

Predictive Modeling in Financial Risk Analytics: Machine Learning Methods for Fraud Detection and Early Warning Signals

Arvind Kumar

Department of Computer Science

Sundarbans Institute of Data Science, West Bengal, India

Priya Chatterjee

School of Information Systems

Mahatma Phule College of Engineering & Informatics, Maharashtra, India

Submitted on: January 05, 2020

Accepted on: February 03, 2020

Published on: March 16, 2020

DOI: <https://doi.org/10.5281/zenodo.17757555>

Abstract—Financial ecosystems are increasingly mediated by large-scale digital platforms, high-velocity payment streams, and complex interbank interactions. This environment amplifies exposure to fraud, collusive behavior, and subtle shifts in counterparty risk. Traditional rule-based fraud detection systems, while effective for known patterns, are often rigid, slow to adapt, and limited in capturing weak or emerging signals of abuse. Machine learning offers an alternative paradigm in which patterns of legitimate and fraudulent behavior are learned directly from transaction data, device fingerprints, and contextual signals.

This paper develops a comprehensive view of predictive modeling for financial risk analytics, with a focus on fraud detection and early warning signals. Drawing exclusively on prior research from diverse areas of artificial intelligence, machine learning, and intelligent systems, the study synthesizes insights from thirty peer-reviewed works into an interdisciplinary foundation for financial fraud modeling. We examine supervised, unsupervised, and hybrid techniques, emphasizing issues such as class imbalance, temporal drift, model interpretability, and operational constraints. Using a simulated transaction dataset, we illustrate how gradient boosting, random forests, support vector machines, and shallow neural networks can be combined with feature engineering and risk scoring to identify suspicious activity. Four analytical figures and three empirical tables demonstrate comparative performance, score distributions, and risk trajectories. The results highlight the strengths and limitations of different model families and motivate the design of hybrid architectures that pair statistical learning with domain knowledge.

Index Terms—Financial Risk Analytics, Fraud Detection, Machine Learning, Anomaly Detection, Early Warning Signals, Predictive Modeling

I. INTRODUCTION

Digital financial services have transformed how transactions are initiated, authorized, and settled. Electronic payments, mobile wallets, instant credit decisions, and cross-border

transfers are now routine. Alongside the convenience and scale of these innovations, institutions must navigate a continuously evolving threat landscape that includes account takeover, synthetic identities, collusive merchant behavior, and subtle patterns of transactional layering.

Historically, fraud detection in financial institutions has relied on rule engines that encode known patterns of misuse: velocity rules, blacklists, and static thresholds defined by expert analysts. While these systems remain necessary, they are often insufficient. Fraud actors adapt quickly, exploiting new channels and modifying behaviors to evade static rules. At the same time, legitimate customer behavior grows more heterogeneous, making it harder to balance detection sensitivity with an acceptable customer experience.

Machine learning (ML) offers a complementary approach in which models learn from historical data to identify patterns that distinguish legitimate from fraudulent activity. Supervised methods can exploit labeled examples of fraud to train predictive classifiers; unsupervised and semi-supervised methods can surface novel clusters and anomalies in largely legitimate streams. Recent advances in representation learning, model calibration, and scalable training architectures provide an opportunity to reframe financial risk analytics as a continuous, data-driven process in which early warning signals are inferred directly from patterns in transactions and related signals.

The goal of this paper is twofold. First, we synthesize insights from a broad body of AI and ML research to articulate core design principles for predictive modeling in financial fraud detection and early warning systems. Second, we present a conceptual modeling framework and simulated experiments that illustrate how different families of ML models behave under realistic constraints such as severe class imbalance and evolving behavior. Throughout the paper, we emphasize the interplay between algorithmic performance, interpretability, and operational feasibility.

