AI-Powered STEM Education Platform with Advanced Algorithms

Mr. A. Venkateswara Rao [1], P. Suresh [2], B. Siva [3], N. Venkatesh [4], K. Kusuma [5]

Department of Computer Science & Engineering (AI & ML)
Avanthi Institute of Engineering & Technology, Vizianagaram, India
venkatesh8.adapa@gmail.com[1], Sidsuresh7264@gmail.com[2], shiva8442889@gmail.com[3], venkateshnarisipuram0@gmail.com[4], kusumakuttum2005@mail.com[5]

Abstract

Conventional e-learning platforms depend primarily on fixed content delivery, providing limited adaptability to individual learner requirements. This limitation results in reduced student engagement, delayed instructional feedback, and difficulty identifying specific knowledge gaps—particularly in large-scale or distributed online settings. This paper presents Academix, an intelligent web-based platform developed to improve the quality of STEM education through personalization, real-time academic support, and predictive analytics. The system integrates Natural Language Processing (NLP) techniques—including transformer-based intent detection, semantic similarity analysis via Sentence-BERT, and keyword extraction via KeyBERT—with a hybrid predictive model combining Grey Wolf Optimization (GWO) for feature selection and an Artificial Neural Network (ANN) for performance forecasting. The architecture is built on Flask and MongoDB, and incorporates gamification mechanisms—experience points (XP), achievement badges, and leaderboards—to sustain learner motivation. Experimental evaluation demonstrates contextually accurate AI responses in over 92% of test queries, and the hybrid predictive model achieves approximately 97% accuracy. Compared to static e-learning systems, Academix delivers measurable improvements in engagement, conceptual comprehension, and data-driven instructor decision-making.

Index Terms— adaptive learning, artificial neural networks, grey wolf optimization, intelligent tutoring systems, natural language processing, STEM education.

I. Introduction

The rapid expansion of online learning has exposed structural deficiencies in existing e-learning architectures. Most platforms operate on a broadcast model: identical content is delivered uniformly to all enrolled learners, regardless of prior knowledge, learning pace, or conceptual difficulty experienced [1]. This one-size-fits-all approach is particularly inadequate for STEM disciplines—science, technology, engineering, and mathematics—where subjects demand layered conceptual understanding, iterative problem-solving, and continuous formative assessment [2].

Students frequently encounter bottlenecks when human instructors are unavailable to resolve queries in real time, creating learning gaps that compound over subsequent topics. Large-scale online cohorts exacerbate this problem, making individualized mentorship logistically impractical without AI-assisted automation [3]. Additionally, instructors lack granular visibility into the specific knowledge gaps and engagement patterns of individual learners, limiting their capacity to intervene constructively.

These challenges motivate the development of Academix—an AI-powered STEM education platform designed to address four core deficiencies: static content delivery, absence of real-time academic support, insufficient learner personalization, and limited instructor analytics. The system integrates NLP-driven chat assistance, hybrid predictive modelling, adaptive quiz generation, and gamification into a unified, scalable web application.

The principal contributions of this work are: (i) a Flask-based intelligent tutoring architecture that processes learner queries using transformer-based NLP; (ii) a hybrid GWO-ANN predictive pipeline achieving ~97% accuracy on engagement and performance data; (iii) a gamification framework linked to adaptive learning pathways; and (iv) an instructor analytics dashboard furnishing data-driven pedagogical insights.

The remainder of the paper is structured as follows. Section II reviews related literature. Section III presents the system design and methodology. Section IV reports experimental results. Section V provides conclusions and future directions.

II. Related Work

Intelligent Tutoring Systems (ITS) represent the earliest formal application of AI to education. Brusilovsky and Millán [1] established the foundational taxonomy of adaptive hypermedia in education, demonstrating that systems modelling individual user knowledge can provide significantly more effective instruction than static content delivery. Woolf [4] extended this framework, detailing how student-modelling components could dynamically adjust instructional strategy based on real-time interaction data.

