Multi-Horizon Short-Term Solar Power Forecasting Using GAN and Neural Network Models

Yen-Hsi Chou, Arundhuti Haldar, Anderson Rodrigo de Queiroz
North Carolina State University, Raleigh, NC, USA
{ychou4, ahaldar, ardequei}@ncsu.edu

***Abstract*—Accurate solar power forecasting is essential for efficient power system planning and operations. This study develops a framework for multi-horizon short-term solar power forecasting that integrates spatiotemporal irradiance prediction and probabilistic modeling. Within this framework, a generative adversarial network (GAN) synthetically generates future global horizontal irradiance maps from past observations, capturing spatiotemporal patterns that influence solar generation. The generated irradiance features are then combined with historical generation and meteorological data to produce quantile-based forecasts using different neural networks. The framework is validated through a case study in California using high-resolution NSRDB and PVOutput datasets for 10-, 30-, and 60-minute horizons. The results show that incorporating GAN-generated irradiance features improves both deterministic and probabilistic forecasts, with notable gains at 30- and 60-minute horizons. The proposed framework offers a practical approach for multi-horizon solar power forecasting and can be extended to other data-driven renewable energy applications.**

***Index Terms*—Short-term solar power forecasting, Time series analysis, Generative adversarial network (GAN), Transformer**

I. Introduction

Accurate solar power forecasting is critical for grid reliability when managing the supply-demand balance in modern power grids [1]. As solar photovoltaic (PV) penetration increases, better short-term forecasts help operators mitigate variability, optimize reserves, and improve market participation. In the short-term horizon, typically from 30 to 360 minutes, solar PV power prediction is particularly challenging because of rapid fluctuations in solar irradiance, local cloud movements, and spatial variability across distributed installations [2]. Recent reviews classify solar PV forecasting methods into three main categories: physical, statistical, and machine-learning approaches [3]. Among these, machine-learning models have demonstrated remarkable capability in capturing complex nonlinear temporal and spatial patterns from diverse data sources, and they represent the most prominent and rapidly advancing direction in the literature.

More recently, generative adversarial networks (GANs) have shown potential in capturing spatial and temporal patterns of irradiance fields. For instance, [4] used a GAN to predict future solar irradiance maps and converted them into power estimates using irradiance-to-power relationships that depend on system parameters such as tilt angle, azimuth, and capacity. However, such detailed information is rarely available for large numbers of small distributed systems. Motivated by this limitation, the present study proposes a GAN-assisted short-term forecasting framework that generates future global horizontal irradiance (GHI) maps and integrates them as auxiliary features within neural-network-based power forecasters. The framework is evaluated at 10-, 30-, and 60-minute horizons to assess how GAN-generated GHI features improve both deterministic and probabilistic forecasts.

II. Methodology

The study develops a short-term forecasting framework as illustrated in Fig. 1. The proposed framework consists of two main components: (1) data preprocessing, which constructs regional GHI maps and integrates spatiotemporal information from multiple PV systems; and (2) forecasting, which combines historical observations with predicted features to generate quantile-based forecasts. The following subsections describe each stage of the methodology in detail.

A. Data Preprocessing

1) Regional GHI Map Construction: The regional GHI maps were constructed using GHI data obtained from the National Solar Radiation Database (NSRDB) [5]. The NSRDB provides regularly gridded data with a spatial resolution of 2 km and a temporal resolution of 10 minutes.

To capture finer spatial variations in solar irradiance and better represent the spatial distribution of irradiance intensity, Kriging interpolation was applied to the original 9×9 NSRDB grids. It exploits spatial correlation in nearby observations to estimate irradiance on a finer grid.

Fig. 2 compares interpolation results at different spatial resolutions. Resolutions finer than 250 m yield only marginal visual gains while substantially increasing computational cost, whereas coarser settings such as 500 m over-smooth irradiance gradients, which may hinder learning. Considering both image quality and efficiency, a target resolution of 250 m was selected. The interpolated 65×65 maps were subsequently cropped to 64×64 as inputs to the GAN model.

