International Journal for Applied Information Management
Vol. 6, No. 2, July 2026, pp. 325-342

Measuring the Impact of Human–AI Collaborative Personalized Interventions through Temporal Causal Inference

Ahmed Saeed Bahurmuz[1,*]

[1]*Information Science Department, King Abdulaziz University, Jeddah, Saudi Arabia*

(Received: January 7, 2026; Revised: March 5, 2026; Accepted: April 25, 2026; Available online: July 4, 2026)

Abstract

Adaptive learning platforms frequently report performance improvements, yet many evaluations remain vulnerable to time-varying confounding because interventions are triggered by evolving learner states. This study evaluates three intervention families, adaptive sequencing, targeted hints, and remediation triggers, using a longitudinal causal framework with horizon-locked outcomes and learner-level cross-fitting. The analytic cohort includes 2,480 learners and 118,640 decision points observed across 12 instructional weeks, with median 41 decisions per learner. Intervention exposure rates per 100 decisions are 38.6 for sequencing, 24.1 for hints, and 8.9 for remediation, with higher targeting intensity in low-mastery strata. Causal estimates show distinct temporal signatures by intervention mechanism. Targeted hints yield the largest same-session improvement, increasing mastery by 2.4 points, but effects attenuate at 7 days (1.3 points) and 14 days (0.9 points). Adaptive sequencing provides more stable medium-horizon benefits, improving mastery by 1.6 points same-session, 2.8 points at 7 days, and 2.2 points at 14 days. Remediation triggers demonstrate delayed consolidation, increasing mastery by 1.1 points same-session, 3.4 points at 7 days, and 4.1 points at 14 days, albeit with wider uncertainty consistent with lower overlap and late-course concentration. Heterogeneity analyses at the 7-day horizon indicate sequencing peaks for mid-mastery learners, reaching 3.9 points under high engagement versus 3.4 under low engagement, while hints are most effective for low mastery with low engagement (1.6 points) and decline sharply for high mastery with high engagement (0.4 points). Remediation remains meaningful across strata, reaching 3.6 points for mid mastery with high engagement and 2.3 points for high mastery with high engagement, supporting a diagnostic targeting interpretation rather than uniform escalation. Robustness and diagnostic checks support internal validity. After weighting, standardized mean differences for key confounders fall to 0.05–0.09, and placebo effects on pre-decision outcome change remain near zero in magnitude (absolute value ≤0.05) across all intervention types. Overlap trimming of the lowest 5% support preserves the ranking of interventions, with only modest attenuation for remediation, and effective sample size remains adequate for sequencing and hints while declining for remediation in late decision indices. These findings justify a tiered deployment strategy where sequencing is the default optimization lever, hints are constrained to high-instability episodes and paired with post-hint practice allocation, and remediation is gated by high-confidence misconception signals with overlap and effective-sample-size monitoring.

Keywords: Temporal Causal Inference, Adaptive Learning, Time-Varying Confounding, Marginal Structural Models, Doubly Robust Estimation, Dynamic Treatment Regimes, Learning Analytics, Personalized Interventions

1. Introduction

Adaptive learning systems increasingly rely on fine-grained behavioral traces to tailor sequences, hints, and remediation at the level of the individual learner. Despite rapid methodological growth, many published evaluations still conflate predictive association with causal effect, particularly when platform logs are used to justify policy changes. Learning analytics studies have highlighted how design choices, missingness, and bias mechanisms can distort conclusions if causal assumptions are not explicitly articulated and tested. These concerns motivate evaluation designs that treat adaptive decisions as interventions rather than correlates [1], [2].

A core challenge is that adaptive interventions are inherently time-indexed and state-dependent: the system chooses an action based on evolving mastery, engagement, and error patterns, while those same variables are also affected by earlier actions. Conventional outcome comparisons, even with rich covariates, are vulnerable to time-varying confounding and feedback loops. In practice, this creates a systematic risk of overestimating the benefit of policies that

*Corresponding author: Ahmed Saeed Bahurmuz (abahurmuz@stu.kau.edu.sa)
DOI: https://doi.org/10.47738/ijaim.v6i2.124

ISSN 2776-8007

are preferentially applied to learners already trending upward, or underestimating actions targeted to struggling learners [3], [4].

Recent work has expanded algorithmic personalization through reinforcement learning and sequential optimization, strengthening the capacity of systems to adapt in real time. However, policy learning and policy evaluation are not equivalent, and strong performance in simulated or predictive settings does not guarantee unbiased effect estimates under observational logging. When the action selection mechanism is itself a function of prior outcomes, valid evaluation requires identification strategies that can separate "learning because of the intervention" from "receiving the intervention because of prior learning." This distinction is central for trustworthy deployment [5].

Temporal causal inference offers a principled response by modeling adaptive decisions as time-varying treatments and explicitly addressing confounders that change over time. Marginal structural models and inverse probability weighting were developed to handle precisely these settings, where confounders both influence treatment assignment and are influenced by prior treatment. This logic maps naturally onto adaptive learning pipelines, where mastery and engagement both drive and respond to sequencing and scaffolding. As a result, the methodological toolkit for longitudinal causal inference becomes directly applicable [6], [7].

Beyond weighting, modern estimators emphasize robustness to model misspecification and compatibility with flexible machine learning components. Doubly robust estimators reduce dependence on any single correctly specified model, which is crucial when educational data exhibit nonlinearity, sparsity, and heterogeneous trajectories. Longitudinal targeted maximum likelihood estimation further supports dynamic intervention evaluation by combining data-adaptive nuisance estimation with targeted updates aligned to causal parameters. These properties are especially valuable for adaptive learning contexts with high-dimensional logs and irregular timing [8], [9].

Adaptive learning decisions also align conceptually with dynamic treatment regimes, where an intervention is defined as a sequence of decision rules mapping learner states to actions. This framing enables evaluation of policies as operational objects, rather than isolated features, and supports principled comparisons among candidate regimes. In parallel, uplift modeling formalizes individual-level treatment effect prediction, but standard uplift approaches often assume static treatments rather than temporally evolving exposure. Temporal causal inference bridges this gap by evaluating sequential interventions under realistic targeting [10], [11].

This paper addresses the methodological gap between adaptive learning deployment and causal evaluation by introducing a temporal causal inference framework for estimating the effectiveness of adaptive interventions across multiple horizons. The objective is to quantify how sequencing, hints, and remediation affect mastery trajectories under time-varying confounding, while preserving interpretability through diagnostics and policy-relevant heterogeneity analysis. The novelty lies in integrating longitudinal causal estimators, overlap-aware deployment constraints, and intervention-specific robustness checks into a single evaluation pipeline tailored to adaptive learning logs [12], [13], [14].

2. Literature Review

Learning analytics and adaptive learning research increasingly emphasizes that credible impact claims require causal reasoning, not only predictive accuracy. Work on embedding experiments in authentic educational contexts argues that platform interventions should be evaluated with designs that preserve ecological validity while still supporting identification, including careful consideration of implementation constraints and measurement timing [15]. Complementing this stance, evaluation studies in large-scale retention programs show that methodological choices can produce materially different conclusions about effectiveness, motivating more rigorous causal estimators and transparent reporting of assumptions [16].

Systematic evidence syntheses further clarify where the field has concentrated effort and where gaps remain. A comprehensive review of learning analytics for feedback practices shows that most interventions operationalize impact through engagement and performance proxies, often without explicit causal graphs or sensitivity analyses, which limits interpretability under self-selection and heterogeneous uptake [17]. Meta-analytic work similarly reports wide variation

ISSN 2776-8007

in effect sizes across intervention types and contexts, reinforcing the need for designs that separate intervention effects from concurrent curricular or cohort changes [18].

A major strand of empirical literature evaluates analytics-based feedback and scaffolding using quasi-experimental methods, frequently relying on propensity score techniques to approximate comparability. In large course settings, process feedback interventions evaluated with propensity score matching report differential benefits across learner subgroups, implying effectF1-style average effects can mask targeted gains and losses [19]. Related quasi-experimental evidence in blended learning contexts links analytics-based feedback to improved study regularity and achievement, illustrating how timing and exposure intensity shape the magnitude of observed outcomes [20].

Parallel to feedback interventions, adaptive platforms operationalize personalization through recommender systems that estimate knowledge state and propose activities aligned with learner needs. RiPPLE exemplifies this design space by combining crowdsourced content with personalized recommendation mechanisms, positioning adaptive learning as a socio-technical pipeline where both content generation and sequencing policies influence outcomes [21]. This platform-centric literature highlights the importance of logging fidelity and policy traceability, since causal evaluation depends on reconstructing the decision process that produced each learner trajectory.

