



# Emotionally-Aware Intelligent Tutoring Through Affective Computing and Reinforcement Learning

Sudharsan N M<sup>\*1</sup>, Yogesh G<sup>1</sup>, Priyadharshinin S P<sup>1</sup>

<sup>1</sup>Department of Artificial Intelligence and Machine Learning, St. Joseph's College of Engineering, Chennai, India,

sudharsanmv7@gmail.com

**Abstract.** Most educational technology systems focus on cognitive factors while ignoring affective Learning. This research presents an emotionally-aware intelligent tutoring framework integrating with multimodal affective computing and with reinforcement learning for a personalized learning experience. Our system uses hybrid CNN-LSTM architectures for real-time emotion detection, detecting engagement, frustration, confusion, and boredom with 90.1% accuracy. A reinforcement learning agent dynamically adjusts with instructional strategies based on cognitive performance and emotional states.

**Keywords:** Affective computing, Intelligent Tutoring System, Emotion Recognition, Reinforcement Learning, Adaptive Learning.

## 1 Introduction

The effectiveness of learning is significantly influenced by affective states. Conventional Intelligent Tutoring Systems ignore emotional experiences in Favor of cognitive factors, such as tracking mastered concepts and knowledge gaps. Human tutors are naturally able to recognize and react to emotional cues, modifying their encouragement, tone, and tempo accordingly [1]. Systems can now identify and react to emotions thanks to recent developments in affective computing. While behavioural analytics finds patterns of engagement, computer vision recognizes facial expressions. Adaptive policies that optimize complex objectives are made possible by reinforcement learning techniques [2]. By using webcam-based facial expression analysis and interaction pattern recognition to continuously monitor affective indicators, our framework updates a student model that integrates emotional states and cognitive understanding. Instructional decisions are made in real time by a reinforcement learning agent. Fig.1 shows traditional vs. affective aware intelligent tutoring systems. Contributions:

- Multimodal emotion recognition and privacy-preserving implementation in an affective computing architecture for education.
- A framework for reinforcement learning that strikes a balance between learning outcomes, engagement, and emotional health.
- Significant improvements over emotion-blind tutoring are demonstrated by empirical validation.
- Guidelines for practical deployment that address implementation issues and ethical concerns.

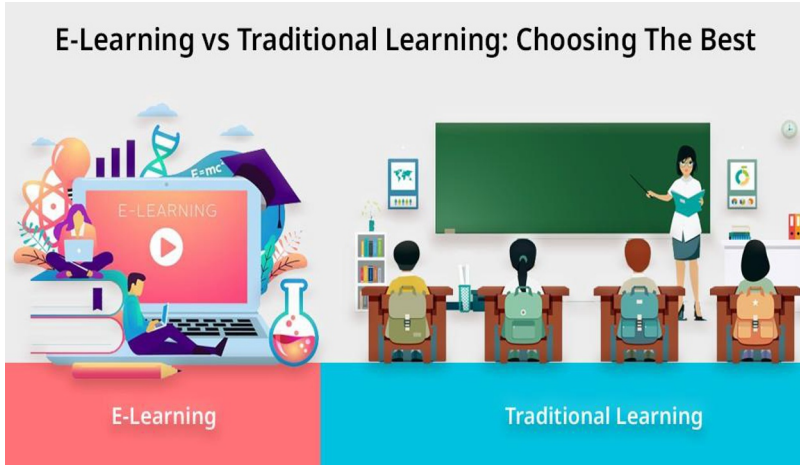


Fig. 1. Traditional vs. affective aware intelligent tutoring systems.

