
The Application of Machine Learning Models in Fraud Detection and Prevention Across Digital Banking Channels and Payment Platforms

OWEN OPEN ACCESS PUBLICATIONS

1-7

©owenpress:

Article reuse guidelines:
<https://owenpress.com/>



Salim Khatib¹

Abstract

This paper presents a comprehensive investigation into the application of machine learning models for fraud detection and prevention across digital banking channels and payment platforms. It examines the unique characteristics of transactional data streams in online banking, mobile banking, peer-to-peer transfers, and emerging payment modalities such as digital wallets and real-time payment rails. We analyze the efficacy of supervised, semi-supervised, and unsupervised learning algorithms under various feature representations, including temporal sequence embeddings, graph-based relational features, and hierarchical behavioral signatures. A detailed exposition of one advanced mathematical formulation frames the detection problem as a stochastic optimization under adversarial perturbations, leveraging concepts from measure-theoretic probability, reproducing kernel Hilbert spaces, and robust statistical decision theory. The proposed framework integrates incremental learning to adapt to concept drift, and meta-learning strategies to transfer insights across heterogeneous channels. Experimental evaluation on large-scale synthetic and anonymized production datasets demonstrates that ensemble architectures combining deep representation learners with probabilistic graphical models can achieve significant improvements in detection latency and false-positive control, while preserving customer experience through adaptive risk-scoring thresholds. The findings underscore the trade-offs between interpretability, computational overhead, and adaptability in real-time fraud prevention systems. The paper concludes with recommendations for deployment architectures, data governance practices, and future research directions toward fully autonomous fraud resilience.

Introduction

The evolution of the digital economy has brought about a profound transformation in the way financial transactions are conducted, especially within the domain of digital banking and payment platforms (1). The ubiquity of smartphones, the widespread availability of high-speed internet, and the increasing adoption of application programming interfaces (APIs) have collectively fueled a surge in mobile wallets, contactless payments, peer-to-peer (P2P) payment applications, and embedded finance services integrated directly within consumer and business-facing platforms. This technological convergence has radically increased the speed, scale, and complexity of financial interactions, giving rise to voluminous transactional data streams that are both rich in behavioral patterns and prone to exploitation by sophisticated threat actors. Consequently, the imperative to develop intelligent, scalable, and proactive fraud detection systems has never been more critical. (2)

Financial institutions now face a dual mandate: to facilitate seamless, real-time transactions while simultaneously ensuring the integrity and security of each interaction. The stakes

are exceedingly high. On one hand, delays or disruptions in processing transactions can result in customer dissatisfaction and attrition (3). On the other hand, insufficient detection of fraudulent activity can lead to financial losses, reputational damage, and regulatory penalties. Traditional fraud detection systems, primarily based on deterministic rules and static thresholds, are increasingly ill-suited to meet these demands. These rule-based engines, though easily interpretable and operationally simple, exhibit significant limitations in adapting to rapidly evolving fraud tactics, which often involve multistep, obfuscated patterns that mimic legitimate behavior to evade detection. (4)

The limitations of traditional systems are particularly pronounced when fraudsters employ techniques such as synthetic identity fraud, where fictitious identities composed of real and fabricated information are used to create new accounts, or account takeovers, in which legitimate accounts are compromised through credential stuffing or

¹Al-Balqa' Applied University, Department of Computer Systems Engineering, Queen Rania Street 45, Salt, Jordan

phishing attacks. Social engineering tactics, often leveraging psychological manipulation and real-time interaction, further complicate the detection process, as do fraud schemes that exploit the inherent latency between transaction initiation and completion. Rule-based systems are inherently reactive and brittle in such contexts; they fail to generalize to new or previously unseen fraud scenarios and require continuous manual tuning by domain experts, which is both labor-intensive and slow. (5)

Machine learning (ML) offers a compelling alternative. By leveraging statistical modeling, pattern recognition, and automated learning from historical data, ML-based fraud detection systems can identify complex and subtle anomalies that may not conform to pre-defined rules. Supervised learning techniques, which rely on labeled datasets, are particularly effective when historical records of fraudulent and legitimate transactions are available (6). These methods can include logistic regression, gradient-boosted decision trees, support vector machines, and deep neural networks, all of which can be trained to discriminate between benign and malicious behavior with high accuracy. In situations where labeled data is scarce or incomplete, unsupervised learning methods such as clustering, density estimation, and autoencoders can be employed to discover anomalous transaction patterns indicative of fraud. Semi-supervised and self-supervised learning further extend these capabilities by leveraging unlabeled data to enhance learning outcomes. (7)