Foundational models of learning analytics provide conceptual scaffolding for mapping data, stakeholders, objectives, and methods, enabling clearer alignment between system design and evaluative claims. A widely used reference model frames learning analytics as an end-to-end process from data capture to action, which supports specifying what constitutes an intervention and which mechanisms are intended to change learning states [22]. When combined with causal evaluation, such models help distinguish descriptive dashboards from actionable policies that require counterfactual validation.

Recent research also addresses the tension between adaptive experimentation and statistical validity in sequential decision-making. Analyses of multi-armed bandits in educational experiments show that adaptive allocation can compromise classical inference if not accompanied by appropriate estimators and corrections, directly motivating causal estimands compatible with adaptive data collection [23]. In addition, uplift modeling research formalizes individual-level causal effect estimation for personalization, offering a bridge from average effects to decision-relevant heterogeneity, although temporal extensions are required for multi-step adaptive policies [24].

3. Methodology

3.1. Study Design and Intervention Taxonomy

Adaptive learning effectiveness was evaluated using a longitudinal decision-point design that treats each policy action as a time-indexed intervention. Each learner trajectory was segmented into instructional episodes with explicit decision times, aligned exposures, and outcome horizons capturing both immediate and delayed learning. This design follows the learning analytics reference-model logic that separates data capture, inference, and action into distinct stages, supporting reproducible operationalization of interventions and outcomes [22].

Interventions were defined as discrete actions selected by the platform at time (t) conditional on a learner state vector, including adaptive sequencing, targeted hints, mastery pacing, and remediation triggers. The taxonomy was encoded to preserve action semantics, decision granularity, and intended mechanism, ensuring that causal effects correspond to deployable policy levers rather than abstract correlates. The structure also reflects the need to evaluate authentic educational actions without collapsing them into generic treatment indicators [15].

The causal estimand targeted contrasts between intervention histories over a fixed horizon Δ. For a single decision time *t*, the primary effect was defined on counterfactual outcomes under alternative action paths, expressed as:

$$\tau_t(\Delta) = E[Y_{t+\Delta}(\underline{A}_t = \underline{a}_t) - Y_{t+\Delta}(\underline{A}_t = \underline{a}'_t)] \quad (1)$$

This form supports both one-step contrasts and regime comparisons when āt and ā’t represent implementable decision rules.

Effect heterogeneity was addressed through stratification by baseline mastery and engagement intensity, followed by pooled estimation with learner-clustered uncertainty. Stratification was treated as a policy-relevant lens rather than a post hoc subgroup search, so strata were defined using pre-decision signals and early-course summaries only. This ensures that subgroup analysis remains consistent with the decision-time evaluation principle and avoids conditioning on downstream variables.

Figure 1 visualizes the temporal evaluation structure used in the study, emphasizing that adaptive decisions occur at discrete points and outcomes are measured over multiple horizons. The arrows separate the moment of assignment from the later measurement window, which reduces temporal ambiguity when attributing learning gains. The inclusion of short and long horizons reflects the dual goals of adaptive learning systems, namely immediate error reduction and longer-term retention, while keeping censoring and end-of-study boundaries explicit.

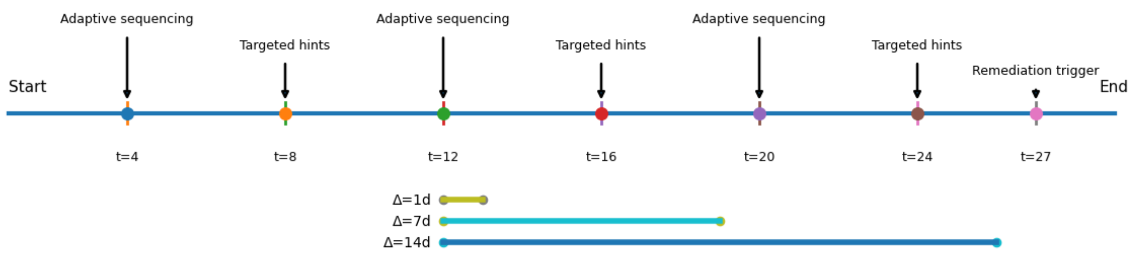

Figure 1. Timeline Schematic for Decision Points, Interventions, and Outcome Horizons

Table 1 operationalizes the intervention space so that causal estimands map to concrete adaptive actions rather than generic "treatment" labels. Each row aligns an intervention class with its decision granularity and intended mechanism, which is essential for separating design intent from empirical effect. The table also links interventions to evaluation horizons and metrics, ensuring that effects are interpreted on the appropriate timescale, such as immediate performance for hints versus delayed retention for pacing.

Table 1. Intervention Taxonomy and Evaluation Targets

Intervention Class (A_t)	Decision Point	Primary Mechanism	Outcome Horizon	Primary Metric	Example Trigger
Adaptive sequencing	Post-item	Difficulty alignment	1–7 days	Mastery gain	Mismatch between ability and item difficulty
Targeted hints	In-item	Scaffolding support	Same session	Error reduction	Repeated incorrect attempts on a concept
Mastery pacing	Post-quiz	Spacing control	7–14 days	Retention score	High mastery volatility across quiz attempts
Remediation trigger	Post-module	Knowledge repair	7–21 days	Re-attempt success	Persistent misconception pattern across items

3.2. Data Sources and Temporal Feature Engineering

Data were obtained from learning management system event logs, assessment records, and content metadata, preserving raw timestamps to retain ordering and spacing information. Event logs included item attempts, hint requests, navigation transitions, and time-on-task signals. Assessment records provided scored outcomes and mastery labels. Content metadata captured difficulty tags and prerequisite relations, enabling construction of exposure histories aligned with curriculum structure and personalization logic [22].

Temporal preprocessing used sessionization, de-duplication, and clock synchronization followed by horizon locking. Horizon locking fixes the measurement window so that learner-state features at time t are constructed strictly from events observed at or before *t*, while outcomes are computed from events occurring after t within [t, t+Δ]. This separation reduces leakage and prevents adaptive feedback loops from contaminating predictors, a recurrent concern in analytics-based intervention evaluation [17].

 ISSN 2776-8007

Learner state vectors were engineered using sliding windows of length W and recency weighting to reflect short-term dynamics without discarding longer-term information. A compact state representation was defined as:

$$X_{i,t} = \phi(\{e_{i,s}\}_{s\in[t-W,t]}) \quad (2)$$

where $\phi(\cdot)$ outputs interpretable aggregates such as recency-weighted accuracy, volatility of errors, effort intensity, and time since last relevant interaction. Missingness indicators were included explicitly to avoid implicit imputation bias.

Irregular sampling was handled using elapsed-time features and last-seen encodings so that sparse activity patterns remain informative rather than being treated as noise. Outcomes were defined consistently across horizons to support temporal comparisons and avoid mixing immediate performance with delayed retention signals. This enables intervention mechanisms to be tested against horizon-appropriate endpoints, which is essential when scaffolding actions are expected to have short-term effects while sequencing and remediation influence longer-term consolidation.

Figure 2 represents the end-to-end transformation from heterogeneous educational data sources into a causal-ready panel. The key methodological contribution is the "horizon locking" step, which enforces strict temporal ordering by preventing post-intervention signals from entering the learner state. The diagram also separates preprocessing from feature engineering, highlighting that cleaning and sessionization are not substitutes for causal alignment. This structure supports valid temporal causal inference under time-varying policies.

