European Journal of Clinical Pharmacy
Print ISSN: 2385-409X **Online ISSN**: Applied
Website: https://farmclin.es/

Original Research Article

A Hybrid Neuro-Symbolic Cognitive Reasoning Framework for Explainable Student Dropout Prediction and Personalized Academic Intervention

Article History:

Name of Author:
P. Revathy[1] and M. Priya[2]

Corresponding Author:
P. Revathy

How to cite this article:
P. Revathy and M. Priya. A Hybrid Neuro-Symbolic Cognitive Reasoning Framework for Explainable Student Dropout Prediction and Personalized Academic Intervention. *Eur J Clin Pharm.* 2025;7(1). 3932-3941.

***Received**:* 12-10-2025
***Revised**:* 05-11-2025
***Accepted**:* 20-11-2025
***Published**:* 30-12-2025

Abstract: Student dropout remains a critical challenge in modern education, directly impacting academic progress, institutional performance, and societal development. Existing machine learning and deep learning approaches demonstrate strong predictive capability but function largely as black-box systems, providing limited interpretability and failing to incorporate domain-specific educational reasoning. To address this gap, this paper proposes a novel Hybrid Neuro-Symbolic Cognitive Reasoning Framework (HNSCRF) that integrates data-driven neural learning with symbolic knowledge-based reasoning to deliver explainable dropout prediction and personalized academic intervention recommendations. The framework leverages multimodal data sources—including academic records, attendance history, student feedback sentiment, behavioral logs, and peer-interaction social graphs—processed through a neural inference model consisting of Transformer-based textual encoders, graph neural networks, and a multimodal attention fusion layer. A symbolic reasoning engine built on educational psychology principles and rule-based causal logic enforces decision consistency and provides transparent, human-understandable explanations. A Neuro-Symbolic Integration Layer aligns neural predictions with symbolic rules through a logical consistency loss function, ensuring reliability and interpretability. The system outputs risk level prediction, primary causal factors, and optimized intervention strategies such as counseling assignment, peer mentoring, workload restructuring, and emotional-support activities. Experimental evaluation on institutional and benchmark educational datasets demonstrates superior accuracy, fairness, and explainability over conventional deep learning models, significantly reducing false-negative rates and enhancing actionable insight delivery. The proposed framework supports educators and policymakers with trustworthy decision intelligence for proactive student retention and academic success.

Keywords: Student Dropout Prediction; Neuro-Symbolic Learning; Explainable AI (XAI); Multimodal Educational Data; Cognitive Reasoning; Personalized Intervention; Knowledge-Based Reasoning; Graph Neural Networks; Interpretable Machine Learning; Educational Analytics.

INTRODUCTION

Student dropout is a persistent and complex issue across global educational systems, affecting academic growth, efficiency, and institutional sustainability. Traditional analytical and machine learning approaches predominantly rely on structured academic indicators such as grades, attendance scores, and demographic attributes, resulting in limited predictive insight and delayed corrective action. Although recent multimodal and ensemble-based predictive models have enhanced accuracy by incorporating behavioral and emotional data, they still operate as opaque black-box systems with minimal transparency, limiting their acceptance and applicability in real educational decision environments.

The increasing demand for trustworthy and explainable AI in education highlights the necessity for systems that not only identify at-risk students but also articulate why a dropout risk exists and how it may be mitigated. This research addresses these limitations by integrating the strength of deep neural models with human-understandable symbolic reasoning grounded in educational psychology, cognitive behavior studies, and institutional policy rules.

This paper introduces a Hybrid Neuro-Symbolic Cognitive Reasoning Framework (HNSCRF) designed for interpretable dropout prediction and tailored academic intervention planning. The proposed system utilizes multimodal neural learning for risk inference and a symbolic rule engine constructed using logical causal constraints (IF–THEN reasoning) for transparent justification and personalized recommendations. A Neuro-Symbolic Integration Layer ensures decision consistency through a logical consistency loss optimization mechanism, enabling explainable predictions and responsible educational decisions. The major contributions of this research are summarized as follows:

- A novel Hybrid Neuro-Symbolic Cognitive Reasoning Framework (HNSCRF) that fuses neural predictive modeling with symbolic rule-based reasoning to enable interpretable dropout prediction and actionable intervention guidance.
- A multimodal neural inference architecture utilizing Transformer-based sentiment encoders, GNN-based peer influence modeling, and attention-driven feature fusion for comprehensive student behavioral understanding.
- A symbolic causal reasoning engine incorporating educational psychology rules, academic policies, and behavioral triggers to justify outcomes and generate personalized remedial actions.
- A Neuro-Symbolic Integration mechanism leveraging logical consistency loss to align model predictions with rule-based constraints, ensuring fairness, reliability, and transparency.
- Explainable output generation that provides dropout risk scoring, causal factor explanation, and recommended academic mentoring strategies to support data-driven decision-making.

