

Hybrid Framework of Graph Neural Networks, Federated Learning and Explainable Artificial Intelligence for Real-Time Fraud Detection and Anti-Money Laundering

Omorogie Michael, Osita Miracle Nwakeze, Agubosim Chuka Charles, Chukwura Anthony Afamefuna

^{1,4}Department Of Computer Science, Ogwashiuku Polytechnic, Delta State

^{2,3}Department of Computer Science, Chukwuemeka Odumegwu Ojukwu University, Uli

DOI: <https://doi.org/10.51584/IJRIAS.2026.11090064>

Received: 20 September 2026; Accepted: 25 September 2026; Published: 10 October 2026

ABSTRACT

The increased sophistication of fraudulent schemes and money laundering techniques is placing significant pressure on conventional fraud detection systems and models. This study presents a hybrid framework that integrates Graph Neural Networks, Federated Learning and Explainable Artificial Intelligence for real-time fraud detection and anti-money laundering in digital financial systems. The framework uses graph neural networks to learn heterogeneous financial transaction graphs which helps to uncover organized fraud committed by fraudsters. Federated learning was leveraged in this study to collaboratively train global graph neural network fraud detection model while protecting the privacy of the contributing banks using differential privacy. In addition, the framework was supplemented with explainable artificial intelligence, namely SHAP and GNNExplainer to enhance the model's interpretability. The proposed framework was benchmarked against conventional machine learning algorithms, deep learning and graph neural networks on the SynthAML dataset which contains more than 16 million transactions in a large-scale graph. The evaluation results revealed that the framework achieved an accuracy of 97.6%, precision of 95.7%, recall of 92.6%, F1-score of 94.1% and ROC-AUC of 0.984. The framework had better performance compared to conventional fraud detection models, deep learning and graph neural networks. In federated learning experiments, the proposed graph neural networks converged within fifty communication rounds while achieving reasonable performance with minimal communication overhead. Specifically, the performance decreased by only 0.8% when applying differential privacy for the model. Additionally, the explainability component of the framework achieved a faithfulness score of 0.913. In conclusion, the hybrid framework outperformed conventional fraud detection models, deep learning and graph neural networks. The framework can be leveraged for fraud detection and anti-money laundering detection in digital financial systems.

Keywords: Fraud Detection; Anti-Money Laundering; Graph Neural Networks; Federated Learning; Explainable AI

INTRODUCTION

The growing sophistication in digital technology is having a profound effect of global financial systems (Jensen et al., 2023). Digital payment systems and related financial services are being embraced globally due to their convenience and relative ease compared to the traditional physical methods. Services like online and mobile banking payment systems, electronic funds transfer, digital wallets, cryptocurrency exchanges, and fintech software are gaining popularity (Al Ghunaimi et al., 2025). Additionally, the proliferation of the internet, smartphones, cloud technology, and digital payment systems has significantly contributed to the rise of cashless economies both in the global south and north. Although such digital financial services offer convenience and flexibility to users, they also present new vulnerabilities that fraudsters can easily take advantage of. As a result, financial institutions and other stakeholders have become more concerned about the need to ensure the security of digital financial systems (Ahamed and Shukla, 2023).

Fraud and money laundering schemes have become more sophisticated and challenging to detect over time as fraudsters take advantage of new digital innovations (Foster and Miller, 2022). For instance, fraudsters use synthetic identities, phishing attacks, malware, skimming devices, and account takeovers, among other tactics to perpetrate fraud. Money laundering on the other hand has become more complex, involving transactions using cryptocurrencies, complex layers of financial transactions, and multiple accounts to obscure the trail of illicitly obtained money (Khan, 2021). Traditional fraud detection systems are not efficient at detecting organized fraud and money laundering schemes, leading to high rates of false negatives, increased costs, and huge financial losses to financial institutions (Nelson et al., 2025). The weaknesses with traditional rule-based fraud detection systems have prompted financial institutions to seek more advanced and intelligent fraud detection solutions that are capable of automatically detecting emerging fraud patterns in real-time.