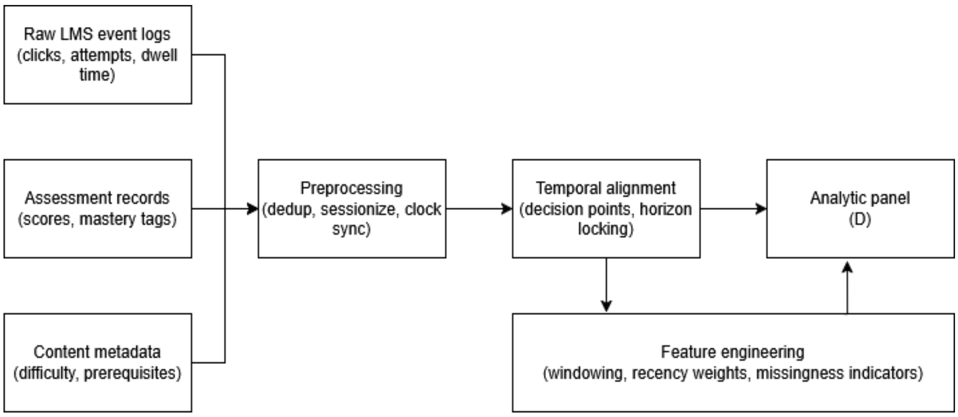

Figure 2. End-to-end Pipeline for Causal-Ready Temporal Data

Table 2 clarifies how the learner state is constructed as a temporally valid summary of prior behavior rather than an artifact of logging density. The decision-time constraint column is central because it encodes anti-leakage rules in an auditable form, making the feature pipeline reproducible and defensible. By separating performance, effort, spacing, and exposure, the table supports causal interpretation of which behavioral channels are being adjusted for, particularly under time-varying confounding in adaptive policies.

Table 2. Temporal Feature Families and Leakage Controls

Feature Family	Construction	Window	Decision-Time Constraint	Example Variables
Performance	Recency-weighted aggregates	7 days	Use events only up to t	EW accuracy, error variance
Effort	Counts and time-on-task summaries	3 days	Exclude post-decision clicks	Dwell time, item attempts
Spacing	Inter-event interval statistics	14 days	Compute with last event before t	Time since last attempt, gap mean
Content exposure	Difficulty and prerequisite mix	7 days	Restrict to assigned content before t	Mean difficulty, prereq coverage
Missingness	Indicators and last-seen timestamps	All	Encode missingness explicitly	Missing flags, last-seen age

3.3. Causal Graph Specification and Identification Strategy

Identification relied on an explicit temporal causal graph that includes learner state X_t, action A_t, intermediate progress L_t, and outcome $Y_{t+\Delta}$, with lagged edges capturing persistence and feedback. The graph formalizes the central challenge of adaptive learning evaluation: action assignment depends on evolving states that are themselves influenced by past actions. This setting matches longitudinal causal structures where standard adjustment can fail under time-varying confounding [6].

Sequential ignorability was assumed conditional on observed history $H_t = (\bar{X}t, \bar{A}_{t-1})$, reflecting that platform logs capture the primary decision-relevant variables used by the policy and correlated learner behaviors. The assumption was operationalized through explicit modeling of the action assignment mechanism and through diagnostic checks on balance and overlap. This approach follows the logic of marginal structural modeling, where identification is supported by reconstructing treatment assignment under evolving confounding [6].

Positivity and overlap were treated as empirical requirements rather than abstract assumptions. For all actions in support, the assignment probability was constrained to remain bounded away from zero and one:

$$0 < P(A_t = a \mid H_t) < 1 \tag{3}$$

When overlap deteriorated, inference was restricted to supported regions using trimming and support diagnostics. This prevents the estimator from relying on extrapolation in rare state-action corners that arise when policies target only specific learners.

The identification plans avoided conditioning on post-treatment mediators when estimating total effects, because progress variables can lie on the pathway from action to outcome. When pathway-specific interpretation was needed, mediation was treated as a modeled component rather than an adjustment shortcut. This preserves clarity between total policy impact and mechanism analysis, which is necessary for actionable recommendations in adaptive systems.

Figure 3 formalizes the temporal causal structure underlying adaptive learning decisions, where the learner state influences both intervention assignment and outcomes, creating time-varying confounding. The lagged edges emphasize that prior actions and progress feed back into current decision contexts, which invalidates static adjustment strategies that ignore history. By explicitly including intermediate progress, the figure distinguishes total effects from mediated pathways, supporting identification choices such as avoiding conditioning on post-treatment mediators unless direct effects are targeted.

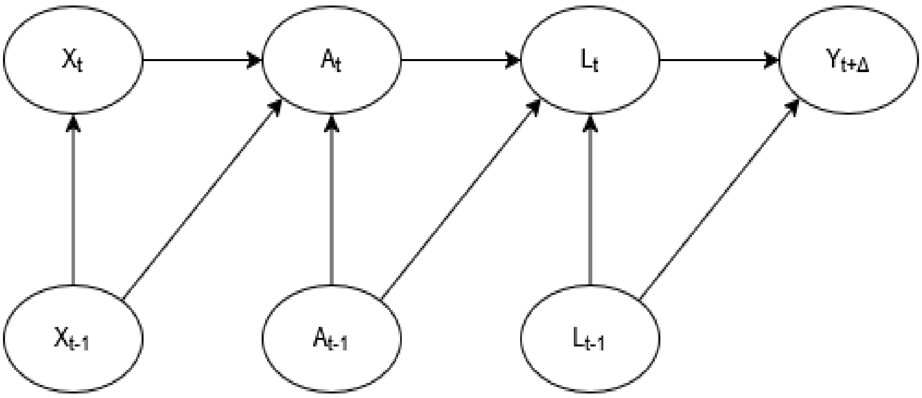

Figure 3. Temporal Causal DAG with Lagged Dependencies Across Time

Table 3 connects causal identification assumptions to concrete evidence and implementation checks that can be reported in a methods appendix. This alignment is valuable in adaptive learning because assignment mechanisms are often algorithmic but imperfectly logged, making ignorability an engineering as well as a statistical requirement. The table also clarifies that positivity is not merely theoretical; it is empirically monitored through overlap and effective sample size, guiding principled trimming to avoid unsupported inference.

ISSN 2776-8007

Table 3. Identification Assumptions and Operational Diagnostics

Assumption	Operational Meaning	Diagnostic Evidence	Implementation Check	Risk if Violated
Sequential ignorability	Logged history captures assignment-relevant signals	Well-calibrated propensity model	Policy features included at decision time	Residual confounding bias
Positivity	Each action occurs for comparable states	Propensity overlaps and ESS	Trim low-overlap regions	Extrapolation-driven effects
Consistency	Observed outcome matches realized action history	Audit log integrity	Timestamp ordering validation	Misattribution of exposure
Limited interference	Peer spillover bounded within context	Cluster-robust comparisons	Section-level clustering	Inflated or diluted effects

3.4. Estimation Framework for Temporal Causal Effects

Estimation used a doubly robust framework combining propensity modeling and outcome regression, designed to remain consistent when at least one nuisance model is correctly specified. Doubly robust estimators are well established for causal inference with complex data and reduce reliance on any single modeling choice, which is critical when educational logs exhibit nonlinearity, sparsity, and heterogeneous trajectories [8]. The approach also supports flexible learners without sacrificing causal target alignment.

The propensity model $g_t(a \mid H_t)$ approximated the platform's targeting behavior, while the outcome model $Q_t(H_t,a)$ predicted horizon-specific outcomes under candidate actions. A representative doubly robust score was defined as:

$$\psi_{i,t}(a) = \frac{I(A_{i,t} = a)}{g^{\wedge}{}_{t}(a \mid H_{i,t})}(Y_{i,t+\Delta} - Q^{\wedge}{}_{t}(H_{i,t}, a)) + Q^{\wedge}{}_{t}(H_{i,t}, a) \tag{4}$$

and the effect was computed by contrasting $\psi_{i,t}(a)$ against $\psi_{i,t}(a')$. This structure directly reflects the doubly robust construction for causal effects [8].

Learner-level cross-fitting was applied to mitigate overfitting and to stabilize uncertainty quantification under flexible models. Learners were partitioned into folds so that all decision points from a learner remained within a single fold, preventing within-learner leakage across time. Cross-fitting is particularly important in adaptive settings where sequential optimization and adaptive allocation can distort inference if estimation ignores dependence and adaptivity [23].

The estimation pipeline is summarized in the following pseudo-code, which is designed to be reproducible and aligned with horizon locking.