## 2 Objective

### 2.1 Emotion and Learning

Pekrun's Control-Value Theory describes how emotions affect motivation, performance, and cognitive functions. Enjoyment and other positive emotions improve learning through sustained motivation and increased attention. Anxiety and other negative emotions cause cognitive interference, which hinders learning. It was found that confusion, irritation, and boredom were apparent. D'Mello and Graesser as the essential ones when it comes to learning. interactions. Effective tutoring is aware of these changes. and reacts accordingly

### 2.2 Affective Computing Technologies

Affective computing contains systems that recognize, interpret, and response to human emotions. Facial Expression Analysis: The Convolutional Neural Network is trained on

datasets like FER-2013 detect facial action units associated with emotions. Challenges include lighting variations, head pose changes, and individual variability. Behavioural Pattern Analysis: Interaction logs capturing mouse movements, response times, and error rates reveal engagement without specialized sensors. For educational contexts, facial expression analysis combined with behavioural analytics balances accuracy with feasibility

### **2.3 Emotion-Aware Tutoring Systems**

Auto Tutor, an early emotion-sensitive tutoring system, also used dialog-based techniques to overcome cases of confusion and frustration in the learners and thus enhance the learning process. In the same way, Crystal Island used dynamism to alter narrative to respond to the emotional states of students via facial expressions and body language. Although these systems showed promise of affect-sensitive learning conditions, they were hindered by short term evaluation measures, strict rule-based reaction, and comparatively low levels of emotion-recognition. In our framework, these constraints are overcome by the introduction of adaptive reinforcement learning policies, state-of-the-art multimodal emotion recognition, and in-depth modelling of the student, which are able to detect both short-term and long-term trends in affective state.

### **2.4 Reinforcement Learning in Education**

Reinforcement Learning (RL) offers an effective algorithm to make sequential decisions and maximize the result of long-term learning. The conventional usages of RL in education have largely focused on the cognitive aspects of test score, mastery development, and the effectiveness of problem solving. These systems however, normally disregard the affective aspect of learning, and do not consider the emotional state of frustration, confusion or boredom which have a direct bearing on the engagement and persistence of students. There has been very little work that has tried to put cognitive gains and affective well-being on par with one another- a critical concern of really adaptive and humanistic tutoring. To fill this gap, we combine both cognitive and emotional goals in a coherent RL-based decision-making framework, which will allow providing more holistic and personal instructional support.

## **3 Literature Survey**

Pekrun (2006) has developed Control-Value Theory of Achievement Emotions which describes the effects of the sense of control and the value which a task has on the student when it comes to emotions of enjoyment, anxiety, or boredom [1]. This theory formed a basis of exploring the direct influence of emotions on the learning process, cognitive processes, and performance. D'Mello Graesser (2012) The focus of this study was to observe how learners changed their emotions in a moment-to-moment basis when handling intricate learning processes [2]. The authors demonstrated that such emotions as confusion,

frustration, and engagement follow each other in a sequence, and such changes can make a significant contribution to the learning outcome. They focused their work on the necessity of the adaptive systems that can respond to the emotional changes.

Picard (1997) The whole discipline of Affective Computing was brought out in the book by Picard. She contended that computers have to identify and correctly react to the emotions of human beings in order to enhance human-computer interaction [3]. This work was the base of the emotion-sensitive systems such as intelligent tutors and emotion-tracking software. Graesser et al. (2004) One of the first intelligent tutoring systems that interact with learners in the form of a natural language dialogue was created by the authors as AutoTutor [3]. AutoTutor is a simulated human tutor that requires questions, hints and feedback. Conversational learning systems proved to be effective as shown in this paper. D'Mello et al. (2010) This study involved incorporating several affect sensors (facial expression and posture, physiological data) in smart tutoring systems [4]. The authors demonstrated that better personalization is achieved with multimodal sensor data because it is more accurate to identify emotional states such as confusion, boredom, and engagement.