2) Feature Integration for Each PV System: PV generation data were collected from PVOutput.org [6], and site-level meteorological variables were sourced from the nearest NSRDB grid points. The Pearson correlation coefficient (PCC) was computed to quantify the relationships between each candidate temporal or meteorological variable and PV output. Based on the approach in [7], variables with an absolute correlation of at least 0.3 were kept as candidates for subsequent selection.

979-8-3315-8138-1

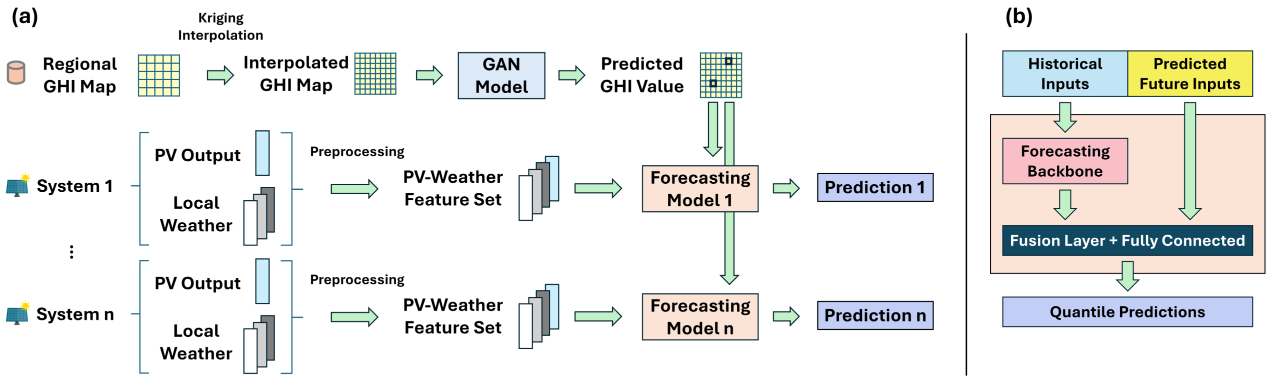

Fig. 1: (a) Overview of the proposed multi-horizon short-term forecasting framework. (b) Structure of the forecasting model.

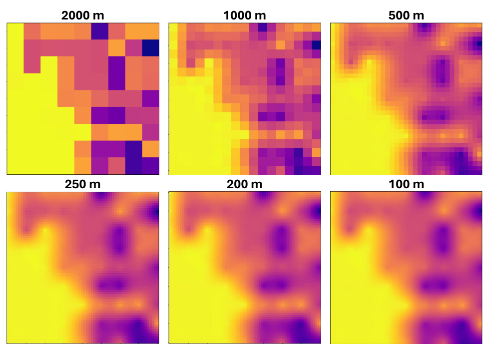

Fig. 2: Effect of spatial resolution on interpolated GHI maps

TABLE I: Network configurations of the Generator (G) and Discriminator (D)

Block	Input	#Conv / Channels	Kernel Sizes
G1	4 × 4	3 / [128, 256, 128]	[3, 3, 3]
G2	8 × 8	3 / [128, 256, 128]	[5, 3, 5]
G3	16 × 16	5 / [128, 256, 512, 256, 128]	[5, 3, 3, 3, 5]
G4	32 × 32	5 / [128, 256, 512, 256, 128]	[7, 3, 3, 3, 7]
G5	64 × 64	6 / [128, 256, 512, 512, 256, 128]	[7, 5, 5, 3, 3, 3]
D1	4 × 4	1 / [64]	[3]
D2	8 × 8	3 / [64, 128, 128]	[3, 3, 3]
D3	16 × 16	3 / [128, 256, 256]	[5, 5, 5]
D4	32 × 32	4 / [128, 256, 512, 128]	[7, 7, 5, 5]
D5	64 × 64	5 / [128, 256, 512, 512, 256]	[7, 5, 5, 3, 3]

B. Neural Network Models

1) GAN for Spatiotemporal GHI Prediction: A hierarchical multi-scale GAN was developed to model the spatiotemporal evolution of GHI maps. The model takes a sequence of five past irradiance frames as input and predicts one, three, or six future frames corresponding to 10-, 30-, and 60-minute forecasting horizons.

