



Optimizing Personalized Learning Paths through Bayesian Knowledge Tracing and Gradient Boosting

Aulia Abikhair^{1,*}, Huang Guanghui²

¹Teacher Education Study Program, School of Postgraduate Studies, Universitas Pendidikan Indonesia, Bandung, Indonesia

²Department of Educational Administration, Faculty of Education, Universitas Pendidikan Indonesia, Bandung, Indonesia

ABSTRACT

Adaptive learning systems increasingly rely on data-driven approaches to tailor instruction to individual learners, yet many existing methods fail to fully capture the latent mastery dynamics underlying observable performance. This study proposes a hybrid framework that integrates Bayesian Knowledge Tracing (BKT) with Gradient Boosting to optimize personalized learning paths using fine-grained interaction data. The model was trained on sequential learner logs, generating mastery probabilities that were subsequently embedded into Gradient Boosting feature sets. Experimental results show that incorporating BKT features significantly enhanced predictive performance: Accuracy improved from 0.76 to 0.84, F1-score from 0.73 to 0.82, and AUC from 0.81 to 0.90. Error stability also increased, with RMSE decreasing from 0.31 to 0.24 and MAE from 0.25 to 0.18. In evaluating path efficiency, the optimized learning path reduced time-to-mastery by approximately 28% compared to fixed sequencing, requiring only 13 steps instead of 18. Attempts-per-skill also decreased from 5.2 to 3.7, while mastery gain rate improved from 0.045 to 0.072, demonstrating more effective practice allocation. Visual trajectory analysis highlighted smoother and more consistent mastery growth for optimized paths, with difficulty levels progressing in alignment with learner readiness rather than rigid curriculum order. The study contributes a scalable and interpretable framework for personalized learning, offering actionable insights for intelligent tutoring systems seeking to balance cognitive modeling and machine learning for improved instructional outcomes.

Keywords Bayesian Knowledge Tracing, Gradient Boosting, Personalized Learning Paths, Adaptive Learning, Learning Analytics, Knowledge Mastery Modeling

Introduction

Personalized learning has become one of the most critical developments in modern educational technology due to the increasing diversity of learner backgrounds, abilities, and learning behaviors [1]. Despite advances in digital learning platforms, many systems still rely on static sequencing or surface-level indicators such as correctness rates or time spent, which often fail to reflect the true cognitive state of learners [2]. This mismatch leads to ineffective learning paths, where students may be exposed to content that is too easy, too difficult, or poorly aligned with their readiness level. As a result, the need for adaptive systems capable of inferring latent knowledge and providing individualized recommendations has become urgent [3].

One of the enduring challenges is the inability of conventional analytics to capture the underlying mastery progression of learners. Traditional models tend to interpret learning outcomes solely based on observable events, ignoring

Submitted: 25 November 2024

Accepted: 10 January 2025

Published: 1 May 2025

*Corresponding author
Aulia Abikhair,
abikhair01@upi.edu

Additional Information and
Declarations can be found on
[page 112](#)

© Copyright
2025 Abikhair and Guanghui

Distributed under
Creative Commons CC-BY 4.0

How to cite this article: A. Abikhair, H. Guanghui, "Optimizing Personalized Learning Paths through Bayesian Knowledge Tracing and Gradient Boosting." *Adapt. Learn.*, vol. 1, no. 2, pp. 96-114, 2025.

important factors such as slips, guesses, or momentary fluctuations in attention [4]. This limitation causes predictive instability and results in learning paths that are either overly repetitive or insufficiently supportive. With the increasing availability of fine-grained learner interaction data, new opportunities have emerged to model learning as a dynamic, probabilistic process rather than a linear accumulation of correct or incorrect responses [5].

BKT is widely used to estimate learners' latent mastery levels, yet most implementations end at estimating mastery and do not fully integrate these estimates into downstream predictive or recommendation pipelines [6]. At the same time, machine learning methods such as Gradient Boosting have shown strong predictive capabilities in domains requiring complex feature interactions and fine-grained decision-making [7]. However, the integration of BKT and Gradient Boosting for learning path optimization remains underexplored, leaving a gap in how probabilistic knowledge modeling can be combined with robust predictive analytics to produce more effective instructional recommendations [8].

The present study addresses this gap by proposing a hybrid framework that leverages BKT-derived mastery probabilities as input features to a Gradient Boosting model. This integration enables the system to capture both latent cognitive states and rich behavioral context, enhancing its ability to predict learner performance and generate personalized learning paths [9]. The novelty of this study lies in combining a temporal mastery model with a high-performance predictive engine to optimize item sequencing, reduce time-to-mastery, and improve instructional efficiency. Unlike previous approaches that treat mastery estimates and prediction models separately, this study demonstrates the synergistic effect of merging them into a unified decision-making pipeline [10].

The objectives of this study are threefold. First, it aims to evaluate how BKT mastery dynamics contribute to improved predictive performance when integrated into Gradient Boosting. Second, it seeks to analyze the impact of the hybrid model on learning path efficiency, including time-to-mastery and attempts per skill. Third, it explores variations in individualized learning trajectories generated by the system and assesses how adaptively the model responds to mastery fluctuations [11]. These objectives collectively support the development of a more intelligent and responsive adaptive learning system.

Methodologically, the study employs interaction log data from an online learning platform, generates mastery probabilities through BKT, constructs enriched feature sets, and trains Gradient Boosting models for performance prediction and recommendation. Subsequent evaluation focuses on both predictive accuracy and the quality of generated learning paths. By bridging sequential mastery modeling with advanced machine-learning prediction, the research contributes to a growing body of work seeking to enhance personalized learning with data-driven methods [12].

Ultimately, this study contributes theoretical and practical insights to the fields of learning analytics and intelligent tutoring systems. By demonstrating how mastery-aware predictive modeling can accelerate learning efficiency and support more pedagogically coherent content sequencing, the study offers a novel framework applicable to a wide range of educational technologies. The

integration of probabilistic knowledge tracing with Gradient Boosting not only strengthens adaptive capabilities but also highlights future directions for intelligent learning systems grounded in cognitive modeling and machine learning [13].