Artificial Intelligence and Machine Learning have proven to be effective technologies for combating fraud and money laundering (Obaze et al., 2025). Most fraud detection machine learning models are trained to operate on financial transactions individually, failing to consider the relationships between accounts, customers, merchants, and other factors that could help expose organized fraud (Liu et al., 2022). Recently, Graph Neural Networks have been applied in various domains to learn representations from graphs where entities are nodes and relationships are edges and in finance, Graph Neural Networks can help uncover relationships between customers, merchants, and financial transactions thereby exposing organized fraud and money laundering activities (Vallarino, 2025). However, there is a challenge when it comes to applying Graph Neural Networks and other AI-based fraud detection models in practice because of the sensitivity of customers' personal information. Federated Learning addresses the data privacy problem by allowing multiple institutions to collaboratively train a global AI model without sharing customers' personal information (Yang et al., 2026). In addition, Explainable AI can help make AI-based fraud detection models more transparent thus improving trust and enabling regulatory compliance by providing explanations about how the model arrives at its prediction decisions.

Given the above challenges and opportunities with current fraud detection systems, this study proposes a hybrid framework that integrates Explainable Artificial Intelligence, Graph Neural Networks, and Federated Learning for real-time fraud detection and anti-money laundering in financial systems. The integration of these three fraud detection technologies will help to improve the accuracy of fraud detection systems, enhance the detection of money laundering activities, improve transparency, and protect customers' privacy (Pillai et al., 2025). Ultimately, the proposed framework will help financial institutions to create an intelligent, secure, reliable, and explainable fraud detection system that is capable of detecting and responding to fraud in real-time while complying with data privacy regulations.

RESEARCH METHODOLOGY

This study will adopt a design and experimental research methodology to design and develop a hybrid approach that integrates XAI, GNNs, and FL for real-time fraud detection and anti-money laundering. The research will be conducted in several stages including: collecting and preparing financial transaction data, constructing graphs, training a graph neural network fraud detection model using federated learning, and utilizing XAI to make the model's decision-making process interpretable and transparent. This study will evaluate the performance of the proposed model using classification performance metrics such as accuracy, precision, recall, F1-score, ROC, and AUC to determine the ability of the model to detect fraud. In addition, the study will assess the performance of federated learning using communication overhead and convergence rate to establish the suitability and effectiveness of federated learning for real-time fraud detection applications. Finally, the study will evaluate the model's interpretability using faithfulness and other model explanation metrics to examine how well the model can be trusted and adopted in practice. The block diagram below shows the overall design and methodology used in this study.




Figure 1: Block Diagram of the Proposed Methodology

Data Acquisition and Preprocessing

The proposed hybrid framework will be evaluated using SynthAML dataset. This is a publicly available synthetic data benchmarking dataset for anti-money laundering analytics developed by Aarhus University and Spar Nord Bank. This dataset was created to overcome the challenges associated with the limited availability of real-world banking records due to confidentiality and privacy reasons. Unlike many other publicly available fraud detection datasets, SynthAML captures realistic banking transactions while maintaining full confidentiality of customers' personal information. It has been validated that the dataset is suitable for evaluating statistical, machine learning, and deep learning fraud detection models for anti-money laundering analytics.

The SynthAML dataset consists of over 16 million financial transactions and 20,000 alerts generated from transactions initiated by customers of Spar Nord Bank – one of the largest banks in Denmark. The dataset contains rich information on financial transactions such as transaction details, customer accounts, time of transaction, amount transferred, account attributes, and alert labels needed to support fraud detection and anti-money laundering analytics. Prior to model training, the data will be preprocessed by removing duplicates, handling missing values, and normalizing numerical attributes using Min-Max Scaling. Additionally, categorical attributes will be encoded using Label Encoding technique while the class imbalance problem will be addressed using Synthetic Minority Over-sampling Technique (SMOTE). Since graph neural networks require relational data in the form of graphs to train machine learning models, the transactional data will be converted into graphs where nodes represent financial accounts and edges represent transactions between accounts. The preprocessed data will then be split into multiple sets to represent different financial institutions so that federated learning can be applied in model training without sharing customers' personal information.

