists in charge [8, 21]. Also though the images are available for download after roughly 20 minutes after being taken, it does no harm to efficiency monitoring and site selection where inspection is usually performed on a daily or annual basis.

Data in SATI-Taiwan are from July 7th, 2015 to March 31st, 2018 and data in SATI-BSRN are from January 1st to December 31st 2017 (for Kwajalein, we took data from September 1st 2016 to August 30th 2017 since the station closed at Sep. 2017). (Please refer to Fig. 7 in supplementary material for more details.)

3.1 Data Preprocessing

We aligned each station to the corresponding pixel on the satellite images. Then we cropped the surrounding 25*25 pixels for both the visible and infrared channels, which were stacked to form a 4-channel patch.

Next, we associated 6 ten-minute-gapped patches in an hour to the corresponding hourly irradiation. For instance, we associated patches of Station A at 09:10, 09:20, ..., 10:00 to the irradiation from 09:00 to 10:00 of Station A. Hence, a (X, Y) pair consists of six four-channel patches as X and the corresponding irradiation as Y .

4 PROPOSED METHOD

We proposed METadata-augmented REcurrent network for Solar irradiation Assessment (MERESA). We separated the model into three parts, namely, irradiation predictor, metadata decoder and spatiotemporal autoencoder. For each part, we will illustrate our logic of design and the insight we uncovered from the datasets. See the framework in Fig. 2.

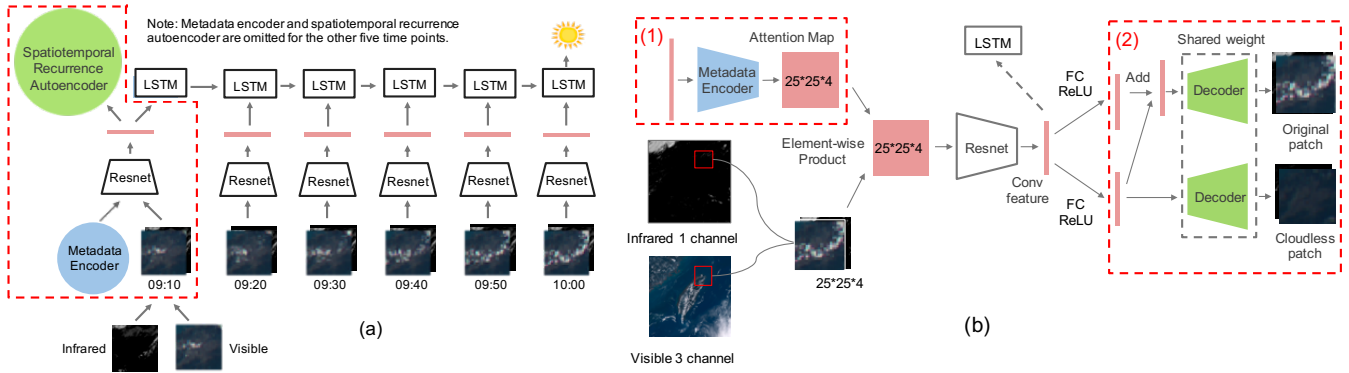


Figure 2: (a) We combine visible and infrared patches. An attention map derived from metadata encoder that captures the rough Sun position is multiplied to highlight the more influential (sunlit) areas, shown in (b)(1). Then a ResNet-LSTM model is deployed to extract spatial and temporal relationships between patches. Also, shown in (b)(2), to aid the training process under limited stations and data, we proposed novel multitasking learning to utilize the repetitiveness of the same stations in the dataset. Such weak supervision is enforced by reconstruction to the original and a cloudless patch of the same place. [Best viewed in color.]




Figure 3: We convert Sun positions (domain knowledge) into attention maps that could capture the most sunlit region on satellite images, where cloud dynamics affect irradiation the most. We then weighted the satellite images accordingly with the attention maps.

