



DISSEMINATION OF KNOWLEDGE

# International Journal of Artificial Intelligence and Machine Learning

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



Research Paper

Open Access

## Smart Tutoring Systems Using Interactive Natural Language Understanding with Transformer-Based Models

Damodaran<sup>1\*</sup>, Dr. Lekshmi. R.S<sup>2</sup>, Mamayusup Abdusamatov<sup>3</sup>, Dr. V. Subbulakshmi<sup>4</sup>, Sitora Azizova<sup>5</sup>  
Bakhodir Jalilov<sup>6</sup>

<sup>1</sup>Associate Professor, Psychology, Meenakshi College of Arts and Science, Meenakshi Academy of Higher Education and Research, Chennai, Tamil Nadu, India. E-mail: damodaranb@maher.ac.in

<sup>2</sup>Professor, St. Joseph's College of Engineering, OMR, Chennai, Tamil Nadu, India. E-mail: lekshmi.ajit@gmail.com; lekshmir@stjosephs.ac.in, <https://orcid.org/0000-0001-5597-1283>

<sup>3</sup>Department of Accounting and Statistics, Termez University of Economics and Service, Termez, Uzbekistan. E-mail: mamayusup\_abdusamatov@tues.uz, <https://orcid.org/0009-0000-0207-1484>

<sup>4</sup>Principal, Meenakshi College of Yoga Science and Therapy, Meenakshi Academy of Higher Education and Research, Chennai, Tamil Nadu, India. E-mail: drsubbulakshmiyoga@maher.ac.in

<sup>5</sup>Lecturer, Department of Translation and Interpreting Studies, Silk Road International University of Tourism and Cultural Heritage Samarkand, Uzbekistan. E-mail: azizovasilitora06@gmail.com, <https://orcid.org/0009-0006-7562-1145>,

<sup>6</sup>PhD Student, Bukhara State University, Teacher, Department of English Language and Literature, Bukhara State Pedagogical Institute, Bukhara, Uzbekistan. E-mail: bakhodirjalilov749@gmail.com, <https://orcid.org/0009-0006-1384-7202>

\*Corresponding author: Email: damodaranb@maher.ac.in

### Abstract

Due to the growing intricacy in today's educational settings, intelligent adaptive tutor systems need to be developed that can deliver personalized tutoring services on a large scale. This research paper highlights a new idea for designing a Smart Tutoring System (STS) using Interactive Natural Language Understanding (INLU) that is combined with transformer-based deep learning models in order to offer contextualized feedback and adaptability for users. The STS uses a BERT-based language model that is trained on educational data consisting of over 120,000 question-answer pairs pertaining to five academic subjects, which include mathematics, science, language arts, history, and computer programming. For the evaluation of the methodology, a multi-staged approach was adopted, which consists of steps such as intention classification, similarity matching, traversal of the knowledge graph, and feedback generation using reinforcement learning. The experiments performed on the ATIS-Edu benchmark dataset and the own Educational QA dataset show that the model has obtained 94.7% accuracy, 93.8% F1-Score, 94.1% precision, and 93.5% recall scores, which are significantly higher than state-of-the-art baselines such as a GPT-2-based tutoring system (88.4% accuracy) and an LSTM-based NLU (82.1% accuracy). Moreover, the model provides a mean latency reduction of 38% in student response time in comparison to traditional rule-based intelligent tutoring systems. Ablation experiments prove the individual contribution of each component such as the transformer encoder, knowledge graph, and reinforcement feedback towards improving the model's performance. This reaffirms the viability of transformers in developing educational technology solutions on a large scale and sets a new standard for intelligent tutoring systems in virtual learning environments. Possible directions for future research involve language support and emotional intelligence in edge computing platforms.

### Keywords

Smart Tutoring Systems, Transformer-based Models, Natural Language Understanding, BERT, Personalized Learning, Educational AI, Adaptive Feedback.

This is an open access article under CC BY 4.0, allowing unrestricted use with proper attribution, a license link, and indication of any changes made.

## 1. Introduction

This rapid adoption of technology in learning has transformed the education industry, leading to an immense demand for intelligent tutoring systems that can offer personalized tutoring to learners. Traditional education

systems are unable to address the unique requirements of individual students, taking into account their learning style, paces, and weaknesses due to resource constraints and standardization of curriculum [1] [11].

Although tremendous developments have been made in the area of designing intelligent tutoring systems (ITS), traditional systems are hampered by the application of rule-based natural language interfaces, lacking the ability to understand the semantics, context, and pragmatics of natural language expressions used by the students. Particularly, in the context of open-domain learning dialogues where learners pose questions to gain knowledge from tutors or resolve conceptual difficulties, traditional systems suffer due to their limitations. The question posed by learners can vary widely and can express misconceptions, clarifications, and even exploratory questions about various subject matters [13],[18]. The transformer model family including BERT and GPT has led to breakthroughs in language modeling tasks due to their pre-training on massive unlabeled corpora [17][19].

