Journal of Computer Science and Artificial Intelligence
ISSN: 3078-8242 (Print), ISSN: 3078-2309 (Online) | Vol. 7, No. 2, 2026

A Time-Series Motif-Based Rule Fusion Method for Interpretable State Prediction of Aluminum Electrolysis Cells

Danyang Cao *, Shuobo Yu, Zifeng Lin, Gao Lei

School of Artificial Intelligence and Computer, North China University of Technology, Beijing 100144, China

Abstract: Accurate state prediction is important for improving the stability and efficiency of aluminum electrolysis cells, but practical prediction remains difficult because the production process contains strongly coupled variables, delayed process responses, and many decisions based on operator experience. To address these issues, this paper presents an interpretable motif-based rule fusion method for cell state prediction. The method first separates industrial variables into decision variables and state variables so that manual control actions and automatically monitored responses can be modeled with clear semantic roles. Multivariate motif groups are then mined from each dataset, after which temporal rules and association rules are extracted. Association rules describe delayed operation-response relationships between decision motifs and subsequent state motifs, while temporal rules describe weekday and monthly occurrence regularities of state motifs. During prediction, the method prioritizes matched association rules and uses temporal rules as a fallback when no reliable association rule is available. Experiments on large-scale data from 3,391 pot controllers show that the combined rule-fusion strategy achieves lower pattern-level prediction error than single-rule strategies and representative forecasting baselines. The method also provides explicit evidence linking interventions to later state evolution, supporting interpretable process monitoring and operator-oriented decision support.

Keywords: Aluminum electrolysis, Process monitoring, State prediction, Rule mining, Motif discovery, Interpretable prediction.

1. Introduction

Aluminum and its alloys are widely used in transportation, construction, and high-end manufacturing because of their favorable mechanical properties and lightweight characteristics [1, 2]. Under the dual pressure of energy conservation and green production, aluminum electrolysis is evolving toward higher efficiency, improved stability, and more intelligent control [3]. However, the production process involves many strongly coupled variables, and their temporal dynamics are highly complex. In practical operation, many control decisions still rely on accumulated operator experience rather than unified data-driven mechanisms. As a result, it remains difficult to predict how a given manual intervention will influence subsequent cell states.

Multivariate time series provide a natural representation for industrial processes because multiple variables are observed simultaneously over time [4]. Among time-series analysis methods, motif discovery is particularly valuable because motifs capture recurrent local patterns that often correspond to typical operating conditions or process events [5]. In aluminum electrolysis plants, large volumes of historical monitoring data have been accumulated, providing an opportunity to uncover interpretable operational patterns from industrial records. However, many existing approaches rely on symbolic discretization, which may reduce computational burden but often causes information loss and weakens prediction accuracy [6]. Another practical limitation is that most existing predictive methods treat all variables equally, although aluminum electrolysis data naturally consist of two distinct categories: manually adjusted control parameters and automatically observed state parameters. Ignoring this distinction obscures the underlying operation-response mechanism that is central to process understanding and prediction. This issue is particularly critical in aluminum electrolysis, where practically meaningful changes are often reflected not by isolated point-wise fluctuations, but by short-term pattern evolution across multiple coupled variables following specific control actions. In other words, what matters in process control is often not whether the value of a single variable can be predicted with minimum point-wise error at one timestamp, but whether a characteristic state pattern is likely to emerge under a given intervention. Meanwhile, although deep learning models can achieve competitive forecasting accuracy in many time-series tasks, they often provide limited interpretability, which reduces their value in process control scenarios where credible decision support is required. Therefore, motif-level and rule-level prediction is particularly attractive for industrial applications such as aluminum electrolysis, where both interpretability and operational credibility are essential.

To address these issues, this paper proposes an interpretable motif-based rule fusion method for aluminum electrolysis cell state prediction. The core idea is to decouple industrial variables into decision and state datasets, mine motif groups separately, and then construct a prediction algorithm by combining temporal rules and association rules. In this way, future state prediction is transformed from black-box numerical extrapolation into explicit rule-based inference.

The main contributions of this work are as follows:

- A parameter-decoupled motif mining strategy is proposed for aluminum electrolysis multivariate time series, where decision and state variables are modeled separately to reveal the "operation-response" structure.
- A unified rule mining algorithm is developed to extract multi-granularity temporal rules and association rules from motif groups.
- An interpretable rule-fusion prediction algorithm is

proposed, where association rules are used first when matched; otherwise, temporal rules are applied.

2. Related Work

The aluminum electrolysis process involves complex physicochemical reactions and strong nonlinear coupling among multiple operational variables [7]. Traditionally, process control and parameter adjustment rely heavily on operator experience, which makes systematic modeling and prediction difficult [8, 9]. With the development of industrial data analytics, time-series analysis has become an important tool for understanding process dynamics.

Motif discovery has been widely used to identify recurring patterns in temporal data. Early work on subsequence matching includes Dynamic Time Warping (DTW)-based methods [10] and Euclidean distance-based approaches [11]. These methods are effective for detecting repeated operational patterns, but they do not explicitly characterize relationships between control actions and subsequent process states.

Rule mining methods have also been extensively studied in data mining. Association rule mining was introduced by Agrawal et al. [12] to discover relationships among variables using support and confidence. Sequential pattern mining algorithms such as SPADE, PrefixSpan, and GSP further incorporated temporal order into rule discovery [13, 14, 15]. These methods provide useful tools for temporal relation extraction, but they were mainly developed for transactional or symbolic data.

To extend rule mining to time-series analysis, several studies introduced discretization or symbolic abstraction strategies [16, 17, 18, 19]. Although such methods improve interpretability and reduce computational complexity, they often suffer from information loss and limited robustness when dealing with noisy industrial data. In aluminum electrolysis, these limitations become more evident because the data contain missing values, gap points, and strong differences between manually controlled variables and automatically observed variables.

In recent years, deep learning has become dominant in time-series forecasting, and models such as TimesNet [20], Autoformer [21], DLinear [22], FEDformer [23], Informer [24], and LightTS [25] have shown strong performance on public benchmarks. However, such models are usually designed as general-purpose forecasters and offer limited interpretability in industrial process analysis.