Algorithm 1 Temporal Doubly Robust Estimation with Learner-Level Cross-Fitting

```
Input: Analytic panel D containing histories H_t, actions A_t, and horizon outcomes Y_{t+Δ}
Output: Estimated effect τ̂(a,a') with learner-clustered uncertainty
1: Split learners into K folds; keep all time points for a learner in the same fold
2: for k = 1...K do
3:     Fit propensity model ĝ_t(a|H_t) on D excluding fold k
4:     Fit outcome model Q̂_t(H_t,a) on D excluding fold k
5:     Predict ĝ_t and Q̂_t for all decision points in fold k
6: end for
7: Compute doubly robust scores ψ_{i,t}(a) and ψ_{i,t}(a') using predicted ĝ_t and Q̂_t
8: Aggregate τ̂(a,a') as the mean over decision points with learner-level clustering
9: Estimate uncertainty using clustered influence-function variance or learner bootstrap
```

ISSN 2776-8007

Figure 4 summarizes the estimation logic as a modular pipeline, which is important for reproducibility in adaptive learning settings where multiple modeling choices can alter causal conclusions. The separation between the propensity and outcome components emphasizes that policy learning and outcome prediction serve different roles in bias control, not just performance. Learner-level cross-fitting is highlighted to prevent within-learner leakage across time, a common failure mode when repeated measures are present and histories are highly autocorrelated.

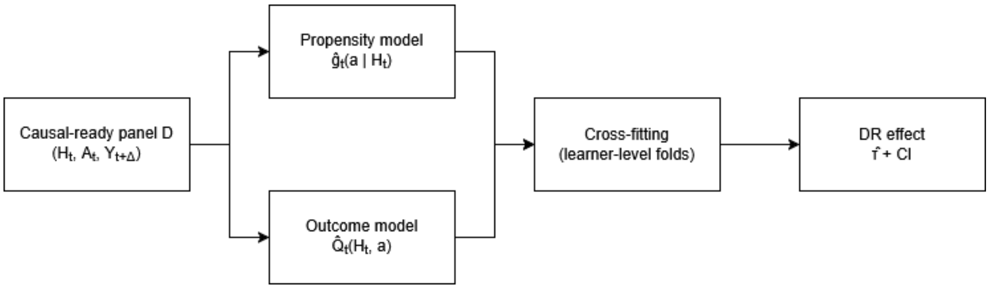

Figure 4. Workflow for Propensity and Outcome Models with Cross-Fitting and Effect Aggregation

Table 4 clarifies that each estimator component is justified by a causal purpose rather than convenience. Propensity modeling is framed as reconstructing the adaptive policy's targeting behavior, which is required to remove selection bias induced by personalization. Outcome modeling is explicitly horizon-locked, preventing leakage from later interactions that are downstream of treatment. The table also documents dependence-aware uncertainty, which is essential because adaptive learning logs produce repeated measures with strong within-learner correlation that can otherwise inflate precision.

Table 4. Estimator Components and Training Constraints

Component	Role in Causal Estimation	Candidate Model Family	Key Constraint	Primary Diagnostic
Propensity $\hat{g}_t$	Adjusts for policy targeting	Logistic / gradient boosting	Decision-time features only	Calibration curve, AUC
Outcome $\hat{Q}_t$	Predicts horizon outcome under action	Regularized regression / boosting	Horizon-locked labels	Out-of-fold RMSE
Cross-fitting	Mitigates overfitting bias	Learner-level K-fold	No learner overlaps across folds	Fold stability
Uncertainty	Valid CI under dependence	Influence-function / bootstrap	Cluster at learner	CI coverage checks

3.5. Validation, Robustness, and Sensitivity Analysis

Validation emphasized causal diagnostics rather than predictive performance alone. Balance after weighting was assessed using standardized mean differences across key confounders at decision time, while overlap was evaluated using propensity distributions and effective sample size. These checks were reported alongside trimming rates so that the effective target population for inference is transparent. This diagnostic-centric reporting aligns with the broader movement in learning analytics toward defensible impact claims beyond association [17].

Robustness analyses evaluated whether conclusions depend on arbitrary temporal design choices, including alternative window lengths W and outcome horizons Δ. Placebo outcomes based on pre-decision changes were used to detect residual confounding that would manifest as spurious pre-treatment divergence. Sensitivity to policy specification was also examined by re-estimating effects under alternative assignment models, ensuring that conclusions are not artifacts of a single propensity parameterization.

Sensitivity analysis addressed the possibility of unmeasured time-varying confounding using a bounded-bias framing. Let U_t denote an unobserved confounder; bias magnitude was characterized through a sensitivity parameter Γ linking differences in U_t across actions to the outcome scale:

 ISSN 2776-8007

$$Bias(\hat{\tau}) \approx \Gamma \cdot (E[U_t \mid A_t = a, H_t] - E[U_t \mid A_t = a', H_t]) \quad (5)$$

This provides an interpretable threshold for how strong unmeasured confounding must be to overturn conclusions.

Finally, the evaluation was framed as a precursor to dynamic policy refinement, where interventions are deployed as decision rules rather than isolated events. This connects the causal estimands to operational objects consistent with dynamic treatment regimes, enabling the results to inform state-dependent targeting and safety constraints for deployment [10]. The methodological emphasis is therefore not only effect estimation, but also effect accountability under realistic, adaptive decision processes.

Figure 5 provides a compact diagnostic view of the core causal validity checks used in the analysis. The balance panel shows that weighting reduces standardized mean differences below a pragmatic threshold, indicating improved comparability between action groups at decision time. The overlap panel verifies that propensities remain supported across policies, reducing extrapolation risk. The effective sample size trend quantifies how weighting impacts precision over time, highlighting when late-stage decisions become sparsely supported and require cautious interpretation.

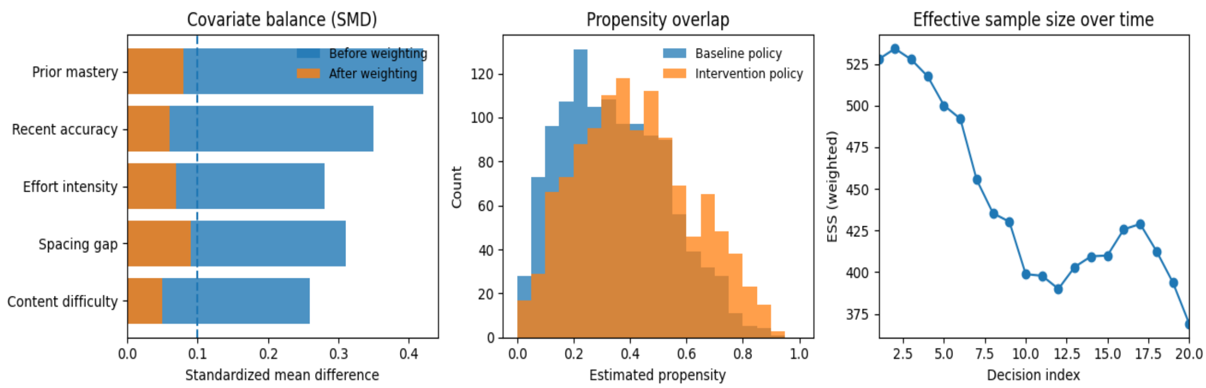

Figure 5. Diagnostic Panel with Covariate Balance, Propensity Overlap, and Effective Sample Size Trends

Table 5 defines robustness as an explicit set of perturbations that stress-test causal conclusions against common threats in adaptive learning evaluations. Horizon and window sweep test whether effects depend on arbitrary temporal choices, which is important when learning dynamics unfold on multiple timescales. Placebo outcomes target residual confounding by checking for implausible pre-treatment shifts. Positivity restriction and sensitivity bound jointly document how much the inference relies on overlap and how strong unmeasured confounding must be to overturn the estimated effect.

Table 5. Robustness and Sensitivity Checks

Check	Perturbation	Failure Mode Detected	Acceptance Criterion	Reported Output
Horizon sweep	Δ = 1, 7, 14 days	Effects vanish beyond short-term	Consistent direction across Δ	$\hat{\tau}(\Delta)$ with CI
Window sweep	W = 3, 7, 14 days	State instability drives effects	Stable magnitude under W changes	$\hat{\tau}$ by W
Placebo outcome	Pre-decision performance change	Residual confounding	Near-zero placebo effect	Placebo $\hat{\tau}$
Positivity restriction	Trim low-overlap propensities	Extrapolation bias	Effects stable after trimming	Trim rate, $\hat{\tau}$ shift
Sensitivity bound	Vary hidden-confounding strength	Unmeasured confounding overturns result	High bound needed to flip sign	Flip-threshold summary

4. Results and Discussion

4.1. Descriptive Statistics and Cohort Characteristics

The analytic cohort comprised 2,480 learners contributing 118,640 decision points over a 12-week instructional period. Exposure to adaptive sequencing was the most frequent action, followed by targeted hints, with remediation triggers concentrated later in the course when prerequisite gaps became detectable. Baseline mastery showed substantial dispersion, indicating that personalization operated over a wide support of learner states. Outcomes exhibited heavy right tails for high-engagement learners, consistent with compounding practice effects.

