
DK-PRACTICE: An Intelligent Platform for Knowledge Tracing and Educational Content Recommendation. A Case Study in Higher Education

[Marina Delianidi](#) , [Konstantinos Diamantaras](#) , [Georgios Kokkonis](#) , [Antonis Sidiropoulos](#) , [Georgios Evangelidis](#) , Dimitrios Karapiperis *

Posted Date: 30 December 2025

doi: 10.20944/preprints202512.2641.v1

Keywords: knowledge tracing; educational recommendations; machine learning; intelligent educational platform; education



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

DK-PRACTICE: An Intelligent Platform for Knowledge Tracing and Educational Content Recommendation. A Case Study in Higher Education

Marina Delianidi ¹, Konstantinos Diamantaras ¹, Georgios Kokkonis ¹,
Antonis Sidiropoulos ¹, Georgios Evangelidis ² and Dimitrios Karapiperis ^{1,*}

¹ International Hellenic University

² University of Macedonia

* Correspondence: karapiperd@ihu.edu.gr

Abstract

This paper introduces DK-PRACTICE, an intelligent educational platform that combines Knowledge Tracing (KT) and recommendation systems to support personalized learning in higher education. The platform utilizes a novel Paired-Bipolar Bag-of-Words (PB-BoW) model to assess students' knowledge states, forecast performance, and offer targeted recommendations. To test its effectiveness in real-world settings, DK-PRACTICE was implemented in the "Computer Organization and Architecture" undergraduate course, involving 138 students in Pre-Test and 106 in Post-Test. Empirical analysis of benchmark datasets and a newly created course dataset showed that the PB-BoW model outperformed an RNN-based KT model in predictive accuracy. Student surveys indicated high levels of satisfaction with usability, relevance of recommendations, and overall learning support, with most participants expressing willingness to reuse the platform in other courses. These results demonstrate the potential of DK-PRACTICE as a scalable and adaptable tool for improving personalized learning and bridging the gap between AI-driven KT research and classroom implementation.

Keywords: knowledge tracing; educational recommendations; machine learning; intelligent educational platform; education

1. Introduction

Addressing the educational challenges, such as accurately assessing student knowledge, predicting performance, and providing tailored educational content, is essential for effective personalized learning. This study introduces DK-PRACTICE, a Knowledge Tracing (KT)-based platform developed by utilizing machine learning (ML) methods to recommend the appropriate educational content to cover students' knowledge gaps.

Machine Learning, as a cornerstone for Artificial Intelligence (AI) is rapidly reshaping educational practices by enabling personalized, adaptive, and interactive learning experiences. In online educational systems, tools based on machine learning methods and, by extension, on artificial intelligence have shown particular progress. As noted in [1], the integration of tools such as PhET [2], Labster [3], ChatGPT [4], Squirrel AI [5], and IBM Watson [6] in the educational process led to a 28% improvement in student performance in many institutions. The study also found increased student motivation, conceptual clarity, and satisfaction, particularly when using simulation-based environments. These results demonstrate that AI can enhance student performance. Learning effectiveness is enhanced, especially when implemented through tools designed to personalize teaching and respond to students' individual needs.

Knowledge Tracing (KT), a modeling technique that assesses students' knowledge evaluation in educational contexts, like skills and concepts, is positioned at the core of many educational intelligent systems. KT allows systems to monitor students' progress and suggest tailored content based on

predicted knowledge gaps. Early KT models, such as Bayesian Knowledge Tracing (BKT) [7] and Item Response Theory (IRT) [8], laid the groundwork for individualized teaching with a possible assessment of learning skills performance. More recent developments, including Deep Knowledge Tracing (DKT) [9] and Graph-based Knowledge Tracing (GIKT) [10], use repetitive neural networks and neural graph networks to better model learning time patterns and relational dependencies between skills and questions.

Despite these technical advances, many AI-powered KT systems remain in a laboratory testing stage and have not been tested to work in real educational environments. Although platforms such as Knewton Alta [11], Carnegie Learning's MATHia [12], and Squirrel AI [5] have demonstrated efficacy in controlled or proprietary environments, there is a lack of research examining how KT systems function when deployed in real-world classrooms with diverse student populations. In addition, advanced models such as AdvKT [13], TrueLearn [14], TutorLLM [15], and ACKT [16] have focused on simulation fidelity or probabilistic feedback, but most lack classroom validation or accessible application frameworks for educators.

To address this gap, our study implements a Case Study using the DK-PRACTICE platform and evaluates the results of applying this knowledge assessment platform in a real educational environment. DK-PRACTICE monitors student knowledge levels and, taking into account student knowledge representation and a recommendation algorithm, recommends personalized educational content. In our case study, DK-PRACTICE was tested with 98 2nd-year university students in the "Computer Organization and Architecture" course.

The present study investigates the following research questions:

RQ1: How can a fast and efficient Knowledge Tracing model, based on BoW-inspired representation of student knowledge, be developed to effectively identify learning gaps through **pre-** and **Post-** test assessments?

RQ2: What is the impact of personalized content recommendations generated by the DK-PRACTICE platform, and how do they improve learning outcomes?

RQ3: How could the DK-PRACTICE platform enhance the educational process, and how could its use be extended to other educational topics?

Key contributions of this work include: (1) the introduction of a new lightweight ML method for Knowledge Tracing based on a Paired-Bipolar Bag-of-Words representation of student's knowledge (2) the design and actual implementation of the DK-PRACTICE platform as a case study in a real-world learning environment, (3) an empirical assessment of the impact of the test operation, and (4) a comprehensive comparison of DK-PRACTICE with existing AI-based KT systems, demonstrating its potential as a scalable and effective tool for adaptive learning.

Section 2 presents the current scientific picture of the evolution of Knowledge Tracing and educational recommendation models. The following Section 3 describes the technologies used for the development of the DK-PRACTICE educational platform and the realization of this study. In Section 4, the platform and its functions are presented. Section 5 describes the case study, with the analysis of the results commented on in Section 6. Section 7 concludes the study and presents the future extensions and improvements of the DK-PRACTICE platform.

2. Related Works

Artificial intelligence (AI) is emerging as a pillar of transformation in education, offering new ways to personalize learning and improve learning outcomes. Examples include AI-powered platforms such as PhET [2] and Labster [3], providing learning experiences based on simulating lessons such as physics or mathematics, while interlocutors such as the ChatGPT [4] support on-demand clarification and problem solving. The integration of AI tools into structured 60-minute virtual physics lessons led to a 28% increase in student performance, along with improvements in engagement, conceptual clarity, and motivation [1], validating the potential of AI-enhanced systems to address pedagogical challenges, including limited interactivity in the classroom.

The common benchmark in many of these AI systems is Knowledge Tracing (KT) – the task of assessing a student’s evolving knowledge of skills or concepts over time based on the results of interaction with educational platforms and observed behavior [7]. Early KT models, such as Bayesian Knowledge Tracing (BKT) [17] and The Item Response Theory (IRT) [8], applied probabilistic frameworks to infer students’ knowledge states. The evolution in KT’s field of research was brought about by the introduction of deep learning, specifically the Deep Knowledge Tracing (DKT) model by [9], which utilized recurrent neural networks (RNNs) to capture time dependencies and non-linear learning patterns. Subsequent improvements were presented in the Dynamic Key-Value Memory Networks (DKVMN) [18], which separated the representation of concepts from the mastery level to enhance interpretability, and Self-Attentive Knowledge Tracing (SAKT) [19], which utilizes transformer-based attention mechanisms to prioritize critical learning events. These innovations allowed for more accurate prediction and knowledge gaps identification.

Further extensions to KT have improved both accuracy and personalization. Exercise-Enhanced Sequential Modeling (EKT) introduced exercise-level features such as difficulty and subject to improve the accuracy of recommendations [20], while the Bidirectional GRU Knowledge Tracing (Bi-GRU) model [21] introduced forward and regressive diagnosis for both predicting performance and estimating students’ knowledge states from previous interactions, combined with a feedback network to predict the answer correctness in upcoming questions.

Online educational systems, along with practical applications on platforms such as ALEKS [22], Knewton [11], and Duolingo [23], highlight the role of KT in providing personalized learning experiences in real-time.

In recent research, [24] introduced two advanced KT models – Knowledge Structure-aware GraphAttention Network (KSGAN) and Knowledge Concept-based Memory Network (KCMN) – that model the complex relationships between exercises and knowledge concepts. By combining graph attention networks and memory mechanisms, their approach enhances interpretability and detail in predicting student performance. In simulation research, the authors of [25] addressed the problem of bias accumulation in KT-based simulators by proposing DSIM. This conditional diffusion model best approximates the actual behavior of the students. This innovation improves the quality of synthetic data used to train recommendation engines, overcoming the limitations of traditional KT models in production tasks. On the student engagement side, [26] conducted a large-scale randomized trial evaluating Squirrel AI, showing significant improvements in student motivation and engagement. The study highlighted that AI-based adaptive systems not only enhance cognitive performance but also support non-cognitive factors vital for learning success.

Within this landscape of increasingly sophisticated models, most remain either research prototypes or validated in synthetic or comparative datasets. This study introduces DK-PRACTICE, a KT-based platform specifically designed for classroom deployment in the real world. DK-PRACTICE assesses students’ knowledge levels and recommends personalized pre-test diagnostic-based content. DK-PRACTICE offers the ability to monitor changes in students’ knowledge state across learning objects or concepts, individually or in groups. Although many advanced models such as AdvKT [13], GIKT [10] or TutorLLM [15], contribute valuable architectural innovations such as multi-step simulation, graph-based integrations, and language model integration, they often lack deployment scenarios that evaluate educational impact, student feedback, and scalability under real-world constraints and have not yet been tested in classroom environments. DK-PRACTICE addresses these gaps by prioritizing functionality aligned with classroom practices, including diagnostic assessments, actionable feedback, and content alignment.

In summary, the literature illustrates a rich and rapidly evolving area of research in knowledge discovery and AI-assisted individualized training with models that have laid the technical foundations, such as DKT, DKVMN, SAKT, etc. However, few systems effectively bridge the gap between research and practical application. DK-PRACTICE contributes to this space by combining a robust machine

learning model with real educational development, proving that adaptive KT platforms can be not only technically innovative but also pedagogically effective.

