
Comparing the carbon costs and benefits of low-resource solar nowcasting

Ben Dixon

University College London (UK)
benjamin.dixon.21@ucl.ac.uk

Jacob Bieker

Open Climate Fix
jacob@openclimategfix.org

María Pérez-Ortiz

AI Centre, University College London (UK)
maria.perez@ucl.ac.uk

Abstract

Mitigating emissions in line with climate goals requires the rapid integration of low carbon energy sources, such as solar photovoltaics (PV) into the electricity grid. However, the energy produced from solar PV fluctuates due to clouds obscuring the sun's energy. Solar PV yield nowcasting is used to help anticipate peaks and troughs in demand to support grid integration. This paper compares multiple low-resource approaches to nowcasting solar PV yield. To do so, we use a dataset of UK satellite imagery and solar PV energy readings over a 1 to 4-hour time range. Our work investigates the performance of multiple nowcasting models. The paper also estimates the carbon emissions generated and averted by deploying models such as these, and finds that short-term PV forecasting may have a benefit several orders of magnitude greater than its carbon cost and that this benefit holds for small models that could be deployable in low-resource settings around the globe.

1 Introduction

Solar photovoltaics (PV) is one of the safest and cleanest sources of energy, as measured in death rates from accidents and air pollution, and greenhouse gas emission [10]. However, the fluctuations in the energy output, or yield, from solar PV makes it challenging for the electricity grid to rely at large scale on solar energy. The yield of solar PV is inherently uncertain because clouds can obscure the sun, and clouds can become thicker, thinner, and their movement across the sky is hard to predict [2]. In order to guarantee stable energy supply, grid operators often have to keep gas spinning-reserve online, releasing carbon dioxide into the atmosphere [6]. Forecasting over less than one day in advance is called 'nowcasting' [12] to emphasize its dynamic use. Nowcasting solar energy input into the grid accurately could reduce the need to keep the reserves on standby, allowing to use every kilowatt efficiently and enhancing the grid resilience to the fluctuations of renewable energy sources.

Since 2019, Open Climate Fix (OCF) has been working with National Grid ESO, the electricity system operator for Great Britain, to build accurate PV nowcasting models that can enhance the resilience of the energy grid [4]. This paper complements the work done by OCF. Specifically, we research alternative low-resource methods (in terms of smaller/shallower neural networks architectures) for predicting PV yield, and analyse their performance, their carbon footprints and aim to estimate, even if roughly, the net emissions that could be averted by their potential use.

As more of the world incorporates solar PV into energy grids, there will be an increased need to develop accurate solar PV forecasts. While a body of recent work has emphasized building more computationally efficient models, most of the machine learning work still focuses on the opposite: building larger models with more parameters to tackle more complex tasks [7]. Recent

work poses the question of how much is the performance gain worth [7], specially in these applications where deployment of a simple technology could already have a big impact. The move towards less computationally intensive models has already happened once in weather forecasting, as deep learning, despite its competitiveness, has often lower computational requirements than numerical weather prediction based on differential equations [11, 12].

Through our work, we study low-resource and shallower predictive models and try to estimate the benefit that they could bring if they were deployed large-scale for solar nowcasting¹. The models, which are build on one Tesla T4 GPU available through Google Colab, could be deployed large-scale around the world independently of hardware limitations. The reason we are interested in researching this low-resource setting, by studying smaller models, is because those are the models that could be deployed to support the grid in emerging economies, which are where the real need is for clean energy solutions, i.e. where the energy demand is rapidly growing.

2 An empirical comparison of PV yield nowcasting: Experimental results

For the task at hand, we study convolutional neural networks (CNNs), Long Short-Term Memory networks (LSTMs) and their intersection [12] (ConvLSTMs). In theory, this comparison may seem futile as LSTMs and ConvLSTMs should outperform CNNs, because of their use of state information. However, Bai et al [1] argue that simpler CNNs can often outperform LSTMs. Moreover, ConvLSTMs require a large amount of memory simultaneously available in order to process and update the state information used in its predictions, which means they often take longer to train, leading to more energy usage and hence higher emissions.

