

基于端到端自监督时序对比学习与异常合成的实时智能电表异常检测框架

王亿鑫¹, 梁高琪¹, 毕霁超², 赵俊华³

(1. 哈尔滨工业大学(深圳)机器人与先进制造学院, 广东 深圳 518055; 2. 浙江省工业和信息化研究院, 杭州 310006;
3. 香港中文大学(深圳)理工学院, 广东 深圳 518172)

摘要: 物联网(Internet of Things, IoT)技术正在通过快速集成和部署智能电表重塑全球能源格局, 这些电表支持高分辨率的能耗监测、双向通信以及先进的计量基础设施服务。然而, 这一数字化转型也使电力系统面临不断演变的威胁, 涵盖网络入侵、电力盗窃以及设备故障等。这些异常行为的不可预测性, 加之标注异常数据的稀缺, 使实时检测异常变得尤为困难。为应对这些挑战, 提出了一种用于智能电表异常检测的实时决策支持框架, 该框架基于滑动时间窗口以及两个自监督对比学习模块。第一个模块通过合成多样化的异常样本来弥补标注的异常数据不足, 第二个模块则捕捉内在的时间模式以增强上下文区分能力。提出的端到端框架可持续地利用滚动更新的电表数据实时更新模型, 及时识别电网中不断演化出现的新兴异常行为。在 8 个公开可用的智能电表数据集上针对 7 种不同的异常模式进行了广泛评估, 结果表明所提的完整框架表现优秀, 平均召回率和 F1 分数均超过 0.85。

关键词: 异常检测; 物理信息安全; 异常合成; 对比学习; 时间序列

Real-Time Smart Meter Abnormality Detection Framework via End-to-End Self-Supervised Time-Series Contrastive Learning with Anomaly Synthesis

WANG Yixin¹, LIANG Gaoqi¹, BI Jichao², ZHAO Junhua³

(1. School of Robotics and Advanced Manufacture, Harbin Institute of Technology, Shenzhen, Guangdong 518055, China;
2. Zhejiang Institute of Industry and Information Technology, Hangzhou 310006, China; 3. School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen, Guangdong 518172, China)

Abstract: The rapid integration of Internet of Things (IoT) technologies is reshaping the global energy landscape by deploying smart meters that enable high-resolution consumption monitoring, two-way communication, and advanced metering infrastructure services. However, this digital transformation also exposes power system to evolving threats, ranging from cyber intrusions and electricity theft to device malfunctions, and the unpredictable nature of these anomalies, coupled with the scarcity of labeled fault data, makes real-time detection exceptionally challenging. To address these difficulties, a real-time decision support framework is presented for smart meter anomaly detection that leverages rolling time windows and two self-supervised contrastive learning modules. The first module synthesizes diverse negative samples to overcome the lack of labeled anomalies, while the second captures intrinsic temporal patterns for enhanced contextual discrimination. The end-to-end framework continuously updates its model with rolling updated meter data to deliver timely identification of emerging abnormal behaviors in evolving grids. Extensive evaluations on eight publicly available smart meter datasets over seven diverse abnormal patterns testing demonstrate the effectiveness of the proposed full framework, achieving average recall and F1 score of more than 0.85.

Key words: abnormality detection; cyber-physical security; anomaly synthesis; contrastive learning; time-series

基金项目: 深圳市自然科学基金稳定支持面上项目(GXWD20231128112434001); 浙江省自然科学基金探索青年项目(LQ24F030015)。

Foundation item: Supported by the Stable Support General Project of Shenzhen Natural Science Fund (GXWD20231128112434001); Zhejiang Provincial Natural Science Foundation of China under Grant(LQ24F030015).

0 Introduction

The global energy sector is rapidly evolving into an intelligent ecosystem to address widespread resource constraints. At the heart of this transformation, Internet of Things (IoT) technologies drive energy infrastructure upgrades, with smart meters serving as the pivotal nexus between consumers and the grid [1]. Smart meters deliver high-resolution consumption and equipment data and enable two-way communication, empowering online monitoring, demand response, and dynamic pricing [2–4]. These devices underpin advanced metering infrastructures (AMI) that bolster grid reliability, lower operational costs, and accelerate the shift toward a resilient, low-carbon energy ecosystem.

