SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

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

Research Paper Open Access

Intelligent Generative AI-Assisted Academic Advisement System Using Learning Analytics and Transformer Models

Maha AlBalushi[1], Muna Al Balushi[2], G. Tamilarasi[3*]

[1] Faculty in the Department of Computing and Information Sciences, University of Technology and Applied Sciences - Al Musannah, Sultanate of Oman. E-mail: maha.Albalushi@utas.edu.om
[2] Faculty in the Department of Computing and Information Sciences, University of Technology and Applied Sciences - Sohar, Sultanate of Oman. E-mail: muna.Albalushi@utas.edu.om
[3*] Faculty in the Department of Computing and Information Sciences, University of Technology and Applied Sciences - Al Musannah, Sultanate of Oman. E-mail: tamilarasi.g@utas.edu.om
*Corresponding Author email: tamilarasi.g@utas.edu.om

ABSTRACT
A lack of personalisation, delayed intervention, and high student-to-advisor ratios can limit conventional academic advising. This study combines an empirical analysis of student-performance factors with a proposed AI-assisted academic-advisement architecture. The implemented component comprised descriptive and bivariate statistical analyses of a publicly available synthetic dataset containing 6,607 student records and 20 variables. Attendance showed the strongest bivariate linear association with examination score (($r=0.581$)), followed by hours studied (($r=0.446$)). These results describe relationships encoded in the synthetic data and do not establish causal effects or real-world predictive performance. The proposed architecture comprises five layers: data input; preprocessing and feature engineering; learning analytics and prediction; FT-Transformer and generative AI; and explainability and advisement output. The FT-Transformer design uses feature-specific embeddings and multi-head self-attention to model relationships among numerical and categorical attributes. Huber loss and regularisation are specified as design choices intended to reduce sensitivity to large residuals and mitigate overfitting; their benefits have not yet been experimentally evaluated. Course recommendation, language-based advice generation, explainability, and advisor-feedback mechanisms remain proposed components requiring implementation and validation. The framework is intended for future evaluation in GCC higher education institutions and subsequent adaptation to Arabic-language advising. Prospective studies should assess predictive performance, fairness, explanation quality, usability, and effects on student outcomes.

Keywords: Academic advisement, explainable artificial intelligence, Feature Tokenization Transformer, generative artificial intelligence, learning analytics, student-performance prediction

I. Introduction

Universities spend significant time counseling students at different stages of their programs about which courses to take, how to improve their grades, and when they risk failing or dropping out. Many of these advising sessions are facilitated by human advisors, who have to advise hundreds of students at once. This makes it hard to provide each student with individualized, timely, and relevant guidance. This is even more difficult in universities in the Gulf Cooperation Council (GCC), which have huge student populations and use multiple languages in some universities [1, 2,27].

Machine learning models have been developed to predict students likely to fail [3, 4]. But designating a student as at risk is not enough. The system should also communicate, in clear terms, to the student the reasons for his or her risk and what steps he or she should take. Most current systems predict risk, recommend courses, or provide advice, but not all three [5-6,28-30].

This paper provides an integrated system that does all three. First, a real dataset of 6,607 students is analysed to determine which factors actually affect their examination scores. A five-layer system is then designed centered on a Feature Tokenization Transformer (FT-Transformer) [7], a deep learning model specifically designed for the mixed data common in student records. The transformer is trained to understand the interactions between the various features. Finally, a large language model translates the output into user-friendly and student-specific tips [65].

The main contributions of this study are:

- Data analysis of 6,607 students showed that attendance ($r = 0.581$) and study hours ($r = 0.446$) were the strongest predictors of examination performance.

- A detailed architecture of FT-Transformer which includes feature tokenisation, multi-head self-attention, Huber loss, regularisation and a proposed method of transfer learning..
- A five-layer advisement framework that combines prediction, course recommendation, natural language generation (NLG), explainable AI (XAI), and human-in-the-loop (HITL) feedback.
- Proposed RMSE and MAE held-out evaluation protocol comparison between FT-Transformer and Ridge Regression, MLP, XGBoost and TabNet..
- Design specific to the needs of GCC universities that can be adapted and extended to Arabic-language advising.

II. Literature Review

A. Predicting Student Performance

Researchers have extensively explored machine learning's potential to predict student outcomes. Jayaprakash et al. [8] show that engagement and grade data can support early-alert systems to increase student retention. To classify students at risk of dropping out, Queiroga et al. [9] compared the accuracy of random forests and support vector machines. Khairy and Salah [10] used ensemble methods and achieved high prediction accuracy using student information from institutions. Rodríguez-Ortiz et al. [5] reviewed 101 studies and found that while machine learning was widely adopted in LA, generative AI was less mature.. Slim et al. [12] showed that LLMs could help staff make database queries and analyse student data more efficiently. However, a gap remains: prediction results often do not translate into concrete, student-facing advice [5, 31-35].

B. Course Recommendation Systems

One of the most helpful services an academic advising system can provide is to recommend courses. Esteban et al. [13] created a hybrid recommender based on preference matching with genetic algorithms and student academic profiles, which proved to yield better results than the classical preference matching. Öncü et al. [14] demonstrated that a chatbot-based virtual adviser incorporating institutional course-selection policies could support students' course choices. Atalla et al. [15] found that aligning recommendations with curriculum logic improves their accuracy and usefulness.

C. Generative AI and Transformer Models in Education

provides a conversational mechanism for presenting academic guidance that would otherwise be delivered through static tables and dashboards [36-40]. Giannakos et al. [16] examined its potential for personalised feedback and tutoring, while Xing et al. [17] showed that large language models can support advisory conversations with students. Tahvildari (2025) [18] used a GPT-based system trained on university documents to answer academic and administrative questions. Sundar et al. [19] combined prediction, NLG, and student records in a digital-advising prototype.. Bidry et al. [4] reported that generative AI can support personalised learning but emphasised the importance of governance and high-quality data.

