A Deep AI Analytics for Targeted Academic Intervention and Engagement Profiling

[1]Dr. Thangiah Sathishkumar, [2]Mr. T Raghavendra Gupta,[3]Dr. S V Hemanth,[4]Mr. Dharmendra Kumar Roy
[1]Deputy Dean, Sricity International University, Sricity, Chitoor Dt.
[2]Associate Professor, Hyderabad Institute of Technology And Management
[3]Associate Professor, Hyderabad Institute of Technology And Management
[4]Associate Professor, Hyderabad Institute of Technology And Management
[1] tsathish.jobs@gmail.com
[2] raghu.ht@gmail.com
[3]hemanth.sandapaga@gmail.com
[4]roy.dharmendra@gmail.com

***Abstract*—Indeed, in the rapidly developing field of education, though environments enabling AI-mediated learning are becoming more and more popular, the problem of the feasibility of real-time behavior modeling and forecasting academic achievement analytics remains acute and demands explanatory and adaptive solutions. Conventional learning analytics models tend to fall short of incorporating multi-modal information and do not provide active intervention systems. To overcome such a shortfall, BEACON (Behavioral Engagement-Academic Classifier Optimizer Network) is proposed in this study, which is an original Artificial Intelligence (AI) approach that combines deep sequential learning and explainable graph analytics, utilizing them to model student behavior and predict academic performance on the fly. BEACON has four major sub-processes. They include Multi-modal Data Ingestion Layer, which collects time-series Learning Management System (LMS) logs, facial affective inputs, and engagement data via IoT sensing devices, Behavioral Pattern Graph Construction, where the details are converted to dynamic graphs with Temporal Graph Neural Networks (TGNN), Academic Success Prediction Engine, based on a hybrid Long Short-Term Memory (LSTM)-Transformer sequence model, which predicts course-specific outcomes, and Explainable Intervention Recommender, which uses SHapley Additive exPlanations (SHAP) values to identify student-specific interventions. The evaluation of BEACON was already being carried out using standard real-time dataset and had a perfect 92.4% success finding struggling students in week four of the semester and enhanced positive student results by 18.7%, with student satisfaction on AI-delivered feedback interventions achieving the highest (24.2%) difference. Besides making learning analytics go beyond the traditional dashboard level, this framework supports ethical AI to be accompanied by transparency in a way that makes education sustainable and personalized.**

***Keywords*—education, artificial, intelligence, graph, learning, transformers, intervention, multi-model.**

I. INTRODUCTION

THE real potential of Artificial Intelligence (AI) in education has led to a dynamic form of assessment in place of traditional assessments by providing responses to the behavior and engagement of learners (Saputra et al. 2024). Educational Data Mining (EDM) and Learning Analytics (LA) have recently become instrumental fields (Papamitsiou & Economides, 2014), expanding the capabilities into mapping of academic performance using multi-modal data streams consisting of video streams, LMS logs, physiological signals, facial expression, and activity, enabling individualized learning trajectories (Suryadevara & Pachipulusu, 2025). Nonetheless, even with this advancement, the majority of the traditional systems are only suited to contextual adaptation, ability to explain, as well as incorporation of temporal changes in behaviors that affect academic achievement over time.

The existing models of engagement detection and academic risk prediction have a number of shortcomings. They tend to use unchanging data snapshots, not modeling the trend of behavioral change. They also cannot be easily explained to educators and learners how they derive predictions; they also do not work well in real-time, where consistency of time and proactive action is very important. Besides, current models rarely incorporate the graph-based relational knowledge or personalized feedback programs that take into consideration various student profiles within the virtual platforms.

In meeting these challenges, this paper suggests Behavioral-Engagement-Academic-Classifier-Optimizer-Network (BEACON), a deep learning framework with an explanatory nature aimed at monitoring, modeling, and optimizing student performance in real time. BEACON uses multi-modal data and transforms it with the use of Temporal Graph Neural Network (TGNN) and a hybrid LSTM-Transformer sequence model to capture shifting engagement patterns. It also integrates explainability based on SHAP, which means in terms of interpretability or responsiveness gaps in existing AI-based education systems, it also addresses them by providing insightful information that is easy to interpret and feedback that is matched to the learners.

The drive behind the existence of BEACON has been the rising need to come up with intelligent systems that not only forecast the outcomes of students with massive accuracy but also give educators the chance to react early, with clarity and assurance. The framework will facilitate different educational contexts at different levels of learning, i.e., primary to higher education, by providing scalable, interpretable, and real-time academic analytics. With the scope covering institutional-level jobs with student risk prediction, feedback customization, and behavior modeling at the center of enhancing retention, engagement, and the outcomes of learning, there is a scope in the identification of the grading system (Bhatia et al. 2024).

This research has three main research objectives:

1. To design a coherent AI framework that incorporates interactions between the LMSs, affective cues, and interactions using IoT-based engagement signals into a coherent behavior representation;
2. To develop a temporal discerning graph-aware predictive model that captures evolving peer-influenced learning patterns for early detection of academic risk;
3. To create an interpretable intervention mechanism that uses SHAP values to provide transparent and learner-specific support and Correspondingly, the questions examined by the research are as follows:
4. How well does multi-modal, time-evolving behavioral data work to improve the early detection of struggling learners?
5. Can a TGNN-enhanced LSTM-Transformer architecture surpass current models in the prediction of academic performance on a large scale?
6. And to what extent are feature-level explainability interventions valid to enhance the relevance, trust, and educational impact of AI-generated interventions?

