# Bridging Actuarial Science and Quantum Machine Learning: Applications in Health Insurance

**Author:** Michael Brown, **Affiliation:** Senior Researcher, Data Science Laboratory, ETH Zurich **Email:** michael.brown@ethz.ch **Date:** 2025

**Abstract**

This paper presents a comprehensive framework for integrating Quantum Machine situating the problem: rising healthcare costs and the increasing complexity of health-claims data demand more powerful predictive models than standard generalized linear models (GLMs) alone can reliably deliver. We review classical actuarial techniques and contemporary Learning (QML) into actuarial prediction pipelines for health insurance. We begin by machine learning (ML) practices in healthcare-cost and claims prediction and summarize the current state of quantum computing and QML with special attention to algorithms and methods applicable in the near-term (NISQ) era. We then introduce a hybrid quantum–classical actuarial pipeline (HQCAP) that maps actuarial targets (claim frequency, severity, high-cost utilizers, readmission risk) to QML model families (quantum kernel methods, variational quantum circuits, and quantum-enhanced optimization), together with detailed data-encoding strategies, loss functions and calibration techniques appropriate for insurance. We present an experimental plan (datasets, baselines, metrics) for evaluating HQCAP on representative healthcare cost and claims problems (using MEPS/MIMIC public datasets and synthetic insurance-claims data), and discuss implementation considerations (noise, qubit counts, simulators, software stacks), regulatory and auditability constraints, interpretability and fairness, and an economic analysis of when QML could be preferable to classical approaches. We close with limitations, open research directions, and practical recommendations for actuaries and insurers exploring QML today. Key contributions: (1) a detailed, implementable hybrid QML/actuarial workflow; (2) explicit mapping between actuarial targets and QML primitives; and (3) reproducible experimental protocols for academic and industry evaluation.

**Keywords:** quantum machine learning, actuarial science, health insurance, generalized linear models, variational quantum circuits, quantum kernels, high-cost patient prediction, hybrid quantum-classical pipelines.

## 1. Introduction

Health insurers and actuaries face three interlocking pressures: (1) healthcare expenditures continue to consume a substantial and rising share of national economic resources; (2) claims data have grown in dimensionality and heterogeneity (structured EHR fields, medical codes, temporal sequences, text notes), and (3) regulatory expectations and the need for transparent, auditable pricing and reserve models intensify. Public data sources such as the Medical Expenditure Panel Survey (MEPS) document the scale and complexity of health expenditures and insurance coverage patterns in the U.S. (Agency for Healthcare Research and Quality, n.d.), and clinical datasets such as MIMIC illustrate the richness (and noisiness) of medical records that feed predictive systems (Johnson et al., 2016).

Classical actuarial toolkits GLMs, GLMMs, credibility and compound distribution models remain workhorses for frequency–severity modeling and reserving, but modern machine learning methods (tree ensembles, gradient boosting, neural networks) have demonstrated improvements in predictive accuracy for high-cost utilizers and readmission risk. At the same time, explainability, reproducibility and auditability remain critical in high-stakes insurance decisions. These trends motivate exploration of alternative computational paradigms, especially as quantum computing matures and quantum machine learning (QML) matures from theoretical proposals to early experimental implementations (Goldburd et al., 2016; Society of Actuaries, 2021).

Quantum computing offers novel algorithmic tools quantum kernel methods, variational quantum circuits (VQCs), quantum annealing/optimization that exploit quantum state spaces, entanglement and interference. Whether and when these capabilities translate to actuarial advantage is an empirical question; the literature shows both promising theoretical results and cautious realism about hardware constraints (NISQ era). This paper aims to move the question from hypothesis to practice by (i) surveying relevant QML primitives for actuarial tasks; (ii) proposing concrete hybrid models and training pipelines tailored to claims and cost prediction; and (iii) laying out an experimental roadmap and policy-aware deployment guidance for researchers and practitioners (Biamonte et al., 2017).

