Theme(s): *Neuromorphic Computing, Next-generation AI/ML algorithms*

Energy-Efficient Solar Forecasting with a Neuromorphic-Bayesian State Estimation Approach

Kumar Anurag[1*], Francesco Sorrentino[1], Wenbin Wan[1]
[1]Department of Mechanical Engineering, University of New Mexico, Albuquerque, NM, USA
Corresponding author: Kumar Anurag[*] (kmranrg@unm.edu)

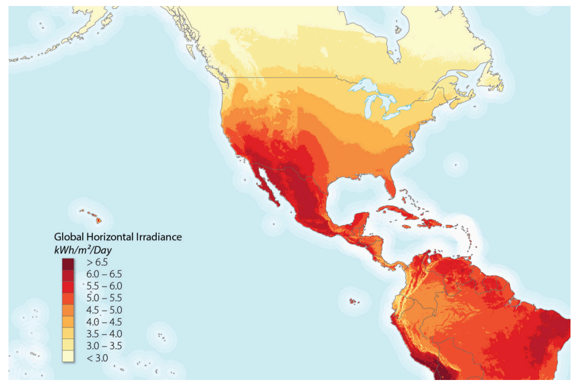

Figure 1: Global Horizontal Irradiance (GHI) from the National Solar Radiation Database (NSRDB).

Overview. The stability of renewable energy grids critically depends on accurate solar irradiance forecasting. While deep learning models offer high performance, their energy-intensive training process poses a significant barrier to large-scale, adaptive deployment. We propose a brain-inspired, energy-efficient approach that integrates reservoir computing (RC) with an unscented Kalman filtering (UKF) [1]. The RC component, which avoids costly backpropagation, serves as a fast-to-train dynamics model. The UKF provides a principled Bayesian method for correcting the model's predictions in real-time using new measurements. Tested on over two decades of solar data from the National Solar Radiation Database (NSRDB) [2], our RCUKF framework achieves superior performance for 30-minute ahead forecasting, outperforming both standalone RC and traditional neural network baselines while requiring significantly less training time. This approach demonstrates a viable pathway for developing low-power, high-performance AI systems for smart grid management.

Background and Motivation. The increasing penetration of photovoltaics into the energy grid presents a major challenge to stability due to the inherent volatility of solar power. Accurate short-term forecasting of Global Horizontal Irradiance (GHI) is therefore essential for reliable grid operation, energy trading, and asset management. The current state-of-the-art often relies on recurrent neural networks (RNNs), such as LSTMs. However, the primary drawback of these models is their immense computational and energy cost for training, which stems from the iterative backpropagation-through-time (BPTT) algorithm. This high training overhead makes frequent model updates or on-device adaptation impractical, creating a critical need for novel algorithms that deliver high accuracy with low energy consumption.

Approach. To address this challenge, we implement a hybrid Neuromorphic-Bayesian approach, RCUKF, which synergistically integrates two powerful components:

1. Reservoir Computer (RC): A brain-inspired RNN where the recurrent hidden layer has fixed, random weights. Only a final linear 'readout' layer is trained. This design replaces computationally expensive gradient descent with a one-shot linear regression solution (ridge regression), drastically reducing the energy budget for model training. The RC learns the underlying system dynamics from historical data and serves as the data-driven process model within the filter.

2. Unscented Kalman Filter (UKF): A robust Bayesian filter for non-linear systems. The trained RC is embedded as the state transition function in the UKF's prediction step. The UKF's measurement

update step then uses incoming ground-truth data to correct the RC's state prediction in real-time. This provides exceptional robustness to noise and system drift without requiring costly retraining.

This two-phase architecture—a computationally efficient training phase followed by a lightweight online correction phase—is ideally suited for energy-constrained forecasting applications.

Results. The RCUKF framework was evaluated on the task of 30-minute ahead GHI forecasting. We used a comprehensive dataset from the National Solar Radiation Database for a single site in the central US, spanning from 1998 to 2024, with a 70/30 chronological train/test split. We compared its performance against a standard RC and a standard multi-layer perceptron (MLP) baseline that used a 24-hour lag window as input.

- Accuracy: As shown in Fig. 2, RCUKF achieves near-perfect tracking of the ground truth. It obtained a test Mean Absolute Error (MAE) of just 4.3 W/m². This dramatically outperformed both the standard RC (MAE of 56.6 W/m²) and the MLP (MAE of 37.4 W/m²).
- Training Efficiency: The RCUKF model's training phase was completed in only 6.6 seconds, which is almost five times faster than the MLP baseline (35.1 seconds), highlighting the significant energy savings from its backpropagation-free learning process.

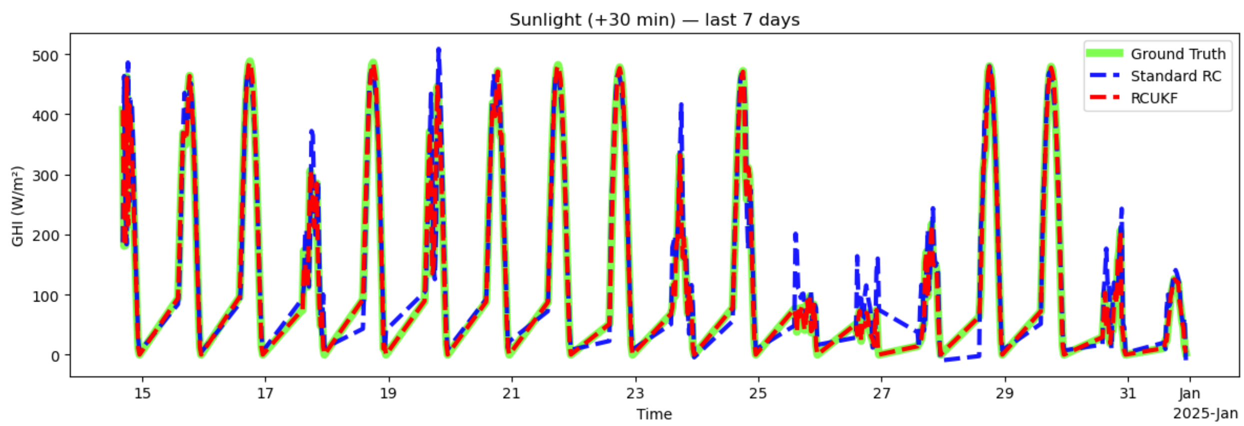

Figure 2: Comparison of 30-minute ahead GHI forecasts on a 7-day test window.

Conclusion. We have successfully presented a Neuromorphic-Bayesian approach that is both highly accurate and computationally efficient for solar forecasting. The key finding is that augmenting a fast-to-train Reservoir Computer with a real-time UKF correction loop yields performance superior to that of traditional, more energy-intensive models. The energy consequence of this work is a significant reduction in the computational barrier for deploying adaptive, high-performance AI in the renewable energy sector. This approach enables the development of practical, large-scale forecasting systems and opens pathways for low-power, on-device model adaptation in smart grid management, contributing to a more efficient and stable energy future.

References

[1] Kumar Anurag, Kasra Azizi, Francesco Sorrentino, and Wenbin Wan. RCUKF: Data-driven modeling meets Bayesian estimation. *The Modeling, Estimation and Control Conference (MECC)*, 2025 (To appear).

[2] Manajit Sengupta, Yu Xie, Anthony Lopez, Aron Habte, Galen Maclaurin, and James Shelby. The National Solar Radiation Data Base (NSRDB). *Renewable and Sustainable Energy Reviews*, 89:51–60, 2018.