Recent studies further show that motif discovery and interpretable time-series modeling remain important for complex multivariate data. For example, LoCoMotif was proposed to discover time-warped motifs in time series and demonstrates that flexible motif discovery is useful for identifying repeated temporal patterns [26]. The ILMGD method further extends motif discovery to multidimensional time series by mining motif groups based on matrix profile structures [4]. In industrial scenarios, interpretable multivariate time-series models have also attracted increasing attention because process engineers need not only prediction results but also understandable evidence for process monitoring and diagnosis [27, 28]. These studies support the necessity of combining temporal-pattern discovery with interpretability in industrial time-series analysis.

Nevertheless, most existing motif-based methods focus mainly on discovering recurrent patterns, while most rule-mining methods focus on symbolic or transactional relationships. They do not explicitly distinguish manually adjusted decision variables from automatically monitored state variables in aluminum electrolysis, nor do they directly model delayed operation-response relationships between decision motifs and state motifs. Motivated by these limitations, this work proposes a parameter-decoupled motif and rule-fusion method that jointly exploits temporal regularities and association rules for interpretable prediction of aluminum electrolysis cell states.

3. Proposed Method

3.1. Terminology and Variable Separation

The method starts with a semantic separation of variables. Decision variables are manually adjusted or control-side variables, while state variables are automatically monitored process-response variables. This separation is essential because it prevents control actions and state reactions from being mixed in a single motif representation. Table 1 lists the main variable groups used in the proposed framework. The variable separation also makes the later rules easier to interpret: the antecedent of a rule is a decision motif, and the consequent is a state motif.

Table 1. Variable grouping used in the proposed method

Variable group	Typical variables	Industrial role
Decision variables	Set Voltage; Number of Fluoride Additions; Feed Amount	Manual or control-side actions that describe operational interventions.
State variables	Electrolysis Temperature; Electrolyte Level; Aluminum Level; Voltage Swing; Average Voltage; Working-Set Voltage Difference; Anode Rod Vibration	Automatically monitored responses that describe cell condition and process evolution.

Let the complete multivariate time-series dataset contain N variables and L time points. The original variable set is divided into a decision dataset and a state dataset. Motif groups mined from the decision dataset are called decision motifs, and motif groups mined from the state dataset are called state motifs. A decision motif is a recurrent control-action pattern. A state motif is a recurrent response pattern. An association rule links a decision motif to a later state motif and represents a delayed operation-response relationship. A temporal rule summarizes the weekday or monthly occurrence regularity of a state motif.

3.2. Parameter-Decoupled Motif Mining

The ILMGD algorithm is used as the motif group mining tool. The method mines motif groups independently from the decision dataset and the state dataset. A motif group consists of multiple similar motif occurrences that belong to the same pattern category. This design improves interpretability because decision motifs describe what the operator or control mechanism does, whereas state motifs describe how the electrolytic cell reacts. Without this separation, the antecedent and consequence of process evolution may be mixed, and the resulting motifs would be difficult to translate into operational evidence.

From an industrial perspective, a useful motif should describe a short-term evolution shape across one or more

variables. For example, a decision motif may include a decrease in feed amount and a change in set voltage. A corresponding state motif may include a later decline in electrolysis temperature and average voltage. By using motif groups instead of isolated points, the method focuses on recurrent evolution patterns, which are more consistent with how engineers interpret process behavior.

3.3. Temporal Rule Extraction

After state motif groups are obtained, temporal rules are extracted to characterize occurrence regularities. Two temporal granularities are considered: weekday and month. Table 2 summarizes the temporal partitioning strategy. For each motif group, the total number of occurrences is counted. For a given temporal stage, such as Sunday or August, the number of motif occurrences in that stage is divided by the total number of occurrences in the motif group. The occurrence frequency is calculated as follows:

$$f_{\text{stage}} = \frac{n_{stage}}{n_{mg}} \quad (1)$$

Where n_stage is the number of occurrences in the given temporal stage and n_mg is the total number of motif occurrences in the motif group.

Table 2. Temporal rule partitioning methods

Time granularity	Corresponding range	Purpose
Weekday	Monday-Sunday	Captures weekly occurrence regularity of state motifs.
Month	January-December	Captures seasonal or monthly occurrence regularity of state motifs.

The temporal rule for each motif group is formed by selecting the stage with the maximum occurrence frequency under each temporal dimension. During prediction, these temporal rules are used as a fallback mechanism when no reliable association rule is matched. This ensures that the prediction process remains rule-based and interpretable even when a direct decision-to-state rule is unavailable.

3.4. Association Rule Mining

After decision motifs and state motifs are obtained, we construct association rules to capture delayed operation-response relationships. In each association rule, the antecedent is a decision motif, and the consequent is a state motif. Therefore, an association rule describes the following logic: when a certain decision motif occurs, a corresponding state motif is likely to appear after a typical lag. Algorithm 1 details the mining process for multivariate time series motif association rules.

For each candidate pair of decision motif group dec_group and state motif group sta_group, we calculate both the confidence and the typical lag as follows:

(1) Let N_d be the total number of occurrences of motifs in dec_group.

(2) For each occurrence of a decision motif at time t_d, search for occurrences of motifs in sta_group within a lag window $[t_d, t_d + max\_lag]$.

(3) Count the number of successful matches N_{ds}, where a successful match occurs if a state motif appears within the valid lag window.

(4) Compute the confidence of the rule as:

$$conf = \frac{N_{ds}}{N_d} \quad (2)$$

Which quantifies how likely the state motif follows the decision motif.

(5) Collect all lag values for successful matches and take the mode as the typical lag of the rule. This reduces the effect of occasional timing fluctuations.

(6) Apply a threshold $conf\_threshold$ to filter rules: only rules with $conf \geq conf\_threshold$ are retained in the association rule set.

These steps produce a set of interpretable rules linking decision motifs to subsequent state motifs with quantified confidence and typical lag, which will be used in the prediction algorithm. By explicitly specifying the antecedent, consequent, confidence, and lag, the resulting rules provide a transparent mechanism for understanding how manual interventions affect future cell states.

