AI, Data Science, and Quantum Neural Networks in E-Commerce: Methods, Applications, Risks, and a Research Roadmap

Author: Sophia Martinez, **Affiliation:** Senior Lecturer, School of AI and Data Science, University of Buenos Aires, Argentina. **Email:** sophia.martinez@uba.ar

Abstract

Ecommerce has matured into a data-intensive ecosystem where consumer behavior, logistics, pricing, personalization, and fraud detection are driven by large, heterogeneous datasets. Classical AI and data-science pipelines have delivered major productivity and revenue gains, yet they face limits in modeling combinatorial recommendation spaces, accelerating molecular-scale cryptography for secure transactions, and solving certain optimization problems at scale. Quantum Neural Networks (QNNs) and hybrid quantum-classical approaches promise novel representational capacity and computational primitives that can enrich recommender systems, optimization for supply chains, privacy-preserving analytics, and next-generation fraud detection. This article synthesizes theory and practice across AI, data science, and QNNs for e-commerce. It offers: (1) a conceptual framework linking data modalities and business problems; (2) detailed methods for modeling, training, deployment, and evaluation; (3) extended technical sections on quantum representations, QNN architectures, and hybrid pipelines with reproducible pseudocode; (4) applied case studies (recommendation, dynamic pricing, inventory optimization, personalization, fraud detection, privacy); (5) engineering and MLOps considerations; (6) an ethics, privacy, and regulatory analysis; and (7) a prioritized research and deployment roadmap. The paper targets researchers and practitioners aiming to integrate quantum-enhanced AI into production e-commerce systems while preserving interpretability, fairness, security, and economic value.

Keywords: e-commerce, recommender systems, data science, artificial intelligence, quantum neural networks, optimization, privacy, fraud detection

1. Introduction

1.1 Motivation and scope E-commerce platforms generate and rely upon massive, multimodal data streams: clickstreams, transaction logs, product images and descriptions, user reviews, inventory feeds, supply-chain telemetry, and third-party signals (weather, events). Traditional AI and data-science methodologies (collaborative filtering, gradient boosting, deep learning, graph learning) have driven personalized recommendations, dynamic pricing, fraud mitigation, and logistics optimization (Ricci et al., 2015; Goodfellow et al., 2016). Nonetheless, specific computational challenges remain: the combinatorial complexity of recommendations under constraints, robust multi-objective optimization under uncertainty, secure and privacy-preserving cross-party learning, and detection of sophisticated fraud that adapts to defenses. Quantum computing and Quantum Neural Networks (QNNs) provide alternative computational models leveraging superposition and entanglement that may offer advantages on certain subproblems (Biamonte et al., 2017; Schuld et al., 2014). This manuscript examines how AI, data science, and QNNs can be integrated into e-commerce workflows to advance personalization, efficiency, and security while considering technical, ethical, and operational constraints.

1.2 Contributions

- A unifying framework mapping e-commerce problem to AI and QNN technical primitives.
- Detailed methods for data ingestion, feature engineering, model design, training, validation, deployment, and monitoring.
- In-depth exposition of QNN architectures applicable to ranking, clustering, anomaly detection, and combinatorial optimization, with reproducible pseudocode.
- Case studies demonstrating hybrid quantum-classical pipelines for recommendation, dynamic pricing, inventory optimization, privacy-preserving cross-platform learning, and fraud detection.
- Comprehensive discussion of MLOps, governance, regulatory compliance, explainability, and ethical considerations.
- A research and deployment roadmap prioritizing near-term hybrid applications and long-term fault-tolerant ambitions.

1.3 Organization Section 2 surveys e-commerce data modalities and business problems. Section 3 reviews classical AI and data-science tools and their limitations. Section 4 develops quantum computing and QNN foundations. Section 5 proposes hybrid architecture and training methods. Section 6 provides technical case studies with pseudocode. Section 7 addresses deployment, MLOps, and performance engineering. Section 8 discusses privacy, security, ethics, and regulation. Section 9 outlines research priorities and a roadmap. Section 10 concludes.