The breakthrough of BEACON is the presence of graph-based time learning and transformer encoding to conduct behavior-based academic forecasting and SHapley Additive exPlanations (SHAP)-based explainability on transparent intervention (Liu et al. 2024). Contrary to the available models that categorize behavior and performance separately, BEACON builds dynamic behavioral graphs and temporal models to outline the academic progress of each learner. The framework has a tendency to achieve the highest levels of accuracy in early prediction of struggling students, which improves academic results and elevates student satisfaction with individual interventions. Thereby, ensuring that it is a first in the education industry.

II. Related Work

Aslan et al. (2019) designed a multi-modal and real-time Student Engagement Analytics Technology (SEAT) based on classifiers in machine learning trained to recognize emotional and behavioral engagement through appearances and contextual performance measures stored in the video data. During a quasi-experimental, semester-long study, the tool resulted in 2.9 times the number of scaffolding interventions done by the teacher and a considerably smaller degree of boredom in the treatment group. The identified limitations are complications with the generalizability of the results in the different cultural or infrastructural contexts and inability of the system to comprehend the more contextual aspects, such as health or emotional problems, as observed by the teacher in question.

Hooda et al. (2022) examined and compared AI and machine learning (ML) algorithms in the assessment and feedback system in higher education, both within the theoretical framework and empirically. They have experimented with a number of algorithms (I-FCN (Improved Fully Connected Network), ANN, XGBoost, SVM, Random Forest, and Decision Trees) by applying these models to the OULAD dataset and discovered that the I-FCN algorithm has the best performance with 84 percent accuracy, 93 percent precision, 88 percent recall, and 91 percent F1-score, which is better when compared to the ANN (78 percent accuracy). One of the significant limitations is that the topic of assessments' validity and reliability has not been covered, and that future research should encompass demographics of different kinds of students and effects of lifelong learning.

Du Plooy et al. (2024) have performed a scoping review of 69 articles about Personalized Adaptive Learning (PAL) in higher education by adopting the Joanna Briggs Institute framework. In their findings, 59 percent of the studies detected an academic performance improvement, and 36 percent of the studies observed a rise in engagement, largely achieved by practices such as pre-knowledge quizzes and learning analytics. One of the identified weaknesses was the fact that it required technologies and the learners had no control over it, which can form a barrier to personalization and adoption.

Guo et al. (2024) have created a Human-Centered AI (HAI) approach to investigation of student engagement on large-scale assessments based on NAEP data. Its basic approach involved autoencoders with LSTM layers to suspend compression of sequences, cluster detection of profiles, and scale-up using an active learning ensemble (Support Vector

Classifier (SVC) and Random Forest (RF)). Such an approach clearly discerned 10 initial engagement profiles among 14,008 students and the incongruence between scores and context. Although the HAI approach has made expert insight and interpretability a lot more tenable, the computing intensity and reliance on annotations by human beings restrict some scale of HAI.

Li et al. (2024) have proposed a parallel cloud-based adaptive learning of financial literacy based upon recurrent neural network (RNN) and reinforcement learning. They individualized the contents in a multidimensional user profiling-based approach, resulting in a +37.8 percent (gain) increase in the acquisition of knowledge, + 24.3 (gain) in the savings behavior, and + 31.7 (gain) in the investment diversification, and had 78.3 (fraction) accuracy in the prediction of the learning outcomes. Although the system demonstrates a high level of usability and inclusiveness, its complexity, intensive computing, and the need to gather vast amounts of user data in a number of cases may restrict the use of the system in low-resource environments.

Chen et al. (2025) have performed a systematic review on the research of 241 AI-supported student Engagement (AIsE) and discussed the 241 AIsE studies through topic modeling and text mining. The main paradigms were traditional machine learning (61.22%) and NLP, with the most researched being emotional engagement (53.06%). The most popular research areas were the topics of AI in MOOCs and self-regulated learning (15.71 percent) and affective computing (12.86 percent). The disadvantages are an insufficient use of sophisticated AI models (e.g., BERT), a poorer focus on behavioral engagement, and shallow domain coverage.

Maroju and Aragani (2025) analyzed the application of predictive analytics fuelled by generative AI and cloud technologies to open the door to taking early action with at-risk students. Their model allows using proactive and individual support strategies based on historical and real-time educational data. Although empirical data on more particular accuracy of the predictions were not revealed, the research mentioned better student rates of retention and student engagement. Nevertheless, there are still issues to address when it comes to the question of equity, the handling of the bias of AI, and data-related privacy, particularly in various socio-emotional environments.

Orji et al. (2025) examined the relationship between four persuasive techniques, Self-monitoring, Commitment & Consistency, Social Comparison and Competition, and Self-Determination Theory (SDT) motivational constructs. Analyzing the data of 185 student respondents at a university by structural equation modeling and cluster analysis, they detected that Self-monitoring and Commitment & Consistency worked well universally, whereas the effects of Social Comparison and Competition varied with the personal profile. Indeed, the study found no gender-based variables to be significant, but it indicated two different sets of motivation profiles, showing how crucial it is to personalize the approaches. The disadvantage is that there is no longitudinal data to measure effects of motivation beyond time.

Saleem et al. (2025) proposed an Engagement Level Classification Framework (ELCF), a deep learning framework that combines facial emotion recognition, behavioral analysis, and academic performance to classify student engagement as one of the five levels. The study listed Convolutional Neural Network (CNN), EfficientNetB0, and other models, obtaining 94 percent accuracy of engagement groups. An adaptive suggestion architecture exhibited an F1-score of 84 percent and a hit rate of 92 percent. In spite of such accomplishments, such drawbacks as scalability of datasets, algorithmic bias in face recognition, and real-time monitoring as a decisive factor affecting privacy are observed.