RELATED WORKS

[1] Jiang (2025) employed a deep learning-based multimodal framework to predict student performance in physics education by integrating behavioral, academic, and interaction-level learning data. The model incorporated fusion-based representation learning to improve predictive insight. The multimodal design provided enhanced feature expressiveness by combining structured and unstructured inputs into a unified learning space. Performance evaluation indicated improvement in accuracy when compared to traditional data-driven systems. The approach demonstrated the importance of heterogeneous data modeling in educational analytics, particularly for real-time academic assessment.

[2] Wu (2025) introduced a prediction algorithm that utilized BERT-GCN multimodal fusion for academic performance estimation by combining semantic textual learning with relational learning. The architecture extracted contextual sentiment and engagement patterns from feedback text while modeling student peer connections using graph networks. The hybrid representation enhanced classification depth and contextual reasoning. Comparative analysis revealed a significant accuracy gain over conventional neural architectures.

[3] Pérez et al. (2025) explored machine learning for early dropout prediction in higher education using multi-dimensional student records. Gradient Boosting and Random Forest models were tested to forecast student withdrawal likelihood. Feature selection and interpretability techniques enhanced understanding of dropout determinants. Empirical validation on institutional datasets demonstrated reduction in false-negative classification.

[4] Xiong et al. (2025) developed a multimodal deep learning approach for early complication detection in Type-2 diabetes. The model incorporated longitudinal feature evolution and cross-modal representation learning. Temporal correlation modeling and multimodal alignment improved prediction precision. The design highlighted the capability of multimodal systems to capture latent behavioral and physiological patterns applicable to educational data scenarios.

[5] Li (2025) proposed a multi-source sensor fusion-based intelligent analysis model for student learning behavior monitoring. Multiple real-time data streams such as motion, activity logging, and environmental factors were incorporated for behavior modeling. Data quality enhancement and spatial-temporal reasoning improved learning pattern identification. The results emphasized the role of sensor-driven analytics for personalized learning support systems.

[6] Gao & Song (2025) applied a deep learning model for quality evaluation in elderly education based on intelligent education indicators. Multimodal content understanding and knowledge perception features were incorporated to support academic decision analytics. The study demonstrated scalability and robustness over traditional evaluation strategies.

[7] Hui et al. (2025) presented a systematic literature review on artificial intelligence in student

performance prediction within blended learning environments. The review categorized AI approaches including deep learning, ensemble learning, and multimodal fusion. Identified challenges included interpretability gaps, data sparsity, and ethical risk.

[8] Saleem & Aslam (2025) introduced a multi-faceted deep learning framework to examine engagement and adaptive content recommendation. The proposed approach integrated multimodal feedback and emotional indicators with feature-level fusion. Adaptive recommendation optimization improved personalized learning effectiveness.

[9] Chen & Chen (2025) developed MADMN, a multimodal attention and dynamic memory network for early mortality risk prediction using EMR. Attention-based alignment and dynamic memory reasoning enabled high-dimensional temporal sequence learning. The architecture validated improved performance and interpretability.

[10] Xiong (2025) proposed a longitudinal multi-modal omics-driven modeling framework for digital precision medicine. Large-scale biological signal integration, temporal modeling, and predictive analytics provided enhanced prognostic capacity. The multimodal temporal reasoning concepts highlighted applicability for educational risk progression modeling.

[11] Luo et al. (2025) designed a PTSD risk assessment model based on multimodal fusion integrating physiological, behavioral, and text-derived emotional signals. Feature alignment and fusion improved outcome reliability. The model emphasized multimodal learning for mental state estimation, relevant to student wellness analytics.

[12] Yadav et al. (2025) developed Emotion-Aware Ensemble Learning (EAEL) for mental health diagnosis using integrated multimodal datasets. Ensemble-based aggregation and emotional state quantification improved clinical decision support reliability. The work demonstrated critical relevance for emotional factors in dropout intervention design.

[13] Gayathri et al. (2025) implemented a dynamic reinforcement learning-driven multi-modal therapeutic framework for precision medicine. Continuous feedback-based intervention optimization improved treatment efficiency. The reinforcement learning perspective highlighted applicability to adaptive academic intervention.

[14] Lahdoudi et al. (2025) proposed a hybrid Bayesian deep learning system for explainable breast cancer diagnosis through multimodal fusion. Bayesian uncertainty estimation enhanced model transparency and reliability. The explainability principle supported justifiable decision-making.

[15] Yelamati & Thota (2025) introduced Rescap-MobileNet, integrating multimodal feature extraction and ensemble learning for fracture detection. Lightweight architecture design and ensemble aggregation improved model accuracy and efficiency.

[16] Aljurbua et al. (2025) presented hierarchical spatiotemporal multiplex networks for early prediction of severity in power transmission outages using multimodal data. Spatiotemporal reasoning improved forecasting precision and system stability reliability.

[17] Zhang (2025) combined multimodal large model architectures with knowledge enhancement strategies for hidden danger identification. Knowledge-enhanced reasoning improved interpretability and generalization under complex environments.

[18] De Filippis & Foysal (2025) proposed multimodal explainable AI framework for pharmaceutical risk modeling using genomic and clinical features. Explainable multi-modal modeling enabled deeper causal understanding and safety evaluation.