Algorithm 1 Association Rule Mining

Input: Decision motif groups dec_mgs, state motif groups sta_mgs, maximum lag max_lag, confidence threshold conf_threshold

Output: Association rule set Mg_Ass, with each rule containing antecedent, consequent, confidence, and typical lag

Initialize empty rule set: $Mg\_Ass \leftarrow \emptyset$
for each decision motif group dec_group in dec_mgs do
  for each state motif group sta_group in sta_mgs do
    $N_d \leftarrow$ total occurrences of dec_group
    $N_{ds} \leftarrow 0$, $lags \leftarrow \emptyset$
    for each occurrence time t_d of dec_group do
      Find occurrences t_s in sta_group such that $t_d < t_s \leq t_d + max\_lag$
      if any t_s exists then
        $N_{ds} \leftarrow N_{ds} + 1$
        Add $t_s - t_d$ to lags
      end if
    end for
    if $N_{ds}/N_d \geq conf\_threshold$ then
      $Mg\_Ass \leftarrow Mg\_Ass \cup$ {antecedent = dec_group, consequent = sta_group, confidence = N_{ds}/N_d, lag = mode(lags)}
    end if
  end for
end for
    return Mg_Ass

3.5. Rule-Fusion Prediction

State prediction of aluminum electrolysis cells is performed through rule-fusion prediction, and Algorithm 2 details the process.

First, given a target decision subsequence, the algorithm compares it with the antecedents of the mined association rules. If one or more association rules are matched, each candidate rule is evaluated by a comprehensive score that combines subsequence similarity and rule confidence:

$$score = w \cdot (threshold - dis) + (1 - w) \cdot conf \quad (3)$$

Where:

• dis is the normalized z-score Euclidean distance between the target decision subsequence and the antecedent decision motif;

• $conf$ is the confidence of the matched association rule;

• $w \in [0,1]$ balances motif similarity and rule reliability;

• $threshold$ is the distance threshold used for motif matching.

The association rule with the highest score is selected, and its consequent state motif is used to construct the predicted state sequence according to the corresponding typical lag. If no reliable association rule is matched, the algorithm uses temporal rules as a fallback mechanism. In this case, the current weekday and month are used to select the state motif with the highest occurrence frequency under the corresponding temporal stage. Figure 1 illustrates the overall workflow of rule-fusion prediction.

The evaluation metric is defined as:

$$E = \frac{\sum_{i=1}^{n} \frac{dis(fseq_i, tseq_i)}{dis(rseq_i, tseq_i)}}{n} \quad (4)$$

Where $fseq_i$ is the predicted sequence, $tseq_i$ is the true sequence, and $rseq_i$ is a random benchmark sequence. A smaller E indicates better predictive performance, and $E <$ 1means that the method outperforms random prediction.

Algorithm 2 Rule-Fusion Prediction

Input: Test decision data test_dec_data, training decision data train_dec_data, test state data test_sta_data, training state data train_sta_data, association rules Mg_{Ass}, temporal rules Temp_Rules, state motif groups sta_mgs, subsequence length L, similarity threshold threshold, weight w

Output: Predicted state sequences

1. Randomly sample subsequences from test_dec_data
2. for each sampled decision subsequence dec_seq do
3. Initialize predicted_seq ← ∅
4. Identify all association rules whose antecedent matches dec_seq within similarity threshold
5. if one or more rules match then
6. for each matched rule do
7. Compute score: $score = w(threshold - dis) + (1 - w)conf$
8. end for
9. Select rule with highest score
10. Use consequent motif and typical lag to construct predicted_seq
11. else
12. Map the current weekday and month to the state motif with the highest occurrence frequency in Temp_Rules
13. Construct predicted_seq from the selected motif
14. end if
15. Combine predicted_seq with historical average of the corresponding state motif
16. end for
17. return predicted sequences

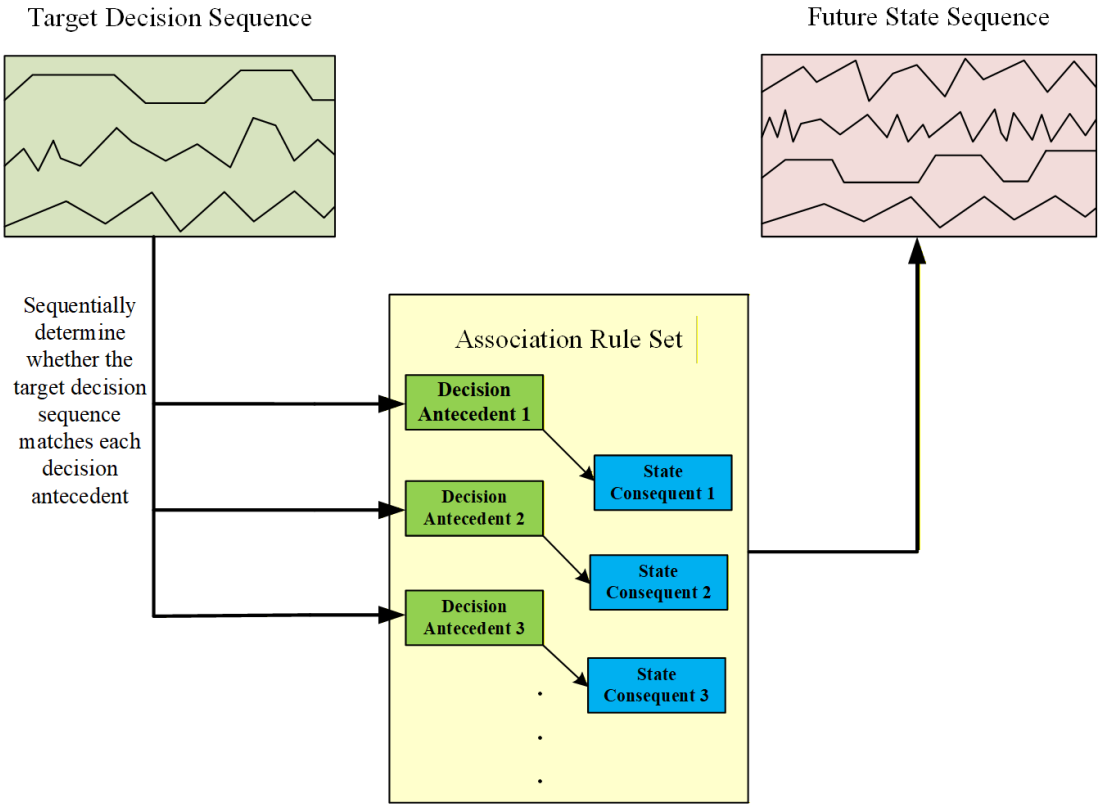

Figure 1. Schematic diagram of the association-rule-based prediction process

4. Experiments

4.1. Data and Preprocessing

The experimental data were collected from 3,391 pot controllers in a large aluminum electrolysis plant from January 1, 2022 to December 30, 2023. The dataset contains ten process dimensions, including temperature, liquid levels, voltage-related variables, feeding behavior, fluoride-addition behavior, and vibration signals. Table 3 summarizes the meaning of each dimension. The dataset was split into a motif mining set and a test set with a ratio of 7:3.