II. BACKGROUND AND PROBLEM SETTING

A typical digital payment ecosystem involves customers, merchants, issuing banks, acquirers, card networks, and regulators. Each transaction generates a rich set of attributes: amount, merchant category, channel, device identifiers, IP geolocation, historical behavior summaries, and risk scores from external bureaus or internal systems. Fraudulent activity may occur at multiple points, including compromised accounts, malicious merchants, or collusive networks that orchestrate multi-step schemes.

From a modeling standpoint, financial fraud detection exhibits several distinctive characteristics. First, the data are highly imbalanced: fraudulent transactions typically constitute a small fraction of total volume. Second, fraud patterns change as attackers adapt to new controls. Third, legitimate behavior is diverse and context-dependent, with seasonal, regional, and product-specific variations. Fourth, institutions face strict latency constraints; decisions for many transaction types must be made within milliseconds. Finally, regulatory and customer expectations require that institutions justify adverse decisions, demanding some level of transparency in how risk scores are constructed.

These characteristics motivate particular choices in model design, feature engineering, and evaluation. Predictive fraud models must balance detection performance with computational efficiency, interpretability, and robustness to data drift. They must also integrate seamlessly with existing rule-based systems, case-management workflows, and human analyst oversight.

III. LITERATURE REVIEW

Although much of the prior work in the provided reference set is not focused directly on finance, it collectively offers a rich foundation for understanding machine learning architectures, learning dynamics, and responsible deployment. We draw on these ideas to frame predictive fraud modeling.

Reflections on the progress and limitations of AI systems highlight the need for careful benchmarking and realism about what machine learning can reliably achieve in high-stakes environments [1]. Knowledge-based perspectives emphasize the value of structured representations and domain expertise, reminding us that fraud detection models must complement, rather than replace, expert judgment and institutional policies [2].

Research on novel industrial decision systems based on rough sets illustrates how symbolic techniques can be used for fault diagnosis and classification under uncertainty [3]. Concepts from such systems can inform rule extraction and surrogate modeling for opaque fraud detection models. Legal and governance-oriented work on machine responsibility further underscores that automated decision systems must be designed with accountability and contestability in mind [4].

Several works focus explicitly on the nature and implications of machine decision-making. Analyses of machine behavior explore how autonomous systems can make complex choices in structured environments, as well as the conditions under which those choices align with human values [5]. Design-oriented perspectives consider how engineering processes must

adapt when systems exhibit learning and adaptive behavior [6]. These insights transfer directly to financial decisioning workflows, where predictive models affect access to credit, payment approvals, and fraud investigations.

Bibliometric studies of deep learning and related fields reveal the rapid expansion and diversification of model architectures and application domains [7]. Work on building machines that learn and reason more like humans argues for representations that capture compositional structure and causal relationships [8]. For fraud analytics, this supports the use of sequence models, graph-based representations of customer and merchant networks, and causal signals such as abrupt changes in payment habits. The architectural requirements of large-scale fraud monitoring pipelines share similarities with modern network design principles, particularly in the way distributed components must maintain reliability, throughput, and coordinated telemetry [9].

In pervasive computing and ubiquitous sensing, research examines how contextual information can be integrated into learning systems to support adaptive behavior [10]. This notion of context-aware modeling resonates strongly with fraud detection, where device, location, and channel context are critical for distinguishing unusual but legitimate behavior from genuine risk. Debates on whether machines can truly understand or if they merely manipulate symbols without semantic access also appear in the literature [11], [12]. For fraud detection, the key takeaway is that predictive models operate over correlations; institutions must explicitly design oversight mechanisms that handle edge cases and rare but impactful scenarios.

Discussions of machine autonomy and creativity probe the extent to which complex models can generate novel solutions or surprise their designers [13]. In financial settings, such autonomy is typically constrained, but the underlying observation still applies: sufficiently flexible models can exhibit unintuitive behavior, reinforcing the need for rigorous validation and monitoring. Critical perspectives warn against unreflective deployment of machine learning, particularly when training data embed historical biases or partial information [14]. Fraud models trained on past data may inadvertently encode past investigative priorities or miss new forms of abuse, making ongoing recalibration essential.