2. E-commerce problem space and data modalities

2.1 Business problems and decision classes E-commerce decision tasks span multiple horizons and decision granularities:

- Instant decisions: click personalization, search ranking, on-page recommendations (milliseconds latency).
- Near-real-time decisions: dynamic pricing, inventory replenishment triggers, fraud scoring at checkout (seconds to minutes latency).
- Batch/strategic decisions: catalog assortment planning, long-horizon demand forecasting, supply-chain network design (hours to weeks latency).

Each class imposes distinct modeling and engineering constraints (latency, interpretability, retrain frequency, economic loss functions).

2.2 Data modalities and schemas

- Clickstream and session logs: high-velocity event sequences with timestamps, session identifiers, and contextual features.
- Transactional data: order line items, prices, promotions, returns, payment methods.
- Content data: product text, categorical attributes, hierarchical taxonomies, images, and video.
- Reviews and social signals: textual reviews, ratings, sentiment, influencer metrics.
- Graph data: user–item bipartite interactions, merchant networks, supply-chain edges.
- Operational telemetry: warehouse sensor feeds, shipment GPS, supplier lead times.
- Exogenous data: weather, holidays, macroeconomic indicators, competitor prices.

2.3 Data quality, labeling, and ground truth challenges Ground truth varies by task: explicit conversion events are clear, but usefulness/lifetime value signals are delayed and censored. For fraud, labeled data are scarce and adversarial. For personalization, offline metrics risk being misaligned with long-term retention and business value. Addressing these requires careful outcome modeling (survival analysis, counterfactual estimands) and labeling strategies (propensity scoring, human adjudication).

3. Classical AI and data-science foundations for e-commerce

3.1 Recommender systems: classes and tradeoffs

- Collaborative filtering: matrix factorization and latent factors capture user–item affinities but struggle with cold start and side information (Koren et al., 2009).

- Content-based and hybrid methods incorporate item attributes and deep content embeddings (e.g., image and text encoders) to address cold start and semantic matching.
- Session-based recommender models (RNNs, Transformers) model sequential dynamics (Hidasi et al., 2016; Kang & McAuley, 2018).
- Graph-based recommenders (graph convolutional networks) exploit multi-relational structures (Wu et al., 2020).

3.2 Ranking, learning-to-rank, and counterfactual learning Learning-to-rank (pointwise/pairwise/listwise) optimizes directly for ranking metrics; counterfactual and bandit frameworks are necessary when evaluating policies from logged bandit data to avoid selection bias (Joachims et al., 2017).

3.3 Pricing and inventory optimization Dynamic pricing leverages demand estimation, price elasticity, and strategic optimization (reinforcement learning or parametric optimization). Inventory optimization involves stochastic optimization under lead times and service constraints; approximate dynamic programming and model predictive control are widely used.

3.4 Fraud detection and anomaly detection Supervised classification works for common fraud types; network analysis, unsupervised anomaly detectors, and adversarial training are crucial for detecting sophisticated or novel fraud patterns.

3.5 Limitations and computational bottlenecks Classical models face challenges in modeling extremely high-dimensional combinatorial action spaces (basket recommendations under constraints), solving large NP-hard optimization problems in near real time, and enabling secure cross-platform learning without raw data sharing. These motivate hybrid quantum-classical approaches for selected subproblems.

4. Quantum computing and Quantum Neural Networks (QNNs): foundations

4.1 Quantum computation primer for practitioners Qubits represent quantum states in Hilbert space, enabling superposition and entanglement. Quantum gates enact unitary transforms; measurement collapses states to classical outcomes. Two practical device families are relevant: gate-model NISQ devices (superconducting, trapped ions) and quantum annealers (optimization-oriented). NISQ devices are noisy; hybrid variational algorithms mitigate noise by shifting classical optimization workloads (Preskill, 2018).

