ADVANCING LEARNING ANALYTICS THROUGH QUANTUM-ENHANCED MACHINE LEARNING: A FRAMEWORK FOR DROPOUT PREDICTION AND INTERVENTION

Muhammad Amin Nadim[1]
Pegaso University, Italy
muhammad.nadim@unipegaso.it

Emanuele Marsico
Pegaso University, Italy
emanuele.marsico@unipegaso.it

Double Blind Peer Review

Citation
Nadim, M. A., & Marsico, E. (2026). Advancing learning analytics through quantum-enhanced machine learning: A framework for dropout prediction and intervention. *Italian Journal of Health Education, Sport and Inclusive Didactics, 10*(4).

Doi: https://doi.org/10.32043/gsd.v10i4.1833

Copyright notice:
© 2023 this is an open access, peer-reviewed article published by Open Journal System and distributed under the terms of the Creative Commons Attribution 4.0 International, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

gsdjournal.it
ISSN: 3103-5744

ABSTRACT
Student dropout remains a persistent challenge in higher education, where timely intervention under limited data is critical. Classical machine learning models underpin most early-warning systems, but their effectiveness declines when behavioural interactions are complex and labelled data are scarce. This study examines Quantum Machine Learning (QML) for dropout prediction, grounded in Tinto's integration model and Zimmerman's self-regulated learning framework. Using OULAD, classical models are benchmarked across three early-semester windows, while QSVMs and hybrid stacking are benchmarked at a fixed Week 2 window across low-sample regimes. XGBoost achieves strong performance as data accumulate (PR-AUC: 0.496–0.592). These exploratory results, obtained on simulated quantum backends and on a single dataset, suggest that QML may act as a regime-specific complement to classical learning analytics rather than as a general-purpose improvement.

KEYWORDS
LA; student dropout; quantum ML; quantum SVM; hybrid models.

Received 2026-07-01
Accepted 2026-10-27
Published 2026-10-30

[1] Muhammad Amin Nadim (sections 1, 4, 5, 6)
Emanuele Marsico (sections 2, 3)

1. Introduction

Student dropout remains a major challenge for higher education systems worldwide, with substantial academic, social, and economic consequences for students, institutions, and policymakers (Tinto, 2012; Sghir et al., 2023). Despite advances in learning analytics, early identification of at-risk students remains difficult due to the dynamic interplay of behavioural engagement, academic performance, and socio-demographic factors (Richardson et al., 2012; Castro-Montoya et al., 2025). Established theories, including Tinto's (2012) academic and social integration model and Zimmerman's (2000) self-regulated learning theory, emphasize that disengagement arises through cumulative, nonlinear processes — not through single abrupt failures — complicating predictive modelling from both a pedagogical and computational standpoint. Classical machine learning approaches such as logistic regression, random forests, and gradient boosting are widely used in early-warning systems for student retention (Sghir et al., 2023). However, empirical evidence indicates limited generalizability across cohorts and reduced capacity to capture high-order interactions in high-dimensional educational data, particularly under conditions of limited labelled data, privacy constraints, and early-semester prediction requirements (Jha et al., 2019). Furthermore, learning analytics systems that incorporate demographic variables without fairness analysis risk amplifying pre-existing structural inequalities (Baker & Hawn, 2022). Quantum Machine Learning (QML) has recently been proposed as a potential alternative for addressing these computational challenges. By exploiting quantum mechanical principles such as superposition and entanglement, QML models can encode complex feature interactions through compact quantum embeddings that are difficult to approximate classically (Biamonte et al., 2017). Quantum kernel methods offer a theoretically

grounded mechanism for representing decision boundaries beyond those accessible to classical kernels of comparable complexity (Schuld & Killoran, 2019). Empirical studies from non-educational domains suggest potential advantages in low-sample, high-dimensional regimes like those observed in educational data (Gujju et al., 2024). Nevertheless, applications of QML in learning analytics remain limited, and its practical value has not been systematically established on real-world educational datasets. To address this gap, a quantum-enhanced learning analytics framework for student dropout prediction is evaluated, integrating classical and quantum models within a unified, psycho-pedagogically grounded pipeline. Quantum Support Vector Machines (QSVMs) are benchmarked against strong classical baselines — Logistic Regression, Random Forests, and XGBoost — under early-semester and low-sample prediction scenarios. Results indicate that the contribution of QML is regime-dependent: classical models dominate at very small sample sizes, QSVMs perform competitively in moderately sized low-sample regimes, and hybrid stacking achieves the strongest performance by exploiting complementary representations. Analysis of fusion weights further clarifies when quantum models become reliable contributors. Computational cost indicators — training time and circuit depth — are reported to contextualize predictive performance within realistic operational constraints. This analysis provides an early systematic evaluation of quantum machine learning for student dropout prediction on real-world educational data and establishes an empirically grounded basis for assessing the feasibility and limitations of quantum-enhanced learning analytics.

2. Related Work and Theoretical Framework

This section reviews the theoretical and methodological foundations underlying student dropout prediction in higher education. We first situate dropout modelling within established psycho-pedagogical theories, then examine classical machine learning approaches and their documented limitations, review emerging work on Quantum Machine Learning and hybrid architectures, and identify unresolved gaps in the literature that motivate the proposed framework.

2.1 Psycho-Pedagogical Perspectives on Student Dropout

Student dropout is a multifactorial phenomenon driven by interacting academic, behavioural, social, and institutional factors (Tinto, 2012; Piepenburg & Beckmann, 2022). Psycho-pedagogical theories provide the conceptual foundation for interpreting digital trace data from learning environments and for ensuring that predictive features are educationally meaningful rather than merely statistical artifacts.

