Cole-Puli
International Conference Publishing Group

Science, Engineering and Technology Proceedings

Article

2025 International Conference on Science Technology, Architecture, Power and Intelligent Information Technology (APIIT 2025)

Leveraging NLP for Semantic and Numerical Inconsistency Detection in Tax Submissions

Jiayu Liang [1,*], Jiayan Fan [2], Zhen Feng [3] and Jing Xin [4]

1 Applied Statistics, Cornell University, NY, USA
2 Information Science, University of Michigan, Ann Arbor, MI, USA
3 Business Analytics, University of Rochester, Rochester, NY, USA
4 Business Analytics, UW Madison, Madison, WI, USA
* Correspondence: Jiayu Liang, Applied Statistics, Cornell University, NY, USA

Abstract: This study presents an innovative method for identifying tax fraud through the application of natural language processing (NLP) to uncover irregularities within tax documents. Departing from conventional approaches that rely primarily on numerical analysis, the proposed framework combines domain-specific BERT embeddings with bidirectional LSTM architectures to effectively capture nuanced contextual information. A hybrid ensemble architecture is developed to process both structured data and free-text components within tax returns, facilitating the identification of semantic associations among financial entities and exposing numerical inconsistencies. The system was evaluated on a dataset comprising 15,000 tax documents, of which 8.5% were identified as fraudulent. The proposed model achieved superior performance, with an F1-score of 0.868 and an AUC of 0.931 — marking a 7.6% enhancement over leading existing models. Detection effectiveness varied by document category: individual income tax filings yielded an F1-score of 0.889, outperforming business-related filings, which scored 0.818. Further examination reveals that semantic features are particularly effective for identifying fraud in corporate tax documents, while numerical coherence indicators are more significant for personal filings. Although the approach requires higher computational resources compared to conventional techniques, its capacity to detect complex fraud schemes — especially those that disguise manipulation within textual content while maintaining plausible numeric data — offers a significant improvement to current tax fraud detection systems.

Keywords: tax compliance monitoring; deep learning; entity recognition; financial document analysis

Received: 01 May 2025
Revised: 06 May 2025
Accepted: 19 May 2025
Published: 17 June 2025

1. Introduction

1.1. Background

Tax fraud remains a major concern for governments and tax authorities globally, leading to significant annual losses in public revenue. The increasing digitization of tax filing systems has added new layers of complexity to detecting fraudulent behavior. Recent estimates suggest that tax evasion accounts for 2–5% of global GDP, translating into a loss of approximately $2–5 trillion in government revenue each year [1]. Tax documents often contain critical indicators of potential fraud, such as discrepancies in declared income, unusual deduction claims, and irregular financial transactions [2]. Detecting these indicators requires advanced analytical tools capable of processing vast amounts of financial data with high precision. Traditional audit methods, which depend heavily on manual review and fixed rule-based systems, have proven inadequate in identifying the evolving tactics used by tax evaders to manipulate their financial obligations [3].

1.2. Barriers to Effective Tax Fraud Detection

Tax fraud detection presents a range of technical and operational hurdles. One primary issue is the severe class imbalance: legitimate tax filings vastly outnumber fraudulent ones, making it difficult to train effective detection algorithms. Additionally, tax evasion tactics are constantly evolving, necessitating continuous updates to detection strategies in order to stay ahead of fraudsters who adapt quickly to circumvent existing safeguards [4]. The diversity and complexity of tax documents — driven by variations in taxpayer profiles, industry-specific practices, and jurisdictional regulations — pose further challenges. Differences in document formats, terminology, and reporting standards hinder the development of standardized detection models [5]. Moreover, the scarcity of labeled fraud data presents a significant limitation, as confirming fraud often requires lengthy investigations and legal validation. These constraints underscore the need for innovative detection methods that can function effectively under conditions of limited ground truth and maintain low false positive rates.

