

Quantum Twin Learning: A Hybrid Quantum-Classical Methodology for Early Prediction of Neurodegenerative Disease Progression using Multimodal Patient Digital Twins

Amaravarapu Pramodkumar¹, Prof. (Dr.) S. K. Singh²

¹ Research Fellow, Lincoln Global Postdoctoral and Research Associate Programme, Lincoln University College, 47301, Petaling Jaya, Selangor Darul Ehsan, Malaysia.

² Professor & Director, Amity Institute of Information Technology, Amity University Uttar Pradesh, Lucknow Campus.

Email ID : amaravarapupramod@gmail.com

Abstract: Neurodegenerative disorders progress silently for up to two decades before diagnosis, yet clinical prediction tools remain snapshot-based, single-modality and population-averaged, so forecasting begins only after irreversible neuronal loss. This paper proposes Quantum Twin Learning (QTL), a methodology that couples a continuously re-synchronised multimodal patient digital twin with a hybrid quantum-classical learner to estimate patient-specific progression trajectories inside the preclinical window. The twin maintains a time-aligned latent state per patient, built by a fusion layer that harmonises imaging, health records, genomics, wearable telemetry and speech across sampling rates and tolerates missing modalities. A compact variational quantum circuit, fed by a classical encoder, operates on the twin state rather than on raw scans and learns cross-modal correlations with few trainable parameters. Three prediction heads output a staged trajectory, a time-to-event risk horizon and calibrated uncertainty, trained with a joint objective and updated each time a new observation arrives. An evaluation protocol on public longitudinal cohorts, with classical baselines and clinically meaningful metrics, is specified so that the method can be validated without hardware-dependent claims. The methodology is designed for hospital triage, remote monitoring and precision therapy timing in next-generation smart healthcare systems.

Keywords: Neurodegenerative disease; Progression prediction; Patient digital twin; Multimodal fusion; Variational quantum circuit; Hybrid quantum-classical learning.

1. Introduction

More than fifty-five million people live with dementia and the figure is projected to exceed one hundred and thirty million by 2050, driven by demographic ageing [1], [2]. In Alzheimer's, Parkinson's, amyotrophic lateral sclerosis (ALS) and Huntington's disease, pathology accumulates for years before the first observable symptom, so a substantial fraction of the affected tissue is already lost when a diagnosis is reached [3], [4]. Clinical practice remains symptom-triggered: a patient reports a change, imaging and cognitive tests follow, and a decision is made where declining neuronal integrity crosses a symptom threshold (Figure 1). Prediction therefore starts late and treatment falls back on generic protocols without an individual timeline.

The problem addressed here is trajectory estimation for one specific patient inside the preclinical window, where imaging, gait, speech and biochemical signals already drift measurably [3], [5]. Existing approaches each solve part of the problem. Deep models extract spatial and temporal features but are data-hungry, opaque and brittle when a modality is absent [12]–[14]. Multimodal fusion improves discrimination but rarely maintains a single evolving per-patient state [5], [6]. Digital twins personalise the unit of analysis but rely on conventional learners with thin predictive intelligence [7], [8]. Quantum machine learning offers expressive, parameter-efficient representations of high-dimensional correlated features but has not been connected to longitudinal patient modelling [9]–[11].

This paper proposes Quantum Twin Learning (QTL), a methodology that integrates these components into one loop: a multimodal fusion layer feeds a patient digital twin, a hybrid quantum-classical model learns from the twin state, and every new observation re-synchronises the twin and refreshes the forecast. The contributions are (i) a formal twin-state formulation with missing-modality tolerance, (ii) a compact hybrid quantum learner operating on the twin state with a joint trajectory, risk-horizon and uncertainty objective, and (iii) a hardware-agnostic evaluation protocol on public longitudinal cohorts.




Figure 1. Neuronal integrity declines for years before the symptom threshold; QTL targets prediction inside the shaded preclinical window.

2. Related Work

Classical models on tabular clinical variables are cheap and interpretable but reach an accuracy ceiling [12]. Convolutional networks dominate MRI and PET analysis [14], and recurrent and attention architectures model longitudinal records [15], [16]; the TADPOLE benchmark showed that ensembles forecast cognitive scores yet remain label-hungry and fragile to missing data [13]. Reviews of imaging-plus-record fusion report consistent gains over unimodal baselines [5], [6], but few architectures hold one time-aligned state per patient. Digital twins have demonstrated a closed monitoring-prediction-decision loop in cardiology [7], while neurological twins remain conceptual [8], [17]. Variational quantum circuits and quantum kernels provide expressive feature maps [9]–[11], [18], hybrid pipelines have been shown on medical images [19], and federated training addresses privacy for distributed cohorts [21]; on noisy intermediate-scale hardware only a classical encoder with a compact quantum head is deployable [20]. Table 1 positions the proposed methodology against these families.