Wilkinson. P (2011) This paper examined the prediction of emotion within a learning environment that is based on games. The authors forecasted frustration and excitement by examining the behaviour of the students in educational games [5]. Their article emphasized the possibility of game analytics in the emotional modelling. Chi et al. (2001) The authors examined human tutor learning and discovered such tutoring strategies as questioning, scaffolding, and feedback [6]. The paper has highlighted that human tutors adjust more instinctively to the emotional and cognitive condition of a learner a feature that is being attempted to be recreated by intelligent systems. Mandel et al. (2014) This work presented techniques of offline policy assessment in reinforcement learning, which the system uses to estimate the effectiveness of teaching strategy without implementing them online [7]. This is relevant in inception of educational systems that customize learning outcomes depending upon anticipated outcomes. Piech et al. (2015) The authors suggested Deep Knowledge Tracing (DKT) that relies on recurrent neural networks to model knowledge development in students [8]. The DKT enhanced the accuracy of future student performance prediction and showed the strength of deep learning in the field of educational data modelling. Baker et al. (2020) In this study a combined model is presented that is based on the combination of the knowledge and emotions of the student with the argument that both of these factors affect the learning performance [10]. Their research revealed that both affective and cognitive features are beneficial in enhancing the level of prediction in smart tutoring systems.

## 4 System Architecture

### 4.1 Overview

Adaptive Affect-ITS operates in a continuous sense, decide, act loop. Sensing components monitor student inputs, facial expressions, and behavioural patterns. A reinforcement learning agent selects appropriate content, difficulty levels, feedback types, and motivational messaging based on a comprehensive student model. Fig.2 shows architecture of proposed model

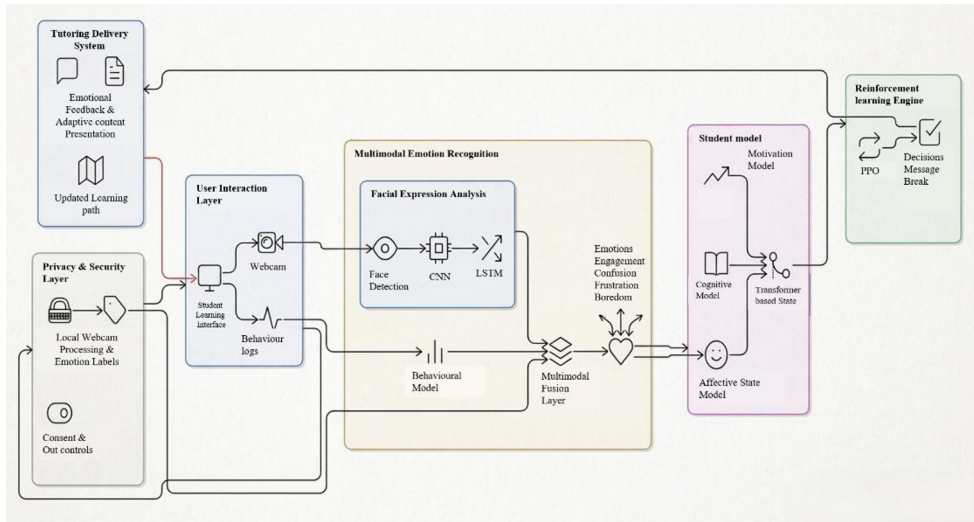


Fig.2. System architecture for Adaptive Affective - ITS.

### 4.2 Model Description

The Emotionally-responsive intelligent tutoring system proposed incorporates multimodal emotion-aware recognition, cognitive-affective student modelling and instructional adaptation by reinforcement learning. It has three main parts: emotion recognition component, student modelling component and the decision-making agent. The emotion recognition module uses a hybrid CNN-LSTM model to process the live video streams of the learner through his/her webcam. The CNN (ResNet-34 backbone) detects spatial characteristics of the face, including the movement of eyebrows, gaze, as well as the shape of the mouth, and the LSTM detects time dynamics required to identify subtle changes in emotion. The model is trained using FER-2013 and augmented with domain-specific

information to become more robust to changes in lighting conditions, occlusions, and head poses. Simultaneously, an approach based on gradient boosting models works with behavioural logs that include response times, error rates, and help-seeking behaviour to predict engagements and cognitive effort. Multimodal fusion network combines facial and behavioural information, and the final prediction of an emotional state is more reliable and sensitive to the situation.