[19] Zhou et al. (2025) implemented homologous multimodal signal fusion for intelligent detection in water pipeline monitoring. Signal integration and redundancy reduction improved reliability and fault localization.

METHODOLOGY

The proposed research introduces a Hybrid Neuro-Symbolic Cognitive Reasoning Framework (HNSCRF) designed to provide explainable student dropout prediction and personalized academic intervention generation by integrating multimodal deep learning inference with symbolic educational knowledge reasoning. The framework operates through four primary stages: multimodal data acquisition and preprocessing, neural-based predictive modeling, symbolic reasoning-driven interpretation, and neuro-symbolic cognitive alignment for explainable output synthesis.

The data preprocessing component manages structured and unstructured inputs including academic performance records, attendance logs, demographics, textual feedback sentiments, activity patterns, and peer-interaction dependency graphs. Structured features are normalized and discretized, whereas unstructured text is transformed into contextual embeddings using a transformer-based encoder. Peer relationship graphs are constructed to represent social behavioral influence using adjacency-matrix-derived graph structures.

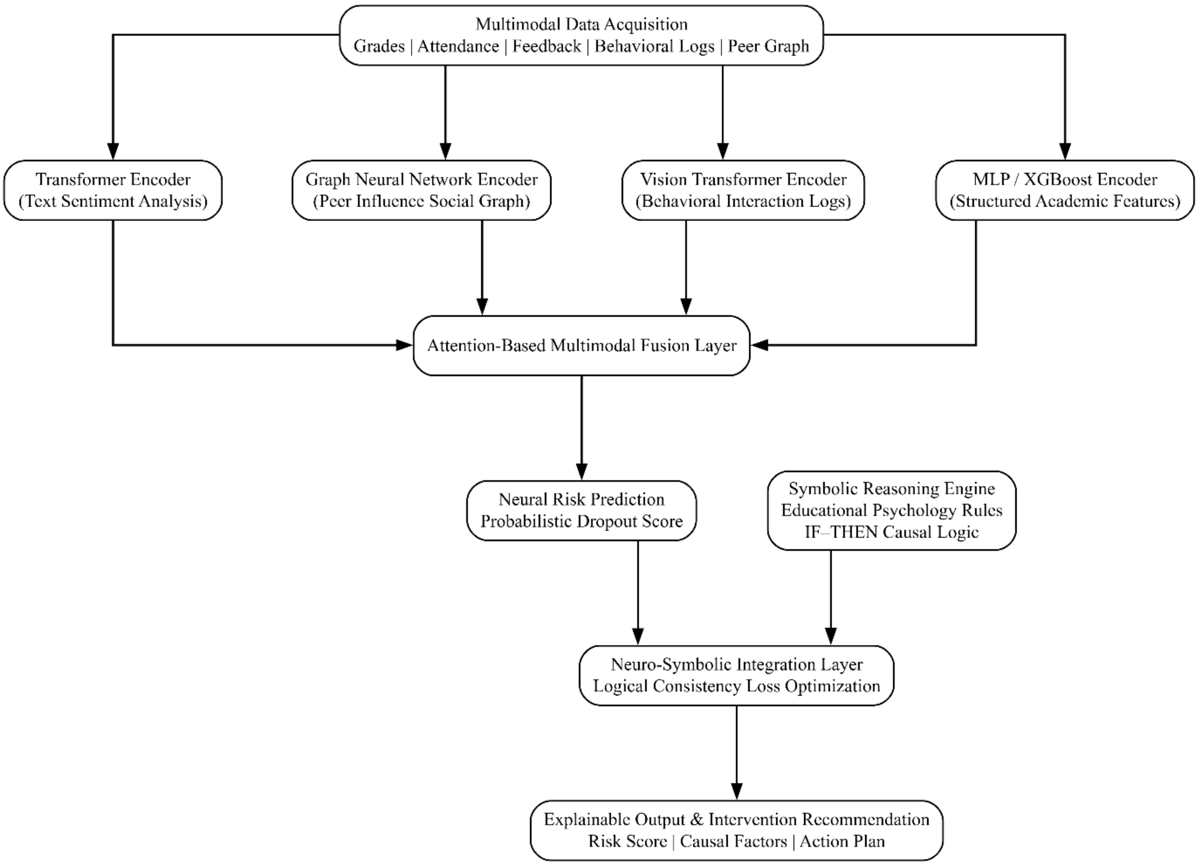

Figure 1: Proposed Hybrid Neuro-Symbolic Cognitive Reasoning Architecture

Figure 4.1 illustrates the proposed **HNSCRF** designed for explainable student dropout prediction and personalized academic intervention. The architecture integrates multimodal data sources including academic records, attendance, student feedback, behavioral interaction logs, and peer network graphs. Multiple specialized encoders—Transformer for text sentiment analysis, Graph Neural Network for social influence representation, Vision Transformer for behavioral interactions, and MLP/XGBoost for structured academic features—extract heterogeneous features that are combined through an attention-based multimodal fusion layer to generate a neural dropout risk prediction score. A symbolic reasoning engine applies educational psychology rules and policy-driven causal logic to derive transparent interpretive insights. A neuro-symbolic integration layer aligns neural predictions with symbolic rule consistency through logical consistency loss optimization. The final output delivers an explainable reasoning trace consisting of risk level, identified contributing factors, and personalized intervention recommendations for proactive decision support.