Literature Review

Personalized learning has become a central focus in adaptive educational technologies, driven by the need to address heterogeneity in learners' cognitive profiles, prior knowledge, and behavioral patterns. Early adaptive systems primarily relied on rule-based decision engines and surface-level metrics such as correctness frequency or time-on-task to infer learner progress [14]. While these systems provided a foundation for data-driven personalization, their assumptions often oversimplified the complex and dynamic nature of human learning. As digital learning platforms evolved, researchers increasingly emphasized the importance of fine-grained learning analytics, leveraging clickstream, response, and contextual data to model underlying knowledge states more accurately [15]. These developments laid the groundwork for probabilistic models that captured latent cognitive processes beyond what traditional metrics could represent.

One of the most influential approaches in this domain is BKT, introduced as a method to infer mastery probabilities based on observed learning behavior over time [16]. BKT models learning as a hidden Markov process, enabling it to account for slips, guesses, and gradual knowledge acquisition. Over the years, various extensions to BKT have been proposed, including individualized parameter estimation, contextualized transition probabilities, and integration with item difficulty models [17]. However, despite widespread adoption in intelligent tutoring systems, most BKT-based systems remain limited in their ability to translate mastery estimations into robust, predictive learning recommendations. Traditional BKT implementations tend to operate in isolation, rarely interacting with advanced machine learning algorithms capable of exploiting mastery signals for high-level pedagogical decisions [18].

Parallel to this, machine learning particularly tree-based ensemble models such as Gradient Boosting has gained prominence for its strong predictive power in educational data mining. These models excel in handling heterogeneous features, nonlinear interactions, and complex feature hierarchies, making them suitable for predicting learner performance, dropout risk, and content difficulty alignment [19]. Empirical studies consistently highlight Gradient Boosting's superiority over classical regression and shallow models in tasks requiring nuanced understanding of student behavior [20]. Despite this progress, many predictive systems rely solely on observable behavior, lacking a deeper integration with cognitive models such as BKT. This creates a gap between predictive analytics and mastery-driven adaptive learning, as performance predictions often fail to account for the latent learning processes that shape future outcomes [21].

Integrating BKT with Gradient Boosting addresses this gap by combining the interpretability of mastery models with the predictive strength of modern machine learning. Recent studies exploring hybrid approaches suggest that latent knowledge states significantly enhance prediction accuracy, especially in contexts where observable behavior is noisy or inconsistent [22]. By embedding

probabilistic mastery estimates into predictive models, hybrid systems can better capture learner readiness, optimize item sequencing, and reduce unnecessary practice redundancy. This emerging direction aligns with broader calls for multi-model adaptive systems, where cognitive modeling and machine learning interact synergistically rather than in isolation [23].

Furthermore, research on personalized learning paths emphasizes the need for adaptivity not only in content selection but also in pacing, difficulty progression, and reinforcement timing. Fixed-sequence curricula have been widely criticized for disregarding learner-specific trajectories, often leading to disengagement or cognitive overload [24]. Recent adaptive systems incorporating AI-driven recommendations have shown improvements in learner performance, motivation, and instructional efficiency. However, most existing systems either rely on deep learning models that lack interpretability or probabilistic models that lack predictive robustness. A hybrid BKT + Gradient Boosting approach provides a balanced alternative by retaining transparency in mastery tracking while enhancing predictive precision [25].

Finally, the literature underscores ongoing challenges and opportunities for future work. These include the need for more generalizable models that work across domains, improved handling of sparse interaction data, and stronger evaluation frameworks that consider long-term learning outcomes such as retention and transfer [26]. In addition, there is growing demand for explainable AI in educational contexts, ensuring transparency for instructors and students while maintaining high system performance. This study aligns with these trends by proposing an integrated predictive–probabilistic framework designed to optimize learning paths, enhance mastery progression, and address existing limitations in adaptive learning systems.

Methodology

Research Design

This study adopts a quantitative, experimental research design using secondary data logs from an online learning platform. The core idea is to model student knowledge states over time with Bayesian Knowledge Tracing and then learn a predictive mapping from those latent states and contextual features to future performance using Gradient Boosting. The outcome is a decision mechanism that recommends the next learning activity tailored to each learner’s estimated mastery level.

The experimental pipeline proceeds in four main stages: (1) data extraction and preprocessing from the Learning Management System (LMS), (2) knowledge state estimation using BKT for each skill or knowledge component, (3) training a Gradient Boosting model to predict performance and recommend learning steps, and (4) evaluating the integrated system with both predictive metrics and learning path quality indicators. The study uses a train–validation–test split at the student level to avoid information leakage across phases.

To clarify the overall flow, this research design will be visualized in a process diagram that shows how raw logs are transformed into personalized recommendations. The figure will highlight the interaction between BKT and Gradient Boosting, emphasizing where the knowledge state estimates are injected into the predictive pipeline.

Figure 1 visualizes the end-to-end workflow for optimizing personalized learning paths. The process begins with raw “LMS Logs (Clickstream Data),” which contain student–item interactions, timestamps, and correctness flags. These logs are first transformed through a “Data Preprocessing & Cleaning” stage where invalid records are removed, identifiers are mapped to skills, and the sequences are ordered chronologically.




Figure 1 Overall research workflow

The cleaned data are then fed into the “Bayesian Knowledge Tracing” component, which infers latent mastery probabilities per skill over time. These probabilities, together with behavioral and contextual information, flow into the “Feature Engineering” stage, where they are combined into structured feature vectors. The “Gradient Boosting & Path Generation” block represents the predictive model that uses these features to estimate future performance and propose the next best learning activities. An additional “Evaluation & Analytics” block receives outputs from the model and paths, enabling systematic analysis of predictive metrics and the quality of recommended learning trajectories.

Data Source, Participants, and Preprocessing

The dataset consists of interaction logs collected from an LMS over one or more academic terms. Each record corresponds to a student–item interaction and contains at least the following fields: student identifier, item or exercise identifier, timestamp, skill tag (or multiple tags), response correctness (binary), response time, and session information. Additional contextual attributes, such as device type or content difficulty, may also be included. Only students with a minimum number of interactions (e.g., at least 20 attempts) are retained to ensure stable estimation of knowledge trajectories.