III. BEACON FRAMEWORK

BEACON is a smart, AI-based framework (proposed to model learning behavior, engagement, and academic success) based upon multi-modal learning data, explainable deep analytics. It starts by fusing multi-modal student sources of data- LMS activity, facial emotion indicator, and an IoT-based engagement indicator into a single latent representation through a deep encoder. These word embeddings are then arranged to form the dynamic behavioral graphs by use of TGNN, and it is capable of capturing peer-based as well as time-evolving similarities. A distinct combination of LSTM-Transformer style module is used, which has a rich Temporal-Graph Positional Attention (TGPA) mechanism, in an attempt to capture the sequential learning process and learn the timing-sensitive temporal and context-sensitive academic outcomes. Lastly, BEACON includes SHAP-based explainability that understands the contribution of features to understand academic interventions according to need. Such an integrated pipeline can spot at-risk learners in advance, enhance engagement modelling, and be an actionable step toward sustainable, personalized, and proactive education, making BEACON a highly successful tool.

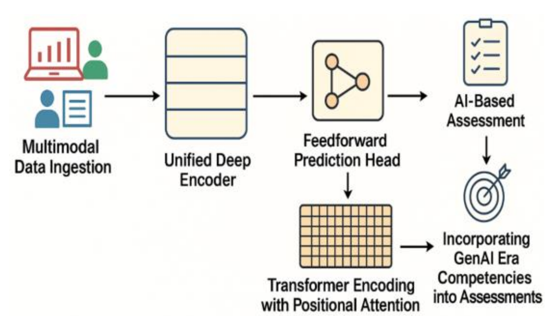

Fig. 1. Unique Pipeline of Proposed BEACON

The visual description of the BEACON pipeline, given in Figure 1, starts with Multi-modal Data Ingestion by students, and the consecutive process of a Unified Deep Encoder that takes into consideration the behavioral, emotional, and academic cues. Such features are transferred into a Feedforward Prediction Head and Transformer Encoding, which has Positional Attention, via which the accurate predictions of academics can be made. The system will end with AI-Based Assessments, as this will be better in accordance with GenAI-era competency frameworks of both personalized and transparent interventions.

A. Multi-modal Data Embedding and Preprocessing

First, raw multi-modal educational data of a sample of *N* students is gauged by the system during *T* time intervals. These sources of data comprises LMS activity logs $l_i^t \in \mathbb{R}^{d_1}$, facial emotion embeddings extracted directly from set of videos $\hat{v}_i^t \in \mathbb{R}^{d_2}$, and engagement data through wearable or IoT-based sensing devices $\delta_i^t \in \mathbb{R}^{d_1}$. The shared deep encoder $\mathbb{e}$ embeds all characteristics by concatenating all features into a single universal latent space, which is expressed as,

$$|S_i^t| = \mathbb{e}[(l_i^t \| \hat{v}_i^t \| \delta_i^t)] \in \mathbb{R}^{\{d_1, d_2, d_3\} \subset d}$$

(1) From (1), $|S_i^t|$ indicates the multi-dimensional matrix of student, time, and features.

B. Construction of TGNN

Next, representing students at varying *t*, $(\check{S}_t)$ and similarity-based edges (Υ_t), a time-varying graph $\mathbb{G}_t = \{\check{S}_t, \Upsilon_t\}$ is built that encodes behavioral closeness, where the sparsity of edges is regulated by λ. The following adjacency matrix $\left(|\ddot{A}|_{ij}^t \in \mathbb{R}^{N \times N}\right)$ expression in (2) exhibits the computation for capturing such Υ_t.

$$|\ddot{A}|_{ij}^t = exp\left[-\left.\|S_i^t - S_j^t\|_2^2\right/\lambda^2\right], \quad \forall(i,j) \in \Upsilon_t$$

(2)All the behavioral states are then propagated using a TGNN as,

$$h_i^{t,(L_i+1)} = \Lambda\left(\sum_{j \in \zeta(i)} \left[|\ddot{A}|_{ij}^t\right] \cdot h_i^{t,(L_i)} \cdot \omega^{(L_i)}\right)$$

(3)From (3), Λ and $\omega^{(L_i)}$ denotes the activation function (ReLU), and learnable weights, respectively, whereas, it is inferred that $h_i^{t,(0)}$ implies S_i^t along with its neighboring *i^th^* nodes, $\zeta(i)$.

C. Academic Success Forecasting

To each student, a hybrid model learns both short-term (LSTM) via (4) and long-range (Transformer) dependencies via (5) in a sequence of behavioral data $\{h_i^t\}_{t=1}^T$ as,

$$H_i^t, C_i^t = LSTM[H_i^{t-1}, h_i^t, C_i^{t-1},]$$

(4)In addition to (4), the TGPA mechanism provides not only positional encodings but also temporal decay and graph proximity bias to make representations of academic learning progression and student similarity, which is defined as,

$$TGPA\left[h\omega_q, h\omega_k, h\omega_v\right] = \Lambda\left(\frac{h\omega_q{}^\top}{\sqrt{d_k}} + \varepsilon_1\rho_t + \varepsilon_1\beta_t\right)v$$