1.3. Utilizing NLP Techniques in Anomaly Detection Frameworks

Natural Language Processing (NLP) offers powerful tools for enhancing tax fraud detection, particularly in the analysis of unstructured and semi-structured components of tax filings [6]. NLP techniques can uncover semantic patterns and contextual cues that are often overlooked by traditional data mining methods. For example, text classification algorithms can assess the risk level of filings, while named entity recognition helps detect inconsistencies in reported entities and business affiliations. Deep learning-based NLP models excel at identifying subtle linguistic anomalies that may suggest fraudulent intent. Transfer learning allows models trained on related financial data to be adapted to tax-specific tasks, helping address the challenge of limited labeled examples [7]. Bidirectional language models provide rich contextual embeddings that capture intricate relationships among financial terms and entities within documents. When combined with conventional numerical anomaly detection, NLP enables the creation of multimodal systems that can identify complex tax evasion schemes more effectively, significantly boosting detection performance compared to single-modality approaches [8,9].

2. Literature Review

2.1. Legacy Methods for Identifying Tax Fraud

Conventional tax fraud detection methods have traditionally relied on rule-based systems and manual audits conducted by tax authorities. These approaches typically involve the use of predefined heuristics and threshold-based criteria to flag suspicious filings. As highlighted by Yan et al., traditional practices encompass manual case selection, whistleblower-driven reporting, and computer-assisted selection — all of which are resource-intensive and time-consuming processes [3,10]. Rule-based systems apply explicitly programmed rules to detect deviations from expected financial behavior, using static thresholds to trigger alerts. For instance, Wu et al. demonstrated the use of association rules within tax databases to enhance the detection of value-added tax (VAT) fraud, achieving moderate gains over manual auditing alone [4,7,11,12].

In addition to rule-based techniques, basic statistical tools such as ratio analysis and correlation evaluation have been used to examine relationships between key financial indicators — for example, the correlation between total GST liabilities and total sales figures [13]. However, these traditional methods are limited in scalability and adaptability, particularly as fraudulent tactics become more sophisticated. They often fail to capture complex fraud schemes embedded in large-scale, high-dimensional datasets. Furthermore, perception-based approaches generally lack the analytical depth required to detect well-disguised fraudulent behavior, especially when perpetrators intentionally fabricate documents to appear compliant [14].

2.2. Overview of Machine Learning Models

Machine learning (ML) has significantly advanced the field of tax fraud detection by enabling systems to uncover complex patterns in taxpayer behavior. Supervised learning models, such as logistic regression and decision trees, have been effectively used to classify compliance based on historical taxpayer data. To address the challenge of limited labeled examples in tax fraud cases, Weng proposed the use of unsupervised conditional adversarial networks for tax default prediction, offering a means to learn from unannotated data [5,15,16].

Clustering techniques have also proven valuable in grouping taxpayers by similar characteristics, thereby facilitating targeted investigations into anomalous subgroups. Liu et al. applied K-means clustering to detect atypical taxpayer profiles by examining consumption behavior and correlation metrics across both smart and non-smart grid networks [6,17]. Ensemble learning methods, which integrate multiple weak classifiers, have demonstrated superior performance. For example, Xu et al. introduced a Transfer Adaptive Boosting (TAB) algorithm that effectively predicts tax compliance outcomes by combining the strengths of individual learners [7,18].

Deep learning models, particularly bidirectional generative adversarial networks (BiGANs), have recently gained attention for their potential in detecting sophisticated fraud patterns. Xu et al. enhanced BiGAN training frameworks to improve anomaly identification in tax data. Their study showed that measuring cosine similarity between original and regenerated data representations can be an effective indicator of fraudulent filings [8,19,20,21].

3. Methodology

3.1. Methods for Collecting and Preparing Tax Filing Data

This study utilizes a dataset comprising 15,000 tax filing documents sourced from regional tax authorities between 2020 and 2023 [22]. These documents fall into three main categories: individual income tax returns, business tax declarations, and value-added tax (VAT) statements. Among them, individual income tax returns represent the largest portion, totaling 9872 documents and accounting for approximately 65.81% of the dataset. Business tax declarations make up 3456 documents, or 23.04%, while the remaining 1672 documents (11.15%) consist of VAT statements.

