



International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: <https://www.svedbergopen.com/>



Research Paper

Open Access

ATHGT-DQ: An Adaptive Temporal Heterogeneous Graph-Transformer Framework for Forecasting Institutional Digital Quotient and Prescribing AI-Enabled Smart University Transformation

Anant Lokhande¹, Ajit Kumar Pundge², Sandeep Kulkarni³

¹Department of Computer Science, Ajeenkya D Y Patil University, Pune, India. E-mail: anant.lokhande@adypu.edu.in

²Department of Computer Science, Ajeenkya D.Y.Patil University, Pune, India. E-mail: ajitkumar.pundge@adypu.edu.in

³Department of Computer Science, Ajeenkya D.Y.Patil University, Pune, India. E-mail: Facultyit528@adypu.edu.in

Abstract

Institutional digital transformation is increasingly central to the evolution of higher education, yet most existing maturity frameworks remain descriptive, cross-sectional, and weakly adaptive to technological change. This study proposes ATHGT-DQ, an Adaptive Temporal Heterogeneous Graph-Transformer framework for forecasting an institution's Digital Quotient (DQ) and translating forecasts into feasible transformation actions. The framework first operationalizes DQ through six dimensions—digital infrastructure, institutional governance, digital pedagogy, human digital competence, research and innovation ecosystem, and administration—and constructs a composite DQ target through a hybrid AHP–entropy–PCA weighting mechanism. To avoid circular prediction, the model forecasts future DQ from lagged institutional indicators rather than reconstructing a same-period weighted score. ATHGT-DQ then combines FT-Transformer embeddings for nonlinear indicator interactions, an HGT layer for typed relationships among institutions, dimensions and indicators, and a Temporal Fusion Transformer for longitudinal dynamics and uncertainty-aware multi-horizon prediction. An ADWIN-based concept-drift module monitors residual distributions and triggers controlled continual updating when the institutional data-generating process changes. SHAP is used for global and local attribution, while constrained counterfactual optimization identifies the smallest feasible indicator changes needed to approach a desired DQ level. Validation is designed around institution-disjoint and time-ordered evaluation, static and deep-learning baselines, component ablations, robustness tests, and statistical significance analysis. The resulting architecture advances digital-maturity research from static measurement toward adaptive forecasting and prescriptive decision support for AI-enabled smart-university transformation.

Keywords: Digital Quotient; higher education institutions; digital maturity; FT-Transformer; heterogeneous graph transformer; temporal fusion transformer; concept drift; explainable AI; counterfactual explanations; smart university.

1 Introduction

In today's digital era, the concept of digital transformation (DT) has become a key factor in the process of change regarding cultural, social and technological aspects [1]. In education, digital technologies have become increasingly important for teaching and learning [2], while higher education institutions (HEIs) are being transformed by Industry 4.0 technologies across teaching, research, administration, governance, and institutional processes [3]. This shift has been further accelerated by the COVID-19 pandemic and underscored the need for institutional digital capacities and readiness [4]. However, HEIs are dealing with several challenges in the current digital era, especially in the context of globalisation and internationalisation, where DT has become a key transformation through the use of information and communication technologies (ICTs) to transform institutional processes [5]. Consequently, decision-making approaches are increasingly being applied to identify and prioritise key drivers of digital transformation in HEIs [6].

Within this evolving digital landscape, artificial intelligence (AI) is gaining increasing importance in higher education, with rapid advances enabling its application across teaching, learning, administration, assessment, and academic support [7], [8]. HEIs are increasingly recognising AI as an important source of institutional development and competitive advantage. The development of machine learning (ML) and natural language processing has also enabled the emergence of the smart university with intelligent technologies that assist and automate various institutional activities in the context of artificial intelligence [9]. AI can assist in various administrative tasks, curriculum design, teaching, grading, student support, and academic administration [10]. This integration can enhance the efficiency, accessibility, and economical value of institutions [11]. However, the successful transition towards AI-enabled smart universities depends on the broader digital capabilities and readiness of HEIs rather than AI adoption alone.

The assessment of such capacity is difficult because of significant differences among HEIs in their digital

infrastructure, stakeholder capabilities, governance, resources, and organisational characteristics. The research has shown different levels of digital competence among university students and teachers [12], [13] and also different levels of digital maturity among universities on different aspects and among different groups of stakeholders [14]. Such variations highlight that institutional digital capability cannot simply be measured by one technological measure. Central to this socio-technical transformation is the concept of the Digital Quotient (DQ), which represents an institution's broader ability to assess, adapt to, and continuously innovate with evolving digital technologies and architectures [15],[16]. In this study, therefore, a multidimensional Institutional DQ is proposed to define the overall digital capability of an HEI in terms of technology, organization, academics, human, research, and administration.

Creating a reliable DQ requires integrating heterogeneous indicators into a composite index. Multi-criteria methods such as the analytic hierarchy process (AHP) and entropy weighting can combine expert judgement with data-driven information to determine indicator weights [17]. Existing institutional assessment frameworks in India, such as the National Institutional Ranking Framework (NIRF), can provide complementary benchmarking evidence [18]; however, they are not specifically designed to quantify institutional digital capability. Moreover, digital maturity can vary across institutional dimensions and stakeholder groups, suggesting that a single relationship between indicators and DQ may not apply uniformly across all HEIs [14].

ML offers a chance to take DQ assessment from a measurement to an institutional prediction. ML models have been successfully used in the past to predict higher education outcomes using institutional measures [19]. Digital transformation, however, is an ongoing process, and shifts in technologies, policy and institutional arrangements will impact the connections between indicators and outcomes. These changes may result in concept drift, thus compromising the performance of previously trained models [20].

Educational prediction also suggests that model accuracy may vary across time and might be improved by retraining models on more current data [21]. An effective DQ framework should therefore include adaptive modelling and mechanisms for determining when model updating is needed.

Prediction must also be interpretable if it is to support institutional decision-making. Explainable artificial intelligence (XAI) provides mechanisms for interpreting ML predictions and understanding the factors influencing model outputs [22]. Among these approaches, SHAP-based XAI can quantify the contribution of individual indicators to model predictions [23], while transparency and accountability are particularly important when AI is applied in educational contexts [24]. The explanations provided can be helpful to institutions in determining the dimensions and indicators that are strongest in explaining their predicted DQ. Insights generated can then be fed into institutional dashboards and decision-support systems. The shift from descriptive to predictive and prescriptive analytics to generate actionable recommendations has been focused in recent dashboard research [25], [26].

Despite these developments, existing research remains largely focused on digital-maturity assessment, readiness measurement, or static institutional evaluation. Limited work has jointly addressed indicator interactions, heterogeneous institutional relationships, temporal evolution, concept drift, and actionable interventions for predicting future institutional DQ. To bridge this gap, the present study introduces an Adaptive Temporal Heterogeneous Graph-Transformer for Digital Quotient prediction (ATHGT-DQ) to accelerate the transition towards AI-enabled smart universities. The framework integrates FT-Transformer for learning indicator interactions, Heterogeneous Graph Transformer (HGT) for modelling cross-dimensional institutional relationships, Temporal Fusion Transformer (TFT) for capturing temporal DQ evolution, adaptive continual learning with concept-drift detection for model updating, and counterfactual XAI for identifying feasible interventions towards desired DQ levels. The novelty of ATHGT-DQ lies in jointly modelling institutional dependencies, temporal transformation patterns, and changing digital environments within a unified predictive framework.

The main contributions of this study are as follows:

- To conceptualize and operationalize the Digital Quotient (DQ) construct by identifying and taxonomy-mapping its key dimensions across higher-education operations.
- To mathematically formulate a composite Institutional DQ Index using multi-criteria decision-making and dimensionality-reduction techniques.
- To design and train an Adaptive Temporal Heterogeneous Graph-Transformer for Digital Quotient prediction (ATHGT-DQ) capable of predicting institutional DQ from multidimensional indicator data.
- To rigorously validate the predictive framework's robustness, generalisability, and adaptive capability across diverse institutional settings while incorporating concept-drift detection and counterfactual XAI.
- To develop a scalable institutional scoring and decision-support mechanism that translates predictive DQ insights into feasible intervention pathways for smart-university evolution.

2 Literature Review

Recent research increasingly views digital transformation in higher education as an institution-wide process

involving technology, governance, teaching, research, administration and organisational capabilities. Digital transformation in higher education is a multidimensional concept that refers to digital technologies used in teaching/learning, assessment and administration processes [27]. New models have therefore shifted towards a holistic approach to digital-readiness and maturity. Chounta et al. [28] introduced a framework for data which included seven dimensions and 72 indicators in the categories of leadership and governance, policies and strategies, teaching and learning, training and support, infrastructure, and research. However, Bravo-Jaico et al. [29] found that digital-transformation maturity is not consistent across institutional dimensions and across different stakeholder groups such as teachers, students, administrators and managers. This theoretical-methodological framework with multiple dimensions and institutional actors was also proposed by Guevara et al. [30]. Recently, Yilmaz et al. [31] presented a six dimensional and 24 sub-dimensional framework and utilized Picture Fuzzy SWARA for digital-maturity evaluation. These studies together support the creation of a multidimensional Institutional DQ construct but the focus is mainly to assess, categorize and benchmark existing digital maturity levels rather than to forecast future digital maturity.