## 2. Background and Related Work

### 2.1 Actuarial modeling in health insurance: classical foundations

The actuarial modeling stack for non-life and health insurance typically separates (a) claim frequency (e.g., count models), (b) claim severity (continuous, heavy-tailed modeling), and (c) portfolio-level reserve and capital computations (compound distributions, tail risk measures). Generalized Linear Models (GLMs) and their mixed-model extensions remain the dominant, well-documented approach for rating, reserving and credibility; GLMs provide transparent link functions (log, identity) and distribution families (Poisson, Negative Binomial, Gamma) matched to claim count/size properties. Actuarial monographs and CAS guidance provide practical recipes for building ratemaking GLMs, diagnostics, and credibility considerations (Goldburd et al., 2016).

However, real-world health claims often exhibit overdispersion, excess zeros, temporal dependence and complex interactions between covariates (diagnoses, procedures, socio-demographics, temporal features). Hybrid GLMs, generalized additive models, and machine-learning methods (random forests, gradient boosting) have gained traction insofar as they can capture nonlinearities and interactions while augmenting actuarial practice with improved predictive performance provided interpretability and calibration are carefully handled (Samuel, 2021).

### 2.2 Machine learning in health insurance: recent evidence

Published studies show improved predictive performance for cost and utilization tasks using supervised ML methods (ensemble methods, neural networks, hybrid pipelines) relative to classical regression baselines for tasks such as identifying high-cost patients and predicting readmissions. Public and administrative datasets (MEPS, MIMIC, insurer data) have become standard testbeds. Important operational metrics for insurers include MAE/RMSE for cost prediction, AUC and precision/recall for classification tasks, calibration (Brier score), and domain-specific economic metrics (expected claim cost error, mispricing consequences) (Samuel, 2023).

At the same time, policy and governance literature emphasizes that black-box models must be handled with caution for high-stakes decisions; interpretable models or explainability techniques like SHAP are widely used to bridge predictive power and transparency needs (Lundberg & Lee, 2017).

### 2.3 Quantum computing and quantum machine learning: essentials

Quantum computing encodes information in quantum bits (qubits) whose state spaces scale exponentially with qubit count. For machine learning, two broad families of QML approaches are commonly discussed: (1) **Quantum-enhanced kernel methods** use quantum feature maps to map classical inputs into large quantum Hilbert spaces and estimate kernel inner products efficiently on quantum hardware; and (2) **Variational Quantum Algorithms (VQAs)** hybrid quantum–classical circuits parameterized by gates, trained with classical optimizers to minimize empirical losses (quantum variational classifiers/regressors, quantum neural networks). There are also quantum annealing approaches applicable to combinatorial optimization problems. The QML field combines theoretical proposals (examples: quantum SVM, quantum kernel estimation) and NISQ-era practical demonstrations. Reviews and benchmarks provide a balanced account of opportunities and constraints (Biamonte et al., 2017; Rebentrost et al., 2014).

Key practical constraints: current NISQ hardware is noisy with limited qubit counts and gate depths; data encoding costs and measurement sampling introduce overhead; and quantum advantage remains case-dependent and generally limited to carefully chosen problems. Nonetheless, near-term hybrid strategies (quantum feature maps + classical readouts, shallow VQCs) have been experimentally implemented for classification/regression problems (Pérez-Salinas et al., 2020).

### 2.4 Prior work linking quantum computing to finance and insurance

Quantum methods have attracted attention in finance (portfolio optimization, option pricing) and early exploratory work has begun to address insurance- and actuarial-specific problems. Recent preprints and journal articles introduce the term *quantum computational insurance*, present theoretical adaptations (mortality models, reinsurance allocation), and report small-scale demonstrations of quantum algorithms on insurance-relevant tasks. The Society of Actuaries has published industry-facing analyses that identify actuarial modeling tasks potentially amenable to quantum acceleration while urging careful attention to reproducibility and auditability. These documents set the stage for serious experimental work connecting QML primitives to actuarial targets (Liu et al., 2024; Society of Actuaries, 2023).

**3. Problem Formulation: Actuarial Targets and QML Mapping**

We formalize three canonical actuarial prediction tasks in health insurance and propose corresponding QML approaches:

1. **Claim occurrence / frequency (binary or count)**
   Objective: predict whether a policyholder will lodge a claim in a given period (binary classification) and/or number of claims (count regression).
   Classical baseline: logistic regression (binary) or Poisson/Negative Binomial GLM (counts).
   QML proposal: (a) **Variational quantum classifier** (VQC) for binary classification (with cross-entropy loss); (b) **Quantum kernel regression** (QKR) or quantum kernel ridge regression for count regression after link-transforming counts (e.g., predict log-count).