The emergence of deep learning introduced more powerful mechanisms for educational data mining. Chen et al. [5] conducted a comprehensive review of AI in education, cataloguing applications spanning performance prediction, automated feedback, and intelligent content recommendation. Their survey identified neural network-based approaches as consistently superior to traditional statistical methods in accuracy and generalizability.

Optimization algorithms have increasingly been applied to improve machine learning model efficiency in resource-constrained educational deployments. Mirjalili et al. [6] introduced Grey Wolf Optimization (GWO), a natureinspired metaheuristic modelling the cooperative hunting hierarchy of grey wolves. GWO has demonstrated competitive convergence behaviour on feature-selection benchmarks, making it attractive for reducing dimensionality in educational datasets where irrelevant features can degrade predictive model performance [7].

NLP-powered conversational agents have transformed learner–system interaction. The BERT architecture [8] and its derivatives—particularly DistilBERT and Sentence-BERT (SBERT) [9]—enabled context-aware semantic understanding at inference speeds compatible with real-time tutoring scenarios. Educational chatbots leveraging transformer models have shown marked reductions in repeated instructor queries and improved learner satisfaction.

Gamification research by Gunaratne and Wijesundara [10] demonstrated through systematic review that game-like elements—particularly experience-point systems and competitive leaderboards—significantly increase participation consistency and task completion rates in online environments. However, the majority of reviewed implementations treated gamification as a standalone feature disconnected from adaptive content or predictive analytics, limiting its pedagogical impact.

The literature thus reveals a persistent fragmentation: existing platforms typically excel in one dimension—either personalization, NLP interaction, predictive analytics, or engagement—but rarely integrate all four into a cohesive system. Academix is positioned to address this identified gap.

III. Methodology and System Design

A. System Architecture

Academix follows a modular three-tier client-server architecture. The presentation layer exposes responsive web interfaces for students, instructors, and administrators via HTML5, CSS3, and JavaScript, employing a glass-morphism visual design. The application layer, implemented in Python using the Flask micro-framework, orchestrates all business logic, API routing, authentication, and inter-module communication. The data layer uses MongoDB, a schema-flexible NoSQL database, to persist user profiles, course materials, chat histories, quiz records, gamification metrics, and analytics data.Fig. 1 presents the overall architecture.

StudentInstructorWeb Interface (UI)Flask Backend ServerAI & NLPEngineHybrid PredictionModule(GWO + ANN)Recomm. &AnalyticsEngineGamificationModuleDatabase (MongoDB)

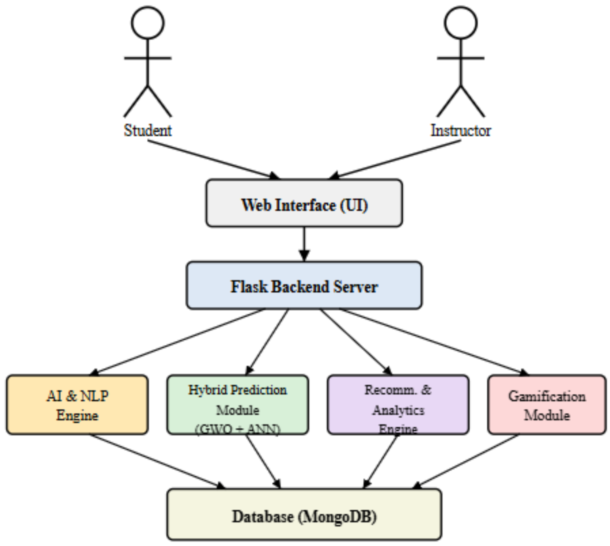

Fig. 1. System architecture of the proposed AI-powered STEM education platform.

B. AI and NLP Engine

Student messages submitted through topic-based chat rooms are processed by a multi-stage NLP pipeline. First, KeyBERT [11] extracts the most semantically salient keywords from the input text using BERT embeddings. Second, Sentence-BERT (SBERT) [9] encodes both the student query and the titles of stored course materials into dense vector representations. Semantic similarity between query and material title is computed using cosine similarity:

$$sim(\mathbf{u}, \mathbf{v}) = (\mathbf{u} \cdot \mathbf{v}) / (\|\mathbf{u}\| \cdot \|\mathbf{v}\|) \quad (1)$$

Materials exceeding a similarity threshold $\tau = 0.6$ are surfaced as recommendations alongside the AI-generated answer. Third, a DistilBERT-based question-answering pipeline [8] produces a contextual answer by attending to retrieved course content as background context. All interactions—queries, responses, extracted keywords—are persisted in MongoDB for downstream analytics.