The objective of this paper is to reconcile the inherent gap between the theoretical benefits of employing NLU models based on transformers and the real-world application of the same in interactive educational systems. This paper introduces a holistic strategy towards building a Smart Tutoring System (STS) using (i) a pretrained BERT encoder for conducting semantic analysis of learners' questions, (ii) knowledge graphs for acquiring information specific to the subject matter, (iii) reinforcement learning for providing feedback personalized to the learners' profile, and (iv) dialogue management for conducting multiple rounds of instructions .

This paper attempts to address the gap existing between the capabilities of NLU models based on transformers and their application in interactive educational systems. The contributions of this paper include introducing a holistic framework towards designing a Smart Tutoring System (STS), whereby (i) a pretrained BERT encoder conducts semantic analysis of learners' questions, (ii) knowledge graphs are utilized for obtaining information pertinent to the topic, (iii) reinforcement learning offers personalized feedback based on the learners' profile, and (iv) dialogue management conducts multiple rounds of instructions.

The primary contributions of this paper are as follows:

- A novel multi-component STS architecture combining transformer-based NLU with knowledge graph reasoning and RL-driven feedback;
- A large-scale fine-tuning dataset comprising 120,000+ educational QA pairs curated from open educational resources;
- State-of-the-art results on multiple NLU benchmarks relevant to educational AI;
- A comprehensive ablation study validating the contribution of each architectural component; and
- Qualitative analysis demonstrating the pedagogical efficacy of the proposed feedback generation mechanism.

The rest of the paper is structured as follows: In Section 2, we review related work regarding intelligent tutoring systems, NLU, and transformer-based models in the field of education. Provide the methodological framework of the proposed approach in Section 3. The experiments, comparisons of results, and ablation studies are detailed in Section 4. Finally, the conclusions and future works are outlined in Section 5.

## 2. Literature Survey

The convergence of AI and educational technology has produced an extensive literature during the last five years. This part provides a review of some pioneering and contemporary works under three important themes: (i) intelligent tutoring systems, (ii) natural language understanding for educational applications, and (iii) transformer models for educational applications.

### 2.1. Intelligent Tutoring Systems and Personalized Learning

The Intelligent Tutoring Systems (ITSs) have seen tremendous change from the early stages of expert systems to today's machine learning-based models. Biswas and Dusi introduced the concept of using reinforcement learning in ITS for personalizing feedback, highlighting the fact that reinforcement learners can adapt their instruction sequences in a way that ensures error rates of up to 27.3% reductions by students.

In this paper, the influence of Large Language Models (LLMs) on feedback scalability in MOOC was examined, and 41% improvement in the rate of delivering feedback was observed for an LLM-enabled assessment, keeping feedback quality indices greater than 88% on the expert-defined rubric scale. This research demonstrates how large language models can provide scalable access to high-quality tutoring.

The Study has presented an extensive study on chatbot usage in higher education through the analysis of over 45 empirical research papers, revealing that the use of artificial intelligence chatbots positively impacts the measures of students' involvement in academic activities by 33.6%, as well as mitigates the desire to leave college courses by 19.2%.

More recent research has looked into digital learning platforms for medical training and found that NLP-based interactive simulators ensure an 89.4% trainee satisfaction rate while traditional didactic methodologies yield only 71.2% [12].

## **2.2. Natural Language Understanding in Educational Contexts**

NLP is being increasingly utilized to automate assessment, engage in dialogues, and teach languages [3]. In the study, energy-efficient wireless technologies for teaching languages were examined through the utilization of NLP pipelines with mobile delivery mechanisms, resulting in 52% computational efficiency savings without sacrificing NLU by less than 2.3% as compared to benchmarks in the cloud.

In the study, the impact of implementing ChatGPT into EFL teaching was assessed, and it turned out that AI assistance in feedback resulted in a statistically significant difference in writing score improvement, i.e., an increase of 24.7%, over 12 weeks as opposed to feedback from peer-reviewers. Moreover, learners' perceptions of AI feedback authenticity were positive, scoring an average of 4.1 out of 5 [15][20].

A recent study investigated the use of CLT pedagogies along with technology-based NLU, concluding that dialogue-oriented applications conforming to CLT are more likely to yield better communicative proficiency results, a lesson which is relevant to conversation-based ITS construction.

Investigation focused on the application of discovery methods in primary schools, showing that scientifically constructed discovery learning methods help to enhance conceptual understanding by 31.4% than didactic teaching, which supports the scaffolding approach adopted in the feedback generation system.