Table 1. Positioning of the proposed methodology against existing method families

Method family	Technique	Multimodal fusion	Temporal modelling	Personalisation	Uncertainty	Hardware ready	Key limitation
[12] Classical ML	SVM, RF, regression	Low	Low	No	No	Yes	Hand-crafted features; accuracy ceiling
[14] CNN imaging	3D CNN on MRI / PET	Low	No	No	No	Yes	Single modality; label hungry
[15], [16] Sequence / attention	LSTM, Transformer fusion	High	High	Low	Low	Yes	Opaque; cohort-level fit
[7] Cardiac digital twin	Mechanistic + ML twin	Medium	High	High	Medium	Yes	Not neuro; thin learning core
[11] Quantum neural nets	Variational circuits	Low	No	No	No	No	NISQ noise; no clinical use
[19] Hybrid QML imaging	CNN + quantum head	Low	No	No	No	Partial	Single modality; no twin
QTL (proposed)	Hybrid QML on twin state	High	High	High	High	Partial	Tied to hardware maturity

3. Proposed Methodology: Quantum Twin Learning

3.1 System Architecture

Figure 2 shows the seven-layer architecture. The patient (L1) is the unit of analysis. Data acquisition (L2) collects MRI and PET, health records, genomics, wearable telemetry, clinical text and speech. Multimodal fusion (L3) produces a common representation. The patient digital twin (L4) maintains a latent state that is simulated forward. The hybrid quantum block (L5) learns from that state, prediction (L6) emits trajectory, risk horizon and uncertainty, and clinical decision support (L7) returns explainable guidance. The feedback path re-synchronises the twin whenever a new observation arrives, so the system is a loop rather than a

pipeline. Interfaces between layers are contracts, allowing any encoder, twin formulation or quantum ansatz to be replaced without redesign.




Figure 2. Seven-layer Quantum Twin Learning architecture with the re-synchronisation loop.

3.2 Multimodal Fusion and Twin State Formulation

Let a patient be observed at irregular times $t_1 < t_2 < \dots < t_k$ with modality set $M = \{\text{MRI, PET, EHR, GEN, WEAR, TEXT, SPEECH}\}$. Each modality m has an encoder f_m that maps its raw input x_m to an embedding $e_m \in \mathbb{R}^{d_m}$ together with a presence mask $\delta_m \in \{0, 1\}$. Embeddings are aligned to a common clock by carrying forward the last available embedding with an age feature Δt_m , and missing modalities are represented by a learned placeholder vector weighted by δ_m , so the twin never requires all streams to be present. A gated attention fusion produces the observation vector $o_t = \sum_m \alpha_m \delta_m W_m e_m$, where attention weights α_m depend on modality age and quality. The twin state $z_t \in \mathbb{R}^d$ is updated by a state-space recurrence $z_t = g(z_{t-1}, o_t, \Delta t)$ with a residual, time-aware gate, and a forward model $h(z_t, \tau)$ simulates the state τ months ahead, enabling what-if perturbation of therapy variables without touching the patient.

3.3 Hybrid Quantum Learning Block

Figure 3 details L5. A classical encoder compresses z_t to n features, which are angle-encoded onto n qubits by $R_y(x_i)$ rotations. L variational layers apply trainable R_r and R_y rotations followed by a ring of CNOT gates, so that entanglement expresses interactions between modalities that were fused into the state. Measuring $\langle Z_i \rangle$ on each qubit yields an n -dimensional quantum feature vector q_t that is passed to three classical heads. With n between 6 and 10 and L between 2 and 4 the circuit remains shallow enough for simulation and for current hardware, and its parameter count grows linearly in n and L rather than quadratically in feature dimension. Data re-uploading is used to raise expressivity without adding qubits, and the parameters θ are updated by an ordinary classical optimiser using the parameter-shift rule, making the block a drop-in replacement for a dense layer.




Figure 3. Hybrid quantum block: classical encoder, angle encoding, variational layers, measurement and three prediction heads.