The preprocessing workflow consists of several sequential steps aimed at standardizing document formats and preparing textual data for downstream analysis. All 15,000 tax documents were first converted into a consistent text format, requiring approximately 8.75 hours to complete. For the subset of 4328 scanned PDF files, optical character recognition (OCR) was applied, with the OCR stage taking around 11.46 hours. Subsequently, noise reduction techniques were implemented across the entire dataset to remove formatting artifacts and extraneous characters, which resulted in a 97.4% error reduction and required about 5.23 hours of processing time. Finally, tokenization was performed on all documents, completing in roughly 3.12 hours.

The text normalization process includes tokenization, stemming, and lemmatization, enhanced with custom extensions designed to handle tax-specific terminology more effectively. To reduce dimensionality and improve model efficiency, a curated stopword list tailored to the tax domain is applied. This filtering process significantly reduces the vocabulary size — from an initial 18,345 unique terms to 10,325 — representing a 43.7% reduction in vocabulary space.

The document processing pipeline is architected as a multi-stage framework specifically optimized for the analysis of tax-related documents. As illustrated in the accompanying diagram, the pipeline progresses sequentially from document ingestion to format normalization, followed by text extraction, content cleaning, and linguistic normalization. Particular emphasis is placed on modules for financial entity recognition and the processing of domain-specific tax terminology. To ensure data integrity and consistency, the architecture integrates quality assurance mechanisms through feedback loops. Automated error detection components monitor processing outcomes, and any instance in which confidence scores fall below predefined thresholds automatically triggers reprocessing of the affected documents.

3.2. NLP-Driven Feature Extraction for Anomaly Detection

The feature extraction process adopts a hybrid strategy that combines traditional statistical NLP techniques with advanced deep learning models. Initially, Term Frequency–Inverse Document Frequency (TF-IDF) vectors are employed to quantify the relative importance of terms within each document, resulting in a high-dimensional representation of 12,456 features, with a memory footprint of 2.34 GB and an average processing time of 0.043 minutes per document [23]. To capture deeper semantic and contextual information, word embeddings are also generated using both Word2Vec and transformer-based models. Word2Vec yields 300-dimensional embeddings, requires 0.87 GB of memory, and processes each document in approximately 0.126 minutes. For more specialized representations, a domain-adapted BERT model — referred to as Tax-BERT — fine-tuned on 2.3 million tax-related documents, is employed [24]. Tax-BERT produces 768-dimensional vectors with a memory requirement of 3.75 GB and a processing time of 0.284 minutes per document. For comparison, Financial-BERT, another transformer model trained on general financial texts, exhibits similar dimensionality (768) but slightly higher memory use (3.82 GB) and marginally longer processing time (0.291 minutes per document).

To enhance information extraction from tax documents, Named Entity Recognition (NER) is applied to identify critical financial entities, including income sources, expense categories, business relationships, and financial institutions. The NER model demonstrates robust overall performance, achieving an F1-score of 0.892 on the validation set. Breakdown by entity type reveals that the model performs particularly well in recognizing financial institutions, with a precision of 0.954, recall of 0.941, and an F1-score of 0.947. Income sources are also accurately identified, yielding an F1-score of 0.900, supported by a precision of 0.913 and recall of 0.887. Slightly lower but still reliable performance is observed in detecting expense categories and business relationships, with F1-scores of 0.868 and 0.838, respectively. Beyond entity recognition, semantic relationship extraction is performed to map inter-entity connections, facilitating the construction of a knowledge graph that captures the transactional and relational structure embedded in the tax data.

The analysis of feature importance across tax document types reveals clear differences in which features contribute most to anomaly detection. Quantitative evaluation shows that entity relationship features, such as inter-company transactions and business affiliations, are particularly influential in the detection of anomalies in business tax declarations. In contrast, numerical consistency features, such as logical correlations between reported income and deductions, play a more dominant role in individual income tax returns. These findings highlight the structural variation in feature relevance between document types and suggest that anomaly detection models benefit from feature prioritization based on document context.

