Two-Stage Prediction Intervals via Spiking Neural Network-Gated State-Specific Residual Models for Solar Power Generation

Thomas Aguilera
Department of Electrical Engineering
University of Chile
Santiago, Chile
thomas.aguilera@ug.uchile.cl

Oscar Cartagena
Department of Electrical Engineering
University of Chile
Santiago, Chile
oscar.cartagena@uchile.cl

Alex Navas-Fonseca
Department of Electrical Engineering
University of Chile
Santiago, Chile
alex.navas@uchile.cl

Alfredo Núñez
Faculty of Civil Engineering and Geosciences
Delft University of Technology
Delft, The Netherlands
A.A.NunezVicencio@tudelft.nl

Doris Sáez
Department of Electrical Engineering
University of Chile and
Instituto Sistemas Complejos de Ingeniería
Santiago, Chile
dsaez@ing.uchile.cl

***Abstract*—A two-stage framework has been developed for reliability-aware day-ahead forecasting. This framework includes a forecasting backbone that generates multi-step point predictions. From these outputs, an aligned residual is defined at a selected horizon and used to identify significant forecast error events. This identification is accomplished through a validation-calibrated threshold that ensures the desired prevalence of these events. In the proposal, an event detector based on a surrogate-gradient spiking neural network (SNN), implemented in snnTorch, estimates the probability of entering a significant error state. To achieve that, it uses features available at the time of the forecast, as the predicted trajectory shape and various calendar descriptors, and combines them with a validation-calibrated operating threshold that converts the detector's probability into a decision regarding the state. The resulting method effectively combines SNN-based state classification with state-dependent to construct prediction intervals for asymmetric residual modeling. The experimental results demonstrate strong performance in event detection and provide practical trade-offs between coverage and width for prediction intervals, supporting the framework's suitability for reliability-aware energy forecasting in microgrids.**

***Index Terms*—Spiking neural networks, surrogate-gradient learning, snnTorch, microgrids, forecast error events, prediction intervals, renewable energy modeling**

I. Introduction

Uncertainty quantification is essential in energy time-series forecasting, including residential demand and photovoltaic (PV) generation [1]. Point forecasts alone are insufficient for risk-aware operation, since occasional large errors can trigger constraint violations, false alarms, or degraded downstream performance. Hence, practical systems such as microgrids, require *reliability awareness*: anticipating when forecast errors are likely to be unusually large and adapting uncertainty estimates accordingly.

Prediction intervals (PIs) address this need by providing tolerance bands around predictions, and have been widely studied for nonlinear systems via fuzzy and neural models [2]. Recent PI methods have been reported for solar [3], [4] and wind forecasting [5]–[7]. Particularly, in PV day-ahead settings, residuals are strongly regime-dependent: low-signal periods and high-variability daytime intervals exhibit distinct error patterns. This motivates a *state-aware* design in which significant-error conditions are detected and the PI model adapts to the inferred state.

In this context, Spiking Neural Networks can be highlighted as a possible alternative for improving PI performance, due to the temporal anomaly/fault detection that has been shown in the literature [8]–[11]. Moreover, a reliability-aware PI that uses SNNs can be relevant for decision-making layers such as MPC under uncertainty, where bounded predictions support chance-constrained operation, such as the work [12]. In this regard, SNNs provide a temporally grounded classifier for such state inference, processing information through spiking dynamics rather than static activations. With Backpropagation Through Time (BPTT) and surrogate-gradient learning, SNNs can be optimized using standard deep learning toolchains while retaining biologically inspired neuron dynamics [8]. SNN learning has historically relied on timing-based plasticity rules (e.g., STDP) and related classifiers such as SWAT [13], SEFRON [14], TMM-SNN [15], and WOLIF [16]. Gradient-based training is challenging because spikes are non-differentiable; early approaches (e.g., SpikeProp [17]) approximated dynamics around firing events. Surrogate-gradient methods [8] replace the spike derivative with a smooth proxy, enabling BPTT and standard optimiza-

This research was funded by ANID PIA AFB230002, ANID/FONDECYT 1220507, ANID/FONDECYT 1260354, ANID/ANILLO ATE230035, ANID FONDECYT Postdoctorado 3250509, ANID/FONDECYT de Iniciación 11240510, and ANID/Vinculación Internacional/FOVI230056.

tion, as commonly implemented in modern pipelines [18]. This choice enables end-to-end training of the SNN detector with standard loss functions and optimizers while preserving spiking dynamics at inference time.