4.1 Irradiation Predictor

For image feature extractor, we simplified off-the-shelf ResNet16 [12] by changing the first 7x7 convolution to 3x3 and cutting the number of channels in each layer to better fit our image size. Also, we removed max pooling, which is potentially harmful for regression tasks on satellite images as discussed in Chen et al. [6]. We trained ResNet from scratch since our target bears a large difference to object recognition task such as ImageNet. After ResNet, we placed a LSTM to capture the temporal relationships between patches, which outputs the predicted hourly irradiation.

4.2 Metadata Encoder

Due to the alternation of day and night and the change of four seasons, time and date largely impacts the irradiation distribution received at a given location. Intuitive solution would be to provide date and time as auxiliary inputs to the model. Yet, one place could still have daylight, while another is already after nightfall even at the same local time due to altitude difference. Therefore, we proposed to use Sun position instead.

Given date, time and geolocation, we computed the corresponding azimuth and elevation, which we then transformed into two one-hot vectors of length 360 and 180 respectively after rounding to the nearest integer. Then, the two vectors were concatenated to a single vector, which was then fused with conv features of the satellite images.

We designed an attention-like early fusion method that guide the network to focus on specific regions of the input patch that are most sunlit. Specifically, We designed an encoder to transform the one-hot vector into a map with the same size as the satellite image, which is then multiplied on the image to obtain the final input to ResNet, as shown in Fig. 2(b)(1) with detail parameters in Fig. 3.

The metadata encoder comprised of one fully connected layer and four transpose convolution layers with batch normalization before the ReLU, whose kernel size are all 3 and stride is 2, 1, 2, 2 respectively.

4.3 Spatiotemporal Recurrence Autoencoder

Since the number of data is limited and only single value irradiation is available for supervision, we proposed a novel method to harness the implicit patterns both spatially and temporally in the dataset for multitask learning that assists model training (cf. Table 5).

The SATI-Taiwan Dataset contains satellite images for two and half years while includes irradiations for the same 30 stations. Based on such spatial and temporal repetitiveness, we could find cloudless patches for a chosen station at a specific time and around an arbitrary day of year. Hence, even without labeled information of cloud, the model could pick up land and partial cloud information by comparing the input patch to a cloudless one.

We hence designed an autoencoder to unsupervisedly leverage this information and guided the network by force of reconstruction from the conv feature to both the clean and the original patch.

Specifically, we split the conv feature from ResNet into two, shown in Fig. 2(b)(2). The first feature noted goes through a decoder and outputs a reconstructed clean patch. Then Mean Square Loss (MSE) is computed with the clean patch. For this feature, we try to capture the surrounding landscape information of an arbitrary station.

The second feature is added with the first part, goes through the same decoder, and outputs a patch, where MSE is computed with the original patch. For this part, we try to catch cloud information.

5 EXPERIMENTS

5.1 Evaluation Metrics

The evaluation was based on Relative Root Mean Squared Error (rRMSE), that quantifies the spread of errors and makes comparisons across seasons and stations more reasonable. And to quantify the deviation from ground truth, we used Mean Absolute Error (MAE).

5.2 Metadata Fusion

We explored the commonly used multi-modality fusion methods on satellite image and metadata features including concatenation, inner and outer product.

The annual rRMSE on SATI-BSRN of different fusion methods are summarized in Table 1. We could observe the performance of early attention outperformed the other commonly used fusion methods, further confirming the suitability of the proposed design. Also, the performance of the models that were trained with Sun position were better than those trained with time, showcasing the necessity and generality of Sun position in cross-location prediction.

In addition, we visualized the attention maps of a validation station from morning to afternoon. We could observe the attention area on the map shifted with time as if capturing the Sun trace. (cf. Fig. 8 in supplementary material.)

	Concat	Inner Product	Outer Product	Attention Map
Time	0.2932	0.2978	0.2996	0.3071
Sun position	0.1860	0.2377	0.1686	0.1627

The rRMSE without metadata is 0.2450.