Table 3. Meanings of dimensions in the dataset

Dimension name	Dimension meaning
Electrolysis Temperature	Temperature of the electrolyte, measured in degrees Celsius.
Electrolyte Level	Height of the electrolyte liquid surface, measured in centimeters.
Aluminum Level	Height of the liquid aluminum surface, measured in centimeters.
Voltage Swing	Amplitude of voltage deviation from the set value, measured in millivolts.
Number of Fluoride Additions	Number of operations that add fluoride salt to the cell.
Set Voltage	Manually set voltage reference value, measured in volts.
Working-Set Voltage Difference	Difference between actual operating voltage and manually set voltage, measured in volts.
Feed Amount	Mass of added materials such as alumina, measured in kilograms.
Anode Rod Vibration	Amplitude of voltage fluctuation between the anode rod and electrolyte, measured in millivolts.
Average Voltage	Average operating voltage of the cell, measured in volts.

Missing values were handled by a stepwise procedure. First, variables with more than 20 percent missing values were removed. For the remaining variables, missing entries were filled using forward-fill and backward-fill methods within each individual pot time series, applied sequentially to maintain temporal consistency. To combine data from multiple cells, each cell time series was aligned by absolute timestamp. Gaps or resets caused by different operational periods were recorded as logical interval points. After motif mining, motif elements that crossed these interval points were removed to avoid artificial patterns caused by concatenating unrelated time segments.

4.2. Experimental Settings

The main parameters are listed in Table 4. The subsequence length was set to 14 time steps to capture short-term state evolution while maintaining sufficient temporal resolution. The maximum lag was set to 7 to cover common response delays without including overly distant relationships. The confidence threshold was set to 0.3 to retain moderately frequent rules and remove weak associations. The motif search upper limit was set to 250,000 to control computational cost during motif mining.

The method was compared with representative forecasting baselines. For all benchmark models, the input length and prediction duration were both set to 14 time steps, which is the same temporal scale as the motif subsequence. All experiments were repeated 100 times, and the average E value was calculated. Because this task emphasizes state-pattern prediction rather than point-wise numerical forecasting, the relative pattern-level metric E was used throughout the evaluation.

Table 4. Experimental parameters

Parameter	Value	Explanation
Subsequence length L	14	Captures short-term evolution patterns of aluminum electrolysis states.
Maximum lag time	7	Allows delayed operation-response matching within a controlled window.
Confidence threshold	0.3	Filters association rules while maintaining coverage.
Motif search upper limit k	250,000	Limits computational cost during motif mining.

5. Results and Discussion

5.1. Motif Groups and Temporal Regularities

After separating the original variables into decision and state datasets, motif groups were mined from each part. The results in Table 5 show that the decision dataset contains 13 motif groups and 60 motifs, whereas the state dataset contains 467 motif groups and 4,581 motifs. This difference is consistent with the industrial meaning of the variables. Decision variables are manually adjusted and therefore contain fewer repeated control patterns. State variables are automatically monitored and respond to many coupled factors, producing richer temporal fluctuations.

Table 5. Statistics of mined motif groups in the decision and state datasets

Dataset	Number of motif groups	Total number of motifs
Decision dataset	13	60
State dataset	467	4,581

Figure 2 presents a representative five-dimensional state motif. The upper curves describe a recurrent multivariate state-evolution pattern involving Electrolysis Temperature, Electrolyte Level, Aluminum Level, Working-Set Voltage Difference, and Average Voltage. The synchronized decline near the end of the subsequence suggests a coordinated state change rather than an isolated fluctuation. The lower histograms show that this motif appears more frequently on Sundays and in August. Therefore, the motif is interpretable at both the process level and the temporal level.

