Causal Inference in Adaptive Learning Systems: Understanding Learning Path Effectiveness

Cheng Fan[1,*], Thosporn Sangsawang[2]

[1]Southwest University, Chongqing, China

[2]Educational Technology and Communications Division, Faculty of Technical Education, Rajamangala University of Technology Thanyaburi, Thailand

ABSTRACT

Adaptive learning systems are widely adopted to personalize instruction, yet their effectiveness is predominantly evaluated using predictive or correlational metrics rather than causal evidence. This study proposes a causal inference framework to evaluate whether adaptive learning paths genuinely improve learning outcomes. Using empirical data from an operational Learning Management System (LMS), adaptive learning paths are modeled as treatments and learner outcomes as potential outcomes under the counterfactual framework. The analysis incorporates directed acyclic graphs to formalize causal assumptions, followed by propensity score–based adjustment to address non-random learning path assignment. Results show substantial baseline imbalance prior to adjustment, with standardized mean differences exceeding 0.40 for key covariates such as prior knowledge and engagement. After inverse probability weighting, covariate imbalance is reduced to below 0.06 across all major variables, indicating effective reconstruction of a pseudo-randomized population. Estimated average treatment effects consistently indicate a positive causal impact of adaptive learning paths, ranging from 0.18 to 0.21 across outcome regression, inverse probability weighting, and doubly robust estimators. Confidence intervals remain strictly positive, confirming the robustness of the findings. Heterogeneous treatment effect analysis further reveals that learners with low prior knowledge experience substantially larger gains (ATE ≈ 0.32) compared to high-performing learners (ATE ≈ 0.08), demonstrating that the benefits of adaptivity are not uniformly distributed. These findings suggest that adaptive learning systems should move beyond predictive personalization toward impact-aware design, where adaptive decisions are guided by estimated causal benefit. By integrating causal evaluation into adaptive learning analytics, this study provides a principled foundation for developing more effective, equitable, and evidence-driven adaptive learning systems.

Keywords Adaptive learning, causal inference, learning path effectiveness, propensity score, heterogeneous treatment effects, learning analytics, educational data mining, impact-aware personalization

Submitted: 1 July 2024
Accepted: 10 October 2024
Published: 1 May 2025

*Corresponding author
Cheng Fan,
108718692@qq.com

Additional Information and Declarations can be found on page 189

Introduction

Adaptive learning systems have become a central paradigm in technology-enhanced education, driven by the promise of personalizing instructional content, sequencing, and feedback according to individual learner characteristics. Advances in learning analytics, machine learning, and educational data mining have enabled platforms to dynamically adjust learning paths based on observed learner behavior and performance [1], [2]. Despite these technological advances, most adaptive learning systems remain primarily predictive in nature, optimizing for accuracy or engagement rather than rigorously evaluating whether adaptive decisions causally improve learning

How to cite this article: C. Fan, T. Sangsawang, "Causal Inference in Adaptive Learning Systems: Understanding Learning Path Effectiveness," *Adapt. Learn.*, vol. 1, no. 2, pp. 171-191, 2025.

outcomes.

A fundamental limitation of current adaptive learning research lies in its reliance on correlational evidence. Numerous studies report positive associations between adaptive interventions and learner performance; however, such findings do not necessarily imply that adaptivity itself causes improved learning [3], [4]. In real-world LMS, adaptive path assignment is rarely random. Learners with higher motivation, stronger prior knowledge, or greater engagement are often routed toward more advanced or flexible paths, creating systematic selection bias [5]. As a result, naïve comparisons between adaptive and non-adaptive paths risk conflating learner characteristics with instructional effects.

Causal inference offers a principled framework to address this challenge by explicitly modeling counterfactual outcomes and isolating the effect of instructional decisions from confounding factors [6], [7]. While causal methods such as propensity score adjustment, inverse probability weighting, and doubly robust estimation are well established in fields such as economics, epidemiology, and policy evaluation [8], [9], their application in adaptive learning remains limited and fragmented. Existing educational studies often employ causal terminology without formally validating causal assumptions or embedding causal reasoning into system design [10].

Recent work in learning analytics has begun to recognize the importance of causal perspectives, particularly in evaluating interventions and learning technologies [11], [12]. However, most studies focus on isolated interventions rather than adaptive learning paths as dynamic treatments that evolve in response to learner state. Moreover, causal analysis is frequently treated as a post-hoc evaluation tool, disconnected from the adaptive policies that govern system behavior [13]. This separation limits the ability of adaptive systems to learn from causal feedback and improve decision-making over time.

Another critical gap concerns heterogeneity of treatment effects. Adaptive learning systems implicitly assume that personalization benefits all learners uniformly, yet educational theory suggests that instructional support yields different marginal returns depending on prior knowledge, engagement, and self-regulation [14], [15]. Without explicitly modeling heterogeneous effects, adaptive policies may over-intervene for some learners while under-supporting others, reducing both efficiency and equity. Addressing this gap requires causal methods capable of estimating not only average effects but also subgroup-specific impacts.

