

Research on Static Vulnerability Detection Model at the Bytecode Level of Blockchain Smart Contracts

Bingyan Hu

School of Economics, Wuhan Donghu University, Wuhan, 430000, China

Keywords: smart contract; bytecode; static analysis; vulnerability detection; control flow graph

Abstract: Once deployed, blockchain smart contracts are characterized by difficulty in rollback, direct asset binding, and publicly reproducible attack chains. Therefore, vulnerability detection must be conducted as early as possible, before deployment. Compared to source code-centric methods, static detection methods targeting bytecode are more suitable for unpublished contracts, contracts compiled and deployed to the blockchain, and cross-platform auditing. This paper, based on research on smart contract vulnerability detection over the past three years, studies multi-view static vulnerability detection at the bytecode level, focusing on low false positives, low false negatives, and low latency. The model uses opcode sequences, control flow graphs, and dangerous semantic triggering patterns as primary inputs, employing lightweight graph representation, sequence context encoding, and multi-branch belief fusion to identify reentrancy, timestamp dependencies, access control, and self-destruct vulnerabilities. The research methodology follows a "current situation analysis - problem summarization - model design - literature evidence comparison" approach, primarily addressing issues such as unobservable bytecode, cross-version EVM semantic drift, high modeling costs for long sequences, and poor performance due to sparse labels. Based on experimental results from publicly available English literature over the past three years, this paper summarizes the representative model data scale, reporting accuracy, and applicability. The results show that bytecode static detection has evolved from rule-based detection to graph neural networks, Transformers, and multimodal fusion, but significant shortcomings remain in terms of unified benchmarks, interpretability, and online deployment latency. This paper proposes a model framework that can provide methodological support for blockchain auditing platforms, pre-transaction risk control, and contract market access review.

1 Introduction

Blockchain smart contracts turn business rules into executable code and deploy them in the form of bytecode to the virtual machine on the chain. Their automatic execution, immutability, and asset custody operation mode improves the efficiency of cross-entity collaboration on the one hand, and directly transforms software defects into on-chain asset risks on the other hand. After the contract

enters the mainnet, any reentrancy, permission mismatch, timestamp dependency, self-destruction, or integer-related errors can be continuously amplified without human intervention, thereby causing rapid and large-scale economic losses. [1][2]

Existing research suggests that smart contract security analysis methods mainly include formal verification, symbolic execution, fuzz testing, rule auditing, and machine learning detection. Among these, static detection at the bytecode layer is more suitable for real-world situations such as unavailable source code, different compiled versions, difficulty in tracing third-party contract libraries, and large-scale access audits in the contract market. Compared to source code analysis, bytecode analysis can analyze the final execution object on the chain without semantic misalignment caused by missing source code, compiler optimization, or changes in intermediate components. [3]

However, bytecode detection is not simply a sinking of source code detection. Low-level representation preserves execution semantics, but loses information such as identifiers, comments, type hints, and high-level control. Furthermore, the EVM instruction sequence length is inconsistent, jump addresses are difficult to recover, and there are significant differences between different compiler versions. Therefore, detection systems are prone to making choices in terms of detection capability, interpretation capability, and inference latency. In the past three years, researchers have used techniques such as CFG reconstruction, graph neural networks, Transformer, and multimodal fusion to solve the above problems, but existing methods do not have a unified paradigm for static bytecode detection. [4][5]

Therefore, based on the need for static vulnerability detection in blockchain smart contract bytecode, this paper proposes a multi-perspective, low-latency, and deployable static vulnerability detection model. The paper unfolds in the order of "analysis of current research status - problem identification - problem solving or strategy - conclusion," first summarizing relevant research from the past three years, then identifying the main shortcomings of bytecode-level static vulnerability detection, providing the relevant model structure and indicator system, and outlining deployment methods. Finally, based on typical results from publicly available English literature, the feasibility and limitations of the proposed approach are discussed.