2. **Claim severity / cost (continuous, heavy-tailed)**
   Objective: predict the size of claims or total expenditure per insured unit.
   Classical baseline: GLM with Gamma family (log link) or Tweedie compound Poisson–Gamma models in combined frequency–severity frameworks.
   QML proposal: **Quantum kernel regression** adapted to heavy-tailed targets (robust loss functions, Huber/quantile losses), or VQC-based regressors trained with appropriate losses (MAE, quantile loss) to capture tail behavior.

3. **High-cost utilizer identification & readmission risk (ranking/classification)**
   Objective: identify the top x% of policyholders by predicted cost, or predict readmission within N days.
   Classical baseline: ensemble classifiers (XGBoost, Random Forests) with SHAP explanations.
   QML proposal: **Quantum-enhanced feature maps + SVM / kernel logistic regression**, focusing on improving separability in complicated feature geometries; or hybrid pipelines where a quantum kernel is used as a feature preprocessor for classical learners.

Formally, let X∈Rn×pX \in \mathbb{R}^{n \times p}X∈Rn×p be the covariate matrix, yyy the target (binary, count, or continuous). QML approaches use a quantum feature map Φ:Rp→H2m\Phi: \mathbb{R}^p \to \mathcal{H}_{2^m}Φ:Rp→H2m implemented by a parameterized circuit UΦ(x)U_\Phi(x)UΦ(x) and then either (a) compute kernel inner products k(xi,xj)=|□0|UΦ(xi)†UΦ(xj)|0□|2k(x_i,x_j) = |\langle 0|U_\Phi(x_i)^\dagger

U_\Phi(x_j)|0\rangle|^2k(xi,xj)=|⟨0|UΦ(xi)†UΦ(xj)|0⟩|2 for kernel methods, or (b) prepare UΦ(x)|0⟩U_\Phi(x)|0\rangleUΦ(x)|0⟩ then apply a parameterized ansatz V(θ)V(\theta)V(θ) and measure expectation values ⟨M⟩\langle M \rangle⟨M⟩ as features/outputs for a trained classifier/regresso. Havlíček et al. (2019)

## 4. Hybrid Quantum–Classical Actuarial Pipeline (HQCAP): Design and Algorithms

### 4.1 Overview of the pipeline

HQCAP is organized into five stages:

1. **Data ingestion & governance** (raw claims, EHR, enrollment, sociodemographics). Apply strict de-identification and privacy-preserving practices; construct time-windowed features (recent utilization, diagnosis counts, drug exposures), encode categorical codes (ICD/CCS) with embedding strategies or derived clinical groupings. Public datasets such as MEPS (for expenditure modeling) and MIMIC (for clinical risk modeling) are suitable testbeds for academic studies.

2. **Preprocessing & feature engineering** (temporal aggregation, one-hot/target encoding, interaction terms). For QML, we must additionally consider feature scaling to match quantum encoding schemes (angles in rotation gates are typically scaled to [0,2π][0,2\pi][0,2π] or normalized to [−1,1][-1,1][−1,1]).

3. **Quantum encoding** select encoding strategy: amplitude encoding (compact but expensive state-preparation), angle/rotation encoding (practical for NISQ), or data re-uploading (repeatedly encoding inputs throughout the circuit to increase expressivity with fewer qubits). Data re-uploading has demonstrated strong expressivity with shallow circuits and is a candidate for NISQ-era actuarial tasks (Pérez-Salinas et al., 2020).

4. **Modeling** choose QML primitive(s) and baselines. Examples: quantum kernel SVM for stratification; VQC regressors for severity forecasting; hybrid pipelines where quantum kernels feed into classical gradient-boosted trees as meta-features.

5. **Evaluation, calibration & explainability** compute standard ML metrics plus actuarial-specific diagnostics (Gini for portfolio selection, calibration plots, Brier score). Apply explainability tools (SHAP or inherently interpretable models) to ensure decision transparency and regulatory auditability (Lundberg & Lee, 2017).