Smart meters provide granular visibility into energy usage, system health, and maintenance needs, etc., which directly reflect to the status and behaviors of the linked systems. Meanwhile, their widespread deployment also introduces new challenges and risks of cyber-physical security, such as vulnerabilities to cyberattacks and incidents of electricity theft [5–7], meter parametrization error and faulty. Real-time detection of these excepted events remains a critical yet complex task. One major challenge is the unpredictability of anomaly patterns, which even evolve with changing attack strategies just like the cyberattacks. Additionally, the scarcity of labeled anomaly data hampers on-the-fly detection capabilities, as manual annotation is costly and impractical for deployment in large-scale application. In this paper, the smart meter anomaly is generally defined as the any unpredictable data deviations from the expected patterns, which directly manifests the abnormal status and behaviors of the linked systems, whether due to cyberattacks, electricity theft, device malfunctions, or other faults. Rapid detection of these smart meter anomalies enables us to pinpoint potential threats and exert interventions promptly of cyber-physical security, thereby preventing further safety, security and economic loss.

Extant detection techniques majorly including measurement-based and machine learning-based approaches. Measurement-based methods explicitly model the statistical characteristics of smart meter data and estimate its coefficients by an optimization algorithm, and ultimately infer the anomalous state by compute the statistical property of the read data. Capozzoli et al. proposed a symbolic aggregate approximation process based for the characterization of energy time series [8]. Yip et al. describeg the smart meters using various state functions and calculateg the discrepancy in meter reading at each time slot [9–10]. Chen et al. designed an error estimation using the truncated singular value decomposition regularization with L-curve optimization, which will output the top N% smart meters in error for onsite calibration [11]. Measurement-based method relies heavily on manually engineered statistical features; however, as user behaviors and attack patterns evolve, these fixed features struggle to adapt to changing data distributions and can thus become ineffective.

Machine learning methods have been widely used in anomaly detection in recent years, including supervised and unsupervised learning. Supervised learning paradigm typically involves to train a prediction model based on given training dataset. Buzau et al. combined all the information of smart meters record and designeg a workflow of feature extraction, distance and density measurement and XGBoost for outlier detection [12]. Takiddin et al. developed an attention-enhanced autoencoder that trained on benign energy consumption data and detected abnormality by assessing the deviation from its learned pattern [13]. Roelofs et al. investigated transfer learning models for wind turbines anomaly and fine-tuning on small abnormous data with the abnormal score being computed based on the root mean square error [14]. Obviously, supervised learning-based anomaly detection models depends on prediction model and residual analysis between the predicted value and real readings.

On the other hand, unsupervised learning relies on collected dataset, which is supposed to include enough abnormal samples to support model training. Lei et al. designed a heuristic optimization-enhanced unsupervised clustering algorithm to evaluate the energy consumption based on annual electricity consumption data of an experimental building [15]. Yin et al. employed a fuzzy C-means to cluster energy consumption into subspaces and connectivity-based outlier factor and mean nearest neighbor distance anomaly factor are applied to the subspace to judge the point anomaly and collective anomaly [16]. Zhang et al. proposed a semi-supervised framework that uses feature augmentation on unlabeled data to generate pseudo-labels, which are then combined with labeled samples to train a predictive model [17]. Unsupervised models perform well when anomalous instances are plentiful, but developing robust detectors for scenarios with scarce and diverse abnormal behaviors remains a critical challenge.

Contrastive learning is a self-supervised paradigm in which a model is trained to map semantically related data pairs close together in embedding space while pushing apart unrelated pairs, thereby discovering meaningful features without manual labels [18–20]. By defining positives (e. g., samples in the same category, different augmentations of the same instance) and negatives (e. g., samples in the different category, different instances), it shapes representations that capture underlying structure and invariances. Chen et al. applied contrastive learning to capture both high- and low-frequency samples of air-conditioning loads, thereby unsupervisedly enhancing smart meter feature representations [21]. Gao et al. first transformed smart meter readings using the Gramian Angular Field and paired these encoded matrices with the original data for contrastive learning; anomalies were subsequently identified by applying the DBSCAN clustering algorithm to the learned representations [22]. These studies highlight the promise of contrastive learning when negative samples are scarce.

tive learning when negative samples are scarce.