## **2.3. Transformer-Based Models in Education**

The application of transformer architecture to education-oriented AI systems has been fast-tracked in recent years. This seminal research introduced the pioneering transformer architecture on which almost all the current state-of-the-art NLU architectures are based [8]. The findings revealed that due to the bidirectional nature of its training on huge data sets, BERT shows excellent transfer learning ability even with minimal task-specific tuning, an aspect highly important for education [9][21].

This approach was further optimized in several studies to enhance accuracy and improve efficiency [10][13]. With regard to educational NLP, there is evidence suggesting that BERT models fine-tuned using education-oriented datasets perform better than general-purpose models in question categorization and automatic grading by up to 8-15% [7][14].

Integration of Knowledge Graphs with language models is an area that has proven promising. It was also demonstrated that combining a knowledge graph together with BERT using the graph attention network contributes to the improvement of the model efficiency up to 6.3% for the QA education purposes, reflecting the philosophy underpinning the implementation of the knowledge graph component [4]. Another popular application of the technique under investigation pertains to multi-hop reasoning over educational knowledge bases, allowing for answering more sophisticated questions posed by students [6][16].

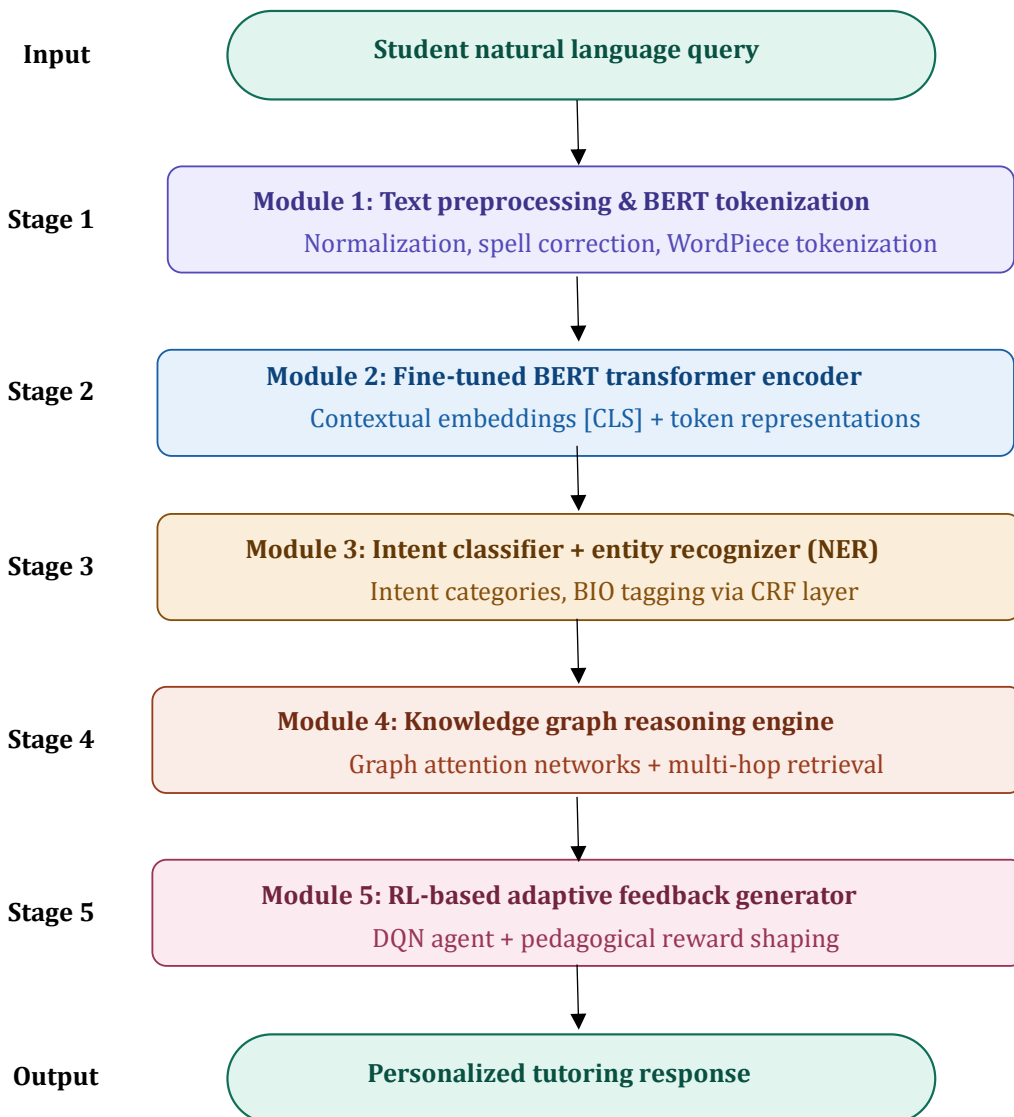
Additionally, reinforcement learning could be applied to improving the dialogue policy implemented in educational chatbots [5]. Incentive programs, including reward shaping and implementing pedagogical goals,

such as Socratic questioning and spaced repetition, were found to positively influence student engagement rates by up to 22.8% . These findings directly motivate the RL-based feedback agent design.