3.3. Conceptual Framework for Anomaly Detection

The anomaly detection framework proposed in this study adopts a multi-layered architecture that integrates both unsupervised and supervised learning techniques. The unsupervised component employs an Isolation Forest algorithm to detect anomalies by analyzing the isolation paths of feature instances. Meanwhile, the supervised component consists of a bidirectional Long Short-Term Memory (Bi-LSTM) network trained on labeled tax document data. Both components take document embedding vectors — generated during the NLP feature extraction stage — as their input.

The architecture is designed to process structured and unstructured tax information in parallel, combining outputs from both detection paths through an ensemble voting mechanism to generate a final anomaly score. The system also incorporates adaptive feedback loops that adjust decision thresholds dynamically based on performance metrics from previous detection rounds, improving model robustness and adaptability over time.

Hyperparameter tuning was conducted using Bayesian optimization, testing a total of 1,143 parameter combinations to identify optimal model settings. The final ensemble configuration achieved a true positive rate of 0.874 and a false positive rate of 0.058 on the validation dataset — representing a 23.6% improvement over established baseline methods [25,26].

4. System Implementation and Evaluation Results

4.1. Testing Arrangement and Data Sources

The experimental evaluation employed computational resources consisting of an NVIDIA A100 GPU with 80GB memory, Intel Xeon Platinum 8380 CPU with 40 cores, and 512GB RAM. The implementation utilized PyTorch 1.12.0 with CUDA 11.6 support for deep learning components and scikit-learn 1.1.2 for traditional machine learning algorithms. Table 1 presents the hardware and software specifications used in the experimental setup.

Table 1. Hardware and Software Specifications.

Component	Specification
CPU	Intel Xeon Platinum 8380, 40 cores, 2.3 GHz
GPU	NVIDIA A100, 80GB VRAM
RAM	512GB DDR4-3200
Operating System	Ubuntu 20.04 LTS
Deep Learning Framework	PyTorch 1.12.0
NLP Libraries	HuggingFace Transformers 4.21.1, spaCy 3.4.1
ML Libraries	scikit-learn 1.1.2, XGBoost 1.6.2

The dataset comprised 15,000 tax documents split into training (60%), validation (20%), and testing (20%) sets, with stratified sampling maintaining consistent class distributions across splits. Within this dataset, 1275 documents (8.5%) were labeled as anomalous based on prior tax audit findings [27]. Table 2 details the dataset partitioning and anomaly distribution across training, validation, and testing subsets.

Table 2. Dataset Partitioning and Anomaly Distribution.

Subset	Total Docs	Normal Docs	Anomalous Docs	Anomaly %
Training	9000	8235	765	8.5%
Validation	3000	2745	255	8.5%
Testing	3000	2745	255	8.5%
Total	15,000	13,725	1275	8.5%

The model training process employed a batch size of 32 with Adam optimization and a learning rate of 3×10^{-5} with cosine annealing. Early stopping with patience of 10 epochs monitored validation loss to prevent overfitting. The tax-BERT model required 14.5 hours for fine-tuning across 25 epochs.

The effectiveness of the model was assessed using several key metrics, including precision, recall, F1-score, area under the ROC curve (ROC-AUC), and precision-recall area under the curve (PR-AUC). Considering the class imbalance typically present in tax fraud detection tasks, PR-AUC serves as a more reliable indicator of model performance [27]. Table 3 summarizes the detailed performance results for different components within the anomaly detection system.

Table 3. Performance Metrics across Detection Components.

Component	Precision	Recall	F1-Score	ROC-AUC	PR-AUC
TF-IDF + Isolation Forest	0.714	0.682	0.698	0.832	0.735
Word Embeddings + LSTM	0.782	0.743	0.762	0.867	0.789
Tax-BERT + BiLSTM	0.835	0.812	0.823	0.904	0.842
Multi-Component Ensemble	0.874	0.863	0.868	0.931	0.879

The framework enhanced by NLP techniques shows notable improvements in detecting anomalies, with the ensemble method reaching an F1-score of 0.868 [27]. Performance varies depending on the tax document category, as detailed in Table 4, where business tax filings exhibit greater difficulty in anomaly identification compared to other types.

Table 4. Performance Variation across Document Types.