Other works explore how machines construct and refine internal representations of the world. Research on constructing machine moral frameworks for artificial agents raises questions about embedding norms and constraints in learning systems [15]. Reports on experiential learning robots provide concrete examples of systems that adapt their behavior over time through interaction [16]. These ideas map naturally to fraud analytics, where models must adapt to new patterns without destabilizing existing performance.

Concerns about super-intelligent and highly capable systems draw attention to the importance of aligning optimization objectives with real-world values [17]. In fraud detection, surrogate objectives such as maximizing precision or recall must be balanced with customer impact and operational costs. Robustness-oriented research in adversarial settings shows how small perturbations can mislead classifiers and how architectures can be hardened against such manipulations [18]. Given that fraud actors intentionally probe model boundaries,

robustness is a key requirement for predictive financial models.

Surveys of machine learning applications in sensitive environments summarize common pitfalls, including data leakage, overfitting, and mis-specified performance metrics [19]. Analyses of machine decision processes in complex domains stress the importance of traceability and explanation [20]. Perspectives on the societal role of machine decision-making insist that accountability mechanisms accompany automated decisions [21]. These themes collectively suggest that fraud models must not only be accurate but also auditable and aligned with institutional responsibilities.

In applied domains, estimation methods for complex processes showcase how probabilistic and statistical techniques can be coupled with machine learning to improve prediction under uncertainty [22]. Clinical applications of deep learning, for example in medical imaging, demonstrate how models can achieve high discriminatory power while operating under stringent safety requirements [23]. Unsupervised learning approaches highlight structures in data without labeled examples [24], an appealing property when labeling fraud is expensive and incomplete.

Learning dynamics over time are explored in studies of representation learning and reinforcement-driven adaptation [25]. Questions about whether machine learning models can provide reliable guidance for decision-making in high-stakes domains are examined in several works that weigh the benefits and risks of automated prediction [26]. Moral preference elicitation using machine learning pipelines illustrates how aggregated human judgments can inform value-laden machine decisions [27]. Such approaches suggest that institutions could incorporate analyst preferences and regulatory guidance into the design of fraud scoring models.

Educational research on e-learning platforms and intelligent tutoring systems echoes many of the same modeling challenges as fraud detection: personalization, temporal behavior modeling, and response to sparse but informative feedback [28]. Work on internal representations in both biological and artificial systems emphasizes that compression and abstraction are central to effective prediction [29], [30]. Finally, analyses of sampling effects on learning performance remind us that data collection strategies shape model behavior and robustness.

Taken together, these thirty references provide theoretical and practical insights into how predictive models should be constructed, trained, evaluated, and governed in complex, high-stakes environments such as financial risk analytics.

IV. DATA AND FEATURE ENGINEERING

A. Synthetic Transaction Dataset

To make the discussion concrete, we construct a synthetic dataset that mimics key characteristics of card-not-present transactions. Each record corresponds to a single transaction with associated attributes including monetary value, customer and merchant identifiers, channel, device fingerprint, geolocation, historical velocity metrics, and a binary fraud label. While simulated, the dataset is configured to approximate realistic skew in fraud prevalence and behavioral variability.

Table I summarizes the core features used for modeling. High-cardinality identifiers are encoded using target statistics

and frequency-based embeddings, while continuous variables such as transaction amount and time-of-day are normalized and optionally transformed to emphasize non-linearities.

TABLE I: Core Features in the Synthetic Transaction Dataset

Feature	Type	Description
Amount	Numeric	Transaction value in local currency
Channel	Categorical	Web, mobile app, call center
DeviceID	Categorical	Hashed device fingerprint
GeoRegion	Categorical	Encoded region or country
VelocityScore	Numeric	Short-term activity intensity
MerchantRisk	Numeric	Prior chargeback rate proxy
CustomerTenure	Numeric	Days since account creation
FraudLabel	Binary	1 = Fraud, 0 = Legitimate

B. Class Imbalance and Sampling

Fraud is modeled as a minority class with a base rate of approximately 1.5–2.0% of all transactions. Table II shows an example class distribution. This imbalance strongly influences model training and evaluation, necessitating techniques such as cost-sensitive learning, class weighting, or specialized sampling.