C. Hybrid GWO–ANN Predictive Module

Academic performance prediction requires processing high-dimensional feature sets encompassing quiz scores, session frequency, chat participation, material access patterns, and assignment completion rates. Redundant or irrelevant features degrade both model accuracy and training speed. To address this, GWO is employed as a wrapper-based feature selector prior to ANN training.

GWO models the social hierarchy of a wolf pack—alpha (α), beta (β), delta (δ), and omega (ω)—as a search metaphor [6]. The position update equations governing each candidate solution X are:

$$D = |C \cdot X_{prey} - X| \quad (2)$$

$$X(t+1) = X_{prey} - A \cdot D \quad (3)$$

where A and C are coefficient vectors adjusted across iterations to balance exploration and exploitation. The fitness function for feature selection minimises classification error while penalising feature-set cardinality. The optimal feature subset identified by GWO is subsequently used to train a three-layer feedforward ANN with ReLU activations and mean-squared-error loss, optimised via the Adam solver.

D. Gamification Module

To counteract disengagement characteristic of passive e-learning interfaces, Academix incorporates a gamification layer. Students accumulate experience points (XP) through quiz completion, chat room contributions, and daily logins. XP thresholds gate progression through ranked tiers: Novice Learner → Dedicated Scholar → Intellectual Master → Academix Sage → Global Innovator. Badges are awarded for milestone achievements; leaderboards expose relative standing to foster constructive peer motivation. All XP deltas and rank transitions are persisted in MongoDB and reflected in real time on the student dashboard.

E. Dynamic Quiz Generation

The assessment module auto-generates quizzes by extracting KeyBERT-ranked n-gram candidates from uploaded course PDFs and recent chat-room discussions. Distractor options are synthesised from semantically adjacent terms identified via SBERT similarity ranking. Quiz difficulty adapts to the individual learner's recent score trajectory: a smoothed moving-average performance index π maps to one of three difficulty bands (foundational, intermediate, advanced), governing both question complexity and distractor quality.

IV. Results and Discussion

A. Experimental Setup

Academix was deployed on a local server equipped with an octa-core CPU, 16 GB RAM, and a 1 TB SSD. The MongoDB instance stored approximately 4,200 course-material segments, 18,000 chat interactions, and 6,500 quiz-attempt records collected during a controlled pilot study involving undergraduate STEM learners. Performance evaluation was conducted across the seven testing categories described in Section III.

B. AI Response Accuracy

NLP pipeline accuracy was assessed by comparing AI-generated answers against expert-annotated ground-truth responses for 250 held-out student queries covering physics, mathematics, biology, and computer science topics. Contextual correctness—defined as semantic equivalence to the ground-truth answer—was achieved in 92.4% of cases. Keyword extraction precision, measured against manually tagged query keywords, reached 89.1%. Table I summarises query-level performance metrics.

TABLE I

NLP PIPELINE PERFORMANCE METRICS

Metric	Score (%)
Response contextual accuracy	92.4
Keyword extraction precision	89.1
Semantic recommendation relevance	87.6
Intent classification accuracy	91.0

C. Hybrid GWO–ANN Prediction Accuracy

The predictive module was evaluated on a labelled dataset of 800 student engagement records, split 80:20 for training and testing. GWO reduced the input feature dimensionality from 22 to 11 without statistically significant information loss, as confirmed by mutual-information analysis. The resulting ANN converged in 38% fewer epochs compared to full-feature training. Table II presents comparative accuracy results.