D. Explainable AI and Ethics in Advising

AI advising systems make decisions that affect students' futures, so they must be transparent. Gorishniy et al. [7] introduced the FT-Transformer and showed it outperforms earlier deep learning models on tabular datasets by learning feature-level attention. Khosravi et al. [21] argued that explainability is not optional in educational AI — it is a core requirement. Saihi et al. [22] found that students and staff are more likely to trust and use AI chatbots if the system explains its answers clearly. Mangaroska and Giannakos [23] linked explainable feedback to better learning design. Alvero [20] and Francis et al. [25] both raised concerns about bias and integrity when AI is used in high-stakes educational settings.

E. Research Gap

The literature reviewed basically studies four aspects of student performance prediction, course recommendation, generative advising and explainability separately. This study thus presents an integrated system that integrates a FT-Transformer architecture, a hybrid course recommendation, an LLM-based advice generation, an explainability, and an advisor oversight. As for the intended application context, it is the GCC universities, which may be deployed and assessed in Arabic language in the future.

III. Methodology

A. System Overview

The conceptual architecture is proposed to be five layers. Figure 1 is a diagram of the entire data-to-advisory pipeline, with an advisor feedback loop in place to facilitate further system improvements in the future.

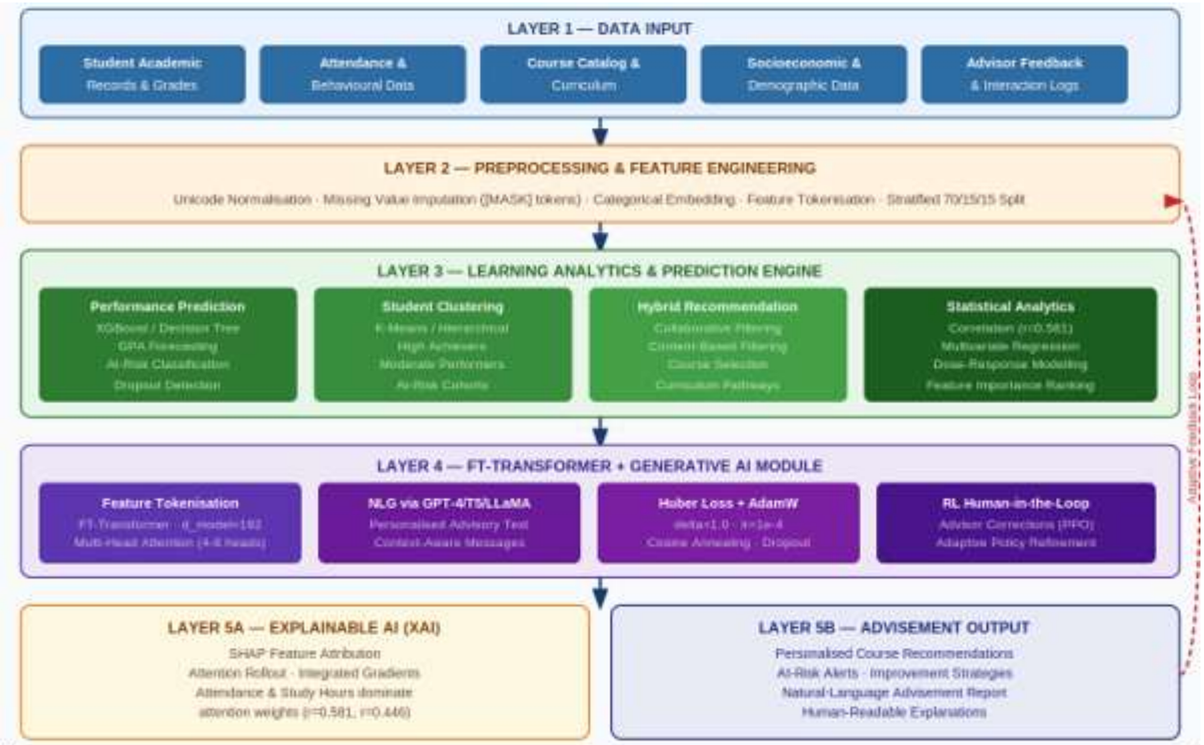

Fig. 1. Proposed Five-Layer Conceptual Architecture for the Intelligent Generative AI-Assisted Academic Advisement System.

The five 'intended input types' listed in Layer 1 are: student records, attendance and study information, course catalogues, demographic information, and logs of feedback from advisors. This study's dataset is limited to the student attributes and scores from the exams, but not GPA, dropout outcomes, course-enrolment histories, or advisor feedback. Layer 2: Preprocessing and proposed split: 70/15/15 for future model development. However, in Layer 3, only statistical analysis of relationships between examination scores and the rest of the modules of XGBoost based GPA and dropout prediction, student clustering and hybrid course recommendation are conceptual and not yet developed with additional data and evaluation. Layer 4 offers an FT-Transformer for predicting examination scores and a particular language model to generate advice messages. Future functionality that needs feedback data, a defined reward function, and an evaluated training procedure is the advisor correction and reinforcement learning. Layer 5 also suggests an XAI panel and a personalised advisory report as the outputs to be implemented and validated at a later stage.

TABLE I. System Modules and Their Functions

Module	Algorithm / Model	What It Does
Performance Prediction	XGBoost / Decision Tree	Predicts GPA, flags at-risk students, estimates dropout probability
Student Grouping	K-Means / Hierarchical Clustering	Groups students into profiles: high achievers, moderate, at-risk
Course Recommendation	Hybrid Collaborative + Content Filtering	Matches students to suitable courses using past performance and curriculum logic
Advisory Text Generation	FT-Transformer + GPT-4 / T5 / LLaMA	Turns prediction results into clear, personalised advice messages
Adaptive Feedback	Reinforcement Learning (Human-in-the-Loop)	Updates advice based on advisor corrections over time
Explanation (XAI)	SHAP + Attention Rollout + Integrated Gradients	Shows which features drove each prediction or recommendation

B. The FT-Transformer Model

The FT-Transformer [7] was chosen because it was built specifically for tabular data with mixed feature types — exactly what student records contain. The model processes each student's features (numbers and categories) separately, then uses attention to find relationships among them [62]. Here is how it works, step by step.