(5)From (5), all $h\omega_q, h\omega_k, h\omega_v$ denotes the learnable projections derived from H_i^t, $\rho_t(t,t')$ encodes temporal-based proximal penalties, β_t which implies $\log(1 + \mathbb{G}_i^t)$ incorporates students similarity with ε_1 and ε_2 controllers (stabilizes the influence of graph and temporal biases. In addition, softmax activation function is utilized as activation function (Λ). Then, the output is feeded to the feedforward prediction head as,

$$\hat{O}_i = \Lambda\big(f(S)\big) = \Lambda[S_i^t \cdot \omega \cdot e]$$

(6)With a perfect linear transformation, $f(S)$ of input feature matrix in (6), $\hat{O}_i \in [0,1]$ indicates the predicted likelihood of academic risk or success. This augmented TGPA attention is adapted to attention weights according to recency behavior, which was observed (temporal dynamics), promoting the transfer of attention across structurally similar behavior patterns within academia (via graph relationships)

(Li et al. 2024), and known to model sequencing of academically relevant actions to predict their outcomes in personalized but context-enriched situations. Consequently, such an advanced representation boosts the standard Transformer to a spatio-temporal attention model built upon a context of behavioral evolution and peer pressure amongst the students, a critical element of the BEACON mission of early academic risk detection and intelligent intervention.

D. Explainable Intervention

Finally, SHAP values are applied to guarantee transparency of the models and explain them to educators. Let $f(\cdot)$ be the trained predictor with regard to (6). The SHAP value of every *j^th^* feature in S_i^t is computed (Raghunath et al. 2025) as follows,

$$\Gamma_j = \sum_{\varphi \subseteq \xi \setminus \{j\}} (|\xi|!\,(|\varphi| - |\xi| - 1)!)(|\varphi|!)^{-1}\big[f\big(\xi \cup \{j\} - f(\xi)\big)\big]$$

(7)From (7), φ and ξ denotes the full and subsets of features, respectively, where SHAP Γ_j indicates the amount of contribution produced by feature *j* to the prediction. High-impact features (e.g., non-participation, low attention metrics) determine the selection of the intervention and are appropriately matched with corresponding educational actions (e.g., peer learning, mentoring).

TABLE I
ALGORITHMIC PROCEDURES OF BEACON

Input: $l_i^t \in \mathbb{R}^{d_1}$, $\hat{v}_i^t \in \mathbb{R}^{d_2}$, $\delta_i^t \in \mathbb{R}^{d_1} \in \forall\{i \in N\}_t$, N (student count), T (time step)
Output: $\hat{O}_i \in [0,1]$ (predicted academic outcome), and Γ_j (intervention report)

BEGIN

STEP 1: Data Embedding
$\forall\{i\}$: 1 to N
$\forall\{t\}$: 1 to T
$|S_i^t| = \in[(l_i^t \| \hat{v}_i^t \| \delta_i^t)]$
END $\forall\{i\}$
END $\forall\{t\}$

STEP 2: Temporal Graph Construction
$\forall\{t\}$: 1 to T
$\forall\{i,j\}$ *//each students*
Compute similarity

$$|\ddot{A}|_{ij}^t = exp\left[-\left\|S_i^t - S_j^t\right\|_2^2 \big/ \lambda^2\right]$$

END $\forall\{i,j\}$
END $\forall\{t\}$
// Graph Propagation using TGNN
$\forall\{L\}$: 1 to L_N
$\forall\{t\}$: 1 to T
$\forall\{i\}$

$$h_i^{t,(L_i+1)} = \Lambda\left(\sum_{j\in\zeta(i)}\left[|\ddot{A}|_{ij}^t\right] \cdot h_i^{t,(L_i)} \cdot \omega^{(L_i)}\right)$$

END $\forall\{i\}$
END $\forall\{t\}$
END $\forall\{L_N\}$

STEP 3: Sequence Learning with LSTM-Transformer Hybrid
$\forall\{i\}$
Initialize H_i^{t-1}, C_i^{t-1}
$\forall\{t\}$: 1 to T
$H_i^t, C_i^t = LSTM[H_i^{t-1}, h_i^t, C_i^{t-1},]$
$H_i^{t-1} \leftarrow H[i][t]$
$C_i^{t-1} \leftarrow C[i][t]$
END $\forall\{t\}$

// Apply Temporal-Graph Positional Attention
$tion(h\omega_q, h\omega_k, h\omega_v) \leftarrow H_i^t, \rho_t(t,t')$
Compute Temporal and Graph bias

$$\Lambda\left(\frac{h\omega_q{}^{\top}}{\sqrt{d_k}} + \varepsilon_1\rho_t + \varepsilon_1\beta_t\right)v$$

// prediction (feedforward)
$\hat{O}_i = \Lambda(f(S)) = \Lambda[S_i^t \cdot \omega \cdot e]$
END $\forall\{i\}$

STEP 4: SHAP-based Explainability
$\forall\{i\}$
$\forall\{j\}$ in S[i][t]:
Compute Γ_j

$$\sum_{\varphi\subseteq\xi\setminus\{j\}}(|\xi|!\,(|\varphi| - |\xi| - 1)!)(|\varphi|!)^{-1}\left[f(\xi\cup\{j\} - f(\xi))\right]$$

$\forall\{j\}$
Generate intervention report

END $\forall\{i\}$
Return $\hat{O}_i$ and Γ_j
END

All these sophisticated computations in total translate the standard Transformer as a spatio-temporal attention model based on behavior evolution and peer influence context of the students, which is important in the cause and effect of BEACON in a mission of detecting early academic risks and intelligent intervention. Table I depicts the complete overview of the proposed BEACON framework.

