Rethinking Time-Series Forecasting as Binary Event Prediction with Spiking Neural Networks

Peng Kang*[1], Hivanshu Tyagi*, Md Tariqul Islam*
*Department of Computer Science, University of Illinois Springfield, Springfield, IL, USA

***Abstract*—Spiking Neural Networks (SNNs) have emerged as an energy-efficient alternative to conventional artificial neural networks (ANNs), offering event-driven computation and temporal sparsity. However, most existing studies evaluate SNNs on regression-based forecasting tasks that require precise value prediction, which may not fully exploit their discrete computation nature. In this paper, we reformulate time-series forecasting as a binary directional prediction problem—predicting whether future values will rise or fall—and systematically evaluate three representative SNN architectures: iSpikformer, Spike-TCN, and Spike-GRU. Experiments on four benchmark datasets show that Spike-TCN achieves the most stable and accurate performance across both periodic and irregular data, while iSpikformer performs well on structured signals and Spike-GRU remains less effective. Interestingly, longer forecasting horizons sometimes yield higher accuracy, reflecting the deterministic nature of coarse-grained binary trends. Moreover, since the model architectures remain unchanged, all methods preserve the energy efficiency benefits demonstrated in the prior work. This study provides a new perspective for evaluating and designing SNNs for event-driven temporal reasoning and energy-efficient forecasting applications.**

***Index Terms*—Spiking Neural Networks, Binary Trend Forecasting, Event-Driven Learning, Time-Series Analysis**

I. Introduction

Recent advances in time-series forecasting have been largely driven by artificial neural networks (ANNs), including recurrent, convolutional, and Transformer-based architectures. These models have achieved impressive performance in predicting continuous-valued signals such as electricity consumption, solar energy production, and traffic flow. However, they rely on dense, real-valued activations and high-precision computations, which limit their energy efficiency and biological plausibility.

Spiking Neural Networks (SNNs) provide a more energy-efficient and event-driven alternative. By processing discrete spikes instead of continuous signals, SNNs can model temporal dynamics using biologically inspired mechanisms such as membrane potentials and firing thresholds. Recent frameworks, such as SeqSNN [1], have demonstrated that SNNs can approach the forecasting accuracy of ANNs while significantly reducing computational cost. Nevertheless, most existing SNN studies focus on **regression-based forecasting**, where the goal is to predict precise continuous future values—a setting that may neither fully exploit the event-driven nature of SNNs nor align with the way biological systems, including the human brain, naturally operate.

In contrast, humans rarely predict precise numerical values in everyday reasoning. When anticipating future events, people tend to make trend-level judgments—such as whether temperature will rise, traffic will worsen, or demand will drop—rather than exact quantitative forecasts. This cognitive tendency suggests that event-driven binary trend prediction may not only be computationally efficient but also more consistent with biological intelligence. Such a formulation naturally matches the discrete and event-based processing of SNNs, making them ideal candidates for this class of tasks.

In this work, we systematically investigate SNNs in the context of binary trend forecasting. We reformulate four benchmark time-series datasets—Electricity, Solar, METR-LA, and PEMS-Bay—into binary classification tasks that indicate whether the next time step represents an increase or a decrease relative to the current state. Using the SeqSNN framework, we evaluate three representative SNN architectures—iSpikformer, Spike-GRU, and Spike-TCN—under a unified experimental setting. Each model is evaluated across multiple forecasting horizons to assess both short-term and long-term trend prediction capability. Our results show that SNNs can effectively capture temporal dependencies and maintain stable accuracy across datasets and horizons, while preserving the inherent energy efficiency of event-driven computation.

The main contributions of this work are as follows:

- Event-driven reformulation: We introduce a unified framework that transforms traditional real-valued time-series forecasting into a binary directional trend prediction task, aligning the objective of forecasting with the event-driven computation of SNNs.
- Binary dataset construction: We generate and release binary directional versions of four widely used benchmark datasets—Electricity, Solar, METR-LA, and PEMS-Bay—facilitating systematic evaluation of event-based prediction models.
- Comprehensive SNN evaluation: We conduct extensive experiments with three representative architectures, including iSpikformer, Spike-GRU, and Spike-TCN, under a consistent SeqSNN setting, and across multiple forecasting horizons.
- Insights into event-driven forecasting: Our study provides quantitative evidence that SNNs can achieve effective temporal inference when the forecasting task is reformulated into a trend-based, event-driven paradigm.