3. Background Technologies

The DK-PRACTICE online education platform was built using machine learning methods. Machine Learning (ML), a prominent subfield of Artificial Intelligence (AI), has been extensively applied in educational contexts. As previously discussed, ML algorithms have played a substantial role in the analysis of student behavior [27], the prediction of academic performance [28], and the provision of personalized learning recommendations [29].

From a theoretical point of view, neural networks constitute a core framework for modeling complex nonlinear mappings in high-dimensional spaces [30,31]. Neural networks have emerged as a foundational paradigm in the advancement of machine learning techniques [32]. Empirical evidence demonstrates that neural networks, in particular, multilayer feedforward architectures, consistently improve performance in domains requiring the extraction of complex feature hierarchies from raw, unstructured data [33]. Their capacity for end-to-end learning demonstrate the ability to autonomously identify and encode hierarchical feature representations from unstructured input data, thereby enabling improved performance across a wide range of complex pattern recognition tasks [34]. These technological advances form the foundation for modern adaptive learning platforms such as DK-PRACTICE.

Educational recommendation systems are smart technologies designed to personalize learning by suggesting relevant content, activities, or learning paths based on a student's level of knowledge and performance history. Their primary role is to support adaptive learning by filling knowledge gaps with educational proposals tailored to individual needs [35]. Unlike traditional static curricula, these systems dynamically adapt in real time, helping students engage with the right content at the right time to maximize efficiency and learning motivation. The advantages of education systems include their contribution to stronger academic outcomes [36]. Leveraging data-driven models, often powered by machine learning or knowledge tracing (KT) techniques. In classroom settings, recommendation systems also help teachers by automating feedback, identifying student weaknesses early, and facilitating data-driven teaching.

As education becomes increasingly digitalized, these systems play a crucial role in scaling up personalized learning experiences. Educational recommendation systems can be broadly categorized into various types based on the algorithms and data sources they use. The most common approaches include content-based filtering, which suggests elements similar to what a student has previously interacted with, and collaborative filtering, which leverages patterns of behavior among like-minded users [37,38]. In the field of education, more advanced categories have emerged. Knowledge-based systems use predefined concept hierarchies or knowledge graphs to align learning content with curriculum structures [39]. In contrast, context-aware systems adjust recommendations based on external variables such as time, location, student goals, or engagement history [40,41] or hybrid systems, which combine multiple techniques to improve accuracy and personalization [42]. These advanced systems are often enhanced with machine learning models and KT algorithms to monitor learning progress and ensure that recommendations evolve along with student development.

4. The DK-PRACTICE Platform

In this study, we introduce DK-PRACTICE (Figure 1), an intelligent online platform designed to conduct adaptive knowledge assessments through a question-answer method, wherein each question is dynamically tailored based on students' prior responses. The intelligent DK-PRACTICE functions as an integrated knowledge tracing and recommendation system and aims to:

- Accurately estimate students' current knowledge states,
- Predict their future performance on subsequent personalized questions,

- Recommend targeted instructional content to address specific learning objectives, concepts, and skills where deficiencies have been identified,
- Support users' roles.

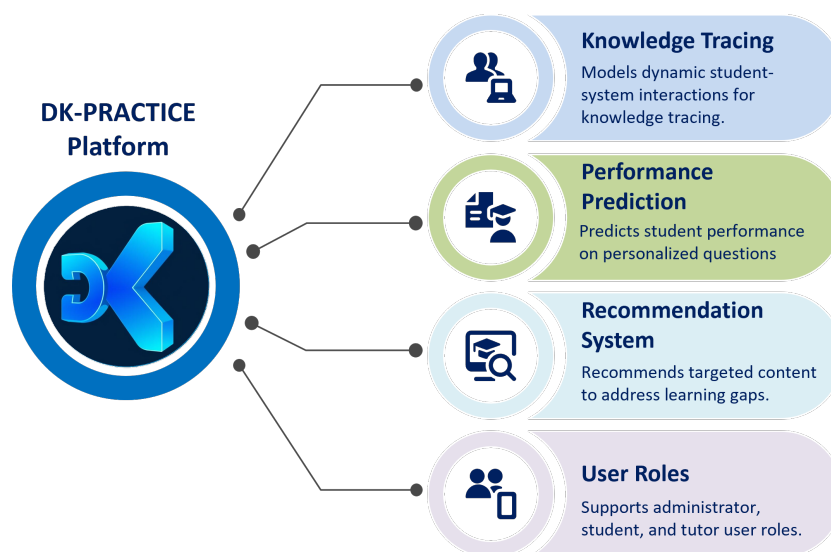


Figure 1. The DK-PRACTICE platform functionalities.

The knowledge tracing component, called PB-BoW (Paired-Bipolare Bag of Words), is inspired by the Bag of Words approach [43] and is implemented using a feedforward neural network, which allows modeling of the dynamic and sequential nature of interactions between students and the platform. Our approach facilitates robust and accurate predictions of students' evolving knowledge states. Educational recommendations are generated based on the predicted probability of correctness for the next unattempted question, thereby guiding the selection of instructional resources aimed at remediating identified knowledge gaps.

DK-PRACTICE supports three distinct user roles: "administrator", "student", and "tutor". Students engage with adaptive assessments, tutors construct and manage test content, and administrators oversee system operations. A comprehensive description of the platform's architecture, functionalities, and implementation details is provided in the following sections.

4.1. The DK-PRACTICE Platform Architecture and Functionalities

The architecture of the DK-PRACTICE platform (Figure 2) includes two main components based on machine learning methods :

- The Knowledge Tracing component to predict student's performance and to estimate the student's knowledge state.
- The Recommendation component to produce educational content recommendations.

The architectural structure of the DK-PRACTICE Platform not only enhances the fidelity of the knowledge state assessment but also makes the model suitable for auxiliary tasks, such as personalized recommendations. Both Knowledge Tracing and Recommendation components constitute core parts of the DK-PRACTICE platform, and their implementations reside within the back-end architecture.

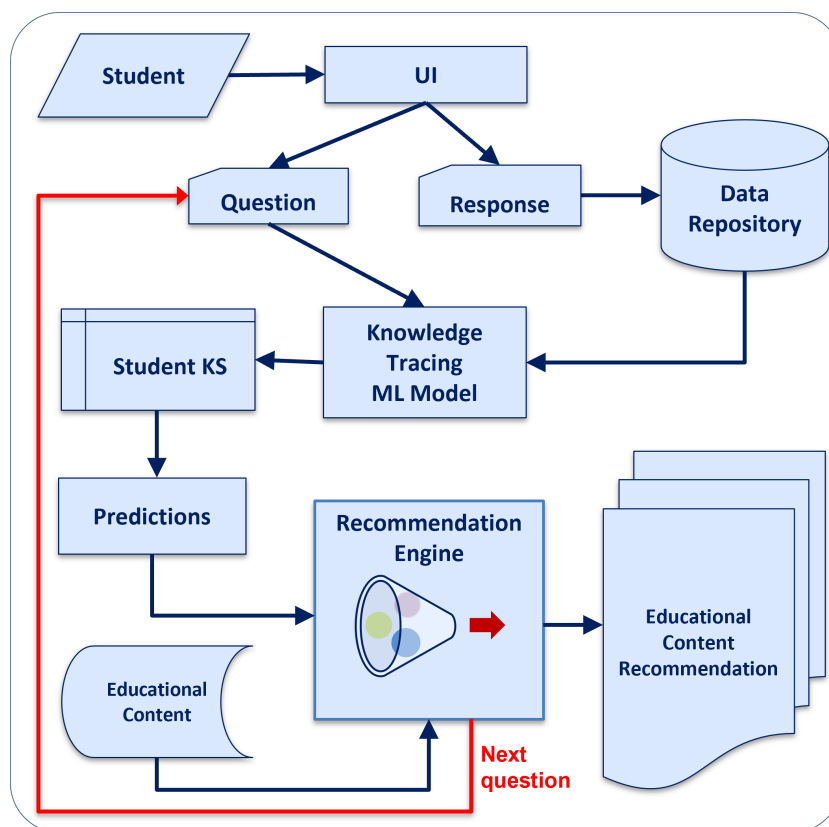


Figure 2. The DK-PRACTICE platform architecture.

4.2. Data Model

Data is one of the most fundamental components of the DK-PRACTICE platform. It is essential for the operation of any application that uses machine learning.

The data model of the platform is shown in Figure 3. The database schema is organized around three primary entities: “Concepts”, “Questions”, and “Exams”. Within the “Questions” entity, the “Chapter” attribute denotes the specific segment of the educational material associated with the question.

The compiled database consists of:

- Examination questions,
- The corresponding underlying concepts,
- The chapters containing the educational content of the course,
- Students’ submitted answers across different examination periods.

The concepts are systematically mapped to both the exam questions and the educational content chapters, thereby establishing a structured linkage between learning materials and the questions. This attribute serves as a key parameter in the recommendation algorithm, enabling the identification and prioritization of relevant educational content when generating personalized recommendation lists.

The dataset derives from the database and consists of interaction records, specifically question–answer pairs (q, r) , recorded within the “Exams” entity. The responses are binary: $r = 0$ for an incorrect answer and $r = 1$ for a correct answer. All personally identifiable information was removed through an anonymization process to ensure compliance with ethical and privacy standards. The constructed dataset was subsequently utilized to train the machine learning model designed to infer the Knowledge Tracing component, described below, was conducted by leveraging students’ historical interaction data.