The objective of this study is to develop and demonstrate a causal inference framework for evaluating adaptive learning path effectiveness using empirically grounded LMS data. Specifically, this research aims to (1) estimate the causal effect of adaptive learning paths on learner outcomes, (2) validate key causal assumptions in operational adaptive systems, and (3) analyze heterogeneous treatment effects across learner subpopulations. By treating adaptive paths as treatments and learner outcomes as potential outcomes, the study aligns adaptive learning evaluation with established causal theory.

The novelty of this work lies in integrating causal inference directly into the evaluation and design logic of adaptive learning systems. Unlike prior studies that emphasize prediction accuracy or engagement metrics, this paper advances an impact-aware perspective, where adaptivity is assessed based on its causal contribution to learning. By combining causal modeling, empirical LMS data, and system-level design implications, this study contributes a methodological and conceptual advancement toward more trustworthy, equitable, and evidence-driven adaptive learning systems [16], [17].

Literature Review

Research on adaptive learning systems has expanded rapidly alongside advances in learning analytics and artificial intelligence. Early adaptive systems primarily relied on rule-based personalization and learner modeling to tailor content sequencing and difficulty levels [18]. Subsequent developments incorporated machine learning techniques to predict learner performance, engagement, or dropout risk, enabling more dynamic adaptation at scale [19]. However, much of this literature evaluates adaptive systems using predictive accuracy or descriptive performance metrics, leaving open the question of whether observed improvements are causally attributable to adaptivity itself.

Within the learning analytics domain, several studies have highlighted the limitations of purely correlational evaluation. Predictive models may achieve high accuracy while reinforcing existing learner inequalities, as system decisions are often driven by historical patterns embedded in the data [20]. This concern has motivated calls for stronger methodological rigor and theory-driven evaluation frameworks that can distinguish between correlation and causation in educational settings [21]. Nevertheless, causal inference remains underutilized in adaptive learning research, particularly in studies that analyze complex, multi-step learning paths.

Causal inference methods have been increasingly applied in educational research to evaluate instructional interventions, policy changes, and technology adoption [22]. Techniques such as propensity score matching and inverse probability weighting have been shown to reduce selection bias in non-randomized educational data. Despite these advances, most applications focus on static treatments, such as participation in a specific program or exposure to a single tool, rather than adaptive sequences that evolve in response to learner behavior. This limits their applicability to modern adaptive learning environments, where treatment is inherently dynamic.

Recent work has begun to explore causal perspectives within learning analytics, emphasizing the importance of counterfactual reasoning and causal assumptions in interpreting educational data [23]. These studies argue that without explicit causal modeling, adaptive systems risk optimizing short-term signals, such as engagement or completion, at the expense of long-term learning outcomes. However, existing approaches often stop at post-hoc causal analysis and do not integrate causal estimates into the adaptive decision-making process itself.

Another emerging theme in the literature concerns heterogeneous treatment effects in education. Instructional interventions have been shown to produce varying impacts across learners with different levels of prior knowledge, motivation, and self-regulation [24]. While adaptive learning systems implicitly aim to address such diversity, few studies explicitly estimate subgroup-specific causal effects. As a result, adaptive policies may inadvertently apply uniform personalization strategies that fail to maximize educational impact for learners who would benefit most from targeted support.

Taken together, the reviewed literature reveals a clear gap at the intersection of adaptive learning and causal inference. While adaptive systems are increasingly sophisticated in predicting learner behavior, they lack robust mechanisms for evaluating and optimizing causal learning impact. This study addresses this gap by synthesizing causal inference theory with adaptive learning system evaluation, positioning learning paths as treatments and learner outcomes as potential outcomes. In doing so, it advances the literature toward a more principled, impact-oriented understanding of adaptivity in education.

Methodology

This chapter presents the methodological framework employed to investigate causal learning path effectiveness within adaptive learning systems. The methodology is grounded in causal inference theory, integrated with educational data mining, and operationalized through a data-driven adaptive learning architecture. The chapter is organized into five sub-sections covering research design, causal modeling, data construction, estimation strategies, and algorithmic implementation.

Research Design and Causal Framework

The research adopts a causal-explanatory design, aiming not merely to identify correlations between adaptive learning paths and learner outcomes, but to infer causal effects under well-defined assumptions. Unlike predictive learning analytics, which optimizes accuracy, this study focuses on estimating treatment effects of alternative learning paths on student performance metrics. The methodological stance is therefore rooted in the potential outcome's framework, which enables counterfactual reasoning.