Considering these antecedents, this paper proposes a two-stage framework that (i) defines significant-error events from an issue-time aligned residual and a validation-calibrated threshold, and (ii) learns an SNN-based event detector to estimate the probability of entering a high-error state using only issue-time features extracted from the backbone forecast curve and calendar descriptors. The predicted state then selects (or softly mixes) *two-sided*, state-specific residual magnitude models to build *asymmetric* PIs calibrated for target coverage. This design enables conservative uncertainty allocation during inferred high-risk conditions, while keeping intervals tighter during low-variability periods.

A. Contributions

The main contributions, which can be highlighted within the entire forecasting-uncertainty pipeline, are the following:

- A reliability-aware two-stage prediction interval method is proposed, combining a point forecasting backbone with state-conditioned residual modeling.
- A significant-error state labeling procedure combined with a spiking neural network is used to classify the prediction error dataset, to disaggregate the prediction interval models for different types of errors.
- The proposal is tested through a case study of a microgrid with renewable generation facilities, where solar power generation and its uncertainty are modeled over a period of 4 years.

The rest of the paper is structured as follows: Section II presents the problem statement regarding the construction of prediction intervals. Section III describes the details of the framework for constructing the two-stage prediction interval proposal. Section IV presents the case study and the simulation results. Finally, Section V concludes the paper with some closing remarks.

II. Problem Statement

Let y_t represent the measured signal sampled every T period. Denote $\hat{y}_{t+h|t}$ as the h-step-ahead prediction made at time t using a backbone forecasting model, with a prediction horizon of day-ahead H. This study focuses on day-ahead uncertainty at a selected horizon H by examining the *issue-time aligned* residual:

$$e_t \triangleq y_{t+H} - \hat{y}_{t+H|t}, \tag{1}$$

and its magnitude defined as $r_t \triangleq |e_t|$. A significant forecast error event for the forecast issued at time t is indicated by the binary index g_t, which is defined as:

$$g_t = \begin{cases} 1, & \text{if } r_t > \mu, \text{ (is a significant event)} \\ 0, & \text{if } r_t \le \mu, \text{ (is not a significant event),} \end{cases} \tag{2}$$

where $\mu > 0$ is calibrated using only the validation split. The objectives of this study are twofold: (i) to estimate the probability $p_{\text{event}}(t) = \mathbb{P}(g_t = 1|\mathbf{F}_t)$ based on the information available at the forecast issue time, and (ii) to construct prediction intervals for y_{t+H} using state-dependent models of residual magnitude, selected based on a state decision derived from $p_{\text{event}}(t)$.

III. Two-Stage SNN-Gated Prediction Intervals

The proposed method is a two-stage, state-aware pipeline designed for uncertainty quantification. Initially, a deep learning backbone generates a day-ahead point forecast along with a feature vector that is extracted from the predicted trajectory. These features are used in two ways: (i) an SNN gate that classifies the forecast into normal and significant-error states, and (ii) state-conditioned residual regressors that estimate the magnitudes of asymmetric errors. The predicted state is used to select (or mix) the appropriate residual model, which helps in constructing asymmetric prediction intervals around the backbone forecast. The offline training and online inference processes are illustrated in Fig. 1 and Fig. 2.

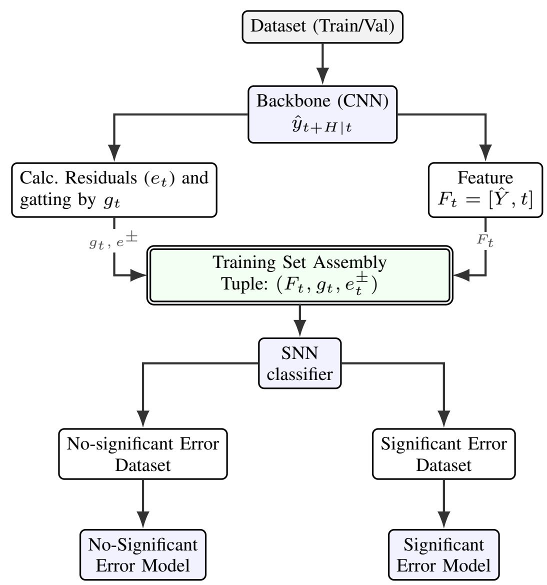

Fig. 1. Offline training pipeline. The system prepares signed targets $e^{\pm}$ and separates regressors by ground truth state g_t for parallel training.