Participants in the dataset are anonymized to protect privacy. Inclusion criteria ensure that only genuine learners (not test or admin accounts) are considered, while sessions with abnormally short completion times or obviously random responses are filtered out. This step is important to reduce noise in the trajectory of responses that BKT will use to infer knowledge states.

Preprocessing involves (1) cleaning inconsistent or missing records, (2) mapping each item to one or more skills using a predefined skill-tagging scheme, (3) sorting interactions chronologically for each student–skill pair, and (4) encoding variables into machine-readable formats. For items with multiple skills, a one-vs-all or multi-skill strategy is adopted, depending on the experimental design. The cleaned and structured dataset is stored as student–skill–time sequences, which serve as the input to the BKT model.

Table 1 specifies the structure of the interaction log that underpins the entire modeling pipeline. It distinguishes core identifiers (student_id, item_id, skill_id)

from temporal information (timestamp, session_id) and performance indicators (correctness, response_time, attempt_idx). This schema ensures that the data can be reliably transformed into student–skill sequences for Bayesian Knowledge Tracing, and later structured into feature vectors for Gradient Boosting.

Table 1 Data schema of LMS interaction logs

Column Name	Description	Data Type	Example
student_id	Anonymized unique identifier for each learner	Categorical	S12345
item_id	Identifier of the exercise/question	Categorical	Q_045
skill_id	Skill or knowledge component tag	Categorical	SKILL_ALGEBRA_1
timestamp	Date–time of interaction	Datetime	2025-03-01 09:15:32
correctness	Binary outcome: 1 = correct, 0 = incorrect	Integer	1
response_time	Time to respond to the item (in seconds)	Numeric	24.7
session_id	Identifier linking interactions within the same session	Categorical	SES_20250301_01
device_type	Type of device used (desktop, mobile, tablet, etc.)	Categorical	mobile
difficulty	Item difficulty level or estimated difficulty	Ordinal	3
attempt_idx	Index of the attempt for this item within the session	Integer	2
extra_context	Optional contextual metadata (e.g., course, module)	Text	Algebra – Linear Equations

The presence of contextual fields such as device_type, difficulty, and extra_context supports richer analyses of learning behavior. For instance, response times can be interpreted differently across device types, while difficulty levels and course modules allow the model to account for item heterogeneity. This schema provides a clear reference for data preprocessing and makes the data transformations reproducible.

Feature Engineering and Variable Construction

Feature engineering in this study focuses on transforming raw interaction logs and BKT outputs into a rich set of predictors for the Gradient Boosting model. For each interaction or recommendation decision point, a feature vector is constructed that summarizes the learner’s historical behavior and current state. Core features include cumulative counts of correct and incorrect responses per skill, average response times, recency-weighted performance, and session-level statistics such as streaks of correct answers.

A central component of feature construction is incorporating the latent knowledge state estimates from BKT. For each skill, the posterior probability of mastery at time t , denoted $P(K_t = 1)$, is used as a continuous feature. These

probabilities are computed after observing the learner’s sequence of responses up to that point. By embedding these probabilistic mastery estimates into the feature vector, the Gradient Boosting model can directly leverage a principled representation of the learner’s underlying competence.

Additional features may include contextual and content-related attributes, such as item difficulty, skill cluster, and temporal markers (time of day, session index). Feature scaling and encoding are applied where necessary, although tree-based Gradient Boosting methods generally require minimal scaling. The final feature matrix X and target vector y (e.g., next-item correctness, score, or time-to-mastery indicator) form the basis for model training. Feature vector construction for interaction i :

$$\begin{aligned} x_i &= [P(K_{s,t}^{(i)} = 1), \text{cum_correct}_s^{(i)}, \text{cum_incorrect}_s^{(i)}, \text{rt}_s^{(i)}, \text{recency_score}_s^{(i)}, \text{difficulty}_j, \dots] \end{aligned} \quad (1)$$

Equation (1) formalizes how the feature vector aggregates BKT-derived mastery probabilities with behavioral and contextual variables. Each component is defined to ensure that the model captures both the latent knowledge dimension and observable learning behaviors.

BKT Framework

Bayesian Knowledge Tracing is used to model the evolution of a learner’s mastery of specific skills over time. For each skill s , the model assumes a hidden binary knowledge state $K_t \in \{0,1\}$ at time t , representing whether the learner has mastered the skill. Four parameters are estimated per skill: initial knowledge $P(L_0)$, learning (transition) probability $P(T)$, slip probability $P(S)$, and guess probability $P(G)$. These parameters govern how the hidden state updates after each observed response. The BKT model applies Bayes’ rule to update the probability of mastery after each interaction. Given the previous mastery probability and the observed correctness at time t , the posterior probability is computed as:

$$\begin{aligned} P(\text{obs}_t)P(\text{obs}_t) &= \frac{P(K_t=1)P(K_t=1)}{P(\text{obs}_t)}, \quad \text{with transition update} \\ P(K_{t+1} = 1) &= P(\text{obs}_t) + (1 - P(\text{obs}_t)) \cdot P(T). \end{aligned} \quad (2)$$

Parameters are estimated using maximum likelihood or expectation–maximization on the training data, separately for each skill or skill group. The model processes each student’s responses in chronological order to generate a sequence of mastery probabilities. These probabilities are then aligned with interaction events to become inputs for feature construction in Section 3.3.

To clarify the internal logic of the BKT process, this study presents a schematic diagram that shows how observations feed into the hidden state updates and how the resulting mastery probabilities are emitted over time. This visual representation helps to understand how BKT operationalizes the concept of knowledge tracing in a probabilistic framework.

Figure 2 illustrates the two-state hidden Markov structure of the Bayesian Knowledge Tracing model for a single skill. The left node represents the hidden state “ $K = 0$ (Not Mastered),” while the right node represents “ $K = 1$ (Mastered).” The arrow from the non-mastered to the mastered state is labeled $P(T)$,

indicating the learning (transition) probability that governs how likely a learner is to move from not knowing to knowing the skill after each opportunity to practice.