Dataset Solar generation data refers to the yield from specific solar panels. The dataset² is composed of satellite imaging and PV readings. The satellite images are provided by EUMETSAT, covering the northern third of the Meteosat disc every five minutes. Open Climate Fix developed the `eumetsat_uk_hrv` dataset [5] which takes the ‘high resolution visible’ (HRV) channel for a reduced time period and geospatial extent. The dataset was reduced in this paper to meet computing resources available. A 64 x 64 crop was randomly chosen to focus on Devon, selecting images taken between 05:00 and 20:00 as there is minimal PV output outside of this window for most of the year. However, this still left many readings in winter when the sun was below the horizon, and the PV yield was zero, so readings when the solar altitude was below 10 degrees were also dropped. Open Climate Fix provided a dataset of 1311 energy yield readings from PV systems in 5 minute increments for most of 2018-2021, originally provided by a PV operator. The size of the dataset varied slightly based on the prediction window. For the shortest window, there were 2018 observations of (12, 64, 64) and (12, 1) to predict a series of 12 readings. For the longest prediction window, there were 659 blocks of the same sizes described earlier to predict a series of 48 readings.

Learning task We formulate our learning task as predicting a sequence of values representing the future PV yield. All models take as input 12 sets of 5 min data, i.e., 1 hour of data as input, and predict forward between 1 and 4 hours. We also consider two learning scenarios: i) learning exclusively from past PV data only, and ii) learning from both past PV and satellite data as inputs.

Models compared For each learning scenario (with and without satellite imaging), we test both CNNs and LSTMs, our objective being comparing their prediction efficiency and carbon footprint at the task at hand. Additionally, LSTMs encode time relationships explicitly, and we are interested in evaluating if this additional complexity increases the prediction performance of our models to a significant extent (and at what potential carbon cost). For the PV only models (CNN and LSTM), which are the simplest models we test, we take in past PV through multiple layers of CNNs, followed by a fully connected layer. The other PV only model has a similar structure but uses LSTMs in place of CNNs. The second group of models take both satellite images and past PV yield as input. The Conv3D is a CNN which uses multiple levels to try to learn abstract features. After each convolutional layer, a MaxPooling layer is used to reduce the dimensionality and learn the features. This is followed by two dense layers to resize the output to predict a sequence of values. The CNN treats the sequence

¹Our low-resource settings are not simulated, this project has been done as part of a MSc thesis dissertation, with no other computational resources.

²Thanks to Open Climate Fix for providing the datasets.

as one vector. The ConvLSTM model (see 2) is similar to the Conv3D but uses recurrent modelling, with ConvLSTMs to process the images and LSTMs to process the PV input sequence.

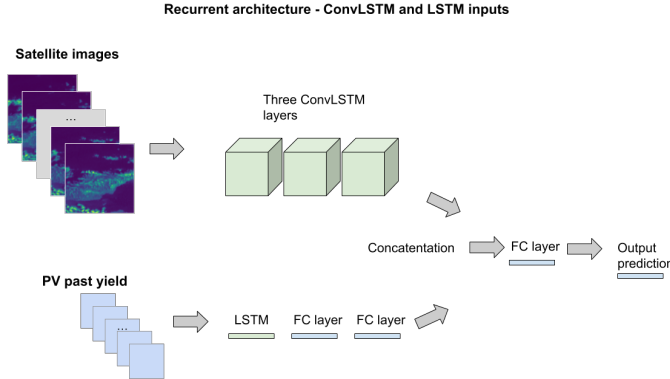


Figure 1: The ConvLSTM model, using both satellite images and past PV yield.