A. Stage 1: Significant-Error Event Labeling

A validation-calibrated threshold is employed to define significant error events while maintaining controlled prevalence. Let π represent the target validation event rate (for example, $\pi = 0.10$). The threshold is determined as follows:

$$\mu = \text{Quantile}\left(\{r_t\}_{t\in\mathcal{V}},\, 1-\pi\right), \tag{3}$$

ensuring that $\mathbb{P}_{t\in\mathcal{V}}(r_t > \mu) \approx \pi$. This process generates ground-truth state labels g_t, which are utilized for training detectors and state-specific residual magnitude models.

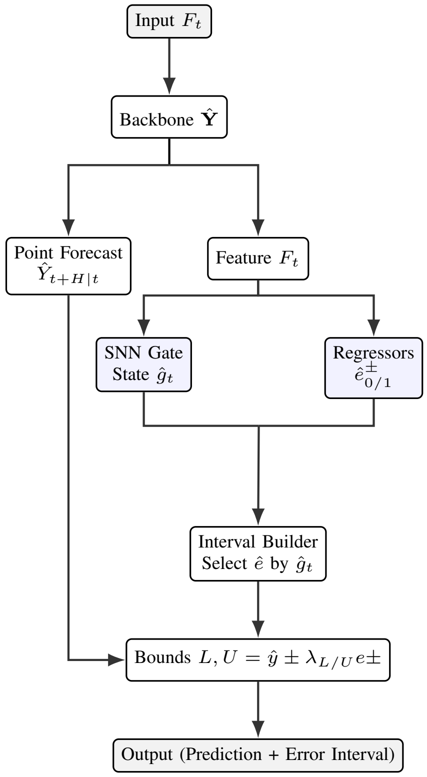

Fig. 2. Online inference flow. Separate branches handle the point forecast and the state-dependent uncertainty estimation, merging to form the final interval.

B. Issue-Time Features from the Backbone Forecast Curve

The state classifier and residual magnitude regressors operate solely with features available at the time of the issue. A compact feature vector $\mathbf{F}_t$ is computed from the backbone multi-step prediction curve $\hat{\mathbf{Y}} = [\hat{y}_{t+1|t}, \ldots, \hat{y}_{t+H|t}]$ and calendar descriptors:

$$\mathbf{F}_t = [\hat{\mathbf{Y}}_{t|t}, t] \in \mathbb{R}^d. \tag{4}$$

The mapping $\mathbf{F}_t$ incorporates summary statistics of the predicted trajectory shape (such as endpoint values, mean, extrema, amplitude, and discrete derivative statistics) as well as periodic encodings of the hour of the day and day of the year using sine and cosine terms. This design guarantees that both state classification and uncertainty estimation can be performed at deployment without the need for future measurements or residual histories.

C. SNN State Classifier and Operating Threshold

A surrogate-gradient Spiking Neural Network (SNN), implemented in snnTorch, is trained as a binary classifier to estimate $p_{\text{event}}(t)$. Inputs are standardized and mapped to $[0, 1]$ for rate coding, generating spike trains over T simulation steps. Leaky Integrate and Fire (LIF) neurons, characterized by a decay factor β, process the spiking inputs, and a membrane-based readout produces logits that are converted to probabilities using softmax. Training is performed using class-weighted cross-entropy, Adam optimization, and early stopping based on validation loss.

A validation-calibrated operating threshold τ translates probabilities into state decisions:

$$\hat{g}_t = \begin{cases} 1, & \text{if } p_{\text{event}}(t) \geq \tau, \\ 0, & \text{if } p_{\text{event}}(t) < \tau, \end{cases} \tag{5}$$

that predicts the significance of the event g_t, previously defined in (2).