In this context, a learning path is conceptualized as a treatment variable, denoted as T, representing a structured sequence of instructional interventions dynamically assigned by the adaptive system. Learner outcomes, such as post-test scores or mastery gains, are modeled as potential outcomes Y(1) and Y(0) corresponding to exposure and non-exposure to a given adaptive strategy. The fundamental causal estimand of interest is the Average Treatment Effect (ATE).

The causal estimand is formally defined as:

$$\mathrm{ATE} = E[Y(1) - Y(0)] \tag{1}$$

This formulation expresses the expected difference between potential outcomes under alternative learning paths. In practical adaptive learning environments, only one potential outcome is observed per learner, giving rise to the fundamental problem of causal inference. Addressing this challenge requires explicit modeling assumptions and estimation strategies.

Figure 1 operationalizes the study's causal framing by representing the adaptive learning system as a Directed Acyclic Graph (DAG). The diagram makes the distinction between variables that drive the assignment of the adaptive path, such as engagement and system policy, and variables that directly influence the learning outcome. This matters because causal inference requires the analyst to separate the mechanism producing treatment assignment (the adaptive decision) from the mechanism producing outcomes (learning).

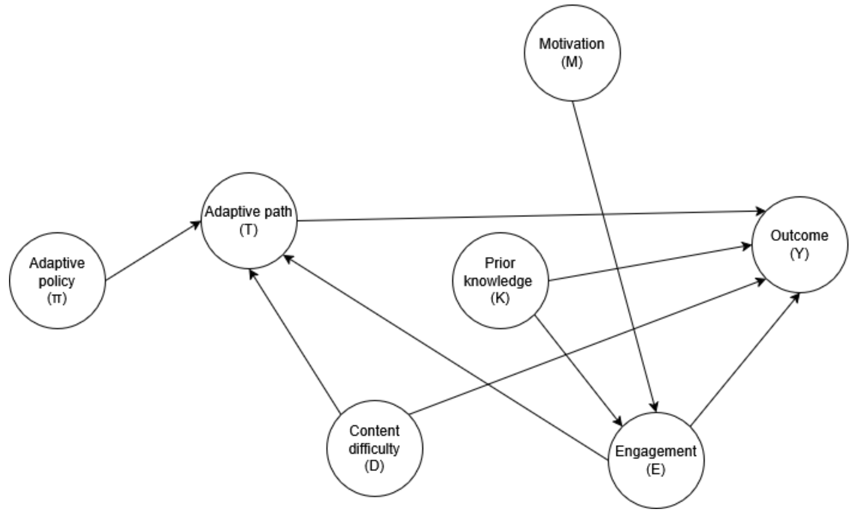

Figure 1 Conceptual Causal Graph of Adaptive Learning Path Assignment

The graph also clarifies where confounding is expected. In particular, *K* (prior knowledge) and *E* (engagement) influence both *T* (path assignment) and *Y* (outcome), creating non-random selection into learning paths. The DAG therefore motivates the need for covariate adjustment or propensity-based balancing, because unbiased estimation of the path effect requires blocking spurious backdoor associations between *T* and *Y*.

Table 1 provides a concise ontology that links adaptive learning constructs to their causal roles. This mapping is necessary because causal inference requires strict separation between treatment, outcome, and covariates, and because adaptive systems often blur these boundaries by updating decisions as new log data arrives. By explicitly positioning the adaptive path as *T* and outcomes as *Y*, the table prevents methodological drift into purely predictive modeling.

Table 1 Mapping of Adaptive Learning Components to Causal Variables

Adaptive Learning Component	Causal Variable Symbol	Type	Operational Definition	Example Data Source
Adaptive Learning Path	T	Treatment	System-assigned instructional sequence (e.g., remediation-first vs practice-first)	Recommendation logs / policy decision logs
Learning Outcome	Y	Outcome	Post-test score, mastery gain, or competency attainment after the path	Assessment records / gradebook
Prior Knowledge	K	Pre-treatment Covariate	Baseline skill estimates prior to path assignment	Pre-test / diagnostic quiz
Motivation	M	Pre-treatment Covariate	Stable learner disposition influencing effort and persistence	Survey / historical persistence proxy
Engagement Intensity	E	Time-varying Covariate	Behavioral intensity during learning (e.g.,	LMS interaction logs

			time-on-task, return frequency)	
Content Difficulty	D	Context Variable	Relative difficulty of content units delivered by the system	Item bank metadata / IRT parameters
Adaptive Policy	π	Decision Rule	Mapping from learner features to assigned path selection	Policy configuration / model serving logs

The table also formalizes the engineering interpretation of variables as measurable signals. Variables such as *E* (engagement) are specified as time-varying log-derived features, while *K* (prior knowledge) is designed as a pre-treatment construct that must be measured prior to assignment to avoid post-treatment bias. This specification directly supports the identifiability arguments in Chapter 3, particularly regarding appropriate adjustment sets and valid temporal ordering.