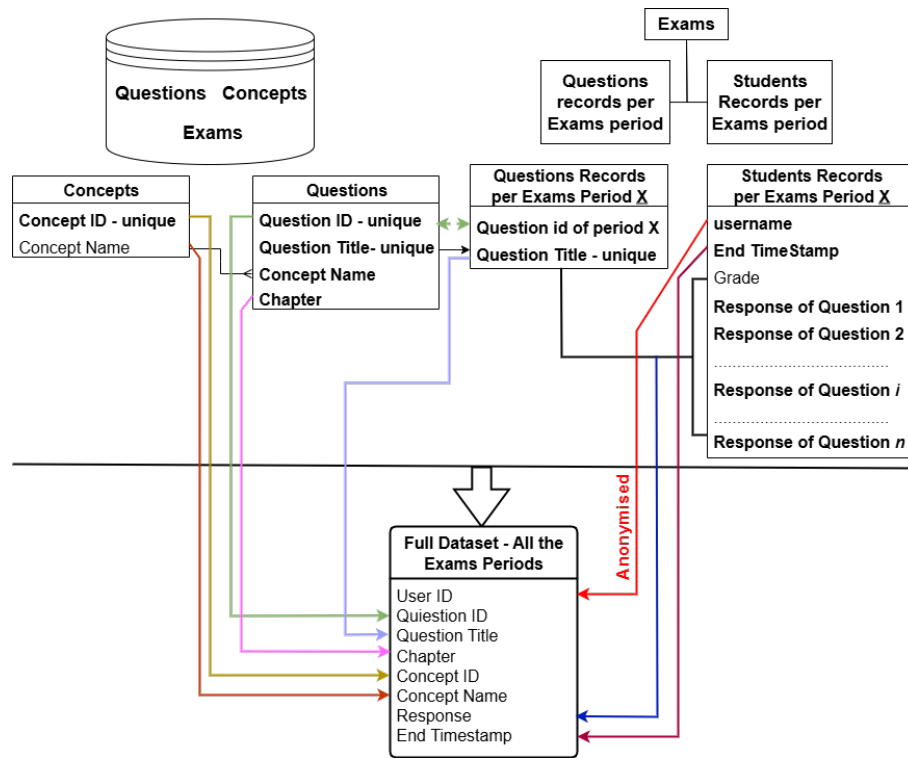


Figure 3. The Data model.

4.3. Knowledge Tracing Component

The idea of [21], in terms of handling question-response interactions, inspires the knowledge tracing model that we introduce and utilize in the DK-PRACTICE platform. Unlike previous deep learning knowledge tracing models [1,7,9,18,19], our model explicitly disconnects the current question from the current representation of knowledge. This leads to the production of the knowledge state representation vector in which the student's learning is incorporated up but not including the current time step τ , based exclusively on the sequence of previous question-answer pairs (q_k, r_k) , $k = 1, \dots, \tau - 1$.

For our Knowledge Tracing model we propose an extension of BoW using paired-bipolar representation of answered questions (Figure 4). The representation vector \mathbf{x} has dimension $2 \cdot N_q$, where N_q is the total number of questions. All elements of \mathbf{x} are zero except that for each question i with an existing response r , the following values are assigned:

- If $r = 0$ (incorrect answer), then $x_{2i-1} = 1$ and $x_{2i} = -1$.
- If $r = 1$, (correct answer) then $x_{2i-1} = -1$ and $x_{2i} = 1$.

The corresponding representation scheme will be called Paired-Bipolar Bag-of-Words (PB-BoW).

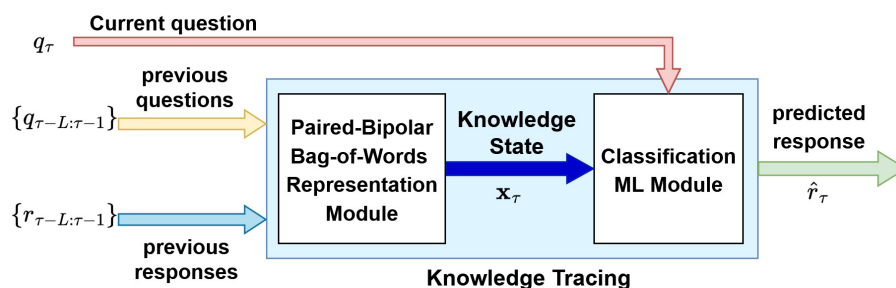


Figure 4. The Paired-Bipolar Bag-of-Words (PB-BoW) Knowledge Tracing model.

A total of $M = 200$ samples are generated for each student through the following procedure:

- A subset of 10 questions is randomly selected to form the **query set** for training. The remaining questions are designated as the **evaluation set**.
- The PB-BoW vector representation, \mathbf{x} , is generated for the query set. This vector has a length of $2 \times N_q$ and serves as the input to the classification module.
- The target vector, \mathbf{t} , for the classification module is constructed using the responses from the evaluation set. All elements of \mathbf{t} are equal to -1 except for each evaluation question q_j with response r_j , where we set $t_j = r_j$. Therefore, unanswered questions have target value $t = -1$, allowing us to exclude them when calculating the loss of the classifier.

Algorithm 1 Generate Training Samples for PB-BoW KT

```

1: procedure GENERATESAMPLES( $X$ )
2:   Input: Questions and responses data  $X = \{(q_1, r_1), (q_2, r_2), \dots, (q_{N_{qs}}, r_{N_{qs}})\}$  for student  $s$ . The
   values  $q_i$  are the indices of the questions answered by the student and  $r_i$  are the corresponding
   responses (0 or 1);  $N_{qs}$  is the number of questions answered by this student.
3:   Output: Dataset  $D_s = \{(\mathbf{x}_i, \mathbf{t}_i)\}_{i=1}^{200}$ 
4:   Let  $N_q$  be the total number of unique questions in the question bank.
5:   Set  $D_s \leftarrow \emptyset$ 
6:   for  $i = 1$  to 200 do
7:     Randomly select a subset of 10 questions to form the query set  $Q \subset X$ .
8:     The remaining questions form the evaluation set  $E = X \setminus Q$ .
9:     Initialize input vector  $\mathbf{x}$  of size  $2 \times N_q$  with zeros.
10:    for each question-response pair  $(q_j, r_j) \in Q$  do
11:      Generate Paired-Bipolar representation:
12:      if  $r_j = 0$  then ▷ Incorrect response
13:         $x[2q_j] \leftarrow 1$ 
14:         $x[2q_j + 1] \leftarrow -1$ 
15:      else ▷ Correct response
16:         $x[2q_j] \leftarrow -1$ 
17:         $x[2q_j + 1] \leftarrow 1$ 
18:      end if
19:    end for
20:    Initialize target vector  $\mathbf{t}$  of size  $N_q$  with  $-1$ .
21:    for each question-response pair  $(q_k, r_k) \in E$  do
22:       $t[q_k] \leftarrow r_k$ 
23:    end for
24:    Add the pair  $(\mathbf{x}, \mathbf{t})$  to  $D_s$ .
25:  end for
26:  return  $D$ 
27: end procedure

```

4.3.1. The Classification Module

The task of the classification module is for every student to predict the responses to all questions, given a history of 10 prior interactions (question/response pairs). The classifier is a feedforward neural network with the following layer specifications:

- **Input Layer:** $2 \times N_q$ neurons
- **Hidden Layers:** Three fully-connected hidden layers.
 - Layer 1: 15 neurons.
 - Layer 2: 10 neurons.
 - Layer 3: 5 neurons.
- **Output Layer:** N_q neurons. Each neuron corresponds to a question, and its value between 0 and 1, predicts the probability of a correct response.

The input to the classifier is the PB-BoW vector representation \mathbf{x} of the student's interactions in the query questions. The target of the classifier is the vector \mathbf{t} corresponding to the evaluation questions.

The set of students is split into training and test subsets. For each training student s , we generate the corresponding dataset D_s^{train} according to Algorithm 1. The overall training set results from the concatenation of all D_s^{train} datasets. In a similar fashion, we generate the test set from the datasets D_s^{test} of the test students.

The classifier loss L , is the mean squared error computed only on the questions with a valid target value, i.e., with $t_i \neq -1$:

$$L = \frac{1}{N} \sum_{i=1}^{N_q} (t_i - \hat{r}_i)^2 I(t_i \neq -1)$$

where \hat{r} is the classifier output, $I(\cdot)$ is the indicator function¹ and $N = \sum_i I(t_i \neq -1)$.

4.4. Recommendation Component

The platform implements two types of recommendation mechanisms: (a) *next-question recommendation* during a knowledge assessment and (b) *post-assessment recommendation* of educational materials.

4.4.1. Next Question Recommendation

Let the query question set for a student be denoted by

$$Q = \{q_1, q_2, \dots, q_L\}.$$

where $L = 10$. Questions are asked sequentially in an order that depends on the students' previous answers. Therefore, questions are not asked in the same order for every student. The function $m(\tau)$ returns the index of the question to be asked at time step τ . Due to lack of prior information, in the beginning of the sequence, all participants start with the same initial question $q_{m(1)}$, referred to as the *cold-start question*.

Subsequent questions are personalized for each student and adapted according to their prior responses. The recommendation model utilizes the knowledge state estimated vector of the student to compute the probability of a correct response for each remaining question in the question bank. At time τ , the question set is split into two subsets:

1. the set of *Potential questions*, Q_{pot}^τ , which remain eligible to be asked as the next question;
2. the set of *Answered questions* Q_{ans}^τ that are ineligible to be asked again.

Initially, all questions, except for the *cold-start question*, are classified as *Potential* while the cold start question is classified as *Answered*.

At each time step $\tau \geq 2$:

1. The asked question q_τ is moved from $Q_{pot}^{\tau-1}$ to Q_{ans}^τ :

$$Q_{ans}^\tau = Q_{ans}^{\tau-1} \cup \{q_\tau\}, \quad Q_{pot}^\tau = Q_{pot}^{\tau-1} \setminus \{q_\tau\}.$$

2. For each $q \in Q_{pot}^\tau$, the recommendation model computes the predicted probability of a correct response based on the student's estimated knowledge state vector \mathbf{x}_τ :

$$p_\tau(q) = \Pr(\hat{r}_\tau(q) = 1 \mid \mathbf{x}_\tau).$$

3. The next question is selected as the one whose probability of correct response is closest to 0.5:

$$q_{\tau+1} = \arg \min_{q \in Q_{pot}^\tau} |p_\tau(q) - 0.5|,$$

thereby targeting questions for which the model is maximally uncertain.