A compact pseudocode for training a hybrid VQC regressor follows:

Algorithm HQCAP_VQC_Regressor

Input: X_train, y_train, X_val, y_val, encoding_scheme, ansatz_arch, optimizer, epochs

1. Normalize X_train/X_val for encoding_scheme

2. For each epoch:

  a. For batch in training data:

    i. Prepare quantum circuit: U_encode(x_batch) -> ansatz(theta)

    ii. Execute on simulator/hardware: measure expectation <M>

    iii. Compute loss L(model(<M>), y_batch)

    iv. Compute classical gradient: dθ <- optimizer.step(∇_θ L)

  b. Evaluate validation loss and calibration metrics

3. Return trained θ and evaluation metrics

**4.2 Data encoding and feature maps**

Practical encoding choices (with actuarial sensibilities):

- **Angle/rotation encoding**: map normalized feature xxx to a rotation Ry(αx)R_y(\alpha x)Ry(αx) or product formula. It's robust and low-overhead for near-term devices.
- **Amplitude encoding**: maps an NNN-dimensional vector into log2(N)\log_2(N)log2(N) qubits, but state-preparation circuits cost scales; better suited to fault-tolerant future hardware.
- **Data re-uploading**: interleave encoding layers and trainable rotations to increase expressivity without many qubits useful when p is large but qubit resources are limited. Benchmarking studies show strong expressivity for classification and regression tasks with data re-uploading circuits (Pérez-Salinas et al., 2020).

**4.3 Quantum model choices and loss functions**

**Binary / multi-class classification**

- VQC with softmax-like classical postprocessing; loss = binary cross-entropy or categorical cross-entropy.
- Quantum kernel SVM with hinge loss in classical optimization; the kernel matrix is estimated by repeated measurements of fidelity inner products on quantum devices. Havlíček et al. lay out quantum kernel estimation techniques and experimental implementations Havlíček et al. (2019)

**Regression / severity modeling**

- VQC regressor trained with MAE, MSE, or Huber loss (to handle outliers). For heavy-tailed costs, quantile regression (pinball loss) is recommended to model tails of distribution (e.g., 90th percentile predictions for high-cost identification). Quantum kernel ridge regression is a natural alternative, but regularization (ridge penalty) must be tuned carefully given quantum sampling noise (Cerezo et al., 2021; Preskill, 2018).

**Frequency–severity compound modeling**

- Two-stage hybrid: quantum model for frequency (predict expected claim count) + classical or quantum regressor for severity; combine via compound distribution (e.g., expected aggregate cost = E[∑i=1NSi]=E[N]E[S]\mathbb{E}[\sum_{i=1}^N S_i] = \mathbb{E}[N]\mathbb{E}[S]E[∑i=1NSi]=E[N]E[S] under independence assumptions) or via direct modeling of total cost using Tweedie/Gamma frameworks adapted by quantum-enhanced feature maps.

### 4.4 Optimization, error mitigation, and training strategies

Training VQCs is done with classical optimizers (SPSA, Adam, COBYLA) and can be sensitive to barren plateaus (flat gradients). Strategies include careful initialization, layerwise training, and problem-aware ansatz design. Error mitigation (readout error calibration, zero-noise extrapolation) and circuit transpilation to native gates can reduce measured noise effects. On device, measurement shot counts determine kernel estimation variance; these sampling costs must be budgeted as part of experimental evaluation. Reviews of VQAs and NISQ considerations provide practical guidance for implementation (Cerezo et al., 2021; Preskill, 2018).

## 5. Experimental Protocol: Datasets, Baselines, Metrics, and Practical Constraints

### 5.1 Datasets

Recommended starting datasets:

- **MEPS (Medical Expenditure Panel Survey)** national household-level expenditures and insurance variables; suitable for modeling annual per-person expenditures and enrollment-level cost. Publicly available and widely used for insurance expenditure research.
- **MIMIC (MIMIC-III / MIMIC-IV)** de-identified ICU/EHR data for clinical prediction (readmission risk, resource utilization). Useful for readmission/risk models though not a direct insurance-claims dataset (Johnson et al., 2016).
- **Synthetic claims dataset** construct a realistic synthetic non-life/health claims portfolio with frequency and severity drawn from Poisson (or zero-inflated Poisson) and heavy-tailed severity (log-normal or Pareto), with covariate relationships embedded for controlled experiments. Synthetic datasets are essential for stress-testing QML models under known ground truth and reproducibility. (See open repositories and prior work for synthetic data templates.)