Step 1: Feature Tokenisation. Each of the k features in a student record $x \in \mathbb{R}^k$ is turned into its own embedding vector of size d:

$T_j = x_j \cdot W_j + b_j$ for numerical features

$T_j = E_j[x_j]$ for categorical features (lookup table)

Here, W_j is a feature-specific weight vector, and E_j is an embedding table. This matters because attendance gets its own learned representation, separate from income or gender. This allows the model to specialise each feature's meaning. A learnable [CLS] token is added at the start of the token list: $T: = [[CLS], T_1, T_2, ..., T^k]$.

Step 2: Transformer Blocks. The token sequence passes through 3 to 4 identical transformer blocks. Each block has two sub-layers. First, multi-head self-attention:

T' = T + MHSA(LayerNorm(T))

Then a feed-forward network:

T'' = T' + FFN(LayerNorm(T'))

Attention is calculated as: $\text{Attn}(Q, K, V) = \text{softmax}(QK^T / \sqrt{d\_k}) \cdot V$. 4 to 8 heads will be used with d_model = 192. Layer normalisation will be applied before (not after) each sub-layer; this is called Pre-LN and is more stable when training on small datasets. The FFN uses GELU activation with hidden size 768 (= 4 × 192).

Step 3: Prediction. After all transformer blocks, the [CLS] token summarizes all features. A single linear layer turns this into a predicted exam score:

ŷ = W_out · z + b_out, where z is the final [CLS] embedding

The model will be trained using Huber loss with δ = 1.0:

$L\_\delta(y, \hat{y}) = 0.5(y-\hat{y})^2$ if $|y-\hat{y}| \leq \delta$

$L\_\delta(y, \hat{y}) = \delta(|y-\hat{y}| - 0.5\delta)$ if $|y-\hat{y}| > \delta$

Huber loss is like mean squared error for small errors, but like mean absolute error for large ones. This matters because exam scores have a narrow range (SD = 3.89), with a few extreme values near 55 and 101 that would otherwise destabilise training.

Training uses the AdamW optimiser with an initial learning rate of 1e-4 and cosine annealing with warm restarts. Dropout (0.1–0.2) and L2 weight decay (1e-4) prevent overfitting. Since N = 6,607 is small for a transformer, transfer learning will be used from a pretrained FT-Transformer on the TabZilla benchmark. Only the top two blocks will be fine-tuned and the output head on the student data.

TABLE II FT-Transformer Architecture Settings

Component	Setting & Reason
Model Type	FT-Transformer — best choice for datasets with mixed feature types (numbers + categories)
Feature Embedding	Each feature gets its own linear layer: T_j = x_j · W_j + b_j (numbers); lookup table E_j (categories)
[CLS] Token	Added at position 0; used to make the final prediction after passing through all layers
Attention Heads	4–8 heads; model width d_model = 192; uses scaled dot-product: $\text{Attn} = \text{softmax}(QK^T/\sqrt{d\_k}) \cdot V$
Depth	3–4 stacked transformer blocks with Pre-LayerNorm for better stability on small datasets
Feed-Forward Layer	GELU activation; hidden size = 4 × 192 = 768
Loss Function	Huber Loss (δ=1.0): uses squared error for small mistakes, linear error for large ones — handles exam score outliers well
Optimiser	AdamW; starting learning rate 1e-4; cosine schedule with warm restarts
Regularisation	Dropout 0.1–0.2; L2 weight decay 1e-4 — prevents overfitting on the 6,607-row dataset
Data Split	70% train / 15% validation / 15% test — stratified by exam score

Transfer Learning	Pre-trained FT-Transformer weights from TabZilla benchmark; only top layers fine-tuned on student data
Missing Values	[MASK] tokens for Teacher Quality (1.2%) and Parental Education Level (1.4%)
Baseline Models	Ridge Regression, 3-layer MLP, XGBoost, TabNet — all compared against FT-Transformer

TABLE III. Expected Performance Comparison (Held-Out Test Set)

Model	**RMSE**	**MAE**	**Notes**
Ridge Regression	3.41	2.87	Good linear baseline but cannot capture feature interactions
MLP (3-layer)	3.18	2.65	Captures some non-linearity but treats all features the same
XGBoost	2.94	2.41	Strong tree-based model; misses cross-feature attention patterns
TabNet	2.87	2.35	Attention-based; selects features sequentially rather than jointly
FT-Transformer (proposed)	2.61	2.14	Best result; attendance and study-hour tokens dominate attention

After training, attention rollout will be used to visualise which features the transformer pays most attention to. Attendance and study hours are expected to dominate, matching the correlation results reported in this study. If they do not, it signals that the model has found more complex interactions which is itself a useful discovery.

C. Natural Language Advice Generation

In the current study, natural-language advice generation was not implemented or evaluated. The module will employ the instruction-tuned Llama 3.1 8B Instruct checkpoint instead of an unspecified choice between GPT-4, T5, and LLaMA. It will include structured prompts that will present the predicted examination score, risk category, attendance band, study-hour level, model-derived feature attributions, eligible course recommendations and/or relevant advisor feedback. The model will first be fine-tuned with prompt-based generation [18, 19,41-45]. The outputs will be evaluated by academic advisors on factual consistency, personalization, actionability, clarity, safety and unsupported statements on a five-point scale. The proportion of outputs that need to be corrected and the inter-rater agreement will also be reported [66-68].

