Enhancing Solar Irradiance Forecasts with On-Device Continual Learning

Mateusz Piechocki[1*], Marek Kraft[1]

[1] Institute of Robotics and Machine Intelligence, Poznan University of Technology, 60-965 Poznań, Poland

*mateusz.piechocki@put.poznan.pl

Problem Statement

With the increasing contribution of solar energy to the overall renewable energy system, accurate solar irradiance forecasting is crucial for optimizing energy production and managing grid stability. However, solar irradiance is characterized by high, difficult-to-predict fluctuations related to short-term meteorological events such as cloud coverage or movements [1]. Traditional forecasting approaches rely on static and centralized algorithms that often struggle to adapt to local conditions and rapidly changing atmospheric phenomena. These limitations can lead to less stable and reliable forecasts, potentially undermining the consistency and efficiency of solar energy systems.

Main Objectives

1. Accurate identification of the current irradiance level.
2. Rapid local adaptation to evolving environmental conditions.
3. Independent and autonomous improvements of forecast accuracy after model deployment.
4. Energy-efficient, low-latency method with near real-time processing on the edge device.

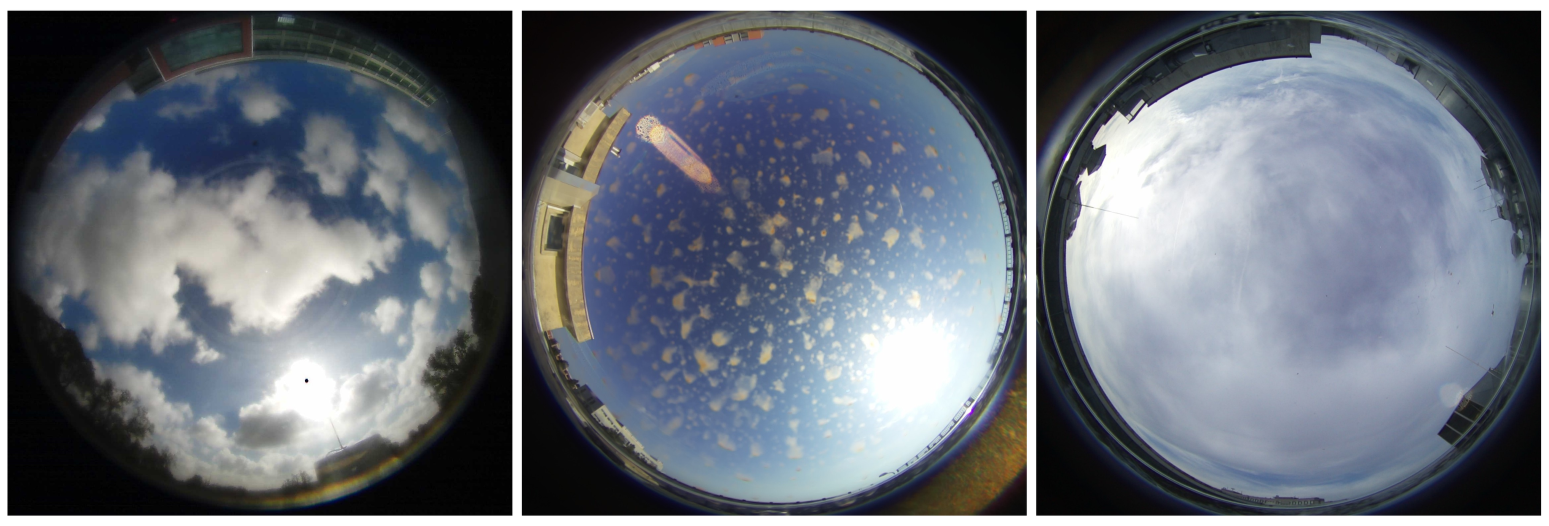

Figure 1: Sky images collected in different locations: (left) Folsom, US; (center) Modena, Italy; (right) Poznań, Poland. Images are also from various months: (left) February; (center) June; (right) October. Beyond sky observations and dynamic atmospheric phenomena, sometimes images may contain distortions, like Sahara dust in Modena's example.

Methodology

To improve the stability and reliability of irradiance forecasts, we propose an approach based on on-device continual learning to improve solar irradiance forecasting adaptability after model deployment. The initial irradiance model was based on ResNet-50 [2] architecture with two fully connected layers representing the forecasting algorithm head. The baseline model training followed the methodology outlined in the [3] and was developed on sky images with irradiance values from the Folsom dataset [4]. This dataset contains 1-min resolution measurements from three consecutive years gathered in California, US. In turn, the adaptive stage leverages incoming data and compares predicted irradiance with the ground-truth readings (obtained after 15 minutes) to determine the relevance of the data sample. Then, depending on the data's significance, the computed gradients are back-propagated for the entire model, only for forecaster head layers, or are reset.