Nonetheless, deploying ML in high-stakes, real-time financial environments introduces a new set of challenges. Among these, the problem of extreme class imbalance is especially acute. In typical financial datasets, fraudulent transactions may constitute less than 0.1% of the total volume (8). This imbalance skews model training and often results in high overall accuracy at the expense of very low precision or recall on the minority class. A high false-negative rate means undetected fraud, while a high false-positive rate leads to unnecessary transaction declines and customer dissatisfaction. To mitigate this, various techniques such as resampling (oversampling the minority class or undersampling the majority), synthetic data generation (e.g., SMOTE), and cost-sensitive learning are employed (9). Anomaly detection models, which treat fraud as deviations from learned norms, offer an alternative that can be particularly effective in handling rare-event scenarios.

In addition to algorithmic considerations, the design of an ML-driven fraud detection system necessitates an architectural approach that integrates data ingestion, feature engineering, model training, online inference, feedback loops, and decision orchestration in a cohesive manner. Data ingestion pipelines must be capable of handling high-velocity, high-volume streams of transaction data from diverse sources including core banking systems, mobile applications, and third-party services (10). These pipelines must ensure data consistency, latency minimization, and real-time availability.

Feature engineering, often involving hundreds or thousands of variables, transforms raw data into meaningful representations that capture transaction semantics, temporal dynamics, behavioral signatures, and contextual metadata. Features such as transaction amount deviation, historical transaction frequency, device fingerprinting, geolocation variance, and customer interaction patterns are commonly used. (11)

Model training involves not just fitting statistical parameters but also hyperparameter tuning, cross-validation, model ensembling, and performance evaluation using metrics such as precision, recall, F1-score, area under the ROC curve (AUC), and detection latency. Importantly, models must be trained on data that reflects the most recent fraud trends, necessitating frequent retraining and validation cycles. Once trained, models are deployed to production environments where real-time inference must occur within stringent latency bounds—often measured in milliseconds (12). This requirement imposes constraints on model complexity, necessitating careful trade-offs between model expressiveness and computational efficiency.

Feedback loops are crucial for system adaptivity. Confirmed cases of fraud, customer complaints, transaction reversals, and manual investigations provide valuable labels that must be reintegrated into the training data to continuously refine model accuracy (13). This closed-loop system enables dynamic adaptation to adversarial behaviors, commonly known as concept drift. Concept drift refers to changes in the statistical properties of the input data over time, which can degrade model performance if not promptly addressed. Techniques such as sliding window retraining, drift detection algorithms, and model ensembles trained on temporally stratified data are employed to combat this phenomenon.

Interpretability and explainability are also paramount, especially in light of regulatory requirements such as the European Union's General Data Protection Regulation (GDPR) and the United States' Fair Credit Reporting Act (FCRA) (14). These regulations mandate that customers be informed of adverse decisions and the rationale behind them. Accordingly, fraud detection models must be auditable and interpretable. Techniques such as SHAP (SHapley Additive exPlanations), LIME (Local Interpretable Model-agnostic Explanations), and monotonic models are increasingly integrated into the model governance pipeline to provide post hoc explanations and transparency into model predictions. (15)

Infrastructurally, deploying ML-based fraud detection systems necessitates a robust and secure data architecture. This includes the use of distributed computing frameworks (e.g., Apache Spark, Flink), cloud-based storage solutions (e.g., Amazon S3, Google Cloud Storage), and model serving frameworks (e.g., TensorFlow Serving, ONNX Runtime). The following table outlines core components and technologies commonly employed: (16)

Table 1. Infrastructure Stack for Machine Learning-Based Fraud Detection Systems

Component	Purpose	Example Technologies
Data Ingestion	Stream processing and event sourcing	Apache Kafka, Amazon Kinesis, Flume
Feature Store	Centralized repository for engineered features	Feast, Tecton, Hopsworks
Model Training	Model development and offline evaluation	TensorFlow, PyTorch, Scikit-learn
Online Inference	Real-time prediction serving	NVIDIA Triton, BentoML, TensorRT
Monitoring	Model performance tracking and drift detection	Prometheus, Grafana, EvidentlyAI
Security and Compliance	Data governance, privacy, and auditing	HashiCorp Vault, Apache Ranger, GDPR Toolkit