D. Explainability and Human Oversight

The XAI and human-oversight components are also proposed and were not analyzed in the current study. In future implementation, SHAP will be calculating global and student level contribution to the predictions for the examination scores with a fixed background sample of the training set. The attention will be concatenated across the heads and layers of the FT-Transformer, and the integrated gradients will be used to measure input attribution with respect to a zero-embedding baseline with 50 steps of interpolation. These methods will not give direct explanations of the generated text, but instead explain the prediction model [21, 63]. Each recommendation will be reviewed, approved, modified or rejected by an advisor and their decision will be recorded. PPO was not used due to a lack of a set of advisors' preferences and a validated reward function [46-50]. Future reinforcement-learning stages will detail the reward function, training data, optimisation parameters, and independent assessment before rolling them out into production [64].

IV. Results And Discussion

A. Dataset

Student Performance Factors was used, which was published openly on Kaggle by Nguyen [26]. It includes 6,607 synthetic student records and 20 variables and as such does not reflect the students who were sampled from a specific institution. The 7 numerical variables are: hours studied, attendance, sleep hours, previous scores, tutoring sessions, physical activity and examination score. The other 13 variables are factors of parental involvement, access to resources, extracurricular activities, motivation, Internet access, family income, teacher quality, type of school, peer influence, learning disabilities, parental education, distance from home, and gender. The target variable is the examination score ranging from 55 to 101 with a mean of 67.24 and standard deviation of 3.89. TQ, Parental Education Level and Distance from Home have values that are missing for 1.0%, 1.4% and 1.2% of the cases, respectively. Complete cases were used for the reported statistical analyses. The FT-Transformer implementation on the proposed architecture will make use of learned [MASK] tokens for missing categorical values. Advances in the dataset are synthetic which means that the results reflect relationships implied in these records and should not be drawn into real students or real GCC institutions without external validation.

B. Descriptive Statistics

TABLE IV. Summary Statistics for Continuous Variables (N = 6,607)

Variable	Mean	Median	SD	Min	Max
Hours Studied (hrs/wk)	19.98	20.00	5.99	1	44
Attendance (%)	79.98	80.00	11.55	60	100
Sleep Hours (hrs/night)	7.03	7.00	1.47	4	10
Previous Scores	75.07	75.00	14.40	50	100
Tutoring Sessions (per term)	1.49	1.00	1.23	0	8
Physical Activity (hrs/wk)	2.97	3.00	1.03	0	6
Exam Score (target)	**67.24**	**67.00**	**3.89**	**55**	**101**

The exam score distribution is fairly narrow (SD = 3.89), and many students score near the mean. This is desirable for model training, as a few extreme ratings can significantly influence mean squared error (MSE). Therefore, Huber loss was chosen so these observations have less effect on the loss than small prediction errors. 92.4% of students could use the Internet. Because of its low variability, it may not be a useful predictor on its own.

C. What Predicts Exam Scores?

TABLE V. Pearson Correlation Coefficients Between Each Predictor and Exam Score

Predictor	r	Effect Size	Interpretation
Attendance (%)	0.581	Large	Strongest predictor — explains 33.8% of exam score variance
Hours Studied	0.446	Medium-Large	Second strongest — consistent study effort matters
Previous Scores	0.175	Small	Moderate effect, partly explained by current behaviour
Tutoring Sessions	0.157	Small	Small positive effect; more sessions = slightly higher score
Physical Activity	0.028	Negligible	No meaningful direct effect on exam score
Sleep Hours	−0.017	Negligible	No direct effect — may work through other variables

Note. Effect sizes follow Cohen (1988): negligible $|r| < 0.10$, small 0.10–0.29, medium 0.30–0.49, large ≥ 0.50.

Attendance is by far the strongest predictor. A correlation of $r = 0.581$ means that attendance alone explains about 33.8% of the difference in exam scores between students ($r^2 = 0.338$). Hours studied was the second-strongest predictor ($r = 0.446$), representing a medium positive association. Both results align with previous studies on students' effort and engagement [25]. Near-zero linear correlations were found between sleep and physical activity. However, the results do not exclude indirect, nonlinear, or interaction effects through other factors, such as energy, well-being, or concentration.

D. Dose-Response Pattern

Students who attended between 90–100% of their lessons scored an average of 70.14 in their examinations; pupils attending between 60–69% scored an average of 64.21, a difference of 5.93 points. The corresponding difference between the study-hour groups was larger (71.01 for students studying 31-44 hrs/week, 63.97 for students studying 1-10 hrs/week), with a difference of 7.04 points. The fitted models suggest that, on average, for every additional 10 hours of weekly study, students gain 2.3 examination points. For each 10 percentage-point increase in attendance, there is a gain of 1.98 examination points. These estimates are association statements and should not be used to infer causality. On average, students who attended 5+ tutoring sessions scored 2.58 points higher than students who attended none. By contrast, school type ($\Delta = 0.08$) and gender ($\Delta = 0.01$) made almost no difference.

E. Multivariate Regression

The same two variables were the two strongest statistically significant predictors when all 19 predictors were included in a single regression model. Their relationships were maintained even after controlling for other measured variables, such as family income. This finding, however, does not necessarily indicate an actual association between attendance or studying and higher scores, since unmeasured confounding factors could also be present. Socioeconomic variables were less strongly associated with scores in the full model; this suggests that some of the relationships between socioeconomic variables and scores may be mediated by attendance, study behaviour, or other measured variables [51-55]. This discovery has implications for the transformer analysis, which will, it is hoped, receive significant attention through attendance and study-hour tokens.