The review reveals that, despite extensive research on smart meter anomaly detection, few methods account for the evolving and scarce nature of abnormal behaviors, potentially allowing unseen anomaly to slip through undetected. To address evolving anomaly in smart meters, we propose a real-time and general detection framework that leverages rolling time windows and contrastive learning with anomaly synthesis to capture dynamic smart meters patterns. We assume that temporally adjacent windows exhibit similar patterns and thus form positive sample pairs. Simultaneously, we synthesize negative samples representing anomalous behaviors from the original windows. Several sequence abnormality synthesis operators are introduced to simulate the nuance between the normal pattern and abnormal ones. By rolling these time windows forward, our framework continually updates the neural network model and performs end-to-end real-time anomaly detection. This work contributions in three aspects:

- 1) We introduce anomaly-synthesis operators to generate generic smart meter abnormalities, effectively mitigating the scarcity of negative samples in real-world data.
- 2) We employ a self-supervised, contrastive-learning paradigm to derive robust representations of smart meter data, thereby enabling the detection of dynamic anomaly.
- 3) This paper proposes an end-to-end decision support framework that ingests rolling time-window data for online detection and periodically updates the neural network model.

1 Problem Statement

Consider smart meters abnormality detection as a decision support system (DSS):

$$F_{\text{DSS}} = \langle O, f_{\theta}, R, A \rangle \quad (1)$$

where F_{DSS} indicates a decision support system and O is the observations space recorded by smart

meters from their environments, f_θ represents representation function that parameters θ used by this system, that maps the observed state into the latent space. R indicates the decision-making function whose output will prescribe the decisions space A implying whether the system is under attack $a_t = 1$ or not $a_t = 0$ at t . Let $x_t \in \mathbb{R}^d$ be the readings at time t by smart meters, such as electricity consumption, temperature, where $d = 1$ for univariate and $d > 1$ for multivariate. Normally, the transition dynamics of smart meters $P_{\text{normal}}(x_{t+1}|x_t)$ is driven by environment factors, i. e., consumer habits. However, in an unusual circumstance, the state probability distribution $P_{\text{abnormal}}(x_{t+1}|x_t)$ imposed by the abnormal events such as cyberattacks and equipment malfunctions. Thus, the environment dynamics of smart meters is given as:

$$x_{t+1} \sim \begin{cases} P_{\text{normal}}(x_{t+1}|x_t) \\ P_{\text{abnormal}}(x_{t+1}|x_t) \end{cases} \quad (2)$$

In our smart meters abnormality detection, the observations space $O = \{x_t\}_{t-w:t}$ is defined as a contiguous span from $t-w$ to t of meter readings. At each time step t , the observed state $o_t \in O$, as known as an observed sample in this paper, will be input to the representation function f_θ that will get the best representation in latent space, whose output is subsequently evaluated by the decision-making function R . Based on the resulting score, the system can

determine and output the behavioral classification of the smart meter. For $t = 1, 2, \dots, T = \infty$, through a rolling screening mechanism, the system continuously monitors and identifies abnormal behaviors in real time by:

$$a_t = R(f_\theta(o_t)) \quad (3)$$

where $o_t \in O$ and $a_t \in A$. By leveraging this smart meter abnormality detection system, utility personnel can precisely manage smart meters, promptly detect anomalies such as electricity theft and irregular consumption, and thereby enhance operational efficiency and system security.

2 Methodology

Inspired by recent the self-supervised learning paradigm of contrastive learning, we devise a decision support framework based on end-to-end contrastive learning for smart meters abnormality as shown in Fig. 1. In our framework, we assume that there are not any anomaly labels for training, which is common in real-world application. Therefore, we introduce pseudo-labels by anomalies synthesis to provide general supervised signals for anomaly-aware representation learning (Sec. 2.1), and followed by context-aware representation learning (Sec. 2.2) driven by similarity measures. The rolling time window are fed into the trained model for anomaly inference (Sec. 2.3) when new data is obtained in real time. Thanks to IoT communication technolo-




