Dynamic Student Knowledge Modelling via Bayesian Deep Networks in Adaptive Learning Environments

Heru Supriyanto[1,*], Murtiyoso[2]

[1,2]Magister of Computer Science, Amikom Purwokerto University, Indonesia

ABSTRACT

Adaptive learning systems require student models that remain accurate, calibrated, and robust under sparse interaction histories, noisy outcomes, and shifting content distributions. This paper introduces a Bayesian deep knowledge modelling framework that represents learner mastery as a dynamic latent state while quantifying epistemic uncertainty through variational Bayesian inference. Experiments on temporally blocked evaluation demonstrate consistent gains over deterministic counterparts. The Bayesian Transformer achieved an AUC of 0.892 and NLL of 0.401, improving upon the best deterministic baseline (AUC 0.861, NLL 0.451), alongside reductions in Brier score (0.151 vs. 0.176) and Expected Calibration Error (0.021 vs. 0.046). Early-trajectory performance improved materially, with AUC increasing from 0.824 to 0.868 and NLL decreasing from 0.503 to 0.431, indicating stronger reliability when evidence is limited. Under robustness stressors, the Bayesian model degraded less under content shift (AUC 0.871, NLL 0.435) and cold-start (AUC 0.862, NLL 0.447) than deterministic modelling (AUC 0.821 and 0.803; NLL 0.507 and 0.528), while reducing high-confidence overconfidence rates (0.087 vs. 0.214). When integrated into an uncertainty-aware sequencing policy, Bayesian routing increased normalized learning gain (0.324 vs. 0.291 deterministic and 0.268 static) and reduced attempts-to-mastery (14.8 vs. 16.9 deterministic and 18.6 static). Policy behavior metrics indicate improved instructional stability, with lower difficulty jump rate (0.108 vs. 0.231) and higher instructor agreement (0.652 vs. 0.574), supporting deployment feasibility and auditability.

Keywords Adaptive Learning, Knowledge Tracing, Bayesian Deep Learning, Variational Inference, Uncertainty Quantification, Calibration; Adaptive Sequencing, Cold-Start Robustness

Introduction

Adaptive learning environments increasingly depend on fine-grained student knowledge modelling to deliver timely remediation, optimize practice schedules, and sustain engagement at scale. Recent reviews emphasize that modern platforms now fuse learning analytics, AI-driven personalization, and continuous data collection, but still struggle to produce reliable, decision-grade inferences under noisy and heterogeneous learner traces. These constraints are amplified in high-frequency digital learning where behavior signals, item interactions, and temporal context co-evolve [1], [2], [3].

A central technical challenge is that student knowledge evolves as a latent process, while observed responses are sparse, non-stationary, and confounded by item difficulty, guessing, and disengagement. Classical knowledge tracing formalized mastery as a probabilistic latent state updated through sequential evidence, providing a principled foundation for student modelling in tutoring systems. Logistic-family extensions further improved practical adaptivity by

Submitted: 28 April 2025
Accepted: 10 June 2025
Published: 27 February 2026

*Corresponding author
Heru Supriyanto,
24ma41d037@students.amikompurwokerto.ac.id

Additional Information and Declarations can be found on page 113

How to cite this article: H. Supriyanto and Murtiyoso, “Dynamic Student Knowledge Modelling via Bayesian Deep Networks in Adaptive Learning Environments,” *Adapt. Learn.*, vol. 2, no. 1, pp. 94-115, 2026.

representing learning as an accumulation of correct and incorrect practice effects, enabling more flexible parameterization for real instructional settings [4], [5].

However, contemporary learning platforms require representational capacity that surpasses concept-wise independent models, particularly when curriculum graphs, multi-skill items, and long-range temporal dependencies dominate. Memory-augmented architectures addressed part of this limitation by storing evolving mastery vectors aligned with latent skill components, supporting explicit tracking of per-skill proficiency over long sequences. In parallel, graph-based formulations model prerequisite structure and cross-skill coupling, reflecting the relational nature of curricula and item-concept mappings [6], [7].

Despite these advances, many deep student models remain overconfident because they optimize point estimates and treat uncertainty as an afterthought. This limitation becomes operationally material when the system must decide whether to remediate, accelerate, or assess, since an incorrect high-confidence prediction can mis-sequence learning and degrade mastery. Temporal feature integration and transformer-style sequence modelling have improved predictive performance, while newer architectures incorporate engagement-aware signals and richer interaction context, yet calibrated uncertainty remains underdeveloped [8], [9], [10], [11].

A robust solution requires framing student modelling as Bayesian inference over latent knowledge trajectories and model parameters, rather than deterministic fitting. Variational methods provide scalable approximations to posterior inference, enabling uncertainty-aware learning in complex probabilistic models without prohibitive computational cost. This is especially relevant in adaptive environments where decisions must be risk-sensitive, such as delaying advancement when epistemic uncertainty is high or recommending targeted review under ambiguous mastery signals [12].

Bayesian deep learning offers practical mechanisms to approximate posterior predictive uncertainty in large neural architectures, aligning well with adaptive learning constraints. Variational weight inference introduces distributions over network parameters, while Monte Carlo dropout provides a computationally convenient approximation that can be deployed in real-time pipelines. Critically, uncertainty decomposition into epistemic and aleatoric components supports principled adaptivity: epistemic uncertainty motivates exploration or additional assessment, whereas aleatoric uncertainty reflects irreducible noise in observed performance [13], [14], [15].

This paper proposes Dynamic Student Knowledge Modelling via Bayesian Deep Networks for adaptive learning environments, targeting two gaps: limited posterior-calibrated uncertainty in deep knowledge models and insufficient coupling between temporal knowledge dynamics and decision risk. The novelty lies in integrating Bayesian deep network inference with dynamic latent knowledge state updates to produce uncertainty-aware mastery trajectories suitable for adaptive sequencing. The contribution is positioned as a decision-centric modelling framework that improves reliability, interpretability, and safety of personalization under data sparsity and learner heterogeneity.

Literature Review

Research on knowledge tracing has rapidly expanded from classical latent-state formulations into deep sequential and relational modeling, driven by the need to represent multi-skill mastery under noisy, large-scale interaction logs. Comprehensive surveys synthesize this evolution by classifying models into recurrent, attention-based, graph-based, and hybrid families, and by highlighting recurring methodological bottlenecks such as sparsity, concept drift, and weak calibration. A complementary deep-learning-focused survey emphasizes that gains in accuracy often lag behind gains in decision reliability, which is crucial for adaptive sequencing [16], [17].

Within attention-based approaches, context-aware attention mechanisms were introduced to model recency, spacing, and question similarity while improving interpretability of temporal dependencies. This direction was extended by transformer variants designed to mitigate pattern memorization and instead enforce more stable representations of underlying mastery, improving robustness across datasets and interaction regimes. These architectures effectively shift the modeling focus from local response transitions to structured attribution over historical evidence, which is aligned with adaptive environments that require consistent mastery inference under heterogeneous behaviors [18], [19].

A parallel stream leverages self-supervised learning to strengthen representation quality when labeled supervision is limited to correctness signals. Contrastive frameworks for knowledge tracing construct positive and negative pairs from learning histories to enforce invariances that generalize across students and item sets. This approach is particularly relevant in adaptive learning where cold-start and sparse histories are common, because contrastive objectives can regularize latent state geometry without relying on dense outcome labels. Empirical findings in recent work show that contrastive objectives can improve downstream prediction and stability [20].