¹ $I(cond) = 1$ if $cond = \text{true}$, and 0 otherwise.

Within the framework of the DK-PRACTICE platform, we have established the personalized question quota at $L = 10$. This means that we can only use 10 questions and responses to determine the knowledge state of the student. The remaining questions are used for evaluating the performance of our prediction model. This quota constitutes a hyperparameter and is subject to occasional adjustments following preferences, experimental findings, course requirements, etc.

4.4.2. Educational Content Recommendation

Following the completion of the test, and based on the predicted knowledge state of the student derived from the test results, the probability of correctly answering each question in the question bank is recalculated. Each question q is associated with a learning object or concept $c(q)$.

The knowledge level for a given learning concept c is computed as:

$$K_c = \frac{1}{|Q_c|} \sum_{q \in Q_c} p_q$$

where Q_c is the set of questions linked to c and $|Q_c|$ is its cardinality. The value K_c represents the average estimated probability of correct responses for questions concerning learning object c , and is interpreted as the student's mastery level for that object.

The recommendation engine selects educational materials corresponding to learning concepts with the lowest K_c values, thereby prioritizing areas in which the student demonstrates the weakest mastery.

All this process is supported by a data model which is organized into a database, a training dataset, and a database of educational content.

4.5. Knowledge Tracing Experiments and Results

In this work, we introduce a new real-world dataset named "COaA course", which was collected from nine examination periods in the undergraduate Computer Organization and Architecture course at the [Department] of [University] via the Moodle electronic learning management system. These data encompassed both the period of the COVID-19 pandemic, during which courses and assessments were conducted under quarantine restrictions and the subsequent post-pandemic period. The database constructed for the case study adheres to the platform's underlying data model.

The "COaA course" dataset comprises a total of 19,339 interactions, encompassing question-answer pairs generated by 1,115 students on 120 multiple-choice questions, 31 concepts, and 20 chapters of educational content. For experimental purposes, the data set was split into training and testing subsets. Consequently, the training set contains 15,555 interactions from 891 unique students, while the testing set consists of 3,784 interactions from 224 unique students. Each student contributes a minimum of 11 interactions (i.e., question-answer pairs), thereby ensuring adequate data density for reliable model evaluation.

In addition to the newly constructed "COaA Course" dataset, developed for evaluating the PB-BoW KT model, we employed two widely used ASSISTments benchmark datasets²: "ASSISTment 2009" and "ASSISTment 2017". The "ASSISTment 2009" dataset comprises 101 concepts, with model inputs represented as corresponding questions. It contains data from 4,151 students, including 1,196 in the training set and 511 in the test set, yielding a total of 274,590 student-question interactions. Similarly, the "ASSISTment 2017" dataset includes 101 concepts with model inputs likewise represented by questions. This dataset encompasses 1,707 students (1,196 in the training set and 511 in the test set), resulting in 864,703 recorded interactions. Both datasets include multiple student responses to identical questions, as well as records with fewer than 11 question-answer interactions per student.

For the implementation of the PB-BoW KT model, during the data preprocessing procedure, we excluded records containing fewer than 11 interactions and retained only the first response to each

² <https://sites.google.com/site/assistmentsdata/>

repeated question per student. The resulting dataset statistics, including the number of students and interactions in both the training and test splits, are summarized in Table 1.

The experiments were conducted on two knowledge tracing models: the proposed PB-BoW KT model and the KT-BiGRU [21] model.

Table 1. Dataset statistics after preprocessing. For each dataset, the number of students and total interactions are reported, along with their distribution across training and test sets.

Dataset	Questions	Students			Interactions		
		Total	Train	Test	Total	Train	Test
ASSISTment 2009	101	818	556	262	23,722	16,511	7,211
ASSISTment 2017	101	1,641	1,144	497	69,712	47,856	21,856
COaA Course	120	1,115	891	224	19,339	15,555	3,784

The results reported in Table 2 demonstrate that the PB-BoW KT model consistently outperforms the KT-BiGRU model across all three datasets in terms of both AUC and Accuracy. On the “ASSISTment 2009” dataset, the PB-BoW KT model achieves an AUC of 0.765 and an Accuracy of 0.760, exceeding the KT-BiGRU scores of 0.753 and 0.713, respectively. A similar trend is observed on the “ASSISTment 2017” dataset, where the PB-BoW KT model attains an AUC of 0.733 and an Accuracy of 0.700, slightly higher than the corresponding values of 0.716 and 0.686 obtained by KT-BiGRU. The most pronounced performance gap appears in the “COaA Course” dataset, where PB-BoW KT achieves the highest results overall with an AUC of 0.805 and an Accuracy of 0.787, compared to 0.736 and 0.706 for KT-BiGRU.

Table 2. The best results of AUC and Accuracy metrics per dataset and per Knowledge Tracing model.

Dataset	PB-BoW KT model		KT-BiGRU model	
	AUC	Accuracy	AUC	Accuracy
ASSISTment 2009	0.765	0.760	0.753	0.713
ASSISTment 2017	0.733	0.700	0.716	0.686
COaA Course	0.805	0.787	0.736	0.706

These findings indicate that the PB-BoW KT model provides more robust predictive performance and generalizes more effectively across different datasets, suggesting its suitability for knowledge tracing tasks. In certain cases, relatively simple machine learning approaches can outperform deep learning methods, while also providing the advantage of substantially reduced training time and eliminating the need for GPU resources.

The PB-BoW KT model was trained on a system equipped with an 11th Generation CPU operating at 2.40 GHz and 16 GB of RAM. For all the datasets, training was conducted for 20 epochs, with an average runtime of approximately 5 seconds per epoch.

4.6. The DK-PRACTICE Platform Implementation

The DK-PRACTICE platform employs a client-server architecture, with the front-end developed in PHP and JavaScript. Data is stored and managed through a MySQL database schema. In the back-end of the platform, a security-enabled Flask^{3,4}-based API (Application Programming Interface) supports bidirectional communication with the machine learning models, providing outputs such as “Next Question” recommendation and the final “Educational Recommendations”, which support the delivery of personalized educational content. In addition, to ensure full functionality, the platform provides user registration and password recovery services in compliance with security policies.

As mentioned above, the DK-PRACTICE platform supports three types of users: “administrator”, “tutor” and “student”. Thus users:

³ Welcome to Flask - Flask documentation, <https://flask.palletsprojects.com/en/3.0.x/>

⁴ Welcome to Flask - Security documentation, <https://flask-security-too.readthedocs.io/en/stable/>

- with the “*administrator*” role (Figure 5). Beyond the general oversight of the platform, the administrator is tasked with the: (a) establishment of courses, (b) the registration of tutors, and (c) the allocation of course management responsibilities to the corresponding tutors.

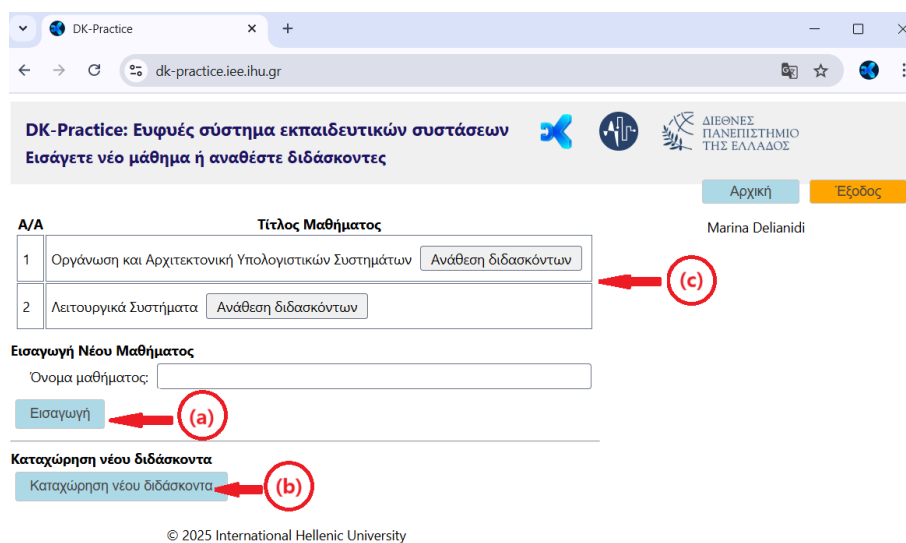


Figure 5. Administrator screen.

- with the “*tutor*” role have the permissions to create a test (Figure 6) by entering all the relevant information and to set parameters such as starting and ending date and time (Figure 7), and monitor the students’ performance individually and in groups per Test (Figure 8). The ability to monitor test results enables the teacher to have an overview of the class’s knowledge state regarding the specific course, and indeed for each educational object or concept.




Figure 6. Tutor screen: The Allocated Courses.

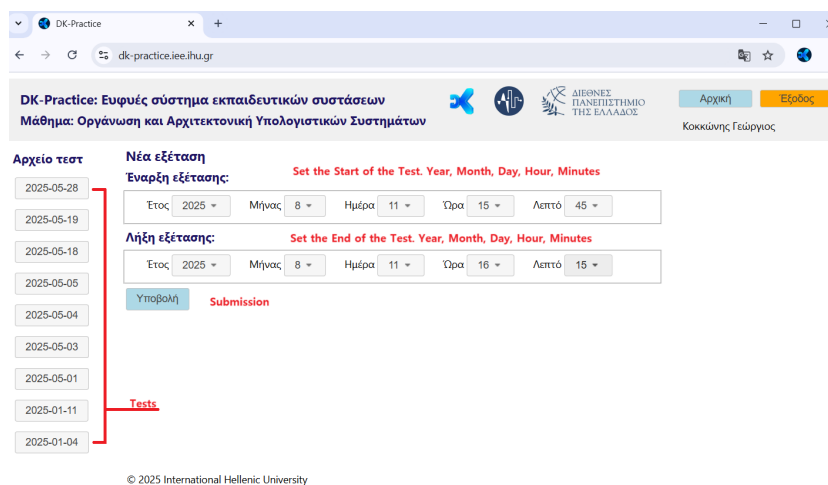


Figure 7. Tutor screen: The Test Settings.

segment of the bar chart visualizes the proportion of knowledge acquired in the concepts taught within the specific course.