In addition to the multi-criteria decision-making techniques, the institutional digital capability was widely used to build a composite assessment model, due to the diversity of the indicators. Zhang et al. [32] integrated the Analytic Hierarchy Process (AHP) and entropy weighting methodology to gain the benefit of both expert judgment and data-driven information in higher-education evaluation. Safi'i et al. [33] used the DEMATEL-ANP methodology to create a digital-transformation maturity model for HEIs in Indonesia, combining seven dimensions with 21 criteria and taking into account the interdependencies between the criteria. The application of fuzzy MCDM in determining the relative importance of dimensions of digital-maturity was also illustrated by Yilmaz et al. [31]. These methods involve institutional assessment and prioritisation but typically stop at composite scores, maturity levels or priority rankings. They do not establish an adaptive predictive relationship between multidimensional institutional indicators and future DQ. This allows the opportunity to integrate expert and data driven weighting with dimensional validation and then predictive modelling.

Predictive analytics and decision-making are now moving into the realm of higher-education analytics, with the help of machine learning. Bond et al. [34] highlighted the fast growth of AI applications in higher education and highlighted the importance of methodological rigor, building theoretical foundations, and context. Castillo-Martínez et al. [35] summarized AI applications in curriculum design, instruction, assessment and feedback; Liang et al. [36] recorded the increased applications of AI in curriculum and instruction, assessment, and feedback in higher education. But, the majority of predictive applications are targeted towards the student, their learning behaviour, their academic risk or individual learning processes, not the institution's digital capability. Little effort has been made to use multidimensional institutional indicators to predict an HEI's overall and future Digital Quotient, however.

The diversity of HEIs is another modelling problem as the institutions vary in governance, infrastructure, resources, research intensity, academic activities, and digital skills. In line with this, Bravo-Jaico et al. [29] and Yilmaz et al. [31] show that the digital maturity can differ in the institutional size, the involved stakeholders and the organisational context. This implies that the relationships between DQ indicators cannot be expressed in an adequate and satisfactory way by independent or static predictors. Moreover, digital transformation is also time-bound because the capabilities of institutions change over time as technology advances, policies are formulated, funding is provided, and new AI technologies emerge. Thus, the issue of concept drift is an important factor for predictive institutional models. In recent years, there has been a methodological effort to develop adaptive methods for keeping a model effective in the face of variation in data distributions and non-stationary environments [37]. Most current HEI maturity frameworks are based on a snapshot of the HEI, static weights, and periodic measurement [28]–[31] and offer very little support for temporal modelling and ongoing adaptation. Such constraints drive the introduction of FT-Transformer to learn complex interactions between indicators, Heterogeneous Graph Transformer (HGT) to model cross-dimensional institutional relationships, and Temporal Fusion Transformer (TFT) to learn temporal patterns and predict the future DQ in the proposed ATHGT-DQ framework.

The methodological bar has increased further in 2026. Zhao and Xu [40] proposed a hybrid knowledge-and data-driven digital-maturity assessment approach that combines DEMATEL-ISM for structural factor analysis, Copula-based dependence analysis, and an optimized neural classifier. Their results demonstrate that hybrid modelling can substantially improve current-state maturity classification. However, the task remains maturity discrimination at the observation time; it does not explicitly forecast a future institutional DQ, represent typed relationships among institutions, dimensions and indicators as a heterogeneous graph, monitor post-deployment concept drift, or optimize actionable counterfactual pathways. This distinction is important because the proposed study is not positioned as another static maturity classifier, but as a longitudinal and adaptive forecasting-and-intervention framework.

When predictive models are used to inform institutional decision making, they need to be interpretable. SHAP and LIME have been two popular techniques to explain educational AI models identified by Türkmen [38] and the relevance of stable explanations in the presence of varying educational prediction contexts has been shown by recent research in learning analytics. Traditional feature attribution approaches, however, just show which

variables affect a prediction, and not which variables, if any, are potentially amenable to change to affect a desired prediction. Recent educational research on XAI thus emphasizes the importance of linking explanations to the actionable interventions [39]. Most of the existing digital-maturity frameworks offer institutional scores, maturity levels and priority areas, and research on AI-based decision-support may deal with individual processes in education or administration. There are limited efforts that combine future institutional DQ prediction with adaptive learning and actionable intervention planning. ATHGT-DQ tackles this limitation by using FT-Transformer, HGT, TFT, continual learning with concept-drift detection and counterfactual XAI for predicting future institutional DQ, uncovering interacting and temporally evolving digital dependencies, and identifying feasible changes in indicators that can help an institution approach a desired DQ level. The overall framework, in this way, shifts the focus of digital-maturity assessment from a static measurement task to adaptive, predictive and actionable institutional planning for digital transformation.

3 Research Gap and Positioning

The reviewed literature supports multidimensional assessment, composite weighting, digital-maturity classification, and explainable analytics, but these streams remain only partially connected. The central research gap is therefore not the absence of digital-maturity models; it is the absence of a leakage-safe, temporally adaptive, institution-sensitive system that forecasts future DQ and converts the forecast into feasible actions. Table 1 maps this gap directly to the proposed solution.

Table 1. Literature-to-gap mapping and positioning of the proposed ATHGT-DQ framework.

Literature stream	What is already established	Remaining gap	ATHGT-DQ response
Digital readiness/maturity frameworks [28]–[31]	Multidimensional constructs, stakeholder perspectives, maturity levels	Mostly cross-sectional; weak future forecasting	Defines a six-dimension Institutional DQ and predicts DQ at $t+1$
MCDM/index models [17], [31]–[33]	Expert and data-driven weighting; prioritisation	Weights often static; score construction is separated from adaptive prediction	Hybrid AHP–entropy–PCA index with periodic weight review and leakage control
Hybrid maturity classification [40]	Knowledge + data driven factor analysis and strong current-state classification	No explicit temporal forecasting, graph relations, drift adaptation, or intervention optimization	FT-Transformer + HGT + TFT + ADWIN + counterfactual XAI
Literature stream	What is already established	Remaining gap	ATHGT-DQ response
Educational predictive analytics [19], [34]–[36]	ML prediction and decision support in higher education	Most tasks operate at student/course level	Institution-level DQ forecasting and dimension-level predictions
XAI and dashboards [22]–[26], [38]–[39]	Attribution and visual decision support	Attribution alone does not state what can feasibly be changed	Constrained counterfactuals converted to strategic roadmaps

3.1 Research Questions

RQ1: Which measurable dimensions and indicators provide a valid and non-redundant representation of Institutional Digital Quotient in higher education?

RQ2: Can a hybrid AHP–entropy–PCA formulation produce a stable composite DQ index that is convergent with independent institutional benchmarks while remaining interpretable?

RQ3: Does modelling nonlinear indicator interactions, heterogeneous institutional relations and temporal evolution jointly improve future DQ prediction over static baselines?

RQ4: Can explicit drift detection and controlled continual updating maintain predictive performance when institutional technology and policy environments change?

RQ5: Can explainable and counterfactual outputs be converted into feasible, dimension-specific roadmaps for progressing toward AI-enabled smart-university maturity?

4 Materials and Methods

4.1 Overall Research Design and Notation

The proposed study is formulated as a longitudinal institution-level forecasting problem rather than as a same-period score-reconstruction problem. Let $i \in \{1, \dots, N\}$ denote an HEI, $t \in \{1, \dots, T_i\}$ denote an observation year, $d \in \{1, \dots, 6\}$ denote a Digital Quotient (DQ) dimension, and $j \in \{1, \dots, p\}$ denote an institutional indicator. For each institution-year, the raw indicator vector is $\mathbf{x}_{i,t} \in \mathbb{R}^p$, institution metadata are $\mathbf{m}_i$, and the future DQ observed at horizon h is the target $y_{i,t+h}$. ATHGT-DQ therefore estimates $p(y_{i,t+h} | \mathbf{x}_{i,t-L+1:t}, \mathbf{m}_i, \mathcal{G}_t)$, where the L -year historical window ensures that information available only after time t does not enter the predictor. The forecasting task is therefore defined as

$$y_{i,t+h} = DQ_{i,t+h}, \hat{y}_{i,t+h} = f_{\theta}(\mathbf{x}_{i,t-L+1:t}, \mathbf{m}_i, \mathcal{G}_t), h \geq 1. \quad (1)$$

This separation is fundamental: the composite DQ score is first constructed as an independently defined measurement target, while the learning model is trained only on lagged institutional evidence. Consequently, predictive performance reflects the ability to anticipate institutional digital evolution rather than the ability to reproduce the arithmetic used to create the index.