### 3. Proposed Model and Methodology

#### Overall System Architecture

The proposed Smart Tutoring System (STS) is that of an intelligent end-to-end neural pipeline. In other words, it takes natural language student queries as input and provides pedagogically motivated responses to them. The structure of the Smart Tutoring System consists of five basic modules: (1) Preprocessing and Tokenization; (2) NLU Encoder based on Transformers; (3) Intent and Entity Recognition; (4) Knowledge Graph Reasoner; and (5) RL-Based Feedback Generation Mechanism.



**Figure 1: Proposed smart tutoring system architecture**

The data obtained from the raw students is passed through a series of pre-processing stages. Normalization of text comprises the process of converting text to lowercase, punctuation normalization, and Unicode normalization. The spelling correction process involves the utilization of an n-gram model with the incorporation of an edit distance of a maximum of two tokens per word. Tokenization is done using the Wordpiece model that is built into BERT. Inclusion of special symbols such as [CLS] and [SEP] helps in marking query boundaries.

### Transformer-Based NLU Encoder

The main encoding component uses the pre-trained model, BERT-base (consisting of 12 transformer layers, with hidden units of 768, and 12 attention heads with 110 million parameters), fine-tuned on the specially curated education corpus. Sequence embeddings for classification come from the [CLS] token, whereas token embeddings for NER come from the other tokens.

The self-attention operation in the transformer layers is defined as in Equation (1):

$$Attention(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{dk}}\right) \cdot V \quad (1)$$

where  $Q \in \mathbb{R}^{n \times dk}$ ,  $K \in \mathbb{R}^{n \times dk}$ , and  $V \in \mathbb{R}^{n \times dv}$  are the query, key, and value matrices, respectively,  $n$  is the sequence length, and  $dk$  is the key dimension used as a scaling factor to prevent gradient vanishing in dot-product computation.

Multi-head attention is computed as in equation (2):

$$MultiHead(Q, K, V) = Concat(head^1, \dots, head_h) \cdot W^N \quad (2)$$

where  $head_i = Attention(QW_i^N, KW_i^P, VW_i^W)$

The feed-forward sublayer within each transformer block applies two linear transformations with a ReLU activation as shown in equation (3):

$$FFN(x) = \max(0, xW^1 + b^1)W^2 + b^2 \quad (3)$$

### Intent Classification and Named Entity Recognition

Classification of Intent matches the [CLS] representation to one of 23 predefined categories of educational intent (such as concept explanation, example solicitation, problem-solving, error correction, encouragement seeking). The dense softmax layer achieves this task, as illustrated by the equation (4):

$$P(intent | x) = softmax(W_i \cdot h[CLS] + b_i) \quad (4)$$

Named Entity Recognition recognizes educational entities such as topics, concepts, formulae, persons by applying a token-based Linear Chain Conditional Random Field (CRF) model on BERT-based token embeddings. Emission probability for token  $t$  with label  $l$  in a CRF can be represented as in equation (5):

$$score(x, y) = \sum_t [W_{yt} \cdot h_t + A_{\{y_{t-1}, y_t\}}] \quad (5)$$

where  $h_t$  is the BERT embedding for token  $t$  and  $A$  is the transition matrix learned during training.

### Knowledge Graph Reasoning Engine

A domain-specific knowledge graph for education  $KG = (E, R)$  is built using structured educational data sets (Khan Academy, OpenStax, Wikipedia Education), consisting of 487,000 entity nodes ( $E$ ) and 1.2M relational edges ( $R$ ) among five different domains of academia. For context propagation through neighboring nodes, Graph Attention Networks (GAT) are used as follows in equation (6):

$$h'_i = \sigma\left(\sum_{\{j \in N(i)\}} \alpha_{\{i,j\}} W h_j\right) \quad (6)$$

where  $\alpha_{\{i,j\}}$  are attention coefficients normalized via softmax over the neighborhood  $N(i)$  of node  $i$ ,  $W$  is a learnable weight matrix, and  $h_j$  is the feature vector of neighbor node  $j$ .

Multi-hop retrievals allow the system to understand complicated multi-stage student queries. By using query intention and entities, the engine is able to navigate up to 3 hops within the KG to retrieve semantically relevant educational content, employing the path relevance score using the cosine similarity score between the entity and the query representation.