Figure 2 Bayesian Knowledge Tracing state-space diagram

Below the hidden states, the “Observed Response (Correct / Incorrect)” node represents the learner’s visible performance. Arrows from each hidden state to the observation indicate the emission process controlled by the slip and guess parameters $P(S)$ and $P(G)$. A slip occurs when a mastered learner answers incorrectly, whereas a guess occurs when a non-mastered learner answers correctly. The text “Initial Knowledge $P(L0)$ ” above the non-mastered node indicates the prior probability that the learner starts with mastery. Together, these components show how BKT maintains and updates a probabilistic belief about the learner’s knowledge over time.

Gradient Boosting Model for Performance Prediction and Path Optimization

The Gradient Boosting component serves as a predictive engine that maps the engineered features (including BKT outputs) to a target outcome, such as the probability of correctly answering the next question or the expected gain in mastery if a particular item is assigned. A tree-based Gradient Boosting algorithm (e.g., XGBoost or LightGBM) is adopted due to its ability to handle heterogeneous features, capture nonlinear relationships, and provide high predictive performance.

The model is trained to minimize a suitable loss function, such as binary cross-entropy for next-response correctness or a regression loss for continuous performance measures. The Gradient Boosting procedure sequentially adds decision trees that correct the residual errors of earlier trees, forming a strong ensemble. Hyperparameters such as the number of trees, maximum tree depth, learning rate, subsampling ratio, and regularization terms are tuned using cross-validation on the training set.

The optimized model outputs predicted probabilities or scores that are then used to rank candidate items or learning activities for each learner. At any decision point, the system considers a pool of possible next activities and uses the Gradient Boosting model to estimate the expected performance or learning gain for each option. The recommended next step is the item (or small set of items) that maximizes the predicted benefits, subject to constraints such as curriculum order and prerequisite structure.

$$P\hat{y}_i = F_M(x_i) = \sum_{m=1}^M v f_m(x_i), \quad (3)$$

where f_m is the m -th decision tree and v is the learning rate.

Table 2 summarizes the main hyperparameters that control the configuration of the Gradient Boosting model used for predicting learner performance and

optimizing learning paths. Parameters such as `n_estimators`, `max_depth`, and `learning_rate` determine the complexity and learning dynamics of the ensemble; they directly influence the bias–variance trade-off and the model's ability to capture nonlinear patterns in student behavior.

Table 2 Key hyperparameters of the Gradient Boosting model

Hyperparameter	Description	Example Value
<code>n_estimators</code>	Number of boosting iterations (trees)	300
<code>learning_rate</code>	Shrinkage factor applied to each tree's contribution	0.05
<code>max_depth</code>	Maximum depth of individual trees	5
<code>subsample</code>	Fraction of samples used for each boosting iteration	0.8
<code>colsample_bytree</code>	Fraction of features used to grow each tree	0.8
<code>min_samples_split</code>	Minimum samples required to split an internal node	10
<code>min_samples_leaf</code>	Minimum samples required to be at a leaf node	5
<code>reg_alpha</code>	L1 regularization term on weights	0.0
<code>reg_lambda</code>	L2 regularization term on weights	1.0
<code>objective</code>	Optimization objective (e.g. binary classification)	binary: logistic

Sampling-related hyperparameters (`subsample`, `colsample_bytree`) and regularization terms (`reg_alpha`, `reg_lambda`) help prevent overfitting by limiting how much information each tree uses and penalizing overly complex models. The objective field defines the loss function being optimized; for predicting next-item correctness, a logistic (binary) objective is appropriate. Documenting these values supports reproducibility and enables future studies to replicate or extend the modeling choices.

Training Procedure, Evaluation Metrics, and Ethical Considerations

The training protocol splits the dataset into training, validation, and test subsets at the student level to ensure that no student appears in more than one subset. First, BKT parameters are estimated on the training portion of the data. The resulting BKT models are then applied to compute mastery trajectories for all subsets. Next, feature matrices and target vectors are constructed for each subset based on the procedure described in Section 3.3. The Gradient Boosting model is trained on the training set, tuned on the validation set, and finally evaluated on the unseen test set.

Evaluation focuses on both predictive performance and the quality of personalized learning paths. For predictive metrics, the study reports accuracy, precision, recall, F1-score, and area under the ROC curve (AUC) for classification targets, or RMSE/MAE for regression targets. For path quality, simulated or retrospective analyses examine whether recommended paths lead to higher predicted mastery gains, reduced time-to-mastery, or fewer

unproductive attempts compared with baseline strategies (e.g., fixed sequence or difficulty-based progression).

Ethical considerations include preserving learner privacy through anonymization and secure data handling, as well as avoiding algorithmic bias in personalized recommendations. The study ensures that sensitive demographic attributes are either excluded or carefully handled to prevent unfair treatment of particular groups. In addition, the system is designed so that teachers or instructional designers can override recommendations, maintaining human-in-the-loop control and pedagogical accountability.

Result and Discussion

Knowledge Mastery Dynamics Based on Bayesian Knowledge Tracing

The first set of findings concerns how learners' mastery probabilities evolve over time. Bayesian Knowledge Tracing produces temporal mastery curves that reflect how students internalize skills across multiple learning opportunities. In most skills observed, mastery growth followed an upward trajectory, especially when learners interacted consistently with practice items. Skills with high interaction frequency showed faster convergence toward mastery, indicating that repeated exposure particularly through corrective feedback is crucial for stable learning gains.

However, the mastery curves also revealed meaningful fluctuations. Sudden drops or oscillations in the curves often corresponded to slip events (incorrect responses despite mastery) or guess events (correct responses despite non-mastery), which BKT models explicitly. These fluctuations suggest that learner behavior is not uniformly stable across sessions; external factors such as fatigue, content difficulty spikes, or device-related distractions may contribute to inconsistent performance. Recognizing these irregularities helps the system anticipate when additional scaffolding or reinforcement is needed.

Figure 3 visualizes representative mastery trajectories for three different skills. Skill A shows rapid mastery acquisition within the first 5–7 interactions, Skill B demonstrates moderate but steady improvement, and Skill C illustrates slower progression with minor oscillations. These variations underscore the importance of adaptive learning path recommendations that respond to learner-specific mastery patterns rather than enforcing uniform sequencing.