3.4 Prediction Heads and Training Objective

The trajectory head predicts the disease stage and a continuous clinical score at horizons $\tau \in \{6, 12, 24, 36\}$ months. The risk-horizon head models time-to-event (for example conversion from mild cognitive

impairment to dementia) with a discrete-time survival likelihood. The uncertainty head outputs a heteroscedastic variance and is calibrated with temperature scaling on a held-out fold. The joint objective is $\mathcal{L} = \lambda_1 \mathcal{L}_{\text{stage}} + \lambda_2 \mathcal{L}_{\text{score}} + \lambda_3 \mathcal{L}_{\text{surv}} + \lambda_4 \mathcal{L}_{\text{unc}} + \lambda_5 \|z_t - h(z_{t-1}, \Delta t)\|^2$, where the final term enforces consistency between the observed and simulated twin state and acts as the re-synchronisation signal. Training proceeds in three stages: modality encoders and fusion are pre-trained with masked-modality reconstruction; the twin recurrence is fitted on longitudinal sequences; and the quantum block and heads are trained end-to-end with the encoders fine-tuned at a low learning rate. The full procedure is summarised in Table 2.

Table 2. Quantum Twin Learning procedure

Step	Operation	Output
1	Encode each available modality with f_m ; build presence masks and age features	$e_m, \delta_m, \Delta t_m$
2	Gated attention fusion over available modalities	Observation vector o_t
3	Time-aware recurrent update of twin state; forward simulation $h(z_t, \tau)$	z_t , simulated states
4	Classical encoder to n features; angle encoding; L variational layers; measure $\langle Z_i \rangle$	Quantum feature vector q_t
5	Trajectory, risk-horizon and uncertainty heads on q_t	Stage, score, time-to-event, σ
6	Joint loss $\mathcal{L}$; parameter-shift gradients for θ ; classical optimiser step	Updated parameters
7	On new observation: re-run steps 1–5, compare with $h(z_{t-1}, \Delta t)$, re-synchronise	Refreshed forecast and attribution

3.5 Explainability and Privacy

Attribution is computed at two levels. Modality-level attribution reads the fusion attention weights α_m and the change in prediction when a modality embedding is replaced by its placeholder. Time-level attribution perturbs past observations in the twin recurrence and records the effect on the forecast, producing a timeline the clinician can inspect. For privacy, the modality encoders and twin recurrence are trained with federated averaging across sites so that raw records never leave the hospital [21], and only encoder gradients and the small quantum parameter vector are exchanged.

4. Evaluation Protocol

The methodology is designed to be validated without hardware-dependent claims. Table 3 specifies the datasets, baselines and metrics. Public longitudinal cohorts provide imaging, cognitive, genetic and, for Parkinson's disease, wearable data with multi-year follow-up. Splits are patient-level with five-fold cross-validation, and a missing-modality stress test removes each modality at inference to measure graceful degradation. Baselines include a tabular gradient-boosted model, a multimodal Transformer without twin state, the twin with a classical dense layer replacing the quantum block, and the quantum block on a state-vector simulator versus a noisy device simulation. The primary comparison is the twin with and without the quantum block under identical parameter budgets, which isolates the contribution of the quantum representation.

Table 3. Datasets, baselines and evaluation metrics

Component	Specification
Datasets	ADNI (Alzheimer's; MRI, PET, cognitive scores, APOE) and PPMI (Parkinson's; DAT imaging, clinical scales, sensor data)
Targets	Stage at 6–36 months; CDR-SB / MDS-UPDRS score; time to MCI-to-dementia conversion
Baselines	Gradient boosting on tabular features; multimodal Transformer; twin + dense layer; twin + quantum block (simulator, noisy)
Metrics	AUROC, balanced accuracy, MAE of score, concordance index, expected calibration error, Brier score
Stress tests	Each modality removed at inference; reduced follow-up length; cross-cohort transfer
Efficiency	Trainable parameters, circuit depth, inference latency per twin update

Discussions

QTL differs from prior work in three ways. It treats the twin state, not the raw scan, as the object of learning, so temporal context and missing modalities are handled before the quantum block sees the data. It restricts the quantum footprint to a shallow circuit that is realistic on present hardware while keeping the rest of the system classical. And it commits to outputs a clinician can act on: a trajectory with a horizon and calibrated uncertainty, with modality- and time-level attribution. Expected advantages are earlier prediction, per-patient personalisation, parameter efficiency and continuous monitoring; these remain hypotheses that the protocol in Section 4 is designed to test. Limitations include dependence on simulator fidelity for circuits beyond ten qubits, the cost of refitting twins continuously at scale, and the need for shared multimodal benchmarks to compare quantum and classical claims fairly.