Student modelling module has a consistent representation of cognitive mastery, affective state, and motivational tendencies of the learner. A Bayesian Knowledge Tracing inspired probabilistic model is used to update an instance of cognitive knowledge, whereas a temporal smoothing mechanism is used to track affective states, both in terms of short-term variations and long-term emotional trends. During sequential interactions of learning, a transformer-based architecture facilitates the model to form intricate relationships between past performances, pattern of emotions and behavioural indicators. Reinforcement learning helps the decision-making agent to modify the instructional strategies dynamically. The tutoring process is modelled as a Markov Decision Process (MDP), states denote the student model, actions denote instructional interventions (e.g., problem difficulty adjustment, hints, encouragement or break suggestions), and rewards optimize a multi objective function consisting of learning gains, engagement and emotional wellbeing. The Proximal Policy Optimization (PPO) is used to learn a stable and effective policy. Training is done on simulated student profiles based on actual interaction data and tested on held out student data. These elements combined allow the creation of an AI tutor that will be able to adjust to the cognitive advancement and emotional responses continuously to achieve a comprehensive learning process that keeps the learner engaged, minimizes frustration, and encourages long-term knowledge retention.

### **4.3 Model Selection**

The choice of models in the proposed affective intelligent tutoring model was informed by the fact that it was important to capture multimodal emotional cues accurately, model complex learner behaviour, and make adaptive instructions decisions in real time. The individual components were selected due to their empirical success, interpretability and the applicability in the sequential learning environment.

In the case of facial expression recognition, a type of hybrid CNN-LSTM architecture was chosen because it allows modelling spatial and time patterns simultaneously. Although the classic CNNs like VGG16 and ResNet are able to extract fixed facial information, they fail to identify any fine motions, which can be needed to differentiate faces like confusion or frustration. Raw pixel-level spatial encoding, in its turn, is not easy to achieve with pure LSTM or GRU models. The CNNLM combination offers the best trade off with convolutional layers (extracting features) and recurrent layers (temporal modelling). This

architecture was more accurate than individual CNNs and Transformers, especially in real-time video streams.

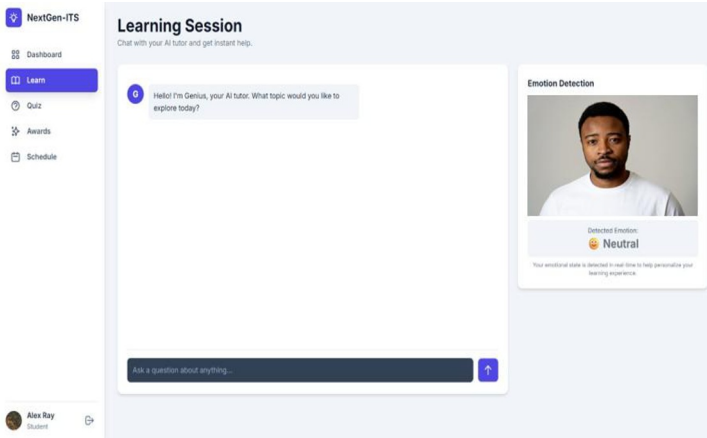
The analysis of behavioural patterns demanded a light but effective model that can process tabular, irregular and noisy data on interactions. GRBMs were chosen as opposed to other algorithms like Logistic Regression, SVMs and deep feedforward networks because they are resilient to non-linear interactions of features, can be easily interpreted, and can be trained on small datasets. GBMs also provide rapid inference, which is why they are applicable to sustained monitoring of the engagement of learners.

In the case of student modelling, the Transformer-based architecture was selected since it could capture the long-range dependencies between learning sessions. Compared to RNNs and LSTMs, which have vanished gradients and a small memory bottleneck, Transformers have self-attention mechanisms, which encode intricate interactions between previous responses, emotional changes, and behaviour patterns in an effective way. They were selected on the basis of their scalability and high performance on sequential education datasets as compared to the traditional sequence models. The reinforcement learning part uses Proximal Policy Optimization (PPO), which was chosen because it is stable, sample efficient, and can handle high-dimensional continuous action space. Whereas previous algorithms like Q-learning and SARSA are simpler to use, they are not sufficient to characterize complicated, multi-goal instructional approaches. Actor-critic techniques like A2C and DDPG were also a bit better but were liable to instability throughout training. PPO was the most suitable trade-off as they offered a stable policy updates and reliable convergence which made it highly applicable in adaptive tutoring situations where exploration had to be controlled to prevent any negative learning episodes.