The generator comprises five sub-modules (G1–G5) operating at spatial resolutions of 4×4, 8×8, 16×16, 32×32, and 64×64 pixels. Each sub-generator receives the upsampled output from the previous scale together with interpolated historical inputs, enabling the network to learn irradiance evolution patterns from large-scale contextual motion to fine-grained spatial detail. To mitigate boundary artifacts, reflective padding is applied before convolution, and each generator block explicitly accounts for the total spatial reduction introduced by its convolutional layers. This approach preserves input–output alignment and reduces edge distortions, mitigating padding artifacts commonly seen in convolutional networks [8]. The discriminator adopts a matching multi-scale design that evaluates realism at corresponding resolutions and provides adversarial feedback that promotes spatial coherence and temporal consistency across scales.

Training employs a composite objective function. The generator combines an adversarial term that encourages realism, an L_1 reconstruction term that preserves numerical fidelity, and a gradient consistency term that promotes smooth transitions while preserving edges. The discriminator is optimized with binary cross-entropy loss.

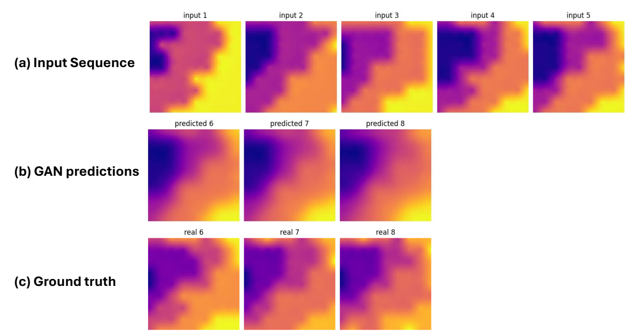

Fig. 3: Illustration of GHI Map forecasting using the proposed multi-scale GAN. (a) Input sequence of five consecutive irradiance maps. (b) Predicted maps generated by the GAN for the next three time steps. (c) Corresponding ground truth maps for comparison.

This design captures multiscale irradiance dynamics and reduces edge artifacts, producing high-resolution GHI maps that serve as future covariates for power forecasting. A detailed configuration of each generator block is summarized in Table I. An illustrative example of the GAN-generated irradiance predictions is shown in Fig. 3.

979-8-3315-8138-1

2) Quantile Forecasting Models with Historical and Predicted Features: A quantile-based formulation was adopted to represent predictive uncertainty in PV generation, consistent with recent probabilistic forecasting studies [9], [10]. This enables interval-oriented evaluation and supports operational decisions beyond point estimates.

Historical inputs, including past power generation and site-level meteorological variables, were encoded using one of several neural forecasting backbones, including Gated Recurrent Unit (GRU), Long Short-Term Memory (LSTM), Bidirectional Long Short-Term Memory (Bi-LSTM), or a Transformer model. These architectures have demonstrated strong capability in capturing nonlinear relationships and complex temporal dynamics in energy time series [11].

Building on these backbones, a fusion-based quantile forecasting framework was developed to jointly utilize historical inputs and predicted future features derived from GAN-generated GHI maps. As illustrated in Fig. 1(b), the encoded historical representation is integrated with projected future covariates and merged through a fusion layer, followed by a fully connected head that outputs three quantile forecasts at the 0.1, 0.5, and 0.9 levels. This design takes advantage of previous dependencies and anticipated conditions to improve short-term accuracy and uncertainty quantification.

C. Evaluation Metrics

Deterministic metrics, including MAE, RMSE, and R^2, were used to evaluate point forecast accuracy [12]. Probabilistic performance was assessed using the average pinball loss (APL) over the 0.1, 0.5, and 0.9 quantiles, and PICP and MPIW. Relative improvement was quantified by a skill score for RMSE and APL [13]

$$\text{Skill Score} = \left(1 - \text{Error}_{\text{model}}/\text{Error}_{\text{baseline}}\right) \times 100\%.$$

III. Case Study

A. Data Description

The study area in California was selected from PVOutput for its density of distributed PV systems. Three representative systems were analyzed as shown in Fig. 4, each providing five-minute PV generation records from 2020 to 2023. Regional GHI maps and site-level meteorological variables were obtained from the NSRDB. Data from 2018–2019 were used exclusively to train the GAN, which was then applied to generate inputs for 2020–2023, while the forecasting models were trained and validated on 2020–2022 and tested on 2023.