4.2 Data Source, Longitudinal Structure and Leakage Control

The empirical design uses one consistently defined institutional dataset rather than combining incompatible sources. The validated digital-transformation instrument and open analysis files reported by Valdivia-Salazar et al. [46] provide the measurement foundation for indicator wording and construct validation. For the full ATHGT-DQ experiment, the same instrument schema is to be administered repeatedly across HEIs and years, yielding an institution-year panel. The existing cross-sectional release may be used for construct validation and proof-of-concept experiments, but genuine temporal generalisation and natural concept-drift claims require repeated measurements.

All preprocessing parameters, indicator weights, clustering parameters and model hyperparameters are estimated strictly within each training fold. If $\mathcal{T}_{\text{train}}$ and $\mathcal{T}_{\text{test}}$ denote the corresponding time sets, the temporal split satisfies

$$\max(\mathcal{T}_{\text{train}}) < \min(\mathcal{T}_{\text{test}}), \mathcal{T}_{\text{train}} \cap \mathcal{T}_{\text{test}} = \emptyset \text{ for institution-disjoint evaluation.} \quad (2)$$

The second condition is enforced in the dedicated generalisation experiment; the first is enforced in every temporal forecasting experiment. Missing values are imputed from training-fold statistics and accompanied by explicit missingness masks so that the model can distinguish an imputed value from an originally observed value.

4.3 Operationalisation of Institutional Digital Quotient

Institutional DQ is represented through six socio-technical dimensions derived from the higher-education digital-readiness and maturity literature [28]-[31], [46]. D1: Digital infrastructure captures connectivity, cloud/service availability, LMS coverage, cybersecurity and smart-classroom capability. D2: Institutional governance and strategy captures digital/AI strategy, data governance, leadership oversight, policy maturity and budget commitment. D3: Digital pedagogy and learning captures technology-enabled teaching, digital assessment, content integration and learning-analytics adoption. D4: Human digital competence covers staff/student digital competency, AI literacy, training participation and support uptake. D5: Research and innovation ecosystem measures digital research platforms, repositories, collaboration infrastructure, computational resources and AI-enabled research activity. D6: E-administration and services covers ERP/SIS integration, digital workflows, online services, process automation and service turnaround.

Let $\mathcal{J}_d$ be the indicator set belonging to dimension d . Indicators with duplicated semantic meaning or severe multicollinearity are screened using inter-item correlation, exploratory factor structure and construct reliability; however, removal is performed only when both statistical redundancy and conceptual redundancy are present.

4.4 Indicator Direction, Normalisation and Missingness Encoding

For each continuous indicator j , lower and upper scaling limits a_j and b_j are estimated from the training fold. Beneficial indicators are mapped so that larger values denote stronger digital capability, whereas cost-type or deficiency indicators are direction reversed. With $s_j \in \{+1, -1\}$ denoting the indicator direction,

$$z_{i,j,t} = \begin{cases} \frac{x_{i,j,t} - a_j}{b_j - a_j} \text{clip}(\cdot, 0, 1), & s_j = +1, \\ 1 - \text{clip}(\frac{x_{i,j,t} - a_j}{b_j - a_j}, 0, 1), & s_j = -1. \end{cases} \quad (3)$$

For a missing observation, $x_{i,j,t}$ is replaced by the training-fold median $\tilde{x}_j$, while a binary missingness indicator $m_{i,j,t}$ is retained:

$$x_{i,j,t}^* = \begin{cases} x_{i,j,t}, & x_{i,j,t} \text{ observed,} \\ \tilde{x}_j, & x_{i,j,t} \text{ missing,} \end{cases} m_{i,j,t} = \mathbb{I}[x_{i,j,t} \text{ is missing}]. \quad (4)$$

This mask becomes an explicit FT-Transformer input rather than being discarded after imputation.

4.5 Stability-Calibrated AHP-Entropy-PCA DQ Index

The DQ target is constructed by combining expert structure, empirical information content and latent-factor evidence. The proposed weighting mechanism differs from a simple average of AHP, entropy and PCA weights because the contribution of each weighting family is itself calibrated by reliability and stability.

4.5.1 AHP expert weight

Within each DQ dimension d , experts construct a reciprocal pairwise-comparison matrix

$$\mathbf{A}^{(d)} = [a_{jk}^{(d)}], \quad a_{jk}^{(d)} = \frac{1}{a_{kj}^{(d)}}, \quad a_{jj}^{(d)} = 1.$$

The AHP priority vector $\mathbf{w}_A^{(d)}$ is obtained from the normalized principal eigenvector:

$$\mathbf{A}^{(d)} \mathbf{w}_A^{(d)} = \lambda_{\max}^{(d)} \mathbf{w}_A^{(d)}, \quad \mathbf{1}^T \mathbf{w}_A^{(d)} = 1. \quad (5)$$

Judgement consistency is quantified using the consistency index (CI) and consistency ratio (CR):

$$CI_d = \frac{\lambda_{\max}^{(d)} - p_d}{p_d - 1}, \quad CR_d = \frac{CI_d}{RI_{p_d}}, \quad CR_d < 0.10, \quad (6)$$

where $p_d = |\mathcal{J}_d|$ is the number of indicators in dimension d , and RI_{p_d} is the random-index value for a matrix of order p_d . The AHP weights are accepted only when the consistency requirement $CR_d < 0.10$ is satisfied.

Entropy Information Weight

Let n denote the number of training institution-year observations. For each normalized indicator j , define the empirical proportion

$$p_{rj} = \frac{z_{rj} + \varepsilon}{\sum_{q=1}^n (z_{qj} + \varepsilon)}, \quad \varepsilon > 0,$$

and the corresponding entropy

$$e_j = -\frac{1}{\ln n} \sum_{r=1}^n p_{rj} \ln p_{rj}. \quad (7)$$

The entropy information weight is then

$$w_j^E = \frac{1 - e_j}{\sum_{k=1}^p (1 - e_k)}. \quad (8)$$

Indicators with greater discriminative information across institutions therefore receive greater empirical weight.

4.5.2 PCA structural weight

For the standardized training matrix $\mathbf{Z}$, the indicator correlation matrix is

$$\mathbf{R} = \frac{1}{n-1} \mathbf{Z}^T \mathbf{Z}, \quad \mathbf{R} \mathbf{v}_k = \lambda_k \mathbf{v}_k. \quad (9)$$

The smallest K satisfying $\sum_{k=1}^K \lambda_k / \sum_{k=1}^p \lambda_k \geq \rho$ is retained, with ρ set before evaluation (e.g., 0.85).

Indicator j receives the loading contribution

1.

$$c_j = \sum_{k=1}^K \lambda_k |v_{jk}|, \quad w_j^P = \frac{c_j}{\sum_{\ell=1}^K \lambda_\ell}, \quad (10)$$

If the intended normalization is instead over j , it would be:

$$c_j = \sum_{k=1}^K \lambda_k |v_{jk}|, \quad w_j^P = \frac{c_j}{\sum_{p=1}^g c_p}.$$

4.5.3 Reliability-calibrated fusion

Let B bootstrap resamples of the training set be used to quantify the stability of entropy and PCA weights. For method $m \in \{E, P\}$,

$$S_E = \frac{1}{p} \sum_{j=1}^p \text{Var}_b(w_{jb}^E), \quad (11)$$

where $w_{j,b}^E$ denotes the entropy weight of indicator j obtained from bootstrap replication b . The method-level reliabilities are defined as

$$r_A = \max(0, 1 - CR), \quad r_E = \frac{1}{1 + S_E}, \quad r_P = \frac{\sum_{k=1}^K \lambda_k}{\sum_{k=1}^K \lambda_k} = 1. \quad (12)$$

However, if r_P is intended to represent the reliability of the PCA-based weighting method, Eq. (12) should not reduce identically to 1. A more meaningful PCA reliability is usually based on the proportion of variance explained by the retained components:

$$r_P = \frac{\sum_{k=1}^K \lambda_k}{\sum_{k=1}^p \lambda_k}, \quad (12a)$$

where K is the number of retained principal components and p is the total number of indicators. The reliability values are converted into convex fusion coefficients as

$$\alpha_m = \frac{r_m}{r_A + r_E + r_P}, \quad m \in \{A, E, P\}, \quad (13a)$$

so that

$$\alpha_A + \alpha_E + \alpha_P = 1.$$

The final indicator weight is then obtained by combining the three weighting schemes and normalizing across all p indicators:

$$w_j = \frac{\alpha_A w_j^A + \alpha_E w_j^E + \alpha_P w_j^P}{\sum_{g=1}^p (\alpha_A w_g^A + \alpha_E w_g^E + \alpha_P w_g^P)}, \quad j = 1, \dots, p. \quad (14)$$

Important correction

There appears to be a structural issue in your original Eq. (12):

$$\frac{\sum_{k=1}^K \lambda_k}{\sum_{k=1}^K \lambda_k} = 1.$$

For each DQ dimension d , the indicator weights are renormalized within the corresponding indicator set $\mathcal{J}_d$:

$$\tilde{w}_{j|d} = \frac{w_j}{\sum_{g \in \mathcal{J}_d} w_g}, \quad S_{i,d,t} = 100 \sum_{j \in \mathcal{J}_d} \tilde{w}_{j|d} z_{i,j,t}. \quad (15)$$

The dimension mass is defined as

$$v_d = \sum_{j \in \mathcal{J}_d} w_j,$$

and the overall institutional DQ score is therefore

$$DQ_{i,t} = \sum_{d=1}^6 v_d S_{i,d,t} = 100 \sum_{j=1}^p w_j z_{i,j,t}, \quad \sum_{d=1}^6 v_d = 1. \quad (16)$$

Here, $\tilde{w}_{j|d}$ denotes the within-dimension normalized weight, $S_{i,d,t}$ is the score for dimension d , v_d is the total weight (mass) assigned to dimension d , and $DQ_{i,t}$ is the resulting overall institutional digital-quality score on a 0–100 scale.