Table 1: The rRMSE of different fusion methods on SATI-BSRN. Performance of the models based on time dropped dramatically since time alone is insufficient to refer Sun position without latitude. Those based on Sun position are maintained well when tested on different geolocations.

5.3 Validation in Different Geo-locations

We reported the seasonal and annual rRMSE of three validation stations of SATI-Taiwan and two stations of SATI-BSRN in Table 2.

The moderate performance decay in cross-geolocation scenarios testified that our model, despite trained in a small geolocation range, could be applied on a wide range of geolocations without tuning.

	SATI-Taiwan			SATI-BSRN	
	Taipei	Taichung	Kaohsiung	Fukuoka	Kwajalein
Spring	0.1503	0.1605	0.0936	0.1393	0.1434
Summer	0.1630	0.1689	0.1255	0.1399	0.1413
Autumn	0.1604	0.1470	0.1096	0.1677	0.2042
Winter	0.1737	0.1278	0.0944	0.2548	0.1221
Annual	0.1609	0.1543	0.1095	0.1557	0.1994

Table 2: The rRMSE of the proposed model on SATI Dataset. The proposed model performed steadily on different geolocations while only trained on SATI-Taiwan.

5.4 Comparison to Related Work

We compared our method to Kosmopoulos et.al [15], which claimed the state-of-the-art performance on BSRN stations and focused on cross-location irradiance prediction. Since the BSRN stations they referred to are different from ours, which are mostly located in Europe and beyond the coverage of Himawari-8, we resorted to the mean errors computed with all stations (which to a certain extent, cancels variations in different stations.) Note that under such constraints of non-overlapping stations, the comparison by no means indicates absolute dominance of one over the other but to provide a rough idea of how well the proposed method performs.

The reported hourly mean difference in Kosmopoulos et.al [15] is -70 and 40 W/m^2 at 25th and 75th percentile, which we converted into MAE under the most optimistic scenario. We assumed no deterioration from both 25 to 0 quantile and 75 to 100 quantile and a zero error from 25 to 75 quantile (Largely in favor of their method). Under such assumption, the MAE is $(|-70| * 25 + |40| * 25) / 100 = 27.5 \text{ W/m}^2$. We then multiplied it by 3600 to obtain the approximate accumulated hourly irradiation.

The comparison is shown in Table 3. We demonstrated with solely satellite images, we performed comparatively with the state-of-the-art meteorological model, which requires multiple data sources like aerosol, ozone column, just to name a few.

5.5 Comparison to Commercialized Product

We compared our model to Solargis, the solar irradiation provider designated by the World Bank Group. It claims state-of-the-art performance combining irradiance from clear-sky model and cloud index calculated with both the visible and infrared channels of satellite images. We purchased annual irradiation data of 2018 at two of the validation stations in SATI-Taiwan and compared our results with theirs. In Table 4, we could observe with solely satellite images, our model achieve lower daily and hourly MAPE compared to Solargis, which requires multiple data sources inclusive of atmospheric parameters, environmental variables and satellite images.

Note that our predictions are currently based on Himawari-8 and we are extending to other satellites around the globe in the future.

	Kosmopoulos et.al [15]	Ours
Annual MAE (MJ/m ²)	0.099 ^a	0.099

^a This is an approximate comparison, due to the non-overlapping data sources.

Table 3: Our method performed comparatively with the state-of-the-art meteorological method, which mostly relies on numerical approaches with a few learning based tools. It shows that we provide a new and yet complementary aspect for solar irradiation prediction based on CNN and satellite images.

	Taipei		KaoHsiung	
	Hourly	Daily	Hourly	Daily
MAPE ↓				
Ours	0.558	0.289	0.294	0.318
Solargis	0.794	0.488	0.444	0.583

Table 4: MAPE comparison to Solargis in both hourly and daily bases. With simply satellite images, our hourly MAPE is lower than Solargis by 25% in absolute scale, which utilizes multiple data sources inclusive of atmospheric parameters, environment variables and satellite images.