Fig. 1 Proposed smart meters abnormality detection framework

gies, we can easily collect and access time series data from different smart meters. We firstly segment the time series data into batches $\{o_i\}$ using predefined stride l and time window length w for representation model training. Before training, necessary data preprocessing is conducted, including missing values filling with zeros and normalization as:

$$x_t = \frac{x_t - x_{\min}}{x_{\max} - x_{\min}} \quad (4)$$

where $x_{\max}$ and $x_{\min}$ stand for the maximum and minimum values in the training data. For the sake of performance and efficiency, we employ a 1D-ResNet with four ResNet blocks shown in Fig. 2 for time-series representation, which has been exhibiting strong performance across different time series tasks [23–24]. Each block contains three sequential convolution layers with various convolutional kernel size of 7, 5 and 3, offering multi-scale feature extraction, which is beneficial for capturing different types of patterns in the input data. To get the next layer's activations, the current activations $o_i^{(i)}$ of layer i are passed through a learnable mapping f_θ , added back into the input, and then a activation function Relu is applied. When this procedure is repeated for $i = 0$ through 3, a four-layer ResNet is produced as:

$$o_i^{(i+1)} = \text{Relu}(f_\theta(o_i^{(i)}) + o_i^{(i)}) \quad i = 0, 1, 2, 3 \quad (5)$$




Fig. 2 1D-ResNet as feed-forward network backbone

2.1 Anomaly Contrastive Learning

At this stage, relying solely on unlabeled training data, we aim to train the model to discriminate between normal and anomalous samples, where a time window of sequence data is defined as a sample and the original samples are seen as the expected data trend in normal operating environments. However, the real-world abnormality typically cannot be known

in advanced, which means that modeling a certain anomaly is arduous. Thus, here we utilize the frequent time-series anomaly types that is commonly adopted to reflect the most general time series abnormal behaviors [25–26]. Through learning the normal pattern of smart meters and comparing with the abnormal ones, the model can identify the potential abnormality. For sure that incorporating certain knowledge of some anomaly types in the real operational environment of smart meters can enhance the performance of tailored models to that anomaly type. In this paper, we adopt five typical anomaly synthesis operators as depicted in Fig. 3 to generate negative samples for training, including: 1) Point outliers, which are individual data points that deviate sharply from the rest of the series; 2) Seasonal anomalies, which occur when observations disrupt the regular, recurring seasonal patterns in the data; 3) Contextual outliers, which are values that appear normal globally but are anomalous within a specific temporal context; 4) Trend anomalies, which arise when the long-term trajectory of the time series is unexpectedly altered or reversed; and 5) Shapelet anomalies, which are atypical subsequences whose shapes significantly deviate from the expected patterns. For each original sample in $\{o_i\}$, a synthetic abnormal sample called as negative sample will be created by randomly selecting an anomaly synthesis operator in 1) —5). Each operator targets one specific signal attribute and the parameter ranges for each operator (e. g., drift magnitude for trend and window size for shape) were determined randomly so that no two operators produce the exact same type of perturbation in our implementation.

For positive samples, we assume a degree of temporal ‘inertia’ in the system state, meaning that the states within a few neighboring time steps tend to exhibit similarity and continuity [27]. For example, with a residential smart meter, electricity consumption typically dips at night due to reduced household activity, while peaks around dinner time as cooking and lighting increase. As a result, the




Fig. 3 Different types of anomaly synthesis

smart meter readings in the same period exhibit a natural consistency, yet reveal diverse trends when comparing across different times of day. Subsequently, a sample is randomly selected within the preceding few time steps of original sample to serve as a positive sample of original sample. Finally, a training batch B consisting of the triplets of original, positive, and negative samples (*org-pos-neg*) is constructed.

The mean square error (MSE) loss function is applied to anomaly contrastive learning on B . For an input triplet consisting of the original observation $o = \{x\}_{t-w:t}$, its positive sample p and negative sample n , we compute the MSE loss $L(\theta)$ for the positive and negative pairs as follows:

$$L(\theta) = \frac{1}{|B|} \sum_{(o, p, n) \in B} (f_\theta(o) - f_\theta(p))^2 - \frac{1}{|B|} \sum_{(o, n) \in B} (f_\theta(o) - f_\theta(n))^2 \quad (6)$$

This objective enables the model to distinguish between normal and synthetic samples, thereby learning a latent space representation M^m from original space D^d , i. e., $f_\theta: D^d \rightarrow M^m$ that supports effective context-based contrastive learning.