Document Type	Precision	Recall	F1-Score	False Positive Rate
Individual Income Tax	0.897	0.881	0.889	0.042
Business Tax	0.823	0.814	0.818	0.076
Value-Added Tax	0.872	0.859	0.865	0.053

4.3. Analysis of Differences from Earlier Techniques

The NLP-enhanced method proposed in this study was evaluated against several established tax fraud detection techniques documented in recent research. Table 5 summarizes a detailed comparison of key performance indicators and computational demands for each method.

Table 5. Comparison of Performance and Resource Usage across Methods.

Method	F1-Score	AUC	Training Time (hours)	Inference Time (ms/doc)	Memory Usage (GB)
Rule-based System	0.683	0.742	N/A	12.4	0.8
K-means Clustering	0.714	0.768	3.2	28.7	2.3
Transfer Learning (IRTED-TL)	0.824	0.882	8.7	42.1	5.6
BiGAN	0.837	0.893	12.3	37.8	7.2
TAB Algorithm	0.846	0.908	10.5	31.2	6.4
Proposed NLP-Enhanced Approach	0.868	0.931	14.5	45.3	8.7

The proposed NLP-driven framework outperforms all compared methods in terms of accuracy metrics, achieving a 7.6% higher F1-score relative to the leading existing technique (TAB Algorithm). While this gain requires greater computational resources and slightly increased inference time, these demands remain manageable within practical tax authority operations.

5. Conclusion

5.1. Summary

This study proposed a novel anomaly detection approach for tax filing documents by leveraging advanced natural language processing techniques, yielding substantial improvements over traditional methods. The multi-component ensemble framework achieved an F1-score of 0.868 and an AUC of 0.931, marking a 7.6% enhancement compared to the leading established approach. The effective combination of tax-domain-specific BERT embeddings with bidirectional LSTM networks enabled robust modeling of contextual relationships within tax documents.

Evaluation across different tax document types revealed that detection accuracy was higher for individual income tax returns (F1-score of 0.889) than for business tax declarations (F1-score of 0.818), reflecting differences in complexity and data characteristics. The NLP-enhanced framework demonstrated strong capabilities in identifying sophisticated tax evasion techniques that manipulate textual content, thus filling a crucial gap in existing fraud detection systems that primarily focus on numerical anomalies.

Moreover, the integration of named entity recognition to extract financial entities facilitated the detection of suspicious transactional relationships, significantly boosting detection performance. The analysis of feature importance underscored that semantic relationship features hold greater relevance in business tax declarations, while numerical consistency features are more critical for individual income tax returns.

5.2. Weaknesses in the Current Methodology

Despite the significant improvements achieved by the proposed NLP-enhanced anomaly detection framework, several limitations must be acknowledged. First, the computational demands are substantially higher than those of traditional methods, with a training duration of 14.5 hours and an inference time of 45.3 ms per document, which may restrict real-time deployment in environments with limited computational resources.

Second, the model faces challenges in domain adaptation when applied across different tax jurisdictions, often requiring retraining or fine-tuning to accommodate changes in regulatory frameworks. Performance degradation observed in business tax declarations indicates difficulties in managing complex document structures and diverse financial reporting formats.

Third, the reliance on tax-domain-specific BERT embeddings necessitates regular updates to the model to keep pace with evolving tax terminology and reporting standards. Geographic disparities in detection accuracy suggest potential biases, as regions underrepresented in the training data exhibit lower performance.

Additionally, the current system is susceptible to adversarial attacks aimed at manipulating linguistic patterns while preserving numerical consistency, posing a risk to detection robustness. Privacy concerns also present barriers, since the detailed textual analysis essential for anomaly detection may conflict with data protection regulations in certain jurisdictions.

Finally, integration with existing tax authority infrastructure poses operational challenges that could hinder broad adoption despite the approach's demonstrated effectiveness.