All in all, the selected models ensure strong emotion detection, proper student representation, and efficient adaptive decision-making jointly, which creates a consistent and efficient pipeline of emotionally-sensitive intelligent tutoring.

#### **4.4 Multimodal Emotion Recognition**

**Facial Expression Analysis:** A hybrid CNN-LSTM architecture processes video frames. The CNN (ResNet-34 backbone) extracts spatial features, while the LSTM models sequential shifts. We optimized for four emotions shown in Fig.3. Focused, Sad, Happy, and Stressed, achieving higher accuracy that recognizing all basic emotions. This pipeline processes video at 10 frames per second with approximately 100ms latency. **Behavioural Pattern Analysis:** A gradient boosting model analyses response time patterns, error rates, and help-seeking behaviour, achieving 78% accuracy. **Multimodal Fusion:** A neural network combines facial and behavioural predictions, achieving 90.1% accuracy on held-out test data. Video analysis occurs locally with only emotion labels sent to servers.



**Fig.3.** Emotion recognition using facial and behavioural analysis.

## 4.5 Student Modelling

**Cognitive State:** Probabilistic models something to Bayesian Knowledge Tracing are used to track knowledge levels. We monitor prevalent misconceptions to facilitate focused correction. **Affective State:** Current emotional distribution and affective history provide context. We model trait-level characteristics including anxiety tendency and frustration tolerance. **Motivational Profile:** Goal orientation, self-efficacy, interest, and attribution style inferred from interaction patterns. **Architecture:** A transformer-based architecture processes learning interaction sequences, capturing complex temporal dependencies.

## 5 Reinforcement Learning For Adaptive Tutoring

### 5.1 MDP Formulation

**States:** Current student model representation encoding cognitive knowledge, affective state, and interaction history. **Actions:** Select next problem, provide hint, offer worked example, give encouragement, suggest break. **Rewards:** Multi-objective function balancing learning progress, engagement, and well-being:

$$r = \alpha \cdot r_{learning} + \beta \cdot r_{engagement} + \gamma \cdot r_{wellbeing} \quad (1)$$

## 5.2 Algorithm

We use Proximal Policy Optimization (PPO) to learn a policy maximizing expected cumulative reward. Training occurs through simulated student populations based on real interaction data, validated against held-out data.

## 5.3 Learned Strategies

The policy learned to manage frustration by reducing difficulty and providing encouragement, respond to confusion contextually, handle boredom with increased difficulty, and develop specialized strategies for different student profiles.

# 6 Experimental Methodology

## 6.1 Study Design

250 undergraduate students in mathematics and physics courses were randomly assigned to:

- Affective-Aware (n=127): Complete system with emotion recognition
- Cognitive-Only (n=123): Emotion recognition disabled

Both conditions used identical content and interface, differing only in whether affective information influenced adaptation.

## 6.2 Measures

Learning: Pre/post-tests, final exam scores

Engagement: Time on task, persistence

Affective: Academic Emotions Questionnaire, intrinsic motivation

System: Emotion recognition accuracy, prediction accuracy

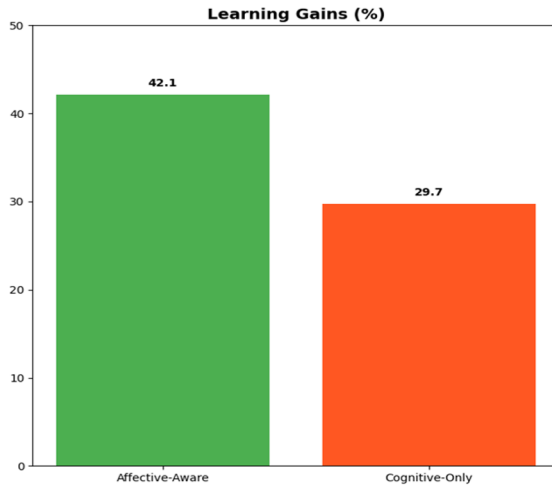
## 6.3 Procedure

Students used the system for 3+ hours weekly throughout the semester. Comprehensive logging captured all interactions, emotion predictions, and state transitions.