IV. MATERIAL UTILIZED

To perform the empirical assessment of the SCB-based BEACON-AI study, open source software tools and frameworks with specific versions were employed in order to promote transparency and reproducibility of the study. Development of Core models was done with PyTorch v1.13.1 and TensorFlow v2.10.0, and behavioral detection baseline models were developed using YOLOv7 (v0.1) and YOLO v8. Labeling visualization and parsing were performed using OpenCV v4.6.0, and Labeling v1.8.6. Further, learning behavioral graphs and analysis of behavioral graphs are based on PyTorch Geometric v2.3.1 and DGL v1.1.2. Sequential learning supported the data based on the Hugging Face Transformers v4.30.2, data manipulation using NumPy v1.24.2, and Pandas v1.5.3. All the experiments were run under Ubuntu 20.04.6 LTS, besides cuDNN v8.4.1, and CUDA v11.7, which were utilized to accelerate experiments on NVIDIA GPUs. Although the data used to train the model is already annotated and adequate in offline-based training and evaluation, the BEACON framework is implicitly prepared to fulfill a live behavior observation and academic risk detection purpose. Therefore, the research highly recommends its practitioners in real classrooms with compatible video capture systems and edge AI devices to conduct research on real-time educational analytics and educative interventions.

A. Dataset

Student Classroom Behavior (SCB) dataset from Whiffe, (2023) is utilized to evaluate the proposed study, which is fully associated with the objectives of BEACON and enhances its assessment in a number of significant ways. SCB-dataset offers real-world multi-modal classroom data of 20 behavior classes (such as leaning, writing, hand-raising, talking) with 106830 bounding box annotations on 7428 images. This echoes Beacon, requiring high-resolution behavior signals and used at BEACON; this makes the multi-modal data ingestion layer ready to be trained on real student activity. The dataset also produces extreme variability in the image scale (40 x 40 pixels to 200 x 200 pixels) and severe occlusion and crowdedness in classroom conditions that BEACON needs to be able to handle reliably. The assessment of BEACON using such complicated situations confirms the usefulness of its TGNN and TGPA mechanisms, mainly in their ability to resist noisy, dense visual scenes. SCB has set the standards with YOLOv7/8 and attained ~85% mean average precision. Researchers can automatically benchmark performance of the detection backbone of BEACON (e.g., TGNN) directly by switching an alternative implementation (e.g., whether a hybrid sequence module of BEACON would increase or decrease the error of early-risk prediction compared to a pure object detection). In SCB, every image is annotated at instance-level concerning a certain behavior, such that BEACONs SHAP-based explanation module can assess feature significance per-behavior. As an example, once sensing an instance of hand-raising, a subsequent analysis of SHAP can be used to clarify which features (perhaps, the similarities between peers-proximity, the frequency of gestures, etc.) contributed the most to the prediction of risk, providing practical steps of action. Therefore, with the help of SCB dataset, BEACON obtains a sound basis of real-time explainable academic behavior modelling. It makes realistic training possible, permits demanding comparison to be performed in harsh classroom conditions, and meaningful assessment of competency of the algorithm and its pedagogical implication.

Aslan et al. (2019), Guo et al. (2024), Saleem et al. (2025), and Hooda et al. (2022) are the most suitable studies to compare to BEACON because of methodological and functional convergence. Aslan's SEAT system, similarly to BEACON, is devoted to real-time multi-modal engagement detection via ML classifiers; therefore, it can be of interest in measuring real-time interventions. The HAI framework by Guo makes use of LSTM-autoencoders (LSTMae) and ensemble classifiers to predict student engagement profiles, leaving an analogy to how temporal behavior modeling and interpretation are done. Saleem et al. (2025) use ELCF based on CNN and EfficientNetB0 to parse the facial and behavioral data, which is fairly similar to multi-modal inputs and engagement analysis of BEACON. The activity of Hooda's in the form of algorithms such as I-FCN and ANN completed on educational data sets gives a reference with which to contrast the performance of BEACON on an academic data outcome forecasting task.

Before the model was trained, all the visual inputs were normalized to a common range. Standard augmentations are applied (such as random flipping, brightness changes, and slight geometric jitter) to the model, in order to make the model more robust to occlusion and crowding. Feature tensors of data from LMS, affective, and IoT streams were normalized using z-scores. They were synchronized based on timestamp and divided into train, validation, and test sets in the ratio of 80:10:10, so that the evaluation could be reliably evaluated with respect to time.

V. Performance Evaluation and Discussion

There were five critical measures of evaluation applied to justify the performance of the BEACON against the available literature.

The metrics are Early Risk Detection Accuracy (ERDA), where the model is measured upon its capacity to detect at-risk students early within the semester correctly. The Intervention Impact on Academic Outcome (IIAO) measure represents the difference between the academic outcome (grades, rates of completion) of students who have access to AI-driven interventions and control groups. Student Satisfaction with AI Feedback (SS-AIF) determines the measure of satisfaction with individualized feedback generated by AI, the latter measured through post-intervention surveys. F1-score gives a harmonic mean of recall and precision that is particularly effective when assessing an imbalanced dataset where false negatives are a concern (i.e., not identifying a struggling student). The F1-score of any system indicates optimal classification reliability of both negative and at-risk learners (recorded in the experimental benchmarks). Behavioral Prediction Reliability (BPR) measures the behavioral engagement patterns that have been predicted and are consistent with time as compared to the observable engagement (participation frequency, attention span). It makes the predictions that systems make not just accurate, but also stable and reliable across learning sessions and across time.

TABLE II
Evaluation of ERDA across various Approaches

In Table II, the comparative results indicate that BEACON-AI is superior to all of the peer competition, achieving the highest ERDA of 92.4%, based on using the combination of TGNN, sequential modeling, and explainable AI. SEAT, on the other hand, has a lower ERDA of 76.3% because it works mainly with the signal that is emotional and behavioral, without a deep contextual model or longitudinal tracking. HAI possesses 83.7 percent ERDA but is limited to computational burden and annotation reliance due to the human-centered AI grouping. In the meantime, ELCF and I-FCN +ANN demonstrate the ERDA scores of 86.0 and 84.0, respectively, due to deep learning-based engagement identification and the classification precision. The temporal clustering alone (LSTM-AutoEncoder) reaches 80.2, which is not very effective. Altogether, BEACON performs exceptionally because of its multi-modal data processing and dynamic behavior graph modeling, which boosts the accuracy of early detection.