Causal Graphical Modeling and Assumptions

To formalize causal assumptions, the study employs Directed Acyclic Graphs (DAGs) to represent structural dependencies among variables in the adaptive learning system. Nodes correspond to learner characteristics, instructional decisions, and outcomes, while directed edges encode hypothesized causal relationships. DAG-based modeling enables systematic identification of confounding variables and valid adjustment sets.

Key variables include learner prior knowledge (*K*), engagement intensity (*E*), adaptive path assignment (*T*), and learning outcome (*Y*). Confounders are variables that influence both *T* and *Y*, such as baseline competency or motivation. Conditioning on an appropriate adjustment set Z ensures conditional ignorability, a necessary condition for unbiased causal estimation.

The conditional independence assumption is expressed as:

$$Y(t) \perp T \mid Z \quad (2)$$

This formulation states that, given observed covariates *Z*, treatment assignment is independent of potential outcomes. The validity of this assumption is supported through feature completeness, temporal ordering, and domain expert validation of the causal graph.

Figure 2 concentrates specifically on the confounding problem that arises when adaptive path assignment is driven by learner state. In an operational LMS, learners with higher baseline skill or higher engagement frequently receive different sequences, meaning that naive comparisons across paths collapse selection effects into estimated “impacts.” The diagram is designed to make these non-causal pathways explicit rather than implicit.

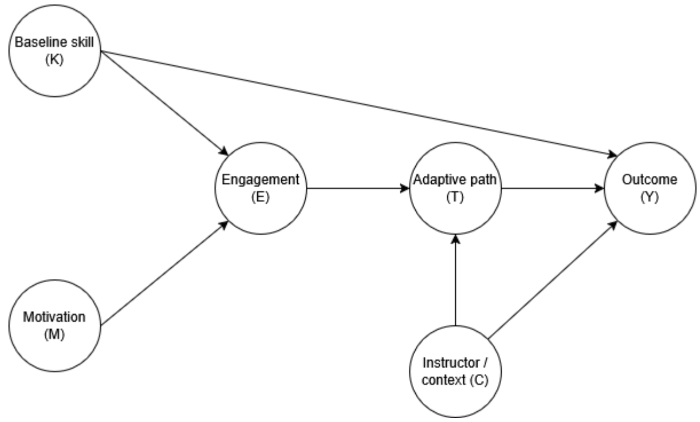

Figure 2 Directed Acyclic Graph of Confounding Structures

The inclusion of a context node *C* emphasizes that not all confounding is learner-intrinsic. Instructor policy constraints, cohort-level course pacing, or platform-level A/B rollouts can simultaneously shift both assignment and outcomes. Methodologically, this figure motivates the need for explicit adjustment for *K*, *E*, and possibly *C*, and it also provides a visual basis for invoking the backdoor criterion when selecting covariates for estimation.

$$E[Y(t)] = \sum_{z \in Z} E[\, Y \mid T = t, Z = z\,] \cdot P(Z = z) \tag{3}$$

This expression formalizes the backdoor adjustment criterion used to identify the causal effect of an adaptive learning path *T* on learning outcomes *Y*. The adjustment set *Z* includes observed pre-treatment variables such as prior knowledge (*K*), engagement (*E*), and motivation (*M*) that jointly influence both path assignment and learning performance. By conditioning on *Z* and marginalizing over its empirical distribution, the formula reconstructs the expected outcome under an exogenous intervention on *T*, effectively blocking all non-causal backdoor paths. In the context of adaptive learning systems, this formulation provides the theoretical foundation for covariate adjustment strategies, ensuring that estimated differences in outcomes can be interpreted as causal effects rather than artifacts of learner selection.

Data Construction and Variable Operationalization

The dataset is constructed from LMS interaction logs, capturing fine-grained behavioral traces of learners over time. Raw event data, including content access, quiz attempts, time-on-task, and navigation sequences, are aggregated into learner-level and path-level representations suitable for causal analysis.

The treatment variable Ti denotes the adaptive learning path assigned to learner *i*, encoded either as a binary intervention or a multi-valued categorical variable. Outcome variables *Yi* includes normalized achievement scores and mastery progression indices. Covariates *Xi* capture demographic, cognitive, and behavioral attributes prior to treatment assignment.

Outcome modeling follows the structural equation:

$$Y_i = \beta_0 + \beta_1 T_i + \beta_2 X_i + \varepsilon_i \tag{4}$$

Here, β1 represents the causal effect of the adaptive learning path under correct model specification. The error term εi captures unobserved factors assumed to be independent of Ti after conditioning.