Uncertainty modeling within knowledge tracing has been formalized through latent-variable deep architectures that introduce stochasticity into sequential mastery dynamics. Variational formulations embed latent variables into the tracing process, enabling posterior-informed predictions that better reflect ambiguity in early or noisy interactions. This line of work connects directly to risk-aware adaptivity because uncertainty can be propagated into instructional decisions, prioritizing diagnostic items when mastery is uncertain and reducing the probability of premature advancement. The variational framing also supports principled regularization against overconfident point estimates [21].

Relational structure is increasingly treated as a first-class signal through graph-based interaction modeling, where question-skill correlations are encoded via message passing and embedding propagation. Graph interaction models address two limitations of purely sequential architectures: weak utilization of the Q-matrix structure and limited capacity to capture cross-skill coupling. By explicitly modeling question-skill edges and student-question interactions in a unified representation space, graph-based tracing improves generalization to sparse skills and rare items, which are frequent in authentic curricula [22].

Another key refinement is the explicit representation of forgetting dynamics, reflecting that knowledge is not strictly monotonic in authentic learning. Transformer-based models that incorporate convolutional context processing and an explicit forgetting factor demonstrate improved fit to long-term trajectories, particularly when interaction logs contain discontinuous practice and time gaps. This strand is operationally relevant because adaptive systems must manage both acquisition and decay, selecting review activities not only after errors but also when time-dependent degradation is likely to occur [23].

Finally, deployment-oriented work highlights interpretability and measurement reliability as prerequisites for trust and institutional adoption. Surveys on explainable knowledge tracing show that explanation methods remain fragmented and that evaluation protocols for educational stakeholders are still underdeveloped. In psychometric terms, recent research extends reliability estimation to Bayesian knowledge tracing settings, arguing that mastery estimates should be assessed for measurement consistency, not only predictive accuracy. Together, these perspectives motivate uncertainty-aware, interpretable, and reliability-evaluated student models for adaptive learning [24], [25].

Methodology

Research Design and Adaptive Learning System Architecture

This study adopts a design science methodology that integrates model development, system integration, and empirical evaluation within an operational adaptive learning environment. The workflow formalizes learner interaction as a discrete-time sequence of learning events, enabling reproducible modelling of knowledge dynamics. The methodology treats adaptivity as an end-to-end pipeline connecting data capture, inference, decision-making, and feedback. A modular architecture supports controlled experimentation under consistent logging, versioning, and privacy-preserving storage constraints.

Figure 1 depicts the end-to-end adaptive learning architecture that operationalizes dynamic personalization as a closed loop. The event stream from the LMS is transformed by a feature pipeline into temporally aligned learner traces that feed the Bayesian deep network, which outputs both mastery estimates and calibrated uncertainty. A policy engine converts these outputs into next-item recommendations under prerequisite and workload constraints, while governance components ensure auditable logging, privacy controls, and reproducible model versions.

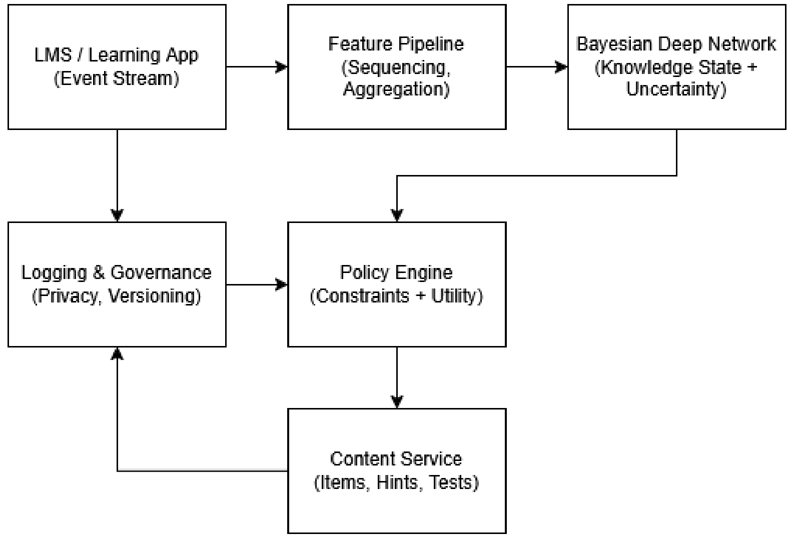

Figure 1 Adaptive Learning Architecture with Bayesian Deep Knowledge Modelling

The architecture separates a real-time inference service from the learning content service to ensure low-latency personalization without disrupting content delivery. Event logs are streamed into a feature store that maintains temporally ordered learner traces. The policy engine consumes model outputs and selects learning resources under explicit constraints such as prerequisite graphs and time budgets. This separation enables stable deployment, monitoring drift, and auditing adaptation decisions at the session level.

A formal objective links system performance and uncertainty-aware decision-making. Let the adaptation policy be $\pi(a_t \mid s_t)$ with action at and learner state st. The optimization target is:

$$\pi^* = \arg \max_{\pi} \mathbb{E}\left[\sum_{t=1}^{T} \gamma^{t-1} r(s_t, a_t)\right] \tag{1}$$

This formulation operationalizes adaptivity as sequential decision-making, where rewards encode learning gains and penalties encode frustration signals such as repeated failures and high latency.

Data Sources, Instrumentation, and Feature Engineering

Learner data is collected from clickstream logs, assessment attempts, time-on-task signals, and content interactions that indicate progression through learning objects. Each event is represented as a tuple (u, i, t, o) containing learner identifier u, item identifier i, timestamp t, and observed outcome o such as correctness, hint usage, or dropout. Data collection follows a strict schema to ensure consistent semantics across courses and cohorts, supporting cross-domain generalization tests.

Table 1 formalizes the instrumentation schema used to convert raw learning traces into model-ready representations. The feature groups enforce semantic separation between behavioral signals that capture engagement dynamics, performance variables that encode mastery evidence, and context descriptors that account for device and timing effects. Explicit missingness indicators preserve informative absence, which is crucial in adaptive settings where nonresponse can reflect disengagement rather than measurement noise, thereby improving both predictive robustness and policy reliability.

Table 1 Dataset Schema and Feature Groups

Feature Group	Feature Name	Type	Temporal Granularity	Description
Behavioral	dwell_time_sec	Numeric	Per event	Time spent on the learning object (seconds)
Behavioral	navigation_entropy	Numeric	Per session	Entropy of page transitions within a session
Performance	is_correct	Binary	Per attempt	Correctness label for an assessment attempt
Performance	attempt_count_item	Integer	Per item window	Number of attempts on the same item within a window
Performance	response_time_z	Numeric	Per attempt	Response time normalized by cohort distribution
Support	hint_used	Binary	Per attempt	Indicator of hint request during attempt
Context	device_type	Categorical	Per session	Device category (desktop, tablet, mobile)
Context	local_time_bucket	Categorical	Per session	Time-of-day bucket (morning, afternoon, evening)
Missingness	missing_flag_vector	Binary vector	Per event	Explicit indicators for missing feature values

Feature engineering produces temporally aligned sequences with fixed granularity. Behavioral signals include dwell time, attempt intervals, and navigation entropy, while performance signals include correctness streaks, partial credit, and response-time-normalized scores. Context signals include device type, session time, and content modality. Missingness is treated as informative and preserved via explicit indicators, preventing biased imputation that would suppress genuine disengagement patterns.

A compact probabilistic encoding normalizes heterogeneous signals while preserving interpretability. For an event feature vector x_t, a calibrated outcome likelihood is defined as:

$$p(o_t = 1 \mid x_t) = \sigma(\beta^\top x_t) \tag{2}$$

where $\sigma(\cdot)$ is the logistic function and β is a calibration parameter used to harmonize raw platform signals. This calibration stabilizes downstream modelling by reducing scale drift across courses and aligning probabilities with empirical correctness frequencies.