2.2 Context Contrastive Learning

The anomaly contrastive learning emphasizes the distinction between the introduced negative samples and the original samples. However, the model is still not guide to capture the disparities among original samples themselves, such as the differences between late evening (low and stable) and dinner hour (high and fluctuating) energy consumption data pat-

terns are still not be learned by model. Different from the anomaly contrastive learning process, this stage is to learn a more comprehensive representation of all available samples, including original and synthetic ones. Firstly, we define the sample in latent space M^m of original observations learned through anomaly contrastive learning as an anchor. Then, we computed the pairwise L2 distances between each sample and all others in M . For each anchor a , the top k most similar samples are selected to form positive set U_a that indicates the data windows with the most similar trend with anchor, i. e., the data in adjacent windows. Meanwhile, the bottom k least similar samples are selected to form the negative set V_a implying the least similar data samples, i. e., the anomaly samples and those data in different intervals. With the anchor and the constructed sets of U_a and V_a , the dot product similarity loss is applied to positive pairs and negative pairs to learn another latent space representation G^g , i. e., $f_\theta: M^m \rightarrow G^g$. Due to there are not anomaly labels with the training data, a cluster head with given hyperparameter of clusters number c is integrated to prescribe the cluster label for input samples, which will help the anomaly inference. To increase the model diversity, an entropy regularization $H(f_\theta(a))$ with control coefficient λ is incorporated into the loss function as:

$$J(\theta) = \frac{1}{k|M|} \sum_{a \in M} \left(\sum_{n \in U_a} f_\theta(a)^\top f_\theta(n) - \sum_{p \in V_a} f_\theta(a)^\top f_\theta(p) + \lambda H(f_\theta(a)) \right) \quad (7)$$

$$H(f_\theta(a)) = -\sum \log(f_\theta(a)) f_\theta(a) \quad (8)$$

2.3 Anomaly Inference

Through the last layer of cluster head, we can assign a cluster label for all original samples in B . The above two contrastive learning modules enable the model to distinguish the synthetic abnormal samples from the original sample, thereby will be categorized in different clusters. To label incoming samples as normal or abnormal, we designate any cluster containing more than 80 % positive samples as

a “major cluster”, denoted C that containing more than 80 % positive sample. Then, the smart meter abnormality decision-making support system infers a newly obtained sample as normal or abnormal according to whether its predicted cluster label belongs to C , i. e. ,

$$\text{anomaly}(o_i) = \begin{cases} 0, & \text{if } f_\theta(o_i) \in C \\ 1, & \text{otherwise} \end{cases} \quad (9)$$

3 Experimental Setting

3.1 Data

Based on the above discussion, anomalous behaviors of smart meters in real production environments may be unknown and unlabeled. Therefore, to validate the effectiveness and robustness of our framework, we acquired eight publicly available smart meter datasets [28], which are assumed to be historical data free of any anomalies in this paper. These datasets exhibit diverse load profiles, such as weekday/weekend contrasts, daily peak/off-peak cycles, and varying weekend behaviors, allowing us to test the model’s ability to generalize across usage types. Each dataset contains totally 17 520 records comprising half-hourly readings collected over the course of one year. Our framework is specifically tailored to capture intraday (24 h) and intraweek (7 d) patterns, which are the dominant drivers of “real-time” anomaly detection in smart-meter streams. Then, we employ the first eleven months of data to train our proposed model and randomly inject seven anomaly events (indexed 0–6) into the last month of data for testing.