2.2 Tinto's Model of Academic and Social Integration

The theory of academic and social integration proposed by Tinto (2012) constitutes the most influential theoretical pillar in the study of student persistence. Tinto's longitudinal model posits that the decision to abandon studies is the result of a cumulative process mediated by the degree of academic integration — encompassing academic performance, faculty interaction, and perceived competence — and social integration, including peer relationships, participation in institutional life, and sense of community belonging (Tinto, 2012). Students whose integration in both dimensions is low

are significantly more likely to abandon their studies, particularly in the first year of enrolment. Recent empirical work has validated this model across diverse national contexts. A longitudinal structural equation modelling study conducted in Colombia with 563 first-year students demonstrated that psychosocial and family-related variables significantly predict initial institutional commitment, which in turn mediates academic and social integration and reduces dropout risk (Castro-Montoya et al., 2025). Similarly, a German factorial survey study confirmed that all subdimensions of both integration types predict dropout intention, with differential weights depending on students' academic family background (Piepenburg & Beckmann, 2022). A systematic review published in Frontiers in Education further confirmed that first-year seminar experiences function as a primary site of integration and retention (Samoila & Vrabie, 2023). In online learning contexts — such as that represented by the OULAD dataset — social integration is largely mediated through asynchronous digital interactions, making VLE behavioural data the primary observable behavioural indicator theoretically associated with institutional engagement (Romdhoni et al., 2026). The transition from residential to online learning environments, however, introduces a structural asymmetry that Tinto's original model did not anticipate. In campus-based settings, social integration is built through face-to-face encounters — shared physical spaces, informal peer interactions, and synchronous faculty contact — all of which generate continuous, low-effort signals of belonging. In distance education, these signals must be actively constructed by the student through deliberate engagement with digital course materials, discussion forums, and assessment activities (Romdhoni et al., 2026). This places a considerably higher self-regulatory burden on the learner and means that the observable behavioural traces left in VLE logs carry dual theoretical weight: they simultaneously index academic integration — the student's engagement with

course content and evaluation structures — and serve as the only available behavioural indicator theoretically associated with the social integration dimension that, in Tinto's (2012) framework, is equally predictive of persistence. This reinterpretation has a direct methodological consequence for the present study. Feature categories derived from VLE interaction data — including interaction volume, active learning days, and assessment submission regularity — are not merely behavioural covariates but theoretically grounded operationalizations of both integration dimensions as they manifest in a distance learning context. The temporal evolution of these features across the first weeks of a course therefore constitutes a longitudinal trace of the student's integration trajectory, consistent with Tinto's (2012) emphasis on the dynamic, processual nature of academic persistence. Students who show declining engagement in the early semester are not simply interacting less with a platform; from a Tinto perspective, they are exhibiting the behavioural signature of eroding academic and social embeddedness — the mechanism that, left unaddressed, culminates in formal withdrawal (Tinto, 2012; Castro-Montoya et al., 2025). This theoretical grounding also clarifies the pedagogical purpose of early-warning systems built on such data. Identifying a student as at-risk on the basis of declining VLE engagement is, within the Tinto framework, tantamount to detecting a student whose integration is deteriorating before it becomes irreversible. The implication for institutional intervention is correspondingly specific: actions that restore the student's sense of academic competence and social belonging — targeted academic advising, peer connection programmes, or proactive faculty outreach — are theoretically more likely to reverse the dropout trajectory than generic administrative alerts (Samoila & Vrabie, 2023; Piepenburg & Beckmann, 2022).

2.3 Self-Regulated Learning Theory and Motivation

Self-regulated learning (SRL) frameworks conceptualize dropout not as a sudden academic failure but as the endpoint of a progressive decline in cognitive, metacognitive, motivational, and behavioural self-regulation (Zimmerman, 2000; Pintrich & De Groot, 1990). Zimmerman's (2000) cyclical model — articulated across forethought, performance, and self-reflection phases — identifies academic self-efficacy as the central motivational construct: students with high perceived self-efficacy set more ambitious goals, persist in the face of academic obstacles, and deploy more effective cognitive strategies, while low self-efficacy is associated with procrastination, avoidance, and progressive disengagement. A comprehensive meta-analysis published in Psychological Bulletin (Richardson et al., 2012), synthesizing 241 datasets across 13 years of research, identified academic self-efficacy, effort regulation, and goal-setting as the three non-intellective constructs most strongly correlated with university GPA, surpassing demographic predictors. Pintrich and De Groot's (1990) foundational study in the Journal of Educational Psychology demonstrated that metacognitive regulation, self-efficacy, and performance anxiety are the strongest within-student predictors of academic performance independently of prior achievement. The distinction between intrinsic goal orientation — studying for mastery and interest — and extrinsic goal orientation is particularly relevant: students with predominantly extrinsic motivation display higher dropout vulnerability in the presence of academic setbacks (Pintrich & De Groot, 1990). Eccles and Wigfield (2020), in elaborating the Situated Expectancy-Value Theory (SEVT) published in Contemporary Educational Psychology, demonstrated that the subjective value attributed to academic tasks is deeply situated in social and institutional contexts; in digital learning environments, the perceived relevance of VLE activities significantly

moderates behavioural engagement, with task value positively predicting LMS resource utilization and final examination performance (von Keyserlingk et al., 2026). These theoretical insights carry direct implications for the design of predictive features in learning analytics systems. If dropout is the observable outcome of deteriorating self-regulation, then behavioural indicators retrievable from VLE logs — such as declining interaction frequency, irregular session timing, or delayed assessment submissions — can be interpreted as early digital manifestations of the self-reflection failures described in Zimmerman's (2000) cyclical model. In this sense, machine learning models trained on VLE behavioural data are not merely statistical classifiers but indirect instruments for detecting disruptions in students' self-regulatory cycles. This interpretive framing is consequential: it shifts the purpose of early-warning systems from prediction-for-prediction's-sake toward pedagogically actionable identification of the specific self-regulatory phase at which a student's engagement begins to deteriorate. The motivational dimension elaborated by Eccles and Wigfield (2020) further refines this picture. In distance learning environments — where institutional social cues are attenuated, and students must sustain engagement autonomously — the subjective task value component of SEVT takes on heightened predictive importance. Students who perceive VLE activities as instrumentally or intrinsically valuable are more likely to engage consistently and submit assessments on time, whereas those who attribute low utility to digital course interactions are at elevated risk of progressive disengagement even before formal academic underperformance becomes visible (Eccles & Wigfield, 2020; von Keyserlingk et al., 2026). This suggests that the temporal regularity of VLE access — a feature prominently used in OULAD-based prediction models — may function as behavioural indicator theoretically associated not only with effort regulation (Pintrich & De Groot, 1990) but also for perceived task value, thereby capturing motivational