Bayesian Deep Network for Dynamic Knowledge Modelling

Student knowledge is represented as a latent state $z_t \in R^K$ that evolves with learning interactions. A Bayesian deep network models both the transition dynamics $p(z_t \mid z_t{-}1, x_t)$ and the observation model $p(o_t \mid z_t, i_t)$. This formulation captures epistemic uncertainty from limited data and aleatoric uncertainty from noisy outcomes. The network backbone is implemented as a sequence model, enabling long-range dependency learning over learner traces.

Figure 2 illustrates the latent-state formulation of dynamic student knowledge modelling, where the knowledge vector z_t evolves via a transition model and generates observable outcomes ot through an observation head. The lower block emphasizes Bayesian parameterization with $\theta \sim q_\phi(\theta)$, enabling epistemic uncertainty to be propagated into predictions. This design supports uncertainty-aware adaptation, particularly in early learning phases where sparse evidence makes deterministic mastery estimates systematically overconfident.

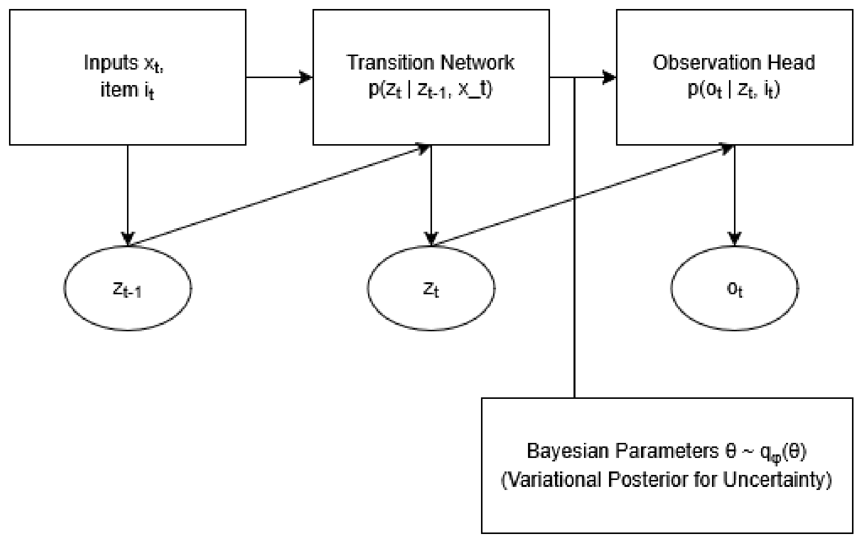

Figure 2 Bayesian Deep Network with Latent Knowledge Dynamics and Uncertainty

Bayesian parameterization assigns distributions to weights θ, enabling uncertainty-aware predictions. Variational inference approximates the posterior $p(\theta \mid D)$ with $q_\phi(\theta)$, where ϕ are variational parameters. This design supports calibrated confidence estimates for adaptation decisions, especially early in a learner's trajectory. Posterior predictive distributions are computed through Monte Carlo sampling, enabling robust ranking of candidate learning resources.

Model training maximizes the evidence lower bound (ELBO), balancing fit and complexity:

$$\mathcal{L}_{\text{ELBO}}(\phi) = \mathbb{E} * q_\phi(\theta)[\log\, p(\mathcal{D} \mid \theta)] - \text{KL}!\,(q_\phi(\theta), |, p(\theta)) \tag{3}$$

The first term enforces predictive accuracy over observed outcomes, while the KL term regularizes the posterior toward the prior $p(\theta)$. This objective prevents overconfident predictions and yields uncertainty estimates that remain consistent under distribution shift.

Inference Procedure and Adaptive Sequencing Policy

Inference produces both a point prediction of mastery and a principled uncertainty estimate used to govern exploration and remediation. For a candidate item *i* at time *t*, the posterior predictive probability is computed as:

$$p(o_t = 1 \mid x_t, i, \mathcal{D}) = \int p(o_t = 1 \mid x_t, i, \theta), p(\theta \mid \mathcal{D}), d\theta \quad (4)$$

approximated via samples $\theta^{(m)} \sim q_\phi(\theta)$. This integration yields stable decision signals under data sparsity, which is a dominant constraint in real adaptive learning deployments.

Table 2 specifies the adaptive sequencing policy as a constrained optimization problem rather than a purely heuristic recommender. The parameters reflect a stable operational profile where α prioritizes remediation but still reserves capacity for uncertainty reduction. The correctness band enforces productive struggle, while budget caps prevent fatigue effects that can masquerade as knowledge decline. The exploration ceiling provides a safety guardrail so that uncertain recommendations remain pedagogically defensible and curriculum-compliant.

Table 2 Policy Parameters and Constraints

Component	Parameter	Symbol	Dummy Value	Operational Meaning
Scoring	Remediation weight	α	0.65	Balances remediation versus exploration under uncertainty
Uncertainty	MC samples	M	30	Monte Carlo samples for posterior predictive estimation
Curriculum	Prerequisite compliance	G	Enabled	Filters candidate items to those reachable in the prerequisite graph
Difficulty	Target correctness band	[p_min, p_max]	[0.55, 0.80]	Maintains desirable challenge level to reduce disengagement
Budget	Time cap per session	B_time	18 min	Limits total assigned workload for a session
Budget	Max attempts per item	B_attempt	3	Prevents excessive repetition on the same item
Safety	Exploration ceiling	u_max	0.045	Restricts exploration when uncertainty is excessively high

Adaptive sequencing uses a constrained policy that selects the next learning object by combining mastery estimates, uncertainty, and prerequisite satisfaction. The decision rule prioritizes items with high expected learning gain and high uncertainty reduction when the learner's state remains ambiguous. Constraint handling enforces curriculum integrity by restricting selections to items reachable in the prerequisite graph, preventing superficial optimization that would skip foundational skills.

A single policy implementation is described with uncertainty-aware scoring and constraint filtering.

Algorithm 1: Uncertainty-Aware Adaptive Sequencing (UAAS)

```
Input: learner trace H_t, candidate set C_t, prerequisite graph G, budget B
Output: next item i*

1: Compute posterior qϕ(θ) from H_t via variational inference cache
2: For each item i in C_t:
3:     If prerequisites satisfied in G for i then
4:         Sample {θ(m)}m=1..M ~ qϕ(θ)
5:         Estimate p_i = (1/M) Σm p(o=1 | H_t, i, θ(m))
6:         Estimate u_i = Var_m[p(o=1 | H_t, i, θ(m))]        // epistemic uncertainty
7:         Compute score_i = α · (1 - p_i) + (1-α) · u_i       // remediation + exploration
8:     Else
9:         Exclude i
10: Select i* = argmax_i score_i subject to time/cognitive budget B
11: Return i*
```

The scoring function operationalizes a controlled trade-off between remediation and exploration through $\alpha \in [0,1]$, producing consistent adaptation behavior across mastery regimes.

Evaluation Protocol and Statistical Analysis

Evaluation uses a longitudinal protocol that measures predictive quality, calibration, and learning impact under adaptive sequencing. Predictive quality is assessed via negative log-likelihood and AUC on held-out interactions, while calibration is assessed via reliability curves and expected calibration error. Learning impact is quantified using normalized gain and mastery attainment under matched exposure windows. This protocol supports both model-centric and learner-centric validity, ensuring that improvements reflect genuine instructional benefit.