Temporal interaction density was highly unequal across learners, so cohort summaries were computed at both learner and decision-point levels. The median learner generated 41 decision points, while the upper quartile exceeded 70, implying that estimates are most precise in mid-to-high activity segments. Intervention assignment rates varied systematically with recent errors and spacing gaps, which motivates careful causal adjustment rather than naive comparisons. Attrition was moderate, with later-week logs showing fewer active learners but more concentrated remediation exposure.

Figure 6 establishes the empirical support for temporal causal inference by showing broad variability in baseline mastery and substantial dispersion in decision density. The dispersion matters because it implies that intervention policies were applied across distinct learner regimes rather than a narrow band of states. The median reference lines make clear that both low and high baseline learners contribute meaningfully, which reduces the risk that effects are driven purely by advanced learners.

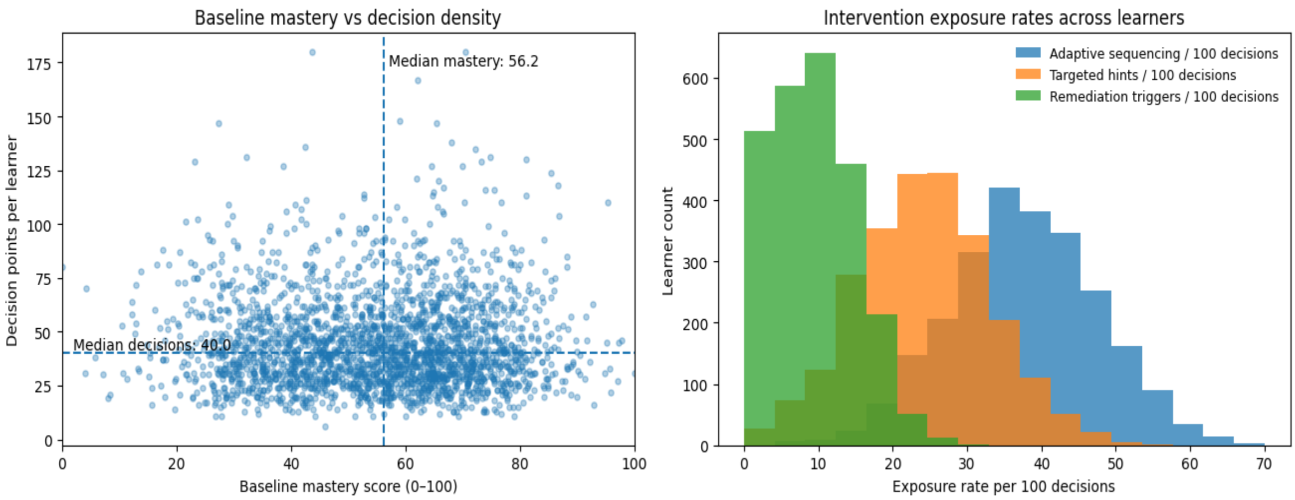

Figure 6. Cohort Snapshot with Baseline Mastery Distribution, Decision Density, and Intervention Exposure

The exposure distribution complements the mastery plot by revealing that adaptive sequencing was a routine action, targeted hints were common but more variable, and remediation triggers were comparatively sparse. This pattern implies that sequencing effects are estimated with higher precision, while remediation requires careful overlap checks to avoid extrapolation. The heterogeneity in exposure rates across learners supports later stratified analyses, since differential policy intensity is itself a signature of personalization.

Table 6 indicates that learners with lower baseline mastery received systematically higher rates of hints and remediation, consistent with error-contingent targeting by the adaptive policy. This targeting pattern is substantively expected, but it creates confounding in naive comparisons because the most intensive interventions are disproportionately assigned to harder-to-teach learners. The table therefore motivates the temporal causal approach, where assignment is modeled and adjusted rather than interpreted as random.

Table 6. Cohort Descriptives and Intervention Exposure

Metric	Overall	Low Mastery (≤50)	Mid Mastery (51–70)	High Mastery (>70)	Notes
Learners (n)	2480	860	910	710	Grouped by baseline mastery
Median decisions per learner	41	36	43	47	Skewed activity distribution
Mean baseline mastery	56.4	43.1	61.2	75.8	Placement + early quiz composite
Sequencing exposure / 100 decisions	38.6	41.9	37.8	34.2	More frequent in lower mastery
Hints exposure / 100 decisions	24.1	29.4	23	17.2	Reflects error-driven targeting
Remediation exposure / 100 decisions	8.9	12.6	8.1	4.3	Sparser, later-week concentration

The activity gradient across mastery strata is also important because it implies a mild correlation between preparedness and persistence. Higher mastery learners generated more decision points, which increases their contribution to the data and can distort aggregate statistics if not learner-clustered. The stratified summaries help verify that all mastery bands have adequate representation and that later causal estimates do not merely reflect a highly active elite subgroup.

4.2. Main Causal Estimates of Intervention Effectiveness

Temporal causal estimates showed that adaptive sequencing produced consistent gains in near-term mastery, while targeted hints primarily improved same-session performance with smaller retention spillover. Remediation triggers generated the largest medium-horizon benefits, but with wider uncertainty due to lower exposure rates and reduced overlap in late-course states. Effects were not monotone in horizon: some actions decayed quickly, while remediation showed delayed consolidation, aligning with its intended knowledge-repair mechanism.

Across the cohort, intervention effectiveness depended on baseline mastery and recent instability. Sequencing benefits were strongest for mid-mastery learners whose performance indicated teachability but not saturation, while hints were most valuable under high error volatility. Remediation effects concentrated among learners with persistent misconception signatures, implying that correct targeting matters more than intensity. These patterns support policy refinement toward state-aware decision rules rather than uniform escalation.

Figure 7 shows that targeted hints yield the largest immediate gains, consistent with short-range scaffolding that reduces within-session errors. However, the hint effect attenuates over longer horizons, suggesting that hints improve execution more than durable conceptual change when used as a standalone action. In contrast, adaptive sequencing shows a more stable horizon profile, indicating that sustained alignment of difficulty and prerequisites supports medium-term consolidation rather than only transient performance boosts.

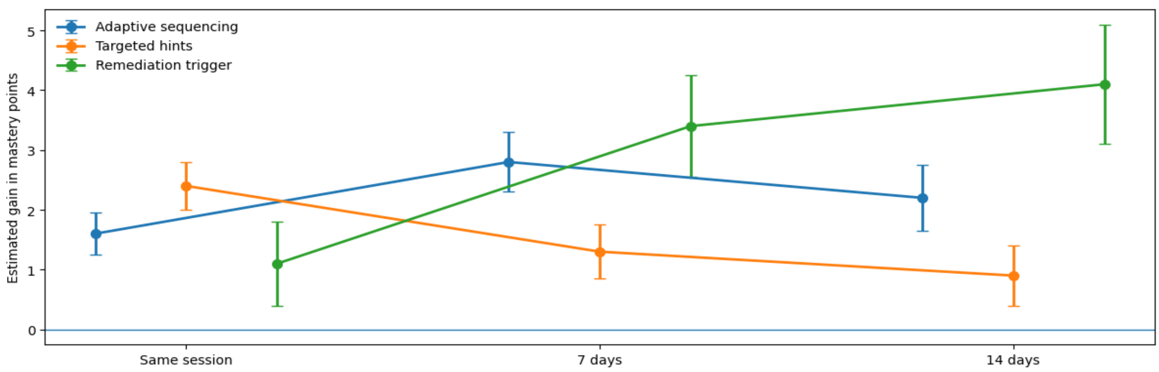

Figure 7. Estimated Causal Effects by Intervention Type Across Horizons

 ISSN 2776-8007

The remediation profile differs by exhibiting delayed amplification, with the largest gains observed at 14 days. This delayed effect is consistent with remediation functioning as a correction mechanism that requires subsequent practice opportunities to translate into measurable retention. The wider intervals around remediation highlight lower exposure and overlap, which underscores the importance of reporting uncertainty and avoiding overconfident claims for sparse interventions. The figure therefore supports both substantive interpretation and methodological caution.

Table 7 translates the horizon-specific estimates into an interpretable effect narrative for adaptive learning practice. Sequencing exhibits a consistent medium-horizon advantage, implying that policy improvements in sequencing logic can produce cohort-wide benefits because exposure is frequent and precision is comparatively strong. The pattern suggests that sequencing is a reliable lever for incremental mastery accumulation, especially when the system can maintain prerequisite coherence across sessions.