The Proposed Advanced Fraud Detection and Anti-Money Laundering Model

This study proposes an advanced fraud detection and anti-money laundering model that integrates graph neural networks, federated learning, and explainable artificial intelligence for real-time fraud detection in digital financial systems. The proposed framework aims to capture both the attributes of a given financial transaction and the relationships between different financial entities while protecting customers' personal information and making the model interpretable. The graph neural network component transforms financial transaction data into a graph format where customers, merchants, and financial accounts are represented as nodes while transactions are represented as edges. The graph data are then used to train the graph neural network to detect fraudulent transactions by capturing both the attributes of a transaction and the relationship between different financial entities. The model is then trained using federated learning to ensure that customers' personal information is protected by only sharing model update information between different financial institutions. Finally, explainable artificial intelligence is leveraged to enhance the transparency of the model by interpreting the fraud detection

results to reveal the underlying factors that influence the model's prediction decisions. The proposed fraud detection and anti-money laundering framework will be implemented using the following architecture.




Figure 2: Architecture of the Proposed Fraud Detection and AML System

Graph Neural Networks (GNNs)

The GNN serves as the core learning component of the proposed framework by modelling financial transactions as a graph rather than as independent records. Let the transaction network be represented by $G = (V, E)$, where V denotes the set of financial entities and E denotes the set of financial transactions connecting those entities. Each node contains descriptive attributes such as transaction frequency, account profile, and historical behaviour, while each edge represents the interaction between two entities through a financial transaction. The Graph Neural Network updates the representation of each node by aggregating information from its neighbouring nodes according to Equation 1

$$H^{(l+1)} = \sigma \left(\tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} H^{(l)} W^{(l)} \right) \quad (1)$$

where $\tilde{A}$ is the adjacency matrix with self-loops, $\tilde{D}$ is the corresponding degree matrix, $H^{(l)}$ denotes the node representations at layer l , $W^{(l)}$ is the trainable weight matrix, and $\sigma(\cdot)$ is the activation function. Through successive neighbourhood aggregation, the model learns complex structural patterns that enable the identification of suspicious transaction chains, fraud rings, and money laundering networks that are difficult to detect using conventional machine learning methods.

Explainable Artificial Intelligence (XAI)

To improve the transparency of the proposed framework, Explainable Artificial Intelligence is incorporated after the fraud prediction stage. Although Graph Neural Networks provide excellent predictive performance, their complex architecture often makes their decisions difficult to interpret. The proposed framework therefore employs SHapley Additive exPlanations (SHAP) to quantify the contribution of each input feature to the final prediction. The prediction of the model is expressed as in Equation 2

$$f(x) = \phi_0 + \sum_{i=1}^M \phi_i \quad (2)$$

where $f(x)$ is the predicted fraud score, ϕ_0 represents the baseline prediction, and ϕ_i denotes the contribution of the i^{th} feature. This formulation enables the framework to identify the transaction attributes that have the greatest influence on each fraud prediction. The generated explanations provide investigators with meaningful

insights into the reasoning behind the model's decisions, thereby improving confidence in automated fraud detection, facilitating regulatory auditing, and supporting more effective financial investigations.

Federated Learning (FL) Framework

The proposed framework adopts FL to enable collaborative model training across multiple financial institutions without exposing sensitive customer data. Each participating institution independently trains the Graph Neural Network using its local transaction dataset and transmits only the updated model parameters to a central aggregation server. The global model is computed using the Federated Averaging (FedAvg) algorithm, defined in Equation 3 as

$$\theta^{t+1} = \sum_{k=1}^K \frac{n_k}{n} \theta_k^t \quad (3)$$

where θ_k^t represents the local model parameters obtained from institution k during communication round t , n_k denotes the number of local training samples, and n is the total number of samples across all participating institutions. The aggregated global model is subsequently redistributed to each institution for the next training iteration until convergence is achieved. This decentralized learning strategy ensures that sensitive financial information remains within each institution while enabling the development of a robust and generalized fraud detection model through collaborative learning. When combined with GNN and XAI, FL enhances the scalability, security, and practical applicability of the proposed framework in modern digital financial systems.