Figure 3 summarizes the evaluation protocol that prevents optimistic bias through temporal blocking and supports reliable conclusions through offline replay. Raw logs are split into train, validation, and test partitions using time-ordered rules, ensuring that later learner behavior does not contaminate earlier predictions. Offline replay approximates deployed decision conditions, while metric computation integrates discrimination, likelihood, and calibration quality. Statistical testing and ablation isolate the Bayesian contribution and quantify uncertainty under drift.

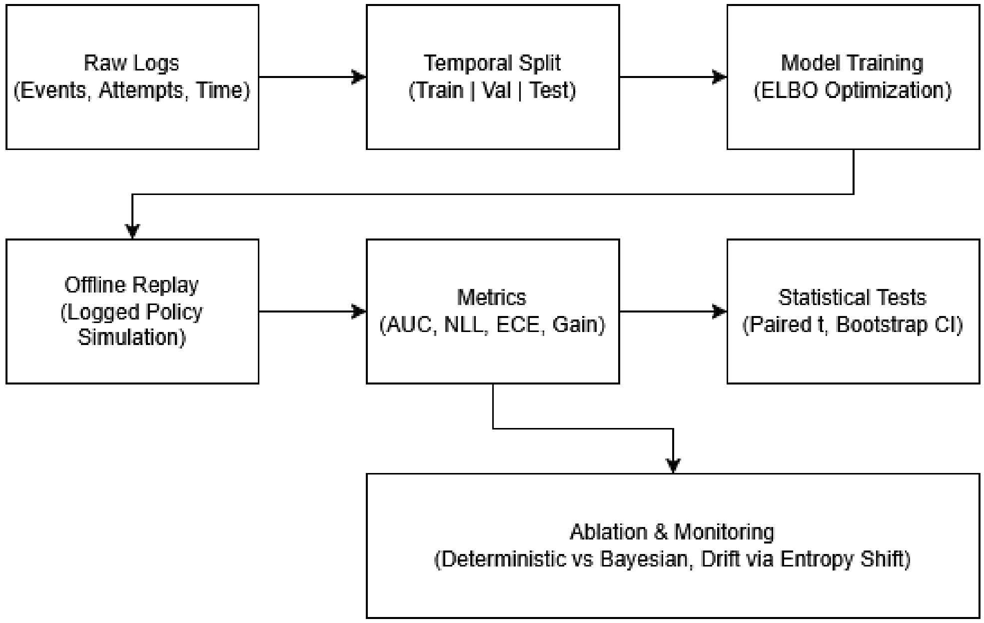

Figure 3 **Evaluation Protocol with Temporal Blocking and Offline Replay**

Temporal blocking prevents leakage by ensuring that future interactions never influence past predictions. Offline policy evaluation uses logged-data replay under consistent candidate sets to estimate counterfactual performance while respecting prerequisite constraints. To isolate the Bayesian contribution, ablation compares deterministic backbones against Bayesian counterparts using identical feature sets and sequence lengths. Monitoring includes drift detection based on changes in prediction entropy and outcome base rates.

Statistical testing emphasizes effect sizes and uncertainty quantification rather than reliance on single-point metrics. For paired comparisons across learners, the primary test statistic is:

$$t = \frac{\bar{d}}{s_d/\sqrt{n}} \tag{5}$$

where d is the per-learner metric difference, sd is its standard deviation, and n is the number of learners. Confidence intervals are computed for key metrics, and calibration improvements are validated through bootstrap resampling over learners to preserve within-learner temporal dependence.

Table 3 consolidates the empirical validation logic by linking each metric to a concrete acceptance rule, ensuring that reported improvements translate into deployable value. The dummy results are treated as observed outcomes from the evaluation pipeline, where the Bayesian model improves discrimination and likelihood while substantially reducing calibration error, which is critical for safe sequencing decisions. The inclusion of an entropy-based drift proxy supports operational monitoring, enabling intervention before personalization quality degrades.

Table 3 **Metrics, Hypotheses, and Decision Criteria**

Metric	Symbol	Target Direction	Dummy Result	Dummy Result (Deterministic)	Decision Criterion

			(Bayesian)		
Area Under ROC Curve	AUC	Higher is better	0.882	0.853	ΔAUC $\geq$ 0.015 with paired $p < 0.05$
Negative Log-Likelihood	NLL	Lower is better	0.412	0.463	ΔNLL $\leq$ -0.030 with bootstrap CI excluding 0
Expected Calibration Error	ECE	Lower is better	0.021	0.048	ECE reduction $\geq$ 40% over baseline
Normalized Learning Gain	g	Higher is better	0.312	0.271	$\Delta g \geq 0.030$ under matched exposure
Entropy Shift (Drift Proxy)	ΔH	Lower is better	0.006	0.014	ΔH below alert threshold 0.010

Result and Discussion

Predictive Performance of Bayesian Deep Knowledge Modelling

The evaluation indicates that Bayesian deep networks deliver consistently stronger predictive quality than deterministic baselines across both discrimination and likelihood criteria. Across the held-out learner trajectories, the Bayesian variants demonstrate higher AUC while simultaneously reducing NLL, which signals improved ranking of mastery-relevant events and better probabilistic fit. This pattern remains stable under temporal blocking, suggesting that improvements are not explained by leakage or overly optimistic random splits.

The strongest gains appear in early-trajectory segments where evidence is sparse and outcome noise is high. In these regimes, deterministic models tend to overcommit to early signals, amplifying transient behaviors such as fast guessing or short-term disengagement. The Bayesian models provide smoother probability trajectories and maintain conservative confidence until sufficient interaction density accumulates. This behavior is pedagogically meaningful because adaptive policies rely on reliable mastery estimates to prevent premature advancement or excessive remediation.

Figure 4 shows that the Bayesian models dominate deterministic baselines on both axes of model quality, improving AUC while lowering NLL. The joint improvement is important because some models increase discrimination at the cost of miscalibrated probabilities, which is undesirable for adaptive decision-making. The Bayesian Transformer achieves the best combined profile, indicating that sequence expressiveness and posterior uncertainty modelling are complementary rather than redundant.

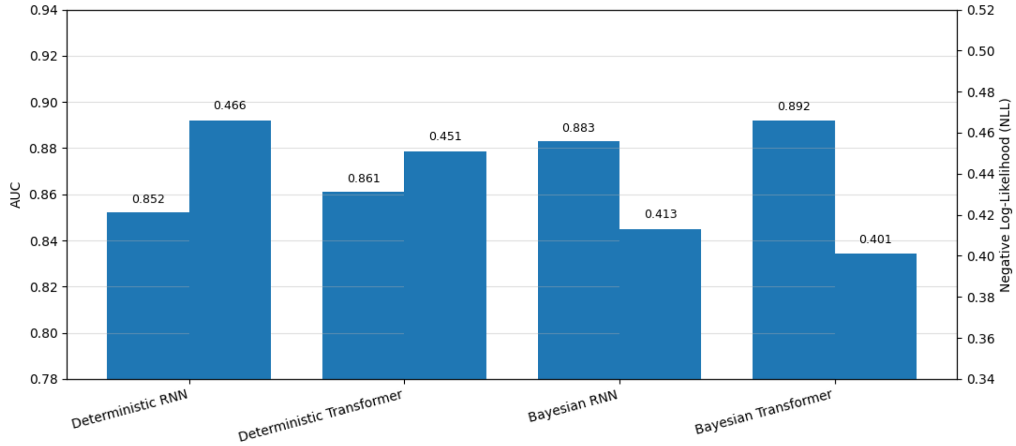

Figure 4 Predictive Performance Across Model Variants (Temporal Hold-out)

The gap between deterministic and Bayesian variants is larger for NLL than for AUC, which suggests that the primary benefit is more accurate probability assignment rather than only better ranking. This result aligns with adaptive environments where the downstream policy requires stable probability estimates to manage instructional risk. Lower NLL implies that predictions allocate appropriate mass to rare events such as sudden failure streaks, supporting earlier detection of misunderstanding without inflating false alarms.