Acknowledgments: I wish to express my deepest appreciation to Yida Zhu, Yining Zhang, and Yuexing Chen for their pioneering work on financial sentiment analysis applied to the detection of abnormal stock market volatility, as detailed in their article entitled "Leveraging Financial Sentiment Analysis for Detecting Abnormal Stock Market Volatility: An Evidence-Based Approach from Social Media Data". Their innovative methodologies and insightful findings have greatly enhanced my comprehension of sophisticated anomaly detection techniques and have served as a valuable source of inspiration for my own investigations into tax fraud detection. I would like to thank the researchers whose work on nursing staff allocation optimization provided valuable insights into data-driven modeling in anomaly detection optimization employing time series data analysis, detailed in their article titled "Optimization of Nursing Staff Allocation in Elderly Care Institutions: A Time Series

Data Analysis Approach". Their in-depth analysis combined with predictive modeling has profoundly enhanced my insights into data-driven anomaly detection and motivated my methodological choices in research.

References

1. L. Yan, S. Zhou, W. Zheng, J. Chen, "Deep Reinforcement Learning-based Resource Adaptive Scheduling for Cloud Video Conferencing Systems," 2024, doi: 10.53469/wjimt.2024.07(06).19.
2. J. Chen, L. Yan, S. Wang, W. Zheng, "Deep Reinforcement Learning-Based Automatic Test Case Generation for Hardware Verification," *J. Artif. Intell. Gen. Sci.*, vol. 6, no. 1, pp. 409–429, 2024, doi: 10.60087/jaigs.v6i1.267.
3. L. Yan, Y. Wang, L. Guo, K. Qian, "Enhanced Spatio-Temporal Attention Mechanism for Video Anomaly Event Detection," 2025, doi: 10.20944/preprints202504.1623.v1.
4. S. Xia, Y. Zhu, S. Zheng, T. Lu, K. Xiong, "A Deep Learning-based Model for P2P Microloan Default Risk Prediction," *Spectrum Res.*, vol. 4, no. 2, 2024.
5. S. Li, H. Xu, T. Lu, G. Cao, X. Zhang, "Emerging technologies in finance: Revolutionizing investment strategies and tax management in the digital era," *Spectrum Res.*, vol. 4, no. 2, 2024.
6. Y. Liu, Y. Xu, S. Zhou, "Enhancing User Experience through Machine Learning-Based Personalized Recommendation Systems: Behavior Data-Driven UI Design," *Appl. Comput. Eng.*, vol. 112, pp. 42–46, 2024, doi: 10.54254/2755-2721/2024.17905.
7. Y. Xu, Y. Liu, J. Wu, X. Zhan, "Privacy by Design in Machine Learning Data Collection: An Experiment on Enhancing User Experience," *Appl. Comput. Eng.*, vol. 97, pp. 64–68, 2024, doi: 10.54254/2755-2721/97/20241388.
8. X. Xu, Z. Xu, P. Yu, J. Wang, "Enhancing user intent for recommendation systems via large language models," arXiv preprint arXiv:2501.10871, 2025. doi: 10.48550/arXiv.2501.10871.
9. L. Li, K. Xiong, G. Wang, J. Shi, "AI-Enhanced Security for Large-Scale Kubernetes Clusters: Advanced Defense and Authentication for National Cloud Infrastructure," J. Theory Pract. Eng. Sci., vol. 4, no. 12, pp. 33–47, 2024.
10. P. Yu, Z. Xu, J. Wang, X. Xu, "The application of large language models in recommendation systems," arXiv preprint arXiv:2501.02178, 2025.
11. J. Yi, Z. Xu, T. Huang, P. Yu, "Challenges and Innovations in LLM-Powered Fake News Detection: A Synthesis of Approaches and Future Directions," arXiv preprint arXiv:2502.00339, 2025.
12. T. Huang, Z. Xu, P. Yu, J. Yi, X. Xu, "A Hybrid Transformer Model for Fake News Detection: Leveraging Bayesian Optimization and Bidirectional Recurrent Unit," arXiv preprint arXiv:2502.09097, 2025.