Furthermore, the diversity of payment channels—ranging from point-of-sale (POS) systems and online banking portals to P2P apps and third-party merchant integrations—necessitates fraud detection systems that are channel-agnostic yet context-aware. For example, what constitutes anomalous behavior in a retail banking scenario (e.g., sudden large withdrawal) may differ significantly from that in a merchant payment context (e.g., rapid refund issuance). Therefore, the system must be capable of incorporating contextual priors and dynamically adjusting detection thresholds and model parameters based on transaction modality, customer profile, and real-time behavioral signals. (17)

To illustrate the variability in algorithmic performance across different contexts, consider the following comparative evaluation of supervised learning models on a real-world banking transaction dataset:

In this table, deep neural networks deliver the highest precision and recall but incur higher inference latency, which may be unsuitable for real-time deployment without dedicated hardware acceleration or model compression techniques. Lighter models like logistic regression and decision trees, while less expressive, offer faster inference and easier interpretability. (18)

Ultimately, the development of a robust, scalable, and compliant fraud detection system is a multidisciplinary endeavor that spans data science, software engineering, cybersecurity, and financial domain expertise. It requires continuous innovation to stay ahead of adversaries who are equally adept at exploiting emerging technologies. As the financial ecosystem continues to digitize and decentralize, the need for intelligent, automated, and context-aware fraud prevention systems will only intensify (19). Research into adversarial machine learning, federated model training, graph neural networks for entity resolution, and real-time behavioral modeling promises to shape the next generation of fraud detection systems that are not only accurate and fast but also secure, explainable, and resilient to manipulation.

Digital Fraud Landscape in Modern Payment Systems

Digital payment ecosystems encompass a diverse array of channels, each presenting distinct risk profiles and data modalities. Mobile banking applications generate rich event logs encompassing session metadata, geolocation signals, biometric authentication metrics, and touch-pattern dynamics (20). Web-based portals record device fingerprints, HTTP header variations, and cookie-based user journeys. Peer-to-peer transfers and digital wallet top-ups introduce graph-structured relationships among accounts, enabling network-centric analyses of transaction flows. Real-time payment networks impose stringent latency constraints on risk assessments, mandating sub-100ms inference pipelines (21). Concurrently, fraud activities manifest across multiple dimensions: velocity attacks exploit rapid successive transactions; account enumeration probes leverage credential stuffing; and mule networks create complex hub-and-spoke transfer patterns to launder illicit proceeds.

Effective countermeasures must integrate heterogeneous data sources, balancing the need for data granularity against privacy regulations and encryption mandates. Feature extraction pipelines must reconcile asynchronous logs, temporal sequences, and relational graphs into unified representations suitable for machine learning (22). Moreover, evolving regulations around data sovereignty and customer consent impose constraints on data retention and model interpretability. Consequently, system architects must embed privacy-enhancing techniques—differential privacy, federated learning, and secure multiparty computation—into the data pipeline to maintain compliance without sacrificing detection efficacy.

Machine Learning Architectures for Fraud Detection

Machine learning solutions for fraud detection span a continuum from classical statistical classifiers to advanced deep

Table 2. Model Evaluation Metrics Across Supervised Learning Approaches

Model Type	Precision	Recall	F1-Score	Latency (ms)
Logistic Regression	0.87	0.76	0.81	3.2
Random Forest	0.91	0.85	0.88	6.5
XGBoost	0.93	0.88	0.90	7.1
Deep Neural Network	0.95	0.92	0.93	12.8
LightGBM	0.92	0.87	0.89	5.9

learning frameworks (23). Supervised approaches such as gradient-boosted decision trees and support vector machines excel in scenarios with abundant labeled data, leveraging handcrafted features and sample reweighting to address severe class imbalance. Unsupervised techniques—including autoencoder-based anomaly detectors and clustering algorithms—enable detection of novel fraud patterns absent historical labels. Graph neural networks (GNNs) have emerged as a powerful paradigm for capturing relational fraud behaviors by embedding account-transaction graphs into continuous vector spaces, facilitating the identification of suspicious subgraph motifs and community anomalies. (24)