variance that static demographic predictors systematically fail to represent (Richardson et al., 2012). Taken together, the SRL and expectancy-value frameworks provide a theoretically coherent basis for treating early-semester behavioural engagement data as meaningful pedagogical signals rather than raw clickstream statistics. They also establish a normative standard for early-warning system design: predictive models should be evaluated not only on statistical performance metrics but on their capacity to generate actionable pedagogical insight — identifying which dimension of self-regulation (effort, planning, metacognitive monitoring, or task valuation) is most implicated in a given student's risk profile, so that intervention strategies can be appropriately differentiated (Zimmerman, 2000; Pintrich & De Groot, 1990). An important theoretical complement to the SRL and expectancy-value tradition concerns the construct of academic buoyancy — students' capacity to successfully manage the routine setbacks, challenges, and motivational dips that characterize everyday academic life, as distinct from major or chronic adversities (Martin & Marsh, 2008). Introduced in the Journal of School Psychology, academic buoyancy is operationalized through five motivational predictors — referred to as the "5Cs": confidence (self-efficacy), coordination (planning), commitment (persistence), composure (low anxiety), and control — each of which maps directly onto observable behavioural patterns within digital learning environments (Martin & Marsh, 2008; Martin et al., 2010). A student with low academic buoyancy is not necessarily performing poorly at any given moment; rather, their motivational architecture is insufficiently robust to absorb the minor but cumulative disruptions — a missed deadline, a disappointing quiz result, a period of low motivation — that collectively precipitate progressive disengagement. Crucially, this construct suggests that early-semester behavioural irregularities in VLE logs may not signal irrecoverable academic failure but rather a temporary deficit in motivational

resilience that is, in principle, addressable through targeted early intervention (Martin et al., 2010). A complementary perspective on the motivational antecedents of disengagement is offered by the study burnout literature. Salmela-Aro and her colleagues have documented, across large-scale longitudinal studies in Finnish higher education, that study burnout — characterized by exhaustion in studying, cynicism toward the meaningfulness of academic work, and a sense of inadequacy as a student — constitutes a distinct psychological pathway to academic withdrawal that unfolds progressively over semesters and is detectable before formal dropout occurs (Salmela-Aro & Upadyaya, 2014). Using the demands-resources model applied to the school context, Salmela-Aro and Upadyaya (2014), publishing in the British Journal of Educational Psychology, demonstrated that study demands directly predicted subsequent burnout, while self-efficacy was positively associated with engagement and negatively associated with burnout — a finding that links the burnout trajectory directly to Zimmerman's (2000) self-efficacy construct and extends it into a temporally sensitive, multidimensional framework. Profile-analytic studies within this tradition have identified a characteristic "engaged-exhausted" subgroup of students — those who remain behaviourally engaged in the early semester but whose motivational resources are progressively depleted — as particularly vulnerable to late-semester withdrawal and as the hardest to identify using standard academic performance metrics alone (Salmela-Aro & Read, 2017). Finally, the meta-analytic evidence assembled by Credé and Kuncel (2008), published in Perspectives on Psychological Science, provides a crucial empirical bridge between the theoretical constructs reviewed above and the practical question of which student characteristics are most amenable to prediction via behavioural indicator theoretically associated variables. Synthesizing 344 datasets (N = 72,431), the authors demonstrated that study motivation and

study skills — both closely related to SRL constructs — exhibit the strongest relationships with both GPA and grades in individual courses, exceeding the predictive validity of standardized admissions tests and prior grades for within-university performance (Credé & Kuncel, 2008). Notably, the authors identified study motivation as a unique predictor that is largely independent of cognitive ability, consistent with the SRL literature's emphasis on effort regulation and goal-setting as malleable, context-sensitive constructs (Richardson et al., 2012). This finding has a direct methodological implication for the present study: since study motivation and study habits are not redundant with demographic or prior-achievement variables, machine learning models that rely exclusively on behavioural interaction data may be capturing motivational variance — in the form of engagement regularity, submission consistency, and access intensity — that is both theoretically meaningful and incrementally predictive beyond what static background features can provide.

2.4 Sociodemographic Factors and Algorithmic Fairness

Dropout risk is not uniformly distributed across the student population. Variables including disability status, geographic origin, family educational background, and age band are associated with differentiated vulnerability trajectories that intersect with motivational and self-regulatory dynamics (Li & Carroll, 2020). A study published in the Journal of Higher Education Policy and Management documented that students from disadvantaged equity groups show significantly higher probabilities of dropout, with health and financial factors as primary determinants (Li & Carroll, 2020). This evidence raises a fundamental psycho-pedagogical concern for automated predictive systems: the inclusion of demographic variables in machine learning models without accompanying fairness analysis risks amplifying pre-existing structural

inequalities (Baker & Hawn, 2022). Baker and Hawn's (2022) review published in the International Journal of Artificial Intelligence in Education documented systematic evidence that educational algorithms trained on demographic data may disadvantage already marginalized groups. Responsible design of early-warning systems therefore requires evaluating predictive performance separately across demographic subgroups.

2.5 Behavioural Digital Indicators as Pedagogical Proxies

VLE interaction data — interaction volume, session duration, temporal access patterns, and assessment submission sequences — are observable proxies for latent psychological processes including effort regulation, academic planning, and institutional integration (Romdhoni et al., 2026). An analysis of 32,593 students across OULAD demonstrated that VLE behavioural variables account for 85.0% of predictive power in early-warning models, with Activity Span — the temporal continuity of engagement — as the dominant predictor (41.3%), confirming that regularity of digital academic behaviour constitutes a robust, cross-module generalizable signal of self-regulation (Romdhoni et al., 2026). This convergence between psycho-pedagogical theory and computational data justifies the theory-driven feature engineering approach adopted in this study.

2.6 Classical Machine Learning for Dropout Prediction

Classical machine learning models dominate dropout prediction research. A decade-long systematic review published in Education and Information Technologies (Sghir et al., 2023) identified gradient boosting approaches — particularly XGBoost — as the dominant high-performance framework across 134 studies from 2012 to 2022, with VLE interaction data consistently yielding the strongest predictive signals. A systematic review of OULAD benchmark

studies confirmed AUC-ROC values up to 0.93 for dropout prediction using behavioural features alone, with gradient boosting demonstrating consistent superiority (Jha et al., 2019). However, classical models typically require substantial labelled data, may overfit context-specific patterns, and remain constrained by predefined feature representations, limiting early-semester effectiveness precisely when interventions are most impactful (Sghir et al., 2023).