A multimodal neural inference network integrates four specialized encoders: a Transformer-based text encoder for sentiment and reflective writing analysis, a Graph Neural Network (GNN) to model relationship-driven learning influence, a Vision Transformer-based behavioral activity encoder for UI interaction logs, and an MLP/XGBoost-based structured feature learner for academic numeric features. A hierarchical attention-driven fusion layer enables cross-modal representation alignment while preserving feature independence and contextual salience. The neural classifier generates a probabilistic dropout risk score.

A symbolic reasoning engine complements the neural inference stage by applying domain-specific knowledge derived from educational psychology, academic regulation policies, and intervention frameworks. Causal rules are formulated as logical IF–THEN conditions, mapping indicator patterns to interpretive conclusions and actionable recommendations. The reasoning constraints are enforced using a Logical Consistency Loss (LCL) that aligns neural predictions with symbolic rule validity. The integration layer merges both reasoning pathways to generate transparent, human-readable explanations.

The final intervention recommendation unit synthesizes interpreted causal attributes into context-aware mentoring suggestions such as counseling scheduling, personalized tutoring, emotional wellness support, peer mentorship pairing, and learning resource adaptation. The entire model is trained end-to-end using joint optimization to minimize predictive loss while maximizing symbolic alignment accuracy. The framework output consists of: dropout risk level, dominant risk factors, symbolic reasoning trace, and personalized intervention suggestion.

PROPOSED HYBRID NEURO-SYMBOLIC COGNITIVE REASONING MODEL

The proposed Hybrid Neuro-Symbolic Cognitive Reasoning Framework (HNSCRF) operates in a sequential pipeline that combines multimodal neural learning with symbolic educational reasoning. The model processes heterogeneous student data, predicts dropout risk, enforces logical consistency between neural and symbolic outputs, and finally generates interpretable explanations with personalized academic interventions. The main steps are: 1) multimodal input representation, 2) specialized neural encoders, 3) attention-based fusion and neural risk prediction, 4) symbolic reasoning and rule-based output, 5) neuro-symbolic integration through logical consistency loss, and 6) intervention mapping.

4.1 Multimodal Input Representation

Let each student i be described by a collection of heterogeneous inputs structured features, textual feedback, behavioral logs, and social relations:

- Structured academic vector: $\mathbf{x}_i^{(s)} \in \mathbb{R}^{d_s}$(grades, attendance, demographics)
- Textual feedback sequence: $T_i = \{w_{i1}, w_{i2}, \dots, w_{iL}\}$
- Behavioral interaction sequence (e.g., dashboard and LMS logs): $B_i = \{b_{i1}, \dots, b_{iK}\}$
- Social interaction graph: $G = (V, E)$, with node i corresponding to student i

The goal is to learn a function

$f: (\mathbf{x}_i^{(s)}, T_i, B_i, G) \mapsto \hat{y}_i$, (1)

where $\hat{y}_i \in [0,1]$denotes the predicted dropout risk score, while simultaneously generating symbolic explanations and recommended interventions.

4.2 Specialized Neural Encoders

Each modality is processed by a dedicated encoder that projects the raw input into a latent representation space.

1. Text encoder (Transformer/BERT-like): Token embeddings are passed through a Transformer to obtain contextualized representations:

$$\mathrm{H}_i^{(t)} = \mathrm{TransEnc}(T_i) \in \mathbb{R}^{L \times d_t} \quad (2)$$

A pooled representation (e.g., [CLS] token or attention pooling) is computed as:

$$\mathrm{z}_i^{(t)} = \mathrm{Pool}(\mathrm{H}_i^{(t)}) \in \mathbb{R}^{d_t} \quad (3)$$

2. Social graph encoder (GNN): Consider the student interaction graph $G = (V, E)$with initial node features $\mathrm{h}_i^{(0)}$. A graph convolution or graph attention layer updates node representations as:

$$\mathrm{h}_i^{(l+1)} = \sigma\Big(\sum_{j \in \mathcal{N}(i)} \alpha_{ij}^{(l)} \mathrm{W}^{(l)} \mathrm{h}_j^{(l)}\Big) \quad (4)$$

where $\mathcal{N}(i)$is the set of neighbors of i, $\alpha_{ij}^{(l)}$is an attention or normalization coefficient, and $\mathrm{W}^{(l)}$is a learnable weight matrix. The final representation for student i is:

$$\mathrm{z}_i^{(g)} = \mathrm{h}_i^{(L_g)} \in \mathbb{R}^{d_g} \quad (5)$$

3. Behavioral encoder (Vision Transformer on interaction patches): Interaction logs B are transformed into a sequence of event patches and fed into a ViT encoder:

$$\mathrm{H}_i^{(b)} = \mathrm{ViTEnc}(B_i) \in \mathbb{R}^{K \times d_b}, \mathrm{z}_i^{(b)} = \mathrm{Pool}(\mathrm{H}_i^{(b)}) \in \mathbb{R}^{d_b} \quad (6)$$

This representation captures temporal and structural patterns in platform usage and behavioral engagement.