6 CONCLUSION

We proposed a cross-location irradiation prediction framework based on off-the-shelf satellite images. The method is shown complementary to traditional meteorological ones.

We identified the significance of domain-specific metadata (e.g., Sun position) and proposed a novel fusion method. Such findings can be extended to other neural network applications where metadata is informative besides raw (and low-level) signals. Observing the spatiotemporal recurrence in satellite patches, we designed a recurrence autoencoder to aid the training process.

We verified our model against two stations of different geolocations and showcased its feasibility as a global solar irradiation predictor. And we achieved a promising performance compared to the state-of-the-art framework with satellite images alone and outperformed Solargis that is widely used by the World Bank Group.

We will release the SATI Dataset for future research. For future work, besides exploration of more satellite and input patch size, we would investigate the possibility of irradiation forecasting based on satellite images with meteorological quantities, such as precipitation, wind, visibility, etc.

7 ACKNOWLEDGMENTS

This work was supported in part by the Ministry of Science and Technology, Taiwan, under Grant MOST 108-2634-F-002-004. We also benefit from the NVIDIA grants and the DGX-1 AI Supercomputer.

REFERENCES

[1] Stefan Achleitner, Ankur Kamthe, Tao Liu, and Alberto E Cerpa. 2014. SIPS: Solar irradiance prediction system. In *Information Processing in Sensor Networks*. IEEE, 225–236.

[2] Matthew Anderson, Ricardo Motta, Srinivasan Chandrasekar, and Michael Stokes. 1996. Proposal for a standard default color space for the internet—srgb.

In *Color and imaging conference*, Vol. 1996. Society for Imaging Science and Technology, 238–245.

[3] Kuk Yeol Bae, Han Seung Jang, and Dan Keun Sung. 2017. Hourly solar irradiance prediction based on support vector machine and its error analysis. *Transactions on Power Systems* 32, 2 (2017), 935–945.

[4] The British Petroleum Company plc BP. 2018. BP Statistical Review of World Energy. (2018). <https://www.bp.com/content/dam/bp/en/corporate/pdf/energy-economics/statistical-review/bp-stats-review-2018-full-report.pdf>

[5] Carlos Campillo, Rafael Fortes, and Maria del Hénar Prieto. 2012. Solar radiation effect on crop production. In *Solar Radiation*. InTech.

[6] Boyo Chen, Buo-Fu Chen, and Hsuan-Tien Lin. 2018. Rotation-blended CNNs on a New Open Dataset for Tropical Cyclone Image-to-intensity Regression. (2018).

[7] Hsu-Yung Cheng and Chih-Chang Yu. 2015. Multi-model solar irradiance prediction based on automatic cloud classification. *Energy* 91 (2015), 579–587.

[8] Amelie Driemel, John Augustine, Klaus Behrens, Sergio Colle, Christopher Cox, Emilio Cuevas-Agulló, Fred M Denn, Thierry Duprat, Masato Fukuda, Hannes Grobe, et al. 2018. Baseline Surface Radiation Network (BSRN): structure and data description (1992-2017). *Earth System Science Data* 10, 3 (2018), 1491–1501.

[9] Lawrence Dunn, Michael Gostein, and Keith Emery. 2012. Comparison of pyranometers vs. PV reference cells for evaluation of PV array performance. In *Photovoltaic Specialists Conference (PVSC), 2012 38th IEEE*. IEEE, 002899–002904.

[10] Galen Margolis Robert Bolinger Mark Chung Donald Fu Ran Seel Joachim Davidson Carolyn Darghouth Naim Wiser Ryan Feldman, David Barbose. 2015. Photovoltaic System Pricing Trends: Historical, Recent, and Near-Term Projections 2015 Edition. (2015).