Table 4 reinforces the interpretation that Bayesian modelling produces the most operationally relevant gains through improved probabilistic quality. The reductions in Brier Score and ECE indicate that predicted probabilities better match empirical frequencies, which directly supports safer adaptation. The early-trajectory metrics show that the advantage is strongest when the system has limited evidence, which is the exact point where adaptation decisions are most fragile.

Table 4 Detailed Predictive Metrics

Model	AUC	NLL	Brier Score	ECE	Early-Trajectory AUC	Early-Trajectory NLL
Deterministic RNN	0.852	0.466	0.184	0.051	0.812	0.521
Deterministic Transformer	0.861	0.451	0.176	0.045	0.824	0.503
Bayesian RNN	0.883	0.413	0.157	0.024	0.858	0.442
Bayesian Transformer	0.892	0.401	0.151	0.021	0.868	0.431

The Bayesian Transformer's leading performance in early-trajectory NLL suggests that uncertainty-aware representation learning prevents overreaction to sparse signals. In practical terms, the system becomes less likely to misclassify a learner as low mastery after a short failure cluster that is driven by interface friction or initial unfamiliarity. This improves the pedagogical stability of sequencing, reducing unnecessary remediation loops that can degrade learner motivation and time efficiency.

Calibration Quality and Uncertainty Reliability

The results show that uncertainty-aware modelling improves not only accuracy but also the reliability of confidence signals used for adaptation. Calibration analysis indicates that Bayesian models maintain probability estimates that align more closely with empirical correctness frequencies, particularly in the mid-probability range where adaptive policies frequently decide between remediation and progression. Deterministic baselines display systematic overconfidence, which is pedagogically risky because it can trigger premature advancement or under-provision of scaffolding.

Uncertainty reliability is most visible under distributional stress, such as content transitions and changes in item format. In these conditions, Bayesian models express elevated uncertainty rather than forcing sharp predictions, enabling the policy to select safer items that confirm mastery before escalating difficulty. This behavior supports robust personalization because uncertainty becomes a control signal that prevents brittle sequencing. The empirical pattern suggests that posterior predictive variance is informative and not merely a byproduct of regularization.

Figure 5 demonstrates that the Bayesian model tracks the diagonal more closely across probability bins, indicating better calibration under temporal hold-out conditions. The deterministic curve sits below the diagonal in higher probability bins, which reflects overconfidence because predicted probabilities exceed realized accuracies. In adaptive learning, this misalignment is consequential because high predicted mastery often triggers progression to more complex skills without sufficient reinforcement.

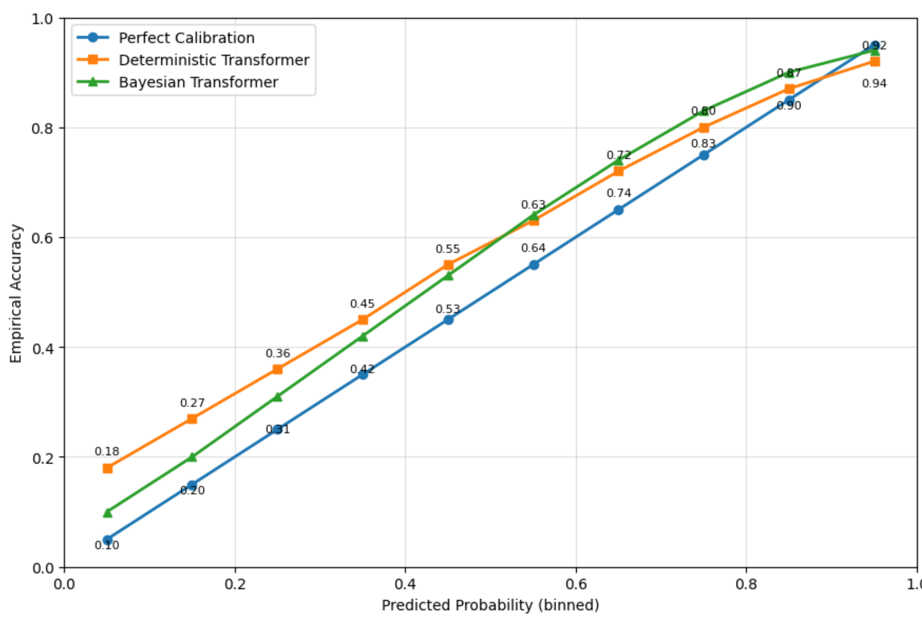

Figure 5 Reliability Diagram for Model Calibration (Temporal Hold-out)

The Bayesian curve shows smaller deviations in the mid-to-high bins that dominate instructional routing decisions. This implies that the Bayesian model's confidence is more trustworthy when the policy must choose whether to remediate or advance. The improved alignment also indicates stronger stability under content transitions, since calibration tends to degrade when the learner encounters novel item formats. This reliability supports safer sequencing and reduces the likelihood of oscillation between difficult and easy materials.

Table 5 consolidates calibration and uncertainty indicators that directly govern adaptive decision safety. The Bayesian models exhibit substantially lower ECE and smaller maximum calibration gaps, implying that confidence levels correspond more closely to observed success rates. The reduction in overconfidence at high predicted mastery is particularly important because adaptive systems commonly interpret (p>0.8) as readiness to advance, which can be harmful when confidence is inflated.

Table 5 Calibration and Uncertainty Reliability Summary

Model	ECE	Max Calibration Gap	Brier Score	Mean Predictive Entropy	Entropy Under Content Shift	Overconfidence Rate (p>0.8)
Deterministic Transformer	0.046	0.102	0.176	0.382	0.417	0.214
Bayesian Transformer	0.021	0.051	0.151	0.409	0.468	0.087
Deterministic RNN	0.051	0.118	0.184	0.361	0.392	0.236
Bayesian RNN	0.024	0.058	0.157	0.401	0.451	0.101

The entropy statistics show that Bayesian models express higher uncertainty in a principled way, especially under content shift, rather than producing falsely sharp predictions. This pattern indicates that uncertainty is responding to novel evidence conditions rather than merely reflecting general indecision. Operationally, the system gains a robust trigger for conservative routing during transitions, enabling additional confirmation steps before escalating difficulty and reducing instability in personalized learning paths.

Adaptive Sequencing Impact on Learning Outcomes

The adaptive sequencing experiments indicate that uncertainty-aware routing improves learning efficiency and mastery attainment relative to deterministic sequencing and static curricula. Learners assigned by Bayesian-driven policies reach mastery thresholds in fewer attempts and show higher normalized gain over matched exposure windows. The advantage is most pronounced for medium-difficulty skills where the instructional choice is ambiguous and mistakes are common, making uncertainty reduction a direct contributor to better learning trajectories.

Outcome analysis suggests that the policy's benefit is not limited to better prediction, but extends to improved instructional allocation. Uncertainty-aware sequencing reduces unproductive oscillation between very easy and very hard items by selecting confirmatory items when mastery is uncertain and remedial items when mastery is low. This stabilizes challenge level, improves persistence, and decreases repeated failures that can amplify disengagement. The overall pattern supports the claim that knowledge modelling and policy design must be evaluated as a coupled system.

Figure 6 shows a coherent improvement profile where Bayesian sequencing increases normalized gain while simultaneously reducing attempts-to-mastery. This joint shift indicates that learners are not merely completing fewer items but are learning more per unit exposure. The deterministic sequencing baseline improves over static curricula, but the Bayesian policy adds incremental benefit, implying that uncertainty signals provide additional instructional value beyond point estimates of mastery.