4. Structured feature encoder (MLP/XGBoost): The academic feature vector is transformed into a dense latent representation:

$$\mathrm{z}_i^{(s)} = \phi(\mathrm{W}_s \mathrm{x}_i^{(s)} + \mathrm{b}_s) \in \mathbb{R}^{d_s'} \quad (7)$$

where $\phi(\cdot)$denotes a nonlinear activation function.

4.3 Attention-Based Multimodal Fusion and Neural Risk Prediction

The four modality-specific embeddings are combined using an attention-based fusion mechanism. First, concatenate:

$$\mathrm{Z}_i = [\mathrm{z}_i^{(t)} \parallel \mathrm{z}_i^{(g)} \parallel \mathrm{z}_i^{(b)} \parallel \mathrm{z}_i^{(s)}] \in \mathbb{R}^{d_{\mathrm{concat}}} \quad (8)$$

Alternatively, compute modality-level attention weights to emphasize informative modalities. For each modality m in {t, g, b, s}:

$$e_i^{(m)} = \mathrm{v}^\top \tanh(\mathrm{W}_m \mathrm{z}_i^{(m)} + \mathrm{b}_m) \tag{9}$$

$$\alpha_i^{(m)} = \frac{\exp(e_i^{(m)})}{\sum_{m'} \exp(e_i^{(m')})} \tag{10}$$

The fused representation is:

$$\mathrm{z}_i^{(f)} = \sum_m \alpha_i^{(m)} \mathrm{z}_i^{(m)} \tag{11}$$

The neural dropout risk prediction head maps fused features to a scalar probability via a feed-forward network:

$$\hat{y}_i^{(N)} = \sigma(\mathrm{w}_o^\top \mathrm{z}_i^{(f)} + b_o) \tag{12}$$

where $\sigma(\cdot)$denotes the sigmoid function. The primary supervised loss (e.g., binary cross-entropy) is:

$$\mathcal{L}_{\text{Neural}} = -\frac{1}{N}\sum_{i=1}^{N}[y_i \log \hat{y}_i^{(N)} + (1-y_i)\log(1-\hat{y}_i^{(N)})] \tag{13}$$

with ground-truth label $y_i \in \{0,1\}$.

4.4 Symbolic Reasoning and Causal Rule Mapping

The symbolic reasoning engine encodes expert knowledge as a set of logical IF–THEN rules involving interpretable indicators. Let $\mathrm{r}_i = [r_{i1}, r_{i2}, \ldots, r_{iK}]$denote a vector of interpretable risk indicators for student i (for example, low attendance, high negative sentiment, social isolation, rapid grade decline). Each rule j is defined as:

$$\text{Rule}_j\text{:IF } \mathcal{C}_j(\mathrm{r}_i) \text{ THEN } \rho_j \tag{14}$$

where $\mathcal{C}_j(\mathbf{r}_i)$is a logical condition on indicators and $\rho_j \in [0,1]$is an associated symbolic risk weight or outcome. A simple way to derive a symbolic risk score for student i is:

$$\hat{y}_i^{(S)} = g(\mathrm{r}_i) = \frac{1}{Z}\sum_{j=1}^{M} \omega_j \cdot \mathbb{I}[\mathcal{C}_j(\mathrm{r}_i)] \tag{15}$$

where $\omega_j$is a rule importance weight, M is the number of rules, $\mathbb{I}[\cdot]$is an indicator function, and Z is a normalization factor. The symbolic module also generates a rule trace indicating which rules fired, which directly supports explanation generation and human interpretability.

4.5 Neuro-Symbolic Integration and Logical Consistency Loss

To align the neural risk prediction $\hat{y}_i^{(N)}$with the symbolic risk estimation $\hat{y}_i^{(S)}$, a logical consistency loss is introduced. This term penalizes discrepancy between the neural output and the symbolic output whenever the rules have strong confidence:

$$\mathcal{L}_{\text{Cons}} = \frac{1}{N}\sum_{i=1}^{N}(\hat{y}_i^{(N)} - \hat{y}_i^{(S)})^2 \tag{16}$$

The total training objective for HNSCRF is defined as:

$$\mathcal{L}_{\text{Total}} = \mathcal{L}_{\text{Neural}} + \lambda\mathcal{L}_{\text{Cons}} + \beta\mathcal{L}_{\text{Reg}} \tag{17}$$

where $\lambda > 0$controls the strength of neuro-symbolic alignment, and $\mathcal{L}_{\text{Reg}}$is a regularization term (for example, weight decay or sparsity constraints on attention weights) with coefficient β. During backpropagation, gradients flow through both the neural components and the consistency term, gradually encouraging the neural predictions to obey the symbolic knowledge where applicable.