At this point, two practical selection modes are available. The *rate* mode sets τ as a validation quantile of p_{event} to maintain a target predicted state frequency ρ (for instance, $\rho = 0.20$). The *Youden* mode selects τ to maximize the objective function

$$J = \text{True Positive Rate} - \text{False Positive Rate} \tag{6}$$

during validation. The rate mode provides explicit control over gating frequency, while the Youden mode prioritizes discriminability.

D. Stage 2: Two-Sided State-Specific Residual Magnitude Models

Prediction intervals require an asymmetric approach to handling positive and negative errors. To achieve this, the residuals e_t are decomposed into non-negative components defined as follows:

$$e_t^+ = \max(e_t, 0), \qquad e_t^- = \max(-e_t, 0). \tag{7}$$

Four non-negative regressors models (which can be based on any suitable neural network model) are trained: one pair for each state $r \in \{0, 1\}$,

$$\hat{e}_r^+(t) = f_r^+(\mathbf{x}_t), \qquad \hat{e}_r^-(t) = f_r^-(\mathbf{x}_t), \tag{8}$$

where the regressors utilize a Softplus output to enforce non-negativity. During training, the state assignments for the residual models use the ground-truth labels g_t defined by μ, rather than the predictions $\hat{g}_t$. This preserves the identifiability of state-conditioned residual distributions.

At inference time, the state signal derived from the SNN determines how the state-specific residual models are applied. In the case of **hard gating**, the estimates are taken as follows:

$$\hat{e}^+(t) = \hat{e}^+_{\hat{g}_t}(t), \qquad \hat{e}^-(t) = \hat{e}^-_{\hat{g}_t}(t). \tag{9}$$

Alternatively, a **soft mixing weight** $w_t \in [0, 1]$ can be formed from $p_{\text{event}}(t)$ (for example, with a sharpening exponent $p > 0$), allowing residual magnitudes to be mixed as follows:

$$\hat{e}^{\pm}(t) = (1 - w_t)\hat{e}_0^{\pm}(t) + w_t \hat{e}_1^{\pm}(t), \qquad w_t \leftarrow w_t^p. \tag{10}$$

This soft mixing approach mitigates brittleness when p_{event} fluctuates across splits while highlighting high-risk conditions.

E. Interval Construction, Calibration, and Robustness Constraints

Asymmetric prediction intervals for the day-ahead target y_{t+H} are constructed around the backbone forecast $\hat{y}_{t+H|t}$:

$$L_t = \hat{y}_{t+H|t} - \lambda_{\text{lw}} \hat{e}^-(t), \qquad U_t = \hat{y}_{t+H|t} + \lambda_{\text{up}} \hat{e}^+(t). \quad (11)$$

The scaling factors λ_{lw} and λ_{up} are calibrated on validation data to achieve a target coverage level α, i.e., $y_{t+H} \in [L_t, U_t]$ in a α fraction of the total cases in the used dataset. Additionally, two-sided calibration is utilized at this point by considering normalized residual ratios, such as selecting high quantiles of the following:

$$\begin{aligned} R_{\text{up}} &= \frac{\max(0, y_{t+H} - \hat{y}_{t+H|t})}{\hat{e}^+(t) + \epsilon} \\ R_{\text{lw}} &= \frac{\max(0, \hat{y}_{t+H|t} - y_{t+H})}{\hat{e}^-(t) + \epsilon} \end{aligned} \quad (12)$$

with a quantile level of $(1 + \alpha)/2$. To enhance conservativeness during event states, an optional **event coverage boost** $\Delta\alpha$ can be applied, recalibrating the event state to $\alpha_1 = \min(0.999, \alpha + \Delta\alpha)$ while retaining $\alpha_0 = \alpha$ for the non-event state.