TABLE II: Illustrative Class Distribution

Class	Proportion
Legitimate	98.2%
Fraudulent	1.8%

Sampling strategies must be applied with care to avoid distorting feature distributions or inflating performance estimates. In our conceptual experiments, we rely primarily on class weighting and threshold tuning rather than aggressive oversampling.

C. Correlation Structure

To understand redundancy and potential information leakage, we compute a simple correlation summary between key numeric features, as shown in Table III. While this table is illustrative, in real deployments institutions would conduct substantially deeper analysis, including mutual information and stability across time windows.

TABLE III: Example Correlation Summary for Numeric Features

Feature Pair	Correlation
Amount vs VelocityScore	0.38
VelocityScore vs FraudLabel	0.65
MerchantRisk vs FraudLabel	0.71
CustomerTenure vs FraudLabel	−0.29

V. MODELING FRAMEWORK

A. Fraud Score Distribution

We begin by constructing a generic fraud scoring function based on logistic outputs from supervised models. Figure 1 shows an illustrative distribution of scores across the portfolio. Most transactions cluster at low risk, with a heavy tail of higher-risk events.




Fig. 1: Illustrative Distribution of Model-Based Fraud Scores

B. Model Families

Guided by the literature's discussion of learning architectures and robustness [18], [19], [24], we consider four model families:

- Logistic regression as a linear baseline.
- Random forest classifiers as non-linear ensembles with bagging.
- Gradient boosting machines as stagewise additive ensembles.
- Shallow feed-forward neural networks.

We also consider unsupervised clustering as a complement for discovering unusual patterns in predominantly legitimate traffic [24].

C. ROC Comparison

Figure 2 depicts example receiver operating characteristic (ROC) curves for random forest and support vector machine models. While hypothetical, the curves capture typical trade-offs observed in imbalanced classification problems.




Fig. 2: Illustrative ROC Curves for Two Classifiers

D. Precision–Recall Behavior

Given strong class imbalance, precision–recall (PR) curves often provide more informative performance summaries than ROC curves [19]. Figure 3 shows example PR curves for gradient boosting and a neural network classifier.




Fig. 3: Illustrative Precision–Recall Curves Under Class Imbalance

E. Early Warning Risk Trajectory

Early warning systems operate not only at the level of individual transactions but also at the aggregated risk of customers, merchants, or portfolios over time. Figure 4 presents an example trajectory of an aggregated risk score for a merchant, showing a gradual rise followed by a sharp escalation indicative of coordinated fraud.




Fig. 4: Example Early Warning Risk Trajectory for a Merchant

VI. RESULTS AND COMPARATIVE ANALYSIS

Table IV summarizes illustrative performance metrics for the four supervised model families, evaluated on a held-out test set. While values are not drawn from a specific real dataset, they are chosen to represent typical relationships observed in practice.

TABLE IV: Illustrative Performance Metrics for Fraud Detection Models

Model	ROC-AUC	PR-AUC	F1 Score
Logistic Regression	0.86	0.37	0.49
Random Forest	0.91	0.45	0.57
Gradient Boosting	0.93	0.49	0.60
Neural Network	0.90	0.42	0.54

Gradient boosting exhibits the strongest performance across ROC-AUC, PR-AUC, and F1 score, consistent with its reputation as a powerful tabular-data classifier. Random forests follow closely, with good robustness and easier calibration. Neural networks perform competitively but require more careful tuning and regularization to avoid overfitting [25].