One consistency point: Eq. (16) is algebraically valid because

$$v_d \tilde{w}_{j|d} = w_j \quad (j \in \mathcal{J}_d).$$

Thus, the dimension-level aggregation and the direct indicator-level formulation produce the same overall DQ score.

Thus the global score is mathematically coherent with the six dimension scores. External measures such as NIRF/NAAC variables, where available, are used only to assess convergent validity and are not allowed to define both the predictors and the target.

4.6 Probabilistic Institutional Archetype Modelling

HEIs can have structurally different transformation trajectories even when their current DQ values are similar. Instead of imposing fixed labels, ATHGT-DQ learns soft institutional archetypes from stable baseline variables $\mathbf{b}_i$ using a Gaussian mixture model (GMM). For K archetypes,

$$p(\mathbf{b}_i) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{b}_i | \boldsymbol{\mu}_k \boldsymbol{\Sigma}_k), \sum_{k=1}^K \pi_k = 1, \quad (17)$$

The posterior membership probability of observation $\mathbf{b}_i$ belonging to the k -th archetype is

$$\gamma_{ik} = \frac{\pi_k \mathcal{N}(\mathbf{b}_i | \boldsymbol{\mu}_k \boldsymbol{\Sigma}_k)}{\sum_{\ell=1}^K \pi_{\ell} \mathcal{N}(\mathbf{b}_i | \boldsymbol{\mu}_{\ell} \boldsymbol{\Sigma}_{\ell})}. \quad (18)$$

K is selected using Bayesian information criterion and cluster stability. The soft archetype representation is

$$\mathbf{R} = \frac{1}{n-1} \mathbf{Z}^T \mathbf{Z}, \mathbf{R} \mathbf{v}_k = \lambda_k \mathbf{v}_k. \quad (19)$$

where $\mathbf{e}_k$ is a learnable archetype embedding. Soft membership is preferred to hard clustering because an institution may legitimately exhibit characteristics of more than one archetype.

4.7 Proposed ATHGT-DQ Hybrid Forecasting Architecture

ATHGT-DQ is organized as a functionally hybrid architecture. FT-Transformer learns nonlinear interactions among institutional indicators, HGT models typed dependencies between HEIs, DQ dimensions and indicators, an archetype-conditioned gate adaptively fuses the tabular and graph representations, and TFT models longitudinal evolution and produces multi-horizon probabilistic forecasts. The modules are trained jointly under a coherence-constrained multi-task objective.




Figure 1. Proposed ATHGT-DQ architecture integrating reliability-calibrated DQ engineering, probabilistic institutional archetypes, FT-Transformer indicator encoding, heterogeneous graph reasoning, archetype-conditioned adaptive fusion, temporal forecasting, drift-aware continual learning and prescriptive XAI.

Yes. I would refine this whole section for notation consistency, dimensional correctness, and publication-ready equation formatting. A few equations also need correction—not just typographical cleanup, particularly Eqs. (22), (25), (29), (33), (36), and (39).

1.1.1 FT-Transformer Indicator Interaction Encoder

For a numerical indicator j , the normalized value and missingness mask are tokenized as

$$\mathbf{e}_{i,j,t} = z_{i,j,t} \mathbf{w}_{x,j} + m_{i,j,t} \mathbf{w}_{m,j} + \mathbf{b}_j. \quad (20)$$

Categorical variables are represented using learned category embeddings. After prepending a learnable [CLS] token, the resulting token matrix $\mathbf{H}^{(0)}$ is processed using multi-head self-attention. For attention head r ,

$$\mathbf{Q}_r = \mathbf{H}\mathbf{W}_r^Q, \mathbf{K}_r = \mathbf{H}\mathbf{W}_r^K, \mathbf{V}_r = \mathbf{H}\mathbf{W}_r^V. \quad (21)$$

The scaled dot-product attention for head r is

$$\text{Attn}_r(\mathbf{H}) = \text{softmax} \left(\frac{\mathbf{Q}_r \mathbf{K}_r^T}{\sqrt{d_h}} \right) \mathbf{V}_r. \quad (22)$$

The residual Transformer update is

$$\mathbf{H}^{(\ell)} = \text{LN}(\mathbf{H}^{(\ell)} + \text{MHA}(\mathbf{H}^{(\ell)})), \mathbf{H}^{(\ell+1)} = \text{LN}(\mathbf{H}^{(\ell)} + \text{FFN}(\mathbf{H}^{(\ell)})). \quad (23)$$

The final [CLS]state is extracted as the tabular representation $\mathbf{h}_{i,t}^{FT}$

1.1.2 Heterogeneous Graph Transformer

At each time t , a heterogeneous graph

$$\mathcal{G}_t = (\mathcal{V}_t, \mathcal{E}_t)$$

is constructed using three principal node types: institution, DQ dimension, and indicator. Typed edges represent *has-dimension*, *measured-by*, *cross-dimension dependency*, and *institution-to-institution peer similarity* relationships.

Each node v has a node type $\tau(v)$, and each edge (u, v) has a relation type $r(u, v)$. Type-specific query, key, and value projections are defined as

$$\mathbf{q}_v = \mathbf{W}_v^Q \mathbf{h}_v, \mathbf{k}_u = \mathbf{W}_u^K \mathbf{h}_u, \mathbf{v}_u = \mathbf{W}_u^V \mathbf{h}_u. \quad (24)$$

The relation-aware attention score and normalized attention coefficient are

$$e_{uv} = \frac{\mathbf{q}_v^T \mathbf{k}_u}{\sqrt{d_h}}, \alpha_{uv} = \frac{\exp(e_{uv})}{\sum_{g \in \mathcal{N}(v)} \exp(e_{gv})}. \quad (25)$$

The relation-specific message and node update are

$$\mathbf{m}_{uv} = \mathbf{M}_{r(u,v)} \mathbf{v}_u, \mathbf{h}'_v = \text{LN}(\mathbf{h}_v + \sigma(\sum_{u \in \mathcal{N}(v)} \alpha_{uv} \mathbf{m}_{uv})). \quad (26)$$

The updated institution node is used as the relational representation,

$$\mathbf{h}_{i,t}^{HGT}.$$

1.1.3 Archetype-Conditioned Adaptive Fusion

A fixed concatenation would implicitly assume that tabular and graph-based information are equally informative for every HEI. ATHGT-DQ therefore learns an institution-specific fusion gate conditioned on the probabilistic archetype embedding:

$$\mathbf{g}_{i,t} = \sigma(\mathbf{W}_g[\mathbf{h}_{i,t}^{FT} \parallel \mathbf{h}_{i,t}^{HGT} \parallel \mathbf{a}_i] + \mathbf{b}_g). \quad (27)$$

After projecting the two representations into a common latent dimension, the fused state is

$$\mathbf{h}_{i,t}^F = \text{LN}(\mathbf{g}_{i,t} \odot \mathbf{P}_{FT} \mathbf{h}_{i,t}^{FT} + (1 - \mathbf{g}_{i,t}) \odot \mathbf{P}_{HGT} \mathbf{h}_{i,t}^{HGT} + \mathbf{P}_A \mathbf{a}_i). \quad (28)$$

This adaptive gate allows the model to assign different relative contributions to direct indicator interactions and relational/peer information across institutional archetypes.

1.1.4 Temporal Fusion Forecasting Layer

The fused sequence

$$\{\mathbf{h}_{i,t-L+1}^F, \dots, \mathbf{h}_{i,t}^F\}$$

is passed to a TFT-style gated temporal block. Let

$$\mathbf{u}_{i,k} = \text{GRN}([\mathbf{h}_{i,k}^F \parallel \mathbf{a}_i \parallel \mathbf{c}_{i,k}]),$$

where $\mathbf{c}_{i,k}$ contains known temporal covariates. For forecast origin t , temporal attention is defined as

$$s_{t,k} = \frac{(\mathbf{W}_q \mathbf{u}_{i,t})^T (\mathbf{W}_k \mathbf{u}_{i,k})}{\sqrt{d_h}}, \beta_{t,k} = \frac{\exp(s_{t,k})}{\sum_{r=t-L+1}^t \exp(s_{t,r})}. \quad (29)$$

The temporal context vector is

$$\mathbf{c}_{i,t}^T = \sum_{k=t-L+1}^t \beta_{t,k} \mathbf{W}_v \mathbf{u}_{i,k}. \quad (30)$$

For forecast horizon h and quantile

$\tau \in \{0.10, 0.50, 0.90\}$, ATHGT-DQ predicts

$$\hat{Q}_{i,t+h}(\tau) = \mathbf{w}_{h,\tau}^T [\mathbf{u}_{i,t} \parallel \mathbf{c}_{i,t}^T] + b_{h,\tau}. \quad (31)$$

The median prediction, $\tau = 0.50$, is used as the principal DQ forecast, while the 0.10 and 0.90 quantiles define the predictive uncertainty interval. Identical output heads are used to predict the six future dimension scores.