Seven anomaly events includes: 1) historical minimum reduction, which forces usage down to the lowest value ever recorded; 2) fixed weekly percentage cut, which applies a uniform percentage reduction during one week in a month; 3) progressive reduction, which gradually tapers consumption over the month; 4) threshold capping, which clips all values above a preset maximum to that maximum; 5) peak-hour progressive reduction, which imposes a taper during high-demand hours; 6) peak-hour

reduction and redistribution, which curtails data during peak periods and redistributes the curtailed volume to off-peak times; and 7) peak-off-peak swap, which exchanges consumption volumes one-for-one between peak and off-peak intervals. Each anomaly event enforces a different constraint or redistribution of smart meter readings reflecting the kinds of meter anomalies that might actually occur. During the training process, each smart meter is equipped with its own dedicated detection model.

3.2 Training Settings

Our method is based on the rolling time windows mechanism, i. e. , as the time goes and the new readings arrive, the time window will move forward as well. Then, the model will immediately evaluate whether the latest time window is normal or not. In testing, every window of data constitutes an individual sample; we then compute our performance metrics by comparing the model’s outputs against the ground truth for each window. In this paper, the length of each time window w in O is set to 144 with a rolling stride of 5. The raw observed time series is encoded in 16-dimensional and 8-dimensional latent space via anomaly contrastive learning and context contrastive learning, i. e. , $M^{(16)}$ and $G^{(8)}$. We train the anomaly contrastive learning module for 50 epochs and the context contrastive learning module for 100 epochs, respectively. In the context contrastive learning stage, we selected $k = 10$ the number of clusters $c = 10$, and set the entropy coefficient λ to 2. Model parameters are optimized with Adam (initial learning rate $lr = 10^{-4}$), and the learning rate is decayed over epochs using a cosine annealing scheduler. All experiments are implemented in Python 3.10 on an Ubuntu 5.4.0 system equipped with 90 GB RAM, a Xeon (R) Silver 4214R CPU (2.40 GHz) and an NVIDIA RTX 3080 Ti (12 GB) GPU.

3.3 Evaluation Metrics

Let N_{TP} , N_{FP} and N_{FN} denote the numbers of true positives, false positives and false negatives at the sample level. To assess the performance of anomaly

detection models, we employ several widely used evaluation metrics, including precision, recall, and F1-score, AUC-PR, AUC-ROC, and the proximity-aware time series anomaly evaluation (PATE) [29]. Precision is defined as $\frac{N_{TP}}{N_{TP} + N_{FP}}$, measuring the proportion of true anomalies correctly identified among all predicted anomalies. Recall is calculated as $\frac{N_{TP}}{N_{TP} + N_{FN}}$, reflecting the proportion of actual anomalies that were successfully detected. The F1-score, given by $\frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$, provides a harmonic mean of Precision and Recall, balancing both metrics. AUC-PR and AUC-ROC, which summarize model performance across varying thresholds. AUC-PR emphasizes the balance between precision and recall. AUC-ROC reflects the trade-off between true positive and false positive rates, indicating the overall discrimination capability of the model. Additionally, we adopt PATE metric, which accounts for temporal proximity when matching predicted anomalies to ground truth. Instead of exact pointwise matching, PATE introduces a tolerance window and penalizes predictions based on their temporal distance from the true anomalies, making it more suitable for real-world time series data where anomalies may not be pinpointed precisely.

4 Result and Analysis

In this section, we present a comprehensive evaluation of the proposed smart meter abnormality detection framework and validate the contributions of its two core modules through targeted ablation studies.

4.1 Overall Performance

Each subplot of Fig. 4 is a radar chart showing the performance of a smart meter across six evaluation metrics, including Precision, Recall, F1, AUC-ROC, AUC-PR and PATE, over seven abnormal test runs. Each colored hexagon corresponds to one run, and the closer a vertex is to the outer circle at 1.0, the better the performance on that

metric. For smart meter abnormality detection task, Recall is usually one of the strongest metrics for performance as missing an actual anomaly can lead to undetected failures, fraud, or security breaches. Across the eight datasets, the proposed model consistently delivers strong detection performance that most radar plots hug the outer ring for Recall, and comprehensive performance metric of F1, PATE, AUC-PR, and AUC-ROC, demonstrating its robust sensitivity to diverse abnormal attacks. It can be observed that our approach achieves more than Recall rate of 0.9 on all trials except SM-MAC004 and SM-MAC006 with abnormality of ‘historical minimum reduction’, approximately 0.8.