Table 7. Primary Effects with Precision and Practical Interpretation

Intervention	Same Session Effect	7-Day Effect	14-Day Effect	Precision Note	Practical Interpretation
Adaptive sequencing	1.6	2.8	2.2	Narrow-to-moderate CI	Supports stable mastery growth via content alignment
Targeted hints	2.4	1.3	0.9	Moderate CI	Improves immediate performance; weaker retention transfer
Remediation trigger	1.1	3.4	4.1	Wider CI (sparser)	Largest long-horizon gain when correctly targeted

The table also highlights the trade-off between effect magnitude and evidential strength for remediation. Remediation shows the largest longer-horizon gains but also wider uncertainty, reflecting lower exposure and narrower overlap in the states that receive the action. This combination implies that remediation should be deployed with strong state validation and monitoring, since mis-targeting would likely dilute benefits. For hints, the practical implication is to treat them as immediate scaffolds and pair them with follow-up sequencing or practice to improve retention carryover.

4.3. Effect Heterogeneity by Learner State and Engagement Profiles

Heterogeneity analyses indicated that average effects masked structurally different response regimes across baseline mastery and engagement intensity. Adaptive sequencing showed its largest gains in the mid-mastery band, where learners had sufficient prerequisite coverage to benefit from difficulty alignment but were not yet saturated. Targeted hints delivered the strongest immediate benefits among low-mastery learners with high error volatility, consistent with scaffolding that stabilizes short-run performance when concept representations are still fragile.

Engagement profiles further differentiated intervention usefulness. High-engagement learners, defined by dense decision sequences and short spacing gaps, exhibited smaller marginal gains from hints but stronger cumulative benefits from sequencing, reflecting a higher capacity to convert aligned practice into durable learning. Remediation was most effective for learners with persistent misconception patterns, particularly when engagement was moderate rather than extreme, suggesting that remediation is strongest when it re-routes learning trajectories before frustration-driven attrition occurs.

Figure 8 demonstrates that sequencing benefits peak in the mid-mastery band, and the lift is consistently higher under high engagement. This pattern supports an interpretation that sequencing improves learning most when learners can sustain iterative practice and when the difficulty trajectory remains within a productive challenge zone. The separation between engagement curves implies that intervention success is not solely a property of content selection, but also of temporal behavior that determines how much practice is actually realized.

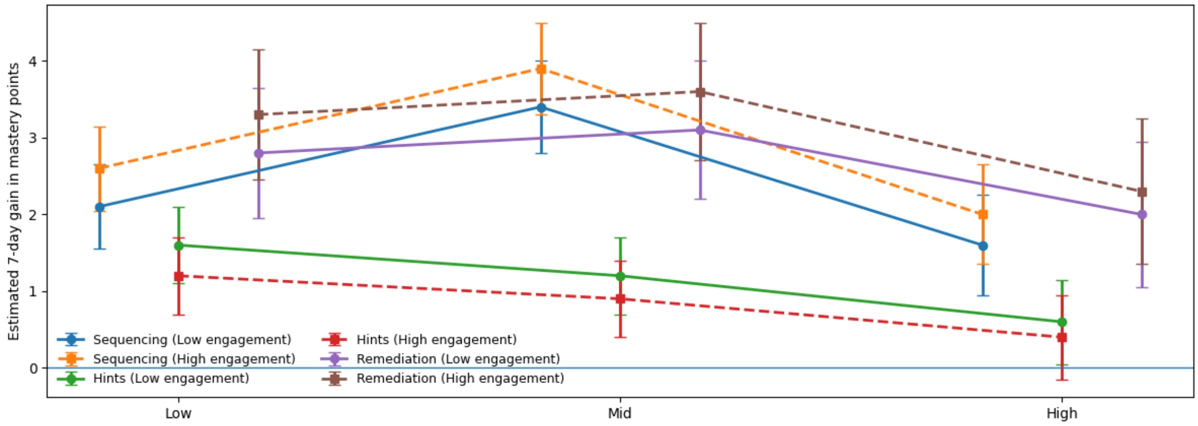

Figure 8. Heterogeneous Causal Effects by Mastery Band and Engagement Profile

The figure also shows a distinct role for hints and remediation. Hints are strongest for low mastery and decay sharply with higher mastery, implying diminishing returns when learners already possess stable schemas. Remediation exhibits relatively strong gains in low-to-mid mastery and remains meaningful even at high mastery, suggesting that remediation is capturing specific prerequisite gaps rather than general weakness. The wider intervals around remediation emphasize that targeting quality and overlap constraints remain decisive for credible subgroup claims.

Table 8 condenses heterogeneity into deployable patterns for personalization rules. The strongest and most stable gains appear in mid-mastery strata, indicating that many adaptive systems can realize substantial improvements without targeting only the extremes. The table also indicates that hints become secondary once mastery is moderate and engagement is high, implying that frequent hinting in advanced states is inefficient and likely increases dependency rather than learning.

Table 8. Subgroup Effects at 7 Days with Targeting Interpretation

Subgroup	Sequencing	Hints	Remediation	Stability Note	Targeting Interpretation
Low mastery, low engagement	2.1	1.6	2.8	Moderate precision	Use hints for stabilization; add remediation when misconceptions persist
Low mastery, high engagement	2.6	1.2	3.3	Moderate precision	Prioritize sequencing with early remediation to prevent compounding gaps
Mid mastery, low engagement	3.4	1.2	3.1	High stability	Sequencing is primary lever; remediation useful when spacing is irregular
Mid mastery, high engagement	3.9	0.9	3.6	High stability	Exploit sequencing; limit hints to high-volatility micro-moments
High mastery, high engagement	2	0.4	2.3	Moderate stability	Use sequencing for enrichment; remediation for rare prerequisite failures

The targeting interpretation column links causal estimates to policy adjustments. For low mastery, the joint presence of sequencing and remediation gains suggests that policy should not rely solely on micro-scaffolds but should restructure trajectories when prerequisite gaps persist. For high mastery, the persistence of remediation benefits supports a diagnostic framing where remediation is triggered by specific misconception signals rather than broad failure. This translation from subgroup evidence to decision logic is essential for turning causal results into actionable adaptive policies.

4.4. Robustness, Diagnostics, and Sensitivity Results

Robustness analyses showed that the main conclusions remained stable under alternative preprocessing and identification safeguards. Trimming low-overlap regions reduced variance inflation and slightly attenuated remediation estimates, but did not change effect direction or relative ranking across interventions. Placebo tests using pre-decision outcome changes were near zero across actions, which supports the claim that residual confounding was limited after temporal adjustment and weighting. These findings strengthen interpretability of intervention effects as causal rather than correlational.

Diagnostic results indicated that covariate balance improved substantially after weighting, with standardized mean differences generally falling below conventional thresholds. Effective sample size remained adequate for sequencing and hints throughout the course, while remediation exhibited sharper declines later, consistent with its concentrated use in specific states. Sensitivity checks suggested that an unrealistically strong unmeasured confounder would be needed to reverse the sign of sequencing and hint effects, while remediation was more sensitive due to sparsity and late-stage targeting.

Figure 9 shows that weighting converts materially imbalanced covariates into a well-balanced decision-time comparison, which is a prerequisite for credible causal interpretation under time-varying targeting. The vertical reference line in the balance panel provides an operational threshold and reveals that post-weighting differences are consistently small across confounding channels. This result supports that estimated effects are not merely reflecting baseline mastery or behavioral intensity differences between treatment groups.