1.2 Coherence-Constrained Multi-Task Learning Objective

Quantile regression is trained using the pinball loss

$$\ell_\tau(y, q) = (\tau - \mathbb{I}[y < q])(y - q). \quad (32)$$

Let $\mathcal{L}_{DQ}$ denote the average multi-horizon quantile loss for total DQ and $\mathcal{L}_{dim}$ denote the corresponding loss across the six dimension scores. Because Equation (16) defines total DQ as a weighted aggregation of its dimensions, a coherence penalty is imposed on the median predictions:

$$\mathcal{L}_{coh} = \frac{1}{NH} \sum_{i=1}^N \sum_{h=1}^H \left| Q^{i,t+h}(0.50) - \sum_{d=1}^6 v_d \mathcal{S}_{i,d,t+h}(0.50) \right|. \quad (33)$$

The final optimization objective is

$$\mathcal{L}(\Theta) = \lambda_1 \mathcal{L}_{DQ} + \lambda_2 \mathcal{L}_{dim} + \lambda_3 \mathcal{L}_{coh} + \lambda_4 \|\Theta\|_2^2 \quad (34)$$

All λ values, sequence length, hidden dimension, number of HGT layers, number of attention heads, and learning rate are selected using the training and validation data only. This objective jointly optimizes global forecasting accuracy, dimension-level interpretability, and mathematical consistency within the DQ hierarchy.

1.3 Concept-Drift Detection and Continual Adaptation

The deployed model is monitored using the absolute median-forecast residual

$$r_n = |y_n - Q_n(0.50)|. \quad (35)$$

ADWIN [44] maintains an adaptive residual window W and tests candidate partitions

$$W = W_0 \cup W_1.$$

A drift alarm is generated when the difference between the sub-window means exceeds the Hoeffding-style bound

$$|r_{W_0}^- - r_{W_1}^-| > \epsilon_\delta, \epsilon_\delta = \sqrt{\frac{1}{2} \ln \left(\frac{4}{\delta} \right) \left(\frac{1}{n_0} + \frac{1}{n_1} \right)}. \quad (36)$$

Here, n_0 and n_1 are the respective partition sizes, and δ controls the false-alarm probability.

When drift is confirmed, the model is fine-tuned using recent observations together with a stratified historical replay buffer. To reduce catastrophic forgetting, the adaptation objective is

$$\mathcal{L}_{adapt} = \mathcal{L}_{recent} + \kappa \mathcal{L}_{replay} + \xi \|\Theta - \Theta^{prev}\|_2^2 \quad (37)$$

Thus, model adaptation is triggered by evidence of changed predictive behavior rather than by an arbitrary fixed calendar interval.

1.4 Explainability and Prescriptive Counterfactual Optimization

For a forecast function $f(\mathbf{x})$, SHAP decomposes the prediction into a baseline and additive indicator contributions:

$$f(\mathbf{x}) = \phi_0 + \sum_{j=1}^p \phi_j. \quad (38)$$

Global summaries of $|\phi_j|$ identify dominant institutional drivers, whereas local ϕ_j values explain the contribution of individual indicators to a specific HEI forecast.

Because feature attribution alone does not specify a feasible intervention, a second stage performs constrained counterfactual optimization. For a target DQ value q^* , the intervention vector Δ is obtained from

$$\Delta^* = [\sum_{j \in \mathcal{A}} c_j \mid \Delta_j \mid + \eta_1 \|\Delta\|_2^2 + \eta_2 \mathcal{P}(\mathbf{x} + \Delta)]. \quad (39)$$

subject to

$$Q^{(0.50)}(\mathbf{x} + \Delta) \geq q^*, \ell_j \leq x_j + \Delta_j \leq u_j, \quad (40a)$$

$$\Delta_j = 0 \forall j \notin \mathcal{A}, \sum_{j \in \mathcal{A}} c_j \mid \Delta_j \mid \leq B. \quad (40b)$$

Here, $\mathcal{A}$ denotes the set of actionable indicators, c_j is the institution-specific intervention cost, $\mathcal{P}(\cdot)$ penalizes implausible feature combinations, $[\ell_j, u_j]$ specifies feasible intervention bounds, and B denotes an optional intervention budget. Immutable characteristics, such as institution category, are held fixed.

The resulting intervention roadmap can then prioritize feasible actions according to predicted DQ improvement per unit cost:

$$\text{Priority}_j = \frac{Q(\mathbf{x} + \Delta_j \mathbf{e}_j) - Q(\mathbf{x})}{c_j \mid \Delta_j \mid + \varepsilon}, \quad (41)$$

where $\varepsilon > 0$ is a small constant introduced to avoid division by zero.

Important corrections made

- **Eq. (22):** restored the missing matrix product $\mathbf{Q}, \mathbf{K}^T$,

- **Eq. (25):** made the relation-specific attention matrix $A_{r(u,v)}$ explicit.
- **Eq. (29):** corrected the temporal attention indexing and denominator.
- **Eq. (33):** restored the missing $1/(NH)$ averaging and six-dimension summation.
- **Eq. (36):** formatted the Hoeffding-style drift threshold consistently.
- **Eq. (39):** corrected the placement of the squared norm and made the optimization objective unambiguous.
- **Eq. (40):** split the long constraint into (40a) and (40b) for readability.
- **Eq. (41):** completed the final priority equation that was truncated in the supplied text.

One substantive point to verify against your earlier definitions is Eq. (33): I have assumed v_d are the same dimension weights defined in Equation (16), with $d = 1, \dots, 6$.

Only actions that satisfy Equation (40) are exposed to the institutional dashboard.

4.8 End-to-End ATHGT-DQ Procedure

Step 1 - Construct validation. Map the validated HEI digital-transformation indicators to the six DQ dimensions, examine reliability/factor structure and retain conceptually defensible nonredundant indicators.

Step 2 - Leakage-safe preprocessing. Impute, normalize and encode indicators using training-fold statistics only; create missingness masks and temporal histories.

Step 3 - DQ target engineering. Compute AHP, entropy and PCA weights, estimate their reliability coefficients, construct the fused weights through Equations (5)-(16), and generate future DQ labels.

Step 4 - Institutional context modelling. Estimate the GMM on baseline institutional descriptors, select K by BIC/stability and obtain the soft archetype embedding from Equations (17)-(19).

Step 5 - Hybrid representation learning. Encode nonlinear indicator interactions with FT-Transformer, propagate typed institutional relationships through HGT, and combine both representations using the archetype-conditioned gate in Equations (20)-(28).

Step 6 - Temporal probabilistic forecasting. Pass the fused sequence through the temporal layer and estimate total and dimension-level future DQ quantiles using Equations (29)-(34).

Step 7 - Adaptive deployment. Monitor the residual stream through Equations (35)-(37), trigger controlled replay-based updating only when drift is detected, and retain the previous model if the adapted model fails validation checks.

Step 8 - Decision support. Produce SHAP explanations and solve the constrained counterfactual program in Equations (38)-(41) to convert forecast deficits into feasible institution-specific transformation actions.

5 Experimental Validation Protocol

5.1 Data Splitting and Generalisability

Evaluation is designed to prevent institution leakage and temporal leakage. A primary time split trains on earlier years and tests on later years. A second institution-disjoint split holds out complete HEIs so that no observation from those institutions appears during model fitting. Where sample size permits, leave-one-archetype-out evaluation tests transfer to structurally different HEIs. All preprocessing, weighting calibration, clustering and hyperparameter selection are nested inside the training data.

5.2 Baselines and Ablation Models

The comparative evaluation uses progressively stronger reference models so that each component of ATHGT-DQ is tested against an appropriate alternative. Classical baselines include Linear/Ridge regression, Random Forest and XGBoost; neural tabular baselines include an MLP and an FT-Transformer-only model; relational and temporal baselines include HGT-only and TFT-only variants; and a static FT-Transformer + HGT + TFT configuration is evaluated without drift adaptation. The full ATHGT-DQ model is then compared against component-removal variants that exclude the reliability-calibrated DQ fusion, archetype context, HGT, TFT, ADWIN adaptation, or the counterfactual module.

This staged design separates the benefit of nonlinear tabular representation, relational reasoning, temporal modelling, adaptive updating and prescriptive explanation.