In the SM-MAC003 radar chart, most runs maintain uniformly high-performance metrics, but we can clearly see three runs (i.e., abnormality of ‘threshold capping’, ‘peak-hour progressive reduction’ and ‘peak-off-peak swap’) whose Precision collapse toward the center, thereby dragging down the scores of other comprehensive metrics. We observed that across these three trials, the model struggled to distinguish abnormal from normal samples, frequently misclassifying normal observations as anomalies. This behavior likely arises from the dataset’s high prevalence of atypical yet legitimate samples that similar with the synthetic abnormalities. One of the future solutions could be to augment “normal” training set with more edge-case examples to tighten the decision boundary and boost Precision without sacrificing Recall.

Overall, the performance of each metric in Fig. 4 directly demonstrates the effectiveness of the proposed framework for smart meter anomaly detection.

4.2 Ablation Experiment

To assess the individual contributions of the anomaly contrastive learning and context contrastive learning modules, we performed ablation experiments. Given their sequential relationship, we designed two experiments: DSS-ccl, which masks context contrastive learning part and adding a cluster-




Fig. 4 Metric performance

ing head directly after anomaly contrastive learning, and DSS, which is the full smart meter abnormality detection framework proposed in this paper.

Table 1 presents a head-to-head comparison between our full DSS model and the DSS-ccl ablation across seven abnormality types on SM-MAC001 dataset. Notably, the DSS-ccl variant maintains competitive performance: Recall exceeds 0.676, and Precision stays above 0.538 in all cases across anomaly types. This indicates that the anomaly contrastive learning backbone alone is effective to distinguish abnormality to some extent, and that context contrastive learning module serves to further refine detection by capturing subtle temporal relationships. Incorporating context contrastive learning module yields clear benefits: recall improves substantially for (0) historical minimum reduction and (1) fixed weekly percentage cut by +0.243 and +0.120, respectively, and precision sees dramatic gains, especially +0.350 for type (1) fixed weekly percentage cut and +0.361 for type (3) threshold capping, indicating that context contrastive learning both recovers more true positives and suppresses false alarms. Consequently, the F1 rises in almost every category, peaking at +0.243 for type (1) fixed weekly percentage cut, while PATE, AUC-PR,

and AUC-ROC all exhibit consistent positive deltas, reflecting better ranking performance and overall detector robustness. These results demonstrate that modeling contextual relationships is essential: without context contrastive learning, the model struggles to distinguish subtle variations in consumption patterns, but the full framework leverages temporal context to achieve stronger and more reliable anomaly detection across diverse scenarios.

5 Conclusion

We have introduced a real-time decision support framework for smart meter anomaly detection that leverages two complementary self-supervised contrastive learning modules. The Anomaly Contrastive Learning module addresses the scarcity of negative samples in practical deployments by synthesizing diverse abnormal patterns, enabling the model to robustly separate normal and anomalous behaviors. Building on this foundation, the contextual contrastive learning module refines the representation space by capturing intrinsic temporal relationships within meter readings. Our end-to-end pipeline ingests rolling time-window data and continuously updates the model, delivering timely detection of emerging anomalies in dynamic energy systems.