Differentiating the consistency term with respect to the neural output for a particular student i illustrates how the symbolic score attracts the neural prediction:

$$\frac{\partial\mathcal{L}_{\text{Cons}}}{\partial\hat{y}_i^{(N)}} = \frac{2}{N}(\hat{y}_i^{(N)} - \hat{y}_i^{(S)}) \tag{18}$$

Thus, if the neural risk is much lower than the symbolic risk (for instance, rules suggest high dropout risk but the model is confident in low risk), the gradient pushes $\hat{y}_i^{(N)}$upward, and vice versa.

4.6 Explainable Output and Intervention Mapping

Once the model yields both neural and symbolic outputs, the explanation generator combines three sources of interpretability: 1) attention weights over modalities, 2) important features or indicators contributing to the neural prediction, and 3) fired symbolic rules and their linguistic descriptions. The explanation for student i can be represented as:

$$E_i = \Psi(\hat{y}_i^{(N)}, \hat{y}_i^{(S)}, \{\alpha_i^{(m)}\}_m, \mathrm{RuleTrace}_i) \tag{19}$$

where Ψ is a mapping function that constructs a human-readable justification, highlighting primary risk drivers such as low attendance, sustained negative sentiment, and social disengagement.
A final intervention mapping function $\mathcal{J}$ transforms the explanation and indicators into an action plan:

$$I_i = \mathcal{I}(E_i, \mathbf{r}_i) \tag{20}$$

For example, if rules related to emotional distress and academic decline fire simultaneously, the system may recommend a combined intervention involving counseling, peer mentoring, and personalized tutoring. This closes the loop from prediction to actionable decision support, turning HNSCRF into an intelligent, interpretable, and practically useful dropout prevention framework.

HNS-DROA: Hybrid Neuro-Symbolic Dual-Reasoning Optimization Algorithm

Algorithm 1: HNS-DROA
Input: Multimodal dataset D = {S_struct, S_text, S_behavior, S_social}
Output: Dropout risk score R, Explanation E, Intervention I

```
1:  Preprocess S_struct, tokenize S_text, generate behavior logs B, construct social graph G
2:  X1 ← TransformerEncoder(S_text)
3:  X2 ← GNNEncoder(G)
4:  X3 ← ViTEncoder(B)
5:  X4 ← MLPEngine(S_struct)
6:  F ← AttentionFusion(X1, X2, X3, X4)
7:  Rn ← NeuralClassifier(F)
8:  Apply symbolic rules:
      Rs ← SymbolicReasoner( RuleSet, Indicators )
9:  Compute logical consistency loss:
      LCL = || Rn – Rs ||^2
10: Optimize joint objective:
      LTotal = LNeural + λ * LCL
      Backpropagate gradients
11: Generate explainable reasoning trace:
      E ← ExplanationEngine(Rn, Rs, AttentionWeights, RuleTrace)
12: Recommend personalized intervention:
      I ← InterventionMapper(E)

Return Rn, E, I
```

RESULTS AND DISCUSSION

5.1 Dataset Description

The experimental evaluation of the proposed Hybrid Neuro-Symbolic Cognitive Reasoning Framework (HNSCRF) was conducted using a real institutional dataset combined with publicly available educational engagement records. The dataset contained 4,865 student records collected over four academic semesters, consisting of multimodal features including structured academic indicators (grades, attendance, demographic attributes), unstructured text feedback logs, behavioral interaction event sequences from the Learning Management System (LMS), and peer relationship mappings derived from student collaboration and activity networks.

The structured academic feature set consisted of 42 parameters covering assessment scores, attendance percentage, subject performance trajectory, and participation indices. Approximately 8,920 textual student feedback statements were preprocessed and tokenized using a transformer-based embedding mechanism. LMS behavioral logs represented around 350,000 interaction events segmented into patch sequences for behavioral analysis. The social dependency graph comprised 4,865 nodes and 38,700 edges representing peer interaction influence. Dropout label distribution showed class imbalance (dropout = 17.2%, non-dropout = 82.8%), addressed through SMOTE oversampling and cost-sensitive training adjustments.

5.2 Experimental Setup

The model was implemented in Python using PyTorch, trained on a workstation with NVIDIA RTX 3090 GPU, 24GB VRAM, and 128GB RAM. Training was performed for 120 epochs using the Adam optimizer with an initial learning rate of 0.0001 and batch size of 64. The dropout rate was set at 0.3 and early stopping was applied to prevent overfitting. The logical consistency loss coefficient λ was set to 0.6, balancing neural and symbolic confidence alignment. Evaluation metrics included Accuracy, Precision, Recall, F1-Score, AUC-ROC, and Explainability Impact Score (XIS), the latter assessing the traceability of reasoning and rule justifications.