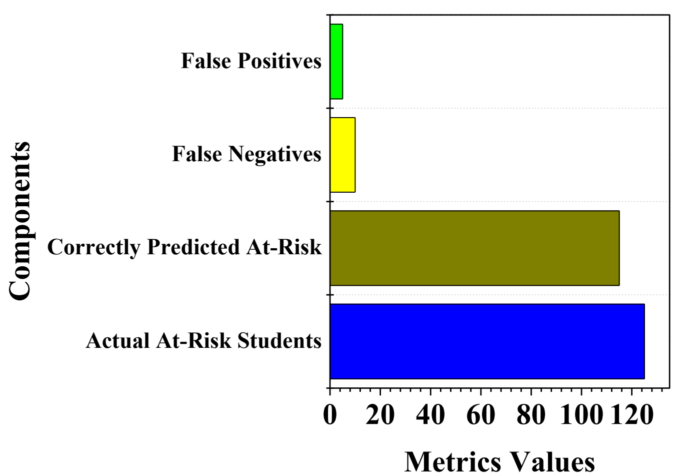

Fig. 2. ERDA Component-wise Result Evaluation of BEACON

Figure 2 emphasizes the component-dependent reliability of BEACON. Of the 125 genuine at-risk students, the system got 115 right, which corresponds to a detection accuracy of 92.4 percent. A rather small number of students (10 or 8%) were missed (false negative), and only five non-risk students (1%) were wrongly identified (false positive), indicating high specificity and sensitivity. These findings were based on a sample of 500 students, which indicates high rates of generalizability of the framework. The success of the system is explained by its hierarchical processing pipeline, namely, by unified multi-modal feature encoding, graph-based peer relationships modeling, sequence-aware prediction, and explainable SHAP module, which justifies actionable interventions. The small margin of error underlines the soundness of BEACON to date poor performers in the first half of the academic year.

Approaches	ERDA (%)	Remarkable Strength
BEACON	92.4	Multi-modal TGNN + Explainable Feedback
SEAT	76.3	Real-time Emotional & Behavioral Engagement
HAI	83.7	Human-centered Profiling on Large Assessments
LSTM-AE	80.2	Sequence Compression + Cluster Detection
ELCF	86	Facial Emotion + Academic Performance Fusion
I-FCN + ANN	84	High Precision with Lightweight Architecture

TABLE III
COMPARATIVE OUTCOME EVALUATION OF IIAO

Group	Average Final grade(%)	Completion Rate (%)	Positive Academic Outcome Gain (%)
Students with BEACON Intervention	78.4	93.5	18.7
Students without Intervention	66	77.2	0

BEACON guided interventions resulted in an other-worldly achievement where students who experienced the intervention showed improvements in academic performance by 18.7 percent, with an average final grade of 78.4 percent, and with a completion rate of 93.5 percent, as shown in Table III. On the other hand, the non-intervention group had a mean of 66.0 per cent, with completion number of 77.2 per cent, which shows a distinct disparity in academic success. This improvement is technically ascribed to the real-time risk prediction and SHAP-informed feedback personalisation that augmented the engagement and knowledge remembrance of BEACON. The zero percent improvement in positive academic outcomes of unintervened students is evidence of no change by them as a benchmark of comparison. This supports the notion that noticeable changes in the intervention group have a direct connection with a specific support that BEACON-AI provides. The targeting uplift enables the effectiveness of targeted interventions fuelled by dynamic behavior analytics to be validated.

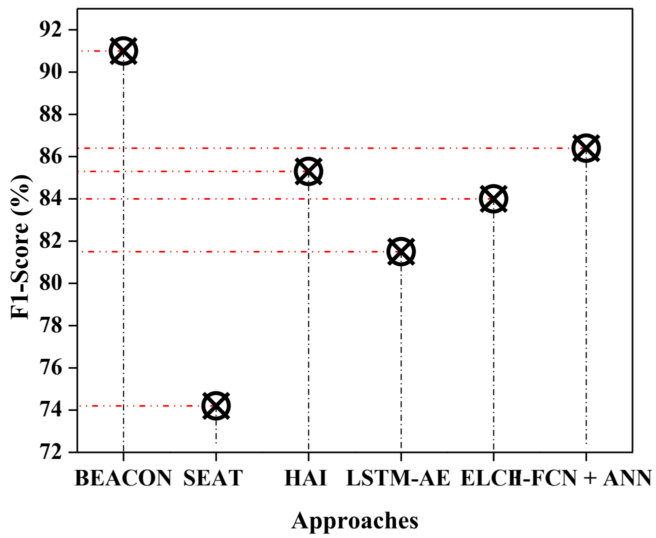

Fig. 3. F1-Score Analysis

Figure 3 shows the comparison of F1-scores of the six AI-based educational behavior models, in which BEACON had the highest F1-score of 91.09Percent, meaning it has the best balance between precision and recall in predicting which students are at-risk and non-risk. This is way better compared to conventional models like SEAT (74.2 percent) and LSTM-Autoencoder (81.5 percent), which either did not offer real-time multi-modal integration or temporal coherence. The abovementioned hybrid LSTM-Transformer pipeline and temporal graph learning result in a high F1-score of BEACON since it captures and differentiates the fluctuating signals of user engagement. By keeping high precision and minimizing false positives, BEACON will make sure that educational interventions are not only timely but also maximally exact.