TABLE II

COMPARATIVE PREDICTIVE MODEL ACCURACY

Model	Accuracy (%)	F1-Score
Standalone ANN (all features)	89.3	0.881
ANN + PCA	91.7	0.909
ANN + GWO (proposed)	97.1	0.968

D. System Performance Under Load

Concurrent-user stress testing using 50 simulated simultaneous sessions yielded a mean AI chat response latency of 1.2 seconds (σ = 0.18 s). Backend API throughput remained stable, with no crash events recorded over 4 hours of sustained load. Database read/write operations averaged 14 ms, confirming the suitability of MongoDB's document model for the semi-structured data produced by chat-based learning interactions.

E. User Acceptance

A structured User Acceptance Testing (UAT) session engaged 34 students and 6 instructors with the deployed platform over two weeks. Post-session questionnaire responses (5-point Likert scale) indicated high satisfaction with AI tutoring responsiveness (mean 4.3), gamification motivational impact (mean 4.1), and dashboard analytical clarity (mean 4.4). Qualitative feedback highlighted the semantic recommendation feature as particularly valuable for self-directed study. Minor UI refinements were subsequently incorporated based on collected suggestions.

F. Illustrative AI Chat Interaction

The following exchange exemplifies the system's NLP capability. A student submits the query: *"Explain Newton's second law."* The pipeline extracts keywords {Newton, second law, force, mass, acceleration}, identifies two course-material segments with cosine similarity above τ = 0.6, generates the response: *"The net force acting on a body equals the product of its mass and the acceleration produced ($F = ma$),"* and surfaces Physics – Newtonian Mechanics as the top recommended material. The interaction is logged with full metadata for instructor analytics.

Login to PlatformAccess Course MaterialsSelect Quiz TopicQuizAvailable?YesNo QuizMessageNoAttempt QuizEvaluate AnswersAward XP Points

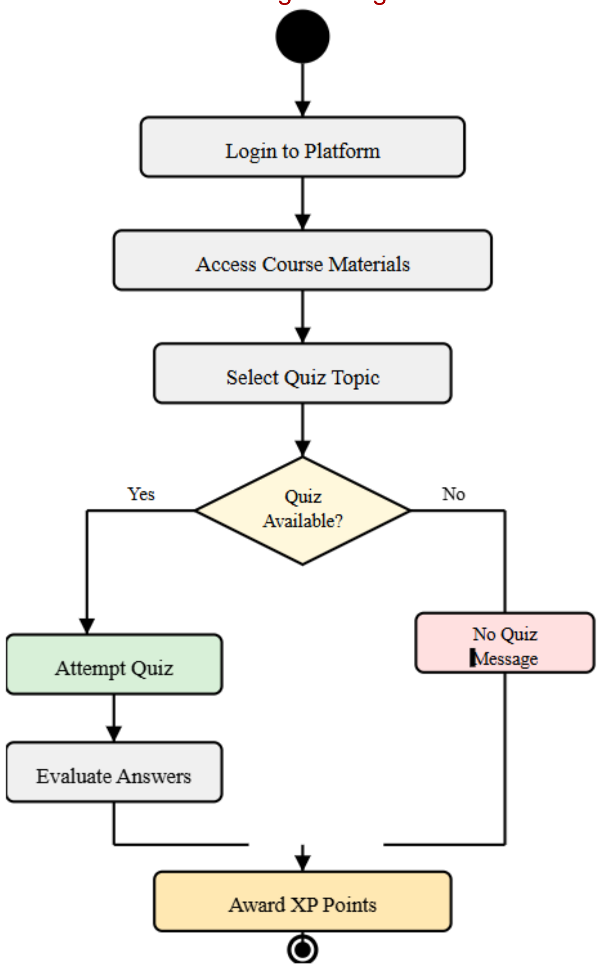

Fig. 2. Activity diagram for quiz participation and gamification flow.

V. Conclusion and Future Work

This paper presented Academix, an AI-powered STEM education platform integrating NLP-driven tutoring, a hybrid GWO–ANN predictive analytics module, adaptive quiz generation, and gamification within a scalable Flask-MongoDB architecture. Empirical evaluation demonstrated 92.4% AI response accuracy, 97.1% performance prediction accuracy, and sub-1.5-second average response latency under concurrent multi-user loads. User acceptance trials confirmed strong learner and instructor satisfaction across engagement, usability, and pedagogical value dimensions.