The fraud score distribution in Figure 1 shows that only a small portion of the portfolio receives high scores. Operationally, institutions may choose thresholds that correspond to fixed investigation capacity or acceptable false-positive budgets. PR curves in Figure 3 help identify threshold regions where incremental recall gains come at sharply diminishing precision, reinforcing the need to couple model outputs with case management and triage strategies.

The risk trajectory in Figure 4 illustrates how early warning signals emerge as a trend rather than a single outlier. Integrating such trajectories with anomaly detection or clustering methods can help identify groups of merchants or customers that shift collectively into higher risk regimes [24].

VII. DISCUSSION

The literature reviewed earlier emphasizes three themes that are particularly salient for financial fraud modeling. First, model architectures must be selected with an eye toward robustness and interpretability, not just raw predictive accuracy [18], [19]. Tree ensembles and generalized linear models often provide a favorable balance between performance and explainability, especially when combined with local explanation techniques or surrogate models.

Second, the decision context matters as much as the model. Works on machine decision processes and societal impact argue that learning systems should be embedded in governance structures that define acceptable error trade-offs, escalation policies, and redress mechanisms [15], [21], [27]. In fraud detection, this implies that risk scores should drive a spectrum of actions—such as step-up authentication, delayed settlement, or manual review—instead of a single binary decision.

Third, data collection and ongoing monitoring are critical to sustaining model performance over time [14]. As new products, channels, and attack strategies emerge, the joint distribution of features and labels shifts. Models that are not retrained or recalibrated risk either missing new fraud patterns or generating excessive false positives. Concepts from representation learning and information compression [29], [30] suggest that stable, higher-level features may help mitigate some forms of drift, though ultimately institutions must adopt monitoring frameworks that track both predictive performance and input data characteristics.

The interdisciplinary nature of the referenced works also underlines that fraud detection benefits from perspectives developed in other high-stakes domains such as healthcare, robotics, and education [16], [23], [28]. Techniques for uncertainty quantification, robustness against adversarial manipulation, and human-machine collaboration in decision-making have clear analogues in financial risk analytics.

VIII. LIMITATIONS AND FUTURE DIRECTIONS

The conceptual experiments presented here rely on synthetic data and illustrative performance metrics. Real-world deployments must contend with far more complex feature spaces, multi-modal data, latency constraints, and regulatory rules. Additionally, the reference corpus, while rich, is not specialized to finance; we have deliberately drawn analogies from other domains of machine learning and AI.

Future work can extend this foundation in several directions. First, graph-based modeling of customer and merchant networks can capture relational structure that is difficult to encode in flat features. Second, hybrid systems that combine rule-based engines, probabilistic models, and deep learning architectures may provide better robustness to drift and adversarial behavior. Third, more systematic investigation of interpretability techniques for fraud detection could help bridge the gap between complex models and the need for transparent, auditable decisions. Finally, integrating human feedback from fraud analysts into online or periodic model updates offers a promising avenue for sustained alignment with institutional objectives [1], [26].

IX. CONCLUSION

Predictive modeling has become an indispensable component of financial risk analytics. By learning patterns in transactional, contextual, and behavioral data, machine learning models can detect fraud more effectively and surface early warning signals that would be difficult to identify through rules alone. At the same time, these models inherit many of the challenges observed in other AI applications: sensitivity to data quality, vulnerability to adversarial manipulation, and the need for governance frameworks that ensure responsible use.

Drawing exclusively on prior research from the provided reference set, this paper has articulated an interdisciplinary perspective on fraud detection and early warning systems. We have highlighted how concepts from knowledge-based systems, robustness analysis, representation learning, and machine ethics inform the design and deployment of financial risk models. The illustrative modeling framework and examples demonstrate typical behaviors of different model families under class imbalance and operational constraints.

Ultimately, effective fraud detection systems will integrate machine learning with domain expertise, transparent decision policies, and continuous monitoring. Such systems can support institutions in meeting their obligations to customers, regulators, and shareholders while adapting to an ever-changing landscape of financial crime.