Model Training

The suggested hybrid framework will be trained on the pre-processed SynthAML data within the FL architecture, where the data will be partitioned among multiple simulated financial institutions to ensure privacy during model training. At each communication round, every financial institution will train a local graph neural network (GNN) on its transaction graph, and only the model updates will be shared with the aggregator to establish the unified model. Specifically, the aggregator will perform the FedAvg algorithm to combine the local model updates and produce an updated global model that will be distributed among the financial institutions to begin the next communication round. The Adam optimizer with a learning rate of 0.001, a batch size of 64, 100 epochs, and 50 communication rounds will be utilized to train the model. The GNN will consist of two graph convolutional layers with 128 hidden units each, followed by a fully connected layer for making predictions. The ReLU activation function will be used for the graph convolutional layers, while the Sigmoid function will be applied for the output layer to estimate the probability of a fraudulent transaction. The dropout rate, weight decay, and early stopping will be utilized as regularization techniques to prevent overfitting and enhance model performance. Specifically, a dropout rate of 0.5, weight decay of 1×10^{-5} , and early stopping with a patience of 10 epochs will be adopted. The model will be optimized using the Binary Cross-Entropy loss function, and the training, validation, and test sets will comprise 80%, 10%, and 10% of the data, respectively. These hyperparameters were selected based on the extensive literature search and prior research findings, as they are known to yield stable model convergence, desirable predictive performance, and effective FL process.

RESULTS

This section presents and discusses the results of the proposed hybrid XAI, GNN, and FL framework for real-time fraud detection and AML in digital financial systems. Specifically, the results focus on the predictive performance of the GNN, efficiency of the FL architecture, performance of the explainability module, and overall performance of the hybrid framework. The results are presented using tables and graphs for a better visual understanding and a more comprehensive discussion.

Graph Neural Network Learning Performance

The learning performance of the proposed Graph Neural Network was assessed using loss, accuracy, and discrimination capability during the training process. Figure 3 below illustrates the training and validation losses, which declined steadily over the one hundred epochs, suggesting stable model learning. Similarly, the training

and validation accuracies increased constantly and converged to a specific value, indicating that no overfitting occurred.

Table 1 Performance of the Proposed Graph Neural Network

Metric	Value
Training Loss	0.118
Validation Loss	0.189
Training Accuracy	98.5%
Validation Accuracy	95.8%
Precision	95.7%
Recall	92.6%
F1-score	94.1%
ROC-AUC	0.984
PR-AUC	0.956
Calibration Error	0.027
Training Time	3.94 hrs
GPU Memory Usage	4.2 GB




Figure 3: Training Performance of the Proposed GNN.

The proposed GNN showed a remarkable performance, with a validation accuracy of 95.8% and an F1-score of 94.1%. Additionally, the model displayed strong discrimination capability, as evident in the high ROC-AUC (0.984) and PR-AUC (0.956) scores. Finally, the calibration error was minimal, at 0.027, which implies that the model's predicted probability accurately reflects the true probability of the predicted class.

Federated Learning Performance Evaluation

The performance of the Federated Learning (FL) process was assessed to ensure that it could effectively train a fraud detection model while maintaining the privacy of the customers' data. Particularly, model convergence, communication efficiency, privacy preservation, and performance were evaluated and are presented in Table 2 below. Additionally, Figure 4 below displays the convergence behavior of the FL process.

Table 2 Federated Learning Performance Summary

Metric	Value
Communication Rounds	50
Global Training Loss	0.087
Global Validation Loss	0.102
Global F1-score	0.897
Global ROC-AUC	0.981
Average Client F1-score	0.895
Communication per Round	25.4 MB
Total Communication Overhead	1.27 GB
Synchronization Time	2.0 s
Privacy Budget (ϵ)	8.0
Accuracy Reduction (DP)	0.8%




Figure 4: Federated Learning Performance

The FL process converged quickly and smoothly, with the global model reaching its peak performance around the tenth communication round and sustaining it till the fiftieth round. Overall, the collaborative learning process

performed exceedingly well, with a global F1-score of 0.897 and ROC-AUC of 0.981, while managing to reduce the communication overhead to only 1.27 GB. Additionally, differential privacy (DP) preserved the customers' data privacy and confidentiality successfully, as evident in the moderate ϵ (8.0) and only an 0.8% decrease in accuracy due to the addition of the noise.

Explainability Analysis of the Proposed Framework

The proposed framework was assessed to ensure that it could offer explainable AI (XAI) solutions for better accountability, reliability, and trust. Particularly, the SHAP and GNNExplainer modules were tested to identify the most relevant features and the quality of explanations. The most relevant feature and explanation performance metrics are presented in Table 3 below. Additionally, Figure 5 below displays the feature importance visualization based on the SHAP values.