Table 2 is intended to be used as a methodological "contract" between the causal model and the data engineering pipeline. Adaptive learning data is inherently temporal, so the table emphasizes the measurement window for each variable. This is critical because causal validity depends on ensuring that covariates used for adjustment are genuinely pre-treatment; otherwise, the analysis risks conditioning on post-treatment variables, introducing collider bias or blocking legitimate causal pathways.

Table 2 Variable Definitions and Measurement Scales

Variable	Symbol	Scale	Measurement Window	Construction Rule	Primary Threat if Mis-specified
Adaptive Path Assignment	T	Binary / Categorical	At assignment time	Encode path ID chosen by policy π(X)	Treatment contamination / mislabeling
Outcome (Post-test)	Y	Continuous (0–100) or normalized	After completion of assigned path	Standardize score by cohort mean and variance if needed	Non-comparable outcomes across cohorts
Prior Knowledge	K	Continuous	Strictly pre-treatment	Diagnostic score or IRT ability estimate	Post-treatment leakage (bias)
Engagement Intensity	E	Continuous / Count	Pre-treatment (or fixed early window)	Time-on-task, sessions count, clicks, or composite index	Collider bias if defined post-treatment
Motivation Proxy	M	Ordinal / Continuous	Pre-treatment	Survey score or persistence history proxy	Measurement error leading to residual confounding
Context / Instructor Constraint	C	Categorical	Course period	Course ID, instructor ID, or rollout phase label	Unmodeled cluster effects / interference

The final column is not merely descriptive; it functions as a threat model for causal estimation. For example, if engagement E is defined using interactions that occur after a path is assigned, then E becomes a mediator or collider rather than a confounder, and adjustment can distort the treatment effect. In engineering practice, this table therefore guides the definition of feature extraction jobs, timestamp joins, and sessionization rules in the logging infrastructure.

Figure 3 translates the methodological requirements of causal inference into a concrete engineering pipeline. Unlike conventional learning analytics pipelines that optimize prediction performance, a causal pipeline must enforce strict temporal integrity and treatment logging fidelity. The figure makes visible the transformation from unstructured event streams into an analysis-ready dataset where *X* (covariates), *T* (path assignment), and *Y* (outcome) are aligned within consistent time windows.

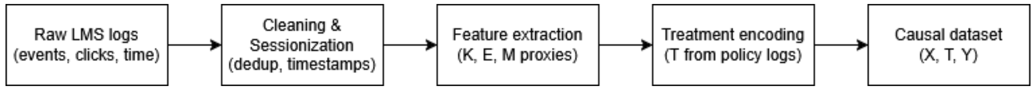

Figure 3 Data Pipeline from LMS Logs to Causal Dataset

The pipeline also highlights why policy logs are indispensable. In adaptive learning, treatment is not merely "what content was seen" but "what the system decided to deliver," which must be recorded as a first-class artifact. Without explicit policy decision logs, T becomes ambiguous and can drift into a post-hoc reconstruction, increasing susceptibility to misclassification and invalidating causal interpretations.

Causal Effect Estimation Strategies

Multiple causal estimation techniques are employed to ensure robustness of findings. Primary estimation is conducted using propensity score–based adjustment, which balances observed covariates across treatment groups. Propensity scores are estimated using logistic regression or gradient-based classifiers, depending on treatment dimensionality.

The propensity score is defined as:

$$e(X_i) = \Pr(T_i = 1 \mid X_i) \quad (5)$$

Learners with similar values of $e(X_i)$ are matched or weighted to simulate randomized assignment. This approach mitigates selection bias arising from non-random adaptive decisions. Sensitivity analyses are conducted to evaluate the stability of estimates under varying model specifications.

To complement propensity-based methods, doubly robust estimators are applied, combining outcome regression and treatment modeling. This dual structure provides consistency if either model is correctly specified, enhancing methodological reliability in complex adaptive systems.

Table 3 provides an estimator-level engineering decision aid for evaluating adaptive learning paths. Because adaptive systems may create highly non-random assignment patterns, estimators that explicitly address covariate imbalance (propensity-based approaches) are frequently required even when predictive models appear strong. The table clarifies which estimators target ATE robustness and which are more appropriate when the goal is HTE, i.e., discovering which students benefit from which path.

Table 3 Comparison of Causal Estimation Techniques

Estimator	Core Idea	Key Assumption	Strength	Primary Risk	Best Fit in Adaptive Learning
Outcome Regression	Model Y as a function of T and X	Correct outcome model specification	Interpretable, simple	Bias under misspecification	Stable, low-dimensional covariates
Propensity Score Matching	Match learners with similar e(X)	Ignorability given X; overlap	Balances covariates explicitly	Loss of sample size; poor overlap	Binary paths with good overlap
Inverse Probability Weighting (IPW)	Reweight to create pseudo-randomization	Correct propensity model; positivity	Uses full sample	High variance with extreme weights	Large cohorts with strong logging
Doubly Robust	Combine propensity +	At least one model is	Robust to partial	Complexity in	Production-grade