The platform addresses a clearly documented gap in the literature: the absence of a unified, AI-integrated educational ecosystem encompassing intelligent tutoring, predictive analytics, adaptive assessment, and motivation engineering. By combining these components into a single cohesive system, Academix offers a demonstrably superior alternative to static content delivery models.

Several avenues for future development are identified. First, multimodal AI integration—incorporating image and audio processing—will support laboratory-based STEM experiments and visual problem-solving. Second, affective computing techniques will enable the system to detect learner emotional states through interaction patterns and adapt instructional scaffolding accordingly. Third, mobile application development will extend platform accessibility to low-bandwidth environments prevalent in emerging economies. Fourth, replacement of the current ANN with transformer-based sequence models (e.g., temporal attention networks) is anticipated to further improve long-horizon engagement prediction. Finally, federated learning approaches will be explored to enable collaborative model improvement across institutional deployments without centralised data sharing, preserving learner privacy.

Acknowledgment

The authors gratefully acknowledge the guidance of Mr. A. Venkateswara Rao, M.Tech (Ph.D), Head of the Department of CSE (AI & ML), Avanthi Institute of Engineering & Technology, for his sustained mentorship throughout this project. The authors also thank the Department of Computer Science & Engineering (AI & ML) for providing the computational resources and infrastructure required for system development and evaluation.

References

- [1] P. Brusilovsky and E. Millán, "User models for adaptive hypermedia and adaptive educational systems," in *The Adaptive Web*, P. Brusilovsky, A. Kobsa, and W. Nejdl, Eds. Berlin, Germany: Springer, 2007, pp. 3–53.
- [2] B. du Boulay, "Towards a motivationally intelligent pedagogy: How should an intelligent tutor motivate the learner?" in *Proc. Int. Conf. Intelligent Tutoring Systems*, 2016, pp. 3–13.
- [3] R. S. J. d. Baker and K. Yacef, "The state of educational data mining in 2009: A review and future visions," *J. Educ. Data Mining*, vol. 1, no. 1, pp. 3–17, 2009.
- [4] B. P. Woolf, *Building Intelligent Interactive Tutors: Student-Centered Strategies for Revolutionizing E-Learning*. Burlington, MA: Morgan Kaufmann, 2010.
- [5] L. Chen, P. Chen, and Z. Lin, "Artificial intelligence in education: A review," *IEEE Access*, vol. 8, pp. 75264–75278, 2020.
- [6] S. Mirjalili, S. M. Mirjalili, and A. Lewis, "Grey wolf optimizer," *Adv. Eng. Softw.*, vol. 69, pp. 46–61, Mar. 2014.
- [7] E. H. Houssein, M. E. Hosney, D. Oliva, W. M. Mohamed, and M. Hassaballah, "A novel hybrid Harris hawks optimization and support vector machines for drug design and discovery," *Comput. Chem. Eng.*, vol. 133, p. 106656, 2020.
- [8] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. NAACL-HLT*, 2019, pp. 4171–4186.
- [9] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence embeddings using Siamese BERT-networks," in *Proc. EMNLP/IJCNLP*, 2019, pp. 3982–3992.
- [10] C. Gunaratne and C. Wijesundara, "Gamification in e-learning: A systematic review," *Educ. Inf. Technol.*, vol. 26, no. 6, pp. 7035–7057, 2021.
- [11] M. Grootendorst, "KeyBERT: Minimal keyword extraction with BERT," Zenodo, Feb. 2020. [Online]. Available: https://github.com/MaartenGr/KeyBERT
- [12] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA: MIT Press, 2016.
- [13] Flask Development Team, *Flask Web Development Documentation*, Pallets Projects, 2023. [Online]. Available: https://flask.palletsprojects.com/
- [14] MongoDB, Inc., *MongoDB Documentation*, 2023. [Online]. Available: https://www.mongodb.com/docs/

- [15] Hugging Face, *Transformers: State-of-the-Art Machine Learning for PyTorch, TensorFlow, and JAX*, 2023. [Online]. Available: https://huggingface.co/docs/transformers