Model training and dataset split The model is trained on 22 months of data from 1/1/2020 to 7/11/2021. To avoid autocorrelation, the datasets are drawn from each month: days 1 to 20 are for training, days 22 to 24 for validation, and days 27 to 29 for testing. For all models, the models are trained to optimise MSE until convergence. Bayesian hyperparameter tuning was done: The learning rate was varied between [0.01, 0.001, 0.0001, 0.00001], and dropout rate was varied between [0.05, 0.5] in steps of 0.05. For the CNNs and the ConvLSTMs, the kernel size was varied between [3,5,7]. For ConvLSTMs the number of filters was varied between [16,32,64]. For the fully connected layers the number of nodes was varied in the set [12,24,48]. The ConvLSTM has three ConvLSTM layers applied to images, and one LSTM layer applied to past PV. The CNN model has three convolutional layers applied to the images, and one convolutional layer applied to past PV readings. For the PV-only models, the LSTM model has three layers of LSTMs, and the CNN model has three layers of CNNs.

2.1 Results

The task is to predict a sequence of values of variable lengths, from 1 hour to 4h. Table 2.1 shows normalised MAE and RMSE by model over different prediction windows. All models outperform the persistence baseline, often used for physical processes. ConvLSTM outperforms other models, especially over longer time horizons. Interestingly, an LSTM model taking PV only delivers results competitive with a Conv3D which also uses satellite data five out of eight times. Does time series modelling improves predictions? The recurrent designs have a lower error metric five out of eight times, but often by less than 0.2%. Do satellite images improve performance? Yes, on six of eight times, and by margins approximately increasing 0.2%-1.2%. Overall, it appears that explicit modelling of temporal information improves the performance to a large extent. This is most significant for the comparison between ConvLSTM and CNN, which is when the inputs include image data. Perhaps this is because image data carries important trend information, such as motion of clouds.

Time	Metric	Baseline	PV only		Sat and PV	
		Persistence	CNN	LSTM	Conv3D	ConvLSTM
60 mins	RMSE	11.60%	10.46%	9.42%	9.44%	9.29%
	MAE	7.91%	7.92%	6.60%	6.81%	6.60%
120 mins	RMSE	15.18%	11.86%	11.93%	11.77%	11.55%
	MAE	11.03%	8.72%	9.07%	8.55%	8.41%
180 mins	RMSE	17.84%	12.83%	12.94%	13.20%	12.26%
	MAE	13.10%	10.13%	10.00%	9.84%	9.15%
240 mins	RMSE	20.32%	13.16%	13.02%	13.75%	12.52%
	MAE	14.95%	10.16%	10.14%	10.08%	9.53%

Table 1: Normalised prediction error over different time horizons, with best results highlighted in bold.

2.2 Carbon emissions

This section estimates the energy requirements of our models. The carbon costs are calculated through the time taken multiplied by estimates of the carbon intensity of time using a GPU. Each of the models was run on a Tesla T4 GPU available through Google Colab, assumed to be in Northern Europe, the region closest to our location, which has a carbon efficiency of 0.21 kgCO₂eq/kWh. Energy usage simulations were conducted using the Machine Learning Impact calculator [8]. Table 2.2 uses the computational time required to train and make inferences over a two-hour prediction window and compare the ConvLSTM and 3D CNN. It is assumed that the predictions are generated for each of the 330 UK grid supply points, and are generated every hour, from 06:00 to 20:00, and for a whole year, giving $330 * (20-6) * 365 = 1686,300$ forecasts per year.

	Conv3D	ConvLSTM
Time to train model (s)	916.39	673.11
Time for inference, one forecast (s)	0.16	2.21
Implies: Total time for year (training + 1686,300 * inferences) (hrs)	75	1024
Emissions generated (tonnes CO₂-equiv) from [8]	0.0108	0.152

Table 2: Estimates of emissions generated through training and deploying Conv3D and ConvLSTM

The carbon reductions through increased model accuracy are harder to quantify. In 2020, the UK energy grid was estimated to produce 51 million tonnes of CO₂-equivalent emissions [9]. Currently, 2% of UK energy currently comes from solar PV, which could rise to 7% by 2050 [3]. Emissions of CO₂ per gigawatt hour of electricity over the lifecycle of solar plant is 5 tonnes [10], where gigawatt hour is the annual electricity consumption of 150 Europeans. Based on this data, and assuming 67.22 million people in UK and a similar electrical consumption than Europeans, a 2% of solar energy would generate annually 45,000 tonnes of emissions, which is 2450 times less than keeping gas-reserves online for generating the same electricity [10]. Even if we assume that the models in Table 2.2 are 0.1% better at estimating solar output than the current standard of energy grid operators (which would be a lower bound as in our analysis these models are still much better than persistence, often a baseline followed for physical processes), and that in this 0.05% of cases we may be able to turn off gas-reserves, we may be looking at a reduction of around 5500 tonnes of CO₂ annually. Given the costs of training the two models and generating the forecasts for a year are between than 1/10th and 1/6th of a tonne of CO₂-equivalent, the benefits far outweigh the costs. This puts in perspective the reduction that could be possible if we were to deploy these low-resource models at large scale.