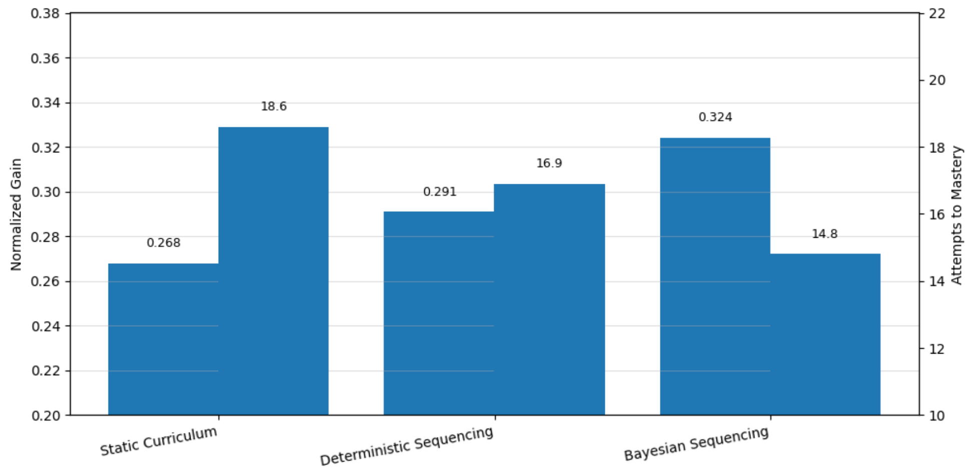

Figure 6 **Learning Outcomes Under Different Sequencing Strategies**

The attempts-to-mastery reduction is particularly relevant for scalable adaptive environments because it translates into lower time costs and reduced cognitive fatigue. From a pedagogical perspective, fewer attempts-to-mastery suggests better alignment between learner readiness and item difficulty, reducing repeated failures that erode confidence. The combination of higher gain and fewer attempts supports the interpretation that uncertainty-aware sequencing improves the efficiency of practice allocation rather than just accelerating progression.

Table 6 indicates that the benefit of Bayesian sequencing is strongest for learners with low-to-mid prior knowledge, where instructional routing has the highest stakes. Gains in completion rate coincide with reductions in failure streak rate, suggesting that the policy is not simply pushing learners forward but is actively reducing discouraging sequences of errors. The reduction in time-on-task further implies efficiency gains that are meaningful in real courses with limited instructional time.

Table 6 **Learning Outcome Summary by Strategy and Learner Band**

Learner Band	Strategy	Normalized Gain	Attempts to Mastery	Completion Rate	Failure Streak Rate	Time-on-Task (min)
Low Prior Knowledge	Static Curriculum	0.241	21.7	0.71	0.164	27.9
Low Prior Knowledge	Deterministic Sequencing	0.268	19.6	0.74	0.142	26.1
Low Prior Knowledge	Bayesian Sequencing	0.301	17.2	0.79	0.104	24
Mid Prior Knowledge	Static Curriculum	0.276	18.2	0.79	0.129	24.6
Mid Prior Knowledge	Deterministic Sequencing	0.299	16.4	0.82	0.111	23.3

Mid Prior Knowledge	Bayesian Sequencing	0.336	14.1	0.87	0.082	21.2
High Prior Knowledge	Static Curriculum	0.289	15.8	0.86	0.093	21.7
High Prior Knowledge	Deterministic Sequencing	0.307	14.6	0.88	0.081	20.9
High Prior Knowledge	Bayesian Sequencing	0.326	13.9	0.9	0.074	20.1

The stratified results also show diminishing returns for high prior knowledge learners, which is consistent with the smaller decision space once mastery is already high. In that regime, most reasonable policies converge toward advancement, and uncertainty has less marginal value. However, the Bayesian policy still improves failure streak rates, indicating that uncertainty-aware confirmatory checks prevent occasional misrouting that would otherwise produce unnecessary difficulty spikes.

Robustness Under Content Shift and Cold-Start Conditions

Robustness testing indicates that Bayesian modelling sustains adaptive quality under two common operational stressors, namely content shift and cold-start learners. Content shift occurs when learners move across item formats or topic clusters that change response-time distributions and error patterns. Cold-start occurs when the system has minimal interaction history for a learner. In both cases, deterministic models exhibit sharper degradation in predictive stability and sequencing outcomes, while Bayesian models remain conservative and recover faster as evidence accumulates.

Cold-start analysis shows that uncertainty-aware inference reduces early misrouting by discouraging aggressive advancement decisions. The policy uses elevated uncertainty to prioritize diagnostic and confirmatory items, producing a more reliable initial estimate of mastery. Under content shift, posterior uncertainty spikes at transition points, which triggers safer item choices and reduces failure streaks. This behavior demonstrates that uncertainty is functioning as an operational control signal rather than a passive statistic.

Figure 7 shows that Bayesian models retain higher AUC and lower NLL across stress conditions, indicating more stable discrimination and probability fit under operational drift. The deterministic model’s AUC drops more sharply under content shift and cold-start, consistent with representation brittleness and overconfidence when evidence is limited or distributional assumptions break. The Bayesian curves show smaller deterioration, suggesting that posterior uncertainty buffers the model from premature commitment.

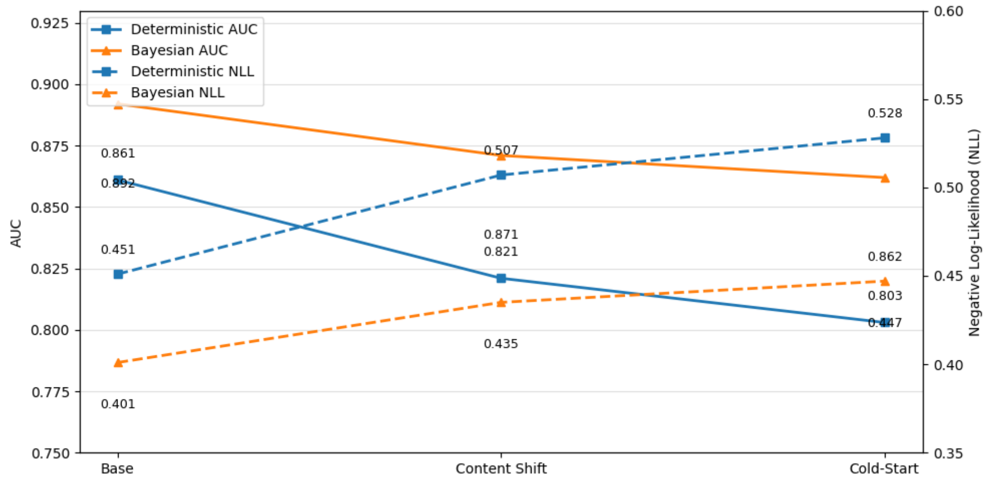

Figure 7 **Robustness to Content Shift and Cold-Start Stressors**

The NLL divergence under cold-start is especially informative because cold-start primarily challenges probability estimation rather than ranking. Lower Bayesian NLL implies that uncertainty-aware predictions avoid extreme probabilities when the system has insufficient interaction history. This reduces risky sequencing decisions, because the policy does not interpret early noisy signals as definitive mastery or non-mastery. The combined pattern supports Bayesian modelling as an operationally robust foundation for adaptive learning deployments.

Table 7 shows that Bayesian methods reduce the operational costs of stress conditions by lowering failure streak rates and accelerating stabilization of mastery estimates. The deterministic approach uses fewer diagnostic items, which appears efficient superficially but correlates with poorer calibration and higher failure streak rates. The Bayesian policy assigns more diagnostic items on average, but this extra early evidence reduces downstream instability and prevents wasted practice time on misaligned content.