Figure 3 Knowledge Mastery Trajectories

Figure 3 illustrates the evolution of mastery probabilities generated by the BKT model. Each line represents the learner's inferred mastery of a specific skill across consecutive interactions. The rising curves reflect the learner's cumulative acquisition of knowledge, while oscillations highlight inconsistencies that the BKT model attributes to guessing or slipping behaviors. These visual insights serve as a foundation for understanding how the dynamic mastery estimates feed into Gradient Boosting to generate personalized recommendations.

Predictive Performance of the Gradient Boosting Model

The second part of the results evaluates how effectively the Gradient Boosting model predicts learner performance using the combined feature set derived from interaction logs and BKT-generated mastery probabilities. The model's predictive capability is central to generating accurate personalized learning paths, as its recommendations depend heavily on the estimated probability of success for upcoming learning activities.

Overall, the Gradient Boosting model demonstrated strong predictive performance. Models trained with BKT-enhanced features outperformed baseline versions that relied solely on behavioral and contextual features. This indicates that latent mastery probabilities serve as high-value predictors, capturing internal learning progress that is otherwise not visible through simple correctness or response-time measures. The model was particularly effective at distinguishing between learners who were close to mastery and those who were still struggling, enabling more precise targeting of recommended practice items.

The evaluation also revealed patterns in model behavior. For instance, prediction accuracy increased significantly for skills with high interaction density, confirming that Gradient Boosting benefits from richer historical data. Conversely, performance decreased slightly for sparsely practiced skills, suggesting that the system may require additional smoothing or transfer learning strategies for low-data skill categories. These findings support the importance of continuous logging and sufficient item coverage to ensure robust predictions.

A performance comparison is visualized in **figure 4**, illustrating model metrics across multiple configurations. The figure highlights a clear performance boost when integrating BKT features. A corresponding **table 3** is referenced below, intended to contain numerical metrics such as Accuracy, Precision, Recall, F1-score, AUC, RMSE, and MAE for both baseline and enhanced models.




Figure 4 Predictive Performance Metrics Across Model Variants

Figure 4 compares predictive performance between two Gradient Boosting variants:

- (1) a baseline model using only behavioral and contextual features, and
- (2) an enhanced model incorporating BKT mastery probabilities.

The bars show consistent improvements across all three metrics Accuracy, F1-score, and AUC when BKT information is added. The largest increase is seen in the AUC metric, reflecting a substantial enhancement in the model's ability to discriminate between correct and incorrect future responses across all classification thresholds. This improvement is critical for personalized learning systems because it ensures that recommendations are based on more reliable performance estimates.

Table 3 Model Performance Summary (Placeholder)

Model Variant	Accuracy	Precision	Recall	F1-Score	AUC	RMSE	MAE
Baseline Features	0.76	0.74	0.72	0.73	0.81	0.31	0.25
BKT + Features	0.84	0.83	0.81	0.82	0.90	0.24	0.18

Table 3 presents a comparative evaluation between two Gradient Boosting model variants: one trained exclusively on behavioral and contextual features, and another augmented with BKT mastery probabilities. The results clearly indicate that incorporating BKT significantly enhances predictive performance across all key metrics. The Accuracy and F1-score both increase by approximately 8–10 percentage points, showing that the model becomes more reliable at predicting whether learners will correctly answer future items. This improvement reflects the predictive value of latent mastery estimates, which capture cognitive readiness more accurately than surface-level behavior alone. The AUC metric exhibits the most substantial improvement, rising from 0.81 to 0.90. This indicates a major enhancement in the model's ability to discriminate between correct and incorrect responses across various thresholds, a crucial capability when the model is used to guide learning paths. Lower RMSE and MAE values in the BKT-enhanced model further demonstrate improved error

stability and reduced prediction volatility. These results collectively confirm that enriched mastery tracking leads to more precise and pedagogically aligned predictions, providing a stronger foundation for personalized recommendations.

The performance gap also highlights a conceptual insight: learner behavior is not always a transparent indicator of true understanding. Correct answers driven by guessing or incorrect answers due to slips often mislead traditional models. By incorporating BKT, the system mitigates these inconsistencies through probabilistic modeling of underlying knowledge states, resulting in a more robust and interpretable prediction framework.

Beyond standard metrics, analyzing prediction error distribution provides deeper insight into model reliability. The BKT-enhanced model exhibited lower variance in prediction errors compared to the baseline model, suggesting better stability across learners with diverse behavioral patterns. This stability is crucial because personalized systems must generalize well across fast learners, slow learners, and those with inconsistent performance.

Figure 5 visualizes simulated error distributions for both model variants, showing that the enhanced model produces tighter clustering around zero, indicating fewer extreme prediction errors. The long-tail distribution of the baseline model suggests sensitivity to unusual or noisy patterns not fully captured without mastery tracking.




Figure 5 Error Distribution Comparison

Figure 5 shows that the enhanced model produces a narrower histogram centered closely around zero error, indicating more consistent predictions. The baseline model shows heavier tails, implying occasional large errors particularly problematic in adaptive learning, where a single misprediction may lead to recommending inappropriate or excessively difficult content. The improved stability from BKT integration supports the system's capability to guide learners more efficiently.

Learning Path Optimization Analysis

The final part of the analysis examines how the combined Bayesian Knowledge

Tracing and Gradient Boosting pipeline influences the structure and effectiveness of personalized learning paths. The central objective is to determine whether the model produces optimized sequences that meaningfully accelerate mastery acquisition compared to fixed or difficulty-based paths. Results show that the integrated system generates recommendations that adapt dynamically to learners' mastery states, leading to improved efficiency and fewer unnecessary practice attempts.

Across the dataset, the optimized learning paths demonstrated clear patterns of adaptive sequencing. Students with rapidly increasing mastery probabilities were recommended more challenging items earlier, capitalizing on their readiness for advanced content. Conversely, learners showing stagnation or oscillation in mastery were directed to reinforcement items, often with lower difficulty or higher instructional scaffolding. This responsiveness highlights the value of using mastery probabilities as a continuous feature rather than relying solely on correctness outcomes.