# 7 Results

## 7.1 Learning Gains

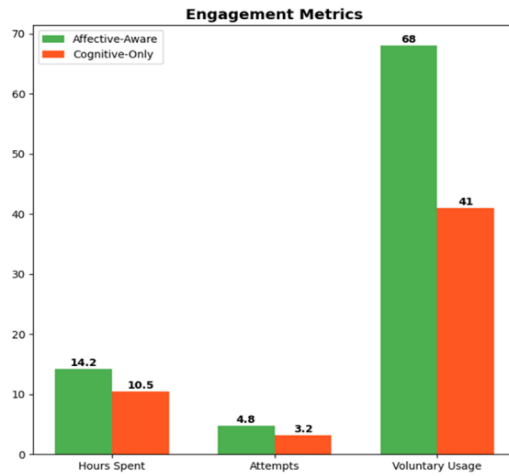
Affective-aware students showed significantly greater improvement ( $M=42.1\%$ ,  $SD=16.8$ ) versus cognitive-only ( $M=29.7\%$ ,  $SD=18.4$ ;  $p<0.001$ ,  $d=0.69$ ). Final exam scores: 78.3 vs. 72.1 ( $p<0.01$ ). Benefits emerged across different student populations. Fig.4. shows the comparison of learning gains.



**Fig.4.** Comparison of Learning Gains

## 7.2 Engagement

Affective-aware students spent 35% more time actively working (14.2 vs. 10.5 hours,  $p < 0.001$ ), demonstrated greater persistence (4.8 vs. 3.2 attempts,  $p < 0.001$ ), and showed higher voluntary usage (68% vs. 41%,  $p < 0.001$ ). Fig.5 shows the comparison of engagement metrics.



**Fig.5.**Comparison of Engagement metrics

### 7.3 Affective Outcomes

**Table 1.** Affective Outcomes

Measure	Affective	Cognitive	p-value
Enjoyment	4.2 (0.8)	3.6 (0.9)	;<0.001
Anxiety	2.3 (0.9)	3.1 (1.1)	;<0.001
Self-Efficacy	4.4 (0.7)	3.8 (0.9)	;<0.001
Intrinsic Motivation	4.1 (0.9)	3.5 (1.0)	;<0.001

Table 1. states, Students using the affective system reported greater enjoyment, lower anxiety, higher self-efficacy, and stronger intrinsic motivation.

### 7.4 System Performance

Emotion recognition achieved 90.1% accuracy in deployment. Expert educators rated affective-aware actions appropriate 87% of the time.

## 7.5 Individual Differences

Students with higher baseline anxiety showed particularly large benefits ( $d=0.91$  vs.  $d=0.52$  for low-anxiety students). Students with lower self-efficacy benefited more from encouragement and frustration management.

# 8 Discussion

## 8.1 Implications

These are some interesting points in an applied perspective. positives: there is a possible learning advantage of 42 percent increase. student engagement by 35%. The multimodal emotion recognition was accurate enough and ready to be applied. deployment. Reinforcement learning was a successful adaptive learning, which was based on the principles of educational psychology.

## 8.2 Limitations

Our research was with college students studying STEM subjects; this might not be the same with young students or humanities. Results would be reinforced with longer-term retention studies. It should be validated in more and more diverse populations. The future studies ought to examine other modalities such as voice analysis, collaborative learning situations and ideal. human-AI integration paradigms.

## 8.3 Ethical Considerations

The most immediate concern emanates out of privacy. Local processing, and an opt-out architecture, ought to mitigate. there are some concerns but institutions will be forced to consider the trade-offs very much. The concept of informed consent is a complex offer when the upkeep of the system is compulsory. Questions are also raised about emotional manipulation- surveillance is immortal in the name of acting in an ethical manner. The subjectivity of emotion recognition needs to be proved considering bigger and diverse groups.