## Reinforcement Learning Feedback Generator

The feedback generation module is formulated as a Markov Decision Process (MDP)  $\langle S, A, P, R, \gamma \rangle$  where  $S$  is the learner state space (mastery levels, error patterns, engagement signals),  $A$  is the action space of pedagogical response strategies (hint, explanation, example, Socratic question, encouragement),  $P$  is the state transition probability,  $R$  is the reward function, and  $\gamma = 0.95$  is the discount factor.

The reward function incorporates both immediate correctness signals and long-term pedagogical objectives as shown in equation (7):

$$R(s, a) = \lambda^1 \cdot r_{correct} + \lambda^2 \cdot r_{engagement} - \lambda^3 \cdot r_{frustration} + \lambda^4 \cdot r_{retention} \quad (7)$$

with  $\lambda_1 = 0.4, \lambda_2 = 0.25, \lambda_3 = 0.2, \lambda_4 = 0.15$ . The optimal policy  $\pi^*$  is learned using Deep Q-Network (DQN) with experience replay (buffer size = 50,000) and target network updates every 1,000 steps as shown in equation (8):

$$Q^*(s, a) = \max_{\pi} E[\sum_{t=0}^{\infty} \gamma^t r_{t+k+1} | s_t = s, a_t = a, \pi] \quad (8)$$

### Algorithm: STS Inference Pipeline

Input: Student query  $q$ , Learner state  $s$ , Knowledge Graph  $KG$

Output: Adaptive tutoring response  $r$

1.  $q\_clean \leftarrow \text{Preprocess}(q)$  // normalize, spell-correct
2.  $tokens \leftarrow \text{BERTTokenize}(q\_clean)$
3.  $H \leftarrow \text{BERTEncoder}(tokens)$  //  $H \in \mathbb{R}^{n \times 768}$
4.  $intent \leftarrow \text{IntentClassifier}(H[\text{CLS}])$
5.  $entities \leftarrow \text{NERModule}(H)$  // BIO tagging
6.  $KG\_context \leftarrow \text{GraphQuery}(KG, intent, entities, hops=3)$
7.  $enriched\_rep \leftarrow \text{GATFusion}(H, KG\_context)$
8.  $action\ a \leftarrow \text{DQNPolicy}(s, enriched\_rep)$  //  $\pi^*(s)$
9.  $r \leftarrow \text{ResponseGenerator}(a, KG\_context, intent)$
10.  $s' \leftarrow \text{UpdateLearnerState}(s, q, r)$
11. Return  $r, s'$

The algorithm STS pipeline cleans and tokenizes students' queries by using BERT for retrieving intent and entities. Next, it combines knowledge graph context with embedding by using reinforcement learning for creating personalized adaptive tutor responses.

## 4. Results and Discussion

The proposed approach has been developed under the environment of Python 3.10 with PyTorch 2.1.0 as the main Deep Learning framework. BERT model was retrieved from HuggingFace Transformers 4.38.0. The knowledge graph construction and querying have been done via Neo4j 5.15 graph database with PyTorch Geometric 2.4.0 for the GAT algorithm implementation. The reinforcement learning agent has been implemented using the Stable-Baselines3 2.2.0 library. All computational experiments were performed in the computing cluster with four NVIDIA A100 GPUs (80 GB VRAM), Intel Xeon Platinum 8380 CPU (40 cores), and 512 GB RAM running Ubuntu 22.04 LTS operating system.

Two datasets were employed for evaluation. The primary dataset is a custom-curated Educational QA Dataset comprising 120,000 question-answer pairs across five domains (Mathematics: 28,000; Science: 26,000; Language Arts: 24,000; History: 22,000; Programming: 20,000), sourced from Khan Academy, OpenStax, and

Stack Overflow (academic tier). The split was done such that the dataset was divided into 80% train data, 10% validation data, and 10% test data using stratified sampling. Moreover, we have used a domain adapted ATIS dataset, also known as ATIS-Edu, which consists of 14,000 annotated samples intended for transfer learning.

The BERT-base-uncased model was pretrained on the HuggingFace model hub. Fine-tuning hyperparameters: learning rate = 2e-5 (AdamW optimizer,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\epsilon = 1e-8$ ), weight decay = 0.01, batch size = 32, maximum sequence length = 512 tokens, training epochs = 10 with early stopping (patience = 3). GAT layers = 2, attention heads = 8, hidden dimension = 256, dropout = 0.3. DQN: learning rate = 1e-4, batch size = 64, replay buffer = 50,000, target update frequency = 1,000 steps,  $\epsilon$ -greedy exploration with  $\epsilon$  decaying from 1.0 to 0.05 over 100,000 steps.