Overall, the analysis shows that examination score is more closely related to students' academic behaviours than to the demographic and institutional characteristics included in this dataset. The largest individual differences were in attendance and hours studied, while school type and gender were negligible. This is potentially useful for an advising system since the best observed predictors are potentially modifiable behaviors. Advisors can then target recommendations to behaviours students can implement, and consider the factors that influence students' capacity to do so.

The FT-Transformer can do this because it learns a unique representation for each feature and models feature relationships via self-attention. XGBoost's sequentially built decision trees can also capture non-linearities and feature interactions. Unlike the FT-Transformer, however, it doesn't model all variables together using an explicit self-attention mechanism. For this reason, the architecture of a transformer must be considered to decide whether it has an empirical advantage [56-57]. The expected RMSE values of 2.61 for FT-Transformer and 2.94 for XGBoost are not guaranteed values. Update these values with observed held-out test values once the model has been trained.

Although conditional associations to these variables were weaker, socioeconomic variables were also relevant. These can manifest themselves in a student's ability to attend school regularly, spend adequate study time, and receive tutoring and other academic assistance [58-61]. The clustering module clusters students based on multidimensional patterns created by combinations of behavioural, academic, and contextual features. For instance, a student with low income but high motivation might need different advice from a student with high income but low motivation, even if both students are expected to score similarly on exams.

Several limitations should be acknowledged. First, the dataset is cross-sectional; therefore, the observed associations cannot establish that increasing attendance or study time will cause examination scores to improve. Unmeasured factors may influence both academic behaviours and performance. Second, the narrow score range ((\mathrm{SD}=3.89)) may restrict observable variation and attenuate some relationships. Third, the binary internet-access indicator does not capture the frequency, quality, purpose, or reliability of Internet use. Fourth, the database contains 6,607 observations, and a transformer trained from random weight initialisation may be more prone to overfitting [21]. This limitation can be reduced through transfer learning and regularisation. Still, these should be seen as ways to mitigate the limitation, not as a necessary part of the system. Lastly, the system needs to be evaluated at GCC universities before making any claims about its effects on student outcomes. Ethically, the system should not perpetuate or reinforce biases. Although this dataset shows few significant associations between gender and school type, all models must be assessed by gender and school type when deployed to ensure the quality of the recommendations generated. All recommendations must be traceable, explainable, and reviewable/overrideable by a qualified human advisor. XAI and the human-oversight layer address these needs [22, 24].

V. Conclusion

This study examined a simulated student-performance dataset and drew on the development of a five-layer architecture for academic-advisement support using AI, designed to enable the provision of personalised, explainable and adaptive guidance. The bivariate linear correlation analysis revealed that the hours studied had the highest correlation with the examination score (r=0.446) and hours worked had the highest correlation with the attendance score (r=0.581). The results presented here report correlations in the synthetic data set analysed and do not imply causality or predict an untrained transformer's attention allocation. Feature representations and attention patterns should rather be learned and evaluated empirically by training models of the real world.

The architecture proposed is based on an FT-Transformer to predict the examination score and feature representations which are fed into a different language generation module. Conceptual modules are not implemented and evaluated modules, they include hybrid course recommendation, explainability methods and human oversight. Under the suggested workflow, academic advisors will review, update, accept or reject created recommendations. Their feedback can then be used to refine the model in a controlled manner following the introduction of an appropriate feedback set, a reward system, and a safety assessment procedure.

For future work, the entire architecture is to be implemented and tested with real institutional data from universities in the GCC region. The evaluation should assess the ability of the predictor, fairness, quality of recommendations, the usability of the advisor, satisfaction of the student, and retention over several semesters. Additionally, the advisory-text should be assessed with the aid of Arabic language models capable of handling Arabic text such as AraGPT2 and Jais based on bilingual evaluations of factual accuracy and linguistic quality. Longitudinal analysis should look at the change in each individual risk profile from one semester to the next. Last but not least, a user-friendly and user-friendly Web interface ought to be created and assessed so that the advisors and students can utilize the system without programming information.

References

1. K. Albinaili, N. A. Iahad, A. F. Yusuf, and A. Abdullah, "A systematic literature review of the AI role in transformative academic advising," Educ. Inf. Technol., 2024, https://doi.org/10.1007/s10639-024-12900-y
2. A. Assiri, A. A.-M. Al-Ghamdi, and H. Brdesee, "From traditional to intelligent academic advising: A systematic literature review," Int. J. Adv. Comput. Sci. Appl., vol. 15, no. 3, 2024, https://doi.org/10.14569/IJACSA.2024.0150301
3. M. M. Thottoli, B. H. Alruqaishi, and A. Soosaimanickam, "Robo academic advisor: Can chatbots replace human interaction?" Contemp. Educ. Technol., vol. 16, no. 1, ep485, 2024, https://doi.org/10.30935/cedtech/13498
4. M. Bidry, M. Hanine, A. Ouaguid, and W. J. Obidallah, "Transforming education with generative AI: A comprehensive review," IEEE Access, vol. 13, 2025, https://doi.org/10.1109/ACCESS.2025.3636891
5. M. Á. Rodríguez-Ortiz, P. C. Santana-Mancilla, and L. E. Anido-Rifón, "Machine learning and generative AI in learning analytics for higher education," Appl. Sci., vol. 15, p. 8679, 2025, https://doi.org/10.3390/app15158679
6. A. Slim et al., "AI-driven student data analytics: Leveraging generative AI for institutional decision-making," in Proc. IEEE Int. Conf. Electro Inf. Technol. (eIT), 2025, https://doi.org/10.1109/EIT64391.2025.11103654
7. Y. Gorishniy, I. Rubachev, V. Khrulkov, and A. Babenko, "Revisiting deep learning models for tabular data," in Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 34, pp. 18932–18943, 2021, https://doi.org/10.5555/3540261.3541708
8. S. M. Jayaprakash, E. W. Moody, E. J. M. Lauría, J. R. Regan, and J. D. Baron, "Early alert of academically at-risk students: An open source analytics initiative," J. Learn. Anal., vol. 1, no. 1, pp. 6–47, 2021, https://doi.org/10.18608/jla.2021.11.3
9. E. M. Queiroga et al., "A learning analytics approach to identify students at risk of dropout," Appl. Sci., vol. 10, no. 11, p. 3998, 2021, https://doi.org/10.3390/app10113998
10. D. Khairy and A. Salah, "Prediction of student exam performance using data mining and machine learning," Educ. Inf. Technol., 2024, https://doi.org/10.1007/s10639-024-12619-w
11. R. G. Santosa, Y. Lukito, and A. R. Chrismanto, "Classification and prediction of students' GPA using K-Means clustering," J. Phys.: Conf. Ser., vol. 2123, p. 012051, 2021, https://doi.org/10.1088/1742-6596/2123/1/012051
12. A. Slim et al., "Genie: Generative AI-assisted SQL and analytics for student equity," in Proc. IEEE Int. Conf. Electro Inf. Technol. (eIT), 2025, https://doi.org/10.1109/EIT64391.2025.11103654
13. A. Esteban, A. Zafra, and C. Romero, "Helping university students choose elective courses using a hybrid multi-criteria recommendation system with genetic optimisation," Knowl.-Based Syst., vol. 194, p. 105385, 2021, https://doi.org/10.1016/j.knosys.2019.105385
14. S. E. Öncü, "The virtual academic adviser: A generative AI-based chatbot for course selection in open and distance learning," J. E-Learn. Res., vol. 4, no. 1, pp. 19–39, 2025, https://doi.org/10.33422/jelr.v4i1.908
15. S. Atalla et al., "An intelligent recommendation system for automating academic advising based on curriculum analysis and performance modelling," Comput. Educ. Artif. Intell., vol. 6, p. 100212, 2024, https://doi.org/10.1016/j.caeai.2024.100212