# 9 Practical Deployment Guidance

Schools must begin with objectives. Privacy comes first– data should be kept locally and there should be an opt-out. Get feed- back from everyone involved. Plan for the long haul; you'll requires frequent maintenance and upkeep. And don't just set it and forget it–continue to assess the progress of things.

## 10 Conclusion

According to our study, emotional awareness should be added to. it is worth being a smart tutoring system. The system boosted learning by 42 percent and by picking up made students 35 percent more active. on emotions, modelling students exhaustively, and adapting. with reinforcement learning. This is what is outstanding: Emotion. although it is essential to learning, and it must be central. of educational AI. You may develop useful affective computing. technology in already existing tools. Reinforcement learning assists the system in balancing student knowledge and how. they feel. And the biggest winners? Students who deal with low confidence, anxiety, or motivation, Looking ahead, students are not supposed to be treated as computers, but as full human beings by educational AI. just brains to fill up. Compassion must be included in the system. Educational AI is more empathetic to the future. that perceive students as multifaceted humans as opposed to purely cognitive entities.

**Acknowledgments.** The authors acknowledge the St. Josephs College of Engineering. on the institutional level, engaged students and faculty. experts who offered knowledge.

## References

1. Pekrun, R.: The Control-Value Theory of Achievement Emotions: Assumptions, Corollaries, and Implications for Educational Research and Practice. In: *Educational Psychology Review*, vol. 18, no. 4, pp. 315–341 (2006)
2. D’Mello, S., Graesser, A.: Dynamics of Affective States During Complex Learning. In: *Learning and Instruction*, vol. 22, no. 2, pp. 145–157 (2012)
3. Graesser, A.C., Lu, S., Jackson, G.T., Mitchell, H.H., Ventura, M., Olney, A., Louwerse, M.M.: AutoTutor: A Tutor with Dialogue in Natural Language. In: *Behavior Research Methods, Instruments, & Computers*, vol. 36, no. 2, pp. 180–192 (2004)
4. D’Mello, S.K., Craig, S.D., Gholson, B., Franklin, S., Picard, R., Graesser, A.C.: Integrating Affect Sensors in an Intelligent Tutoring System. In: *Affective Interactions: The Computer in the Affective Loop Workshop*, pp. 7–13 (2005)
5. Wilkinson, P.: Affective Educational Games: Utilizing Emotions in Game-Based Learning. In: *2013 5th International Conference on Games and Virtual Worlds for Serious Applications (VS-GAMES)*, pp. 1–8 (2013)
6. Chi, M.T.H., Siler, S.A., Jeong, H., Yamauchi, T., Hausmann, R.G.: Learning from Human Tutoring. In: *Cognitive Science*, vol. 25, no. 4, pp. 471–533 (2001)
7. Vanitha, V., Joe, S.B., Krishnan, R., Fletcher, A.S.A., Anju, M., Akila, V.: Cognitive Threats Detection Model using Nature Inspired Chimpanzee Optimization for IoT Networks (CCM-COM). In: *Atlantis highlights in engineering/Atlantis Highlights in Engineering*. pp. 629–637 (2025). [https://doi.org/10.2991/978-94-6463-754-0\\_55](https://doi.org/10.2991/978-94-6463-754-0_55).

8. Piech, C., Bassen, J., Huang, J., Ganguli, S., Sahami, M., Guibas, L.J., Sohl-Dickstein, J.: Deep Knowledge Tracing. In: *Advances in Neural Information Processing Systems*, vol. 28 (2015)
9. Wu, C., Shi, S., Hu, J., Huang, H.: Knowledge-Enriched Joint-Learning Model for Implicit Emotion Cause Extraction. In: *CAAI Transactions on Intelligence Technology*, vol. 8, no. 1, pp. 118–128 (2023)
10. Baker, R.S., Corbett, A.T., Koedinger, K.R.: Detecting Student Misuse of Intelligent Tutoring Systems. In: *International Conference on Intelligent Tutoring Systems*, pp. 531–540 (2004)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