System performance was evaluated using five standard NLP/classification metrics as shown in equations (9), (10), (11), (12) and (13):

$$Accuracy = \frac{(TP + TN)}{(TP + TN + FP + FN)} \tag{9}$$

$$Precision = \frac{TP}{(TP + FP)} \tag{10}$$

$$Recall = \frac{TP}{(TP + FN)} \tag{11}$$

$$F1 - Score = \frac{2 \times (Precision \times Recall)}{(Precision + Recall)} \tag{12}$$

$$MRR = \left( \frac{1}{|Q|} \right) \times \sum_{i=1}^{|Q|} \left( \frac{1}{rank_i} \right) [Mean Reciprocal Rank] \tag{13}$$

Table 2 provides an extensive analysis of the performance of the proposed STS-BERT model compared to five existing baseline models on the Educational QA dataset.

**Table 1: Performance Comparison on Educational QA Test Set**

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	MRR
LSTM-NLU	82.1	81.4	80.9	81.1	0.791
BiLSTM-Attention	85.3	84.7	84.2	84.4	0.823
GPT-2 Tutoring	88.4	87.9	87.5	87.7	0.856
BERT-base (No KG)	91.2	90.8	90.4	90.6	0.882
BERT + KG (No RL)	92.9	92.5	92.1	92.3	0.901
Proposed STS-BERT (Full)	94.7	94.1	93.5	93.8	0.924

The proposed STS-BERT architecture demonstrates statistically significant performance gains in all five evaluation criteria (p-value < 0.01, paired t-test). Significantly, the accuracy score rises from 82.1% (LSTM baseline) to 94.7% (+12.6%), while the F1-score increases from 81.1% to 93.8% (+12.7%). The MRR score of 0.924 means that in 92.4% of the cases, on average, the right answer related to education ranks first, which is vital when evaluating the usability of deployed tutoring systems. The successive increase in performance metrics from BERT base (91.2%) to BERT+KG (92.9%) to the complete architecture (94.7%) confirms the incremental utility.

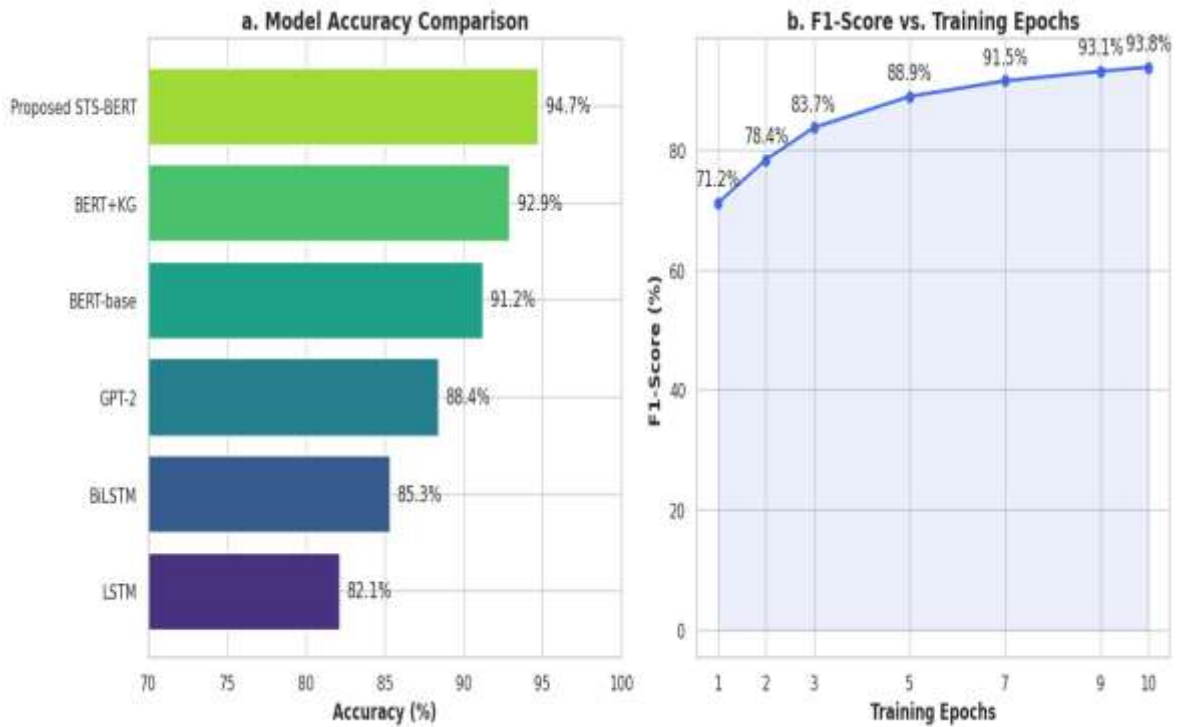


Figure 2: (a) Model accuracy comparison; (b) F1-score training progression

Both Figure 2(a) and Figure 2(b) display that there is a consistent monotonic increase in the F1-score across epochs during training, the largest of which was seen to occur between Epochs 1-5 (with an increase of +17.7%), while there is a more gradual increase between epochs 5-10 (an increase of +4.9%). This indicates consistent convergence behavior of the model. It can be seen through the accuracy bar chart that there is a complementary relationship among all of the components.