[1]Peng Kang is the corresponding author. Email: pkang8@uis.edu

979-8-3315-7111-5/26/$31.00 ©2026 IEEE

II. Related Work

A. Time-Series Forecasting with ANNs

Traditional time-series forecasting has been dominated by artificial neural networks (ANNs) such as convolutional neural networks (CNNs) [2], [3], [4], recurrent neural networks (RNNs) [5], [6], and graph neural networks (GNNs) [7], [8]. Recently, Transformer-based architectures have achieved state-of-the-art results in a variety of forecasting tasks [9], [10]. These models predict continuous-valued signals by optimizing regression objectives, typically requiring dense real-valued computation at each time step. While effective, such approaches are energy intensive and less aligned with the sparse, event-driven nature of biological computation.

B. Time-Series Forecasting with SNNs

Spiking Neural Networks (SNNs) represent an emerging paradigm that processes information using discrete spikes rather than continuous activations [11], [12]. Through temporal integration of membrane potentials and threshold-based firing, SNNs can capture temporal dynamics with low power consumption, offering a biologically inspired alternative to ANNs. Recent work has demonstrated that SNNs can perform competitively on sequential event-based learning tasks such as speech recognition [13], gesture classification [14], tactile recognition [15], and event-based vision [16].

While SNNs have shown strong performance in event-driven perception tasks, their potential for modeling continuous temporal dynamics remains comparatively underexplored. Unlike static classification or short-duration event recognition, time-series forecasting requires capturing long-range temporal dependencies and producing temporally consistent predictions. This challenge has motivated several early efforts to apply SNNs to forecasting tasks across specific domains, such as financial data [17], wind power data [18], and electricity demand [19]. However, these studies either focused primarily on deploying SNNs to neuromorphic hardware for proof-of-concept demonstrations or achieved limited forecasting accuracy due to the simplicity of their network architectures.

More recently, the SeqSNN framework [1] has introduced a unified method for adapting modern sequence models such as GRU, TCN, and Transformer to spiking computation. By incorporating temporal alignment, hierarchical spike encoding, and surrogate-gradient learning, SeqSNN enabled SNNs to perform regression-based long-term forecasting on continuous-valued datasets with performance close to their ANN counterparts. Despite these advances, existing SNN forecasting studies remain limited to precise real-value prediction tasks. Such regression objectives may not fully leverage the event-driven processing capability of SNNs or reflect how biological systems reason about temporal information. Our work differs by reformulating time-series forecasting into a binary directional event prediction problem, thereby aligning the forecasting task more closely with both the computational characteristics of SNNs and the qualitative trend prediction observed in human cognition.

III. Methods

A. Event-Driven Binary Trend Reformulation

Traditional time-series forecasting tasks aim to predict continuous future values given a historical input window. In this work, we reformulate forecasting as a *binary directional prediction* problem, where the model predicts whether the next observation will increase or decrease relative to the current state. Given an input sequence $\{x_t\}_{t=1}^{T}$, the binary label y_t is defined as:

$$y_t = \begin{cases} 1, & \text{if } x_{t+1} - x_t > 0, \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

This reformulation emphasizes qualitative temporal trends rather than exact magnitudes, aligning the task with the event-driven processing nature of SNNs and with how humans naturally perceive and predict future changes.

B. Dataset Transformation

We apply this binary conversion to four commonly used benchmark datasets: **Electricity** [20], **Solar** [20], **Metr-la** [21], and **Pems-Bay** [21]. Each dataset contains multivariate continuous-valued sequences originally used for regression-based forecasting. After transformation based on Section III-A, each sample is represented by a fixed-length binary input window $\mathbf{X} \in \mathbb{R}^{W \times D}$ and its corresponding binary target $\mathbf{y} \in \{0,1\}^H$, where W denotes the input length, D the number of variables, and H the forecasting horizon. We generate directional labels for multiple horizons $H \in \{1, 2, 4, 6, 24, 48, 96\}$ to analyze both short-term and long-term prediction performance.

C. Model Architectures

We evaluate three representative spiking models implemented within the SeqSNN framework [1]:

- **iSpikformer**: a spiking adaptation of the iTransformer architecture [10] with spike-based self-attention and feed-forward modules;
- **Spike-GRU**: a recurrent SNN variant derived from gated recurrent units [22];
- **Spike-TCN**: a temporal convolutional SNN that captures local dependencies via spiking convolutions, inspired by Temporal Convolutional Network (TCN) [2].

All models share the same encoder–decoder backbone as SeqSNN, using a temporal alignment factor of $T_s = 4$ for each original time step. The membrane potentials are updated within each sub-step and reset at time-step boundaries to ensure training stability. The output layer aggregates spike responses and produces a probability distribution via a Sigmoid activation [23] for binary classification.