2. Current Status Analysis of the Research Topic

In terms of the development history of research, smart contract vulnerability detection has evolved from the original rule matching and symbolic execution to the current automated detection dominated by learners. Systematic reviews show that in recent years, the research focus has shifted from "whether it can be detected" to "whether it can be generalized, whether it can be expanded, and whether the false positive rate can be controlled". After systematically reviewing smart contract vulnerability detection technologies, Vidal et al. concluded that due to the different data set construction methods, the different sources of vulnerability labels, and the differences in evaluation metric standards, different methods cannot be compared horizontally. Crisóstomo et al. further classified the application of feature engineering, graph representation, and deep models in this field from the perspective of machine learning. [1], [2]

One type of research in the direction of bytecode static analysis focuses on low-level semantic recovery. Pasqua et al. proposed to use quantization abstraction and EVM operand stack symbolic execution to recover the goal of indirect jumps in order to build an accurate control flow graph. This shows that the quality of CFG directly affects the identifiability of vulnerability paths in subsequent static detection; inaccurate jump recovery will cause graph neural networks and sequence models to have inherent limitations, even though they have strong representation capabilities. [3]

Another type of research focuses on bytecode detection from a learning perspective. Liu et al.

proposed using the DL4SC method to model the opcode level with deep learning, combining sequence context and convolutional features to improve the recognition of complex vulnerability patterns. [4] Wang et al. believe that Contract Sentry combines static analysis process and toolchain, and pays more attention to rules, graph structure, and engineering deployability. [5] This indicates that the bytecode level has begun to change from the original simple rule reasoning to the current parallel technical route of "rules + learning".

Multimodal fusion further expands the boundaries of bytecode detection capabilities. VulnSense uses source code, opcode sequence and CFG obtained from bytecode for joint learning, and achieved an average accuracy of 77.96% in a set of 1769 smart contracts, showing that the fusion of cross-modal features is effective. [6] Zhang et al. used heterogeneous semantic graphs and pre-training to achieve detection accuracy of 90.96%, 89.57%, 88.46% and 87.34% on four vulnerabilities, indicating that the combination of graph representation and pre-training has good cross-scene adaptability. [7]

In recent years, research has focused on two aspects: strong structure awareness and low inference cost. Gautrin et al. combined masked attention with CFG analysis, emphasizing the use of structural priors to reduce false detections in complex bytecode scenarios. [8] Balcı et al. proposed VASCOT, which directly performs Transformer sequential scanning on EVM bytecode/instructions, using a sliding window to process long sequences, and conducted research on 16,469 verification contract data. [9] Yang et al. proposed ByteEye, which uses edge-enhanced CFG and graph neural networks to improve bytecode-level vulnerability detection. On the ESC_P dataset, it achieved an F1 score of 65.98% on the timestamp dependency task and 100% success, which is important for engineering auditing. [10]