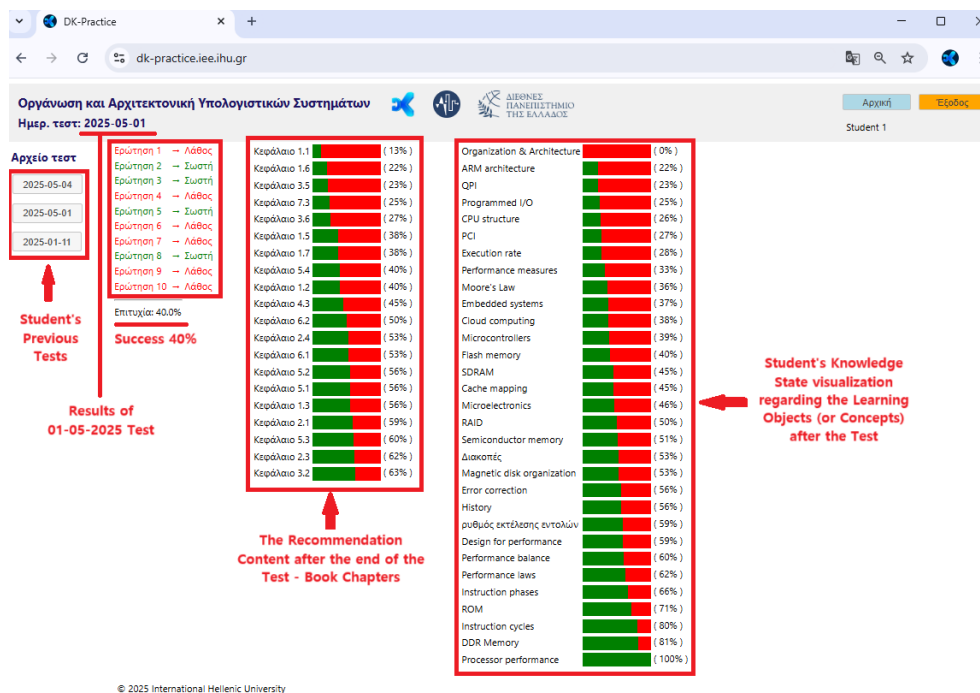


Figure 11. Student screen: Test Results, Personalized Educational Content Recommendations, Personalized Knowledge State According to the Learning Objects or Concepts of the Course.

5. The Case Study

The objective of this case study is to examine and assess the DK-PRACTICE platform and its functionalities within a real learning environment in higher education. To this end, the platform was implemented in the undergraduate course “*Computer Organization and Architecture*”, offered in the second year of study within the Department of Information and Electronic Engineering at the International Hellenic University in Greece.

As part of the case study design, two diagnostic assessments were implemented, targeting primarily second-year students enrolled in the course and scheduled to participate in the subsequent official examination period. The first assessment, hereafter referred to as the “**Pre-test**”, was administered approximately one month before the official examination. Its purpose was to evaluate students’ knowledge readiness and to generate personalized recommendations for targeted study based on identified areas of need. Following a two-week interval, a second assessment, hereafter referred to as the “**Post-test**”, was conducted. This assessment was based on the assumption that the students had engaged in the recommended learning resources and had studied the instructional material proposed through the DK-PRACTICE platform.

In both experimental tests, students responded to **10** questions among the set of 120 questions, the outcomes of which were used to assess their knowledge state concerning the 31 concepts covered in the 20 course chapters (Table 3).

Based on this assessment, personalized content recommendations were generated to facilitate further study and enhance knowledge acquisition. Apart from the initial question, each subsequent question asked to a student was adaptively determined by the history of previous question responses. This process was supported by the PB-BoW Knowledge Tracing model, which operates in the backend of the platform.

Student participation in both assessments was strictly voluntary and based on informed consent. Moreover, participation did not influence students’ grading outcomes, nor did it interfere with the formal assessment procedures of the official examination period.

Table 3. Concepts and the Corresponding Chapters.

Concept_name	Chapter	Concept_name	Chapter
CPU structure, Organization & Architecture.	1.1	Interrupts,	3.2
Microelectronics, Moore’s Law.	1.2	Instruction cycles, Instruction phases.	
History.	1.3	QPI.	3.5
Embedded systems, Microcontrollers.	1.5	PCI.	3.6
ARM architecture.	1.6	Cache mapping.	4.3
Cloud computing.	1.7	Semiconductor memory, ROM.	5.1
Design for performance, Performance balance.	2.1	Error correction.	5.2
Performance laws.	2.3	DDR Memory, SDRAM.	5.3
Execution rate,	2.4	Flash memory.	5.4
Processor performance,		Magnetic disk organization.	6.1
Performance measures,		RAID.	6.2
Order execution rate.		Programmed I/O.	7.3

5.1. Pre-Test Description

The first, “Pre-Test”, knowledge tracing test was to examine the knowledge state of the participating students. A total of 138 students participated, all of whom had previously registered and activated their accounts. The test was administered in groups, a procedure determined by the availability of laboratory facilities and computing resources. Two laboratories were used, which contained 25 and 27 internet-enabled workstations, respectively. Throughout testing sessions, the DK-PRACTICE platform demonstrated reliable performance, supporting up to 50 simultaneous users without malfunctions in its services. The tutor monitored the progress of the students in real time through the user interface of the “tutor” of the platform (Figure 12).

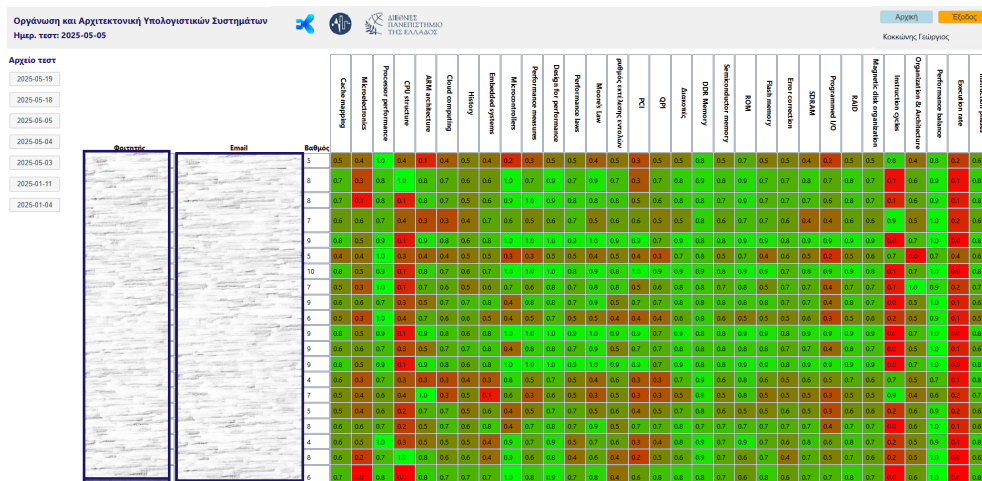


Figure 12. A part of the tutor’s user interface during the “Pre-Test” process on the DK-PRACTICE platform, showing student outcome monitoring.

In addition, the participants were asked to complete a questionnaire to record the user experience regarding the functionality, usability, and environment of the platform. 31 of the 138 participants completed and submitted the questionnaire shared online. Their responses provide valuable insight into the user experience and are further analyzed below in Chapter 5.3.

The questions of the 1st questionnaire, along with demographic data such as year of study and experience using e-learning platforms (where 65% of students responded positively), were divided into the following categories:

- (e) I felt that the platform understood my learning needs
 - (f) I feel that the platform has contributed to my progress
 - (g) I have a better picture of my weaknesses after this experience
 - (h) The personalization of the proposals was evident and useful
 - (i) I feel that I am better prepared for the exam
2. **Effectiveness of Recommendations**
- (a) The content was relevant to the points I was lacking
 - (b) The suggested material helped me to better understand the subject matter
 - (c) The test of the second experiment recorded my progress
 - (d) The proposals met my real needs
 - (e) My progress from the first to the second test was evident
 - (f) I would use the platform again to test my knowledge
3. **Willingness to use the platform by students**
- (a) I would continue to use the platform in other courses
 - (b) I would recommend the platform to other students

Both questionnaires were created to encompass a wide range of factors influencing student experience, system reliability, quality of personalized recommendations, and perceived learning progress. This multi-layered approach ensures valid and reliable evaluation outcomes, providing evidence-based insights for improving the platform and integrating it effectively into educational practices.

5.3. Evaluation

The evaluation of the platform is conducted according to two primary axes: User Experience and Platform Effectiveness.

1. **User Experience:** In terms of user experience, the evaluation was conducted in two phases. The first phase was implemented during the final stage of system development, during which the platform was subjected to testing to identify usability-related improvements. The evaluation highlighted issues concerning the visual presentation of questions, the display of results upon test completion, and the functionality of the password recovery mechanism. Following the incorporation of the necessary modifications, these issues were resolved, thereby ensuring the platform's readiness for deployment in the "Computer Organization and Architecture" course. The second phase of the usability evaluation was carried out by the students who engaged with the platform during the two testing sessions. In both evaluations, students reported a positive user experience with the DK-PRACTICE platform.
- The survey results of the "Pre-Test", presented in the chart (Figure 14) demonstrate an overall positive evaluation of the intelligent knowledge tracing platform by higher education students. The highest levels of agreement were recorded for the statement "Using the platform was easy and intuitive", with 48.39% agreeing and 38.71% strongly agreeing, indicating that the platform's usability is one of its strongest features. Similarly, the layout of the environment was well received, with 41.94% agreeing and 38.71% strongly agreeing that it was logical and clear.
- The design of the platform was also evaluated positively, with 29.03% agreement and 38.71% strong agreement, though this item received the highest proportion of neutral responses (29.03%). Regarding functionality, 35.48% of students agreed and another 35.48% strongly agreed that the platform's features met their expectations, although 9.68% disagreed, suggesting some room for improvement. Navigation was perceived as smooth and seamless by most respondents (32.26% agreed, 38.71% strongly agreed), although 25.81% remained neutral. In summary, the findings show that the DK-PRACTICE platform is largely intuitive, user-friendly, and effective, with only a small minority reporting dissatisfaction. These results highlight strengths in ease of use and

logical design, while pointing to opportunities for further refinement of features and overall design appeal.