Ensemble methods that combine heterogeneous base learners can harness complementary strengths: for instance, coupling a light-weight tree-based model for preliminary screening with a deep sequence model for in-depth analysis of flagged events. Meta-learning approaches further enhance adaptability by learning to update model parameters rapidly in response to concept drift, using few-shot adaptation on recent labeled feedback. Reinforcement learning can optimize the orchestration of risk policies by modeling the trade-off between intervention costs and expected fraud losses, framing the problem as a Markov decision process where actions correspond to hold, reject, or require step-up authentication. (25)

Mathematical Modeling of Fraud Detection Dynamics

Let $\{X_t\}_{t \geq 0}$ denote the multivariate stochastic process representing transactional feature vectors observed in real time, where $X_t \in \mathbb{R}^d$. Define $Y_t \in \{0, 1\}$ as the indicator of fraudulent activity at time t . The detection problem can be formulated as minimizing the expected risk functional

$$\mathcal{R}(f) = \mathbb{E}[\ell(f(X_t), Y_t)] + \lambda \mathcal{C}(f),$$

where $f: \mathbb{R}^d \rightarrow [0, 1]$ is a probabilistic scoring function, ℓ is a convex surrogate loss (e.g., logistic loss), and \mathcal{C} is a regularization term capturing model complexity. To account for adversarial perturbations δ within an ℓ_p -ball of radius ϵ , we consider the robust counterpart

$$\min_f \mathbb{E} \left[\max_{\|\delta\|_p \leq \epsilon} \ell(f(X_t + \delta), Y_t) \right] + \lambda \mathcal{C}(f).$$

When f resides in a reproducing kernel Hilbert space with kernel k , representer theorems guarantee a solution of the form (26)

$$f(x) = \sum_{i=1}^n \alpha_i k(x, X_i),$$

where $\{\alpha_i\}$ are coefficients optimized under a distributionally robust optimization framework. To model temporal correlation and concept drift, augment the feature space with time-decay functions $w(t, i) = \exp(-\gamma|t - t_i|)$, leading to weighted empirical measures. The optimization can be solved via stochastic gradient descent with adversarial training steps:

$$\alpha^{(m+1)} = \alpha^{(m)} - \eta \nabla_{\alpha} \left[\ell(f(X_t + \delta^*), Y_t) \right],$$

In the limit of continuous-time observations, the process can be described by a stochastic differential equation

$$dX_t = \mu(X_t, t) dt + \sigma(X_t, t) dW_t + dJ_t, \quad (27)$$

where W_t is a Brownian motion capturing background noise and J_t is a jump process modeling abrupt attack events. The optimal scoring function f^* solves a Hamilton–Jacobi–Bellman equation in the space of value functions, yielding a dynamic programming formulation for real-time risk assessment under resource constraints.

Data Preprocessing and Feature Engineering

Effective fraud detection hinges on transforming raw transactional logs into discriminative features (28). Temporal aggregation windows capture spending velocity and periodicity, while sliding-window sketches approximate transaction frequency distributions. Behavioral biometrics—keystroke dynamics, touchscreen pressure profiles, and mouse movement statistics—provide additional anomaly detection signals. Graph-based features derive from constructing bipartite graphs between accounts and merchants, computing metrics such as personalized PageRank, eigenvector centrality, and motif-based anomaly scores (29). Feature normalization and outlier clipping ensure numerical stability for gradient-based learners. Feature selection can leverage mutual information estimators and Bayesian optimization to prune redundant attributes, reducing model inference latency. Streaming feature pipelines employ approximate algorithms—Bloom filters for cardinality estimation and Count–Min sketches for

frequency counts—to sustain high-throughput environments. (30)

Implementation and System Integration

Deployment of machine learning models in production fraud systems requires careful orchestration of microservices, message queues, and caching layers. Feature computation services must be horizontally scalable, with idempotent processing guarantees to handle replayed or late-arriving events. Inference microservices expose low-latency APIs capable of handling thousands of requests per second, leveraging model quantization and hardware acceleration where applicable (31). Feedback loops integrate human analyst labels and customer dispute resolutions to update model weights on a continuous basis. The risk orchestration tier fuses model scores with business rules and external threat intelligence, producing final actions such as transaction approval, challenge prompts, or manual review flags. Secure key management, data encryption at rest and in transit, and access controls enforce compliance with financial regulations (32). Chaos engineering practices validate system resilience under simulated service disruptions and adversarial load.