To improve robustness, predicted residual magnitudes are constrained via validation quantiles (a process known as winsorization):

$$\hat{e}^\pm(t) \leftarrow \min\left(q^\pm_{\text{cap}}, \max(\hat{e}^\pm(t), q^\pm_{\text{floor}})\right), \quad (13)$$

where $q^\pm_{\text{floor}}$ and $q^\pm_{\text{cap}}$ are estimated from validation predictions (e.g., at the 5% and 99.5% levels). The floors prevent degenerate widths in low-variance conditions, while the caps reduce the impact of extreme regressor outputs that could otherwise produce implausible interval spikes.

Finally, for non-negative processes, such as photovoltaic generation signals, the lower bound can be optionally clipped as follows:

$$L_t \leftarrow \max(L_t, 0). \quad (14)$$

With the proposed methodology already described, the next section presents the case study and simulation results achieved in this work.

IV. Case Study and Results

A. Dataset and Case Study

The case study is based on the residential UK energy dataset available in IEEE Dataport [19]. The data is sampled at 15-minute resolution over a multi-year period and includes residential electricity load, electricity price, renewable generation, and weather variables. In the present study, the target variable is the photovoltaic generation signal `Residential_solar_generation` (kW), which is non-negative.

The experimental protocol follows chronological splitting into train/validation/test sets to reflect deployment under distribution shift. The backbone provides wide-format multi-step predictions (e.g., `pred_h*` and `true_h*`). Residual alignment is performed at $H = 96$ (prediction one day-ahead), and the labeling threshold μ is calibrated on validation to enforce a target event prevalence π. The SNN detector is trained to predict g_t, and the operating threshold τ is calibrated on validation using a specified mode. Residual magnitude regressors are trained on issue-time features, and prediction intervals are calibrated using validation ratios and then evaluated on train/validation/test.

Two configurations are compared:

- **No gating (2 models):** a single pair of residual magnitude regressors models (e^+ and e^-) is trained on the full dataset, and a single pair of calibration factors $(\lambda_{\text{up}}, \lambda_{\text{lw}})$ is used.
- **With SNN gating (4 models):** the dataset is partitioned by state and four regressors models are trained (2 states $\times$ 2 signs), using an SNN-derived state decision $\hat{g}_t$ (hard or soft gating) at inference time.

B. Metrics

Two categories of metrics are reported.

State detection: Ranking performance is measured using ROC-AUC (area under the Receiver Operating Characteristic curve: True Positive Rate vs. False Positive Rate across thresholds) and PR-AUC (area under the Precision–Recall curve), the latter being more informative under class imbalance. Operating-point behavior is characterized at the selected threshold τ.

Prediction intervals: Interval reliability and sharpness are assessed via empirical coverage (PICP) and normalized width (PINAW), respectively.

The prediction-interval metrics are defined as

$$\text{PICP} = \frac{1}{N} \sum_{k=1}^{N} c_k, \quad (15)$$

$$c_k = \begin{cases} 1, & \text{if } L_k \le y_{k+H} \le U_k, \\ 0, & \text{otherwise}, \end{cases} \qquad c_k \in \{0, 1\}. \quad (16)$$

$$\text{PINAW} = \frac{1}{N\,(y_{\max} - y_{\min})} \sum_{k=1}^{N} (U_k - L_k), \quad (17)$$

where $(y_{\max} - y_{\min})$ is the empirical range on the evaluated split. Since the PV target is non-negative, the lower bound is clipped as $L_k \leftarrow \max(L_k, 0)$ to enforce physical plausibility.

C. Significant-error Detection and Gating Behavior

The surrogate-gradient SNN detector provides a probability estimate $p_{\text{event}}(t)$, which is converted into a state decision via the operating threshold τ. The detector achieved strong ranking performance on the test split (ROC-AUC $= 0.8888$, PR-AUC $= 0.5375$). Under a validation-calibrated operating

TABLE I
PREDICTION-INTERVAL AND GATING MODEL PERFORMANCE ($H = 96$, $\alpha = 0.90$, RMSE$= 0.2465$, MAE$= 0.1403$). FORMAT: *train* (*test*).