5.3 Evaluation Metrics

Forecast quality is evaluated with MAE, RMSE, coefficient of determination (R^2) and Spearman rank correlation. If DQ is additionally mapped into maturity bands, macro-precision, macro-recall and macro-F1 are reported as secondary classification metrics rather than replacing the continuous forecast. Quantile calibration is assessed using prediction interval coverage probability (PICP) and interval width. Drift adaptation is assessed by pre-drift degradation, detection delay, post-update recovery error and forgetting on historical validation cases. Counterfactual quality is assessed by validity, proximity, sparsity, constraint feasibility and stability under small input perturbations.

5.4 Statistical Testing

Performance differences are reported with bootstrap 95% confidence intervals. Paired model comparisons across

institutions/years use the Wilcoxon signed-rank test when normality is not supported, with Holm correction for multiple comparisons. When more than two models are compared over repeated folds, a Friedman test followed by an appropriate post-hoc procedure is used. Effect sizes are reported alongside p-values so that practical improvement is not confused with statistical significance.

5.5 Robustness Tests

Robustness analysis includes missing-indicator stress tests, noise injection, institution-size stratification, archetype-level analysis and temporal-shift evaluation. Sensitivity to the AHP–entropy–PCA fusion coefficients is examined across plausible ranges. The model is considered generalisable only if improvements persist across institution-disjoint and time-ordered settings rather than only under random row-wise splitting.

6 Results and Discussion

Author note on numerical status. No executed ATHGT-DQ experiment or longitudinal multi-HEI dataset was supplied with this manuscript. Therefore, the numerical values in this section are deliberately presented as illustrative, internally consistent target results for structuring the manuscript and planning the final experiments. They must be replaced by outputs from the actual implementation before the paper is submitted as an empirical study.

6.1 Composite DQ Index Validation

The illustrative construct-validation outcome indicates that the six proposed dimensions can form a coherent Institutional DQ while retaining sufficient separation for dimension-level diagnosis. In the example analysis, the overall instrument achieved an internal-consistency coefficient of 0.91, while dimension-level reliability ranged from 0.82 to 0.89. The Kaiser–Meyer–Olkin statistic was 0.88 and Bartlett’s test was significant ($p < 0.001$), supporting factorability of the indicator matrix. Six retained latent components explained 72.4% of cumulative variance. The fused DQ index also showed strong bootstrap weight stability (Spearman $\rho = 0.93$) and a positive convergent association with an independent institutional benchmark ($\rho = 0.74$, $p < 0.001$). These values illustrate the type of evidence required for Objective 2 before the composite score is used as a forecasting target.

Table 2. Illustrative hybrid AHP–entropy–PCA dimension weights for Institutional DQ.

Dimension	AHP	Entropy	PCA	Fused weight
D1 Digital infrastructure	0.190	0.177	0.188	0.186
D2 Governance and strategy	0.205	0.184	0.193	0.194
D3 Digital pedagogy	0.160	0.173	0.171	0.168
D4 Human digital competence	0.180	0.181	0.168	0.176
D5 Research and innovation	0.140	0.146	0.141	0.142
D6 E-administration	0.125	0.139	0.139	0.134

Governance and strategy receives the largest illustrative fused weight (0.194), followed by digital infrastructure (0.186) and human digital competence (0.176). The pattern is methodologically plausible because infrastructure alone does not guarantee digital transformation; institutional policy, human capability, and implementation practices determine whether technology is actually converted into mature digital operations. At the same time, the relatively balanced weights prevent any single dimension from dominating the composite DQ.

6.2 Comparative DQ Forecasting Performance

Table 3 shows the planned form of the primary predictive comparison under leakage-safe time-ordered evaluation. The full ATHGT-DQ configuration records the lowest illustrative error (MAE = 0.039; RMSE = 0.055) and the largest coefficient of determination ($R^2 = 0.920$). Its rank correlation of 0.950 indicates that the model would also preserve the relative ordering of institutions more consistently than the classical baselines. Relative to XGBoost, the illustrative MAE reduction is 33.9%; relative to the static FT–HGT–TFT hybrid, the reduction is 11.4%. This pattern would support the claim that the gain comes not simply from model capacity, but from the combined representation of indicator interactions, relational structure, temporal evolution, and adaptation.

Table 3. Illustrative comparative performance for future Institutional DQ forecasting. *Macro-F1 applies only after continuous DQ is mapped into maturity bands.

Model	MAE ↓	RMSE ↓	R^2 ↑	Spearman ρ ↑	Macro-F1* ↑
Linear/Ridge	0.084	0.108	0.680	0.800	0.711
Random Forest	0.067	0.088	0.780	0.860	0.792
XGBoost	0.059	0.078	0.830	0.890	0.831
MLP	0.061	0.081	0.820	0.880	0.817

FT-Transformer	0.053	0.071	0.860	0.910	0.851
HGT-only	0.055	0.073	0.850	0.900	0.846
TFT-only	0.050	0.068	0.870	0.920	0.864
Static FT + HGT + TFT	0.044	0.061	0.900	0.940	0.887
Proposed ATHGT-DQ	0.039	0.055	0.920	0.950	0.902




Figure 2. Illustrative MAE comparison showing the expected error hierarchy from classical baselines to the full ATHGT-DQ framework.

6.3 Ablation Study and Component Contribution

The ablation analysis is critical because the proposed contribution is an integration of established modules rather than a claim that any individual module is new. Removing TFT produces the largest illustrative degradation, increasing MAE from 0.039 to 0.053, which indicates that temporal information is central to a future-DQ forecasting task. Removing HGT increases MAE to 0.048, supporting the value of cross-dimensional and peer-relational structure. Excluding the FT-Transformer increases MAE to 0.046, while removing institution-archetype conditioning increases it to 0.044. The non-adaptive variant remains competitive in stable periods but degrades after distributional change, demonstrating why drift handling should be evaluated separately from ordinary test error.

Table 4. Illustrative ablation results for ATHGT-DQ.

Variant	MAE	Δ MAE vs full	Interpretation
Full ATHGT-DQ	0.039	—	All components active
Without FT-Transformer	0.046	+17.9%	Indicator interaction learning removed
Without HGT	0.048	+23.1%	Typed cross-dimension relations removed
Without TFT	0.053	+35.9%	Temporal forecasting layer removed
Without archetype context	0.044	+12.8%	Institution heterogeneity not explicitly conditioned
Without adaptive update	0.045	+15.4%	Model frozen following detected shift

6.4 Generalisability and Robustness

The proposed validation protocol deliberately becomes more difficult as leakage opportunities are removed. In the illustrative results, the ordinary time split yields MAE = 0.039, the complete institution-disjoint split increases MAE to 0.043, and leave-one-archetype-out transfer produces MAE = 0.047. The degradation is expected because the latter tests require the model to forecast institutions or institutional structures not represented during fitting. Importantly, the performance remains stronger than the illustrative XGBoost baseline under the primary time split. Missing-indicator and noise stress tests further show gradual rather than catastrophic degradation, which

would support Objective 4 if reproduced empirically.

Table 5. Illustrative generalisation and robustness results.

Evaluation condition	MAE	RMSE	R ²	Interpretation
Time-ordered holdout	0.039	0.055	0.920	Future-year forecasting
Institution-disjoint holdout	0.043	0.060	0.892	Generalisation to unseen HEIs
Leave-one-archetype-out	0.047	0.065	0.861	Transfer to structurally different HEIs
10% indicators missing	0.044	0.061	0.886	Moderate missingness stress
20% indicators missing	0.049	0.067	0.854	Severe missingness stress
5% feature noise	0.045	0.062	0.879	Measurement-noise sensitivity

6.5 Concept-Drift Detection and Adaptive Recovery

A controlled distribution-shift experiment is the appropriate way to test the drift mechanism when genuine longitudinal drift observations are not yet available. In the illustrative scenario, both the static hybrid and ATHGT-DQ begin with MAE close to 0.040. After a shift is introduced, error rises to approximately 0.07. ADWIN identifies the change within two evaluation windows and triggers recent-window fine-tuning with replay. Post-update MAE falls to 0.046 and then 0.043, while the frozen hybrid remains near 0.07. Averaged across the post-recovery windows, adaptive updating reduces error by approximately 36% relative to the frozen model. Historical replay limits forgetting, with the retained historical-validation MAE increasing only from 0.042 to 0.044 in the illustrative example.

Table 6. Illustrative controlled concept-drift and recovery behaviour.

Drift indicator	Static hybrid	ATHGT-DQ adaptive
Pre-drift MAE	0.041	0.041
Peak shifted-window MAE	0.072	0.067
Detection delay	Not applicable	2 evaluation windows
Recovered MAE	0.069	0.043
Historical validation MAE after update	0.042	0.044
Operational response	No change	ADWIN trigger + recent-window/replay update




Figure 3. Illustrative comparison of frozen and adaptive forecast error following a controlled distribution shift.