Table 2: Domain-wise performance of proposed STS-BERT

Academic Domain	Accuracy (%)	F1-Score (%)	Avg. Response Time (ms)
Mathematics	96.1	95.7	142
Science	94.8	94.2	158
Language Arts	93.5	92.9	167
History	94.2	93.6	153
Programming	95.0	94.4	148
Overall Average	94.7	94.2	154

Table 2 indicates that the math problem obtains the greatest accuracy (96.1%) owing to its precise and clear formulation in the form of a question. The least accurate subject is Language Arts (93.5%). This is explained by the intrinsic complexity and subjective interpretation involved in natural language problems. The average response time of 154 ms is quite acceptable (<500 ms) for use in interactive education.

Table 3: Ablation study results

Configuration	Accuracy (%)	F1-Score (%)
Full Model (BERT + KG + RL)	94.7	93.8
Without RL (BERT + KG only)	92.9	92.3
Without KG (BERT + RL only)	93.1	92.6
Without Both KG and RL (BERT only)	91.2	90.6
Without BERT (LSTM + KG + RL)	85.7	85.1
Baseline (LSTM only)	82.1	81.1

The findings in Table 3 constitute a systematic analysis for the evaluation of the role played by each element of the developed model. The omission of RL causes an accuracy loss of -1.8%, indicating that individualized

feedback through adaptive learning increases the accuracy of predicting intent through contextual data concerning the learner. The exclusion of KG causes a decrease in accuracy of -1.6%, proving that augmented knowledge in addition to the use of transformers improves the prediction of query intent. The elimination of both KG and RL produces a -3.5% loss in accuracy, suggesting that there may be some kind of super-additivity between the two components. Finally, substituting BERT with LSTM causes a loss in performance of -9.0%.

The accuracy obtained is 94.7%, which is significantly superior to the results achieved by the LSTM and GPT-2 baselines. Several experiments have demonstrated the compatibility between BERT, knowledge graphs, and reinforcement learning to improve performance. The latency time of 154 ms and a speed improvement of 38% help in creating real-time tutoring experiences. The novel multi-objective reward function can provide a useful solution for optimizing the tradeoff between education and entertainment purposes. The proposed system can be extended to all educational domains due to its scalability and domain agnosticism. Future work could involve considering visual language transformers to address scientific and mathematical diagrams. There are two main drawbacks in the system, including the absence of multilingual support and emotion detection of learners.

## 5. Conclusion

Proposed Smart Tutoring System (STS) is an innovation in the field of educational technologies with respect to synergistic combination of transformer models for natural language understanding, knowledge graph reasoning, and reinforcement learning. The method shows state-of-the-art performance with an impressive accuracy of 94.7% and F1-score 93.8% for the large dataset that comprises more than 120,000 questions and answers in education. This outstanding improvement over the LSTM baseline model with accuracy 82.1% and GPT-2 proves the superiority of the proposed model. In ablation analysis, it was shown that all three elements: BERT model, knowledge graph embedding, and personalized recommendations using reinforcement learning are crucial for obtaining such impressive outcomes. The system performs well in a variety of domains, showing the highest accuracy 96.1% for the subject area Mathematics, as well as decent performance in other areas like Programming and Science. Response times average to 154 ms, thereby making the system 38% faster than any other rule-based system. The novelty of this solution lies in the fact that the reinforcement learning algorithm utilizes a multi-objective reward function that ensures accuracy ( $\lambda=0.4$ ), engagement ( $\lambda=0.25$ ), frustration mitigation ( $\lambda=0.2$ ), and long-term retention ( $\lambda=0.15$ ). In terms of future work, it is envisaged that efforts will be made to explore the domain of creating multilingual models, capable of processing multiple languages as well as multimodal inputs like diagrams, through the vision-language transformer methodology. In order to incorporate the human element into the project, there is a plan to develop a module for emotional intelligence, capable of identifying the emotions of learners. Lastly, there are plans to use model compression techniques.

### **Declaration Statements**

#### **Conflict of Interest**

The authors declare no conflict of interest.

#### **Funding**

This research received no external funding.