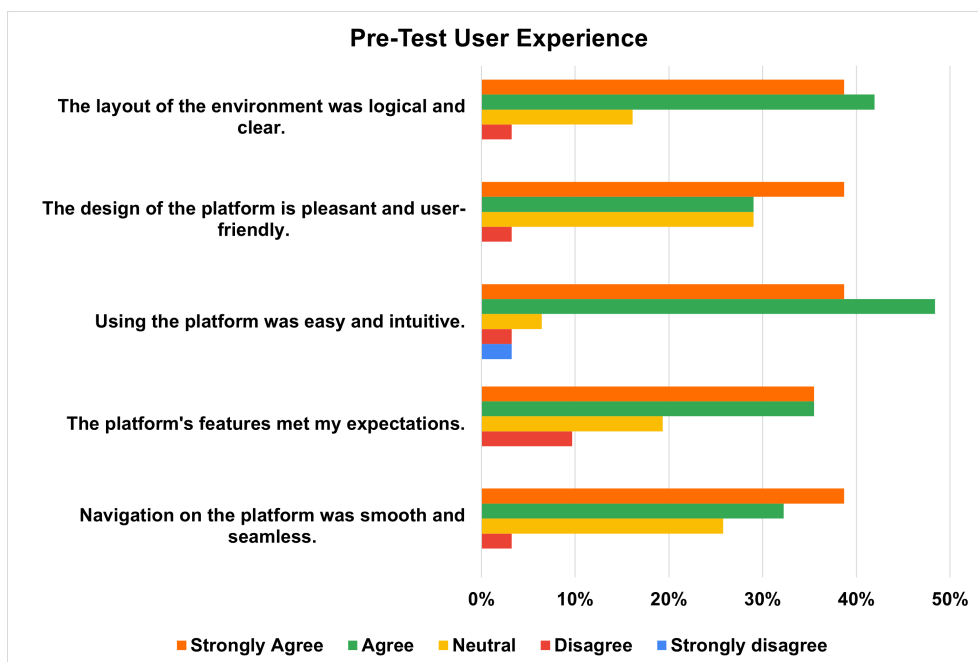


Figure 14. Pre-Test: User Experience Evaluation of using DK-PRACTICE platform by students.

Based on questionnaire answers after the “Post-Test” (Figure 15, most respondents expressed satisfaction, with notable percentages agreeing or strongly agreeing with key statements about platform usefulness. Specifically, 72.06% felt the material was well-organized, understandable, and enjoyable to use. Additionally, students overwhelmingly believed the platform was beneficial for their learning; a combined 73.53% of users agreed or strongly agreed that it contributed to their progress and helped them identify their weaknesses more effectively. The most positive feedback centered on the intention to reuse the platform, as 70.59% of respondents would use it again to test their knowledge, demonstrating strong support for its effectiveness as a learning tool.




Figure 15. Post-Test: User Experience Evaluation of using DK-PRACTICE platform by students.

2. **Platform Effectiveness:** This axis relates to the effectiveness of the services offered, specifically the assessment of knowledge state and the educational recommendations during Pre-/Post- tests performed by students.

The findings of the quality assurance questionnaire after “Pre-Test” indicate that students mostly viewed the intelligent knowledge tracing platform DK-PRACTICE primarily positively (Figure 16). Most respondents chose “Agree” on all five items, with especially high agreement for the statements “This experience helped me understand where I needed improvement” and “I trust the knowledge assessment process followed by the platform”, both surpassing 60%. A significant number of students also strongly agreed, especially regarding the platform’s ability to identify individual strengths and weaknesses and accurately assess their knowledge. Neutral responses were moderate, while reports of disagreement or strong disagreement were minimal. These results show that students generally see the platform as effective, reliable, and helpful for their learning, highlighting its potential value in higher education.

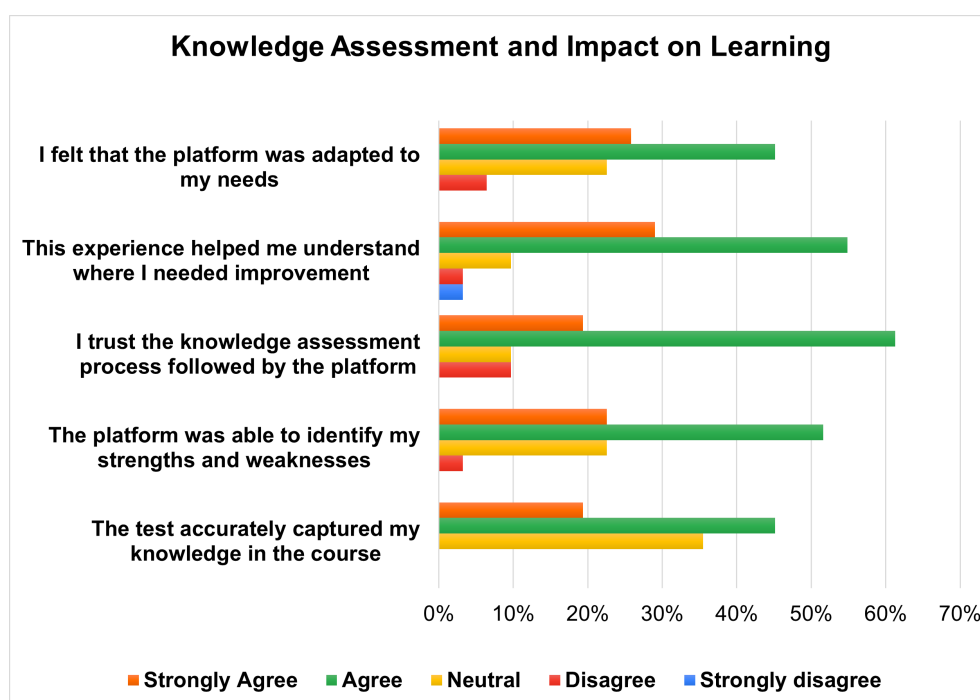


Figure 16. Pre-Test: Knowledge Assessment and Impact on Learning through the DK-PRACTICE platform.

The “Post-Test” quality assurance questionnaire, given to higher education students, highlights both the high use and perceived effectiveness of the DK-PRACTICE platform’s recommendation system. As shown in Figure 17, most respondents actively engaged with the recommendations. Specifically, 36.76% reported a moderate level of use (score of 3 on a 5-point scale), while a significant number of students indicated a high degree of use, with 30.88% selecting a score of 4 and 16.18% choosing the highest score of 5. These results suggest that the platform’s recommendations were considered valuable enough to encourage active engagement, supporting the platform’s role as a tool for personalized learning.

Further analysis of the survey data, as shown in Figure 18, confirms the high effectiveness of these recommendations from the students’ point of view. Most students agreed or strongly agreed with the statements, indicating a positive user experience. Notably, over 50% of respondents agreed that “The content was relevant to the points I was lacking” and that “The proposals met my real needs”. Additionally, a large percentage of students felt that the “suggested material helped me to understand the subject matter better” and that “The test of the second experiment recorded my progress”.

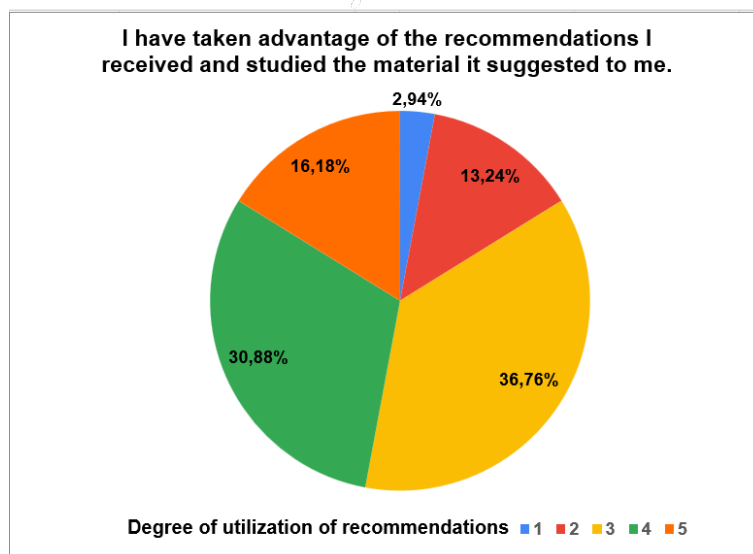


Figure 17. Post-Test: The Degree of utilization of recommendations.

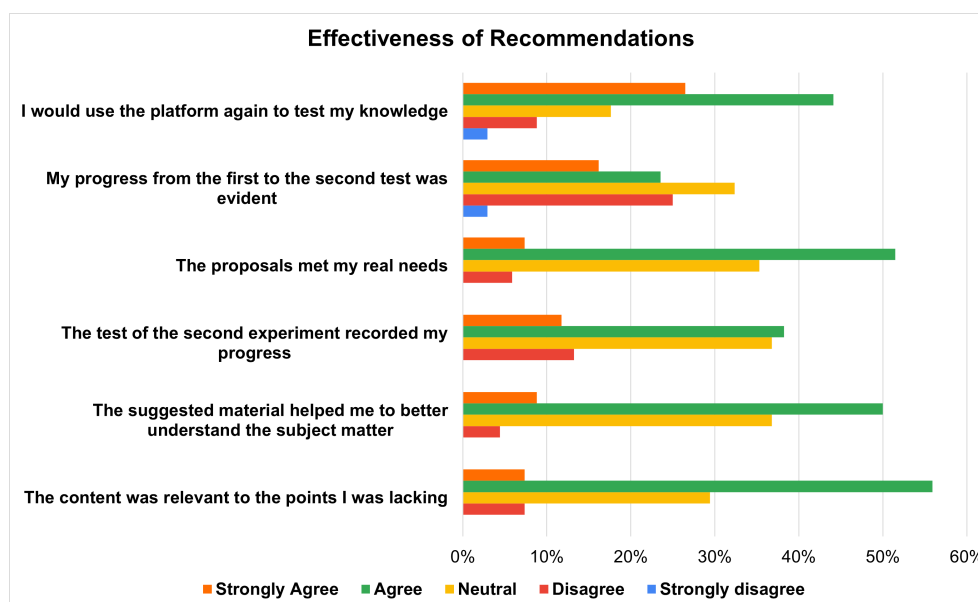


Figure 18. Post-Test: Effectiveness of Recommendations.