IV. Experiments

A. Experimental Setup

We evaluate our proposed event-driven binary trend forecasting task on four benchmark datasets widely used in time-series research:**Electricity** [20], **Solar** [20], **Metr-la** [21], and

TABLE I
RESULTS ON THE ELECTRICITY DATASET. BEST VALUES PER HORIZON ARE HIGHLIGHTED IN BOLD.

Model	Metric	1	2	4	6	24	48	96
iSpikformer	ACC	**0.798**	**0.820**	**0.844**	**0.863**	**0.874**	**0.869**	**0.862**
	AUC	**0.886**	**0.908**	**0.931**	**0.946**	**0.955**	**0.952**	**0.947**
	AP	**0.873**	**0.899**	**0.928**	**0.945**	**0.955**	**0.952**	**0.948**
Spike-TCN	ACC	0.750	0.765	0.798	0.820	0.839	0.840	0.838
	AUC	0.838	0.854	0.890	0.911	0.930	0.931	0.929
	AP	0.814	0.835	0.882	0.909	0.929	0.929	0.928
Spike-GRU	ACC	0.577	0.543	0.527	0.532	0.638	0.638	0.625
	AUC	0.600	0.517	0.526	0.546	0.693	0.693	0.678
	AP	0.543	0.474	0.493	0.519	0.670	0.680	0.674

TABLE II
RESULTS ON THE METR-LA DATASET. BEST VALUES PER HORIZON ARE HIGHLIGHTED IN BOLD.

Model	Metric	1	2	4	6	24	48	96
iSpikformer	ACC	0.583	0.583	**0.712**	0.582	0.575	0.572	0.548
	AUC	0.500	0.505	**0.789**	0.501	0.505	0.503	0.503
	AP	0.417	0.419	**0.702**	0.419	0.426	0.430	0.441
Spike-TCN	ACC	**0.723**	0.704	0.707	**0.705**	**0.686**	**0.693**	**0.712**
	AUC	**0.806**	**0.781**	0.782	**0.779**	0.742	**0.755**	**0.783**
	AP	**0.716**	**0.690**	0.691	**0.688**	0.647	**0.695**	**0.730**
Spike-GRU	ACC	0.706	**0.705**	0.704	0.700	0.683	0.674	0.648
	AUC	0.785	0.783	0.781	0.775	**0.743**	0.728	0.702
	AP	0.689	0.686	0.686	0.677	**0.650**	0.650	0.642

TABLE III
RESULTS ON THE PEMS-BAY DATASET. BEST VALUES PER HORIZON ARE HIGHLIGHTED IN BOLD.

Model	Metric	1	2	4	6	24	48	96
iSpikformer	ACC	0.529	0.620	0.621	0.623	0.610	0.671	0.672
	AUC	0.523	0.669	0.670	0.675	0.653	0.742	0.742
	AP	0.493	0.644	0.647	0.652	0.650	0.738	0.751
Spike-TCN	ACC	**0.632**	**0.636**	**0.642**	**0.650**	**0.669**	**0.681**	**0.743**
	AUC	**0.680**	**0.686**	**0.694**	**0.704**	**0.734**	**0.761**	**0.832**
	AP	**0.651**	**0.656**	**0.669**	**0.677**	**0.713**	**0.759**	**0.838**
Spike-GRU	ACC	0.590	0.591	0.598	0.597	0.603	0.615	0.645
	AUC	0.624	0.629	0.634	0.635	0.645	0.664	0.705
	AP	0.589	0.593	0.597	0.596	0.620	0.663	0.719

Pems-Bay [21]. Each dataset is reformulated into a binary directional prediction problem, indicating whether the next observation increases or decreases relative to the current value. Forecasting horizons are set to $\{1, 2, 4, 6, 24, 48, 96\}$ to explore both short-term and long-term dynamics.

All models share the same SeqSNN backbone. We train all models using the Adam optimizer by setting different learning rates and mini-batch sizes for different datasets. Binary cross-entropy loss is employed as the training objective. Each model is trained for a maximum of 1000 epochs with early stopping based on validation loss. To assess performance across horizons, we compute classification metrics including **Accuracy**, **AUC**, and **AP**.

B. Results and Analysis

Tables I–IV present detailed results of the three spiking models—**iSpikformer**, **Spike-TCN**, and **Spike-GRU**—on all four datasets across multiple horizons. The best performance for each horizon and metric is highlighted in **bold**.