Experimental Framework and Evaluation

The evaluation framework employs a multi-phase approach. In the offline phase, models are trained and tuned using time-aware cross-validation to respect chronological order and simulate concept drift (33). Performance metrics include area under the precision–recall curve to account for extreme class imbalance, time-to-detection measured in milliseconds, and cost-weighted error rates reflecting differential fraud loss and customer friction. Ablation studies quantify the contribution of individual feature groups and modeling paradigms. In the online phase, shadow deployments route live traffic to the new model in parallel with the incumbent system, enabling unbiased A/B testing (34). Drift detection mechanisms monitor changes in feature distributions and model score histograms, triggering retraining pipelines when threshold deviations occur. Observability dashboards track key performance indicators, latency percentiles, and system health metrics.

Discussion and Future Directions

The integration of machine learning into fraud detection systems has yielded substantial gains in adaptability and precision (35). However, challenges persist in balancing interpretability with complexity, particularly as deep learning components become more prevalent. Explainable AI techniques—such as Shapley value approximations and counterfactual explanations—offer pathways toward transparent decisioning but incur additional computational overhead.

Moreover, the rise of privacy regulations necessitates exploration of federated learning architectures that allow models to benefit from cross-institution data without exposing sensitive customer information (36). Adversarial robustness remains an active frontier, requiring the development of certification methods for model behavior under worst-case perturbations. Finally, fully autonomous fraud resilience systems will likely incorporate continual learning paradigms, blending unsupervised anomaly detection with human-in-the-loop verification to maintain vigilance against unseen attack vectors.

Conclusion

This paper has detailed a holistic framework for deploying machine learning models in fraud detection across digital banking channels and payment platforms (37). We surveyed the fraud landscape, delineated the requirements for high-velocity inference, and compared diverse algorithmic families spanning supervised, unsupervised, and graph-based methods. A rigorous mathematical model was presented, formulating fraud detection as a robust optimization under adversarial dynamics and stochastic processes. We described practical considerations for data preprocessing, feature engineering, system integration, and continuous evaluation (38). Experimental design strategies ensure unbiased performance assessment and timely adaptation to concept drift. The discussion highlighted the imperative of interpretability, privacy preservation, adversarial robustness, and fully autonomous learning. Future work will advance federated and continual learning frameworks, integrate real-time explainability, and develop formal robustness guarantees. By synthesizing theoretical foundations with engineering best practices, the proposed framework aims to fortify financial systems against an evolving threat landscape while maintaining seamless customer experiences. (39)

References

1. S. Angelopoulos, E. Bendoly, J. Fransoo, K. Hoberg, C. Ou, and A. Tenhiälä, “Digital transformation in operations management: Fundamental change through agency reversal,” *Journal of Operations Management*, vol. 69, pp. 876–889, 8 2023.
2. S. Sahoo, S. Kumar, U. Sivarajah, W. M. Lim, J. C. Westland, and A. Kumar, “Blockchain for sustainable supply chain management: trends and ways forward,” *Electronic Commerce Research*, vol. 24, pp. 1563–1618, 5 2022.
3. G. K. Jagarlamudi, A. Yazdinejad, R. M. Parizi, and S. Pouriye, “Exploring privacy measurement in federated learning,” *The Journal of Supercomputing*, vol. 80, pp. 10511–10551, 12 2023.
4. M. Pešić, D. Egamberdieva, B. Kolodziejczyk, S. J. Elsässer, V. S. Neerghen, and A. Kagansky, “Towards policies that capture the expected value of biomolecular diversity for drug