[11] D Feuermann and A Zemel. 1993. Dust-induced degradation of pyranometer sensitivity. *Solar Energy* 50, 6 (1993), 483–486.

[12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*. 770–778.

[13] International Energy Agency IEA. 2017. Renewables2017 – Analysis and Forecasts to 2022. (2017).

[14] Qiang Ji and Si-Chee Tsay. 2000. On the dome effect of Eppley pyrgometers and pyranometers. *Geophysical Research Letters* 27, 7 (2000), 971–974.

[15] Panagiotis G Kosmopoulos, Stelios Kazadzis, Michael Taylor, Panagiotis I Raptis, Iphigenia Keramitsoglou, Chris Kiranoudis, and Alkiviadis F Bais. 2018. Assessment of surface solar irradiance derived from real-time modelling techniques and verification with ground-based measurements. *Atmospheric Measurement Techniques* 11, 2 (2018), 907–924.

[16] et al. Ohmura, Atsumu. 1998. Baseline Surface Radiation Network (BSRN/WCRP): New precision radiometry for climate research. *Bulletin of the American Meteorological Society* 79, 10 (1998), 2115–2136.

[17] Richard Perez, Sergey Kivalov, James Schlemmer, Karl Hemker Jr, David Renné, and Thomas E Hoff. 2010. Validation of short and medium term operational solar radiation forecasts in the US. *Solar Energy* 84, 12 (2010), 2161–2172.

[18] Xiangyun Qing and Yugang Niu. 2018. Hourly day-ahead solar irradiance prediction using weather forecasts by LSTM. *Energy* 148 (2018), 461–468.

[19] Arun Pratap Singh Rathod, Poornima Mittal, and Brijesh Kumar. 2016. Analysis of factors affecting the solar radiation received by any region. In *Emerging Trends in Communication Technologies (ETCT), International Conference on*. IEEE, 1–4.

[20] Ray Melendez Aldo P Hunt WD Sullivan, Greg Pugh. 2010. *Operations & Maintenance Best Practices-A Guide to Achieving Operational Efficiency (Release 3)*. Technical Report. Pacific Northwest National Lab.(PNNL), Richland, WA (United States).

[21] Sheng-Hsiang (Carlo) Wang. 2017. Solar radiation measurement networks in Taiwan. ftp://ftp.pmodwrc.ch/stealth/ipc-xii/Seminar/surface_radiation_in_tw.ppt

A APPENDICES

B IMPLEMENTATION DETAILS

B.1 Training Details

All trainings were carried out on a Nvidia GeForce GTX 1080 Ti, and optimized by Adam with betas (0.9, 0.999), learning rate $1e^{-4}$ and zero weight decay. The weight initialization was Xavier for both convolution and fully connect layer. Batch size is 128. All satellite images were normalized with subtraction of 127 followed by division of 127 into (-1, 1) range. Training loss is based on Mean Squared Error for both irradiation prediction and image reconstruction. All experiments and framework are developed and deployed with pytorch 0.4.1, CUDA 8 and cuDNN 5.1 on Ubuntu 16.04.

B.2 Geographic Coordinate to Pixel

To map a pair of latitude and longitude to the corresponding pixel on satellite images, we first manually labeled four anchors by looking for landmarks that are not only searchable online for latitude and longitude but also lie in discriminative pixel on satellite images that is easy to spot (e.g. lighthouse located at the border of land and ocean).

Then, with either three of the four anchors, we could obtain the number of pixels in width and height covered by one degree of latitude or longitude (e.g. 1° increase in latitude maps to 100 pixel change in height and 10 pixel change in width). Obtaining the mapping scale between latitude/longitude and pixel, we could get the pixel coordinate by converting the latitude/longitude difference to pixel difference between the target location and one of the three anchors that we chose. Since we could choose between four different anchor sets, we averaged the four obtained pixel coordinates to obtain the final one. Note due to projection distortions, the above method could only be used locally. That is, one will have to find different anchors when applying the aforementioned method onto different regions across the globe.