4.2 Quantum algorithms of interest to e-commerce

- Quantum optimization (QAOA, quantum annealing) maps combinatorial optimization problems to quantum cost Hamiltonians (Farhi et al., 2014).

- Quantum simulation and amplitude amplification can accelerate certain subroutines (Grover search) relevant to database search and inner-product computations.
- Quantum machine learning (QML), including QNNs and quantum kernels, offers alternative feature maps and model classes (Biamonte et al., 2017; Schuld et al., 2014).

4.3 Quantum Neural Networks: architectures and properties QNNs are parameterized quantum circuits (PQCs) with trainable gates acting as layers; outputs are expectation values of observables measured on circuit states. Key elements: encoding/feature maps (how classical data are embedded into quantum states), ansatz/variational layers (parameterized unitaries), measurement schemes, and classical optimizers (Guerreschi & Smelyanskiy, 2017). Prospective advantages include high-dimensional feature embeddings accessible via quantum Hilbert space and potential sample complexity benefits in some kernels.

4.4 Practical constraints and NISQ considerations

- Circuit depth and qubit connectivity limit model expressivity on current hardware.
- Barren plateaus (vanishing gradients) and noise-induced errors complicate training; hardware-aware ansatz design and error mitigation are required (Cerezo et al., 2021).
- For near-term e-commerce use, hybrid quantum-classical approaches run small QNNs for representation learning or use quantum annealing for combinatorial searches while delegating bulk training and inference to classical infrastructure.

5. Hybrid architectures: integrating classical AI, data science, and QNNs

5.1 Design patterns for hybrid systems

- Preprocessing and feature engineering remain classical; quantum modules are invoked for targeted subproblems (ranking re-ranking stages, combinatorial optimization, kernel estimation).
- Two-stage recommendation: classical candidate generation (fast recall) followed by QNN-based re-ranking for hard tradeoffs (diversity, inventory constraints).
- Quantum-assisted optimization: use QAOA or annealing to propose near-optimal bundles or shipping allocations, validated with classical simulation.

5.2 Data representation and quantum encoding choices

- Angle encoding: map scalar features to rotation angles; efficient for small feature vectors.
- Amplitude encoding: compact superposition of normalized vectors; potentially exponential compression but requires complex state preparation.
- Hybrid embeddings: classical neural encoders reduce raw modalities to compact vectors (e.g., 16–64 dims) that are then amplitude/angle encoded into QNNs.

5.3 Training regimes and loss designs

- Hybrid training with classical optimizers (SPSA, Adam) updating quantum parameters via gradient estimates (parameter-shift rule) for QNNs; use classical losses aligned with business objectives (ranking loss, NDCG, profit-weighted objectives).
- Meta-learning and transfer: pretrain classical encoders on abundant data and fine-tune QNN readouts on specific constrained optimization tasks.

5.4 System orchestration and latency considerations

- For user-facing low-latency tasks, QNN modules must be confined to milliseconds to seconds. Practical deployments use quantum accelerators asynchronously or in cached re-ranking contexts where latency budgets are larger.
- Batch and overnight optimization tasks (assortment planning, large-scale combinatorial allocation) are excellent QNN/annealer candidates.

6. QNN architectures and algorithms applied to e-commerce tasks

6.1 QNN design for re-ranking and personalization Objective: given candidate items and contextual features, produce a personalized re-rank that maximizes expected lifetime value subject to constraints (inventory, margin, fairness).

Architecture:

- Classical encoder E(x_user, x_context, x_item) → z ∈ R^d (d small, e.g., 16).
- Quantum feature map U(φ(z)) embeds z into qubits; variational layers V(θ) implement parameterized unitaries; measurement returns expectation values ŷ.
- Loss: cross entropy on click/conversion labels weighted by expected revenue; additional constraint penalties (Lagrangian) appended for margin and fairness.