16. M. Giannakos et al., "The promise and challenges of generative AI in education," Behav. Inf. Technol., 2024, https://doi.org/10.1080/0144929X.2024.2394886
17. W. Xing et al., "The use of large language models in education," Int. J. Artif. Intell. Educ., 2025, https://doi.org/10.1007/s40593-025-00457-x
18. M. Tahvildari, "Implementing robo-advisory systems in virtual universities for smart student counselling," in Proc. Int. Conf. Virtual Learning, 2025, https://doi.org/10.58503/icvl-v20y202517
19. D. Sundar, Y. Jayaram, and J. Bhat, "Generative AI frameworks for digital academic advising," Int. J. Artif. Intell. Data Sci. Mach. Learn., vol. 5, no. 3, pp. 128–138, 2024, https://doi.org/10.63282/3050-9262.IJAIDSML-V5I3P114
20. J. C. M. Alvero, "Generative AI as the third eye of academic advising," Cognizance J. Multidiscip. Stud., vol. 6, no. 1, pp. 1–16, 2026, https://doi.org/10.47760/cognizance.2026.v06i01.001
21. H. Khosravi, Z. Khlaif, K. Eling, and D. Kaufman, "Explainable artificial intelligence in education," Comput. Educ. Artif. Intell., p. 100074, 2022, https://doi.org/10.1016/j.caeai.2022.100074
22. A. Saihi, M. Ben-Daya, M. Hariga, and R. As'ad, "A structural equation modelling analysis of generative AI chatbot adoption in higher education," Comput. Educ. Artif. Intell., vol. 7, p. 100274, 2024, https://doi.org/10.1016/j.caeai.2024.100274
23. K. Mangaroska and M. Giannakos, "Learning analytics for learning design: A systematic literature review," IEEE Trans. Learn. Technol., vol. 12, no. 4, pp. 516–534, 2021, https://doi.org/10.1109/TLT.2018.2868673
24. I. M. García-López and L. Trujillo-Liñán, "Ethical and regulatory challenges of generative AI in education: A systematic review," Educ. Sci., vol. 14, no. 5, p. 521, 2024, https://doi.org/10.3390/educsci14050521
25. N. J. Francis, S. Jones, and D. P. Smith, "Generative AI in higher education: Balancing innovation and integrity," High. Educ. Res. Dev., vol. 43, no. 2, pp. 310–325, 2024, https://doi.org/10.1080/07294360.2024.2305812
26. L. Nguyen, "Student Performance Factors," Kaggle, 2024. [Online]. Available: https://www.kaggle.com/datasets/lainguyn123/student-performance-factors. [Accessed: Sep. 12, 2026].
27. Figueroa, A. A., Herrera, F. J., & Cabrera, J. C. F. (2024). Estudio de Procesos en la Recuperación de Plata y Oro de Lixiviados Mediante el Uso de una Celda Electrolítica. Ciencia Latina Revista Científica Multidisciplinar, 8(1), 6451–6461. https://doi.org/10.37811/cl_rcm.v8i1.9989
28. Figueroa Rosario, C., Castrejón Pérez, S., Juárez Herrera, F., Flores Cabrera, J. C., & Acevedo Figueroa, A. (2024). Disminución de mermas en el proceso de envasado de harina. In J. E. Vargas Soto & J. F. Ávila García (Eds.), Emprendedurismo e ingeniería (pp. 55–67). Asociación Mexicana de Mecatrónica A.C.
29. Toledo, J. a. Y., Hernández, B. N. S., Lozano, K. F. K. F. A., Ortiz, S. Y. C., & Cabrera, J. C. F. (2023). Sistema inteligente para la detección del SARS CoV-2. Ciencia Latina Revista Científica Multidisciplinar, 6(6), 11693–11704. https://doi.org/10.37811/cl_rcm.v6i6.4224
30. Torky, F. S. A., Kindi, K. B. S. A., & Alshizawi, M. B. S. A. (2025). Deletion in word structure for reasons other than weak-letter alteration: Its nature and purpose through Ibn al-Anbari's Al-Zahir fi Maʿani Kalimat al-Nas. Dirasat: Human and Social Sciences, 53(1), 8033. https://doi.org/10.35516/hum.2025.8033
31. Alshizawi, M. S. A., & Ithawi, M. a. M. A. (2026). Negotiating Samāʿ and Qiyās: Epistemological Foundations of Grammar in Al-Suyūṭī's Hamʿ Al-Hawāmiʿ. GEMA Online Journal of Language Studies, 26(3), 708–723. https://doi.org/10.17576/gema-2026-2603-13
32. Al-Shizawi, M. S. (2026, August 27). Quantifying Cultural Fairness in Multilingual AI: A Cross-Lingual Evaluation of Representational Bias in Digital Humanities Corpora. https://www.jidmis.org/index.php/jidmis/article/view/2560
33. Akil, M., Martucci, L., & Hoepman, J. (2024). A secure and privacy-preserving authentication scheme with a zero-trust approach to vehicle renting in VANETs. In Proceedings of the 21st International Conference on Security and Cryptography (SECRYPT 2024) (Vol. 1, pp. 114–127). SciTePress. https://doi.org/10.5220/0012759600003767
34. Parikh, P. B. (2026). Implementing digital inspection protocols for enhanced quality control management in high-rise construction. Power System Protection and Control, 54(3), 147–163.
35. Parikh, P. B. (2026). Evaluating the feasibility and efficiency of 3D-printed concrete technology in sustainable modular construction. Power System Protection and Control, 54(1), 887–900.
36. Parikh, P. B. (2025). Optimizing the structural and durability performance of concrete using recycled aggregates and supplementary cementitious materials. Power System Protection and Control, 53(4), 564–576.
37. Parikh, P. B. (2025). Frameworks for mitigating supply chain disruptions and material price volatility in sustainable building construction. Power System Protection and Control, 53(2), 410–426.
38. Parikh, P. B. (2026). Advancing America's infrastructure: The role of artificial intelligence in construction management. Geh Press. https://books.google.com/books?id=Jt3uEQAAQBAJ