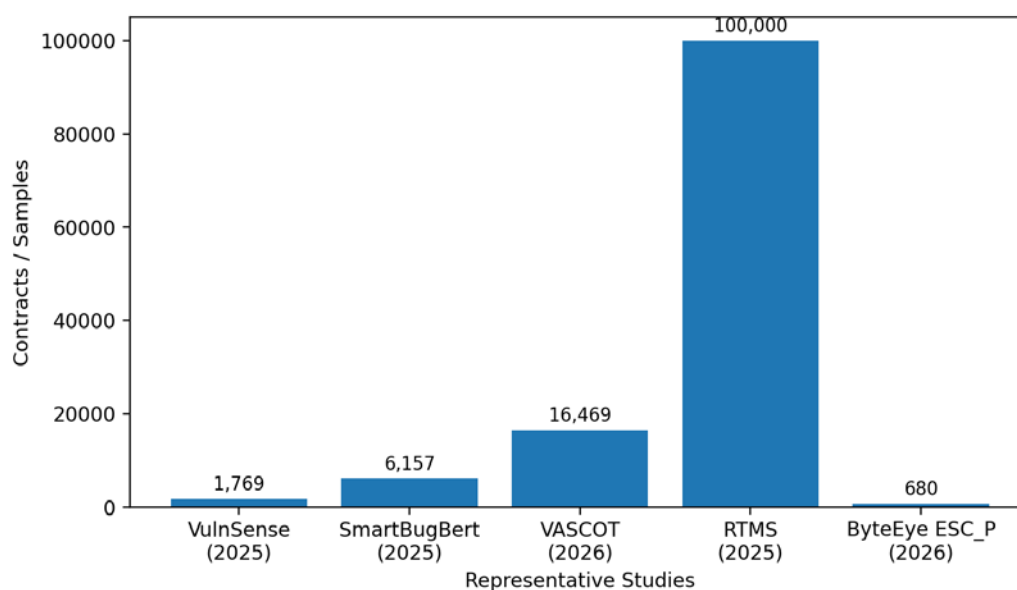


Figure 1 Comparison of data size from representative studies in recent years

Figure 1 shows the results of statistical analysis of the data scale information available in publicly available research over the past three years. The figure illustrates the number of contract samples covered by various models. Figure 3.1 shows that recent methods such as RTMS and VASCOT tend to expand the sample coverage to improve label sparsity and scenario bias. ByteEye primarily focuses on highly targeted modeling for specific subsets of vulnerabilities. While increased data volume can improve model generalization, it also introduces problems such as class imbalance, label noise, and increased inference costs. Therefore, sample size alone cannot

determine detection capability; it needs to be considered in conjunction with label quality and input representation.

Table 1 Comparison of Representative Smart Contract Vulnerability Detection Studies in the Past Three Years

Study	Level	Core Input	Model Family	Dataset Size	Reported Metric
DLASC (2024)	Opcode	Opcode sequence	Transformer+C NN	N/A	Improved deep detection
ContractSentry (2025)	Static	Rule+CFG	Static analysis tool	N/A	Tool-oriented detection
VulnSense (2025)	Hybrid	Source+Opcode +CFG	BERT+BiLSTM+GNN	1769	Avg Acc. 77.96%
HCSG+PT (2024)	Graph	Semantic graph	Pre-training graph model	N/A	Avg Acc. 89.08%
VASCOT (2026)	Bytecode	Opcode sequence	Transformer	16469	Recent large-scale scan
ByteEye (2026)	Bytecode	Edge-enhanced CFG	GNN	680 (ESC_P subset)	TS-Dep F1 65.98%

Table 1 compares representative studies using analysis hierarchy, core input, model family, and publicly available metrics. The early single-input, single-model approach is gradually giving way to composite structures based on graph structures, sequence context, and rule priors. Especially in bytecode detection, the quality of CFG recovery and the representation of opcode context have become key factors determining model performance. Furthermore, different studies report varying metrics, indicating that a unified and stable evaluation standard covering bytecode scenarios has not yet been established in the field.

3. Raise questions

Existing research indicates that static vulnerability detection at the bytecode layer has the following four structural flaws.

First, there is an input bottleneck between semantic loss and structural recovery. Bytecode strips away variable names, comments, explicit control semantics, and the developer's intent. Many high-level security constraints need to be re-established by low-level jumps, storage reads, and call relationships. The CFG establishment phase cannot accurately find the reachable edges, function entry points, and abnormal exit paths of JUMP/JUMPI, making it difficult for subsequent models to capture important features such as reentrant call chains, permission gating, and state update order. [3][10]

Second, the coexistence of long sequences and sparse dangerous patterns makes the model unable to maintain robustness under low latency requirements. In scenarios such as financial transactions, DeFi risk control and contract market access, the detection system needs to complete the scan in a short time, but the opcode sequence is long, while dangerous semantic segments only account for a small proportion. Directly using a full sequence deep model will lead to excessively high inference costs, and a large number of useless instructions will be submerged in a small number of dangerous patterns. [4], [9]

Third, the datasets are heterogeneous and the labeling standards are inconsistent. Most of the current smart contract vulnerability datasets come from automated tool annotation, injection benchmarks, manual review, or on-chain tracing. The use of multiple standards will lead to high-

precision results not necessarily equaling high-reliability results. In particular, at the bytecode level, the final presentation results will fluctuate due to differences in compiler version, optimization level, and library inlining method, thereby reducing its applicability across different datasets. [1][2] None

Fourth, the detection results are neither interpretable nor usable. Simply providing the probability of vulnerability types does not support secondary verification by audit engineers; however, over-reliance on complex path constraints and full-graph reasoning leads to excessive analysis latency. Finding a balance between sufficient interpretation, controllable false positives, and online usability remains a problem that bytecode static vulnerability detection models need to solve.