Pseudocode (QNN re-ranking training loop):

text

```
Input: dataset D = {(user, context, candidates, label_relevances, revenue)}
for batch in D:
   z_batch = E(batch)
   for i in candidates:
      state = PrepareQuantumState(z_batch[i])  # feature map
      for l in 1..L:
         state = V_layer(state, theta_l)
      y_hat[i] = MeasureExpectation(state)
   loss = WeightedRankingLoss(y_hat, label_relevances, revenue_weights) + ConstraintPenalties(y_hat)
   theta = ClassicalOptimizerStep(theta, loss)
```

6.2 Quantum kernels for similarity and cold-start bridging Quantum kernel methods compute inner products in implicitly high-dimensional Hilbert space via quantum circuits. For cold start, map user/item meta features into quantum kernels and use kernelized ridge regression or SVMs to estimate affinities with small labeled data (Havlíček et al., 2019).

6.3 QAOA for combinatorial bundling and assortment optimization Encode discrete selection (which items to bundle given constraints) into cost Hamiltonian; run QAOA to sample low-energy bitstrings representing candidate bundles. Combine with classical heuristics (local search) to refine selections.

Pipeline (QAOA integrated):

- Formulate QUBO representing objective (expected revenue minus substitution costs) and constraints (inventory, shipping).
- Translate QUBO → Ising Hamiltonian; instantiate QAOA circuit depth p.
- Optimize variational angles γ, β classically; sample candidate solutions; rank solutions by expected profit; feed top K to classical validation.

6.4 Quantum anomaly detection for fraud QNN-based autoencoder analogs: encode transactional patterns into quantum latent codes; reconstruct and measure reconstruction error as anomaly signal. Hybrid approach: classical encoder → small QNN latent transform → classical decoder to reconstruct; use reconstruction loss as fraud score.

6.5 Complexity, sample efficiency, and expected benefits Theoretical and empirical evidence for QNN advantage is nascent; selected use cases where QNNs may offer value include high-dimensional kernel embeddings for small-label regimes, sampling from complex combinatorial posterior distributions, and leveraging amplitude encoding to compactly represent structured item features. Practical value depends on problem structure, hardware maturity, error rates, and integration costs (Biamonte et al., 2017; Schuld et al., 2014).

7. Case studies: hybrid quantum-classical applications in e-commerce

7.1 Case study 1 — Candidate generation + QNN re-ranking for personalization (production-style pipeline)

Problem: Improve conversion lift and long-term revenue by reranking classical recall candidates under margin and inventory constraints.

Pipeline:

- Stage A: Candidate generation (classical): fast nearest neighbors via approximate embeddings (FAISS) or item popularity heuristics.
- Stage B: Feature fusion (classical): compute E(user, item, session) vectors.
- Stage C: QNN re-rank: embed vectors into QNN, run inference to obtain scores with constraint penalties, apply post-processing (deduplication, inventory filter).
- Stage D: A/B test served via canary rollout; monitor conversion, revenue per session, and margin.

Evaluation: offline NDCG, expected revenue uplift in simulation with counterfactual corrections (logged bandit data methods); online randomized experiments for causal validation.

7.2 Case study 2 — QAOA for Bundling and Assortment under Shipping Constraints

Problem: select assortments per region to maximize expected margin net of shipping costs and stockouts under stochastic demand.

Approach: formulate binary decision variables for item inclusion; build QUBO objective with penalty terms. Use QAOA to sample high-quality assortments, refine via tabu search, and compare against classical integer programming baselines (Gurobi) on medium-scale instances.

Outcomes: evaluate solution quality (objective gap), runtime (wall clock and quantum queue time), and economic impact via offline simulations.

7.3 Case study 3 - Quantum-enhanced fraud detection with QNN embeddings

Problem: detect coordinated checkout fraud where attackers simulate legitimate purchase patterns with minor perturbations.

Hybrid method: construct per-user session embeddings; train QNN encoder to map embeddings into a latent Hilbert space where malicious distributions separate; use classical classifier on QNN outputs.

Evaluation: AUC, precision@k at investigator workload constraints, and adversarial robustness under simulated attacker strategies (poisoning, evasion).