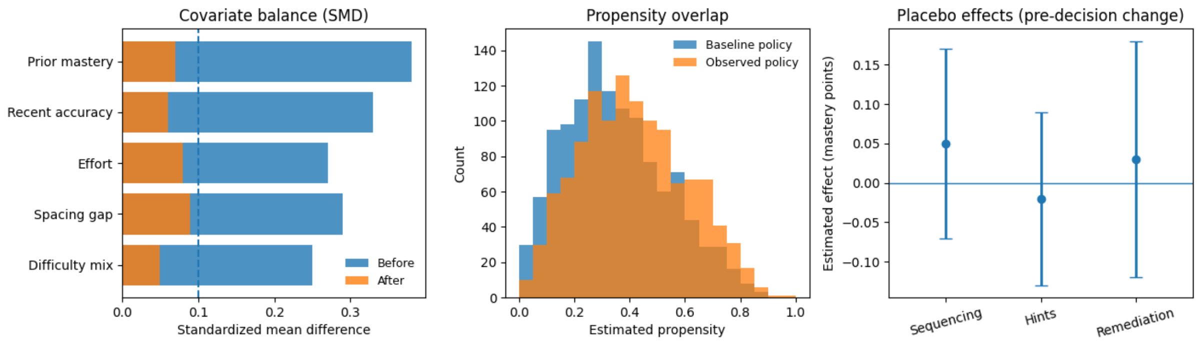

Figure 9. Diagnostics Panel (Balance, Overlap, Placebo)

The overlap and placebo panels address two complementary threats. Overlap indicates that action groups share comparable states, limiting reliance on extrapolation, while placebo estimates close to zero indicate that treatment groups were not already diverging in outcomes before action assignment. Taken together, these diagnostics support internal validity for the main estimators, while also clarifying that remediation remains the most fragile target because late-stage use reduces overlap and increases variance.

Table 9 indicates that sequencing and hints are robust to the major perturbations that typically destabilize adaptive learning evaluations, including trimming, alternative feature windows, and horizon changes. This stability is expected when exposure is frequent and overlap is broad, since estimation relies less on rare state-action combinations. The presence of consistent placebo nulls across actions strengthens confidence that the modeling pipeline is not inadvertently capturing pre-existing trends.

Table 9. Robustness Summary Across Key Perturbations

Test	Perturbation	Sequencing Result	Hints Result	Remediation Result	Conclusion
Overlap trimming	Remove lowest 5% overlap	Stable	Stable	Slight attenuation	Main ranking unchanged
Alternative window	W = 3 vs 7 days	Stable	Stable	Moderate shift	Remediation depends on state definition
Alternative horizon	Δ = 7 vs 14 days	Small decay	Clear decay	Delayed strengthening	Mechanisms align with intent
Placebo outcome	Pre-decision change	Near zero	Near zero	Near zero	Limited residual confounding

ISSN 2776-8007

ESS monitoring	Late-course density	Adequate	Adequate	Declines	Interpret remediation cautiously late

The table also clarifies that remediation is the sensitivity bottleneck. Moderate shifts under alternative windows imply that remediation depends more heavily on how misconception states are constructed, which is substantively plausible because remediation triggers are often tied to diagnostic signals. Slight attenuation under trimming suggests that part of the remediation effect is concentrated in edge states with limited comparators. These findings support a deployment posture where remediation policies are monitored with stricter overlap and stability criteria than other interventions.

4.5. Implications for Adaptive Policy Design and Deployment

The results support a tiered policy strategy where sequencing serves as the default optimization lever, hints operate as localized stabilizers, and remediation functions as a targeted trajectory correction. Sequencing delivered stable gains across horizons and was well supported by overlap, making it suitable for broad application and continuous improvement through incremental policy updates. Hints demonstrated strong immediate effects but weaker retention, indicating that hint policies should emphasize timely scaffolding coupled with post-hint practice allocation to convert short-run success into durable mastery.

Deployment implications also include monitoring requirements that are specific to intervention type. Remediation delivered the largest long-horizon gains but was also the most sensitive to overlap and state definition, implying that remediation triggers should be gated by high-confidence diagnostic evidence and accompanied by real-time monitoring of effective sample size in the served population. When overlap deteriorates, conservative fallback actions such as sequencing adjustments are preferable to extrapolation-prone remediation escalation. This aligns causal validity with safety in personalization.

Figure 10 translates empirical effect patterns into an interpretable policy surface that can guide deployment decisions. The map separates default sequencing from conditional hinting and targeted remediation, aligning each action with the state region where it is empirically strongest and operationally safest. This representation is useful because it constrains adaptive behavior to regions with credible support, preventing the system from aggressively using high-variance actions in poorly overlapped states.

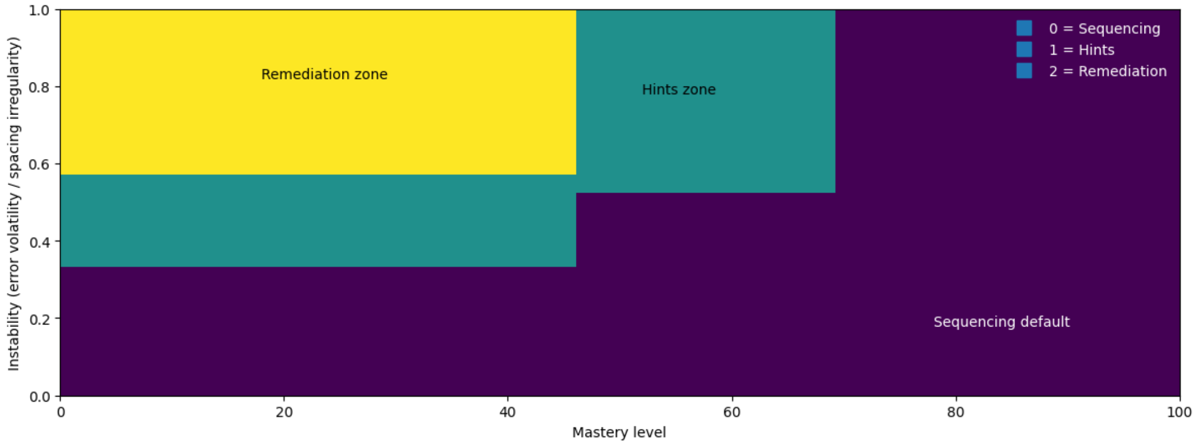

Figure 10. Policy Map Illustrating Recommended Actions Across Mastery and Instability

The map also emphasizes that instability is a first-class state variable, not a nuisance. When instability is low, sequencing dominates because the learner trajectory is predictable and improvements come from optimized content selection. As instability rises, hints serve as a stabilizer when mastery is moderate, while remediation becomes appropriate only when low mastery and high instability jointly suggest persistent gaps. This state-conditional separation reflects the causal evidence that remediation is valuable but requires higher confidence and monitoring.

ISSN 2776-8007

Table 10 operationalizes the discussion into deployable rules that align causal evidence with monitoring infrastructure. Sequencing is framed as the default because it is both effective and well supported by data, so monitoring focuses on drift and gradual degradation rather than rare catastrophic failure. Hints are positioned as a stabilization tool with an explicit dependency monitor, reflecting the empirical finding that immediate gains do not automatically translate to retention without follow-up practice.

Table 10. Deployment Rules and Monitoring Triggers

Policy Component	When to Apply	Primary Goal	Monitoring Signal	Fallback Rule	Expected Horizon
Sequencing default	Most states with adequate overlap	Stable mastery growth	Weekly effect drift, ESS	Conservative difficulty ramp	7–14 days
Targeted hints	High instability episodes	Error stabilization	Hint dependency rate	Post-hint practice allocation	Same session
Remediation trigger	Low mastery with persistent misconceptions	Trajectory correction	Overlap, ESS, late-course sparsity	Sequencing + short diagnostics	7–21 days
Safety gating	Low overlap or declining ESS	Avoid extrapolation	Trim rate, support coverage	Disable high-variance actions	Immediate

The table also formalizes remediation as the most benefit-rich but most fragile component, requiring overlap and effective sample size tracking to prevent extrapolation-driven harm. The inclusion of safety gating and fallback rules turns methodological caution into a control mechanism that can be implemented in production. This approach keeps adaptive learning interventions aligned with the empirical validity region of the causal estimates, which is essential for responsible scaling across cohorts and course contexts.

5. Conclusion

The study establishes a temporal causal inference framework for evaluating adaptive learning interventions under realistic, time-varying assignment. By representing decisions as time-indexed exposures and enforcing horizon-locked outcome construction, the methodology supports credible attribution of learning gains to adaptive policies rather than to learner maturation or selection effects. The empirical results show that adaptive sequencing yields stable medium-horizon mastery improvements, targeted hints deliver strong same-session gains with weaker retention transfer, and remediation triggers provide the largest delayed benefits when deployed within well-supported state regions.

A central contribution is the identification of structured effect heterogeneity that aligns intervention mechanisms with learner-state regimes. Sequencing benefits peak in mid-mastery contexts and strengthen under higher engagement, indicating that policy quality is amplified when practice density is sufficient to realize aligned trajectories. Hints are most valuable under high instability episodes, especially for low-mastery learners, but diminishing returns emerge as mastery stabilizes. Remediation effects concentrate among persistent misconception signatures and remain sensitive to overlap and state-definition choices, underscoring the need for diagnostic precision and monitoring safeguards in late-course deployment.