B.3 Search algorithm for Cloudless Patch

To obtain a supporting clean patch for an arbitrary day and time, we applied an heuristic based on our observation that clouds are always whitish. This means the average brightness of a patch with more clouds is often higher if all the comparing patches are of the same place and of the same day and time.

For the visible part, we first converted RGB into intensity following the method proposed in Matthew et al. (1996) [2]. We then grouped the patches into 8760 bins ($365 * 24$, based on a hourly resolution omitting minute and 29th Feb. in leap years). Then given an arbitrary input patch, we searched from the bins of the same hour in a date range, specifically 10 days, before and after the given day of year. The patch with the lowest mean brightness is selected as the visible part of the clean patch. And the cloudless patch for the infrared channel is simply an one-channel image with nothing but zeros.




Figure 4: Data augmentation. (a) The original satellite image. (b) The naive flipping breaks the mutual relationship between Sun and clouds. (c) Flipping with Sun position. [Best viewed in color]

C OTHER EXPERIMENTS

C.1 Ablation Studies

We performed a series of ablation studies to further verify the effectiveness of the major components, namely, the metadata attention map, the spatiotemporal recurrence autoencoder and the metadata based data augmentation. We sequentially removed the components and the continual drop in performance could be noticed in Table 5.

Note that fusion of Sun position also facilitated data augmentation in the task. While techniques like horizontal and vertical flipping are broadly used in various image tasks, they are ineffective and even harmful to solar irradiation prediction since relative positions of Sun and clouds matter and naively applying these techniques disrupts the mutual relationships. Yet, if we alter the sun position correspondingly, we could safely perform data augmentation, as illustrated in Fig. 4. We could also see in Fig. 5 that if we performed data augmentation without changing metadata correspondingly, the performance is even slightly inferior to the model with no augmentation.

Infrared Channel	Metadata Attention	Auto Encoder	Augmentation	rRMSE	MAE
-	-	-	-	0.2604	0.1305
V	-	-	-	0.2635	0.1247
V	V	-	-	0.1487	0.0859
V	V	V	-	0.1466	0.0718
V	V	V	V	0.1403	0.0706
V	V	V	w/o Metadata	0.1484	0.0736

Table 5: Ablation studies on SATI-Taiwan. The improving performances demonstrate the effectiveness of each proposed component.

C.2 Performance in Real Scenarios

We have deployed our framework onto several running solar plants where pyranometers were instrumented and regularly cleaned and calibrated. We referred to these pyranometers for ground truth and validated our prediction against the readings. Note that in practice, pyranometers are not always available and in regular maintenance. We set up a baseline by directly using the readings from the closest weather station for comparison with the proposed method.

The sites we referred to are located in the southern part of Taiwan for its stable and ample sunlight. The geographic coordinates

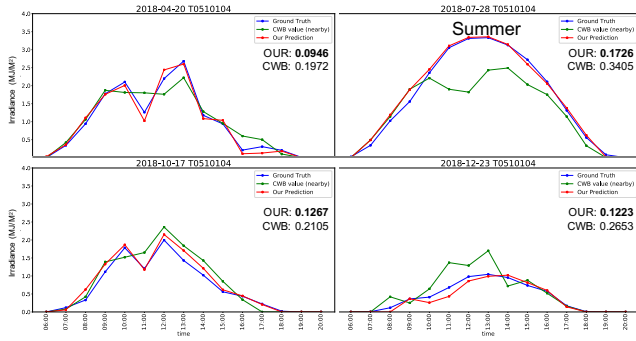


Figure 5: Real cases with seasonal rRMSE shown aside. Our predictions fit well to ground truth irradiances compared to readings from even the closest weather station, which is 28.74 km away. It demonstrates the deficit of referring to the closest weather station and the necessity of our method concerning weather locality. [Best viewed in color.]

of the site is (23.61762, 120.1842), and the closest weather station is 28.74 km away.