Across all datasets, **Spike-TCN** consistently yields the strongest and most stable results, maintaining high performance on periodic datasets such as Electricity and Solar and showing clear advantages on irregular datasets like Metr-la and Pems-Bay. This confirms that convolutional temporal processing synergizes effectively with binary event-driven dynamics, allowing Spike-TCN to capture both short-term and long-term dependencies with robustness. **iSpikformer** demonstrates competitive results on structured and periodic signals but remains less stable for noisy, nonperiodic traffic data, while **Spike-GRU** performs the weakest across all datasets and horizons, reflecting the limitations of simple recurrent integration under sparse spike activation.

Although the dataset reformulation converts the task from real-valued forecasting to binary directional prediction, the intrinsic challenge remains nontrivial. The models must still

TABLE IV
RESULTS ON THE SOLAR DATASET. BEST VALUES PER HORIZON ARE HIGHLIGHTED IN BOLD.

Model	Metric	1	2	4	6	24	48	96
iSpikformer	ACC	0.816	0.813	0.809	0.805	0.774	0.731	0.677
	AUC	0.793	0.807	0.808	0.803	0.784	0.776	0.770
	AP	0.434	0.469	0.485	0.481	0.522	0.581	0.634
Spike-TCN	ACC	**0.929**	**0.924**	**0.920**	**0.920**	**0.933**	**0.946**	**0.952**
	AUC	**0.976**	**0.974**	**0.972**	**0.973**	**0.982**	**0.989**	**0.991**
	AP	**0.891**	**0.883**	**0.878**	**0.887**	**0.938**	**0.970**	**0.982**
Spike-GRU	ACC	0.825	0.820	0.815	0.810	0.773	0.730	0.678
	AUC	0.805	0.785	0.767	0.764	0.702	0.675	0.687
	AP	0.574	0.545	0.528	0.520	0.484	0.477	0.512

learn to discriminate subtle variations and noise-driven sign changes, especially in nonstationary environments. Further fine-tuning of models such as iSpikformer and Spike-GRU could help them better adapt to the event-driven binary formulation. Interestingly, there is no clear monotonic relationship between forecasting horizon and performance: in some cases, longer horizons even yield higher accuracy and AUC. This behavior is likely a consequence of the binary nature of the prediction task, where binary longer-horizon trends become smoother and more deterministic, while binary short-horizon fluctuations remain dominated by stochastic noise.

It is also important to note that since we do not modify the underlying network architectures, the **energy efficiency advantages of SNNs remain preserved**. As demonstrated in *SeqSNN* [1], spiking operations (ACs) are substantially more energy-efficient than the multiply-accumulate (MAC) operations of ANNs. Therefore, even under the proposed binary reformulation, all three models maintain the same computational and energy benefits reported in prior work.

V. CONCLUSION

This work presents a systematic study of spiking neural networks for time-series forecasting under a binary event-driven reformulation. Instead of predicting precise future values, we focus on directional trend prediction, which aligns more closely with the discrete and energy-efficient nature of SNNs as well as with human perceptual reasoning. Through experiments on four benchmark datasets, we show that spiking models can effectively capture temporal dependencies in both periodic and irregular data. Among the tested architectures, Spike-TCN demonstrates the best overall performance and stability, while iSpikformer and Spike-GRU require further tuning to handle noisy and nonstationary dynamics. The results also reveal that the binary formulation remains a nontrivial challenge, with no consistent improvement across horizons. Since the model architectures are preserved, all methods maintain the energy efficiency advantages previously established in *SeqSNN*. Overall, this study provides a new perspective on evaluating SNNs for practical forecasting tasks and lays the groundwork for future exploration of event-driven learning frameworks for trend analysis.

ACKNOWLEDGMENT

This work was supported in part by the *UIS AI Campus Server* provided by the Information Technology Services (ITS) at the University of Illinois Springfield. We gratefully acknowledge the computing resources and technical assistance that made this research possible.

REFERENCES

[1] Changze Lv, Yansen Wang, Dongqi Han, Xiaoqing Zheng, Xuanjing Huang, and Dongsheng Li, "Efficient and effective time-series forecasting with spiking neural networks," *arXiv preprint arXiv:2402.01533*, 2024.

[2] Shaojie Bai, J Zico Kolter, and Vladlen Koltun, "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," *arXiv preprint arXiv:1803.01271*, 2018.

[3] Minhao Liu, Ailing Zeng, Muxi Chen, Zhijian Xu, Qiuxia Lai, Lingna Ma, and Qiang Xu, "Scinet: Time series modeling and forecasting with sample convolution and interaction," *Advances in Neural Information Processing Systems*, vol. 35, pp. 5816–5828, 2022.