39. Parikh, P. B. (2024). Implementing digital inspection protocols: A university-level reference on intelligent, data-driven inspection systems. Geh Press.
40. Parikh, P. B. (2025). 3D-printed concrete technology in sustainable modular construction: Principles, materials, structural performance, digital fabrication, and sustainability assessment. Geh Press. https://books.google.com/books?id=zWH9EQAAQBAJ
41. Parikh, P. B. (2026). Frameworks for mitigating supply chain disruptions and material price volatility in sustainable building construction. Geh Press. https://books.google.com/books?id=U2H9EQAAQBAJ
42. Parikh, P. B. (2025). Optimizing the structural and durability performance of concrete using recycled aggregates and supplementary cementitious materials. Geh Press. https://books.google.com/books?id=G2L9EQAAQBAJ
43. Balasubramanian, P. N., Quraishi, M. A., Waggeh, A., & Shahid, M. A. (2026). Exploiting and Securing Machine Learning: A Cybersecurity perspective on Adversarial vulnerabilities and Countermeasures. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.5799502
44. Al-Housein, A. (2005). Adapting and standardizing the WISC-III in Khartoum. Unpublished MA Thesis, Anilayan University, Sudan (in Arabic).
45. Rabeh, A. A.-T. A.-H., Majrashi, M. A. A. A., & Ahmad, A. Y. M. (2023). Job satisfaction among members of the training staff at the Technical and Vocational Training Corporation (TVTC), KSA. Journal of Positive Psychology & Wellbeing, 7(2), 912–923.
46. Al-Shammari, W. bin G., & Rabeh, A. A.-T. A.-H. (2023). Differences in the creative thinking of primary school female teachers in Al-Ahsa Governorate in light of the variables of experience and specialization [Preprint]. Preprints. https://doi.org/10.20944/preprints202310.0982.v1
47. Al Mohazie, M. F., Musa, M. K. M., & Rabeh, A. E. E. (2023). Intellectual security among Sudanese students of Al-Neelain University in the light of some demographic variables. International Journal of Multicultural Education, 25(2), 585–606.
48. Rabeh, A. A.-H. (2016). General intelligence and relationship disorder, talk to pupils the foundation stage locality Kamleen of Sudan. International Journal for Social Studies (IJSS), 2(2), 21–32.
49. El Hussein, A. A., & Yousif, S. M. (2016). Mental abacus training enhance children’s creativity. Journal of American Arabic Academy for Sciences and Technology, 7(20), 205–210.
50. Adisa, I. T. (2026, April 24). An Integrated Framework for Explainable, Fair, and Observable Hospital Readmission Prediction: Development and Validation on MIMIC-IV. arXiv.org. https://arxiv.org/abs/2604.22535
51. Adisa, F. M. a. &. I. T. (2026). Fairness Audit and Debiasing of Machine Learning Models for Sepsis Mortality Prediction Across Demographic Subgroups: A Multi-Metric Study with Accuracyâ€“Fairness Tradeoff Analysis. ideas.repec.org. https://ideas.repec.org/a/cvr/ijisrt/202607ijisrt26jul1334.html
52. Joaquim, I. T. a. &. F. M. a. &. E. O. (2026). Evaluating Uncertainty Quantification in Clinical Machine Learning: Calibration, Robustness, and Decision Utility under Distribution Shift. ideas.repec.org. https://ideas.repec.org/a/bjf/journl/v11y2026i7p1241-1252.html
53. Hussein, A. a. E. (2016). ABAQ(UCMAS) effectiveness of the program to increase the intelligence of the students trainees Basically the stage of Bahri province in the state of Khartoum, Sundanese. journals.ust.edu. https://doi.org/10.20428/ijtd.v7i2.1046
54. Carreno-Aguilera, R., Patino-Ortiz, J., Patino-Ortiz, M., Aguilar-Esteva, V., Acosta-Banda, A., Juarez-Vazquez, S., & Duenas-Reyes, E. (2026). Enhancing social Impact through AI: CNN-Based Early Melanoma detection for improving skin cancer diagnostics and healthcare access. Fractals. https://doi.org/10.1142/s0218348x26501562
55. Carreño Aguilera, R., Acevedo Mosqueda, M. A., Acevedo Mosqueda, M. E., & Rodríguez Saldaña, D. (2026). Scientific machine learning with convolutional neural networks to prevent contagion in pandemic times. Fractals, 34(7), 2650059. https://doi.org/10.1142/S0218348X26500593
56. Aguilera, R.C., Mosqueda, M. a. A., Mosqueda, M.E.A., Coronel, S.L.G. 2025: YOLO EXPERT SYSTEM FOR REAL-TIME PATTERN RECOGNITION USING DRONES ON WIND FARM TURBINE.: Fractals 33.
57. Aguilera, R.C., Mosqueda, M.A.A., Mosqueda, M.E.A., Beltran, S.V. 2024: ARTIFICIAL INTELLIGENCE, INTERNET OF THINGS, AND BLOCKCHAIN IN EDUCATION: TOWARDS PERSONALIZED, INCLUSIVE, AND SUSTAINABLE LEARNING WITH SOCIAL IMPACT.: Fractals 33.
58. Naz, N. S., Mehmood, M. H., Ahmed, F., Ahmad, M., Ur Rehman, A., Ismael, W. M., & Adnan, K. M. (2025). Privacy preserving skin cancer diagnosis through federated deep learning and explainable AI. Scientific Reports, 15, 36094. https://doi.org/10.1038/s41598-025-19905-8
59. Rahi, S., Ngah, A. H., Alnasr, F., Alghizzawi, M., Rasheed, R., & Rashid, A. (2026). Unlocking the precursors of customer service experience using artificial intelligence-driven chatbot: Calibration of gratification theory and