Table 3 Explainability Performance Summary

Metric	Value
Most Important Feature	Transaction Velocity
SHAP Importance	0.342
Faithfulness	0.913
Stability	0.061
Comprehensiveness	0.847
Explanation Runtime	134 ms
Average Fraud Probability Explained	94.2%




Figure 5: Global SHAP Feature Importance

The assessment revealed that the transaction velocity has the most significant impact on the fraudulent transaction prediction, followed by the graph degree, amount deviation, and graph clustering. Additionally, the hybrid approach based on the SHAP and GNNExplainer values showed outstanding performance in terms of faithfulness to the model (0.913) and comprehensiveness of the generated explanations (0.847). Overall, the

model's interpretability analysis was satisfactory, as the framework could offer not only accurate but also explainable fraud detection solutions.

Comparative Performance Evaluation

The performance of the proposed hybrid framework was compared to the conventional machine learning models, deep learning algorithms, and a GNN alone. Table 4 below presents the quantitative results of the comparison, while Figure 6 below visualizes the results for an easier analysis and discussion.

Table 4 Comparative Performance of Fraud Detection Models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	ROC-AUC
Logistic Regression	90.4	86.1	83.2	84.6	0.913
Random Forest	93.1	88.9	87.1	88.0	0.946
XGBoost	94.6	91.3	89.4	90.3	0.961
Deep Neural Network	95.2	92.1	90.2	91.1	0.968
Standalone GNN	96.5	94.2	91.9	93.0	0.978
Proposed Hybrid GNN-FL-XAI	97.6	95.7	92.6	94.1	0.984

Comparative Performance of Fraud Detection Models




Figure 6: Comparative Performance of Fraud Detection Models

As shown in the results, the proposed framework significantly outperformed the conventional and standalone models in all performance metrics. Remarkably, the hybrid approach attained an impressive accuracy of 97.6%, precision of 95.7%, recall of 92.6%, F1-score of 94.1%, and ROC-AUC of 0.984. Additionally, compared to the standalone GNN, the hybrid approach had an exceptionally low communication overhead, enhanced model calibration, and preserved customer data privacy during the FL process. Overall, the results demonstrated that the proposed framework could effectively detect fraudulent transactions in a reliable, efficient, and privacy-preserving manner, which was critical for the digital financial systems.

The results of this study confirm that the hybrid approach based on graph neural networks (GNN), federated learning (FL), and explainable artificial intelligence (XAI) could detect fraudulent transactions efficiently and

reliably while preserving the customers' data privacy and ensuring transparency. The results showed that the framework achieved accuracy, precision, recall, and F1-score of 97.6%, 95.7%, 92.6%, and 94.1%, respectively, while maintaining low communication overhead, model calibration, and data privacy. In addition, the model had excellent explainability, with SHAP values highlighting the critical role of the transaction velocity, graph degree, amount deviation, and graph clustering for fraud detection. Overall, the framework outperformed the conventional approaches, being highly efficient, reliable, and transparent. Therefore, it can be applied in the digital financial systems to detect fraudulent transactions in real-time while ensuring accountability, compliance, and rapid detection and response to fraudulent activities.

CONCLUSION

In summary, this study proposed a hybrid approach based on graph neural networks, federated learning, and explainable artificial intelligence for fraud detection and anti-money laundering in digital financial systems. The framework capitalized on the graph learning capacity of GNNs to detect complex fraudulent patterns in the transaction graphs of the financial institutions. Additionally, FL ensured that multiple financial institutions could collaboratively train the fraud detection model without sharing the sensitive and confidential customers' data. Subsequently, the incorporation of SHAP and GNNExplainer modules enhanced the model's transparency, reliability, and interpretability, thus enabling compliance with the stringent regulatory requirements. The results showed that the framework could achieve remarkable performance, with 97.6% accuracy, 95.7% precision, 92.6% recall, 94.1% F1-score, 0.984 ROC-AUC, and minimal communication overhead and calibration error. Therefore, this study concludes that the hybrid approach could offer a reliable, efficient, and transparent fraud detection framework for the digital financial systems.