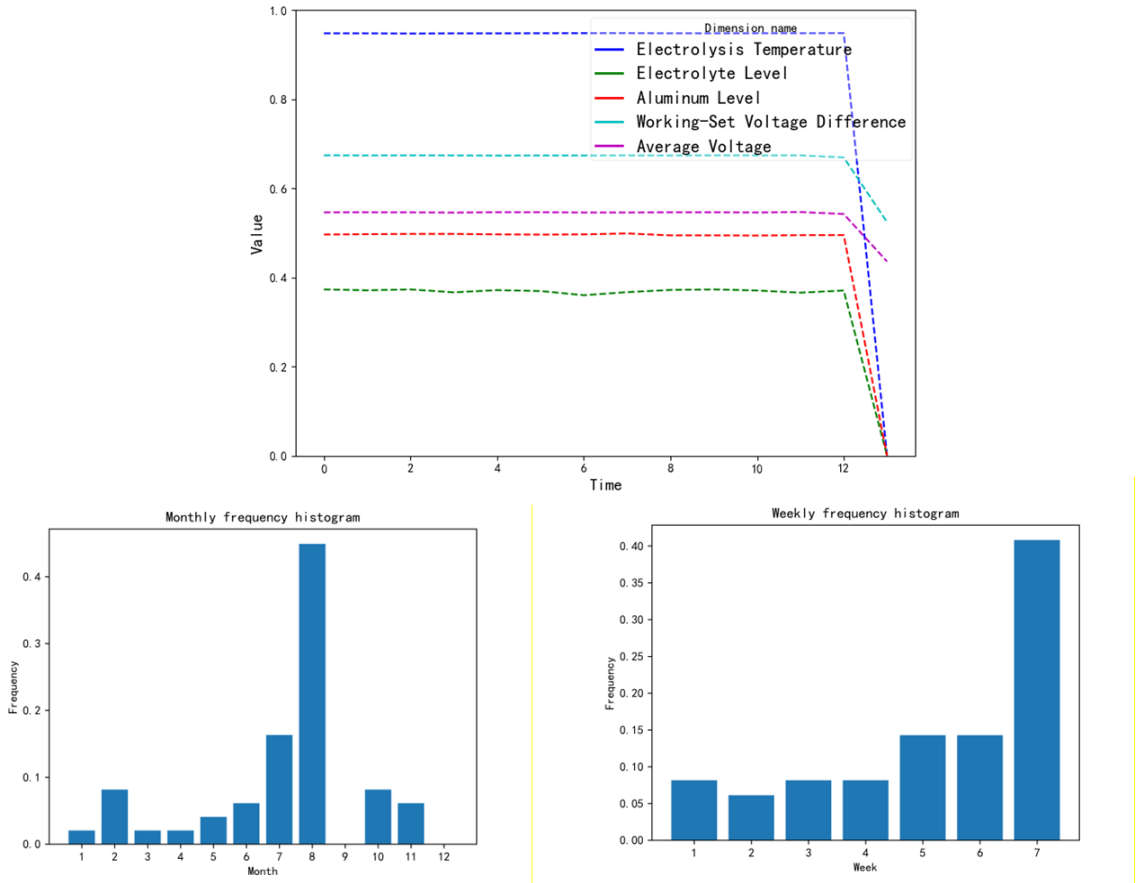

Figure 2. Representative sequence and occurrence frequency histogram of a five-dimensional state motif

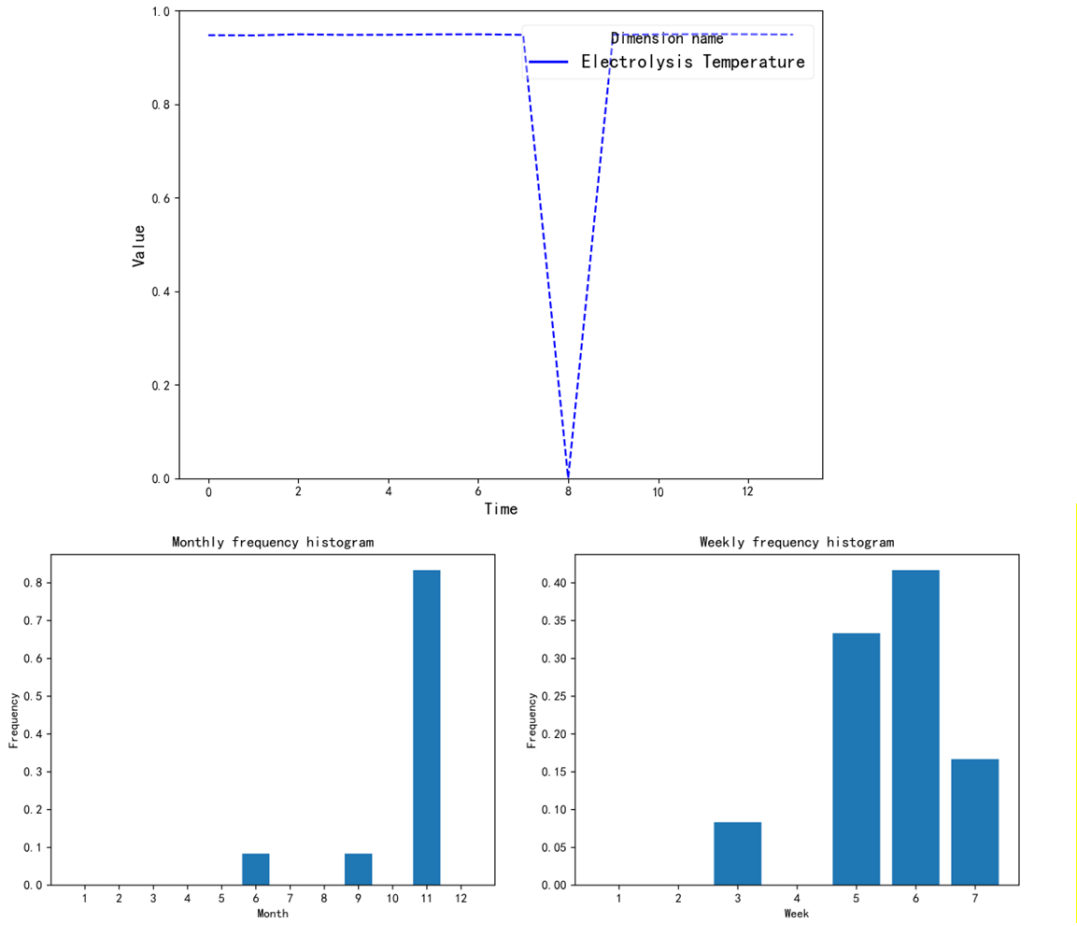

Figure 3. Representative sequence and occurrence frequency histogram of a one-dimensional state motif

Figure 3 shows a representative one-dimensional state motif. This motif focuses on the local evolution of

Electrolysis Temperature and captures a sudden drop followed by rapid recovery. The pattern may correspond to a short-term process disturbance followed by compensation or adjustment. The histograms indicate that this motif occurs frequently on Saturdays and in November. This example demonstrates that even a single-variable motif can carry interpretable temporal information.

The temporal rules for the representative five-dimensional motif are summarized in Table 6. The motif occurs most frequently on Sunday with a frequency of 0.4082 and in August with a frequency of 0.4490. These values mean that 40.82 percent of all occurrences of this motif appear on Sunday and 44.90 percent appear in August. Such rules provide compact prior information that can be used when association rules are unavailable.

Table 6. Temporal rules for a representative five-dimensional motif

Time dimension	Most frequent value	Frequency
Weekday	Sunday	0.4082
Month	August	0.4490

5.2. Association Rules and Industrial Interpretation

Association rules describe how decision motifs are followed by state motifs. The number of mined rules under different maximum lag settings is shown in Figure 4. As the maximum lag increases, more decision-state pairs can be matched, and more rules are retained. However, a larger lag window can also introduce indirect relationships that are harder to explain. Therefore, the chosen maximum lag should balance rule coverage and interpretability. In this experiment, the maximum lag was set to half of the subsequence length.