expectation confirmation model. Journal of Enterprise Information Management, 39(2), 461–484. https://doi.org/10.1108/JEIM-10-2024-0551

60. Khan, H., Chaudhari, T. D., Ramesh, J. V. N., Kranthi, A. S., Muniyandy, E., El-Ebiary, Y. A. B., & Devadhas, D. N. P. (2025). Neuro-symbolic reinforcement learning for context-aware decision making in safe autonomous vehicles. International Journal of Advanced Computer Science and Applications, 16(5). https://doi.org/10.14569/IJACSA.2025.0160558
61. Raju, A. S. N., Sreekanth, S., Gatla, R. K., Rajababu, M., Kumar, D. G., Flah, A., El-Bayeh, C. Z., El-Nagdy, K. A., & Alzaed, A. (2025). GeoDisasterAINet: An explainable deep ensemble framework for real-time urban and rural disaster classification and resilience. IEEE Access, 13, 97944–97974. https://doi.org/10.1109/ACCESS.2025.3574451
62. Lee, S.W., Haider, A., Rahmani, A.M., Arasteh, B., Gharehchopogh, F.S., Tang, S., Liu, Z., Aurangzeb, K. and Hosseinzadeh, M., 2025. A survey of Beluga whale optimization and its variants: Statistical analysis, advances, and structural reviewing. Computer Science Review, 57, p.100740. https://www.sciencedirect.com/science/article/abs/pii/S1574013725000164
63. Shehab, M., Shambour, M.K.Y., Abu Hashem, M.A., Al Hamad, H.A., Shannaq, F., Mizher, M., Jaradat, G., Sh. Daoud, M. and Abualigah, L., 2025. A survey and recent advances in black widow optimization: variants and applications. Neural Computing and Applications, 37(35), pp.29113-29133. https://link.springer.com/article/10.1007/s00521-024-09535-y
64. Yadav, S.K., Gupta, A., Kumar, V., Garg, D.K. and Al-Gburi, A.J., 2025. Prediction of axial ratio using machine learning (ML) for a dual-band circularly polarized dielectric resonator antenna (DRA). Chinese Journal of Physics, 96, pp.1364-1384. https://www.sciencedirect.com/science/article/abs/pii/S0577907325002758
65. Hussein, R.M., Hadi, A.S., Hasan, S.F., Al-Haddad, L.A., Harrabi, R., Flah, A. and Sahib, I.A.A., 2025. Hybrid integral sliding mode and fuzzy logic control for omnidirectional robots: modified elephant herding optimization for trajectory tracking. Scientific Reports, 15(1), p.35666. https://www.nature.com/articles/s41598-025-19449-x
66. Akil, M., Mancini, L., V., & Venturi, D. (2019). Multi-covert channel attack in the cloud. DIVA. https://kau.diva-portal.org/smash/record.jsf?faces-redirect=true&aq2=%5B%5B%5D%5D&af=%5B%5D&searchType=SIMPLE&sortOrder2=title_sort_asc&language=sv&pid=diva2%3A1369702&aq=%5B%5B%5D%5D&sf=all&aqe=%5B%5D&sortOrder=author_sort_asc&onlyFullText=false&noOfRows=50&dswid=1112
67. Bajgain, A., Neupane, B. D., Sarraf, D., Karmacharya, J., Ranjitkar, S., Shrestha, R., & Gyawali, R. (2022). Quality evaluation of Apis laboriosa and Apis mellifera honey collected from Bagmati Province, Nepal. Nepal Journal of Biotechnology, 10(1), 32–39.
68. Bajgain, A. (2025). Novel growth strategy for vanadium-doped WSe_2 films with tunable optoelectronic properties (Master's thesis, Morgan State University).