B. Data Preparation

A fixed daytime window of 07:20–16:30 was used for all days to ensure temporal consistency, with daylight-saving shifts corrected to maintain continuity. In line with [12], the dataset was split chronologically, with 2023 reserved for testing and 2020–2022 used for training and validation. Two input configurations were considered for each forecast horizon: historical inputs only, and historical inputs augmented with GAN-generated GHI. PV generation was scaled using Min–Max normalization and meteorological variables were standardized. Building on [14], the candidate features filtered by PCC were added iteratively and MAE was monitored to select season-specific inputs.

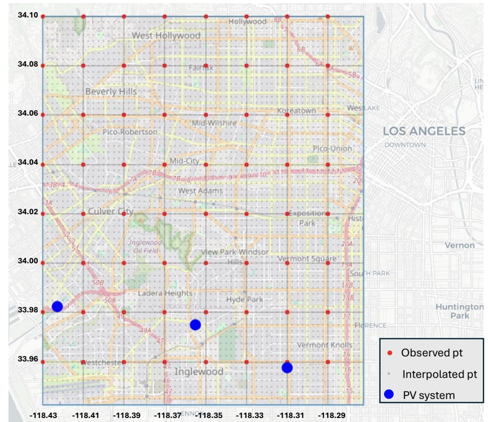

Fig. 4: Overview of study area

TABLE II: Average inference latency of each component in the proposed forecasting pipeline (excluding training). K, G, and F denote Kriging interpolation, GAN inference, and the forecasting model, respectively.

Component	Latency	Notes
Kriging interpolation (per map)	140 ms	CPU
GAN inference (per map)	~ 2 ms	GPU, bs = 32
Forecasting model (per sample)	< 1 ms	GPU, bs = 16
End-to-end (first prediction of the day)	703 ms	5K+G+F
End-to-end (subsequent predictions)	143 ms	K+G+F

C. Experimental Settings

All models were implemented in PyTorch and trained on an NVIDIA RTX A6000. The multi-scale GAN was trained using stochastic gradient descent with a momentum of 0.9, cosine annealing for learning-rate scheduling, and gradient clipping. The generator and discriminator losses followed the definitions provided in Section 2. Forecasting models were trained using the Adam optimizer. A grid search was conducted to determine the optimal network architecture and training hyperparameters, and early stopping was applied based on validation performance. Each model used an input sequence of five consecutive 10-minute intervals, and the testing coverage differed slightly across horizons due to the fixed daytime window. Component-wise and end-to-end inference latency for online forecasting is also reported in Table II.

D. Experimental results

The experimental analysis evaluates the effectiveness of incorporating GAN-generated GHI maps as auxiliary predictive information for short-term solar power forecasting across three PV systems at 10-, 30-, and 60-minute horizons. Models trained without GAN-generated GHI maps are treated as baselines, enabling a direct assessment of the proposed framework. In addition to recurrent neural network baselines,

979-8-3315-8138-1

TABLE III: Comparison of forecasting model performance for 10-min ahead

System	Model	MAE	RMSE	R^2	$Skill_{RMSE}$
	BiLSTM	14.71	29.83	0.9731	–
1	CNN-LSTM	15.08	29.86	0.9731	-0.10%
	BiLSTM(GAN)	14.69	29.68	0.9734	0.50%
	GRU	18.24	38.69	0.8944	–
2	CNN-LSTM	18.41	39.00	0.8927	-0.80%
	GRU(GAN)	18.60	38.61	0.8948	0.21%
	BiLSTM	17.03	38.05	0.9678	–
3	CNN-LSTM	17.30	38.69	0.9667	-1.68%
	BiLSTM(GAN)	17.16	38.46	0.9671	-1.11%
System	**Model**	**APL**	**PICP**	**MPIW**	$\mathbf{Skill_{APL}}$
	BiLSTM	4.98	0.8284	44.62	–
1	CNN-LSTM	5.07	0.8407	48.61	-1.81%
	BiLSTM(GAN)	4.98	0.7946	43.05	-0.05%
	GRU	6.11	0.8028	53.12	–
2	CNN-LSTM	6.15	0.8453	57.78	-0.65%
	GRU(GAN)	6.24	0.7885	54.54	-2.08%
	BiLSTM	5.75	0.8201	51.49	–
3	CNN-LSTM	5.82	0.8330	52.41	-1.22%
	BiLSTM(GAN)	5.83	0.8175	50.43	-1.39%