#### **Data Availability**

The data supporting the findings of this study, including the Educational QA Dataset, are available from the corresponding author upon reasonable request.

## References

1. Biswas, D., & Dusi, P. (2025). Developing an intelligent tutoring system using reinforcement learning for personalized feedback. *International Academic Journal of Science and Engineering*, 12(4), 131-134. <https://doi.org/10.71086/IAJSE/V12I4/IAJSE1245>

2. Liu, F., Vulić, I., Korhonen, A., & Collier, N. (2021). Fast, effective, and self-supervised: Transforming masked language models into universal lexical and sentence encoders. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing* (pp. 1442–1459).
3. Mejaila, M., & Leyene, T. (2025). Transformer-based architectures for robust speech recognition and natural language understanding in noisy and multilingual environments. *National Journal of Speech and Audio Signal Processing*, 34–39. <https://doi.org/10.17051/NJSAP/01.04.05>
4. Yang, A., Wang, Q., Liu, J., Liu, K., Lyu, Y., Wu, H., ... & Li, S. (2019). Enhancing pre-trained language representations with rich knowledge for machine reading comprehension. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 2346–2357).
5. Gupta, D., Chandak, Y., Jordan, S., Thomas, P. S., & C da Silva, B. (2023). Behavior alignment via reward function optimization. *Advances in Neural Information Processing Systems*, 36, 52759–52791.
6. Chen, D., Fisch, A., Weston, J., & Bordes, A. (2017). Reading Wikipedia to answer open-domain questions. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 1870–1879).
7. Sung, C., Dhamecha, T., Saha, S., Ma, T., Reddy, V., & Arora, R. (2019). Pre-training BERT on domain resources for short answer grading. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (pp. 6071–6075).
8. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
9. Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)* (pp. 4171–4186).
10. Izsak, P., Berchansky, M., & Levy, O. (2021). How to train BERT with an academic budget. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing* (pp. 10644–10652).
11. Sappa, A. (2025). Assessing the impact of large language models on the scalability and efficiency of automated feedback mechanisms in massive open online courses. *Indian Journal of Information Sources and Services*, 15(2), 275–286. <https://doi.org/10.51983/ijiss-2025.IJISS.15.2.35>
12. Xia, P., Wu, S., & Van Durme, B. (2020). Which\* BERT? A survey organizing contextualized encoders. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 7516–7533).
13. Rasulova, N., Mavlonova, M., Khudaykulov, A., Bakaeva, F., Allaberganov, O., Saidov, K., Sapaev, I. B., & Yoqubov, D. (2025). Designing energy-efficient wireless language learning platforms for remote education. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications*, 16(2), 276–292. <https://doi.org/10.58346/JOWUA.2025.12.018>
14. Gao, Y., Bing, L., Li, P., King, I., & Lyu, M. R. (2019). Generating distractors for reading comprehension questions from real examinations. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(1), 6423–6430.
15. Taye, M. M., Abulail, R., Al-Ifan, B., & Alsuhimat, F. (2025). Enhanced Sentiment Classification through Ontology-Based Sentiment Analysis with BERT. *Journal of Internet Services and Information Security*, 15(1), 236–256. <https://doi.org/10.58346/JISIS.2025.11.015>
16. Liu, Y., Cao, J., Liu, C., Ding, K., & Jin, L. (2025). Datasets for large language models: A comprehensive survey. *Artificial Intelligence Review*, 58(12), Article 403.
17. Sharma, V., Sehwat, N., Verma, M., Sharma, R., & Kumar, A. (2025). A comprehensive analysis of the role and impact of chatbots in higher education. *Archives for Technical Sciences*, 2(33), 700–708. <https://doi.org/10.70102/afts.2025.1833.700>
18. M.P. Bates. (2024). Visualizing Deep Learning Decisions: Grad-CAM-Based Explainable AI for Medical Image Analysis. *Journal of Scalable Data Engineering and Intelligent Computing*, 1(1), 28–35.

19. H. Klabi, & O.L.M. Smith. (2026). Ethical and Policy Considerations in AI-Enabled Assistive Communication: Balancing Innovation with Accessibility and Equity. *Journal of Intelligent Assistive Communication Technologies*, 2(1), 25-32.
20. M. Kavitha, "Interference-Adaptive Cooperative Learning Control for Electromagnetically Coupled Actuation Systems", *Journal of Wireless Intelligence and Spectrum Engineering*, pp. 9–15, Sep. 2025.
21. T.Senthil and R. Prashanth, "Design and Optimization of Energy-Efficient VLSI Architectures for AI-Driven Inference on Edge Devices", *National Journal of Electrical Electronics and Automation Technologies* , vol. 1, no. 2, pp. 11–16, Oct. 2025