2.7 Quantum Machine Learning and Quantum Kernel Methods

QML enables alternative data representations by encoding classical features into quantum states, allowing exploration of high-dimensional Hilbert spaces through superposition and entanglement (Biamonte et al., 2017). Biamonte et al. (2017), in a landmark review published in Nature, articulated the theoretical potential of quantum algorithms for pattern recognition in feature spaces inaccessible to classical computation. Schuld and Killoran (2019), in a paper published in Physical Review Letters, demonstrated that quantum feature maps in Hilbert space can represent decision boundaries inaccessible to classical kernels of comparable complexity. A comprehensive review of quantum machine learning from NISQ-era devices (Gujju et al., 2024) confirmed that quantum advantages are most reliably demonstrated in data-scarce settings with complex feature interactions — precisely the conditions characterizing early-semester educational prediction. However, reported advantages remain largely demonstrated on non-educational benchmarks, leaving their generalizability to real-world educational data unresolved.

2.8 Hybrid Quantum–Classical Learning Architectures

An emerging consensus in the QML literature is that near-term quantum systems are unlikely to replace classical machine learning; instead, hybrid quantum–classical architectures represent the most viable path to practical impact (Biamonte et al., 2017; Gujju et al., 2024). Stacking is particularly principled: it allows a meta-learner to assess complementarity dynamically without imposing a fixed fusion rule, and the resulting learned weights constitute interpretable evidence of regime-specific quantum contribution — essential in educational settings where explainability is an ethical requirement (Baker & Hawn, 2022).

2.9 Research Gap and Positioning of This Study

Prior work has not systematically examined when quantum models outperform classical baselines on real educational data, whether quantum and classical representations are complementary, or how performance gains trade off against computational cost on near-term simulators. Moreover, psycho-pedagogical grounding of feature construction has received insufficient attention, with most studies treating behavioural features as atheoretical proxies. This study addresses these gaps through theory-driven feature construction anchored in Tinto (2012) and Zimmerman (2000), systematic benchmarking of QSVMs against strong classical models across early-semester windows, and a stacked hybrid architecture to empirically test quantum–classical complementarity.

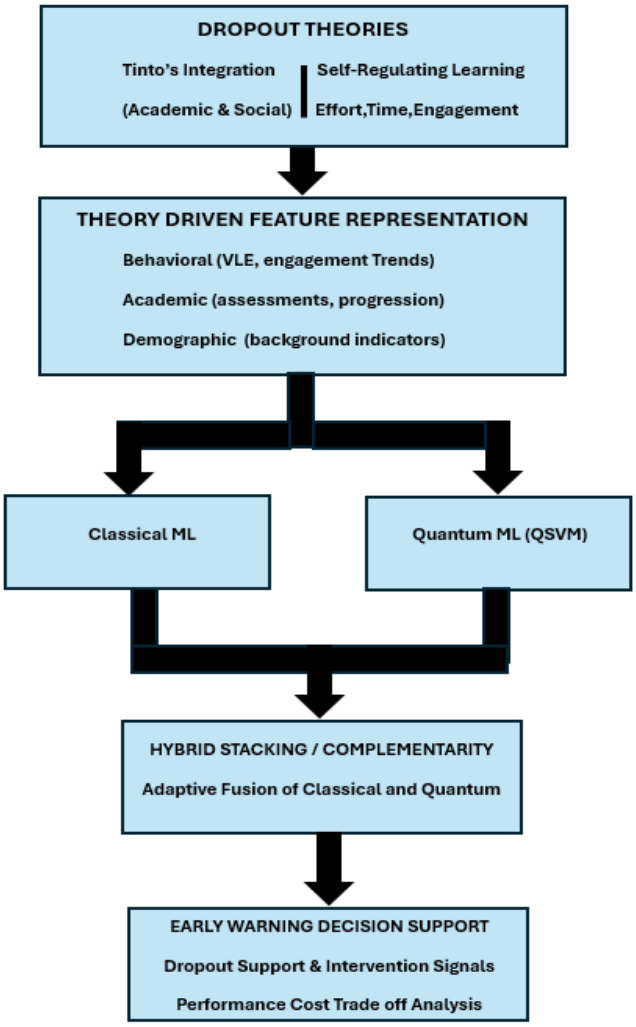


Fig. 1. Theory-driven quantum–classical framework for student dropout prediction and early-warning decision support.

3. Methodology

A theory-informed and temporally grounded methodology is adopted, covering dataset preparation, feature engineering, preprocessing, and experimental evaluation, with emphasis on reproducibility and fair model comparison.

3.1 Dataset

The study utilizes the Open University Learning Analytics Dataset (OULAD), a large-scale, publicly available dataset containing longitudinal records of student behaviour, assessment activity, and demographic characteristics in distance higher education. OULAD comprises 32,593 student–module-presentation registrations across 22 module presentations of seven modules, corresponding to 28,785 unique students; approximately X% of students are registered in more than one presentation. The unit of analysis in this study is the student–presentation registration: each row represents one student's participation in one module presentation, and a student enrolled in two presentations contributes two independent rows. Dropout is operationalised as withdrawal from the presentation prior to completion, identified by a non-missing date_unregistration and the binary target distinguishes withdrawn registrations from those completed or failed, yielding a positive-class prevalence of 31.16% in the analysed sample (Romdhoni et al., 2026; Jha et al., 2019). Dropout is operationalized as withdrawal from a course prior to completion (Sghir et al., 2023). The binary target variable distinguishes withdrawn students from those who completed or failed, yielding a class imbalance of 31.16% withdrawn students.