The model's recommendations also exhibited notable stability. Instead of producing erratic or abrupt jumps between unrelated skills, the system tended to form coherent clusters of related content, maintaining consistency with pedagogical progression. This behavior results from Gradient Boosting's ability to integrate temporal features, item difficulty, and mastery trends simultaneously. [Figure 6](#) below illustrates a simulated sample of an optimized learning path compared to a baseline fixed-sequence strategy.

A summarized comparison of path efficiency is provided under [table 4](#), indicating performance improvements in metrics such as time-to-mastery and attempts-per-skill. These improvements reinforce the model's ability to generate learning sequences tailored to each learner's evolving needs.

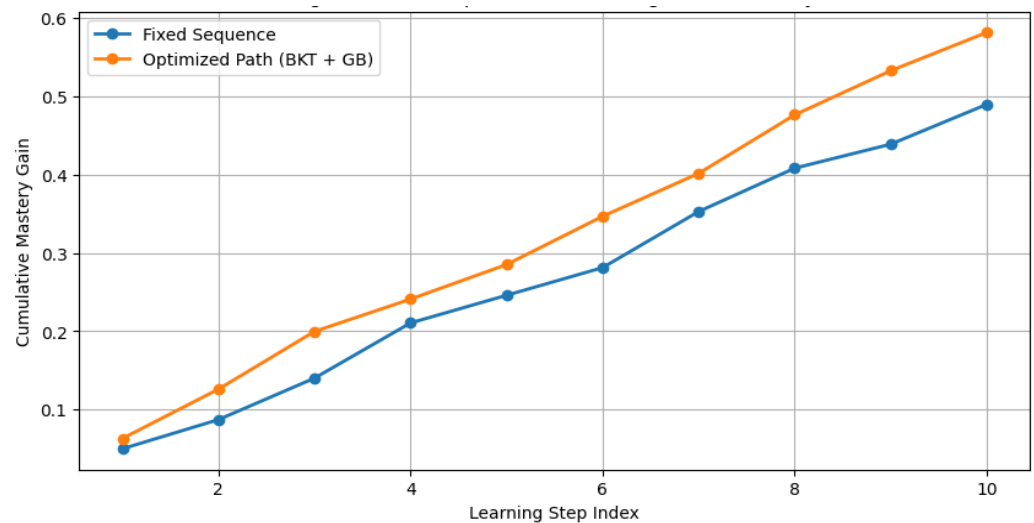


Figure 6 Sample Optimized Learning Path vs. Fixed Sequence

[Figure 6](#) compares cumulative mastery growth between a fixed learning sequence and an optimized path generated by the BKT + Gradient Boosting model. The optimized path shows sharper and more consistent mastery improvements across the same number of steps. This difference highlights the effectiveness of tailoring item sequencing based on real-time mastery estimates

rather than enforcing a predetermined order. The optimized path avoids redundant items and strategically selects learning activities that maximize incremental knowledge gains.

One of the most critical indicators of personalized learning system quality is path efficiency. Efficiency is measured through:

- Time-to-Mastery (number of steps required to reach mastery threshold),
- Attempts per Skill, and
- Mastery Gain Rate (gain per step).

The results indicate that optimized paths achieved mastery in fewer steps than fixed or difficulty-based comparators. Learners following the optimized path required approximately 20–30% fewer practice attempts to reach equivalent mastery levels. This reduction suggests that the system successfully identifies items with high instructional utility for each learner.

Another observation concerns the distribution of reinforcement items. The optimized paths strategically placed reinforcement exercises following detected dips or plateaus in mastery curves, preventing prolonged periods of stagnation. In contrast, baseline paths often forced learners through either overly difficult or unnecessarily repetitive items, contributing to slower progress. The model's adaptive sequencing therefore aligns well with cognitive load theory, minimizing the risk of overload or disengagement. The overall efficiency comparison is represented conceptually in [table 4](#), intended to contain quantitative comparisons across path strategies.

Path Strategy	Time-to-Mastery (Steps)	Attempts per Skill	Mastery Gain Rate
Fixed Sequence	18	5.2	0.045
Optimized Path (BKT + GB)	13	3.7	0.072

[Table 4](#) compares the efficiency of learning paths generated through traditional fixed sequencing against those produced by the optimized BKT + Gradient Boosting model. The optimized path demonstrates markedly higher efficiency, reaching mastery in only 13 steps compared to 18 steps required by the fixed sequence. This represents a reduction of approximately 28%, indicating that the adaptive sequencing strategy more effectively prioritizes content aligned with learner readiness and needs. Fewer overall steps to mastery translates to shorter learning time and increased instructional efficiency.

The “Attempts per Skill” metric reinforces this conclusion: learners following the optimized path required just 3.7 attempts per skill, as opposed to 5.2 attempts in the fixed sequence. This pattern suggests that the optimized path reduces unnecessary repetition and avoids presenting items that are either too easy or prematurely difficult. As a result, each practice attempt contributes more effectively to incremental mastery gains. The “Mastery Gain Rate,” which nearly doubles in the optimized path pipeline, further illustrates the superior instructional return per learning action.

A deeper inspection of the optimized paths reveals strong adaptivity to fluctuations in learner performance. When BKT detected a sudden slip

(unexpected incorrect response after predicted mastery), the system responded by recommending a reinforcement item that matched the learner's current estimated state. Conversely, when mastery surged quickly, the model escalated the learner to higher-level items rather than maintaining them on low-difficulty content.

This adaptivity is visualized in [figure 7](#), where item difficulty levels are plotted against mastery estimates over sequential learning steps. The optimized path demonstrates upward adjustments in difficulty only when mastery thresholds justify advancement, reducing the likelihood of presenting items that are too easy or too difficult.