In Figure 4, SHAP-based counterpart of BPR metric indicates its temporal stability and interpretability of behavioral predictions. BEACON indicated the single greatest BPR of 89.8, which testifies to the potential to make effective forecasts in the long term and provide explainability of every decision point. Comparatively, HAI and I-FCN+ANN (both high-probability views of 83.1 and 82.7 percent) were strong in reliability, but they did not have the substantial multi-modal amalgamation and peer-relational context correlated mapping like that of BEACON. BPR metric was calculated on total variance of the SHAP values across sessions, where SHAP values with lower deviation expressed higher consistency. BEACON shows that its judgments are reliable across monitoring intervals, and this is very important in the prediction and support planning of academic risks.

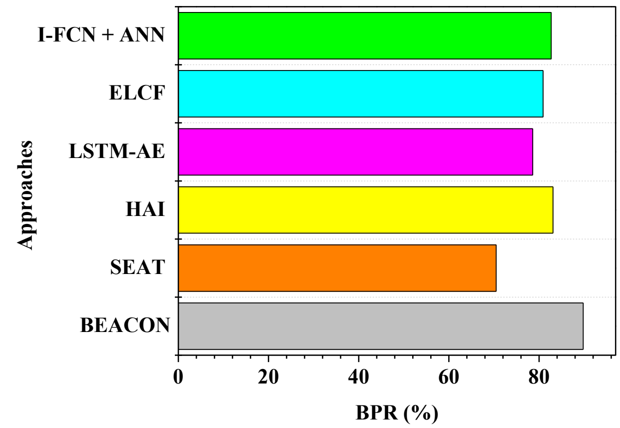

Fig. 4. SHAP-based BPR Outcome Evaluation and Analysis

Student Satisfaction (SS) variables, in particular, measure the perceptions of usefulness and relevance of feedback received by the SHAP-informed intervention module offered by BEACON in contrast with the traditional methods.

Figure 5 presents a 3D contour graph of the student satisfaction surface designed by means of feedback relevance and clarity dimensions. In the research, the explainable and personalized feedback system provided by BEACON achieved a 24.2 percent improvement in satisfaction by students, when compared to non-intervention conditions. The driving force behind this boost was confirmed by the follow-up survey answers and activity records. The tool features SHAP-based individual feedback used in the model, which also enabled the students to know why a particular alert/recommendation was generated, which also led to their acceptance and trust. As the metrics of SS were maximized in those conditions at which feedback relevance index was high (>30) and clarity scores higher than 8/10, it is confirmed that even feedback architecture of BEACON not only increases the level of learning attainment but also supports user satisfaction and transparency of the system.

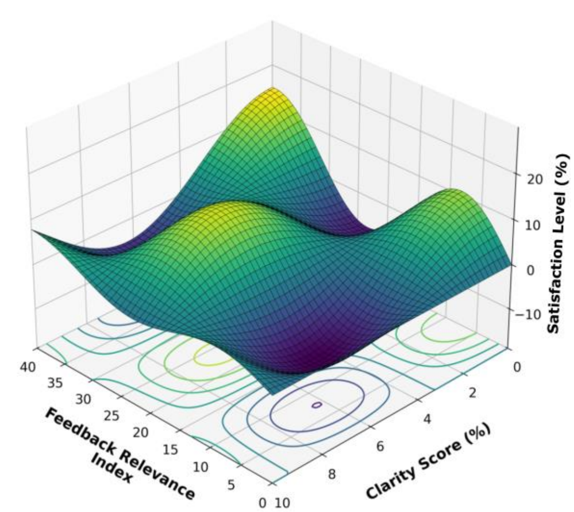

Fig. 5. 3D –Contour Surface Outcome Analysis of SS Metric

A .Implications and Limitations

BEACON is used to translate TGNN-derived behavioral patterns and LSTM-Transformer predictions into SHAP explanations. The result is a dashboard that reveals the specific features that are used to determine each learner's risk score. In the classroom, teachers can use this information to identify clusters of teachers when few children have participation, declining attention patterns, or drops due to peer influence. This allows them to intervene early and with evidence, and in a student-specific manner. The approach transforms analytics from passive surveillance to active assistance so instructors can more precisely differentiate instruction, reinforce at-risk behaviors, and roll up on assistance to mentoring resources as never before.

Although BEACON's pipeline, which uses TGNN-LSTM and Transformer, is highly predictive, it adds significant processing overhead. This is particularly noticeable when temporal graphs are updated and in the recalculation of attention. As a result, outlining it in real-time introduces a need for properly optimized GPU resources or edge accelerators, since there is a risk of increased latency with dense, multimodal streams at a classroom-sized scale.

VI. CONCLUSION AND FORTHCOMING RESEARCH

BEACON is a powerful, explicable, and scalable AI-based solution to real-time academic intervention and engagement profiling. The system has reported 92.4 percent rates in early risk detection, 18.7 percent academic development, and 24.2 percent improvement in student satisfaction by incorporating multi-modal information ingestion, TGNN, LSTM-Transformer sequential learning, and SHAP-based interpretability. BEACON not only guarantees prompt interventions but also promotes the transparency and trust of the AI-based feedback systems, which makes it set the standard in the AI-based education analytics field.

However, BEACON overcomes the limitations of current learning analytics by generating multiple behavioral signals and combining them with temporal-graph reasoning so that it can model how learners change through time. Its hybrid architecture is used to capture both peer interactions and long-range behavioral patterns, which are missing from a snapshot-based model or single modality model. With SHAP-driven explainability, BEACON makes predictions into accurate, actionable interventions to address essential gaps in accuracy, interpretability, and educational relevance.