- discovery, human health, and well-being,” *Biologia futura*, vol. 72, pp. 119–125, 9 2020.
5. M. Khodabakhshi, M. Asgharian, and G. N. Gregoriou, “An input-oriented super-efficiency measure in stochastic data envelopment analysis: Evaluating chief executive officers of us public banks and thrifts,” *Expert Systems with Applications*, vol. 37, pp. 2092–2097, 3 2010.
 6. P. J. Buckley and M. Casson, “Decision-making in international business,” *Journal of International Business Studies*, vol. 50, pp. 1424–1439, 5 2019.
 7. X.-P. Zhang, M. Ou, Y. Song, and X. Li, “Review of middle east energy interconnection development,” *Journal of Modern Power Systems and Clean Energy*, vol. 5, pp. 917–935, 11 2017.
 8. B. Zakeri, M. Khashehchi, S. Samsam, A. Tayebi, and A. Rezaei, “Solving partial differential equations by a supervised learning technique, applied for the reaction–diffusion equation,” *SN Applied Sciences*, vol. 1, pp. 1–8, 11 2019.
 9. D. Aerts, M. S. de Bianchi, S. Sozzo, and T. Veloz, “Modeling human decision-making: An overview of the brussels quantum approach,” *Foundations of Science*, vol. 26, pp. 1–28, 7 2018.
 10. C. E. Carroll and R. Olegario, “Pathways to corporate accountability: Corporate reputation and its alternatives,” *Journal of Business Ethics*, vol. 163, pp. 173–181, 6 2019.
 11. S. Barak and N. Parvini, “Transfer-entropy-based dynamic feature selection for evaluating bitcoin price drivers,” *Journal of Futures Markets*, vol. 43, pp. 1695–1726, 8 2023.
 12. J. Maric, C. Galera-Zarco, and M. Opazo-Basáez, “The emergent role of digital technologies in the context of humanitarian supply chains: a systematic literature review,” *Annals of operations research*, vol. 319, pp. 1–42, 5 2021.
 13. A. Ahl, M. Goto, and M. Yarime, “Smart technology applications in the woody biomass supply chain: interview insights and potential in japan,” *Sustainability Science*, vol. 15, pp. 1531–1553, 9 2019.
 14. J. Machireddy, “Customer360 application using data analytical strategy for the financial sector,” *Available at SSRN 5144274*, 2024.
 15. J. D. Farmer, M. Gallegati, C. Hommes, A. Kirman, P. Ormerod, S. Cincotti, A. Sánchez, and D. Helbing, “A complex systems approach to constructing better models for managing financial markets and the economy,” *The European Physical Journal Special Topics*, vol. 214, pp. 295–324, 12 2012.
 16. B. Al-Khateeb and G. Kendall, “Effect of look-ahead depth in evolutionary checkers,” *Journal of Computer Science and Technology*, vol. 27, pp. 996–1006, 11 2012.
 17. J. R. Machireddy, “Data quality management and performance optimization for enterprise-scale etl pipelines in modern analytical ecosystems,” *Journal of Data Science, Predictive Analytics, and Big Data Applications*, vol. 8, no. 7, pp. 1–26, 2023.
 18. Y. Ziser, B. Webber, and S. B. Cohen, “Rant or rave: variation over time in the language of online reviews,” *Language Resources and Evaluation*, vol. 57, pp. 1329–1359, 3 2023.
 19. B. Xu, L. Zuo, J. Jin, L. Han, and K. Hu, “Cab: a combinatorial-auction-and-bargaining-based federated learning incentive mechanism,” *World Wide Web*, vol. 26, pp. 2351–2372, 3 2023.
 20. A. I. Çetin and S. E. Ahmed, “Determinants of credit ratings and comparison of the rating prediction performances of machine learning algorithms,” *E3S Web of Conferences*, vol. 409, pp. 5013–05013, 8 2023.
 21. I. Aleksander, “Information technology and the management of change,” *Journal of Information Technology*, vol. 1, pp. 7–13, 2 1986.
 22. I. Ghosh, R. K. Jana, and M. Z. Abedin, “An ensemble machine learning framework for airbnb rental price modeling without using amenity-driven features,” *International Journal of Contemporary Hospitality Management*, vol. 35, pp. 3592–3611, 3 2023.
 23. N. Liu, P. Shapira, and X. Yue, “Tracking developments in artificial intelligence research: constructing and applying a new search strategy,” *Scientometrics*, vol. 126, pp. 3153–3192, 2 2021.
 24. C. Hesselman, P. Grosso, R. Holz, F. A. Kuipers, J. H. Xue, M. Jonker, J. de Ruitter, A. Sperotto, R. van Rijswijk-Deij, G. C. M. Moura, A. Pras, and C. de Laat, “A responsible internet to increase trust in the digital world,” *Journal of Network and Systems Management*, vol. 28, pp. 882–922, 9 2020.
 25. S. L. Piano, “Ethical principles in machine learning and artificial intelligence: cases from the field and possible ways forward,” *Humanities and Social Sciences Communications*, vol. 7, pp. 1–7, 6 2020.
 26. S. A. M. Selamat, S. Prakoonwit, and W. A. Khan, “A review of data mining in knowledge management: applications/findings for transportation of small and medium enterprises,” *SN Applied Sciences*, vol. 2, pp. 1–15, 4 2020.
 27. U. Y. Nafizah, S. Roper, and K. Mole, “Estimating the innovation benefits of first-mover and second-mover strategies when micro-businesses adopt artificial intelligence and machine learning,” *Small Business Economics*, vol. 62, pp. 411–434, 5 2023.
 28. N. Naeem, I. A. Rana, and A. R. Nasir, “Digital real estate: a review of the technologies and tools transforming the industry and society,” *Smart Construction and Sustainable Cities*, vol. 1, 10 2023.
 29. W. F. Clocksin, “Giving ’em the business,” *Nature*, vol. 312, pp. 81–81, 11 1984.
 30. A. Bhargava, M. Bester, and L. E. Bolton, “Employees’ perceptions of the implementation of robotics, artificial intelligence, and automation (raia) on job satisfaction, job security, and employability,” *Journal of Technology in Behavioral Science*, vol. 6, pp. 106–113, 8 2020.
 31. J. Tuckett, “Spirituality and intersubjective consensus: A response to ciocan and ferencz-flatz,” *Human Studies*, vol. 41, pp. 313–331, 2 2018.
 32. N. Aggarwal and L. Floridi, “Towards the ethical publication of country of origin information (coi) in the asylum process,”