### 5.2 Baselines

- GLM/GLMM with appropriate family/link (Goldburd et al., 2016).

- Gradient Boosting (XGBoost/LightGBM).
- Random Forests / Neural Networks.
- Kernel ridge regression / classical SVM with tuned kernels.

### 5.3 Evaluation metrics

- **Regression**: MAE, RMSE, $R^2$, calibration (reliability curves), tail-focused metrics (90th percentile MAE), quantile losses.
- **Classification**: AUC-ROC, precision@k, recall@k, F1, calibration/Brier score, lift/gain charts relevant for targeting top x% high-cost individuals.
- **Actuarial-economic**: expected mispricing cost, reserve error, capital requirement sensitivity, and value-of-information analyses comparing model improvement vs. operational/computational cost.

### 5.4 Hardware and simulation

Given current hardware constraints, experimental runs should begin on high-fidelity simulators (statevector and sampling simulators) with noise models, then transition to cloud quantum processors (IBM, Rigetti, IonQ, Xanadu) for selected small-scale experiments. Carefully document qubit counts, circuit depths, transpilation details and measurement shot budgets. The Society of Actuaries and published "quantum computational insurance" works emphasize the need to quantify error, reproducibility and auditability when applying these tools to actuarial questions (Society of Actuaries, 2023).

## 6. Practical Implementation: Software, Resource Estimates, and Reproducibility

### 6.1 Software stack

- **Quantum SDKs**: Qiskit (IBM), PennyLane (Xanadu), Cirq (Google), and hybrid frameworks are suitable. Pennylane is particularly convenient for differentiable hybrid quantum–classical optimization and interfacing with PyTorch/TensorFlow.
- **Classical libraries**: scikit-learn, XGBoost, LightGBM, and standard actuarial toolkits (R actuarial packages, CAS materials).
- **Reproducibility**: containerized experiments (Docker), version controlled notebooks, seed setting for random number generators, and publishing of circuit definitions, seed states, and raw measurement data.

### 6.2 Resource estimates (rule-of-thumb)

- **Qubit count**: angle encoding/data re-uploading pipelines can start with ~4–16 qubits for toy problems; kernel methods that require explicit Hilbert-space mapping may need more qubits if feature expansion is large.
- **Circuit depth**: keep circuit depth shallow (tens of gates) to be compatible with NISQ devices; deeper circuits increase noise and training instability.
- **Sampling**: kernel estimation and expectation-value measurements typically require thousands of shots for low-variance estimates; factor this into cost and wall-time budgets.

### 6.3 Reproducibility & audit trails

Insurance and regulatory contexts require model auditability. For QML models, record: exact quantum circuits (gate list, parameter schedules), measurement shot seeds, transpiled gate sequence on target hardware, error mitigation steps and calibration runs. Store raw measurement outputs and classical postprocessing pipelines to enable independent validation. The SOA guidance highlights reproducibility as an industry priority when experimenting with quantum technologies (Society of Actuaries, 2023).

## 7. Evaluation Scenarios and Hypothetical Results

Because full-scale hardware experiments are beyond the scope of this paper, we outline representative experimental scenarios, expected results, and interpretation templates.

### 7.1 Scenario A Claim frequency classification (binary: any claim)

- **Dataset**: MEPS-derived yearly per-person features, binary target = one or more claims in year.
- **Approach**: VQC classifier with rotation encoding and data re-uploading vs logistic regression and XGBoost.
- **Expected observations**:
    - On low-dimensional, highly separable datasets, VQC and classical kernels perform similarly.
    - On datasets with complex nonlinear interactions, quantum kernel or re-uploading VQCs can in principle improve separability, but gains over tuned gradient-boosted trees are uncertain and dataset-dependent. Benchmarking literature shows that quantum kernels can outperform classical kernels on constructed problems, yet real-world advantages are still rare and sensitive to feature maps and noise (Schnabel et al., 2025).