Tab. 1 Performance comparisons of ablation experiment on SM-MAC001

Metrics	Abnormality	(0)	(1)	(2)	(3)	(4)	(5)	(6)
Recall	DSS-ccf	0.676	0.769	0.949	1.000	0.981	1.000	0.886
	DSS	0.919 (+0.243)	0.889 (+0.120)	0.943 (−0.006)	0.966 (−0.034)	0.865 (−0.116)	0.911 (−0.089)	0.981 (+0.095)
Precision	DSS-ccf	0.962	0.610	0.794	0.639	0.704	0.687	0.778
	DSS	1.000 (+0.038)	0.960 (+0.350)	0.937 (+0.143)	1.000 (+0.361)	0.971 (+0.267)	0.993 (+0.306)	0.791 (+0.013)
F1	DSS-ccf	0.794	0.680	0.865	0.780	0.819	0.814	0.828
	DSS	0.958 (+0.164)	0.923 (+0.243)	0.940 (+0.076)	0.983 (+0.203)	0.915 (+0.095)	0.950 (+0.136)	0.876 (+0.047)
PATE	DSS-ccf	0.850	0.610	0.912	0.501	0.805	0.488	0.913
	DSS	0.982 (+0.132)	0.975 (+0.365)	0.986 (+0.074)	0.995 (+0.494)	0.971 (+0.166)	0.987 (+0.499)	0.808 (−0.105)
AUC-PR	DSS-ccf	0.840	0.538	0.903	0.484	0.794	0.474	0.907
	DSS	0.979 (+0.139)	0.966 (+0.427)	0.984 (+0.082)	0.995 (+0.511)	0.963 (+0.169)	0.985 (+0.511)	0.799 (−0.108)
AUC-ROC	DSS-ccf	0.907	0.657	0.804	0.195	0.652	0.015	0.789
	DSS	0.994 (+0.087)	0.953 (+0.296)	0.958 (+0.154)	0.990 (+0.795)	0.900 (+0.247)	0.959 (+0.944)	0.684 (−0.105)

Extensive experiments confirm the strength of our approach. The full framework consistently obtains excellent performances over seven abnormalities across Recall, Precision, F1, AUC-PR, AUC-ROC and PATE metrics. In future work, we plan to incorporate richer, real-world anomaly priors and explore finer-grained detection strategies to further enhance resolution and adaptability in evolving smart meter environments. In addition, it is also interesting to explore a united framework that does not need to acquire each smart meter historical data for privacy protection.

References

- [1] ORLANDO M, ESTEBSARI A, PONS E, et al. A smart meter infrastructure for smart grid iot applications [J]. IEEE Internet of Things Journal, 2022, 9(14): 12529 – 12541.
- [2] WANG F, LU X, CHANG X, et al. Household profile identification for behavioral demand response: A semi-supervised learning approach using smart meter data [J]. Energy, 2022(238): 121728.1 – 121728.12.
- [3] SĂNDULEAC M, CIORNEI I, TOMA L, et al. High reporting rate smart metering data for enhanced grid monitoring and services for energy communities [J]. IEEE Transactions on Industrial Informatics, 2022, 18(6): 4039 – 4048.
- [4] MENG F, MA Q, LIU Z, et al. Multiple dynamic pricing for demand response with adaptive clustering-based customer segmentation in smart grids [J]. Applied Energy, 2023(333): 120626.1 – 120626.12.
- [5] LIANG G, WELLER S R, ZHAO J, et al. The 2015 ukraine blackout: implications for false data injection attacks [J]. IEEE Transactions on Power Systems, 2017, 32(4): 3317 – 3318.
- [6] LIANG G, ZHAO J, LUO F, et al. A review of false data injection attacks against modern power systems [J]. IEEE Transactions on Smart Grid, 2017, 8(4): 1630 – 1638.
- [7] XIA X, XIAO Y, LIANG W, et al. Detection methods in smart meters for electricity thefts: a survey [J]. Proceedings of the IEEE, 2022, 110(2): 273 – 319.
- [8] CAPOZZOLI A, PISCITELLI MS, BRANDI S, et al. Automated load pattern learning and anomaly detection for enhancing energy management in smart buildings [J]. Energy, 2018(157): 336 – 352.
- [9] YIP S C, WONG K, HEW W P, et al. Detection of energy theft and defective smart meters in smart grids using linear regression [J]. International Journal of Electrical Power & Energy Systems, 2017(91): 230 – 240.
- [10] YIP S C, TAN W N, TAN C, et al. An anomaly detection framework for identifying energy theft and defective meters in smart grids [J]. International Journal of Electrical Power & Energy Systems, 2018(101): 189 – 203.
- [11] CHEN L, LAO K W, MA Y, et al. Error modeling and anomaly detection of smart electricity meter using TSVD+L method [J]. IEEE Transactions on Instrumentation and Measurement, 2022(71): 1 – 14.
- [12] BUZAU M M, TEJEDOR-AGUILERA J, CRUZ-ROMERO P, et al. Detection of non-technical losses using smart meter data and supervised learning [J]. IEEE Transactions on Smart Grid, 2019, 10(3): 2661 – 2670.

(下转第 89 页 Continued on Page 89)