- Minds and Machines*, vol. 30, pp. 247–257, 3 2020.
33. N. G. Taylor, P. Grillas, H. A. Hreisha, Özge Balkız, M. Borie, O. Boutron, A. Catita, J. Champagnon, S. Cherif, K. Çiçek, L. Costa, M. Dakki, M. Fois, T. Galewski, A. Galli, N. M. Georgiadis, A. J. Green, V. Hermoso, R. Kapedani, M. A. Lange, Z. Mateljak, M. Osta, E. Papastergiadou, C. Papazoglou, S. Sabater, B. Samraoui, F. Samraoui, A. S. Bachir, E. Tankovic, M. Thévenet, A. Troya, and W. J. Sutherland, “The future for mediterranean wetlands: 50 key issues and 50 important conservation research questions.,” *Regional environmental change*, vol. 21, pp. 33–33, 3 2021.
 34. L. P. Ruster and G. Brown, “Termination for cultural misalignment: Setting up contract terms to ensure community well-being in the development of ai,” *International Journal of Community Well-Being*, vol. 3, pp. 523–537, 10 2020.
 35. Y. Zeng, P. Doshi, Y. Chen, Y. Pan, H. Mao, and M. Chandrasekaran, “Approximating behavioral equivalence for scaling solutions of i-dids,” *Knowledge and Information Systems*, vol. 49, pp. 511–552, 12 2015.
 36. J. Moosavi, L. M. Naeni, A. M. Fathollahi-Fard, and U. Fiore, “Blockchain in supply chain management: a review, bibliometric, and network analysis,” *Environmental science and pollution research international*, pp. 1–15, 2 2021.
 37. G. A. Vasilakis, K. Theofilatos, E. F. Georgopoulos, A. Karathanasopoulos, and S. Likothanassis, “A genetic programming approach for eur/usd exchange rate forecasting and trading,” *Computational Economics*, vol. 42, pp. 415–431, 10 2012.
 38. A. M. Yaakob, S. Shafie, A. Gegov, and S. F. A. Rahman, “Z-hesitant fuzzy network model with reliability and transparency of information for decision systems,” *International Journal of Computational Intelligence Systems*, vol. 14, pp. 176–, 10 2021.
 39. J. Zhou, L. Yin, X. Wei, S. Zhang, Y. Song, B. Luo, J. Li, L. Qian, L.-G. Cui, W. Chen, C. Wen, Y.-L. Peng, Q. Chen, M. Lu, M. Chen, R. Wu, W. Zhou, E.-S. Xue, Y.-J. Li, L. Yang, C. Mi, R. Zhang, G. Wu, G. Du, D. Huang, W. Zhan, S. Organ, and B. Ultrasound, “2020 chinese guidelines for ultrasound malignancy risk stratification of thyroid nodules: the c-tirads.,” *Endocrine*, vol. 70, pp. 256–279, 8 2020.