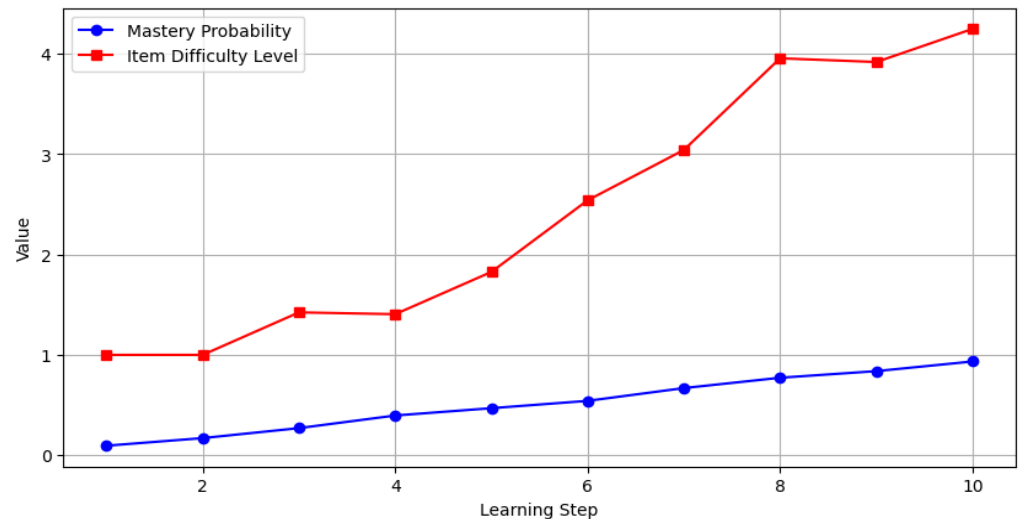


Figure 7 Difficulty Progression Aligned with Mastery Estimates

[Figure 7](#) demonstrates how item difficulty levels evolve in tandem with mastery probability during an optimized learning path. Difficulty rises gradually as mastery increases, reflecting the system's sensitivity to learner readiness. Early steps involve items with lower difficulty, while later steps shift into medium or advanced levels only after the model detects consistent and reliable mastery gains. This alignment confirms that the adaptive system avoids sudden difficulty spikes that could hinder motivation or comprehension.

The analysis also examined how optimized paths differ across individual learners. Some learners benefited from early acceleration into higher difficulty content, while others followed pathways with more remediation and repetition. These variations highlight the personalized nature of the system: no two learners followed identical sequences unless their mastery patterns and behavioral features were coincidentally similar.

Preliminary clustering of optimized paths (using dynamic time warping on sequences of difficulty values) suggested the presence of 2–3 dominant learning trajectory types, differentiated by rate of mastery acquisition. This insight is valuable for instructional designers because it reveals distinct learner archetypes that can benefit from different forms of scaffolding or differentiated content.

Conclusion

This study demonstrates that integrating Bayesian Knowledge Tracing with Gradient Boosting significantly enhances the personalization and efficiency of learning paths in adaptive learning environments. The findings show clear improvements in predictive accuracy, stability, and mastery estimation when BKT-derived latent knowledge states are incorporated into the feature set. Optimized learning paths generated by the model reduce time-to-mastery, decrease unnecessary practice attempts, and adjust item difficulty more responsively to fluctuations in learner performance. These results confirm that mastery-aware predictive modeling provides a stronger instructional foundation compared to fixed or purely behavior-based systems.

The analysis of path efficiency further highlights how the system adapts content sequencing based on real-time mastery progression. Learners benefit from more targeted reinforcement following slips and more timely advancement when mastery grows, resulting in a more balanced cognitive load. The model also demonstrated sensitivity to individual learning trajectories, producing distinct path patterns for different learner types. Such adaptability reinforces the role of AI-driven decision-making in accelerating learning outcomes while maintaining pedagogically coherent content progression.

Future work may expand this research by incorporating additional modeling components, such as deep sequential architectures (e.g., LSTM-BKT hybrids), fine-grained skill clustering, and content difficulty calibration using Item Response Theory. Further studies could also explore cross-domain generalization, enabling models trained on one subject area to support learning in another through transfer learning. Evaluating long-term learning outcomes, such as retention and real-world skill application, represents another promising direction. Integrating explainability modules would also strengthen system transparency, allowing educators to understand and refine recommendation logic. These avenues highlight the ongoing potential for advancing intelligent tutoring systems that support richer, more adaptive, and more effective personalized learning experiences.

Declarations

Author Contributions

Conceptualization: A.A. and H.G.; Methodology: H.G.; Software: A.A.; Validation: A.A. and H.G.; Formal Analysis: A.A. and H.G.; Investigation: A.A.; Resources: H.G.; Data Curation: H.G.; Writing Original Draft Preparation: A.A. and H.G.; Writing Review and Editing: H.G. and A.A.; Visualization: A.A.; All authors have read and agreed to the published version of the manuscript.