7.4 Case study 4 - Privacy-preserving cross-platform recommendations

Problem: collaborative learning across multiple retailers increases recommendation quality but sharing raw user identifiers violates privacy.

Solution: federated learning augmented with quantum kernels for local feature transformation; encrypt model updates and aggregate with secure multiparty computation. Samuel (2022) and Samuel (2024) demonstrate similar cloud-native architectures for secure data exchange and privacy-focused analytics in energy systems; analogous architectures apply to e-commerce cross-platform collaborations (Samuel, 2022; Samuel, 2024).

8. Training, optimization, and evaluation protocols

8.1 Training QNNs: optimizer choices and gradient estimation

- Use gradient-free optimizers (COBYLA, SPSA) and gradient-based with parameter-shift rule depending on circuit differentiability and noise levels.
- Employ minibatch classical training for outer loops; within each minibatch, estimate quantum expectation via repeated shots; amortize measurement budgets.

8.2 Regularization, constraints, and constrained losses

- Incorporate Lagrangian multipliers for inventory, fairness, or margin constraints; update multipliers via classical dual ascent during training.
- Use constrained policy gradients for RL-type optimization in dynamic pricing.

8.3 Evaluation metrics and offline/online alignment

- Offline proxies: NDCG, MRR, expected revenue, calibration metrics, AUC for binary tasks.
- Value-oriented metrics: revenue lift, margin per session, customer lifetime value uplift.

- Use contextual bandit evaluation or off-policy estimators for logged data to estimate causal effects (Dudík et al., 2014).

8.4 Robust validation: adversarial and distribution shift testing

- Simulate seasonal shifts, catalog changes, and adversarial manipulations.
- Use robust optimization and distributionally robust training where suitable.

9. MLOps, orchestration, and system engineering

9.1 Data pipelines and feature stores

- Continuous ingestion pipelines handle events at scale; event time semantics and watermarking ensure correctness.
- Feature stores maintain offline/online parity, freshness guarantees, and lineage (Sculley et al., 2015).

9.2 Model registry, experimentation, and deployment patterns

- Model registry stores artifacts, metrics, and model cards.
- Canary and shadow deployments provide safety; rollback policies and automated monitors mitigate regressions.

9.3 Monitoring: technical signals and business KPIs

- Technical signals: prediction latency, error rates, model drift, feature distribution changes.
- Business KPIs: conversion, average order value, churn, fraud rates.
- Explainability monitoring: track SHAP distributions and alerts for unexpected feature shifts.

9.4 Cost, scaling, and hybrid compute economics

- Quantum runtime costs include queue time and access fees; hybrid designs minimize quantum calls and use classical proxies where possible. Cost–benefit analyses compare improved business metrics against quantum access and integration costs.

10. Privacy, security, and legal considerations

10.1 Data privacy and differential privacy for e-commerce

- Apply differential privacy for sensitive aggregates (user profiles), and use local DP for telemetry where possible. Federated learning reduces raw data sharing (McMahan et al., 2017).

10.2 Secure model serving and model inversion risks

- Protect model APIs with rate limits and authentication; monitor for model extraction and membership inference. Quantum-era cryptography and post-quantum secure channels should be planned, though current threats are limited (NIST PQC initiatives).

10.3 Regulatory landscape and consumer rights

- GDPR and consumer protection laws require explainability and right to objection for automated profiling; ensure human-in-loop and recourse mechanisms.

10.4 Ethical considerations: fairness, manipulation, and consumer welfare

- Personalization must balance relevance and fairness; avoid exploitative pricing. Implement fairness constraints and human oversight for price personalization.

11. Fraud, adversarial resilience, and economic defenses

11.1 Adversarial threat taxonomy in e-commerce

- Account takeover, synthetic identity creation, coupon abuse, return fraud, triangulation fraud; attackers adapt to detection through adversarial mimicry.

11.2 Defensive strategies and game-theoretic modeling

- Use multi-agent simulations to model attacker–defender dynamics; design detection systems robust to adaptive attackers; incorporate economic costs of false positives into detection thresholds.