We randomly picked a day from each of the four seasons and plotted the irradiation from the pyranometer (ground truth), the closest station and our framework. As shown in Fig. 5, our prediction curves not only lie closer to the ground truth ones but the seasonal rRMSE are smaller compared to that of the closest station in all cases. Fig. 5 also demonstrates the locality of weather and showcases the potential issues of naively referring to the closest weather station since weather status could vary in even several kilometers.

D REAL-SCENARIO AND POTENTIAL CASES

D.1 Site Selection

Our method could assist PV site selection where accumulated irradiation over a period of time of arbitrary locations are required. In these scenarios, no prior pyranometers could be referred to and using readings from the closest weather station is prone to large errors (cf. Subsec. C.2). With the proposed method, a user could specify an arbitrary location with a query time range, and obtain the corresponding accumulated irradiation.

D.2 Efficiency Monitoring

We calculated the PR value by dividing the actual generation output of the PV panels reported from the inverters to the theoretical output calculated with the predicted irradiation from our method. When PR is below a preset threshold for too long, we issue a warning to our clients, recommending them to schedule an examination for possible stains, cracks and malfunctions. With the warning mechanism, we could largely prevent possible financial loss due to unnoticed efficiency degradation.

The mechanism could not only be used on sites that lacks of pyranometer instrumentation (prevalent in many roof-top and community sites), but also as a complement to pyranometers for efficiency trend monitoring and possible indication of potential pyranometer malfunction.

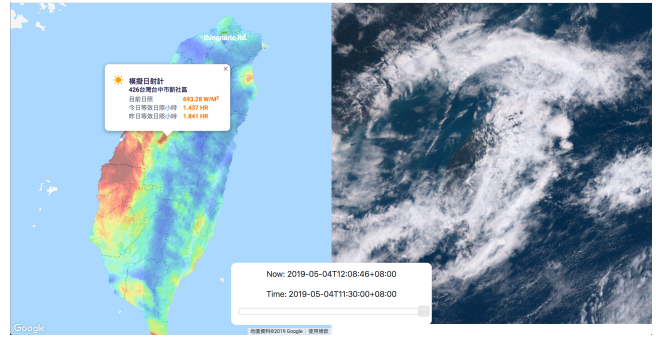


Figure 6: Irradiation heatmap produced with the proposed method and the corresponding satellite image at the same timestamp. We could see the cloud distribution is highly correlated to that of irradiation. [Best viewed in color.]




Figure 7: (a) We took the corresponding patches from the full-disk satellite image where SATI stations lie in, noted as b, c and d. We further cropped a 25*25 surrounding region for each station, shown as yellow rectangle in (b), (c) and (d). (b) The training (green) and the three validation stations (red) in SATI-Taiwan. (c)(d) The two BSRN stations. [Best viewed in color.]




Figure 8: The weighted area shifts with Sun from morning to afternoon and reflects the Sun incident angles from East to West in the metadata-encoded attention maps (Taichung station on July 9th, 2015.). Note the corresponding elevation and azimuth are listed aside.

D.3 Environmental and Agricultural Impact

Solar radiation has an important factor on crop production [5], as it affects other factors like soil moisture and plant's rate of evaporation. Hence, potentially with the proposed method, agricultural decisions like where and when to plant certain plants could be made with ease as different plants require different amount of irradiation.

D.4 Irradiation Web Service

We have set up a website (<http://thingnario-fs.synology.me:39080/?email=reviewers@easychair.org&password=admin>) that provides irradiation on every single point of Taiwan with a rough 30 minute delay. A screenshot of the website and its corresponding satellite

image are shown in Fig. 6. The colors on the screenshot represent different quantity of irradiation. We could observe that the distribution of visual cloud thickness in the satellite image highly matches with the intensity of the heatmap. Users could click on an arbitrary location to view the current irradiation, sunshine hour of today and yesterday, enabling other possible irradiation related applications.