TABLE IV: Comparison of forecasting model performance for 30-min ahead

System	Model	MAE	RMSE	R^2	$Skill_{RMSE}$
	LSTM	31.58	58.37	0.8955	–
1	CNN-LSTM	32.52	59.71	0.8907	-2.30%
	LSTM(GAN)	31.05	56.12	0.9034	3.85%
	GRU	26.60	51.59	0.8106	–
2	CNN-LSTM	26.60	51.88	0.8083	-0.56%
	GRU(GAN)	26.08	49.25	0.8273	4.54%
	LSTM	36.44	66.78	0.9015	–
3	CNN-LSTM	35.89	65.32	0.9058	2.19%
	LSTM(GAN)	36.34	65.36	0.9057	2.13%
System	**Model**	**APL**	**PICP**	**MPIW**	$\mathbf{Skill_{APL}}$
	LSTM	10.38	0.7820	91.41	–
1	CNN-LSTM	10.76	0.8079	97.71	-3.66%
	LSTM(GAN)	10.34	0.7593	87.39	0.39%
	GRU	8.79	0.8032	76.36	–
2	CNN-LSTM	8.78	0.7935	78.64	0.11%
	GRU(GAN)	8.66	0.8119	79.96	1.48%
	LSTM	12.09	0.7373	103.38	–
3	CNN-LSTM	11.80	0.7445	105.54	2.40%
	LSTM(GAN)	12.02	0.7341	100.16	0.58%

TABLE V: Comparison of forecasting model performance for 60-min ahead

System	Model	MAE	RMSE	R^2	$Skill_{RMSE}$
	BiLSTM	44.37	80.21	0.8014	–
1	CNN-LSTM	44.73	79.68	0.8040	0.66%
	BiLSTM(GAN)	41.46	73.18	0.8347	8.76%
	GRU	32.78	60.65	0.7391	–
2	CNN-LSTM	32.90	61.66	0.7303	-1.67%
	GRU(GAN)	31.70	57.63	0.7644	4.98%
	LSTM	47.33	83.25	0.8498	–
3	CNN-LSTM	46.28	78.84	0.8653	5.30%
	LSTM(GAN)	45.75	79.11	0.8644	4.97%
System	**Model**	**APL**	**PICP**	**MPIW**	$\mathbf{Skill_{APL}}$
	BiLSTM	14.65	0.7438	111.78	–
1	CNN-LSTM	14.39	0.8301	131.04	1.77%
	BiLSTM(GAN)	13.88	0.7305	111.17	5.26%
	GRU	10.72	0.7411	93.34	–
2	CNN-LSTM	10.60	0.7899	92.13	1.12%
	GRU(GAN)	10.49	0.7401	91.79	2.15%
	LSTM	15.92	0.6968	123.21	–
3	CNN-LSTM	15.05	0.7939	136.68	5.47%
	LSTM(GAN)	15.58	0.6952	122.60	2.14%

a convolutional neural network (CNN)–LSTM hybrid model, which has been recently adopted for capturing spatiotemporal information through convolutional processing and subsequent sequence modeling in solar forecasting studies [15], is included for comparison. Here, interpolated GHI maps are first processed by the CNN module, and the resulting representations are fused with other historical inputs before being passed to the LSTM for forecasting. For each horizon and system, the best-performing baseline model is compared with its GAN-assisted counterpart, and the results are summarized using both deterministic and probabilistic metrics. Specifically, RMSE-based skill scores and APL-based skill scores are reported to quantify improvements in point and quantile forecasts, respectively. A qualitative comparison across different horizons and seasons is further illustrated in Fig. 5.