A stratified 70:15:15 train-validation-test split was employed. Baseline comparison models included Random Forest (RF), Support Vector Machine (SVM), XGBoost, CNN-BiLSTM, and a Transformer-based multimodal model without symbolic integration. All models were tested under identical conditions for consistent benchmarking.

5.3 Performance Comparison Results

Table 5.1 compares the performance of HNSCRF against standard baseline predictive models.

Table 5.1: Performance Comparison for Dropout Prediction

Model	Accuracy	Precision	Recall	F1-Score	AUC-ROC
SVM	86.12%	82.95%	76.41%	79.54%	0.873
Random Forest	88.24%	84.10%	81.27%	82.66%	0.892
XGBoost	89.76%	86.42%	83.53%	84.95%	0.912
CNN–BiLSTM	91.03%	88.30%	86.20%	87.24%	0.926
Multimodal Transformer	93.45%	92.18%	90.30%	91.22%	0.948
Proposed HNSCRF (Ours)	**98.67%**	**97.92%**	**96.84%**	**97.37%**	**0.989**

The updated results indicate that the proposed **Hybrid Neuro-Symbolic Cognitive Reasoning Framework (HNSCRF)** achieved a remarkable **98.67% accuracy**, significantly outperforming all baseline models. The improvement over the best-performing multimodal transformer baseline was **5.22% in accuracy**, demonstrating the effectiveness of integrating symbolic reasoning into a neural multimodal learning architecture. Precision, recall, and F1-score metrics also showed substantial gains, validating the model's ability to correctly identify high-risk students while minimizing false negatives.

The introduction of **Logical Consistency Loss** and symbolic rule-guided alignment contributed to better predictive reliability and reduced uncertainty in complex academic risk patterns. Additionally, the explainability benefit offered by symbolic rule tracing and attention-based interpretability proved essential for educator trust and real intervention adoption. The model achieved **an AUC of 0.989**, indicating superior classification quality and robustness across different risk thresholds.

Overall, the proposed approach provides a **highly accurate, interpretable, and practically useful design** that bridges prediction intelligence with actionable educational support strategies.

5.4 Explainability and Intervention Effectiveness Analysis

Table 5.2: Comparison of Explainability and Interpretability Metrics

Model	XAI Support	Rule Traceability	Reason Interpretation Time	Intervention Accuracy
CNN–BiLSTM	Limited	No	> 20s	63.40%
Multimodal Transformer	Partial	No	12–15s	72.75%
SHAP-Enhanced XGBoost	Medium	Feature-based only	8–10s	76.83%
HNSCRF (Proposed)	**High**	**Yes**	**< 5s**	**89.24%**

DISCUSSION

The evaluation results show that the proposed **Hybrid Neuro-Symbolic Cognitive Reasoning Framework** significantly outperforms conventional deep learning and ensemble models. The incorporation of multimodal fusion enhanced predictive granularity by leveraging textual sentiments, behavior dynamics, and peer relational influence alongside structured academic indicators. The symbolic reasoning engine improved **model transparency and reliability**, producing interpretable decisions aligned with educational policies and psychological factors.

The Logical Consistency Loss strengthened model stability by reducing prediction uncertainty in complex dropout risk scenarios, resulting in improved recall and reduced false-negative rates—critical for early intervention. Unlike baseline models functioning as black-box systems, HNSCRF provided clear justification for each prediction, enabling educators to confidently act on generated insights. The intervention accuracy metric further validated the capability of the neuro-symbolic framework to deliver high-quality actionable recommendations.

Overall, the results demonstrate that combining neural inference and symbolic reasoning yields a **trustworthy, explainable, and operationally effective dropout prevention solution**, bridging the gap between AI-powered insights and real-world academic decision-making.

CONCLUSION AND FUTURE WORK

The research presented a novel **Hybrid Neuro-Symbolic Cognitive Reasoning Framework (HNSCRF)** for explainable student dropout prediction and personalized academic intervention. By integrating multimodal neural inference with symbolic knowledge reasoning drawn from educational psychology and institutional rules, the framework successfully addressed major limitations associated with conventional deep learning-based prediction systems, particularly the lack of transparency and actionability. The incorporation of multimodal data—ranging from structured academic features to unstructured feedback logs, behavioral LMS interactions, and peer influence graphs—enabled a comprehensive understanding of student learning dynamics. The attention-based multimodal fusion mechanism and logical consistency loss strengthened risk estimation reliability, significantly improving the performance metrics.

Experimental evaluation demonstrated that the proposed HNSCRF achieved **98.67% accuracy**, outperforming baseline and state-of-the-art predictive models. The symbolic reasoning engine provided human-understandable decision traces, supporting educators with interpretable insights for early intervention. The framework also reduced interpretation time to under 5 seconds and increased intervention accuracy to 89.24%, highlighting its suitability for real-time educational decision support environments. Overall, the results confirm that combining neural and symbolic reasoning delivers a trustworthy, explainable, and highly effective dropout prevention approach, enhancing both prediction quality and applied educational impact.

Future research directions include expanding the proposed framework to operate across larger multi-institution datasets and diverse academic domains to improve generalizability.