11.3 Quantum-enabled cryptographic primitives and secure exchanges

- Quantum computing may enable new cryptographic protocols; concurrently, post-quantum cryptography is needed to protect data against future quantum adversaries. Samuel (2022, 2024) discuss secure cloud-native AI transfer protocols and privacy concerns in energy systems applicable to e-commerce secure exchanges (Samuel, 2022; Samuel, 2024).

12. Interpretability, transparency, and human factors

12.1 Auditable model cards and documentation

- Produce model cards for each production model capturing intent, datasets, performance, fairness tests, and limitations (Mitchell et al., 2019).

12.2 Human interface design for model explanations

- Give investigators concise prioritized rationales; provide actionable insights rather than raw feature attributions.

12.3 Organizational change management

- Training for product, customer service, and legal teams on AI workflows; establish protocols for contested decisions and appeals.

13. Evaluation of quantum readiness and practical adoption criteria

13.1 Problem selection heuristic for quantum pilots

- Prioritize problems where: (a) classical baselines struggle or are computationally intensive; (b) labeled data are scarce but structure is amenable to kernel methods; (c) problem scales are moderate enough for NISQ exploration; (d) economic upside justifies integration costs.

13.2 Benchmarks and success metrics for pilots

- Define success via improvement in business KPIs, solution quality gaps, total cost of ownership, and operational reliability.

13.3 Integration maturity model

- Stages: exploration (simulators), hybrid pilots (small quantum tasks + classical validation), scaled hybrid (repeatable processes), and long-term fault-tolerant integration.

14. Research agenda and roadmap

14.1 Near-term research priorities (0–2 years)

- Develop robust hybrid architectures and benchmark datasets for re-ranking and combinatorial recommendation tasks.
- Advance QNN ansatz design resilient to barren plateaus and noise.
- Empirical studies comparing quantum kernels to classical kernels in small-label, high-dimensional regimes.

14.2 Medium-term priorities (2–5 years)

- Federated quantum-classical learning frameworks for privacy-sensitive cross-platform analytics.
- QAOA and heuristic benchmarking for real-world assortment and routing problems.
- Adversarial robustness methods combining quantum encodings and classical defenses.

14.3 Long-term ambitions (5+ years)

- Fault-tolerant quantum models for large combinatorial optimizations at scale.
- Integration of quantum cryptographic primitives into e-commerce transaction systems.
- Institutional adoption pathways and regulatory standards for quantum-assisted AI systems.

15. Limitations, risks, and open challenges

15.1 Technical limitations and uncertainty of quantum advantage

- Evidence for practical quantum advantage in ML is still limited; careful benchmarking and realistic evaluations are necessary (Preskill, 2018).

15.2 Integration, operational, and human factors risks

- Integration complexity, personnel training, and organizational inertia impede adoption. Economic viability depends on both performance gains and total cost of integration.

15.3 Ethical risks and societal impact

- Personalization and pricing algorithms may exacerbate inequality or manipulate consumer choice if unchecked; design for consumer welfare is essential.

16. Conclusion AI and data science will continue to drive e-commerce innovation. QNNs and hybrid quantum-classical approaches offer promising new tools for specific, computationally demanding subproblems in recommendation, combinatorial optimization, and adversarial detection. Responsible integration demands rigorous method design, realistic evaluation, privacy and security protections, and ethical governance. The roadmap herein guides researchers and practitioners through piloting, validating, and scaling quantum-enhanced AI in e-commerce with scholarly rigor and industrial pragmatism.