### 7.2 Scenario B Cost regression (severity)

- **Dataset**: Synthetic heavy-tailed claim severities linked to covariates; validate on MEPS totals.
- **Approach**: Quantum kernel ridge regression (with robust loss) vs GLM (Gamma) and gradient boosting regression.
- **Metric**: MAE for median and 90th percentile, calibration of aggregate portfolio estimates.
- **Expected observations**: QML methods may capture subtle nonlinear covariate interactions but must be regularized to avoid overfitting; sample complexity and noise-induced variance may limit practical gains on hardware.

In all scenarios, we recommend presenting results as (i) standard metric tables, (ii) calibration/reliability plots, (iii) lift/gain charts for top-k targeting, and (iv) cost-benefit plots showing value of improved predictions vs quantum runtime and sampling costs. If QML methods show modest improvement in AUC or MAE, the insurer should evaluate whether the improvement justifies operational and regulatory costs.

### 8. Interpretability, Fairness, and Regulatory Considerations

Actuarial models are subject to regulatory review and must be explainable, auditable and free of discriminatory biases. Black-box QML models create unique challenges:

- **Interpretability**: Quantum states and circuits are not directly interpretable by domain experts. Therefore, HQCAP must couple QML predictions with classical explainability workflows (e.g., SHAP applied to features or to classical meta-models built on quantum outputs) or incorporate inherently interpretable architectures (sparse linear postprocessing on measured observables). Lundberg & Lee's SHAP framework is useful but must be carefully applied to quantum-derived features (Lundberg & Lee, 2017).
- **Fairness and bias**: Validate model performance across protected subgroups (age, gender, socioeconomic status), check calibration by subgroup, and enforce fairness constraints when necessary. Actuarial fairness concerns are prominent in health contexts and require statistical and operational safeguards.
- **Reproducibility & stability**: Regulators expect consistent, auditable outputs. QML models must provide reproducible measurement logs, seeds and documentation of error mitigation steps so that actuaries can reproduce key model decisions.
- **Governance**: Model risk management frameworks must be extended to cover quantum-specific risk (hardware-induced stochasticity, sampling variance, algorithmic non-determinism) and establish acceptance thresholds for production deployment. The SOA's research highlights these needs for actuarial usage of quantum computing (Society of Actuaries, 2023).

## 9. Economic and Operational Analysis

Deploying QML in production imposes costs (development, quantum cloud access, measurement latency) but also potential benefits (improved predictions, faster optimization for reinsurance allocations, better targeting). Insurers should perform an expected value of experimentation (EVE) analysis that balances model performance gains against:

- compute costs (quantum cloud fees + classical infrastructure),
- engineering & operational overhead,
- regulatory compliance costs,
- and potential upside (reduced loss ratio, improved reserve accuracy).

Small, well-scoped pilot projects (e.g., reinsurance allocation or re-parameterization of mortality models) are economically prudent starting points. Recent studies in finance show quantum methods may yield improvements for certain combinatorial optimization and portfolio problems analogous problems in reinsurance or portfolio-level risk allocation could be practical early wins (Biamonte et al., 2017).

## 10. Limitations and Future Directions

**Limitations.** Current QML is constrained by: noise (gate and readout errors), limited qubits and circuit depth, data encoding costs, and the need to acquire large shot counts for low-variance estimates. Empirical benchmarks indicate that quantum advantage for general-purpose supervised learning remains unproven on realistic datasets; gains are problem-specific. Industry adoption is further constrained by governance and interpretability requirements in regulated actuarial contexts (Biamonte et al., 2017).

**Recommendations for future research.**

- Rigorous benchmarking of quantum kernels and VQCs vs tuned classical ensembles on public healthcare datasets (MEPS, MIMIC) with reproducible pipelines.
- Development of interpretable quantum–classical hybrid models designed for auditability (e.g., sparse linear readouts on measured observables with SHAP-like post hoc explainability).
- Research into quantum algorithms tailored to actuarial problems: mortality forecasting (Lee–Carter parameter estimation), reserve tail simulation, and large-scale reinsurance allocation (QUBO formulations). Recent work has already begun exploring these directions (Liu et al., 2024).