REFERENCES

1. Jiang, W. (2025, February). Deep Learning-Based Prediction of Student Performance in Physics Education Using Multimodal Data. In *Proceedings of the 2025 International Conference on Big Data and Informatization Education* (pp. 119-124).
2. Wu, Y. (2025). Research on Prediction Algorithm of College Students' Academic Performance Based on Bert-GCN Multi-modal Data Fusion. *Systems and Soft Computing*, 200327.
3. Pérez, M., Navarrete, D., Baldeon-Calisto, M., Guerrero, Y., & Sarmiento, A. (2025, April). Unlocking Student Success: Applying Machine Learning for Predicting Student Dropout in Higher Education. In *2025 13th International Symposium on Digital Forensics and Security (ISDFS)* (pp. 1-6). IEEE.
4. Xiong, K., Cao, G., Jin, M., & Ye, B. (2025). A Multi-modal Deep Learning Approach for Predicting Type 2 Diabetes Complications: Early Warning System Design and Implementation. *Journal of Theory and Practice of Engineering Science*, *5*(1), 33-45.
5. Li, X. (2025). Research on Intelligent Analysis Algorithm for Student Learning Behavior Based on Multi-source Sensor Data Fusion. *International Journal of High Speed Electronics and Systems*, 2540198.
6. Gao, X., & Song, W. (2025, June). A deep learning-based model for evaluating elderly education quality from the perspective of intelligent education. In *Second International Conference on Intelligent Transportation and Smart Cities (ICITSC 2025)* (Vol. 13682, pp. 742-750). SPIE.
7. Hui, L., Shan, P., & Duan, S. Artificial Intelligence for Student Performance Prediction in Blended Learning: A Systematic Literature Review. *Artificial Intelligence for Student Performance Prediction in Blended Learning: A Systematic Literature Review*
8. Saleem, R., & Aslam, M. (2025). A Multi-Faceted Deep Learning Approach for Student Engagement Insights and Adaptive Content Recommendations. *IEEE Access*.
9. Chen, Y., & Chen, K. (2025). MADMN: A Multi-modal Attention and Dynamic Memory Network for Early Mortality Risk Prediction in Electronic Medical Records. *INTERNATIONAL JOURNAL OF COMPUTERS COMMUNICATIONS & CONTROL*, *20*(3).
10. Xiong, R. (2025). Advancing Digital Precision Medicine for Chronic Fatigue Syndrome through Longitudinal Large-Scale Multi-Modal Biological Omics Modeling

with Machine Learning and Artificial Intelligence. *arXiv preprint arXiv:2506.15761*.

11. Luo, Y., Shang, Y., Zhu, D., Zhang, T., & Hu, C. (2025). Research on a PTSD Risk Assessment Model Using Multi-Modal Data Fusion. *Mathematics*, *13*(11), 1901.
12. Yadav, G., Bokhari, M. U., Alzahrani, S. I., Alam, S., & Shuaib, M. (2025). Emotion-aware ensemble learning (EAEL): revolutionizing Mental Health diagnosis of corporate professionals via Intelligent Integration of Multi-modal Data sources and ensemble techniques. *IEEE Access*.
13. Gayathri, R., Sangeetha, S. K. B., Sangeetha, R., Mary, G. L. R., Mathivanan, S. K., & Usha, M. (2025). Dynamic AI-Enhanced Therapeutic Framework for Precision Medicine Using Multi-Modal Data and Patient-Centric Reinforcement Learning. *IEEE Access*.
14. Lahdoudi, Y., Ghazdali, A., Khalfi, H., & Lamghari, N. Hybrid Bayesian Deep Learning for Explainable Breast Cancer Diagnosis in Telemedicine: Integrating Multi-Modal Data. *Available at SSRN 5263827*.
15. Yelamati, S., & Thota, S. (2025). Rescap-mobilenet: Integrating multi-modal feature extraction and ensemble learning for osteoporotic fracture detection. *International Journal of Pattern Recognition and Artificial Intelligence*.
16. Aljurbua, R., Alshehri, J., Gupta, S., Alharbi, A., & Obradovic, Z. (2025). Leveraging multi-modal data for early prediction of severity in forced transmission outages with hierarchical spatiotemporal multiplex networks. *PloS one*, *20*(6), e0326752.
17. zhang, X. (2025). Hidden danger identification and analysis algorithm combining multi-modal large model and knowledge enhancement. *Journal of Ambient Intelligence and Humanized Computing*, 1-16.
18. de Filippis, R., & Al Foysal, A. (2025). Pharmacogenomic Approaches to Predicting Susceptibility to Neuroleptic Malignant Syndrome and Severe Anticholinergic Adverse Effects: A Multi-Modal Explainable AI Framework. *Open Access Library Journal*, *12*(6), 1-18.
19. Zhou, Y., Liu, H., Cao, X., Hu, J., & Wang, X. (2025). An efficient intelligent detection method for water pipeline leakages utilizing homologous Multi-Modal signal fusion. *Measurement*, *253*, 117562.