References

1. Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. Nature, 549(7671), 195–202. https://doi.org/10.1038/nature23474
2. Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., McClean, J. R., Mitarai, K., Yuan, X., & Yamamoto, N. (2021). Variational quantum algorithms. Nature Reviews Physics, 3, 625–644. https://doi.org/10.1038/s42254-021-00348-9
3. Dudík, M., Langford, J., & Li, L. (2014). Doubly robust policy evaluation and learning. In Proceedings of the 31st International Conference on Machine Learning (ICML), 1097–1105.
4. Farhi, E., Goldstone, J., & Gutmann, S. (2014). A quantum approximate optimization algorithm. arXiv:1411.4028.
5. Fatunmbi, T. O. (2022). Deep learning, Artificial Intelligence and machine learning in cancer: Prognosis, diagnosis and treatment. World Journal of Advanced Research and Reviews, 15(02), 725–739. https://doi.org/10.30574/wjarr.2022.15.2.0359
6. Fatunmbi, T. O. (2024). Advanced frameworks for fraud detection leveraging quantum machine learning and data science in fintech ecosystems. World Journal of Advanced Engineering Technology and Sciences, 12(01), 495–513. https://doi.org/10.30574/wjaets.2024.12.1.0057
7. Fetterly, D., & others. (2016). Scalable candidate generation and ranking with biased sampling. Proceedings of the RecSys.
8. Guerreschi, G. G., & Smelyanskiy, V. N. (2017). Practical optimization for hybrid quantum-classical algorithms. arXiv:1701.01450.
9. Havlíček, V., Córcoles, A. D., Temme, K., Harrow, A. W., Kandala, A., Chow, J. M., & Gambetta, J. M. (2019). Supervised learning with quantum-enhanced feature spaces. Nature, 567, 209–212. https://doi.org/10.1038/s41586-019-0980-2
10. Hidasi, B., Karatzoglou, A., Baltrunas, L., & Tikk, D. (2016). Session-based recommendations with recurrent neural networks. Proceedings of ICLR.
11. Joachims, T., Swaminathan, A., & Schnabel, T. (2017). Unbiased learning to rank with biased feedback. Proceedings of the 10th ACM International Conference on Web Search and Data Mining, 781–789.
12. Kang, W. C., & McAuley, J. (2018). Self-attention for sequential recommendation. Proceedings of the 2018 IEEE International Conference on Data Mining (ICDM).

13. Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. Computer, 42(8), 30–37.

14. Lim, B., Arık, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. International Journal of Forecasting, 37(4), 1748–1764. https://doi.org/10.1016/j.ijforecast.2020.11.011

15. Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. Decision Support Systems, 50(3), 559–569. https://doi.org/10.1016/j.dss.2010.08.006

16. Perdomo-Ortiz, A., Fluegemann, J., Narayanan, S., & others. (2019). Opportunities and challenges in quantum-assisted machine learning for cancer. Journal of Biomedical Informatics, 94, 103196. https://doi.org/10.1016/j.jbi.2019.103196

17. Preskill, J. (2018). Quantum computing in the NISQ era and beyond. Quantum, 2, 79. https://doi.org/10.22331/q-2018-08-06-79

18. Ricci, F., Rokach, L., & Shapira, B. (Eds.). (2015). Recommender systems handbook (2nd ed.). Springer.

19. Samuel, A. J. (2022). AI and machine learning for secure data exchange in decentralized energy markets on the cloud. World Journal of Advanced Research and Reviews, 16(2), 1269–1287. https://doi.org/10.30574/wjarr.2022.16.2.1282

20. Samuel, A. J. (2024). Optimizing energy consumption through AI and cloud analytics: Addressing data privacy and security concerns. World Journal of Advanced Engineering Technology and Sciences, 13(2), 789–806. https://doi.org/10.30574/wjaets.2024.13.2.0609

21. Schuld, M., Sinayskiy, I., & Petruccione, F. (2014). The quest for a quantum neural network. Quantum Information Processing, 13(11), 2567–2586. https://doi.org/10.1007/s11128-014-0809-8

22. Topol, E. (2019). Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again. Basic Books.

23. Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Philip, S. Y. (2020). A comprehensive survey on graph neural networks. IEEE Transactions on Neural Networks and Learning Systems, 32(1), 4–24. https://doi.org/10.1109/TNNLS.2020.2978386