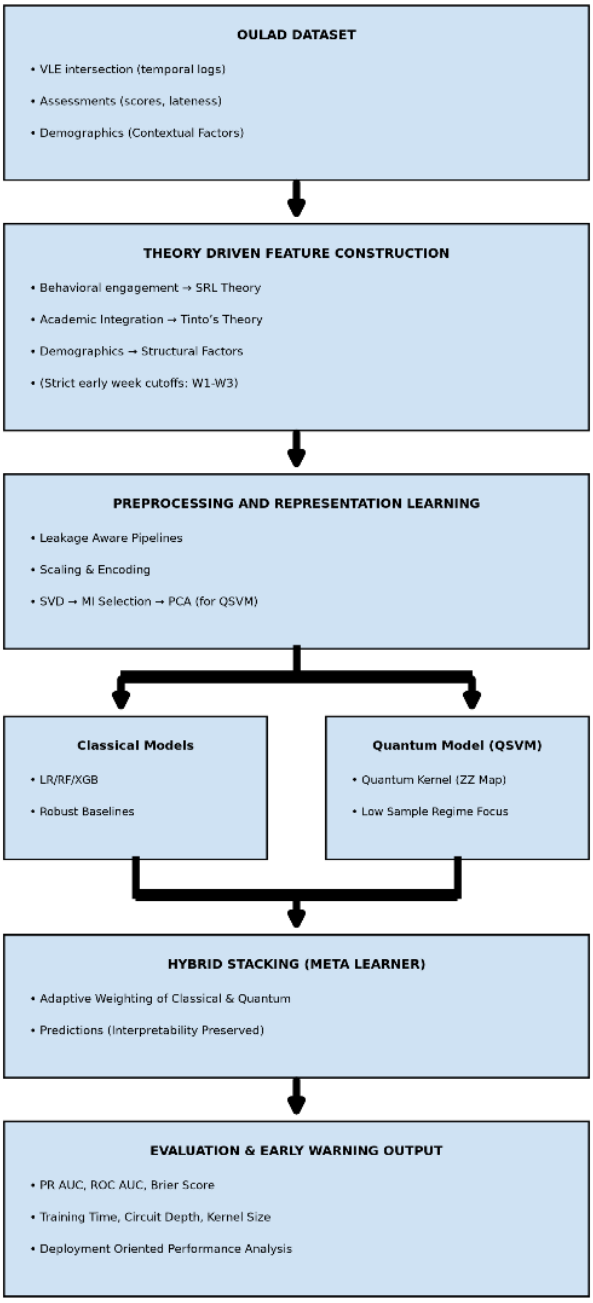


Fig. 2. End-to-end theory-driven quantum–classical pipeline for early student dropout prediction.

3.2 Theory-Driven Feature Construction

Feature construction is *informed by* Tinto's (2012) integration model and Zimmerman's (2000) SRL framework, which guided the selection of which behavioural traces to aggregate. It is important to state the epistemic status of these features: they are observable behavioural indicators derived from VLE and assessment logs, and they do not measure self-regulation, motivation or institutional integration directly. No self-report or validated instrument was available in OULAD to establish construct validity, and the theoretical labels below should be read as design rationale rather than as measurement claims. Three feature categories are constructed:

- Engagement-behaviour features (theoretically motivated by the effort-regulation and planning dimensions of SRL, but not validated against SRL measures): aggregated early-semester VLE metrics including interaction volume, active learning days, and normalized engagement intensity, capturing the effort regulation and temporal planning dimensions of Zimmerman's (2000) SRL model.

- Assessment-behaviour features (theoretically motivated by Tinto's academic-integration dimension but not validated as a measure of it): assessment-related indicators including submission counts, lateness measures, and mean scores, reflecting performance quality and course structure compliance (Tinto, 2012).

- Sociodemographic features: age band, prior education level, geographic region, and disability status, included to account for structural influences on persistence (Li & Carroll, 2020), while ensuring they do not dominate behavioural signals.

All features are constructed using strict temporal cutoffs to ensure that only information available up to a given prediction week is used, preventing temporal data leakage.

3.3 Temporal Framing and Data Splitting

Classical baselines are evaluated using data available up to Weeks 1, 2, and 3 of each course presentation, enabling analysis of how predictive performance evolves as behavioural evidence accumulates (Sghir et al., 2023). Quantum and hybrid models, owing to the computational cost of kernel evaluation on simulated backends, are evaluated at a single temporal window (Week 2) across increasing training-set sizes; their behaviour in Weeks 1 and 3 was not assessed and remains an open question. Evaluation employs group-aware train–test partitioning at the course presentation level to prevent information leakage across cohorts. Two distinct experimental conditions are used. In the *full-sample condition* (Section 4.1), all models are trained on the complete, imbalanced training partition (31.16% positive class) with class-weighted losses. In the *balanced low-sample condition* (Sections 4.2–4.3), all models — classical, quantum and hybrid alike — are trained on class-balanced subsamples of 50, 100, 200 and 300 instances drawn from the same training partition, so that classical and quantum learners are compared under identical data conditions. Test sets remain fixed and unmodified in both conditions.

3.4 Preprocessing and Evaluation Protocol

All preprocessing is implemented within leakage-aware pipelines. Numerical features are z-score standardized; categorical variables are one-hot encoded. For quantum models, supervised dimensionality reduction is applied via truncated SVD, mutual information–based feature selection, and PCA.

Precision–Recall AUC (PR-AUC) is the primary metric, given its superiority over ROC-AUC under class imbalance (Sghir et al., 2023), complemented by ROC-AUC and Brier scores. Computational cost is reported as training time, circuit depth, and kernel matrix size (Gujju et al., 2024).

3.5 Classical and Quantum Modelling Approaches

Classical baselines — Logistic Regression, Random Forests, and XGBoost — represent the dominant approaches in the dropout prediction literature (Sghir et al., 2023; Jha et al., 2019). All classical models are trained using class-weighted loss functions and conservatively selected hyperparameters. Quantum models are implemented as Quantum Support Vector Machines (QSVMs) using a ZZFeatureMap to introduce entanglement and nonlinear interactions, with the quantum kernel defined as the fidelity between encoded quantum states — consistent with the theoretical framework of Schuld and Killoran (2019). Configuration: n_qubits = 4, feature map reps = 1, entanglement = linear. Training subset sizes are progressively increased from 50 to 300 samples.