4. Problem Solving/Strategies

To address the aforementioned shortcomings, this paper proposes a static vulnerability detection method oriented towards the bytecode layer, i.e., a vulnerability detection method that uses the bytecode layer as its view. The model uses three features as input: opcode sequences, edge-enhanced control flow graphs, and dangerous semantic triggering patterns. Vulnerability detection is achieved through a three-step procedure: lightweight sequence encoding, local structure convergence, and confidence fusion. The goal is not to pursue the most complex structure, but to improve the accuracy of identifying critical vulnerabilities and cross-version stability while ensuring controllable pre-deployment scanning latency.

The bytecode is standardized and segmented, and environment-sensitive instructions such as PUSH constants, CALL calls, SSTORE/SLOAD state access, and TIMESTAMP/BLOCKHASH are separately annotated. Local windows are constructed centered on dangerous semantic segments. To reduce noise in long sequences, the model does not apply equal weights to the entire sequence, but instead uses a risk score function to weight different segments.

$$S_i = \lambda_1 O_i + \lambda_2 C_i + \lambda_3 D_i \quad (1)$$

In equation (1), S is the overall risk score of the fragment, O is the opcode statistical feature, C is the CFG local structural feature, and D is the trigger strength of the dangerous instruction; λ_1 , λ_2 , and λ_3 are fusion weights that can be learned or set empirically. After prioritizing the fragments, the model will first analyze the local areas that may generate vulnerabilities, thereby reducing unnecessary computation.

Secondly, a lightweight context encoder is used to learn representations for high-risk segments in the sequence branch, while in the graph branch, only edges related to external calls, storage updates, conditional jumps, and abnormal terminations are retained, constructing an edge-enhanced subgraph to avoid heavy inference on the complete large graph. The outputs of the two branches are then combined through a gating fusion layer for judgment.

$$P(v) = \sigma(W_s h_s + W_g h_g + b) \quad (2)$$

In equation (2), $P(v)$ is the probability that sample v belongs to the target vulnerability category. The sequence branch output and graph branch output are divided into two parts: context information and structural information. The mapping matrix W and the bias term b appear together, and their results are processed using the Sigmoid activation function. This equation reflects the process of the model operating simultaneously in both bytecode context and structural information channels.

To improve the stability of detection in multi-vulnerability scenarios, this paper uses a joint optimization of primary task loss and secondary consistency loss. The primary task handles the determination of vulnerability labels, while the secondary loss is used to control the consistency of different view outputs near high-risk segments, thereby reducing false positives caused by distortion

in individual views.

$$L = \alpha L_{cls} + \beta L_{con} + \gamma R(\Theta) \quad (3)$$

In equation (3), the classification loss term learns the missing labels, the consistency loss term ensures that the output angles of each viewpoint are consistent, $R(\Theta)$ is the parameter regularization term, and α , β , and γ are the weights of the three parts, respectively. Using this objective function, we can simultaneously obtain the missing labels and suppress the complexity of the parameters, so that the representation differences between the various modalities can be controlled to a certain extent.

In the context of financial transactions and on-chain risk control, this paper continues to use average detection latency as an engineering metric to reflect the throughput capacity that the model can achieve in the batch access review process, that is, the model's performance when processing large amounts of data.

$$Lat_{avg} = \left(\frac{1}{N}\right) \sum_{i=1}^N t_i \quad (4)$$

Average detection latency is the average time from bytecode parsing to risk result for a single contract, where t is the execution time of a single sample and N is the number of samples. Because pre-deployment risk control prioritizes batch processing efficiency, a lower metric makes it more suitable for applications such as pre-match review, contract market access screening, and automated asset routing filtering.

This paper mainly uses precision, recall, and F1 score as evaluation criteria. Among these evaluation criteria, F1 score places more emphasis on the balance between false positives and false negatives than precision and recall, and is therefore more suitable as the evaluation result.

$$F1 = 2 \times \text{Precision} \times \text{Recall} / (\text{Precision} + \text{Recall}) \quad (5)$$

Equation (5) defines the F1 score. For smart contract vulnerability detection, if only high recall is pursued, the audit system will generate a large number of false positives; if only high accuracy is pursued, some important risk samples will be missed. Therefore, F1 is particularly suitable as a performance indicator for bytecode static detection models before deployment.