13. J. Wang, X. Xu, P. Yu, Z. Xu, "Hierarchical Multi-Stage BERT Fusion Framework with Dual Attention for Enhanced Cyberbullying Detection in Social Media," in *Proc. 2024 4th Int. Conf. Artif. Intell., Robot. Commun. (ICAIRC)*, 2024, pp. 86–89, doi: 10.1109/ICAIRC64177.2024.10900203.
14. T. Huang, J. Yi, P. Yu, X. Xu, "Unmasking Digital Falsehoods: A Comparative Analysis of LLM-Based Misinformation Detection Strategies," arXiv preprint arXiv:2503.00724, 2025.
15. J. Weng, X. Jiang, "Research on movement fluidity assessment for professional dancers based on artificial intelligence technology," *Artif. Intell. Mach. Learn. Rev.*, vol. 5, no. 4, pp. 41–54, 2024, doi: 10.69987/AIMLR.2024.50404.
16. C. Jiang, G. Jia, C. Hu, "AI-driven cultural sensitivity analysis for game localization: A case study of player feedback in East Asian markets," *Artif. Intell. Mach. Learn. Rev.*, vol. 5, no. 4, pp. 26–40, 2024, doi: 10.69987/AIMLR.2024.50403.
17. D. Ma, "AI-driven optimization of intergenerational community services: An empirical analysis of elderly care communities in Los Angeles," *Artif. Intell. Mach. Learn. Rev.*, vol. 5, no. 4, pp. 10–25, 2024, doi: 10.69987/AIMLR.2024.50402.
18. D. Ma, Z. Ling, "Optimization of Nursing Staff Allocation in Elderly Care Institutions: A Time Series Data Analysis Approach," *Ann. Appl. Sci.*, vol. 5, no. 1, 2024.
19. S. Zheng, Y. Zhang, Y. Chen, "Leveraging Financial Sentiment Analysis for Detecting Abnormal Stock Market Volatility: An Evidence-Based Approach from Social Media Data," *Acad. Nexus J.*, vol. 3, no. 3, 2024.
20. C. Zhang, W. Lu, C. Ni, H. Wang, J. Wu, "Enhanced user interaction in operating systems through machine learning language models," in *Proc. Int. Conf. Image, Signal Process. Pattern Recognit. (ISPP)*, vol. 13180, 2024, pp. 1623–1630, doi: 10.1117/12.3033610.
21. H. Wang, J. Wu, C. Zhang, W. Lu, C. Ni, "Intelligent security detection and defense in operating systems based on deep learning," *Int. J. Comput. Sci. Inf. Technol.*, vol. 2, no. 1, pp. 359–367, 2024.
22. W. Lu, C. Ni, H. Wang, J. Wu, C. Zhang, "Machine learning-based automatic fault diagnosis method for operating systems," 2024, doi: 10.53469/wjimt.2024.07(02).12.
23. C. Zhang, W. Lu, J. Wu, C. Ni, H. Wang, "SegNet network architecture for deep learning image segmentation and its integrated applications and prospects," *Acad. J. Sci. Technol.*, vol. 9, no. 2, pp. 224–229, 2024, doi: 10.54097/rfa5x119.
24. J. Wu, H. Wang, C. Ni, C. Zhang, W. Lu, "Data Pipeline Training: Integrating AutoML to Optimize the Data Flow of Machine Learning Models," in *Proc. 2024 7th Int. Conf. Adv. Algorithms Control Eng. (ICAACE)*, 2024, pp. 730–734, doi: 10.1109/ICAACE61206.2024.10549260.
25. J. Wu, H. Wang, C. Ni, C. Zhang, W. Lu, "Case Study of Next-Generation Artificial Intelligence in Medical Image Diagnosis Based on Cloud Computing," J. Theory Pract. Eng. Sci., vol. 4, no. 02, pp. 66–73, 2024.

26. B. Wu, C. Shi, W. Jiang, and K. Qian, "Enterprise Digital Intelligent Remote Control System Based on Industrial Internet of Things," *World J. Innov. Manag. Technol.*, vol. 7, no. 2, 2024, doi: 10.53469/wjimt.2024.07(02).09.
27. C. Fan, Z. Li, W. Ding, H. Zhou, K. Qian, "Integrating artificial intelligence with SLAM technology for robotic navigation and localization in unknown environments," Appl. Comput. Eng., vol. 77, pp. 245–250, 2024.