3.6 Hybrid Quantum–Classical Stacking

A stacked hybrid meta-learner architecture combines independently trained classical and quantum predictors via a regularized logistic regression meta-learner trained on held-out predictions. This enables adaptive weighting, provides interpretable meta-coefficients quantifying each model's contribution, and avoids imposing a fixed fusion rule across sample-size regimes. The learned fusion weights indicate how the meta-learner allocates reliance between base models in each regime. They are not a direct measure of representational complementarity: meta-coefficients in a regularised logistic

meta-learner depend on the scale and calibration of base-model outputs, on the collinearity between their predictions, and on the regularisation strength. Any interpretation of these weights in terms of distinct learned representations is therefore hypothetical and would require direct tests, such as the correlation and disagreement structure of base-model predictions on held-out cases, or kernel-alignment analyses between the quantum and classical representations. Such tests were not performed in this study.

3.7 Computational Considerations

All quantum experiments are conducted on simulated backends to reflect current NISQ-era hardware limitations (Gujju et al., 2024). Computational cost is explicitly co-reported with predictive performance, consistent with responsible benchmarking practice.

4. Results

4.1 Classical Baselines (Weeks 1–3)

Classical performance improves monotonically as behavioural evidence accumulates — consistent with Zimmerman's (2000) SRL framework, which predicts growing observability of self-regulatory processes over time. Using PR-AUC as the primary metric (class imbalance: 31.16% withdrawn), XGBoost is the strongest classical model: 0.496 (Week 1), 0.533 (Week 2), 0.592 (Week 3); ROC-AUC: 0.694, 0.711, 0.752. Logistic Regression: 0.476 → 0.508 → 0.559. Random Forest: 0.467 → 0.505 → 0.571. Brier score (XGBoost): 0.173 → 0.167 → 0.156, indicating improving calibration. Training time: XGBoost 2.11–2.75 s, Logistic Regression 0.18–0.58 s, Random Forest ~19–24 s. These results are consistent with OULAD benchmark studies (Jha et al., 2019).

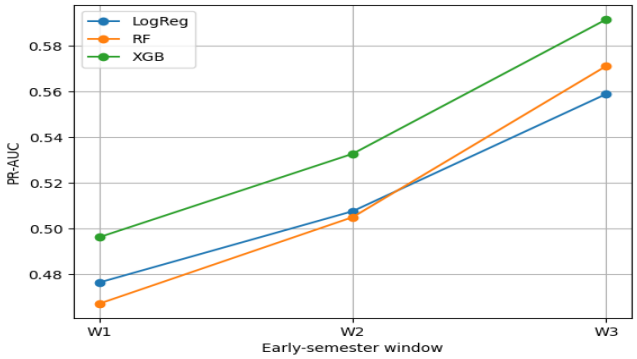


Fig. 3 Classical baselines across early-semester windows (OULAD)

4.2 QSVM Performance under Balanced Low-Sample Regimes

All values reported in this and the following subsection refer to the balanced low-sample condition and are therefore not directly comparable with those of §4.1. Because PR-AUC has a chance-level baseline equal to the positive-class prevalence, the shift from an imbalanced training distribution (31.16% positive) to balanced subsamples (~50% positive) mechanically raises the achievable PR-AUC range. The XGBoost Week 2 values of 0.533 (§4.1) and 0.651–0.660 (this section) thus reflect two different evaluation settings rather than an inconsistency; within-condition comparisons across models are the only valid comparisons. QSVM PR-AUC increases with training size: 0.590 (50 samples), 0.616 (100), 0.663 (200), 0.683 (300). QSVM underperforms XGBoost at minimal samples (0.590 vs. 0.660 at 50) but exceeds XGBoost from 200 samples onward (0.663 vs. 0.656 at 200; 0.683 vs. 0.660 at 300), indicating a regime-dependent performance crossover. This pattern is theoretically consistent with Schuld and Killoran (2019): quantum feature maps encode high-dimensional interaction structures that become reliably exploitable once sufficient samples stabilize the kernel matrix.

4.3 Hybrid Stacking and Complementarity Analysis

Under the balanced low-sample condition, the hybrid meta-learner attains PR-AUC = 0.681 at 200 samples and 0.682 at 300 samples. At 200 samples the hybrid model outperforms both QSVM (0.663) and XGBoost (0.656). At 300 samples, however, QSVM attains the nominally highest value (0.683 vs. 0.682 for the hybrid model and 0.660 for XGBoost); a difference of 0.001 PR-AUC is not interpretable as a performance ranking, and the two models are best described as empirically equivalent in this regime. Overall, the hybrid architecture is therefore characterised not as uniformly superior, but as the configuration that matches the best single model across regimes without requiring prior knowledge of which model will dominate — a robustness property rather than a performance advantage. Learned meta-coefficients reveal regime-dependent complementarity: QSVM weight increases from 0.328 (50 samples) to 0.750 (100). This shift describes how the meta-learner reallocates reliance as sample size grows; interpreting it as a transition between learning regimes remains a hypothesis not tested here and consistent with theoretical predictions from Biamonte et al. (2017) and Schuld and Killoran (2019).

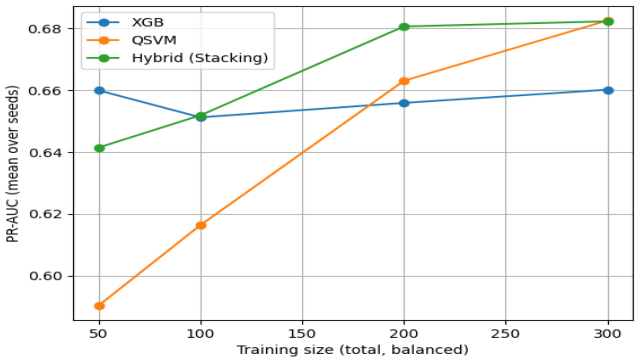


Fig. 4. Sample-efficiency comparison under the balanced low-sample condition (Week 2; all models trained on class-balanced subsamples)

4.4 Computational Cost

QSVM runtime scales sharply: 39.7 s (50 samples), 91.2 s (100), 219.1 s (200), 385.2 s (300) — reflecting quadratic kernel-matrix evaluation costs. XGBoost requires 2.11–2.75 s across all conditions. These results are consistent with NISQ-era scalability limitations (Gujju et al., 2024) and motivate selective hybrid deployment.