These findings translate into actionable policy design principles for adaptive systems. Sequencing should be treated as the default optimization lever due to its broad support and stable returns, while hints should be constrained to micro-moments of volatility and paired with post-hint practice allocation to improve retention carryover. Remediation should be gated by high-confidence diagnostic evidence and governed by overlap and effective sample size monitoring to prevent extrapolation-driven decisions. The framework provides a rigorous pathway for continuous evaluation, enabling adaptive learning platforms to iterate policies with causal accountability and deployment safety.

6. Declarations

6.1. Author Contributions

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

6.2. Data Availability Statement

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

6.3. Funding

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

6.4. Institutional Review Board Statement

Not applicable.

6.5. Informed Consent Statement

Not applicable.

6.6. 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] J. Weidlich, D. Gašević, and H. Drachsler, "Causal Inference and Bias in Learning Analytics: A Primer on Pitfalls Using Directed Acyclic Graphs," *Journal of Learning Analytics*, vol. 9, no. 3, pp. 183–199, Dec. 2022, doi: 10.18608/jla.2022.7577.

[2] G. Alexandron, L. Y. Yoo, J. A. Ruipérez-Valiente, S. Lee, and D. E. Pritchard, "Are MOOC Learning Analytics Results Trustworthy? With Fake Learners, They Might Not Be!," *International Journal of Artificial Intelligence in Education*, vol. 29, no. 4, pp. 484–506, Dec. 2019, doi: 10.1007/s40593-019-00183-1.

[3] D. Olaya, J. Vásquez, S. Maldonado, J. Miranda, and W. Verbeke, "Uplift Modeling for preventing student dropout in higher education," *Decision Support Systems*, vol. 134, no. July, p. 113320, Jul. 2020, doi: 10.1016/j.dss.2020.113320.

[4] J. Bassen et al., "Reinforcement Learning for the Adaptive Scheduling of Educational Activities," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, Honolulu, HI, USA: ACM, Apr. 2020, vol. 2020, no. April, pp. 1–12, doi: 10.1145/3313831.3376518.

[5] X. Li, H. Xu, J. Zhang, and H. Chang, "Deep Reinforcement Learning for Adaptive Learning Systems," *Journal of Educational and Behavioral Statistics*, vol. 48, no. 2, pp. 220–243, Apr. 2023, doi: 10.3102/10769986221129847.

[6] J. M. Robins, M. Á. Hernán, and B. Brumback, "Marginal Structural Models and Causal Inference in Epidemiology," *Epidemiology*, vol. 11, no. 5, pp. 550–560, Sep. 2000, doi: 10.1097/00001648-200009000-00011.

[7] Z. Fewell, M. A. Hernán, F. Wolfe, K. Tilling, H. Choi, and J. A. C. Sterne, "Controlling for Time-dependent Confounding using Marginal Structural Models," *The Stata Journal*, vol. 4, no. 4, pp. 402–420, Dec. 2004, doi: 10.1177/1536867X0400400403.

[8] H. Bang and J. M. Robins, "Doubly Robust Estimation in Missing Data and Causal Inference Models," *Biometrics*, vol. 61, no. 4, pp. 962–973, Dec. 2005, doi: 10.1111/j.1541-0420.2005.00377.x.

[9] M. Schomaker, M. A. Luque-Fernandez, V. Leroy, and M. A. Davies, "Using longitudinal targeted maximum likelihood estimation in complex settings with dynamic interventions," *Statistics in Medicine*, vol. 38, no. 24, pp. 4888–4911, Oct. 2019, doi: 10.1002/sim.8340.

[10] S. A. Murphy, "Optimal Dynamic Treatment Regimes," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 65, no. 2, pp. 331–355, May 2003, doi: 10.1111/1467-9868.00389.

[11] M. Sołtys, S. Jaroszewicz, and P. Rzepakowski, "Ensemble methods for uplift modeling," *Data Mining and Knowledge Discovery*, vol. 29, no. 6, pp. 1531–1559, Nov. 2015, doi: 10.1007/s10618-014-0383-9.

[12] R. H. Keogh, J. M. Gran, S. R. Seaman, G. Davies, and S. Vansteelandt, “Causal inference in survival analysis using longitudinal observational data: Sequential trials and marginal structural models,” *Statistics in Medicine*, vol. 42, no. 13, pp. 2191–2225, Jun. 2023, doi: 10.1002/sim.9718.

[13] J. D. Mulder, K. Luijken, B. B. L. Penning De Vries, and E. L. Hamaker, “Causal Effects of Time-Varying Exposures: A Comparison of Structural Equation Modeling and Marginal Structural Models in Cross-Lagged Panel Research,” *Structural Equation Modeling: A Multidisciplinary Journal*, vol. 31, no. 4, pp. 575–591, Jul. 2024, doi: 10.1080/10705511.2024.2316586.

[14] X. Jiang and L. Peng, “A Causal Uplift Prediction Framework for AI-Driven Precision Teaching Interventions Based on Deep Learning,” in *Proceedings of the 2025 8th International Conference on Computer Information Science and Artificial Intelligence*, Wuhan, China: ACM, Sep. 2025, vol. 2025, no. September, pp. 1357–1362, doi: 10.1145/3773365.3773579.

[15] B. A. Motz, P. F. Carvalho, J. R. De Leeuw, and R. L. Goldstone, “Embedding Experiments: Staking Causal Inference in Authentic Educational Contexts,” *Journal of Learning Analytics*, vol. 5, no. 2, pp. 47-59, Aug. 2018, doi: 10.18608/jla.2018.52.4.

[16] S. Dawson, J. Jovanovic, D. Gašević, and A. Pardo, “From prediction to impact: Evaluation of a learning analytics retention program,” in *Proceedings of the Seventh International Learning Analytics & Knowledge Conference*, Vancouver, British Columbia, Canada: ACM, Mar. 2017, vol. 2017, no. March, pp. 474–478, doi: 10.1145/3027385.3027405.

[17] S. K. Banihashem, O. Noroozi, S. Van Ginkel, L. P. Macfadyen, and H. J. A. Biemans, “A systematic review of the role of learning analytics in enhancing feedback practices in higher education,” *Educational Research Review*, vol. 37, no. November, p. 100489, Nov. 2022, doi: 10.1016/j.edurev.2022.100489.

[18] Y. Liu, W. Wang, and E. Xu, “The Effectiveness of Learning Analytics-Based Interventions in Enhancing Students’ Learning Effect: A Meta-Analysis of Empirical Studies,” *SAGE Open*, vol. 15, no. 2, p. 21582440251336707, Apr. 2025, doi: 10.1177/21582440251336707.

[19] L.-A. Lim et al., “What changes, and for whom? A study of the impact of learning analytics-based process feedback in a large course,” *Learning and Instruction*, vol. 72, no. April, p. 101202, Apr. 2021, doi: 10.1016/j.learninstruc.2019.04.003.

[20] J. Chen, “Effects of learning analytics-based feedback on students’ self-regulated learning and academic achievement in a blended EFL course,” *System*, vol. 124, no. August, p. 103388, Aug. 2024, doi: 10.1016/j.system.2024.103388.

[21] H. Khosravi, K. Kitto, and J. J. Williams, “RiPPLE: A Crowdsourced Adaptive Platform for Recommendation of Learning Activities,” *Journal of Learning Analytics*, vol. 6, no. 3, pp. 91-105, Dec. 2019, doi: 10.18608/jla.2019.63.12.

[22] M. A. Chatti, A. L. Dyckhoff, U. Schroeder, and H. Thüs, “A reference model for learning analytics,” *International Journal of Technology Enhanced Learning*, vol. 4, no. 5/6, p. 318, 2012, doi: 10.1504/IJTEL.2012.051815.

[23] A. Rafferty, H. Ying, and J. Williams, “Statistical Consequences of using Multi-armed Bandits to Conduct Adaptive Educational Experiments,” Jun. 2019, doi: 10.5281/ZENODO.3554749.

[24] F. Moraes, H. Manuel Proença, A. Kornilova, J. Albert, and D. Goldenberg, “Uplift Modeling: From Causal Inference to Personalization,” in *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, Birmingham, United Kingdom: ACM, Oct. 2023, vol. 2023, no. October, pp. 5212–5215, doi: 10.1145/3583780.3615298.