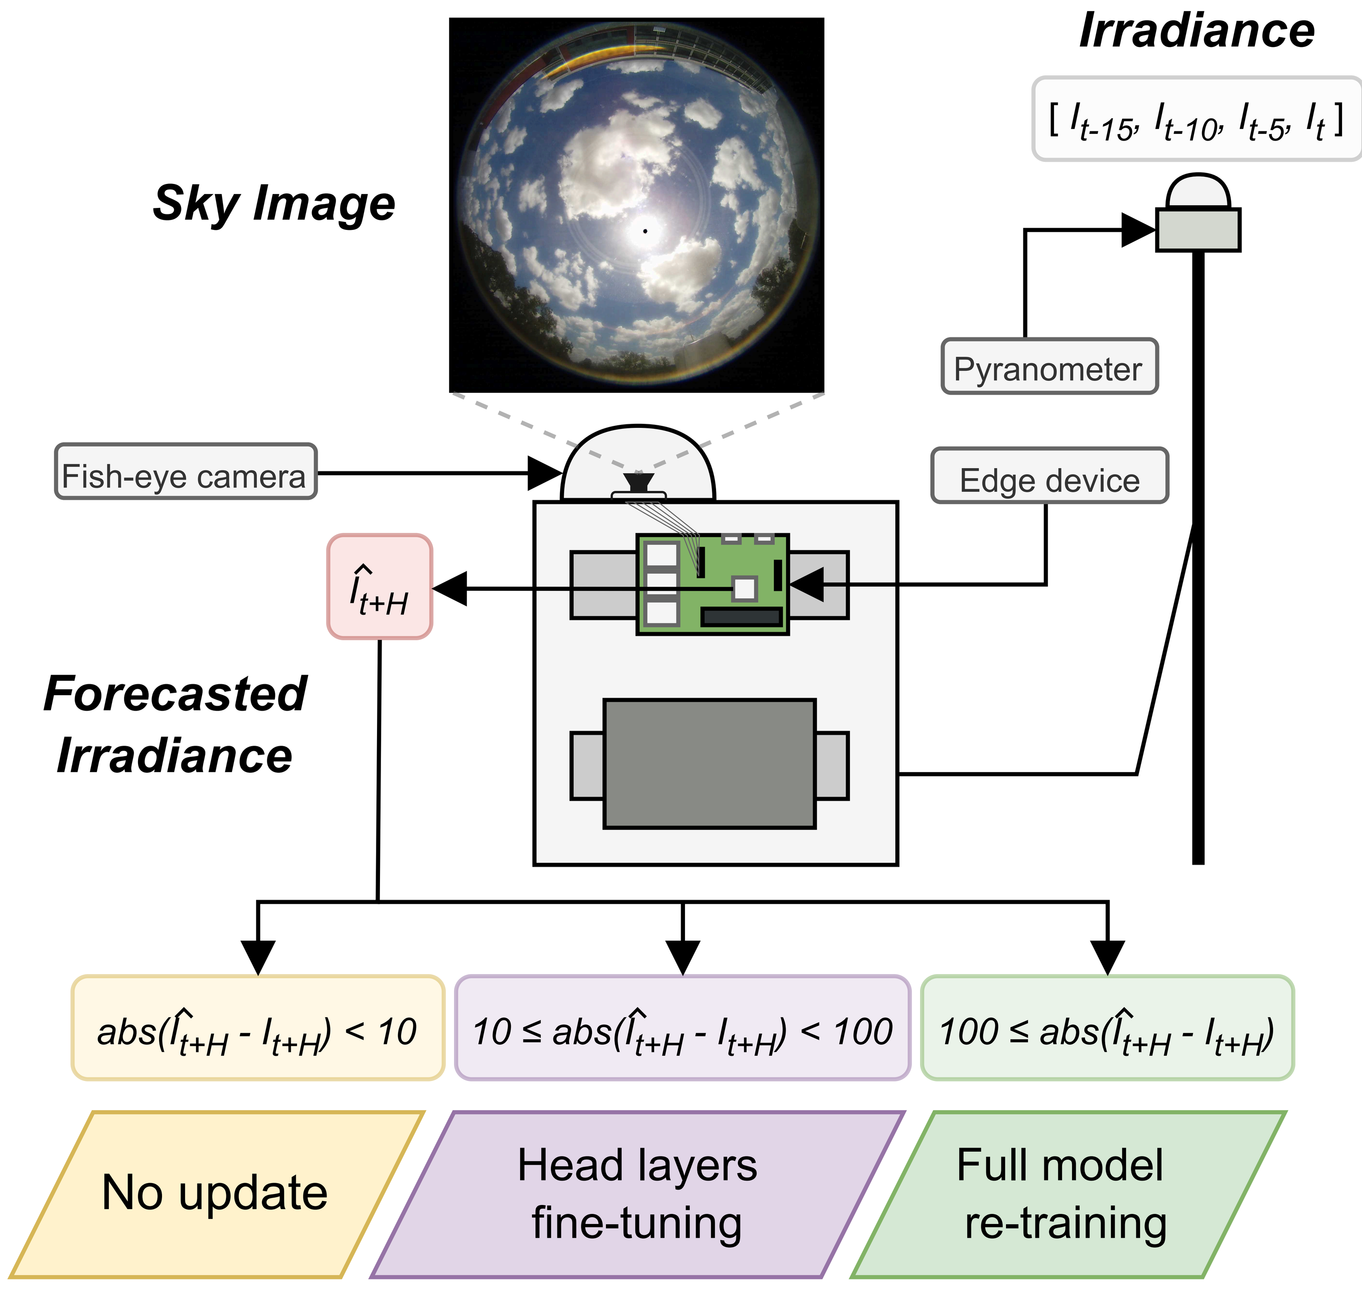

Figure 2: Illustrative diagram of proposed on-device adaptive pipeline.

Results

By utilizing new, relevant data samples our on-device continual learning pipeline can rapidly adjust to evolving environmental conditions, ensuring that the forecasting model remains accurate and responsive to local atmospheric changes. To confirm our assumptions, we selected two sub-sets, collected in different locations (Modena and Poznań) and various months (June and October, respectively). Tab. 1 shows achieved metrics, while Fig. 3 presents a visual comparison of the forecasting results of offline and adaptive methods. The presented results highlight the potential of on-device continual learning to advance solar irradiance forecasting, providing a scalable and adaptive solution to enhance energy management and facilitate more effective grid integration in the renewable energy sector.

Table 1: Metric results.

Dataset	Method	Metrics	
		MAE $[W/m^2]\downarrow$	RMSE $[W/m^2]\downarrow$
Folsom	-	34.895	70.353
Modena	Offline model	78.857	125.340
	Proposed	47.185	112.298
Poznań	Offline model	35.340	46.190
	Proposed	11.256	28.350