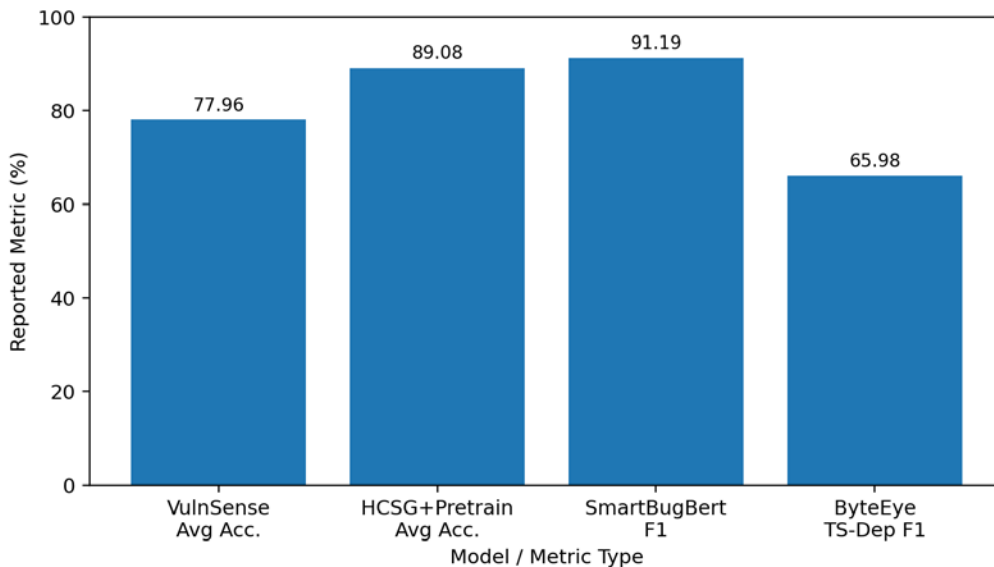


Figure 2 Comparison of indicators from publicly available representative research reports

Figure 2 compares publicly available metrics that can be directly extracted from representative studies in recent years. It should be noted that the types of metrics differ across reports. The figure presents both average accuracy and F1 score, thus more closely resembling a "research progress profile" than a strictly standardized ranking. However, two trends are still observed: first, publicly available results from methods combining graph structures and pre-training are significantly better than those using a single representation; second, companies like ByteEye place greater emphasis on bytecode engineering exploitability solutions. While these solutions do not achieve optimal unified metrics, they demonstrate higher success rates on specific vulnerability tasks and exhibit targeted deployment.

Table 2 Key Design Points for Static Vulnerability Detection at the Bytecode Layer

Vulnerability	Bytecode Clue	Static Challenge	Recommended Signal
Reentrancy	CALL before SSTORE	Cross-path reasoning	Call edge + state-write order
Timestamp Dep.	TIMESTAMP / NUMBER	Context ambiguity	Environment opcode + branch guard
Access Control	CALLER / ORIGIN checks	Implicit policy loss	Condition chain + modifier proxy
Selfdestruct	SELFDESTRUCT opcode	Rare pattern sparsity	Rare opcode boost + path reachability
Integer Issue	Arithmetic + JUMPI	Version-dependent semantics	Opcode neighborhood + compiler tag

Table 2 summarizes the key considerations for model design from four aspects: vulnerability type, bytecode clues, difficulties in static analysis, and suggested signals. The table shows that bytecode layer detection is not simply a bag-of-words analysis of opcodes; it must include environment instructions, path reachability, state write order, and compilation semantics within the modeling scope. Especially in access control and integer issues, many risks are not directly caused by a single instruction, but rather by the combination of multiple instructions in certain control paths. Therefore, multi-view modeling is particularly important.

5. Model Applicability and Discussion

In terms of engineering implementation, the model presented in this paper has the following three advantages. First, it can directly target bytecode, covering contracts that are not open source and post-deployment auditing. Second, it adopts a risk fragment priority strategy and a local graph reasoning strategy, which can reduce reasoning time without significantly increasing reasoning latency. Third, it incorporates explicit risk clues, allowing audit results to be traced back to dangerous fragments, important call edges, and environment-sensitive instructions, thereby improving interpretability.