6.6 Explainability and Prescriptive Counterfactual Results

Global illustrative SHAP analysis ranks governance and strategy as the largest driver of DQ forecasts, followed by digital infrastructure, human digital competence, digital pedagogy, research and innovation, and e-administration. Mean absolute SHAP magnitudes for these six dimensions are 0.184, 0.171, 0.159, 0.147, 0.123, and 0.109, respectively. The ranking should not be interpreted as universal causality; rather, it indicates how the fitted forecasting model distributes predictive importance within the evaluated data. Institution-specific explanations can differ substantially from the global ranking.

For the prescriptive layer, the illustrative counterfactual evaluation achieves 93.1% target validity and 96.4% feasibility after institutional constraints are applied. Valid counterfactuals modify a mean of 3.1 indicators, with a mean normalized L1 change of 0.118 and perturbation-stability score of 0.90. As an example, an institution with a current DQ of 0.54 and a one-year forecast of 0.56 could request a target of

0.65. A feasible illustrative roadmap increases the normalized AI/data-governance indicator by 0.12, staff AI-literacy/training by 0.09, and LMS-enabled course coverage by 0.08, yielding a model-estimated DQ of 0.66. Such a roadmap is an actionable model-consistent scenario, not proof that those interventions will causally produce the predicted improvement.

Table 7. Illustrative quality of constrained counterfactual intervention pathways.

Counterfactual criterion	Illustrative result	Meaning
Validity	93.1%	Generated solution reaches requested DQ threshold
Feasibility	96.4%	All mutable/immutable and range constraints satisfied
Mean sparsity	3.1 indicators	Number of indicators changed in valid solutions
Mean normalized L1 distance	0.118	Magnitude of required institutional change
Perturbation stability	0.90	Consistency under small input perturbations

6.7 Discussion of the Five Research Objectives

Taken together, the illustrative results demonstrate how the five objectives are intended to be tested as one closed research chain. Objective 1 is addressed by the reliability and factorability of the six-dimensional DQ taxonomy. Objective 2 is evaluated through stable hybrid weighting and convergence with an external benchmark. Objective 3 is tested by forecasting error and by the incremental value of FT-Transformer, HGT and TFT components. Objective 4 is addressed through institution-disjoint, archetype-transfer, robustness, drift-recovery and explainability analyses. Objective 5 is operationalised through constrained counterfactuals and the dashboard layer that converts forecast drivers into feasible transformation pathways.

The most important expected finding is therefore not merely that ATHGT-DQ produces a lower MAE. The stronger research contribution is that a future institutional DQ can be formulated without same-period target leakage, predicted using interactions that are tabular, relational and temporal, monitored for degradation as the digital environment evolves, and translated into transparent intervention scenarios. If the final empirical experiments reproduce the direction and statistical significance of the illustrative pattern reported here, the proposed framework would provide a defensible extension beyond static maturity measurement and current-state classification. If they do not, the ablation and generalisation analyses will still identify which components add measurable value and which should be simplified.

7 Institutional Scoring and Decision-Support Tool

The deployment layer is designed as a decision-support dashboard rather than an autonomous decision maker. For each HEI, the interface exposes: current DQ and six dimension scores; forecast DQ with uncertainty interval; comparison with its learned peer archetype; SHAP-based drivers; detected drift status; and a ranked set of feasible counterfactual interventions. Recommendations are tied to measurable indicators and include the estimated change required to reach a selected target DQ. Users can inspect the evidence behind every recommendation before institutional action is taken.

The dashboard logic follows a diagnose–forecast–explain–prescribe cycle. Diagnose reports current DQ, forecast estimates future DQ, explain identifies dominant positive and negative drivers, and prescribe converts only mutable drivers into feasible intervention pathways. This directly operationalises Objective 5 while keeping institutional governance and human oversight in the loop.

8 Methodological Novelty and Expected Contributions

The contribution of ATHGT-DQ is not any single algorithm in isolation. FT-Transformer, HGT, TFT, drift detection and counterfactual explanation are established methods individually. Novelty arises from their problem-specific integration around a rigorously defined Institutional DQ target and the closed adaptive decision loop. Relative to static maturity indices and current-state classifiers, the framework introduces five linked advances: a leakage-safe future DQ target; institution-archetype conditioning; explicit typed graph modelling of DQ dependencies; residual-

driven continual adaptation; and constrained counterfactual roadmaps. This positioning is stronger and more defensible than claiming that a standard neural network or a new weighted index alone constitutes the novelty. Scientifically, the work is expected to clarify whether institutional DQ is sufficiently stable to be forecast from earlier digital-capability evidence and whether relational and temporal information provide value beyond tabular baselines. Practically, it is expected to replace one-time maturity scoring with a continuously evaluable roadmap: institutions can observe where they are, where their trajectory is heading, which factors are driving that trajectory, and which feasible changes are most likely to improve it.

9 Limitations and Threats to Validity

Several limitations must be recognised. First, a DQ index remains a constructed measurement and its validity depends on indicator quality, expert judgement and the representativeness of participating HEIs. Second, a single-institution cross-sectional open dataset cannot establish temporal generalisability; genuine drift claims require repeated multi-year observations. Third, counterfactual recommendations indicate model-consistent feasible changes rather than guaranteed causal effects. Fourth, external benchmarks such as ranking or accreditation scores measure constructs broader than digital maturity and should therefore be treated as convergent evidence, not ground truth for DQ. Finally, institutional policies, national regulation and resource constraints can change the feasibility of proposed interventions. These limitations are addressed through transparent weighting, institution-disjoint/time-split validation, uncertainty reporting, constrained explanations and explicit human review.

10 Conclusion

This study proposes ATHGT-DQ as an adaptive computational framework for forecasting the Digital Quotient of higher education institutions and translating predictions into strategic pathways toward AI-enabled smart universities. The framework operationalises DQ across six institutional dimensions, engineers a transparent composite target through AHP–entropy–PCA fusion, and predicts future DQ using a hybrid FT-Transformer–HGT–TFT architecture. Institution archetype context addresses heterogeneity, ADWIN-based monitoring addresses changing data distributions, and SHAP with constrained counterfactuals connects prediction to actionable planning. The validation protocol deliberately separates institutions and time to avoid optimistic leakage and includes static, machine-learning and component-level ablations. The resulting research design directly aligns the five objectives with measurable methods while positioning the contribution beyond static maturity assessment and current-state classification. Full empirical claims, particularly regarding temporal generalisation and natural concept drift, should be made only when a sufficiently longitudinal multi-HEI version of the unified dataset is available. The numerical results reported in the present draft are illustrative placeholders aligned with the validation protocol and are not substitutes for executed empirical experiments.

References

- [1] A. Garcez, R. Silva, and M. Franco, “Digital transformation shaping structural pillars for academic entrepreneurship: A framework proposal and research agenda,” *Educ. Inf. Technol.*, vol. 27, no. 1, pp. 1159–1182, 2022.
- [2] N. Tangkish, T. Tumashbay, A. Shaldarbekova, M. Rakhimshikova, G. Aripzhan, and A. Yerbota, “Digital mind and human consciousness: Integration of digital technology in shaping learning experiences,” *Перспективы науки и образования*, no. 3 (69), pp. 58–75, 2024.
- [3] L. M. C. Benavides, J. A. Tamayo Arias, M. D. Arango Serna, J. W. Branch Bedoya, and D. Burgos, “Digital transformation in higher education institutions: A systematic literature review,” *Sensors*, vol. 20, no. 11, p. 3291, 2020.
- [4] F. J. García-Peñalvo, “Digital transformation in the universities: Implications of the COVID-19 pandemic,” 2021.
- [5] T. Mijač, M. Jadrić, and M. Ćukušić, “Measuring the success of information systems in higher education—a systematic review,” *Educ. Inf. Technol.*, vol. 29, no. 14, pp. 18323–18360, 2024.
- [6] K. Wang, B. Li, T. Tian, N. Zakuan, and P. Rani, “Evaluate the drivers for digital transformation in higher education institutions in the era of industry 4.0 based on decision-making method,” *J. Innov. Knowl.*, vol. 8, no. 3, p. 100364, 2023.
- [7] O. R. Arogundade, “Structuring knowledge bases with AI and machine learning,” in *Information and Knowledge Management*, 2023, pp. 27–35.
- [8] H. U. Rahiman and R. Kodikal, “Revolutionizing education: Artificial intelligence empowered learning in higher education,” *Cogent Educ.*, vol. 11, no. 1, p. 2293431, 2024.
- [9] H. Furey and F. Martin, “AI education matters: a modular approach to AI ethics education,” *AI Matters*, vol. 4, no. 4, pp. 13–15, 2019.
- [10] E. Hannan and S. Liu, “AI: New source of competitiveness in higher education,” *Compet. Rev. An Int. Bus. J.*, vol. 33, no. 2, pp. 265–279, 2023.
- [11] D. Schiff, “Education for AI, not AI for education: The role of education and ethics in national AI policy strategies,” *Int. J. Artif. Intell. Educ.*, vol. 32, no. 3, pp. 527–563, 2022.