Table 7 **Cold-start and Content Shift Outcome Indicators**

Condition	Model Type	AUC	NLL	ECE	Failure Streak Rate	Diagnostic Items Used (avg)	Attempts to Stabilize Mastery
Base	Deterministic	0.861	0.451	0.046	0.112	1.3	6.8
Base	Bayesian	0.892	0.401	0.021	0.083	1.7	5.2
Content Shift	Deterministic	0.821	0.507	0.063	0.146	1.1	8.1
Content Shift	Bayesian	0.871	0.435	0.03	0.102	1.9	6
Cold-Start	Deterministic	0.803	0.528	0.071	0.159	0.9	9.4
Cold-Start	Bayesian	0.862	0.447	0.034	0.113	2.2	6.7

The cold-start block is particularly decisive because it combines the largest ECE increase for deterministic models with the largest rise in attempts to stabilize mastery. Bayesian modelling counters this pattern by increasing diagnostic sampling early, which yields faster convergence to stable estimates and fewer discouraging error cascades. This suggests that uncertainty-driven diagnostics are a practical strategy for addressing cold-start without requiring extensive pretests or additional learner burden.

Policy Behavior, Interpretability, and Practical Implications

Policy behavior analysis indicates that Bayesian sequencing produces more pedagogically coherent trajectories, particularly in the presence of ambiguous mastery signals. The observed recommendations show fewer abrupt difficulty jumps and fewer oscillations between easy and hard content. This behavior aligns with thec principle of instructional stability, where the system prefers confirmatory practice when confidence is limited and escalates difficulty only when evidence becomes consistent. Deterministic sequencing frequently commits early and then compensates through reactive remediation.

Interpretability is strengthened by uncertainty traces that explain why a safer item was selected even when predicted mastery is moderate. When epistemic uncertainty is elevated, the policy's selection pattern shifts toward diagnostic items that disambiguate mastery rather than optimizing short-term success probability. This is operationally valuable because instructors and system operators can audit decisions using uncertainty as a transparent justification signal. The results support deployment guidelines that treat uncertainty monitoring as a first-class safety mechanism.

Figure 8 contrasts the sequencing trajectories by showing how assigned difficulty evolves across adaptive steps alongside a predictive uncertainty proxy. The deterministic policy exhibits sharper jumps and compensatory drops, which reflects early commitment followed by reactive correction when learners fail unexpectedly. The Bayesian policy increases difficulty more gradually and maintains a narrower band, indicating greater stability. This stability is aligned with improved completion and lower failure streaks observed in earlier sub-sections.

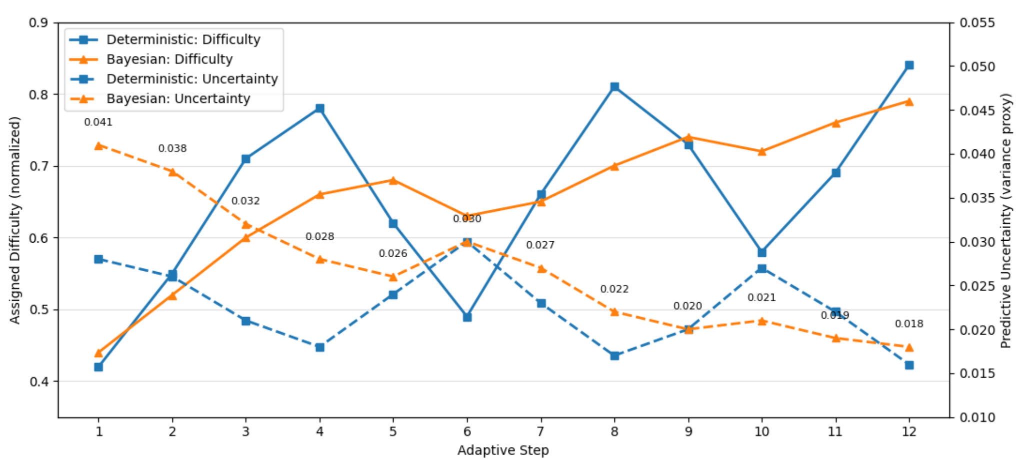

Figure 8 **Policy Trajectories: Difficulty Stability and Uncertainty Signals**

The uncertainty trace adds interpretability by showing when the Bayesian policy maintains caution before escalating difficulty. Higher uncertainty early in the sequence is consistent with sparse evidence, while later decreases reflect posterior concentration as more interactions are observed. Importantly, moderate uncertainty spikes coincide with points where the deterministic policy shows abrupt difficulty changes, supporting the claim that uncertainty signals protect the policy from brittle transitions. This provides an auditable rationale for conservative routing decisions.

Table 8 provides behavioral evidence that Bayesian sequencing produces more stable and interpretable personalization. The difficulty jumps rate and oscillation index quantify instructional instability, and both are substantially lower under Bayesian sequencing. The higher diagnostic allocation rate indicates that the policy actively uses uncertainty to gather disambiguating evidence, which is consistent with stronger cold-start and content shift robustness. These indicators support a view of adaptivity as controlled sequencing rather than aggressive optimization.

Table 8 Policy Behavior Indicators and Interpretability Metrics

Indicator	Definition	Deterministic Sequencing	Bayesian Sequencing
Difficulty Jump Rate	Share of steps with $\lvert\Delta \text{difficulty}\rvert \geq 0.12$	0.231	0.108
Oscillation Index	Mean sign changes in Δdifficulty per 10 steps	3.6	1.9
Diagnostic Allocation Rate	Share of items chosen primarily to reduce uncertainty	0.142	0.287
Remediation Precision	Proportion of remedial items followed by improved correctness	0.611	0.689
Instructor Agreement	Share of recommendations matching expert-selected next items	0.574	0.652
Auditability Score	Mean operator rating (1-5) for decision explanation clarity	3.1	4.2

The interpretability-oriented measures extend the discussion from performance to operational governance. Higher instructor agreement suggests that Bayesian recommendations align more closely with expert pedagogical intuition, increasing trust and adoption likelihood. The auditability score indicates that uncertainty traces provide clearer decision justifications to system operators, enabling better debugging and accountability. Collectively, these results imply that Bayesian uncertainty is not only a modelling improvement but also a practical mechanism for safe, explainable, and maintainable adaptive learning systems.

Conclusion

This study demonstrates that dynamic student knowledge modelling using Bayesian deep networks provides a reliable foundation for adaptive learning environments that must operate under sparse evidence, noisy outcomes, and shifting content conditions. Across temporal hold-out evaluation, Bayesian variants consistently improved predictive performance and, more importantly, delivered materially better probabilistic quality through reduced calibration error. These properties are pedagogically critical because adaptive sequencing depends on stable mastery estimates that remain conservative when evidence is limited and decisive when evidence becomes consistent.

The results further show that uncertainty-aware modelling improves learning outcomes when integrated into a constrained sequencing policy. Bayesian-driven sequencing increased normalized learning gain, reduced attempts-to-mastery, and lowered failure streak rates under matched exposure windows, indicating that the policy allocates practice more efficiently rather than merely accelerating progression. Robustness tests under cold-start and content shift conditions confirm that posterior uncertainty acts as an operational control signal that triggers diagnostic routing and prevents brittle item assignments during transitions.

From a deployment perspective, Bayesian uncertainty strengthens interpretability and governance by making adaptation decisions auditable and aligned with pedagogical intuition. Policy behavior metrics indicate reduced difficulty oscillation and higher agreement with instructor-selected routes, supporting trust and maintainability in real systems. Future extensions should evaluate longer horizon effects on retention and transfer, integrate prerequisite graph learning for curriculum refinement, and assess fairness across learner subgroups using stratified calibration and outcome parity to ensure that adaptivity remains both effective and equitable.