3 Conclusions

This paper includes an estimate of carbon emissions for training and running the inference of low-resource shallow neural architectures for solar photovoltaic nowcasting for a year. We attempt to estimate as well the emissions that would be averted by having a better integration of solar in the energy grid, showing that even under very pessimistic assumptions there is an important benefit in terms of total carbon footprint averted from integrating solar nowcasting machine learning models into the grid. Our analysis focuses on low-resource models that could be integrated into emerging economies, where energy requirements may be growing quickly. Further analyses are needed to understand the benefit that larger and deeper models could bring. A more in depth analysis of the current strategies implemented by grid operators used to meet energy demands in case of disruption to the renewable solar supply would be needed to better put in perspective the climate change mitigation potential of nowcasting models.

References

- [1] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. *An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling*. URL: <https://arxiv.org/abs/1803.01271>.
- [2] Olivier Boucher and David Ransall. "Clouds and Aersols". In: *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I. to the 5th Assessment Report of the Intergovernmental Panel on Climate*. Cambridge University Press, 2013, pp. 571–657.

- [3] National Grid ESO. *Future Energy Scenarios: July 2022*. 2022. URL: <https://www.nationalgrideso.com/future-energy/future-energy-scenarios>.
- [4] National Grid ESO. *The Road to Zero Carbon*. 2022. URL: <https://www.nationalgrideso.com/future-energy/our-progress/road-zero-carbon/report>.
- [5] Open Climate Fix. *EUMETSAT Dataset*. 2022. URL: <https://huggingface.co/datasets/openclimatefix>.
- [6] David Wenzhong Gao. "Chapter 2 - Applications of ESS in Renewable Energy Microgrids". In: *Energy Storage for Sustainable Microgrid*. Ed. by David Wenzhong Gao. Oxford: Academic Press, 2015, pp. 35–77. ISBN: 978-0-12-803374-6. DOI: <https://doi.org/10.1016/B978-0-12-803374-6.00002-0>. URL: <https://www.sciencedirect.com/science/article/pii/B9780128033746000020>.
- [7] Peter Henderson et al. "Towards the Systematic Reporting of the Energy and Carbon Footprints of Machine Learning". In: *J. Mach. Learn. Res.* 21.1 (Jan. 2020). ISSN: 1532-4435.
- [8] Sasha Luccioni et al. "Quantifying the Carbon Emissions of Machine Learning". In: (2019). URL: <https://www.climatechange.ai/papers/neurips2019/22>.
- [9] UK Office for National Statistics. *2020 UK greenhouse gas emissions*. 2021. URL: <https://www.gov.uk/government/statistics/provisional-uk-greenhouse-gas-emissions-national-statistics-2020>.
- [10] Hannah Ritchie. "What are the safest and cleanest sources of energy". In: *Our World Data*. <https://ourworldindata.org/safest-sources-of-energy>. Accessed 22 (2021).
- [11] Martin Schultz et al. "Can deep learning beat numerical weather prediction?" In: *Philosophical Transactions of The Royal Society A* 379 (Feb. 2021). DOI: <https://doi.org/10.1098/rsta.2020.0097>.
- [12] Xingjian Shi et al. "Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting". In: *Advances in Neural Information Processing Systems*. Ed. by C. Cortes et al. Vol. 28. Curran Associates, Inc., 2015. URL: <https://proceedings.neurips.cc/paper/2015/file/07563a3fe3bbe7e3ba84431ad9d055af-Paper.pdf>.