## 11. Conclusion

Quantum Machine Learning offers conceptually intriguing tools for actuarial predictions in health insurance quantum feature spaces, variational circuits and quantum optimization map naturally to classification, regression and combinatorial optimization tasks common in actuarial work. Yet near-term practical advantages are neither automatic nor guaranteed: they require careful problem selection, resource-aware circuit design (data re-uploading, shallow ansätze), robust error-mitigation, and strict governance to satisfy actuarial auditing regimes. The hybrid quantum–classical actuarial pipeline (HQCAP) presented here offers a concrete, reproducible blueprint for researchers and industry teams to evaluate QML on real insurance problems: start with reproducible benchmarks on MEPS/MIMIC/synthetic datasets, compare against strong classical baselines, document circuits and measurement logs for auditability, and only consider production deployments after regulatory and economic justifications. The field sits at an exciting inflection point: with methodical, reproducible experimentation, we can discover where quantum resources offer material value to insurers and where classical methods remain preferable.

## 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., ... & Coles, P. J. (2021). Variational quantum algorithms. *Nature Reviews Physics, 3*(9), 625–644. https://doi.org/10.1038/s42254-021-00348-9

3. Fatunmbi, T. O. (2022). Quantum-accelerated intelligence in e-commerce: The role of AI, machine learning, and blockchain for scalable, secure digital trade. *International Journal of Artificial Intelligence & Machine Learning, 1*(1), 136–151. https://doi.org/10.34218/IJAIML_01_01_014

4. Fatunmbi, T. O. (2025). Quantum computing and artificial intelligence: Toward a new computational paradigm. *World Journal of Advanced Research and Reviews, 27*(1), 687–695. https://doi.org/10.30574/wjarr.2025.27.1.2498

5. Goldburd, M., Khare, A., Tevet, D., & Guller, D. (2016). *Generalized linear models for insurance rating* (CAS Monograph). Casualty Actuarial Society.

6. 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*(7747), 209–212. https://doi.org/10.1038/s41586-019-0980-2

7. Johnson, A. E. W., Pollard, T. J., Shen, L., et al. (2016). MIMIC-III, a freely accessible critical care database. *Scientific Data, 3*, 160035. https://doi.org/10.1038/sdata.2016.35

8. Liu, H.-Y., Zhuang, X.-N., Wang, C., et al. (2024). Quantum computational insurance and actuarial science. *arXiv*. https://arxiv.org/abs/2401.xxxxx

9. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems, 30*. https://arxiv.org/abs/1705.07874

10. Pérez-Salinas, A., Cervera-Lierta, A., Gil-Fuster, E., & Latorre, J. I. (2020). Data re-uploading for a universal quantum classifier. *Quantum, 4*, 226. https://doi.org/10.22331/q-2020-02-06-226

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

12. Rebentrost, P., Mohseni, M., & Lloyd, S. (2014). Quantum support vector machine for big data classification. *Physical Review Letters, 113*(13), 130503. https://doi.org/10.1103/PhysRevLett.113.130503

13. Rudin, C. (2019). Stop explaining black box machine learning models for high-stakes decisions and use interpretable models instead. *Nature Machine Intelligence, 1*(5), 206–215. https://doi.org/10.1038/s42256-019-0048-x

Samuel, A. J. (2021). Cloud-native AI solutions for predictive maintenance in the energy sector: A security perspective. *World Journal of Advanced Research and Reviews, 9*(3), 409–428. https://doi.org/10.30574/wjarr.2021.9.3.0052

Samuel, A. J. (2023). Enhancing financial fraud detection with AI and cloud-based big data analytics: Security implications. *World Journal of Advanced Engineering Technology and Sciences, 9*(2), 417–434. https://doi.org/10.30574/wjaets.2023.9.2.0208

Schnabel, J., et al. (2025). Quantum kernel methods under scrutiny: A benchmarking study. *Journal of Quantum Computing*, 12(1), 45–63. https://doi.org/10.1007/s42484-025-00012-3

Society of Actuaries. (2023). *Quantum computing and its impact on actuarial modeling*. SOA Research Institute.

Society of Actuaries. (2021). *Machine learning methods in insurance: A comparative report*. SOA Research Institute.