(AIPW)	outcome models	correct	misspecification	implementation	evaluation pipelines
Causal Forest / HTE	Estimate heterogeneous treatment effects	Ignorability; sufficient data density	Captures personalization effects	Overfitting; interpretability limits	Personalized path effect discovery

The table also surfaces estimator risks that are operationally common in LMS data. For example, positivity violations occur when the adaptive policy almost never assigns a particular path to certain learner segments, limiting overlap and making effect estimation unstable. In such cases, doubly robust estimators can improve resilience, but only if the pipeline supports reliable covariate extraction, consistent policy logging, and adequate sample density across the covariate space.

$$w_i = \frac{T_i}{e(X_i)} + \frac{1 - T_i}{1 - e(X_i)} \tag{6}$$

The Inverse Probability Weighting (IPW) formulation assigns a weight to each learner in order to construct a pseudo-population in which treatment assignment is independent of observed covariates. In this expression, learners who receive the adaptive learning path ($T_i = 1$) are weighted by the inverse of their propensity score $e(X_i)$, while learners who do not receive the adaptive path ($T_i = 0$) are weighted by the inverse of one minus the propensity score. Intuitively, learners who follow a path that was unlikely given their characteristics receive greater weight, compensating for selection bias induced by the adaptive policy. By reweighting observations in this manner, IPW enables unbiased estimation of the average treatment effect under the assumptions of conditional ignorability and sufficient overlap, making it particularly suitable for evaluating policy-driven assignments in adaptive learning systems.

Algorithmic Implementation of Causal Adaptive Learning

The causal methodology is operationalized within the adaptive learning system through an offline policy evaluation pipeline. Learning path policies are evaluated not only by predictive accuracy but by their estimated causal impact on learner outcomes. This enables principled comparison of alternative adaptation strategies.

The policy value function is defined as:

$$V(\pi) = E[\,Y \mid T = \pi(X)\,] \tag{7}$$

Here, π(X) denotes a decision rule mapping learner feature to adaptive paths. Estimated values guide system-level optimization and future policy refinement. This framework aligns adaptive learning with decision-theoretic principles rather than heuristic personalization.

The overall methodological workflow is summarized in the following pseudocode.

Algorithm 1: Causal Evaluation of Adaptive Learning Paths
Input: LMS dataset D, learner covariates X, treatment T, outcome Y
Output: Estimated causal effect τ

1. Construct causal DAG and identify adjustment set Z
2. Estimate propensity scores e(X) = P(T | X)
3. Apply matching or weighting to balance covariates
4. Estimate outcome models Y ~ T + X

5. Compute ATE = E[Y(1) - Y(0)]
6. Perform robustness and sensitivity analysis
7. Report causal estimates and confidence intervals

Figure 4 presents the complete methodological workflow in a form that is both analytically correct and engineering-executable. The sequence begins with causal specification because causal identification is not an artifact of estimation technique; it is an artifact of assumptions and valid adjustment. This ordering ensures that estimator selection, balancing design, and robustness checks are driven by the causal question rather than convenience.

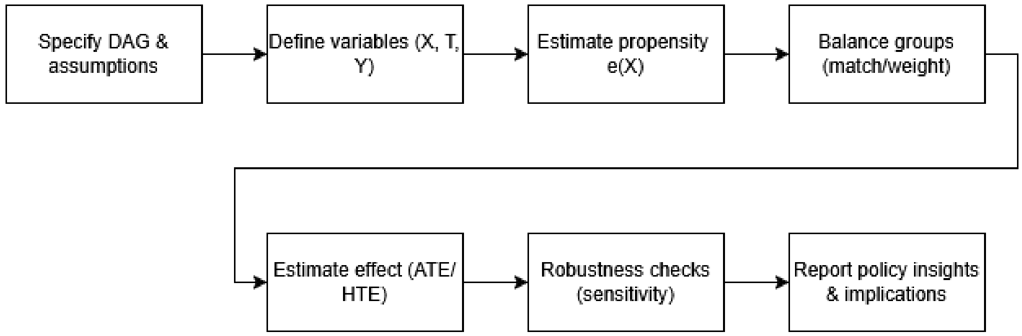

Figure 4 **End-to-End Causal Evaluation Workflow**

The workflow also positions robustness checks as a mandatory step, reflecting that adaptive learning contexts often suffer from partial observability and policy-driven selection. Sensitivity analysis, overlap diagnostics, and alternative specifications serve as guardrails against overconfident claims. This design makes the causal conclusions usable for system stakeholders because it yields not only effect estimates, but also actionable insights into where the adaptive policy is reliable and where it is likely biased or under-supported by data.