Gating Method	ROC-AUC	F1-score	PICP train (test)	PINAW train (test)
No gate	-	-	0.948 (0.792)	0.175 (0.204)
SNN	0.889	0.808	0.943 (0.900)	0.261 (0.326)
MLP	0.811	0.833	0.952 (0.883)	0.193 (0.242)
XGBoost	0.843	0.861	0.955 (0.890)	0.200 (0.255)
LSTM	0.839	0.858	0.956 (0.886)	0.193 (0.243)

point (rate mode), the predicted activation rate matches the target on validation but may drift on test under score-distribution shift, leading to a more conservative activation of the event-conditioned uncertainty model. This behavior is consistent with reliability-aware operation, where false negatives on significant-error events are typically more detrimental than false positives.

D. Comparative analysis of gating alternatives

Table I compares four gating alternatives against a no-gating baseline. Across all configurations, only the Stage 1 gate changes, while the CNN backbone (test RMSE $= 0.2465$, MAE $= 0.1403$) and the Stage 2 MLP model structure remain fixed. The calibration factors λ_{up} and λ_{lw} are estimated on the training set to enforce the nominal coverage target $\alpha = 0.90$, whereas validation and test are used only for evaluation under chronological distribution shift.

The no-gating baseline achieves near-nominal training coverage (PICP $= 0.900$) but drops to PICP $= 0.808$ on test, indicating that a single global residual model does not generalize well under temporal shift. Among the gated models, the SNN yields the highest test coverage (PICP $= 0.869$), consistent with its stronger gating performance (ROC-AUC $= 0.889$, F1-score $= 0.808$), although at the cost of wider intervals (PINAW $= 0.284$). By contrast, the MLP, XGBoost, and LSTM gates produce narrower intervals (PINAW $= 0.208$, 0.217, and 0.205, respectively) but lower test coverage (PICP between 0.836 and 0.842). Based on this comparison, it is noteworthy that the SNN-based approach achieves more consistent performance in terms of interval coverage, demonstrating less overfitting than the other methods.

E. Qualitative Interval Behavior (Spring/Autumn Examples)

Fig. 3 illustrates a representative spring test window with SNN gating, showing the day-ahead forecast, the constructed prediction interval, and the gate panel ($p_{\text{event}}(t)$, τ, and $\hat{g}_t$). The intervals widen around high-variability periods (e.g., ramps and peak-generation transitions) while remaining tighter during low-signal periods.

Fig. 4 presents an autumn example under SNN gating, highlighting the seasonal variability and the corresponding interval adaptation. In both cases, the predicted state modulates the interval widths via state-specific residual magnitude models.

For comparison, Fig. 5 shows the same spring window without gating (2-model configuration). In this baseline, uncertainty is modeled globally and cannot explicitly allocate additional width to inferred high-risk conditions, which is consistent with the observed under-coverage on the test split.

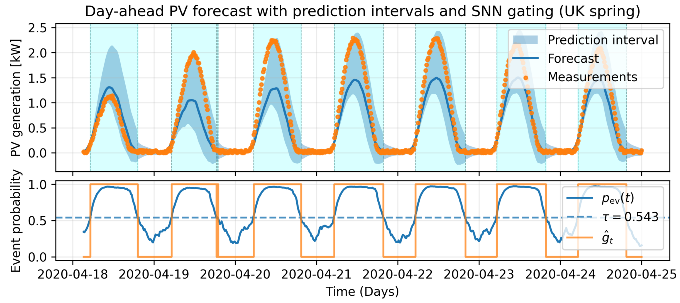

Fig. 3. Day-ahead PV forecast with prediction intervals and SNN gating (UK spring). The top panel shows forecast, measurements, and the prediction interval; the bottom panel shows $p_{\text{event}}(t)$, the operating threshold τ, and the resulting state decision $\hat{g}_t$.

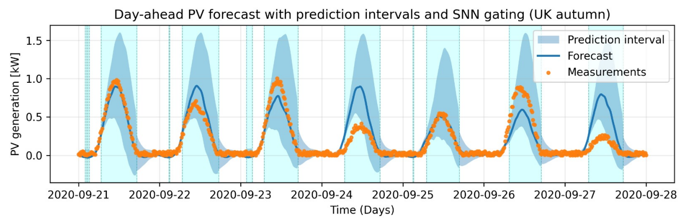

Fig. 4. Day-ahead PV forecast with prediction intervals and SNN gating (UK autumn). The interval adapts to a entire year variability and inferred high-risk conditions.