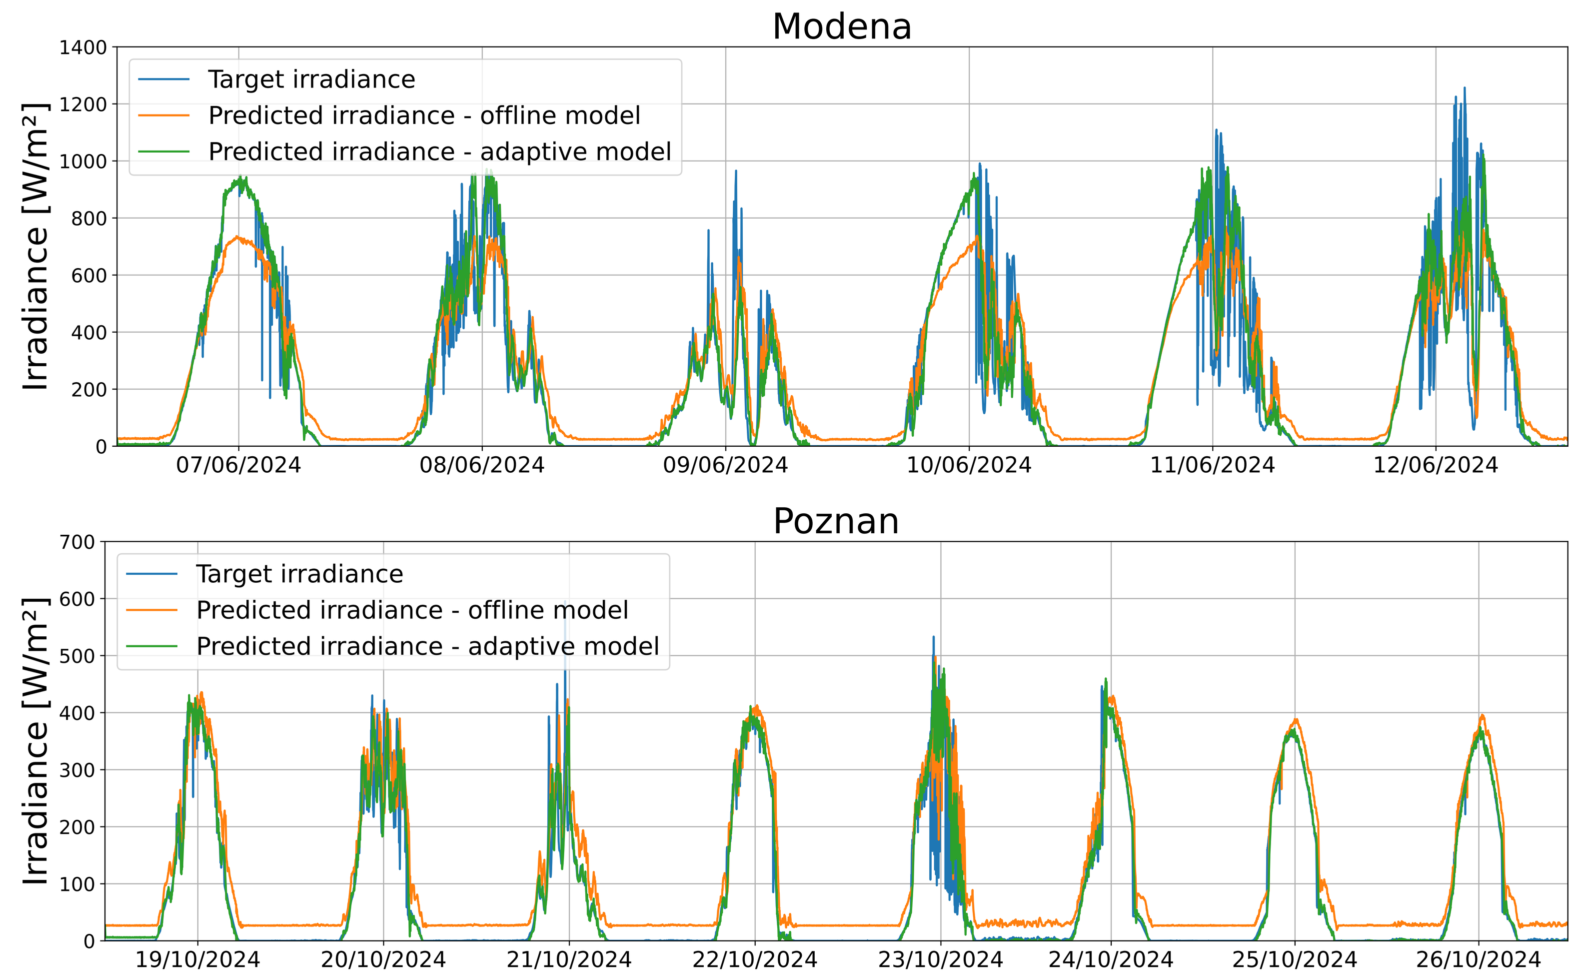

Figure 3: Visual comparison of irradiance forecasts for measurement stations located in Modena (top) and Poznań (bottom) and readings collected in June and October, respectively. Note the different irradiance ranges on the y-axis.

The proposed adaptive pipeline is intended to operate efficiently on remote resource-constrained hardware. Therefore, the entire processing loop was benchmarked on Raspberry Pi 5 and Jetson Orin Nano, and performance results, divided into individual processing options, are presented in Tab. 2.

Table 2: Performance results on selected edge devices. Presented times refer to a single data sample processing, consisting of 128×128 sky image and a vector of 4 historical irradiance values.

Device	Avg. inference time [ms]	Avg. training time (full) [ms]	Avg. training time (head) [ms]
Raspberry Pi 5	166.292 ± 9.193	526.145 ± 21.263	174.039 ± 6.737
Jetson Orin Nano	21.634 ± 0.114	79.515 ± 15.936	26.192 ± 0.161

Conclusions

- The developed pipeline improves the deployed model performance, improving intra-hour forecasts.
- Compared to the non-adaptive model, the online approach enhances adaptation to dynamically changing local weather conditions.
- Measured times demonstrate the feasibility of on-device adaptation, relying only on the available constrained resources, for edge devices like Raspberry Pi 5 or Jetson Orin Nano.

References

[1] N. Krishnan, K. R. Kumar, and C. S. Inda, "How solar radiation forecasting impacts the utilization of solar energy: A critical review," *Journal of Cleaner Production*, vol. 388, p. 135860, 2023.

[2] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.

[3] E. A. Papatheofanous, V. Kalekis, G. Venitourakis, F. Tziolos, and D. Reisis, "Deep learning-based image regression for short-term solar irradiance forecasting on the edge," *Electronics*, vol. 11, no. 22, 2022.

[4] H. T. C. Pedro, D. P. Larson, and C. F. M. Coimbra, "A comprehensive dataset for the accelerated development and benchmarking of solar forecasting methods," *Journal of Renewable and Sustainable Energy*, vol. 11, 06 2019. 036102.

Acknowledgement

This work has been supported by the Empiria and Knowledge Foundation under the "Talents of Tomorrow" program (Grant Number: 18/TJ_2/2023).