1) 10-Minute-Ahead Forecasting Performance: Table III summarizes the 10-minute-ahead results, showing that incorporating GAN-generated GHI maps leads to only marginal changes in forecasting performance. Across the three systems, the CNN-LSTM model does not exhibit clear advantages over the best-performing recurrent baseline, indicating that at very short lead times the recent temporal continuity already dominates the predictive signal. For System 1, BiLSTM(GAN) achieves a slight reduction in MAE and RMSE (skill score of 0.50%), while Systems 2 and 3 show only negligible changes with mixed gains depending on the system. The probabilistic metrics (APL, PICP, and MPIW) exhibit similarly small variations, suggesting that neither convolutional processing nor GAN-derived spatial information substantially alters the predictive distribution at this horizon.

2) 30-Minute-Ahead Forecasting Performance: Table IV reports consistent improvements at the 30-minute horizon, where the benefit of incorporating GAN-generated GHI features becomes more evident. Across the three systems, GAN-assisted models reduce RMSE with skill gains of approximately 2–5%, indicating that the predicted irradiance maps provide useful spatial context as the horizon extends. Compared with the CNN-LSTM model, which yields mixed results depending on the system, the GAN-assisted framework provides more consistent improvements overall, particularly in the probabilistic metrics where lower APL is often achieved without widening the prediction intervals. The probabilistic results further confirm this trend, where the GAN-assisted models generally exhibit lower APL and comparable or slightly reduced MPIW, reflecting more confident yet accurate interval predictions.

3) 60-Minute-Ahead Forecasting Performance: At the 60-minute horizon, Table V shows that the performance gap widens further. While the CNN-LSTM model provides moderate gains for some systems, the GAN-assisted models achieve substantially larger improvements, with RMSE-based skill gains reaching 4–9%. Both deterministic (MAE, RMSE, R^2) and probabilistic (APL) metrics show consistent enhancement,

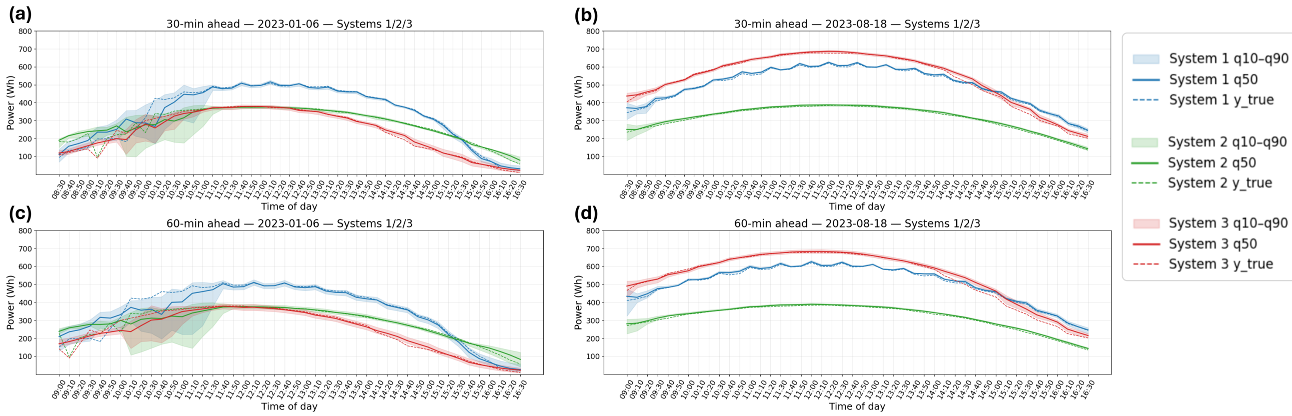

Fig. 5: Qualitative examples of forecasting performance across horizons and seasons for three PV systems. (a) 30-min ahead on 2023-01-06; (b) 30-min ahead on 2023-08-18; (c) 60-min ahead on 2023-01-06; (d) 60-min ahead on 2023-08-18. Each panel shows predicted (solid) and observed (dashed) power for three systems; all panels share the same vertical scale.

especially for System 1, where the RMSE decreased by nearly 9%. These results suggest that as the horizon increases, it becomes increasingly valuable to incorporate forward-looking spatial irradiance information. Unlike CNN-LSTM, which extracts representations from historical interpolated maps, the proposed framework leverages GAN-predicted future GHI maps that encode evolving cloud-driven irradiance transitions, leading to more accurate and better-calibrated forecasts at longer lead times.