Conclusions

This paper proposed Quantum Twin Learning, a methodology for early, personalised prediction of neurodegenerative progression that couples a multimodal patient digital twin with a hybrid quantum-classical learner. A twin-state formulation tolerant to missing modalities, a compact variational block operating on that state, a joint trajectory, risk-horizon and uncertainty objective, and a re-synchronisation loop were specified, together with explainability and federated privacy mechanisms and a hardware-agnostic evaluation protocol on public cohorts. Neurodegeneration is a process rather than an event, so prediction must also be a process; QTL operationalises that principle. Future work will implement the protocol, extend the twin to ALS and Huntington's cohorts, and explore quantum federated learning and real-time bedside synchronisation.

References

- [1] World Health Organization, "Global status report on the public health response to dementia," Geneva, 2021.
- [2] G. Livingston et al., "Dementia prevention, intervention, and care: 2020 report of the Lancet Commission," *The Lancet*, vol. 396, pp. 413–446, 2020. DOI: 10.1016/S0140-6736(20)30367-6
- [3] C. R. Jack Jr. et al., "Tracking pathophysiological processes in Alzheimer's disease: an updated hypothetical model of dynamic biomarkers," *The Lancet Neurology*, vol. 12, no. 2, pp. 207–216, 2013. DOI: 10.1016/S1474-4422(12)70291-0
- [4] R. A. Sperling et al., "Toward defining the preclinical stages of Alzheimer's disease," *Alzheimer's & Dementia*, vol. 7, no. 3, pp. 280–292, 2011. DOI: 10.1016/j.jalz.2011.03.003
- [5] S.-C. Huang et al., "Fusion of medical imaging and electronic health records using deep learning: a systematic review and implementation guidelines," *npj Digital Medicine*, vol. 3, art. 136, 2020. DOI: 10.1038/s41746-020-00341-z
- [6] J. N. Acosta et al., "Multimodal biomedical AI," *Nature Medicine*, vol. 28, pp. 1773–1784, 2022. DOI: 10.1038/s41591-022-01981-2
- [7] J. Corral-Acero et al., "The 'Digital Twin' to enable the vision of precision cardiology," *European Heart Journal*, vol. 41, no. 48, pp. 4556–4564, 2020. DOI: 10.1093/eurheartj/ehaa159
- [8] T. Sun et al., "Digital twin in healthcare: Recent updates and challenges," *Digital Health*, vol. 9, 2023. DOI: 10.1177/20552076221149651
- [9] J. Biamonte et al., "Quantum machine learning," *Nature*, vol. 549, pp. 195–202, 2017. DOI: 10.1038/nature23474
- [10] M. Cerezo et al., "Variational quantum algorithms," *Nature Reviews Physics*, vol. 3, pp. 625–644, 2021. DOI: 10.1038/s42254-021-00348-9
- [11] A. Abbas et al., "The power of quantum neural networks," *Nature Computational Science*, vol. 1, pp. 403–409, 2021. DOI: 10.1038/s43588-021-00084-1
- [12] S. Rathore et al., "A review on neuroimaging-based classification studies for Alzheimer's disease and its prodromal stages," *NeuroImage*, vol. 155, pp. 530–548, 2017. DOI: 10.1016/j.neuroimage.2017.03.057
- [13] R. V. Marinescu et al., "The TADPOLE Challenge: Results after 1 year follow-up," *Journal of Machine Learning for Biomedical Imaging*, vol. 1, 2021.
- [14] G. Litjens et al., "A survey on deep learning in medical image analysis," *Medical Image Analysis*, vol. 42, pp. 60–88, 2017. DOI: 10.1016/j.media.2017.07.005
- [15] B. Shickel et al., "Deep EHR: A survey of recent advances in deep learning techniques for electronic health record analysis," *IEEE Journal of Biomedical and Health Informatics*, vol. 22, no. 5, pp. 1589–1604, 2018. DOI: 10.1109/JBHI.2017.2767063
- [16] A. Vaswani et al., "Attention is all you need," *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [17] K. Bruynseels et al., "Digital twins in health care: Ethical implications of an emerging engineering paradigm," *Frontiers in Genetics*, vol. 9, art. 31, 2018. DOI: 10.3389/fgene.2018.00031
- [18] V. Havlíček et al., "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, pp. 209–212, 2019. DOI: 10.1038/s41586-019-0980-2
- [19] L. Wei et al., "Quantum machine learning in medical image analysis: A survey," *Neurocomputing*, vol. 525, pp. 42–53, 2023. DOI: 10.1016/j.neucom.2023.01.049
- [20] J. Preskill, "Quantum computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, 2018. DOI: 10.22331/q-2018-08-06-79
- [21] N. Rieke et al., "The future of digital health with federated learning," *npj Digital Medicine*, vol. 3, art. 119, 2020. DOI: 10.1038/s41746-020-00323-1