The model presented in this paper is better able to discover different manifestations of the same type of vulnerability than tools that rely entirely on rules. Compared with heavy full-graph learning models, the framework presented in this paper is more suitable for rapid screening and batch access control of a large number of users before deployment. In financial, asset-related, and high-frequency interaction applications, if risk filtering can be uniformly input using bytecode before the contract goes live, it is possible to greatly improve the consistency and proactiveness of the audit process. [5][9][10]

Of course, the model established in this paper still has some limitations. First, bytecode layer detection is limited by the lack of high-level semantics; even with enhancements using CFG and

dangerous modes, complex business logic cannot be reconstructed. Second, different studies use different datasets, label sources, and vulnerability classification standards, meaning that cross-document comparisons can only be used to judge trends, and cannot be directly equated to performance ranking in the same experimental environment. For proxy contracts, inline assembly, cross-contract cooperative calls, and novel EVM extension instructions, existing static modeling still has the potential for identification errors.

Therefore, future research should focus on these three aspects: first, establishing a unified and traceable bytecode-level benchmark dataset to identify vulnerability labels and compilation environments; second, enhancing the ability to jointly model bytecode, CFG, call graphs, and transaction contexts to improve the ability to identify cross-contract risks; and third, combining interpretable outputs with automatic remediation suggestions so that the model can not only know whether a risk exists, but also provide information on where the risk is located, why it occurs, and how to mitigate it.

6. Conclusion

This paper focuses on static vulnerability detection at the bytecode level of blockchain smart contracts. Based on publicly available English research results in recent years, it analyzes the current research status from three aspects: technological development, major bottlenecks, and model design. Building upon this, a multi-view static vulnerability detection model for pre-deployment scenarios is proposed. This model combines opcode sequences, edge-enhanced CFG, and dangerous semantic triggering patterns to improve the ability to identify critical vulnerabilities while maintaining low analysis costs.

Research indicates that the performance ceiling for static vulnerability detection at the bytecode layer is not solely determined by the complexity of the deep model; it is also influenced by factors such as whether the input semantics can be recovered, whether local high-risk segments can be identified, and the reliability of the evaluation data. Current representative research shows a general trend towards rule engines moving towards graph neural networks, Transformers, and multimodal fusion. However, achieving a unified benchmark, interpretability, and low-latency implementation remain significant challenges that must be overcome in the future.

In summary, static vulnerability detection at the bytecode level has significant engineering implications and promising future prospects. The model proposed in this paper can provide a feasible technical solution for on-chain auditing platforms, DeFi access review, contract market governance, and pre-trade risk control.

References

- [1] Liu, X., & Yang, D. (2025, March). *LLM Data Strategy: Improving Data Availability and Efficiency*. In *Doctoral Symposium on Computational Intelligence* (pp. 425-437). Singapore: Springer Nature Singapore.
- [2] Zhang, Q. (2025, October). *Application of Reinforcement Learning in Dynamic Advertising Content Generation*. In *2025 2nd International Conference on Software, Systems and Information Technology (SSITCON)* (pp. 1-5). IEEE.
- [3] Zhang, Q. (2026). *Security Improvement and Application of Identity and Access Management in SaaS Platform*.
- [4] Wu, Y. (2025, October). *Multi-Level Belief Rule Base Modeling Architecture and Intelligent Optimization Technology for Decision Support Systems*. In *2025 2nd International Conference on Software, Systems and Information Technology (SSITCON)* (pp. 1-8). IEEE.