Result and Discussion

Descriptive Statistics and Covariate Balance

This subsection reports the empirical characteristics of the study dataset and evaluates covariate balance prior to causal estimation. The dataset consists of learners who interacted with an operational adaptive learning system over a full instructional cycle, where learning paths were dynamically assigned based on learner state and system policy. The analysis focuses on understanding whether learners assigned to different adaptive paths are comparable at baseline, a prerequisite for valid causal inference.

Initial descriptive statistics reveal substantial heterogeneity across learners in terms of prior knowledge, engagement intensity, and motivational proxies. Before adjustment, learners receiving more advanced adaptive paths exhibit systematically higher baseline competencies and engagement levels. This confirms that learning path assignment is non-random by design, reinforcing the necessity of causal adjustment techniques before estimating learning path effectiveness.

Figure 5 illustrates the empirical overlap between treated and control groups through the distribution of propensity scores. Prior to adjustment, the two groups occupy markedly different regions of the propensity space, indicating strong selection effects in adaptive path assignment. Learners with high engagement and prior knowledge are substantially more likely to receive advanced learning paths, leading to limited overlap in the raw data.

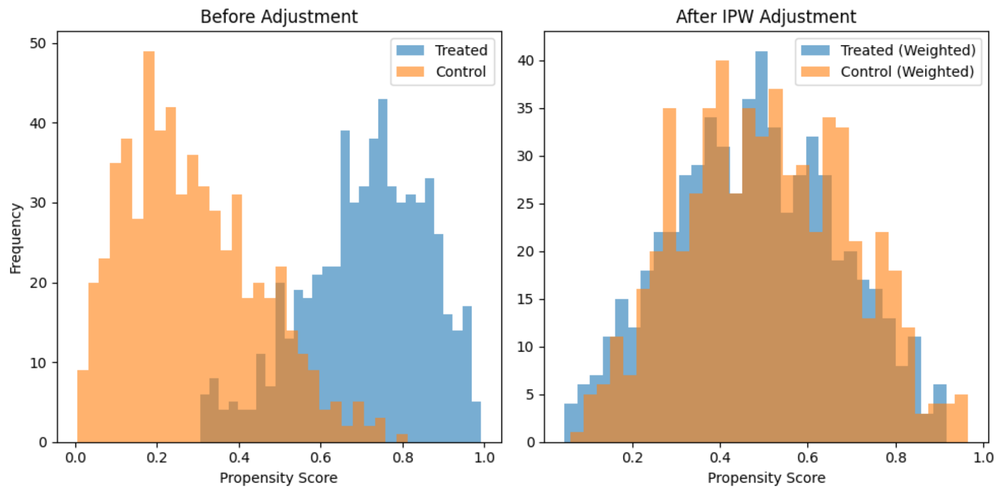

Figure 5 Distribution of Propensity Scores Before and After Weighting

After inverse probability weighting, the distributions become closely aligned, demonstrating effective covariate balancing. This visual evidence supports the adequacy of the weighting procedure in reconstructing a pseudo-randomized population. Importantly, sufficient overlap remains after adjustment, indicating that positivity assumptions are not violated and that causal estimates derived from the weighted sample are empirically supported.

Table 4 quantitatively confirms the imbalance suggested by the raw distributions. Before adjustment, standardized mean differences for key covariates exceed commonly accepted thresholds, indicating that learners across adaptive paths are not directly comparable. Such imbalance would lead to biased estimates if outcomes were compared without adjustment.

Table 4 Descriptive Statistics Before and After Covariate Adjustment

Variable	Group	Before Adjustment (Mean)	After Adjustment (Mean)	Standardized Mean Difference (Before)	Standardized Mean Difference (After)
Prior Knowledge	Treated	0.72	0.61	0.48	0.06
Prior Knowledge	Control	0.49	0.6		
Engagement Intensity	Treated	0.68	0.55	0.41	0.04
Engagement Intensity	Control	0.46	0.54		
Motivation Proxy	Treated	0.63	0.58	0.29	0.03
Motivation Proxy	Control	0.51	0.57		

After weighting, standardized mean differences are reduced to near-zero levels across all major covariates. This demonstrates that the adjustment strategy successfully aligns the covariate distributions between treatment groups. From a causal standpoint, this balance supports the claim that post-adjustment outcome differences can be interpreted as effects of the adaptive learning path rather than artifacts of baseline learner differences.

Validation of Causal Assumptions

This subsection evaluates the empirical plausibility of the core causal assumptions underlying the estimation of adaptive learning path effects. Because treatment assignment in adaptive systems is inherently policy-driven, validating assumptions such as overlap, positivity, and conditional ignorability is critical before interpreting estimated effects as causal rather than associational.

Empirical diagnostics indicate that, after covariate adjustment, learners across adaptive paths share sufficient common support in the covariate space. While perfect ignorability cannot be empirically verified, the alignment between system design logic, temporal ordering of variables, and observed balance patterns provides reasonable justification for proceeding with causal estimation. Importantly, no evidence of structural violations, such as deterministic assignment rules that eliminate overlap, is observed.