These results collectively highlight the platform's success in providing targeted and useful content that genuinely helps students improve academically.

The findings of the quality assurance questionnaire after both tests show a strong willingness among students to continue using the DK-PRACTICE] platform and to recommend it to others.

As illustrated in the Figure 19, most respondents chose higher ratings (4 and 5 on a five-point scale), with 40% indicating that they would continue using the platform in other courses and 36% stating that they would recommend it to fellow students. Furthermore, 22.58% of students reported the highest level of willingness (5) to continue using the platform, while 27.1% showed the highest intention to recommend it. Lower ratings (1 and 2) were minimal, indicating very little resistance to future adoption or endorsement of the system. In general, the students valued its usability, design, and potential as a learning tool.

Despite positive feedback, students identified areas for improvement. In "Pre-Test", technical issues such as system crashes were noted, but local network problems caused these. In "Post-Test", participants recommended features such as a timer, question navigation, a skip option, and semester

access for continuing practice. Overall, the findings indicate that the platform effectively fulfills student learning needs and is seen as a valuable and efficient educational resource.

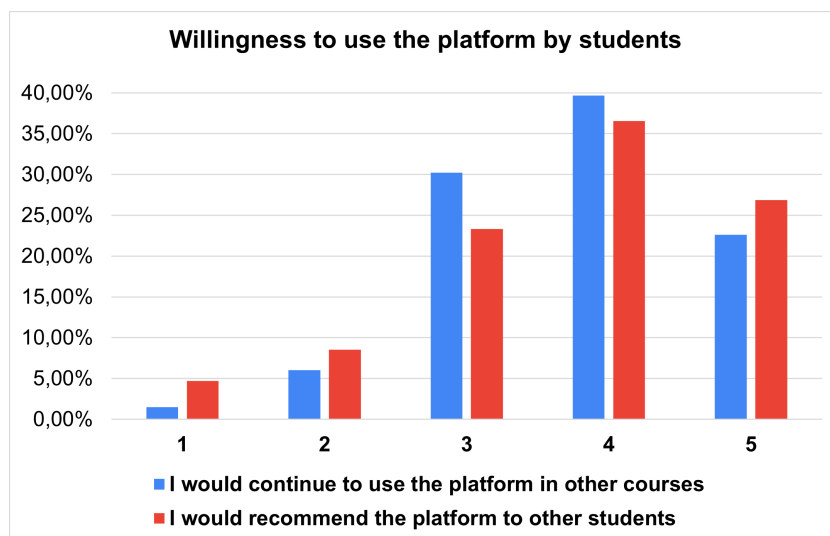


Figure 19. Willingness to use the platform DK-PRACTICE platform by students.

6. Results and Discussion

The deployment of DK-PRACTICE illustrates the practical potential of integrating machine learning-based knowledge tracing with personalized recommendation mechanisms in higher education. By accurately identifying individual learning gaps and providing tailored instructional content, the platform supports both academic achievement and self-regulated learning. Its lightweight yet effective PB-BoW model not only outperforms deeper architectures in predictive accuracy but also offers efficiency in training and scalability for large classrooms. The evaluation of DK-PRACTICE focused on two interrelated dimensions: the technical performance of its knowledge tracing model and the platform's educational effectiveness in a real classroom environment. The predictive capacity of the proposed PB-BoW model was benchmarked against the KT-BiGRU model using three datasets: "ASSISTment 2009", "ASSISTment 2017", and the newly constructed "COaA Course" dataset. In all cases, the PB-BoW approach consistently outperformed KT-BiGRU across both AUC and accuracy metrics. The greatest performance gap appeared in the COaA dataset, where PB-BoW recorded an AUC of 0.805 and an accuracy of 0.787, outperforming KT-BiGRU's 0.736 and 0.706. These results suggest that the PB-BoW model provides more robust predictions, generalizes effectively across datasets of varying size and complexity, and does so with considerably lower computational requirements. Training was conducted efficiently on a standard CPU without the need for GPUs, requiring only around five seconds per epoch, which underscores the model's suitability for practical classroom applications where resources may be limited.

The case study in the "Computer Organization and Architecture" course provided an opportunity to evaluate how effectively the platform identifies knowledge gaps and supports learning progress. A total of 138 students participated in the "Pre-Test" phase, which aimed to establish baseline knowledge levels across 31 concepts. The results of the "Pre-Test" confirmed that the system could identify strengths and weaknesses with a high degree of accuracy. Survey responses indicated that more than sixty percent of students trusted the system's ability to reflect their actual knowledge state, and a majority also reported that the assessment helped them to recognize areas in which improvement was necessary. These findings highlight the role of the platform not only as a diagnostic tool but also as a means of raising students' awareness of their own learning needs.

As shown in Figure 20, the results of the post-test are consistently higher than those of the pre-test across most course concepts. Gains are particularly pronounced in topics such as "QPI", "DDR Memory", "Interrupts", and "Magnetic disk organization", where the post-test results clearly exceeded the

pre-test baselines. This improvement demonstrates that the system's recommendations were not only relevant but also effective in guiding students toward addressing their weaker areas. Some concepts, such as "Execution rate" and "Instruction phases", showed relatively smaller improvements, suggesting areas where further refinement of the recommendation mechanism or expansion of content resources may be required. Overall, the comparative analysis indicates that students benefited significantly from the adaptive learning process, validating the platform's ability to translate diagnostic assessments into tangible learning outcomes.




Figure 20. Comparison of "Pre-Test" (blue line) and "Post-Test" (orange line) performance across 31 course concepts. The results illustrate consistent improvement in student knowledge following engagement with the personalized recommendations provided by DK-PRACTICE.

The results from student questionnaires further support the effectiveness of the platform. During the "Pre-Test", more than 60% of respondents expressed trust in the system's ability to accurately assess their knowledge and identify strengths and weaknesses. Many students reported that the test helped them gain a clearer understanding of where improvement was needed. In the "Post-Test" survey, approximately three-quarters of the participants stated that the platform contributed positively to their progress, with most agreeing that the recommendations were both relevant and useful. Additionally, a large majority expressed willingness to continue using the platform in other courses and to recommend it to peers.

Taken together, these findings highlight the dual strengths of DK-PRACTICE: its robust predictive accuracy and its capacity to generate meaningful educational improvements. The chart comparing Pre- and Post- tests results provides direct evidence that the system not only diagnoses knowledge gaps but also helps students to address them, leading to measurable learning gains within a short timeframe. Importantly, the PB-BoW model's simplicity and computational efficiency show that advanced deep learning architectures are not always necessary to achieve strong results in knowledge tracing, especially when the goal is practical deployment in educational contexts.

The students' positive evaluations further reinforce the value of integrating knowledge tracing with personalized recommendations. The willingness of students to continue using the platform in future courses indicates both satisfaction with its current design and confidence in its potential for broader application. At the same time, the limited improvements in certain topics reveal opportunities for refining the recommendation algorithms and enriching the content base to ensure consistent benefits across all course areas.

In summary, the results demonstrate that DK-PRACTICE successfully bridges the gap between AI-based research and classroom practice. By providing accurate assessments, effective recommendations,

and measurable learning gains, the platform emerges as a scalable and practical tool for advancing personalized learning in higher education.

7. Conclusions and Future Works

This study presented DK-PRACTICE, an intelligent platform that integrates knowledge tracing with personalized recommendations to enhance learning in higher education. The proposed PB-BoW model consistently outperformed a recurrent neural network baseline, achieving higher AUC and accuracy across benchmark datasets and the newly developed “COaA Course” dataset. Its efficiency and low computational requirements demonstrate that robust predictive performance can be achieved without complex deep learning architectures, making the model practical for real classroom deployment.

The case study in the “Computer Organization and Architecture” course confirmed the platform’s educational impact. **Pre-** and **Post-** tests comparisons revealed measurable improvements in student knowledge in most of the concepts evaluated, while surveys showed that students trusted the assessments of the system, valued the relevance of its recommendations, and expressed a strong willingness to reuse them in future courses. These findings validate DK-PRACTICE as both technically reliable and pedagogically effective, successfully bridging the gap between AI-based research and classroom practice.

The findings of this study also provide clear answers to the research questions posed. The PB-BoW model demonstrated higher predictive accuracy and superior efficiency in detecting and predicting student knowledge gaps compared to existing KT models. The personalized recommendations of the platform had a direct impact on student performance, as evidenced by measurable learning gains between the pre- and post-tests, and were reinforced by student reports of improved preparedness. Finally, positive reception, coupled with low computational demands and stable performance in practice, suggests that DK-PRACTICE can be adapted across disciplines and scaled to broader institutional use. Although smaller gains in certain concepts indicate areas for refinement, the results confirm that the platform is accurate, impactful, and scalable in addressing the challenges of personalized learning.

Future work will focus on three directions. First, enhancing the recommendation engine and expanding the content database will ensure more consistent improvements across all concepts. Second, the integration of advanced models, such as attention-based or hybrid approaches, may further improve predictive accuracy while preserving efficiency. Third, scaling the platform across multiple disciplines and institutions, complemented by longitudinal studies, will provide evidence of its broader applicability and long-term educational benefits. Student feedback also points to practical improvements, such as timers, navigation features, and extended access, that could improve usability and engagement.

In summary, DK-PRACTICE demonstrates the potential of combining knowledge tracing with recommendation systems to deliver adaptive, personalized learning. With further refinement and broader implementation, it represents a scalable tool capable of transforming higher education into a more data-driven, student-centered environment.