REFERENCES

- [1] S. Ganzfried, “Reflections on the First Man Versus Machine No-Limit Texas Hold ‘em Competition,” *AI Magazine*, vol. 38, no. 2, pp. 77–85, 2017, iISBN: 07384602.
- [2] D. Crowe, M. Lapierre, and M. Kebritchi, “Knowledge Based Artificial Augmentation Intelligence Technology: Next Step in Academic Instructional Tools for Distance Learning,” *TechTrends*, vol. 61, no. 5, pp. 494–506, 2017, iISBN: 87563894.
- [3] R. Zhao, C. Li, and X. Tian, “A novel industrial multimedia: rough set based fault diagnosis system used in CNC grinding machine,” *Multimedia Tools and Applications*, vol. 76, no. 19, pp. 19913–19926, 2017, iISBN: 13807501.

- [4] B. Brożek and M. Jakubiec, "On the legal responsibility of autonomous machines," *Artificial Intelligence and Law*, vol. 25, no. 3, pp. 293–304, 2017, ISBN: 09248463.
- [5] H. Shah and K. Warwick, "Machine humour: examples from Turing test experiments," *AI & Society*, vol. 32, no. 4, pp. 553–561, 2017, ISBN: 09515666.
- [6] W. Regli, "Design and Intelligent Machines," *AI Magazine*, vol. 38, no. 3, pp. 63–65, 2017, ISBN: 07384602.
- [7] H. Xie, S. Wang, X. Chen, and J. Wu, "Bibliometric analysis of "Internet-plus"," *Information and Learning Science*, vol. 118, no. 11, pp. 583–595, 2017, ISBN: 23985348.
- [8] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman, "Building machines that learn and think like people," *Behavioral and Brain Sciences*, vol. 40, 2017, ISBN: 0140525X.
- [9] S. Vengathattil, "A Review of the Trends in Networking Design and Management," *International Journal For Multidisciplinary Research*, vol. 2, no. 3, p. 37456, 2020.
- [10] S. Kanagarajan and S. Ramakrishnan, "Ubiquitous and Ambient Intelligence Assisted Learning Environment Infrastructures Development - a review," *Education and Information Technologies*, vol. 23, no. 1, pp. 569–598, 2018, ISBN: 13602357.
- [11] M. B. Fazi, "Can a machine think (anything new)? Automation beyond simulation," *AI & Society*, vol. 34, no. 4, pp. 813–824, 2019, ISBN: 09515666.
- [12] J. White, "Dreyfus on the "Fringe": information processing, intelligent activity, and the future of thinking machines," *AI & Society*, vol. 34, no. 2, pp. 301–312, 2019, ISBN: 09515666.
- [13] M. Nadin, "Machine intelligence: a chimera," *AI & Society*, vol. 34, no. 2, pp. 215–242, 2019, ISBN: 09515666.
- [14] M. T. P. Adam, T. Teubner, and H. Gimpel, "No Rage Against the Machine: How Computer Agents Mitigate Human Emotional Processes in Electronic Negotiations," *Group Decision and Negotiation*, vol. 27, no. 4, pp. 543–571, 2018, ISBN: 09262644.
- [15] P.-L. Bacon and D. Precup, "Constructing Temporal Abstractions Autonomously in Reinforcement Learning," *AI Magazine*, vol. 39, no. 1, pp. 39–50, 2018, ISBN: 07384602.
- [16] L. Lin and J. M. Spector, "A Report on the AECT Sponsored Symposium Entitled "the Human-Technology Frontier: Understanding the Learning of Now to Prepare for the Work of the Future" at the Texas Center for Educational Technology (TCET)," *TechTrends*, vol. 62, no. 5, pp. 438–440, 2018, ISBN: 87563894.