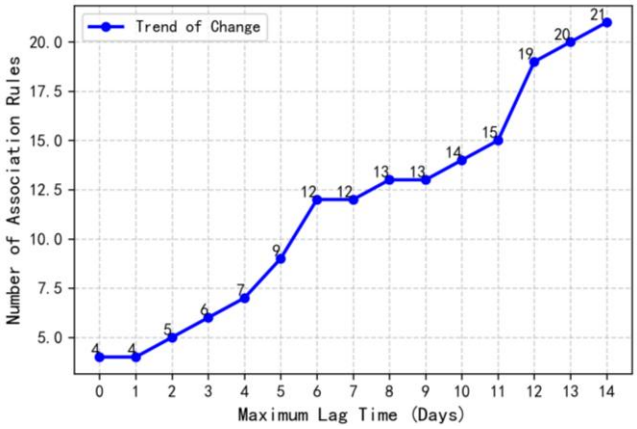

Figure 4. Number of association rules mined under different maximum lag settings

Table 7 lists statistical information for representative association rules. The lag time is the most frequent delay between the decision motif and the corresponding state motif. A lag time of 0 indicates that the state motif appears almost immediately after the decision motif, while a positive lag indicates a delayed process response. Confidence quantifies the proportion of decision motif occurrences that are followed by the corresponding state motif within the lag window. For example, Rule 5 has a confidence of 1.00 and a lag time of 3 days, meaning that all observed occurrences of this decision motif were followed by the corresponding state motif after a typical delay of 3 days.

Table 7. Statistical information of representative association rules

Index	Lag time (days)	Confidence
1	0	0.33
2	5	0.33
3	0	0.33
4	6	0.40
5	3	1.00
6	5	0.33
7	6	1.00
8	4	0.40
9	0	0.33
10	2	0.33

Figure 5 shows a representative association rule with index 1. In the antecedent decision motif, the Number of Fluoride Additions and Feed Amount drop sharply to zero at the beginning of the subsequence, indicating that regular feeding and fluoride-addition operations were temporarily interrupted. In the consequent state motif, Electrolysis Temperature, Aluminum Level, and Anode Rod Vibration also decrease significantly. From a process-mechanism perspective, the sudden stop of feeding may reduce alumina supply and disturb the bath-metal thermal balance; if it continues, the cell can show a cold-cell tendency, unstable alumina concentration, or early risk features related to anode-effect development. Therefore, the rule is not only a curve-shape match, but also an interpretable warning pattern that links intervention stagnation with subsequent state deterioration. Later, feeding behavior recovers, and the state variables rise and stabilize. This recovery stage is consistent with gradual restoration of material feeding and thermal compensation, so the rule links interruption-like decision behavior to a recovery-type state response.

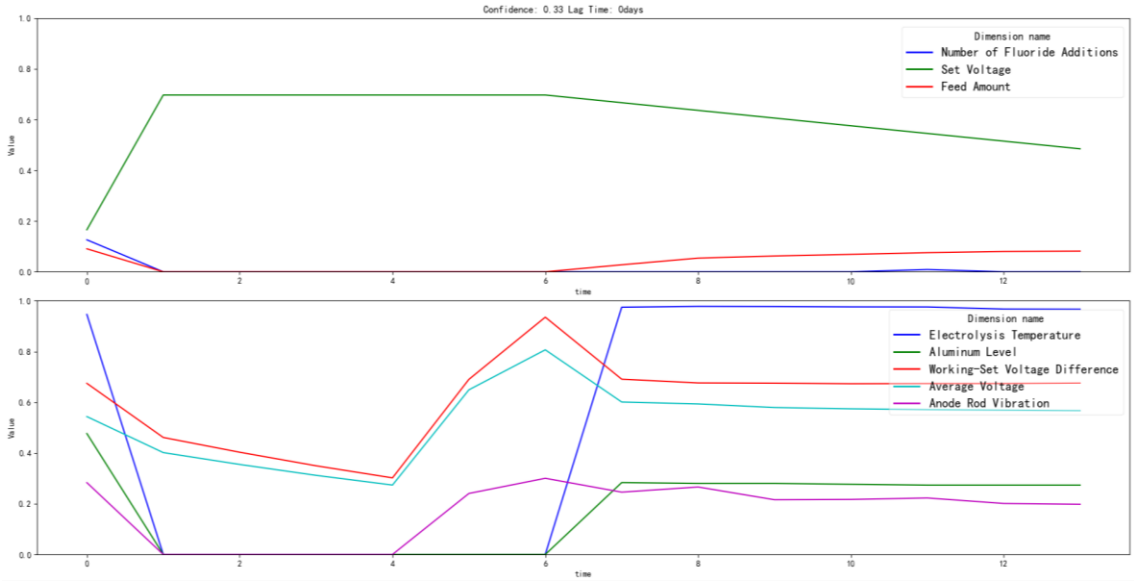

Figure 5. Representative association rule with index 1

Figure 6 shows another representative rule with high confidence. In this rule, Set Voltage increases, while the Number of Fluoride Additions and Feed Amount are reduced to zero shortly afterwards. The state variables remain relatively stable at first and then show a concentrated decline after a short delay. This is industrially meaningful because the cell may temporarily maintain apparent stability using existing thermal and material reserves, but the lack of continuous feeding can later amplify imbalance in bath temperature, metal level, and voltage-related variables. The three-day lag of Rule 5 therefore helps operators understand why a state deviation may occur several days after a strong operational adjustment, rather than immediately at the time of intervention. Such a rule should be interpreted as process-oriented evidence for risk screening, and final diagnosis should still be combined with plant logs and operator records.