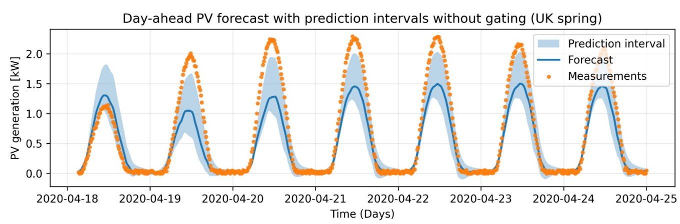

Fig. 5. Day-ahead PV forecast with prediction intervals without gating (UK spring). A single global residual model is used (2-model configuration), limiting state-adaptive uncertainty allocation.

F. Discussion

Three practical aspects are emphasized. First, the *labeling threshold* μ and the *gating threshold* τ serve different purposes: μ defines significant-error events (with controlled validation prevalence) to support identifiable state-conditioned residual models, whereas τ controls how often the event-conditioned uncertainty model is activated at deployment. Under chronological shift, the score distribution of $p_{\text{event}}(t)$ may drift, so the realized test activation rate can deviate from its validation target.

Second, two-sided residual magnitude modeling is key for asymmetric uncertainty. Modeling e^+ and e^- separately avoids symmetric-band assumptions and enables distinct upper/lower calibration, which is especially relevant for bounded and skewed energy signals such as PV generation near zero output and saturation.

Third, robustness constraints improve operational behavior. Floors on $\hat{e}^{\pm}$ prevent degenerate widths, caps reduce implausible spikes from extreme outputs, and an event coverage boost $\Delta\alpha$ adds conservativeness in event states to partially compensate for imperfect gating.

V. Conclusion

This paper presented a two-stage, reliability-aware framework for day-ahead uncertainty quantification. Significant-error events are defined using a training-calibrated threshold μ on the issue-time aligned residual, and a surrogate-gradient SNN estimates $p_{\mathrm{event}}(t)$ from issue-time features derived from the backbone forecast curve and calendar descriptors. Non-negative two-sided residual magnitude models generate asymmetric prediction intervals with separate upper and lower calibration, offering optional conservativeness during event states for non-negative solar generation targets. In the reported configuration, the detector achieved ROC-AUC $= 0.8888$ and PR-AUC $= 0.5375$ on test, and the SNN-gated configuration attained the highest test coverage among all gated benchmarks (PICP $= 0.869$) under chronological distribution shift.

A. Limitations and Implications

A main limitation is sensitivity to distribution shift in the gating scores: τ enforces a target activation rate on training but may drift on unseen splits, affecting the coverage–width trade-off. Additionally, state-dependent calibration factors may become suboptimal when state prevalence changes between training and deployment.

B. Future Work

Future work will address shift-robust deployment through online recalibration of τ via score-quantile tracking, domain adaptation for the SNN detector to reduce seasonal drift, and integration of conformal calibration layers to improve finite-sample reliability without unnecessary width inflation.

References

[1] T. Hong, P. Pinson, S. Fan, H. Zareipour, A. Troccoli, and R. J. Hyndman, “Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond,” *International Journal of Forecasting*, vol. 32, no. 3, pp. 896–913, 2016.

[2] O. Cartagena, S. Parra, D. Muñoz-Carpintero, L. G. Marín and D. Sáez, ”Review on Fuzzy and Neural Prediction Interval Modelling for Nonlinear Dynamical Systems,” in IEEE Access, vol. 9, pp. 23357-23384, 2021, doi: 10.1109/ACCESS.2021.3056003.

[3] Y. Jiang, X. Wang, D. Yang, R. Cheng, Y. Zhao, D. Liu, ”An adaptive photovoltaic power interval prediction based on multi-objective optimization,” in Computers and Electrical Engineering, vol. 120, Part A, 109717, 2024. https://doi.org/10.1016/j.compeleceng.2024.109717.

[4] W. Yuan, J. Ding, L. Zhang, J. Ni, and Q. Zhang, “Short-Term Probabilistic Prediction of Photovoltaic Power Based on Bidirectional Long Short-Term Memory with Temporal Convolutional Network,” Energies, vol. 18, no. 20, p. 5373, Oct. 2025, doi: https://doi.org/10.3390/en18205373.