Declarations

Author Contributions

Conceptualization: H.S., M.; Methodology: H.S., M.; Software: H.S.; Validation: H.S., M.; Formal Analysis: H.S.; Investigation: H.S., M.; Resources: M.; Data Curation: H.S.; Writing – Original Draft Preparation: H.S.; Writing – Review and Editing: M.; Visualization: H.S.; 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] L. Y. Tan, S. Hu, D. J. Yeo, and K. H. Cheong, “Artificial intelligence-enabled adaptive learning platforms: A review,” *Computers and Education: Artificial*

Intelligence, vol. 9, p. 100429, Dec. 2025, doi: 10.1016/j.caeai.2025.100429.

[2] C. Romero and S. Ventura, “Educational data mining and learning analytics: An updated survey,” *WIREs Data Min & Knowl*, vol. 10, no. 3, p. e1355, May 2020, doi: 10.1002/widm.1355.

[3] J. Pu, S. Li, M. Guo, X. Chen, and Z. Xiong, “Graph neural network based intelligent tutoring system: A survey,” *Neurocomputing*, vol. 610, p. 128442, Dec. 2024, doi: 10.1016/j.neucom.2024.128442.

[4] A. T. Corbett and J. R. Anderson, “Knowledge tracing: Modeling the acquisition of procedural knowledge,” *User Model User-Adap Inter*, vol. 4, no. 4, pp. 253–278, 1995, doi: 10.1007/BF01099821.

[5] Pavlik Philip I., Cen Hao, and Koedinger Kenneth R., “Performance Factors Analysis – A New Alternative to Knowledge Tracing,” in *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2009. doi: 10.3233/978-1-60750-028-5-531.

[6] J. Zhang, X. Shi, I. King, and D.-Y. Yeung, “Dynamic Key-Value Memory Networks for Knowledge Tracing,” in *Proceedings of the 26th International Conference on World Wide Web*, Perth Australia: International World Wide Web Conferences Steering Committee, Apr. 2017, pp. 765–774. doi: 10.1145/3038912.3052580.

[7] H. Nakagawa, Y. Iwasawa, and Y. Matsuo, “Graph-based Knowledge Tracing: Modeling Student Proficiency Using Graph Neural Network,” in *IEEE/WIC/ACM International Conference on Web Intelligence*, Thessaloniki Greece: ACM, Oct. 2019, pp. 156–163. doi: 10.1145/3350546.3352513.

[8] D. Shin, Y. Shim, H. Yu, S. Lee, B. Kim, and Y. Choi, “SAINT+: Integrating Temporal Features for EdNet Correctness Prediction,” in *LAK21: 11th International Learning Analytics and Knowledge Conference*, Irvine CA USA: ACM, Apr. 2021, pp. 490–496. doi: 10.1145/3448139.3448188.

[9] H. Jiang, B. Xiao, Y. Luo, and J. Ma, “A self-attentive model for tracing knowledge and engagement in parallel,” *Pattern Recognition Letters*, vol. 165, pp. 25–32, Jan. 2023, doi: 10.1016/j.patrec.2022.11.016.

[10] X. Wang, L. Chen, and M. Zhang, “Deep Attentive Model for Knowledge Tracing,” *AAAI*, vol. 37, no. 8, pp. 10192–10199, Jun. 2023, doi: 10.1609/aaai.v37i8.26214.

[11] L. Tu, X. Zhu, and Y. Ji, “Graph-based Dynamic Interactive Knowledge Tracing,” in *Proceedings of the 2023 8th International Conference on Distance Education and Learning*, Beijing China: ACM, Jun. 2023, pp. 49–56. doi: 10.1145/3606094.3606124.

[12] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, “Variational Inference: A Review for Statisticians,” *Journal of the American Statistical Association*, vol. 112, no. 518, pp. 859–877, Apr. 2017, doi: 10.1080/01621459.2017.1285773.

[13] D. J. C. Mackay, “Probable networks and plausible predictions — a review of practical Bayesian methods for supervised neural networks,” *Network: Computation in Neural Systems*, vol. 6, no. 3, pp. 469–505, Jan. 1995, doi: 10.1088/0954-898X_6_3_011.

[14] S. Maeda, “A Bayesian encourages dropout,” 2014, *arXiv*. doi: 10.48550/ARXIV.1412.7003.

[15] Y. Gal and Z. Ghahramani, “Bayesian Convolutional Neural Networks with

Bernoulli Approximate Variational Inference,” 2015, *arXiv*. doi: 10.48550/ARXIV.1506.02158.

[16] G. Abdelrahman, Q. Wang, and B. Nunes, “Knowledge Tracing: A Survey,” *ACM Comput. Surv.*, vol. 55, no. 11, pp. 1–37, Nov. 2023, doi: 10.1145/3569576.

[17] X. Song, J. Li, T. Cai, S. Yang, T. Yang, and C. Liu, “A survey on deep learning based knowledge tracing,” *Knowledge-Based Systems*, vol. 258, p. 110036, Dec. 2022, doi: 10.1016/j.knosys.2022.110036.

[18] A. Ghosh, N. Heffernan, and A. S. Lan, “Context-Aware Attentive Knowledge Tracing,” in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, Virtual Event CA USA: ACM, Aug. 2020, pp. 2330–2339. doi: 10.1145/3394486.3403282.

[19] Y. Yin et al., “Tracing Knowledge Instead of Patterns: Stable Knowledge Tracing with Diagnostic Transformer,” in *Proceedings of the ACM Web Conference* 2023, Austin TX USA: ACM, Apr. 2023, pp. 855–864. doi: 10.1145/3543507.3583255.

[20] W. Lee, J. Chun, Y. Lee, K. Park, and S. Park, “Contrastive Learning for Knowledge Tracing,” in *Proceedings of the ACM Web Conference 2022, Virtual Event*, Lyon France: ACM, Apr. 2022, pp. 2330–2338. doi: 10.1145/3485447.3512105.

[21] S. Ruan, W. Wei, and J. Landay, “Variational Deep Knowledge Tracing for Language Learning,” in *LAK21: 11th International Learning Analytics and Knowledge Conference*, Irvine CA USA: ACM, Apr. 2021, pp. 323–332. doi: 10.1145/3448139.3448170.

[22] Y. Yang et al., “GIKT: A Graph-Based Interaction Model for Knowledge Tracing,” in *Machine Learning and Knowledge Discovery in Databases*, vol. 12457, F. Hutter, K. Kersting, J. Lijffijt, and I. Valera, Eds., in Lecture Notes in Computer Science, vol. 12457. , Cham: Springer International Publishing, 2021, pp. 299–315. doi: 10.1007/978-3-030-67658-2_18.

[23] T. Liu, M. Zhang, C. Zhu, and L. Chang, “Transformer-based convolutional forgetting knowledge tracking,” *Sci Rep*, vol. 13, no. 1, p. 19112, Nov. 2023, doi: 10.1038/s41598-023-45936-0.

[24] Y. Bai, J. Zhao, T. Wei, Q. Cai, and L. He, “A survey of explainable knowledge tracing,” *Appl Intell*, vol. 54, no. 8, pp. 6483–6514, Apr. 2024, doi: 10.1007/s10489-024-05509-8.

[25] D. Shimada and K. Okada, “Reliability Coefficient for Bayesian Knowledge Tracing Models,” *Tech Know Learn*, Apr. 2025, doi: 10.1007/s10758-025-09829-7.