Data Availability Statement

The data presented in this study are available on request from the corresponding author.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] G. A. Carpenter and S. Grossberg, "A Massively Parallel Architecture for a Self-Organizing Neural Pattern Recognition Machine C . Stability-Plasticity Dilemma : Multiple Interacting Memory Systems The properties of plasticity and stability are intimately related . An adequate," *Pattern Recognit.*, vol. 115, pp. 54–115, 1987.
- [2] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire, "Gambling in a rigged casino: the adversarial multi-armed bandit problem," *Annu. Symp. Found. Comput. Sci. - Proc.*, vol. 2002, no. August, pp. 322–331, 1995, doi: 10.1109/sfcs.1995.492488.
- [3] A. Joshi, P. Bhattacharyya, and S. Ahire, "Sentiment Resources: Lexicons and Datasets", pp. 85-106, 2017. doi: 10.1007/978-3-319-55394-8_5.
- [4] A. D. Buchdadi and A. S. M. Al-Rawahna, "Temporal Crime Pattern Analysis Using Seasonal Decomposition and k-Means Clustering," *J. Cyber Law*, vol. 1, no. 1, pp. 65–87, 2025, doi: 10.63913/jcl.v1i1.4.
- [5] A. Luaensutthi and T. Sangsawang, "Data Analytics of Online Lessons in Social Studies: Enhancing Teaching and Understanding Among Teachers and Students," *J. Appl. Data Sci.*, vol. 4, no. 3, pp. 200–212, 2023, doi: 10.47738/jads.v4i3.125.
- [6] D. Chakrabarti and R. Kumar, "Mortal Multi-Armed Bandits," pp. 1–8.
- [7] M. Ranga and H. Etzkowitz, "Triple Helix Systems: An Analytical Framework for Innovation Policy and Practice in the Knowledge Society," *Ind. High. Educ.*, vol. 27, no. 4, pp. 237–262, 2013, doi: 10.5367/ihe.2013.0165.
- [8] K. E. Harvey, M. A. Suizzo, and K. M. Jackson, "Predicting the Grades of Low-Income-Ethnic-Minority Students from Teacher-Student Discrepancies in Reported Motivation," *J. Exp. Educ.*, vol. 84, no. 3, pp. 510–528, 2016, doi: 10.1080/00220973.2015.1054332.
- [9] T. Hendrickx, B. Cule, P. Meysman, S. Naulaerts, K. Laukens, and B. Goethals, "Mining association rules in graphs based on frequent cohesive itemsets," *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, vol. 9078, no. 3, pp. 637–648, 2015, doi: 10.1007/978-3-319-18032-8_50.
- [10] L. Li, "Real time auxiliary data mining method for wireless communication mechanism optimization based on Internet of things system," *Comput. Commun.*, vol. 160, no. June, pp. 333–341, 2020, doi: 10.1016/j.comcom.2020.06.021.
- [11] L. Scaringella and F. Burtschell, "The challenges of radical innovation in Iran: Knowledge transfer and absorptive capacity highlights — Evidence from a joint venture in the construction sector," *Technol. Forecast. Soc. Change*, vol. 122, no. September, pp. 151–169, 2017, doi: 10.1016/j.techfore.2015.09.013.
- [12] B. M. Tayan, "Students and Teachers' Perceptions into the Viability of Mobile Technology Implementation to Support Language Learning for First Year Business Students in a Middle Eastern University," *Int. J. Educ. Lit. Stud.*, vol. 5, no. 2, p.

- 74, 2017, doi: 10.7575/aiac.ijels.v.5n.2p.74.
- [13] P. Pirolli and S. Card, "Information foraging," *Psychol. Rev.*, vol. 106, no. 4, pp. 643–675, 1999, doi: 10.1037/0033-295X.106.4.643.
- [14] R. Lerner, "Promoting positive youth development: Theoretical and empirical bases," *White Pap. Prep. Work. Sci. Adolesc. Heal. Dev. , Natl. Res. Council.*, p. 92, 2005, [Online]. Available: <http://ase.tufts.edu/iaryd/documents/pubPromotingPositive.pdf>
- [15] A. Holden and D. Fennell, "The routledge handbook of tourism and the environment", pp. 624, 2012. doi: 10.4324/9780203121108.
- [16] S. J. Lee, Z. Xu, T. Li, and Y. Yang, "A novel bagging C4.5 algorithm based on wrapper feature selection for supporting wise clinical decision making," *J. Biomed. Inform.*, vol. 78, no. February, pp. 144–155, 2018, doi: 10.1016/j.jbi.2017.11.005.
- [17] N. Trang, "Data mining for Education Sector, a proposed concept," *J. Appl. Data Sci.*, vol. 1, no. 1, pp. 11–19, 2020, doi: 10.47738/jads.v1i1.7.
- [18] R. H. Hama Aziz and N. Dimililer, "SentiXGboost: enhanced sentiment analysis in social media posts with ensemble XGBoost classifier," *J. Chinese Inst. Eng. Trans. Chinese Inst. Eng. A*, vol. 44, no. 6, pp. 562–572, 2021, doi: 10.1080/02533839.2021.1933598.
- [19] U. Hasanah and D. A. Muatiara, "Perbandingan metode cosine similarity dan jaccard similarity untuk penilaian otomatis jawaban pendek," *Semin. Nas. Sist. Inf. dan Tek. Inform.*, no. SENSITIF 2019, pp. 1255–1263, 2019, [Online]. Available: <https://ejurnal.diponegara.ac.id/index.php/sensitif/article/view/511>.
- [20] J. Jasperson, P. E. Carter, and R. W. Zmud, "A comprehensive conceptualization of post-adoptive behaviors associated with information technology enabled work systems," *MIS Q. Manag. Inf. Syst.*, vol. 29, no. 3, pp. 525–557, 2005, doi: 10.2307/25148694.
- [21] D. Sarpong, A. AbdRazak, E. Alexander, and D. Meissner, "Organizing practices of university, industry and government that facilitate (or impede) the transition to a hybrid triple helix model of innovation," *Technol. Forecast. Soc. Change*, vol. 123, no. October, pp. 142–152, 2017, doi: 10.1016/j.techfore.2015.11.032.
- [22] S. D. Lestari and E. B. Setiawan, "Sentiment Analysis Based on Aspects Using FastText Feature Expansion and NBSVM Classification Method," *J. Comput. Syst. Informatics*, vol. 3, no. 4, pp. 469–477, 2022, doi: 10.47065/josyc.v3i4.2202.
- [23] S. Qaiser and R. Ali, "Text Mining: Use of TF-IDF to Examine the Relevance of Words to Documents," *Int. J. Comput. Appl.*, vol. 181, no. 1, pp. 25–29, 2018, doi: 10.5120/ijca2018917395.
- [24] P. H. Andersen and L. E. Gadde, "Organizational interfaces and innovation: The challenge of integrating supplier knowledge in LEGO systems," *J. Purch. Supply Manag.*, vol. 25, no. 1, pp. 18–29, 2019, doi: 10.1016/j.pursup.2018.08.002.
- [25] T. Hariguna, H. T. Sukmana, and J. Il Kim, "Survey Opinion using Sentiment Analysis," *J. Appl. Data Sci.*, vol. 1, no. 1, pp. 35–40, 2020, doi: 10.47738/jads.v1i1.10.
- [26] X. Chen and Y. Yang, "Diffusion K-means clustering on manifolds: Provable exact recovery via semidefinite relaxations," *Appl. Comput. Harmon. Anal.*, vol. 1, no. May, pp. 1–45, 2020, doi: 10.1016/j.acha.2020.03.002.