Author Contributions: Conceptualization, M.D. and K.D., investigation M.D., software M.D., methodology M.D. and K.D., resources M.D., K.D. and A.S., data curation M.D., validation M.D., K.D. and G.K., writing—original draft preparation M.D., supervision K.D., writing—review and editing K.D., A.S., G.K., G.E. and D.K.. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding

Institutional Review Board Statement: Ethical review and approval were waived for this study due to it being conducted within the context of a course

Informed Consent Statement: Verbal informed consent was obtained from all participants involved in the study. Prior to participation, respondents were informed about the purpose of the research, the voluntary nature of their participation, and the confidentiality of their responses. The study involved no sensitive personal data and was conducted in full compliance with the ethical principles outlined in the Declaration of Helsinki (2013 revision).

Data Availability Statement: The GitHub link will be available after the review policy

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

KT	Knowledge Tracing
PB-BoW	Paired-Bipolar Bag-of-Words
BoW	Bag of Words
RNN	Recurrent Neural Network
ML	Machine Learning
AI	Artificial Intelligence
BKT	Bayesian Knowledge Tracing
IRT	Item Response Theory
DKT	Deep Knowledge Tracing
DKVMN	Key-Value Memory Networks
SAKT	Self-Attentive Knowledge Tracing
GIKT	Graph-based Knowledge Tracing
KSGAN	Knowledge Structure-aware Graph Attention Network
EKT	Exercise-Enhanced Sequential Modeling
KT-BiGRU	Knowledge Tracing Bidirectional Gated Recurrent Unit
GPU	Graphics Processing Unit

References

- Anbu, K. Enhancing physics education through artificial intelligence tools. *Scientia: The International Journal for Research* **2025**, *2*, 25–32. <https://doi.org/10.54646/SAPARS.2025.13>.
- PhET Interactive Simulations. <https://phet.colorado.edu>, 2018. University of Colorado Boulder, Accessed July 2025.
- Labster Virtual Labs. <https://www.labster.com>, 2025. Accessed July 2025.
- ChatGPT by OpenAI, 2022. Accessed July 2025.
- Squirrel AI Learning. <https://www.squirrelai.com>, 2025. Accessed July 2025.
- IBM Watson Education. <https://www.ibm.com/watson-education>, 2025. Accessed July 2025.
- Corbett, A.T.; Anderson, J.R. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction* **1995**, *4*, 253–278.
- Baker, F.B. The basics of item response theory. *ERIC Clearinghouse on Assessment and Evaluation* **2008**.
- Piech, C.; Bassen, J.; Huang, J.; Ganguli, S.; Sahami, M.; Guibas, L.; Sohl-Dickstein, J. Deep knowledge tracing. In Proceedings of the Advances in Neural Information Processing Systems, 2015, Vol. 28, pp. 505–513.
- Yang, Y.; Shen, J.; Qu, Y.; Liu, Y.; Wang, K.; Zhu, Y.; Zhang, W.; Yu, Y. GIKT: A graph-based interaction model for knowledge tracing. *arXiv preprint arXiv:2009.05991* **2020**.
- Knewton Alta. <https://www.knewton.com>, 2020. Accessed July 2025.
- Carnegie Learning – MATHia. <https://www.carnegielearning.com>, 2025. Accessed July 2025.
- Fu, S.; et al. AdvKT: An Adversarial Multi-Step Training Framework for Knowledge Tracing. *arXiv preprint arXiv:2401.12578* **2025**.
- Ghosh, S.; Ranjan, P.; Drachsler, H.; Iqbal, Q.; Chakraborty, S.; Yau, J.Y.K. TrueLearn: A family of Bayesian algorithms to match lifelong learners to open educational resources. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2020, Vol. 34, pp. 3555–3562.
- Li, Z.; Yazdanpanah, V.; Wang, J.; Gu, W.; Shi, L.; Cristea, A.I.; Kiden, S.; Stein, S. TutorLLM: Customizing learning recommendations with knowledge tracing and retrieval-augmented generation. *arXiv preprint arXiv:2502.15709* **2025**.
- Wang, H.; Wu, Q.; Bao, C.; Ji, W.; Zhou, G. Research on knowledge tracing based on learner fatigue state. *Complex & Intelligent Systems* **2025**, *11*, 226. <https://doi.org/10.1007/s40747-025-01831-x>.
- Corbett, A.T.; Anderson, J.R. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction* **1994**, *4*, 253–278.

18. Zhang, J.; Shi, X.; King, I.; Yeung, D.Y. Dynamic key-value memory networks for knowledge tracing. In Proceedings of the Proceedings of the 26th International Conference on World Wide Web (WWW). ACM, 2017, pp. 765–774. <https://doi.org/10.1145/3038912.3052581>.
19. Pandey, S.; Karypis, G. Self-attentive knowledge tracing. In Proceedings of the Proceedings of the 12th International Conference on Educational Data Mining (EDM), 2019, pp. 384–389.
20. Tong, S.; Zhang, Y.; Chen, E.; Nie, J.Y. EKT: Exercise-enhanced sequential modeling for student performance prediction. *IEEE Transactions on Knowledge and Data Engineering* **2020**, *34*, 2340–2353. <https://doi.org/10.1109/TKDE.2020.2973984>.
21. Delianidi, M.; Diamantaras, K.I. KT-Bi-GRU: Student Performance Prediction with a Bi-Directional Recurrent Knowledge Tracing Neural Network. *Journal of Educational Data Mining* **2023**, *15*, 1–21.
22. Reddy, A.A.; Harper, M., ALEKS-based Placement at the University of Illinois. In *Knowledge Spaces: Applications in Education*; Falmagne, J.C.; Albert, D.; Doble, C.; Eppstein, D.; Hu, X., Eds.; Springer Berlin Heidelberg: Berlin, Heidelberg, 2013; pp. 51–68. https://doi.org/10.1007/978-3-642-35329-1_3.
23. Settles, B.; Meeder, B. A Trainable Spaced Repetition Model for Language Learning. In Proceedings of the Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Erk, K.; Smith, N.A., Eds., Berlin, Germany, 2016; pp. 1848–1858. <https://doi.org/10.18653/v1/P16-1174>.
24. Mao, S.; Zhan, J.; Deng, Y.; Qin, Y.; Jiang, Y. Improving exercise-level knowledge tracing via knowledge concept-based memory network. *Expert Systems with Applications* **2025**, *284*, 127825. <https://doi.org/10.1016/j.eswa.2025.127825>.
25. Long, T.; Yin, L.; Chang, Y.; Xia, W.; Yu, Y. Simulating Question-answering Correctness with a Conditional Diffusion. In Proceedings of the Proceedings of the ACM on Web Conference 2025, New York, NY, USA, 2025; p. 5173–5182. <https://doi.org/10.1145/3696410.3714508>.
26. Zhou, C.; Liu, X.; Tang, M.; Li, X. The impact of AI-based adaptive learning technologies on motivation and engagement of higher education students. *Computers & Education* **2025**.
27. Hegde, V.; Vishrutha, M.; Shanthappa, P.M.; Bhat, R.; Raveendran, N.; Roshin, C. Analysing learning behaviour: A data-driven approach to improve time management and active listening skills in students. *MethodsX* **2025**, *14*, 103262.
28. Kotsiantis, S.; Pierrakeas, C.; Pintelas, P. Predicting Students' Performance in Distance Learning Using Machine Learning Techniques. *Applied Artificial Intelligence* **2004**, *18*, 411–426. <https://doi.org/10.1080/08839510490442058>.
29. Holmes, W.; Bialik, M.; Fadel, C. *Artificial intelligence in education promises and implications for teaching and learning*; Center for Curriculum Redesign, 2019.
30. Cybenko, G. Approximations by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems* **1989**, *2*, 183–192.
31. Hornik, K.; Stinchcombe, M.; White, H. Multilayer feedforward networks are universal approximators. *Neural Networks* **1989**, *2*, 359–366. [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8).
32. LeCun, Y.; Bengio, Y.; Hinton, G. Deep learning. *nature* **2015**, *521*, 436–444.
33. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* **2012**, *25*.
34. Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*; MIT Press: Cambridge, MA, USA, 2016.
35. Chen, W.; Yang, T. A recommendation system of personalized resource reliability for online teaching system under large-scale user access. *Mobile Networks and Applications* **2023**, pp. 1–12.
36. Hukkeri, G.S.; Goudar, R. Machine Learning-Based Personalized Recommendation System for E-Learners. In Proceedings of the 2022 Third International Conference on Smart Technologies in Computing, Electrical and Electronics (ICSTCEE). IEEE, 2022, pp. 1–6.
37. Da Silva, F.L.; Slodkowski, B.K.; Da Silva, K.K.A.; Cazella, S.C. A systematic literature review on educational recommender systems for teaching and learning: research trends, limitations and opportunities. *Education and information technologies* **2023**, *28*, 3289–3328.
38. Deschênes, M. Recommender systems to support learners' Agency in a Learning Context: a systematic review. *International journal of educational technology in higher education* **2020**, *17*, 50.
39. Troussas, C.; Krouska, A. Path-Based Recommender System for Learning Activities Using Knowledge Graphs. *Information* **2023**, *14*. <https://doi.org/10.3390/info14010009>.
40. Tang, T.Y.; McCalla, G. Smart recommendation for an evolving e-learning system: Architecture and experiment. In Proceedings of the International Journal on E-Learning, 2005, Vol. 4, pp. 105–129.

41. Manouselis, N.; Drachsler, H.; Verbert, K.; Duval, E. *Recommender Systems for Learning*; Springer, 2012. <https://doi.org/10.1007/978-1-4614-4361-2>.
42. Roy, D.; Dutta, M. A systematic review and research perspective on recommender systems. *Journal of Big Data* **2022**, *9*, 59.
43. Qader, W.A.; Ameen, M.M.; Ahmed, B.I. An Overview of Bag of Words;Importance, Implementation, Applications, and Challenges. In Proceedings of the 2019 International Engineering Conference (IEC), 2019, pp. 200–204. <https://doi.org/10.1109/IEC47844.2019.8950616>.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.