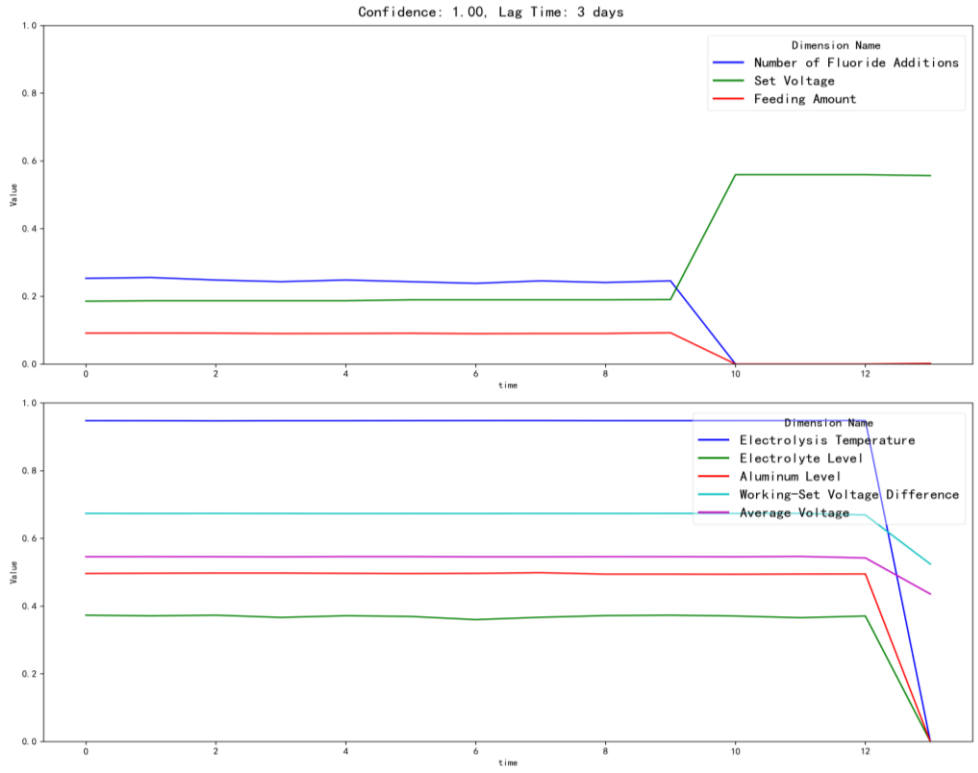

Figure 6. Representative association rule with index 5

Overall, the association-rule results indicate that the proposed method is not limited to statistical pattern matching. Each rule has an industrial interpretation: the antecedent explains what operational behavior occurred, the consequent explains what process state followed, the lag explains when the response typically appeared, and the confidence quantifies how reliable the relationship is. This makes the rules suitable for process engineers who need traceable decision evidence.

5.3. Prediction Performance

Three rule-based prediction strategies were compared first: association rules only, temporal rules only, and the combined rule-fusion strategy. The results are shown in Table 8. All three strategies achieved E values below 1, which means that they outperformed the random benchmark. The association-rule-only strategy achieved an E value of 0.78, showing that decision-state operation-response rules are effective for state

prediction. The temporal-rule-only strategy achieved an E value of 0.83, indicating that occurrence time also contains useful information. The combined strategy achieved the lowest error value of 0.75, which confirms that association rules and temporal rules are complementary.

Table 8. Comparison of rule-based prediction strategies

Prediction strategy	E value
Association rules only	0.78
Temporal rules only	0.83
Association + temporal rules	0.75

The proposed method was also compared with representative deep learning forecasting baselines. Table 9 shows that the baseline models produced E values close to 1 under the pattern-level evaluation metric. TimesNet, Autoformer, and Informer achieved E values of 0.99, FEDformer achieved 0.98, and DLinear and LightTS achieved 0.96. These values indicate that the baselines were only slightly better than the random benchmark for the motif-level prediction task. In contrast, the proposed method achieved an E value of 0.75, indicating a clearer advantage in predicting recurring state-evolution patterns.

Table 9. Prediction effectiveness of different models

Model	E value
TimesNet	0.99
Autoformer	0.99
DLinear	0.96
FEDformer	0.98
Informer	0.99
LightTS	0.96
Proposed method	0.75

The improvement can be explained by the task definition. The proposed method predicts patterns rather than isolated future values. Motif mining captures typical local evolution shapes, association rules model delayed operation-response mechanisms, and temporal rules describe when certain state motifs are likely to occur. This combination is better aligned with aluminum electrolysis process behavior than point-wise numerical forecasting alone. Moreover, the proposed method provides interpretability by showing which rule was matched, which state motif was selected, and how confidence and lag supported the prediction.

The results demonstrate three practical advantages. First, parameter decoupling improves the semantic clarity of discovered motifs by separating control actions from system responses. Second, the rule-fusion strategy produces better prediction performance than either association rules or temporal rules alone. Third, compared with black-box forecasting models, the method gives explicit evidence linking manual interventions to later state evolution. These advantages are valuable for industrial process monitoring, diagnosis, and operator-oriented decision support.

6. Conclusion

This paper presented an interpretable motif-based rule-fusion method for aluminum electrolysis cell state prediction. By separating variables into decision and state datasets, mining motif groups, and integrating temporal rules with association rules, the method transforms state prediction from black-box numerical forecasting into explicit rule-based inference. Experiments on large-scale industrial data demonstrated that the proposed method achieved better pattern-level prediction performance than representative forecasting baselines and that combining association rules with temporal rules produced the best overall result.

The method also has practical value for aluminum electrolysis production. Mined association rules link interventions such as changes in Set Voltage, Feed Amount, and Number of Fluoride Additions to subsequent state changes such as Electrolysis Temperature, Aluminum Level, Average Voltage, and Anode Rod Vibration. These rules provide process engineers with interpretable evidence for understanding delayed operation-response relationships. Future work may extend the method toward online monitoring and closed-loop decision support by dynamically updating rule sets or combining rule-based prediction with downstream optimization modules.

Acknowledgments

This work was supported by the Ministry of Education Humanities and Social Science Research Project [grant number 22YJA870003]; the Beijing Natural Science Foundation [grant number L242153]; the 2024 Graduate Education Reform Project of North China University of Technology [grant number YJS2024JG02]; and the Organized Research Project of NCUT [grant number 2024NCUTYXCX110].