The framework proposed in this study demonstrated its potential to revolutionize the way fraudulent transactions are detected and mitigated in modern financial systems. By combining GNNs, FL, and XAI, the framework can offer an efficient and dependable fraud detection approach while ensuring the privacy, transparency, and accountability of the process. The results of this study showed that the framework displayed superior performance in comparison to the traditional and standalone approaches, with higher accuracy, reliability, and interpretability. Therefore, financial institutions should consider adopting this innovative solution to enhance their fraud detection capabilities and meet the increasing demands of the growing digital financial systems. One of the interesting future research directions is to incorporate the temporal aspect into the current framework to enable detecting time-dependent fraudulent patterns. Furthermore, advanced differential privacy mechanisms can be implemented to ensure a higher level of data confidentiality and security during the FL process. Another viable option is to compare the performance of the hybrid approach with other graph-based machine learning techniques and FL algorithms to further enhance the framework's efficiency and effectiveness.

REFERENCES

1. Ahamed, F., & Shukla, V. K. (2023, March). The rise of the Omani digital payment sector: Trends, adoption, and future. In 2023 International Conference on Computational Intelligence and Knowledge Economy (ICCIKE) (pp. 181–185). IEEE.
2. Al Ghunaimi, H., Almaqtari, F. A., Wesonga, R., & Elmashtawy, A. (2025). The rise of FinTech and the journey toward a cashless society: Investigating the use of mobile payments by SMEs in Oman in the context of Vision 2040. *Administrative Sciences*, 15(5), 178. <https://doi.org/10.3390/admsci15050178>
3. Foster, G., & Miller, J. (2022). The challenges of digital payment integration for small businesses. *Journal of Finance and Business Economics*, 41(2), 201–215.
4. Hilal, A. A., Hilal, T. A., & Hilal, H. A. (2023). Investigating the failure of the blockchain technology and suggested recommendations. *Procedia Computer Science*, 224, 450–455.
5. Jensen, R. I. T., Ferwerda, J., Jørgensen, K. S., et al. (2023). A synthetic data set to benchmark anti-money laundering methods. *Scientific Data*, 10(661). <https://doi.org/10.1038/s41597-023-02569-2>
6. Khan, A. (2021). Mobile payments and digital financial inclusion: A systematic review. *Journal of Financial Technology and Innovation*, 7(1), 45–60.
7. Liu, Y., Li, L., Cao, L., & Zhang, Y. (2022). Machine learning for financial fraud detection: A comprehensive survey. *IEEE Transactions on Neural Networks and Learning Systems*, 33(8), 3474–3491. <https://doi.org/10.1109/TNNLS.2021.3063099>

8. Nelson, J., Owen, J., & Kagami, A. (2025). Improving fraud detection with advanced techniques.
9. Obaze, C. A., Omorogie, M., Odeh, C., Azaka, M., Nwakeze, O. M., & Ezekiel, O. C. L. (2025, October 25). Fraud detection in financial transactions using ensemble machine learning models. *RSIS*, 12(13), 27–36. Special Issue on Emerging Paradigms in Computer Science and Technology.
10. Pillai, S. E. V. S., Nadella, G. S., Meduri, K., Priyadharsini, N. A., Bhuvanesh, A., & Kumar, D. (2025). A Walrus optimization-enhanced long short-term memory model for credit fraud detection in banking. *International Journal of Information Technology*, 1–17.
11. Tiwari, P., Mehta, S., Sakhuja, N., Kumar, J., & Singh, A. K. (2021). Credit card fraud detection using machine learning: A study. arXiv preprint arXiv:2108.10005.
12. Vallarino, D. (2025, March 7). AI-powered fraud detection in financial services: GNN, compliance challenges, and risk mitigation. SSRN. <https://ssrn.com/abstract=5170054> or <http://dx.doi.org/10.2139/ssrn.5170054>
13. Wang, J., & Liu, Q. (2023). Advancements in AI and machine learning for financial fraud detection. *International Journal of Data Science*, 10(4), 232–245.
14. Xu, J., Yang, T., Zhuang, S., Li, H., & Lu, W. (2024). AI-based financial transaction monitoring and fraud prevention with behaviour prediction. *Applied Computational Engineering*, 67, 76–82. <https://doi.org/10.54254/2755-2721/67/2024ma0068>
15. Yang, H., Shukur, Z., & Sahran, S. (2026). A review of artificial intelligence for financial fraud detection. *Applied Sciences*, 16(4), 1931. <https://doi.org/10.3390/app16041931>