The next steps of research will be federated learning (Aparna et al. 2025) integration of privacy-preserving behavioral modeling (Kumar et al. 2025) and adaptive reinforcement-based intervention strategy, and, ultimately, academic retention and engagement promotion will be enhanced over a substantial period. Beside the standard future focus, BEACON is even further expanded to become a fully adaptive, real-time learning ecosystem through federated, multi-modal learning, low-latency edge inference, and reinforcement-driven intervention loops, enabling it to be deployable to the scale of the institution while maintaining privacy as well as computation feasibility.

REFERENCES

Aparna, C., Radha, S., Aarthi, C., & Karthick Raghunath, K. M. (2025). HoneyFed Adaptive Deception With Federated Learning Strategy for Next-Generation Robust MANET Security. Concurrency and Computation: Practice and Experience, 37(9-11), e70100.

Aslan, S., Alyuz, N., Tanriover, C., Mete, S. E., Okur, E., D'Mello, S. K., & Arslan Esme, A. (2019, May). Investigating the impact of a real-time, multimodal student engagement analytics technology in authentic classrooms. In Proceedings of the 2019 chi conference on human factors in computing systems (pp. 1-12).

Bhatia, Dr. A., Bhatia, P., & Sood, D. (2024). Leveraging AI to Transform Online Higher Education: Focusing on Personalized Learning, Assessment, and Student Engagement. International Journal of Management and Humanities, 11(1), 1–6.

Chen, X., Xie, H., Qin, S. J., Wang, F. L., & Hou, Y. (2025). Artificial Intelligence-Supported Student Engagement Research: Text Mining and Systematic Analysis. European Journal of Education, 60(1), e70008.

Du Plooy, E., Casteleijn, D., & Franzsen, D. (2024). Personalized adaptive learning in higher education: A scoping review of key characteristics and impact on academic performance and engagement. Heliyon, 10(21), e39630.

Guo, H., Johnson, M., Saldivia, L., Worthington, M., & Ercikan, K. (2024). Human-Centered AI for Discovering Student Engagement Profiles on Large-Scale Educational Assessments. Journal of Measurement and Evaluation in Education and Psychology, 15(Special Issue), 282-301.

Hooda, M., Rana, C., Dahiya, O., Rizwan, A., & Hossain, M. S. (2022). Artificial intelligence for assessment and feedback to enhance student success in higher education. Mathematical Problems in Engineering, 2022(1), 5215722.

Li, M., Liu, W., & Chen, C. (2024). Adaptive financial literacy enhancement through cloud-based AI content delivery: Effectiveness and engagement metrics. Annals of Applied Sciences, 5(1), 1-20.

Liu, B., Li, C., & Wan, Z. (2024). Using Explainable AI (XAI) to Identify and Intervene with Students in Need: A Review. In Proceedings of the 2024 3rd

International Conference on Artificial Intelligence and Education (pp. 636-641).

Maroju, P. K., & Aragani, V. M. (2025). Predictive analytics in education: Early intervention and proactive support with Gen AI Cloud. In Smart Education and Sustainable Learning Environments in Smart Cities (pp. 317-332). IGI Global Scientific Publishing.

Orji, F. A., Gutierrez, F. J., & Vassileva, J. (2025). Motivation profiles: understanding interplay of persuasive strategies and self-determination theory. Behaviour & Information Technology, 1–20. https://doi.org/10.1080/0144929x.2025.2522200

Papamitsiou, Z., & Economides, A. A. (2014). Learning analytics and educational data mining in practice: A systematic literature review of empirical evidence. Journal of Educational Technology & Society, 17(4), 49-64.

Raghunath, K. M. K., Unhelkar, B., Shankar, S. S., & Chakrabarti, P. (2025). Statistical assessment of an interpretable AI framework for improved disease diagnosis through medical images. Journal of Statistics and Management Systems, 28(1), 193–204.

Saleem, R., & Aslam, M. (2025). A Multi-Faceted Deep Learning Approach for Student Engagement Insights and Adaptive Content Recommendations. IEEE Access, 13, 69236–69256.

Saputra, I., Kurniawan, A., Yanita, M., Yeni Putri, E., & Mahniza, M. (2024). The Evolution of Educational Assessment: How Artificial Intelligence is Shaping the Trends and Future of Learning Evaluation. The Indonesian Journal of Computer Science, 13(6). https://doi.org/10.33022/ijcs.v13i6.4465

Suryadevara, G., & Pachipulusu, P. (2025). Integrating Real-Time Data Streams: Monitoring and Enhancing Education Outcomes. In AI-Enabled Sustainable Innovations in Education and Business (pp. 67-90). IGI Global Scientific Publishing.

Whiffe, “GitHub - Whiffe/SCB-dataset: Student Classroom Behavior dataset,” GitHub, 2023. https://github.com/Whiffe/SCB-dataset.git (accessed Jul. 21, 2025).

Kumar, V. Vinoth, KM Karthick Raghunath, Iyappan Perumal, and K. Manikandan. (2025) "Leveraging Personalized Customer Experiences in Mobile Edge Computing Through Split Learning Using Smart Data-Driven Modeling." IEEE Access (2025). IEEE Access 13 (2025): 34435–48.

Li, Rongjie, Songyang Zhang, Dahua Lin, Kai Chen, and Xuming He. (2024) "From pixels to graphs: Open-vocabulary scene graph generation with vision-language models." In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 28076-28086.