[5] Y. Chen, S. S. Yu, C. P. Lim and P. Shi, ”Multi-Objective Estimation of Optimal Prediction Intervals for Wind Power Forecasting,” in IEEE Transactions on Sustainable Energy, vol. 15, no. 2, pp. 974-985, April 2024, doi: 10.1109/TSTE.2023.3321081

[6] Z.-F. Liu, Y.-y. Liu, X.-R. Chen, S.-R. Zhang, X.-F. Luo, L.-L. Li, Y.-Z. Yang and G.-D.You, “A novel deep learning-based evolutionary model with potential attention and memory decay-enhancement strategy for short-term wind power point-interval forecasting,” Applied Energy, vol. 360, pp. 122785–122785, Apr. 2024, doi: https://doi.org/10.1016/j.apenergy.2024.122785.

[7] G. Li, C. Lin, and Y. Li, “Probabilistic Forecasting of Provincial Regional Wind Power Considering Spatio-Temporal Features,” Energies, vol. 18, no. 3, p. 652, Jan. 2025, doi: https://doi.org/10.3390/en18030652.

[8] M. Eshraghian *et al.*, “Training Spiking Neural Networks Using Lessons From Deep Learning,” *Proceedings of the IEEE*, vol. 111, no. 9, pp. 1016–1054, 2023.

[9] S. Wu, Y. Zhang, and C. S. Chin, “Deep Spiking Neural Networks with Bowman’s Capsule-Alike Structure for Fault Diagnosis of Rotating Machinery,” *IEEE Transactions on Industrial Electronics*, vol. 70, no. 8, pp. 8295–8305, 2023.

[10] Z. Chen, J. Teng, W. Zhang, and Y. Liu, “A Deep Spiking Neural Network With Developmental Plasticity for Rotating Machinery Fault Diagnosis,” *IEEE Transactions on Instrumentation and Measurement*, vol. 71, pp. 1–12, 2022.

[11] Y. Luo, M. Osborne, and S. M. Hedge, “Spiking Neural Networks for Anomaly Detection in Smart Grids,” in *Proceedings of the IEEE Power & Energy Society General Meeting (PESGM)*, Denver, CO, USA, 2022, pp. 1–5.

[12] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. Scokaert, “Constrained model predictive control: Stability and optimality,” *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[13] J. J. Wade, L. J. McDaid, J. A. Santos, and H. M. Sayers, “SWAT: A spiking neural network training algorithm for classification problems,” *IEEE Transactions on Neural Networks*, vol. 21, pp. 1817–1830, 2010.

[14] A. Jeyasothy, S. Sundaram, and N. Sundararajan, “SEFRON: A new spiking neuron model with time-varying synaptic efficacy function for pattern classification,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 30, pp. 1231–1240, 2019.

[15] S. Dora, S. Sundaram, and N. Sundararajan, “An Interclass Margin Maximization Learning Algorithm for Evolving Spiking Neural Network,” *IEEE Transactions on Cybernetics*, vol. 49, no. 3, pp. 989–999, Mar. 2019, doi: 10.1109/TCYB.2018.2791282.

[16] I. Hussain and D. M. Thounaojam, “WOLIF: An efficiently tuned classifier that learns to classify non-linear temporal patterns without hidden layers,” *Applied Intelligence*, vol. 51, pp. 2173–2187, 2021.

[17] S. M. Bohte, J. N. Kok, and H. L. Poutré, “Error-backpropagation in temporally encoded networks of spiking neurons,” *Neurocomputing*, vol. 48, pp. 17–37, 2002.

[18] W. Phusakulkajorn, J. Hendriks, Z. Li, and A. Núñez, “Spiking neural network with time-varying weights for rail squat detection,” *Applied Soft Computing*, vol. 184, p. 113689, 2025, doi: 10.1016/j.asoc.2025.113689.

[19] M. Wijesooriya, ”Residential Load Consumption, Electricity price, Renewable Energy Generation DATA Set in the UK ”, IEEE Dataport, January 6, 2025, doi:10.21227/mbtn-zt12