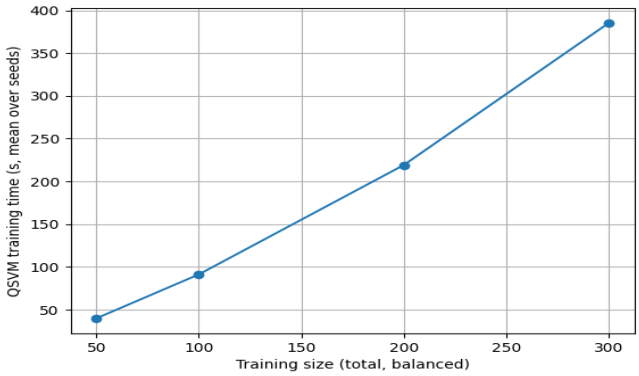


Fig. 5. QSVM computational cost scaling (simulated backend)

5. Discussion

The following interpretations should be read as hypotheses generated by an exploratory comparison, not as validated findings. All quantum results were obtained on classical simulations of quantum circuits; no physical quantum hardware was used.

Within the single temporal window and the sample-size range examined here, the relative performance of the quantum model varies with training-set size rather than being uniformly higher or lower than that of classical baselines. This pattern is compatible with theoretical arguments on the expressivity of quantum feature maps (Schuld & Killoran, 2019), but the present design does

not test that theory: the observed crossover is a single empirical observation on one dataset, one feature map and one simulated backend, and its replication is required before it can be regarded as a regularity. Classical models — particularly XGBoost — remain strong and computationally efficient baselines consistent with the decade-long systematic review by Sghir et al. (2023). From a psycho-pedagogical standpoint, the monotonic improvement of classical models from Week 1 to Week 3 is interpretable through Zimmerman's (2000) SRL framework: as students accumulate observable behavioural traces, the model's ability to capture underlying self-regulatory dynamics improves. Lower Week 1 performance is consistent with two non-mutually-exclusive explanations that the present design cannot disentangle: genuine information latency in early-semester behaviour, and the reduced availability of assessment-derived features in the first week. Distinguishing between them would require an ablation contrasting behavioural-only feature sets across the three windows, which we did not perform. The regime-dependent advantage of QSVMs in data-scarce conditions may reflect their ability to extract structural patterns from sparse behavioural data before classical models have access to sufficient signal (Schuld & Killoran, 2019). The shift in fusion weights is compatible with, but does not establish, the hypothesis that QSVMs and XGBoost learn partly non-redundant representations of student engagement. Testing this hypothesis directly — e.g. via prediction-disagreement or kernel-alignment analysis — is left to future work. This finding aligns with the consensus that near-term quantum systems are most valuable as complementary modules within hybrid architectures (Biamonte et al., 2017; Gujju et al., 2024). On algorithmic fairness, the inclusion of demographic variables (disability status, prior education, geographic region, age band) as predictive features necessitates caution. Baker and Hawn (2022) documented systematic evidence that machine learning systems trained on demographic

data may disadvantage already marginalized groups, consistent with the equity concerns raised by Li and Carroll (2020). Future implementations of this framework must include formal fairness audits disaggregated by demographic subgroup before any institutional deployment. A further discussion dimension concerns the operationalization of intervention: identifying at-risk students is a necessary but insufficient condition for reducing dropout. Early-warning systems must be embedded within broader psycho-pedagogical support frameworks — including academic advising, tutoring, and motivational support aligned with SRL principles (Zimmerman, 2000; Pintrich & De Groot, 1990) — to translate predictive accuracy into actual retention gains.

6. Conclusion

This study examined the role of QML in student dropout prediction using a theory-driven, leakage-aware, and psycho-pedagogically grounded framework applied to the OULAD dataset. Classical machine learning models — particularly XGBoost — remain robust and computationally efficient baselines whose accuracy improves progressively as early-semester behavioural data accumulate, consistent with self-regulated learning theory (Zimmerman, 2000). In the balanced low-sample condition examined, QSVMs achieved PR-AUC values comparable to or slightly above those of the strongest classical baseline at 200 and 300 training samples. Whether this reflects a distinctive encoding of feature interactions by the quantum kernel, differences in effective regularisation, or the specific dimensionality-reduction pipeline applied to the quantum inputs, cannot be determined from these data. Hybrid quantum–classical stacking models achieve the best overall performance by adaptively combining complementary representations, with learned fusion weights providing transparent, interpretable evidence of regime-dependent quantum

contribution. Within the limits of a single dataset, a single temporal window, one feature map and simulated backends, these results provide preliminary evidence that quantum kernel methods can reach performance comparable to strong classical baselines in balanced low-sample settings. They do not establish an advantage of QML for learning analytics, nor its value under operational conditions.

6.1 Limitations

First, all quantum experiments were conducted on simulated backends, which impose substantial computational overhead and may not reflect the behaviour or noise characteristics of near-term quantum hardware; validation on physical NISQ devices remains essential (Gujju et al., 2024). Second, the evaluation is limited to a single dataset from UK distance higher education (OULAD), limiting generalizability to traditional campus-based or culturally different institutions. Third, quantum feature maps were intentionally constrained (4 qubits, linear entanglement, 1 repetition), potentially underestimating the expressive capacity of more complex circuit architectures. Fourth, the study does not include a fairness analysis across demographic subgroups — a significant ethical limitation given the use of disability, age, and geographic region as predictive features (Baker & Hawn, 2022). Fifth, the predictive window is restricted to the first three weeks of the semester, leaving open questions about quantum advantages at later stages or in courses with atypical temporal structures. Sixth, Using OULAD, classical models are benchmarked across three early-semester windows, while QSVMs and hybrid stacking are benchmarked at a fixed Week 2 window across low-sample regimes. Sixth, quantum and hybrid models were evaluated at a single temporal window (Week 2), so the temporal generalisability of the quantum results is untested. Seventh, the behavioural

features used are not validated measures of self-regulation, motivation or institutional integration, and their theoretical labelling rests on design rationale rather than on construct-validity evidence. Eighth, quantum inputs were subjected to a dimensionality-reduction pipeline not applied to the classical baselines, so the comparison isolates the effect of the full quantum pipeline rather than of the quantum kernel alone.