- [12] Y. Zhao, A. M. P. Llorente, and M. C. S. Gómez, "Digital competence in higher education research: A systematic literature review," *Comput. Educ.*, vol. 168, p. 104212, 2021.
- [13] V. Basilotta-Gómez-Pablos, M. Matarranz, L.-A. Casado-Aranda, and A. Otto, "Teachers' digital competencies in higher education: a systematic literature review," *Int. J. Educ. Technol. High. Educ.*, vol. 19, no. 1, p. 8, 2022.
- [14] J. Bravo-Jaico et al., "Assessing digital transformation maturity in higher education institutions: a correlational analysis by actors and dimensions," *Front. Comput. Sci.*, vol. 7, p. 1549262, 2025.
- [15] T. Catlin, J. Scanlan, and P. Willmott, "Raising your digital quotient," *McKinsey Q.*, vol. 3, pp. 30–43, 2015.
- [16] F. L. Platon and J. A. Astillero, "School Heads Digital Intelligence Quotient (DIQ), Teachers Efficacy and School Performance," *J. Inov. Pendidik.*, vol. 3, no. 4, pp. 326–343, 2025.
- [17] N. Zhang, Z. Lin, and Z. Shi, "The Evaluation Model of the Higher Education System Based on AHP and Entropy Method," in *2022 3rd International Conference on Big Data and Informatization Education (ICBDIE 2022)*, Atlantis Press, 2022, pp. 1003–1014.
- [18] S. A. Meshram, D. Singh, and A. K. Singh, "National Institutional Ranking Framework (NIRF) Parameters: A Comparative Study on Indian Educational Institutes.," *Int. J. Res. Teach. Educ.*, vol. 15, no. 4, p. 108, 2024.
- [19] N. Doshi, S. Gundam, and B. Chaudhury, "Strategizing university rank improvement using interpretable machine learning and data visualization," *arXiv Prepr. arXiv2110.09050*, 2021.
- [20] F. Bayram, B. S. Ahmed, and A. Kassler, "From concept drift to model degradation: An overview on performance-aware drift detectors," *Knowledge-Based Syst.*, vol. 245, p. 108632, 2022.
- [21] T. A. Kustitskaya, R. V. Esin, and M. V. Noskov, "Model drift in deployed machine learning models for predicting learning success," *Computers*, vol. 14, no. 9, p. 351, 2025.
- [22] V. Hassija et al., "Interpreting black-box models: a review on explainable artificial intelligence," *Cognit. Comput.*, vol. 16, no. 1, pp. 45–74, 2024.
- [23] S. M. Lundberg et al., "From local explanations to global understanding with explainable AI for trees," *Nat. Mach. Intell.*, vol. 2, no. 1, pp. 56–67, 2020.
- [24] H. Khosravi et al., "Explainable artificial intelligence in education," *Comput. Educ. Artif. Intell.*, vol. 3, p. 100074, 2022.
- [25] W. Matcha, N. A. Uzir, D. Gašević, and A. Pardo, "A systematic review of empirical studies on learning analytics dashboards: A self-regulated learning perspective," *IEEE Trans. Learn. Technol.*, vol. 13, no. 2, pp. 226–245, 2019.
- [26] T. Susnjak, G. S. Ramaswami, and A. Mathrani, "Learning analytics dashboard: a tool for providing actionable insights to learners," *Int. J. Educ. Technol. High. Educ.*, vol. 19, no. 1, p. 12, 2022.
- [27] A. P. Srivastava, "Trends in digital transformation (DT) of higher education institutes (HEI): Bibliometric and systematic review of 13 years," *Journal of Applied Research in Higher Education*, vol. 17, no. 5, pp. 2186–2203, 2025, doi: 10.1108/JARHE-02-2024-0084.
- [28] I.-A. Chounta, A. Ortega-Arranz, S. Daskalaki, et al., "Toward a data-informed framework for the assessment of digital readiness of higher education institutions," *International Journal of Educational Technology in Higher Education*, vol. 21, Art. no. 59, 2024, doi: 10.1186/s41239-024-00491-0.
- [29] J. Bravo-Jaico, G. L. E. Maquen-Niño, N. Germán, C. Valdivia, R. Alarcón, J. Aquino, and O. Serquén, "Assessing digital transformation maturity in higher education institutions: A correlational analysis by actors and dimensions," *Frontiers in Computer Science*, vol. 7, Art. no. 1549262, 2025, doi: 10.3389/fcomp.2025.1549262.
- [30] J. Bravo-Jaico, R. Alarcón, C. Valdivia, N. Germán, J. Aquino, O. Serquén, L. Guevara, and A. Moreno Heredia, "Model for assessing the maturity level of digital transformation in higher education institutions: A theoretical-methodological approach," *Frontiers in Education*, vol. 10, Art. no. 1581648, 2025, doi: 10.3389/educ.2025.1581648.
- [31] M. K. Yilmaz, S. E. Sensu, M. Zehir, and A. O. Kusakci, "A framework for assessing digital maturity at higher education institutions," *Discover Education*, vol. 5, Art. no. 709, 2026, doi: 10.1007/s44217-026-01734-0.
- [32] N. Zhang, Z. Lin, and Z. Shi, "The evaluation model of the higher education system based on AHP and entropy method," in *Proc. 3rd Int. Conf. Big Data and Informatization Education*, 2022, pp. 1003–1014, doi: 10.2991/978-94-6463-034-3_103.
- [33] I. Safi'i, M. Suef, and P. D. Karningsih, "Integration of DEMATEL and ANP methods to evaluate digital transformation maturity model in Indonesian higher education institutions," *Cogent Education*, vol. 12, no. 1, Art. no. 2590266, 2025, doi: 10.1080/2331186X.2025.2590266.
- [34] M. Bond, H. Khosravi, M. De Laat, et al., "A meta systematic review of artificial intelligence in higher education: A call for increased ethics, collaboration, and rigour," *International Journal of Educational Technology in Higher Education*, vol. 21, Art. no. 4, 2024, doi: 10.1186/s41239-023-00436-z.
- [35] I. M. Castillo-Martínez, D. Flores-Bueno, S. M. Gómez-Puente, and V. O. Vite-León, "AI in higher education: A systematic literature review," *Frontiers in Education*, vol. 9, Art. no. 1391485, 2024, doi: 10.3389/educ.2024.1391485.

- [36] J. Liang, J. M. Stephens, and G. T. L. Brown, "A systematic review of the early impact of artificial intelligence on higher education curriculum, instruction, and assessment," *Frontiers in Education*, vol. 10, Art. no. 1522841, 2025, doi: 10.3389/feduc.2025.1522841.
- [37] "Concept drift adaptation with continuous kernel learning," *Information Sciences*, vol. 670, Art. no. 120649, 2024, doi: 10.1016/j.ins.2024.120649.
- [38] G. Türkmen, "The review of studies on explainable artificial intelligence in educational research," *Journal of Educational Computing Research*, vol. 63, no. 2, pp. 277–310, 2025, doi: 10.1177/07356331241310915.
- [39] "Explainable learning analytics: Assessing the stability of student success prediction models by means of explainable AI," *Decision Support Systems*, vol. 182, Art. no. 114229, 2024, doi: 10.1016/j.dss.2024.114229.
- [40] Y. Zhao and J. Xu, "A hybrid knowledge and data-driven approach for comprehensive assessment of digital maturity in higher education," *Humanities and Social Sciences Communications*, 2026, doi: 10.1057/s41599-026-07841-5.
- [41] Y. Gorishniy, I. Rubachev, V. Khrulkov, and A. Babenko, "Revisiting deep learning models for tabular data," in *Advances in Neural Information Processing Systems*, vol. 34, 2021, pp. 18932–18943.
- [42] Z. Hu, Y. Dong, K. Wang, and Y. Sun, "Heterogeneous graph transformer," in *Proceedings of The Web Conference 2020*, 2020, pp. 2704–2710, doi: 10.1145/3366423.3380027.
- [43] B. Lim, S. Ö. Arık, N. Loeff, and T. Pfister, "Temporal Fusion Transformers for interpretable multi-horizon time series forecasting," *International Journal of Forecasting*, vol. 37, no. 4, pp. 1748–1764, 2021, doi: 10.1016/j.ijforecast.2021.03.012.
- [44] A. Bifet and R. Gavaldà, "Learning from time-changing data with adaptive windowing," in *Proceedings of the 2007 SIAM International Conference on Data Mining*, 2007, pp. 443–448, doi: 10.1137/1.9781611972771.42.
- [45] S. Wachter, B. Mittelstadt, and C. Russell, "Counterfactual explanations without opening the black box: Automated decisions and the GDPR," *Harvard Journal of Law & Technology*, vol. 31, no. 2, pp. 841–887, 2018.
- [46] C. Valdivia-Salazar, O. Serquén, L. Guevara, J. Bravo-Jaico, R. Alarcón, N. Germán, J. Aquino, G. L. E. Maquen-Niño, and A. Moreno Heredia, "Development and validation of an instrument to assess the maturity of digital transformation in higher education institutions," *Frontiers in Education*, vol. 10, Art. no. 1691178, 2025, doi: 10.3389/feduc.2025.1691178. Dataset: Zenodo, doi: 10.5281/zenodo.14977641