Overall, the results indicate a clear relationship between forecasting horizon and the effectiveness of incorporating GAN-generated GHI maps. While the improvement is negligible at 10 minutes ahead, the benefit becomes increasingly apparent at 30 and 60 minutes. Although CNN-LSTM provides a competitive hybrid baseline by exploiting historical spatiotemporal patterns, the proposed GAN-assisted framework yields more consistent and larger gains at longer horizons by integrating forward-looking irradiance information derived from predicted GHI maps. This suggests that the spatially enhanced irradiance information produced by the GAN model provides a valuable context for longer-term forecasts, where temporal features alone become less informative. The findings highlight the potential of combining generative spatial features with sequential models to enhance the robustness of short-term solar forecasting.

IV. Conclusion

This study developed and evaluated a hybrid forecasting framework that integrates generative irradiance features with deep learning models for solar power prediction. Experimental results across three systems and multiple forecasting horizons demonstrate that the proposed approach improves accuracy, particularly for the 30-minute and 60-minute horizons, by leveraging the spatiotemporal patterns captured by the GAN-generated GHI maps. The findings highlight the importance of spatial information for extending predictive lead times and provide practical guidance for the deployment of data-driven forecasting systems in distributed PV networks. While this framework was validated in California's high-irradiance climate, its integration of GAN-based spatial features offers a scalable foundation for capturing cloud dynamics across different micro-climates. Future work will focus on adapting the GAN component with region-specific irradiance data to evaluate the framework's generalizability and robustness under more variable and volatile climatic conditions.

References

[1] H.-O. Pörtner *et al.*, Eds., *Climate Change 2022: Impacts, Adaptation, and Vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, 2022.

[2] R. Ahmed *et al.*, "A review and evaluation of the state-of-the-art in pv solar power forecasting: Techniques and optimization," *Renewable and Sustainable Energy Reviews*, vol. 124, p. 109792, 2020.

[3] M. Gupta *et al.*, "A review of pv power forecasting using machine learning techniques," *Progress in Engineering Science*, p. 100058, 2025.

[4] H. Wen *et al.*, "A regional solar forecasting approach using generative adversarial networks with solar irradiance maps," *Renewable Energy*, vol. 216, p. 119043, 2023.

[5] M. Sengupta *et al.*, "The national solar radiation data base (nsrdb)," *Renewable and Sustainable Energy Reviews*, vol. 89, pp. 51–60, 2018.

[6] "PVOutput.org," https://pvoutput.org/.

[7] E. Kim *et al.*, "Designing solar power generation output forecasting methods using time series algorithms," *Electric Power Systems Research*, vol. 216, p. 109073, 2023.

[8] B. Alsallakh *et al.*, "Mind the pad – cnns can develop blind spots," 2020. [Online]. Available: https://arxiv.org/abs/2010.02178

[9] J. Hu *et al.*, "Probabilistic net load forecasting based on transformer network and gaussian process-enabled residual modeling learning method," *Renewable Energy*, vol. 225, p. 120253, 2024.

[10] J. Hu *et al.*, "A novel wind power interval prediction method based on neural ensemble search and dynamic conformalized quantile regression," *Applied Soft Computing*, p. 113476, 2025.

[11] Y.-H. Chou *et al.*, "Day-ahead solar power forecasting in a large-scale system using statistical and neural network models," in *IEEE Power & Energy Society General Meeting (PESGM)*, 2025.

[12] Q. Paletta *et al.*, "Advances in solar forecasting: Computer vision with deep learning," *Advances in Applied Energy*, vol. 11, p. 100150, 2023.

[13] D. Yang *et al.*, "Verification of deterministic solar forecasts," *Solar Energy*, vol. 210, pp. 20–37, 2020, special Issue on Grid Integration.

[14] M. Abuella and B. Chowdhury, "Solar power forecasting using artificial neural networks," in *2015 North American Power Symposium*, 2015.

[15] A. Agga *et al.*, "Cnn-lstm: An efficient hybrid deep learning architecture for predicting short-term photovoltaic power production," *Electric Power Systems Research*, vol. 208, p. 107908, 2022.

979-8-3315-8138-1