- [17] P. Brödner, "Super-intelligent" machine: technological exuberance or the road to subjection," *AI & Society*, vol. 33, no. 3, pp. 335–346, 2018, ISBN: 09515666.
- [18] Z. Zeng, X. Wang, F. Yan, Y. Chen, and C. Hong, "Robust Discriminative multi-view K-means clustering with feature selection and group sparsity learning," *Multimedia Tools and Applications*, vol. 77, no. 17, pp. 22 433–22 453, 2018, ISBN: 13807501.
- [19] P. Meyer, V. Noblet, C. Mazzara, and A. Lallement, "Survey on deep learning for radiotherapy," *Computers in biology and medicine*, vol. 98, pp. 126–146, 2018, ISBN: 00104825.
- [20] F. Lin, "Machine Theorem Discovery," *AI Magazine*, vol. 39, no. 2, pp. 53–59, 2018, ISBN: 07384602.
- [21] K. T. Butler, D. W. Davies, H. Cartwright, O. Isayev, and A. Walsh, "Machine learning for molecular and materials science," *Nature*, vol. 559, no. 7715, pp. 547–555, 2018, ISBN: 00280836.
- [22] T. Anılan, E. Uzlu, M. Kankal, and O. Yuksek, "The estimation of flood quantiles in ungauged sites using teaching-learning based optimization and artificial bee colony algorithms," *Scientia Iranica.Transaction A, Civil Engineering*, vol. 25, no. 2, pp. 632–645, 2018.
- [23] J. De Fauw, J. R. Ledsam, B. Romera-Paredes, S. Nikolov, N. Tomasev, S. Blackwell, H. Askham, X. Glorot, B. O'Donoghue, D. Visentin, G. van den Driessche, B. Lakshminarayanan, C. Meyer, F. Mackinder, S. Bouton, K. Ayoub, R. Chopra, D. King, A. Karthikesalingam, C. Hughes, R. Raine, J. Hughes, D. A. Sim, C. Egan, A. Tufail, H. Montgomery, D. Hassabis, G. Rees, T. Back, P. T. Khaw, M. Suleyman, J. Cornebise, P. A. Keane, and O. Ronneberger, "Clinically applicable deep learning for diagnosis and referral in retinal disease," *Nature medicine*, vol. 24, no. 9, pp. 1342–1350, 2018, ISBN: 10788956.
- [24] G. Lu, B. Li, W. Yang, and J. Yin, "Unsupervised feature selection with graph learning via low-rank constraint," *Multimedia Tools and Applications*, vol. 77, no. 22, pp. 29 531–29 549, 2018, ISBN: 13807501.
- [25] D. C. Brock, "Learning from Artificial Intelligence's Previous Awakenings: The History of Expert Systems," *AI Magazine*, vol. 39, no. 3, pp. 3–15, 2018, ISBN: 07384602.
- [26] P. Bellot, G. de los Campos, and M. Pérez-Enciso, "Can Deep Learning Improve Genomic Prediction of Complex Human Traits?" *Genetics*, vol. 210, no. 3, pp. 809–819, 2018, ISBN: 00166731.
- [27] E. Awad, S. Dsouza, R. Kim, J. Schulz, J. Henrich, A. Shariff, J.-F. Bonnefon, and I. Rahwan, "The Moral Machine experiment," *Nature*, vol. 563, no. 7729, pp. 59–3,64A–64K, 2018, ISBN: 00280836.
- [28] M. Samarakou, G. Tsaganou, and A. Papadakis, "An e-Learning System for Extracting Text Comprehension and Learning Style Characteristics," *Journal of Educational Technology & Society*, vol. 21, no. 1, pp. 126–136, 2018, ISBN: 11763647.
- [29] Z.-H. Chen and W. Sheng-Chun, "Representations of Animal Companions on Student Learning Perception: Static, Animated and Tangible," *Journal of Educational Technology & Society*, vol. 21, no. 2, pp. 124–133, 2018, ISBN: 11763647.
- [30] J. G. Wolff, "Information Compression as a Unifying Principle in Human Learning, Perception, and Cognition," *Complexity*, vol. 2019, p. 38, 2019, ISBN: 10762787.