[4] Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long, "Timesnet: Temporal 2d-variation modeling for general time series analysis," *arXiv preprint arXiv:2210.02186*, 2022.

[5] Xu Zhang, Furao Shen, Jinxi Zhao, and GuoHai Yang, "Time series forecasting using gru neural network with multi-lag after decomposition," in *International conference on neural information processing*. Springer, 2017, pp. 523–532.

[6] Sima Siami-Namini, Neda Tavakoli, and Akbar Siami Namin, "The performance of lstm and bilstm in forecasting time series," in *2019 IEEE International conference on big data (Big Data)*. IEEE, 2019, pp. 3285–3292.

[7] Bing Yu, Haoteng Yin, and Zhanxing Zhu, "Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting," *arXiv preprint arXiv:1709.04875*, 2017.

[8] Yuchen Fang, Kan Ren, Caihua Shan, Yifei Shen, You Li, Weinan Zhang, Yong Yu, and Dongsheng Li, "Learning decomposed spatial relations for multi-variate time-series modeling," in *Proceedings of the AAAI conference on artificial intelligence*, 2023, vol. 37, pp. 7530–7538.

[9] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang, "Informer: Beyond efficient transformer for long sequence time-series forecasting," in *Proceedings of the AAAI conference on artificial intelligence*, 2021, vol. 35, pp. 11106–11115.

[10] Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long, "itransformer: Inverted transformers are effective for time series forecasting," *arXiv preprint arXiv:2310.06625*, 2023.

[11] Wolfgang Maass, "Networks of spiking neurons: the third generation of neural network models," *Neural networks*, vol. 10, no. 9, pp. 1659–1671, 1997.

[12] Kaushik Roy, Akhilesh Jaiswal, and Priyadarshini Panda, "Towards spike-based machine intelligence with neuromorphic computing," *Nature*, vol. 575, no. 7784, pp. 607–617, 2019.

[13] Jibin Wu, Emre Yılmaz, Malu Zhang, Haizhou Li, and Kay Chen Tan, "Deep spiking neural networks for large vocabulary automatic speech recognition," *Frontiers in neuroscience*, vol. 14, pp. 199, 2020.

[14] Yannan Xing, Gaetano Di Caterina, and John Soraghan, "A new spiking convolutional recurrent neural network (scrnn) with applications to event-based hand gesture recognition," *Frontiers in neuroscience*, vol. 14, pp. 590164, 2020.

[15] Peng Kang, Srutarshi Banerjee, Henry Chopp, Aggelos Katsaggelos, and Oliver Cossairt, "Boost event-driven tactile learning with location spiking neurons," *Frontiers in Neuroscience*, vol. 17, pp. 1127537, 2023.

[16] Lin Zhu, Xiao Wang, Yi Chang, Jianing Li, Tiejun Huang, and Yonghong Tian, "Event-based video reconstruction via potential-assisted spiking neural network," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 3594–3604.

[17] David Reid, Abir Jaafar Hussain, and Hissam Tawfik, "Financial time series prediction using spiking neural networks," *PloS one*, vol. 9, no. 8, pp. e103656, 2014.

[18] Juan Manuel González Sopeña, Vikram Pakrashi, and Bidisha Ghosh, "A spiking neural network based wind power forecasting model for neuromorphic devices," *Energies*, vol. 15, no. 19, pp. 7256, 2022.

[19] Santosh Kulkarni, Sishaj P Simon, and Kinattingal Sundareswaran, "A spiking neural network (snn) forecast engine for short-term electrical load forecasting," *Applied Soft Computing*, vol. 13, no. 8, pp. 3628–3635, 2013.

[20] Guokun Lai, Wei-Cheng Chang, Yiming Yang, and Hanxiao Liu, "Modeling long-and short-term temporal patterns with deep neural networks," in *The 41st international ACM SIGIR conference on research & development in information retrieval*, 2018, pp. 95–104.

[21] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu, "Diffusion convolutional recurrent neural network: Data-driven traffic forecasting," *arXiv preprint arXiv:1707.01926*, 2017.

[22] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio, "Empirical evaluation of gated recurrent neural networks on sequence modeling. arxiv 2014," *arXiv preprint arXiv:1412.3555*, vol. 1412, 2014.

[23] Shiv Ram Dubey, Satish Kumar Singh, and Bidyut Baran Chaudhuri, "Activation functions in deep learning: A comprehensive survey and benchmark," *Neurocomputing*, vol. 503, pp. 92–108, 2022.