- [5] Chen, M. (2026). *Research on Privacy-Preserving AI Model Training and Validation Methods Based on Federated Learning*.
- [6] Yiting Hong. *Differentially Private High-Dimensional Business Data Publishing and Analysis Algorithm*. *International Journal of Business Management and Economics and Trade* (2026), Vol. 7, Issue 1: 28-35.
- [7] Xu, D. (2026). *Analysis of the impact of video infrastructure optimization on large-scale content quality improvement*.
- [8] Pan, H. (2025, March). *Research on Efficient Computing Model of Hartree Fock and Density Functional Theory Based on GPU Acceleration*. In *Doctoral Symposium on Computational Intelligence* (pp. 485-496). Singapore: Springer Nature Singapore.
- [9] Wu, W. (2025, June). *Construction and optimization of intelligent gateway software management platform based on jenkins cluster management under cloud edge integration architecture in industrial internet of things*. In *International Conference on 6G Communications Networking and Signal Processing* (pp. 633-645). Singapore: Springer Nature Singapore.
- [10] Wu Y. *Software Engineering Practice of Microservice Architecture in Full Stack Development: From Architecture Design to Performance Optimization*[J]. 2025.
- [11] Wu Y. *Optimization of Generative AI Intelligent Interaction System Based on Adversarial Attack Defense and Content Controllable Generation*[J]. 2025.
- [12] Sun J. *Quantile Regression Study on the Impact of Investor Sentiment on Financial Credit from the Perspective of Behavioral Finance*[J]. 2025.
- [13] Wang Y. *Application of Data Completion and Full Lifecycle Cost Optimization Integrating Artificial Intelligence in Supply Chain*[J]. 2025.
- [14] Chen M. *Research on Automated Risk Detection Methods in Machine Learning Integrating Privacy Computing*[J]. 2025.
- [15] Wu Y. *Software Engineering Practice of Microservice Architecture in Full Stack Development: From Architecture Design to Performance Optimization*[J]. 2025.
- [16] Zhou, Y. (2026). *Energy efficiency and sustainability strategies for data centers*. *European Journal of Engineering and Technologies*, 2(1), 46-53.
- [17] Yanchun Wang. (2025) *Research on Enhancing ERP System Efficiency Through AI in Cross-border Supply Chain Environments*. *Advances in Computer and Communication*, 6(5), 268-273.
- [18] Yanchun Wang. (2025) *Research on Enhancing ERP System Efficiency Through AI in Cross-border Supply Chain Environments*. *Advances in Computer and Communication*, 6(5), 268-273.
- [19] Sun, J. (2025). *Research on Business Data-driven Risk Prediction Methods Based on Machine Learning*. *Advances in Computer and Communication*, 6(4).
- [20] Wu, Y. (2026). *Federated Learning-based Algorithm Design for Privacy Preservation in Cross-domain Data Sharing*. *Engineering Advances*, 6(1).
- [21] Ding, J. (2025). *Exploration of Process Improvement in Automotive Manufacturing Based on Intelligent Production*. *International Journal of Engineering Advances*, 2(2), 17-23.
- [22] Ding, J. (2025). *Research On CODP Localization Decision Model Of Automotive Supply Chain Based On Delayed Manufacturing Strategy*. *arXiv preprint arXiv:2511.05899*.
- [23] Zhang, C., Han, J., Zou, Y., Dong, K., Li, Y., Ding, J., & Han, X. (2024, April). *Detecting the anomalies in LiDAR pointcloud*. In *WCX SAE World Congress Experience*. SAE Technical Paper.

- [24] Hou, Y. (2026). *Research on Heterogeneous Server Upgrade Strategies and Resource Utilization Efficiency Oriented Toward Green Computing Objectives*. *Advances in Computer and Communication*, 7(1).
- [25] Hou, Y. (2026). *Exploration of Data Center Cost Optimization Pathways Under Multi-generation CPU and GPU Collaborative Architectures*. *Engineering Advances*, 6(1).
- [26] Huang, J. (2025, September). *Performance Evaluation Index System and Engineering Best Practice of Production-Level Time Series Machine Learning System*. In *2025 International Conference on Intelligent Communication Networks and Computational Techniques (ICICNCT)* (pp. 01-07). IEEE.
- [27] Huang, J. (2025, August). *Research on Multi-Model Fusion Machine Learning Demand Intelligent Forecasting System in Cloud Computing Environment*. In *2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS)* (pp. 1-7). IEEE.
- [28] Zou, H., & Hu, H. (2026). *Privacy-preserving P2P energy trading for economic optimization in urban microgrids*. *Applied Energy*, 412, 127646.