6.2 Future Work

Future research should prioritize four directions. First, computational efficiency: reducing QSVM runtime through quantum kernel approximation, Nyström methods, and quantum-inspired algorithms to enable operational viability at institutional scale (Gujju et al., 2024). Second, temporal and sequential modelling: extending the framework to quantum recurrent or sequential architectures that directly model the temporal evolution of student engagement rather than aggregated weekly snapshots. Third, fairness and ethics: formal fairness audits disaggregated by gender, disability status, and socioeconomic background before any institutional deployment, drawing on established frameworks for algorithmic fairness in education (Baker & Hawn, 2022). Fourth, multi-institutional validation: replication across diverse institutional contexts — campus-based universities, MOOCs, and blended learning environments — to establish external validity of the regime-dependent quantum advantage identified here.

6.3 Acknowledgments

This study uses the Open University Learning Analytics Dataset (OULAD), publicly available via the UCI Machine Learning Repository (https://archive.ics.uci.edu/dataset/349/open+university+learning+analytics+dataset). The dataset was accessed and used in accordance with its stated terms. The code used for data preprocessing, model training, and evaluation is available from the authors upon reasonable request.

References

Baker, R. S., & Hawn, A. (2022). Algorithmic bias in education. International journal of artificial intelligence in education, 32(4), 1052-1092. https://doi.org/10.1007/s40593-021-00285-9

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

Castro-Montoya, B., Vélez-Gómez, P., Segura-Cardona, A., & French, B. F. (2025). A cultural adaptation of Tinto's student integration theory in undergraduate students of a private university in Colombia. Cogent Education, 12(1), 2479384. https://doi.org/10.1080/2331186X.2025.2479384

Credé, M., & Kuncel, N. R. (2008). Study habits, skills, and attitudes: The third pillar supporting collegiate academic performance. Perspectives on psychological science, 3(6), 425-453. https://doi.org/10.1111/j.1745-6924.2008.00089.x

Eccles, J. S., & Wigfield, A. (2020). From expectancy-value theory to situated expectancy-value theory: A developmental, social cognitive, and sociocultural perspective on motivation. Contemporary educational psychology, 61, 101859. https://doi.org/10.1016/j.cedpsych.2020.101859

Gujju, Y., Matsuo, A., & Raymond, R. (2024). Quantum machine learning on near-term quantum devices: Current state of supervised and unsupervised techniques for real-world applications. Physical Review Applied, 21(6), 067001. DOI: https://doi.org/10.1103/PhysRevApplied.21.067001

Jha, N. I., Ghergulescu, I., & Moldovan, A. N. (2019). OULAD MOOC dropout and result prediction using ensemble, deep learning and regression techniques. In Proceedings of the 11th International Conference on Computer Supported Education (CSEDU 2019), pages 154-164. DOI: 10.5220/0007767901540164

Li, I. W., & Carroll, D. R. (2020). Factors influencing dropout and academic performance: An Australian higher education equity perspective. Journal of Higher Education Policy and Management, 42(1), 14-30. https://doi.org/10.1080/1360080X.2019.1649993

Martin, A. J., & Marsh, H. W. (2008). Academic buoyancy: Towards an understanding of students' everyday academic resilience. Journal of school psychology, 46(1), 53-83. https://doi.org/10.1016/j.jsp.2007.01.002

Martin, A. J., Colmar, S. H., Davey, L. A., & Marsh, H. W. (2010). Longitudinal modelling of academic buoyancy and motivation: Do the '5Cs' hold up over time? British journal of educational psychology, 80(3), 473-496. https://doi.org/10.1348/000709910X486376

Piepenburg, J. G., & Beckmann, J. (2022). The relevance of social and academic integration for students' dropout decisions. Evidence from a factorial survey in Germany. European Journal of Higher Education, 12(3), 255-276. https://doi.org/10.1080/21568235.2021.1930089

Pintrich, P. R., & De Groot, E. V. (1990). Motivational and self-regulated learning components of classroom academic performance. Journal of educational psychology, 82(1), 33. https://doi.org/10.1037/0022-0663.82.1.33

Richardson, M., Abraham, C., & Bond, R. (2012). Psychological correlates of university students' academic performance: A systematic review and meta-analysis. Psychological bulletin, 138(2), 353. https://doi.org/10.1037/a0026838

Romdhoni, R. D., Majid, N. W. A., & Permatasari, A. F. (2026). Machine learning-based early warning for student dropout: Evidence from LMS behavioral engagement patterns in online higher education. Journal of Informatics and Vocational Education, 9(2), 55-66. https://doi.org/10.20961/joive.v9i2.3508

Salmela-Aro, K., & Read, S. (2017). Study engagement and burnout profiles among Finnish higher education students. Burnout research, 7, 21-28. https://doi.org/10.1016/j.burn.2017.11.001

Salmela-Aro, K., & Upadyaya, K. (2014). School burnout and engagement in the context of demands–resources model. British journal of educational psychology, 84(1), 137-151. https://doi.org/10.1111/bjep.12018

Samoila, M. E., & Vrabie, T. (2023). First-year seminars through the lens of Vincent Tinto's theories of student departure. A systematic review. Frontiers in Education, 8, 1205667. doi: 10.3389/feduc.2023.1205667

Schuld, M., & Killoran, N. (2019). Quantum machine learning in feature Hilbert spaces. Physical Review Letters, 122(4), Article 040504. https://doi.org/10.1103/PhysRevLett.122.040504

Sghir, N., Adadi, A., & Lahmer, M. (2023). Recent advances in predictive learning analytics: A decade systematic review (2012–2022). Education and Information Technologies, 28(7), 8299–8333. https://doi.org/10.1007/s10639-022-11536-0

Tinto, V. (2012). Leaving college: Rethinking the causes and cures of student attrition. University of Chicago Press.

von Keyserlingk, L., Schwerter, J., Kunina-Habenicht, O., & Lauermann, F. (2026). Motivation and self-regulated learning as predictors of learning resource use in learning management systems in higher education. Educational Psychology, 46(3), 437-457. https://doi.org/10.1080/01443410.2025.2564891

Zimmerman, B. J. (2000). Self-efficacy: An essential motive to learn. Contemporary Educational Psychology, 25(1), 82–91. https://doi.org/10.1006/ceps.1999.1016