References

[1] Sun, Y., Chen, X., Cen, L., et al. (2024). A dynamic graph structure identification method of spatio-temporal correlation in an aluminum electrolysis cell. Applied Soft Computing, 157, 111536. https://doi.org/10.1016/j.asoc.2024.111536

[2] Zhou, L., Yao, Z., Sun, K., et al. (2024). Methodological review of methods and technology for utilization of spent carbon cathode in aluminum electrolysis. Energies, 17(19), 4866. https://doi.org/10.3390/en17194866

[3] Wang, J., Xie, Y., Xie, S., et al. (2024). Development of data-knowledge-driven predictive model and multi-objective optimization for intelligent optimal control of aluminum electrolysis process. Engineering Applications of Artificial Intelligence, 134, 108664. https://doi.org/10.1016/j.engappai.2024.108664

[4] Cao, D., & Lin, Z. (2024). Multidimensional time series motif group discovery based on matrix profile. Knowledge-Based Systems, 304, 112509. https://doi.org/10.1016/j.knosys.2024.112509

[5] Lin, J., Keogh, E., Lonardi, S., & Chiu, B. (2002). Finding motifs in time series. In Proceedings of the 2nd Workshop on Temporal Data Mining (pp. 23–26).

[6] Pradhan, G. N., & Prabhakaran, B. (2017). Association rule mining in multiple, multidimensional time series medical data. Journal of Healthcare Informatics Research, 1(1), 92–118. https://doi.org/10.1007/s41666-017-0008-6

[7] Yi, J., Bai, J., Zhou, W., et al. (2018). Operating parameters optimization for the aluminum electrolysis process using an improved quantum-behaved particle swarm algorithm. IEEE Transactions on Industrial Informatics, 14(8), 3405–3415. https://doi.org/10.1109/TII.2018.2820099

[8] Wang, J., Xie, Y., Xie, S., et al. (2024). Operational decision-making optimization of aluminum electrolysis process based on health evaluation of aluminum electrolytic cell. In 2024 IEEE International Conference on Cybernetics and Intelligent

Systems and IEEE International Conference on Robotics, Automation and Mechatronics (pp. 156–161). https://doi.org/10.1109/CIS-RAM61216.2024.10634521

[9] Wang, J., Xie, S., Xie, Y., et al. (2024). A general knowledge-guided framework based on deep probabilistic network for enhancing industrial process modeling. IEEE Transactions on Industrial Informatics, 20(3), 3050–3059. https://doi.org/10.1109/TII.2023.3338896

[10] Berndt, D. J., & Clifford, J. (1994). Using dynamic time warping to find patterns in time series. In Proceedings of the 3rd International Conference on Knowledge Discovery and Data Mining (pp. 359–370).

[11] Keogh, E., Chakrabarti, K., Pazzani, M., & Mehrotra, S. (2001). Locally adaptive dimensionality reduction for indexing large time series databases. In Proceedings of the 2001 ACM SIGMOD International Conference on Management of Data (pp. 151–162). https://doi.org/10.1145/375663.375679

[12] Agrawal, R., Imieliński, T., & Swami, A. (1993). Mining association rules between sets of items in large databases. In Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data (pp. 207–216). https://doi.org/10.1145/170035.170072

[13] Zaki, M. J. (2001). SPADE: An efficient algorithm for mining frequent sequences. Machine Learning, 42(1), 31–60. https://doi.org/10.1023/A:1007638616156

[14] Pei, J., Han, J., Mortazavi-Asl, B., et al. (2001). PrefixSpan: Mining sequential patterns efficiently by prefix-projected pattern growth. In Proceedings of the 17th International Conference on Data Engineering (pp. 215–224). https://doi.org/10.1109/ICDE.2001.914830

[15] Srikant, R., & Agrawal, R. (1996). Mining sequential patterns: Generalizations and performance improvements. In International Conference on Extending Database Technology (pp. 1–17). https://doi.org/10.1007/BFb0054687

[16] Das, G., Lin, K. I., Mannila, H., et al. (1998). Rule discovery from time series. In Proceedings of the Fourth International Conference on Knowledge Discovery and Data Mining (Vol. 98, pp. 16–22).

[17] Harms, S. K., Deogun, J., & Tadesse, T. (2002). Discovering sequential association rules with constraints and time lags in multiple sequences. In International Symposium on Methodologies for Intelligent Systems (pp. 432–441). https://doi.org/10.1007/3-540-45813-1_41

[18] Hoeppner, F. (2002). Learning dependencies in multivariate time series. In Proceedings of the ECAI (Vol. 2, pp. 25–31).

[19] Sacchi, L., Larizza, C., Combi, C., & Pozzi, G. (2007). Data mining with temporal abstractions: Learning rules from time series. Data Mining and Knowledge Discovery, 15(2), 217–247. https://doi.org/10.1007/s10618-007-0066-7

[20] Wu, H., Hu, T., Liu, Y., et al. (2023). TimesNet: Temporal 2D-variation modeling for general time series analysis. In International Conference on Learning Representations.

[21] Wu, H., Xu, J., Wang, J., et al. (2021). Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. Advances in Neural Information Processing Systems, 34, 22419–22430.

[22] Zeng, A., Chen, M., Zhang, L., & Xu, Q. (2023). Are transformers effective for time series forecasting? In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 37, No. 9, pp. 11121–11128). https://doi.org/10.1609/aaai.v37i9.26333

[23] Zhou, T., Ma, Z., Wen, Q., et al. (2022). FEDformer: Frequency enhanced decomposed transformer for long-term series forecasting. In International Conference on Machine Learning (pp. 27268–27286).

[24] Zhou, H., Zhang, S., Peng, J., et al. (2021). Informer: Beyond efficient transformer for long sequence time-series forecasting. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 35, No. 12, pp. 11106–11115). https://doi.org/10.1609/aaai.v35i12.17325

[25] Campos, D., Zhang, M., Yang, B., et al. (2023). LightTS: Lightweight time series classification with adaptive ensemble distillation. Proceedings of the ACM on Management of Data, 1(2), 1–27. https://doi.org/10.1145/3588772.3594269

[26] Van Wesenbeeck, D., Yurtman, A., Meert, W., & Van den Berghe, S. (2024). LoCoMotif: Discovering time-warped motifs in time series. Data Mining and Knowledge Discovery, 38(4), 2276–2305. https://doi.org/10.1007/s10618-024-01030-9

[27] Tang, C., Xu, L., Yang, B., et al. (2023). GRU-based interpretable multivariate time series anomaly detection in industrial control system. Computers & Security, 127, 103094. https://doi.org/10.1016/j.cose.2023.103094

[28] Rizzi, W., Comuzzi, M., Di Francescomarino, C., & Maggi, F. M. (2024). Explainable predictive process monitoring: A user evaluation. Process Science, 1(1), 3. https://doi.org/10.1007/s44388-024-00003-8