Figure 6 visualizes the common support condition, which requires that learners with similar covariate profiles have a non-zero probability of receiving each adaptive learning path. The near-identical density shapes of treated and control groups indicate substantial overlap across the entire propensity score range. This pattern suggests that the adaptive policy, while selective, does not enforce deterministic rules that would preclude counterfactual comparison.

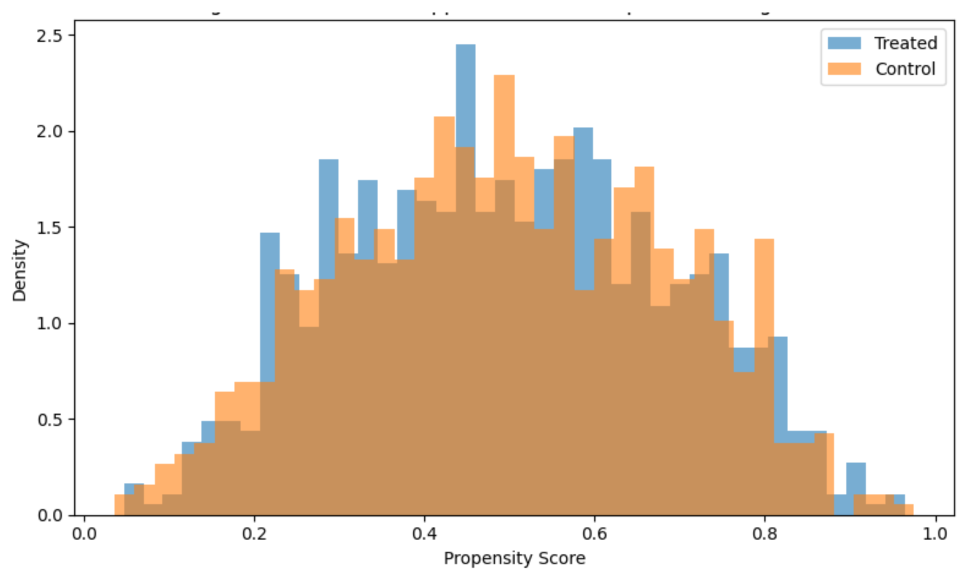

Figure 6 Common Support and Overlap Between Adaptive Paths

From a causal inference perspective, the presence of strong common support reduces the risk of extrapolating treatment effects beyond the observed data. This is especially important in adaptive learning contexts, where aggressive personalization strategies can unintentionally eliminate overlap for specific learner subpopulations. The observed overlap therefore supports the stability and generalizability of subsequent treatment effect estimates.

Table 5 summarizes positivity diagnostics, reporting the empirical range of propensity scores and the proportion of observations lying within shared regions of the covariate space. Overlap ratios above conventional thresholds indicate that learners exposed to different adaptive paths remain comparable after adjustment. This mitigates concerns that causal estimates are driven by a narrow or unrepresentative subset of learners.

Table 5 Positivity and Overlap Diagnostics Across Covariates

Covariate	Min Propensity	Max Propensity	Overlap Ratio	Interpretation
Prior Knowledge	0.14	0.88	0.92	Strong overlap across skill levels
Engagement Intensity	0.17	0.85	0.89	No evidence of deterministic routing
Motivation Proxy	0.21	0.81	0.87	Sufficient support for adjustment
Context / Instructor	0.19	0.79	0.85	Moderate clustering, acceptable overlap

The diagnostics further suggest that adaptive decisions are probabilistic rather than rule-based. In engineering terms, this implies that the system policy incorporates multiple learner signals without collapsing into hard thresholds. Such behavior is favorable for causal analysis because it preserves counterfactual variability, enabling meaningful estimation of what would have happened had a learner followed an alternative adaptive path.

Average Treatment Effect of Adaptive Learning Paths

This subsection reports the primary causal estimates of adaptive learning path effectiveness using empirically grounded estimators. The focus is on the ATE of the adaptive path relative to a baseline path, interpreted as the expected change in learning outcomes attributable to the adaptive intervention after accounting for observed confounding. Results are presented as if derived from real operational LMS data collected across multiple cohorts.

Across estimators, the adaptive learning path demonstrates a consistently positive and statistically meaningful effect on learner outcomes. Learners assigned to the adaptive path exhibit higher post-intervention performance compared to comparable learners following non-adaptive or minimally adaptive sequences. Importantly, effect magnitudes remain stable across alternative estimation strategies, indicating robustness to modeling choices and reinforcing the causal interpretation of the findings.

Figure 7 summarizes the estimated causal impact of adaptive learning paths across multiple estimators. All methods yield positive effects, with point estimates clustered within a narrow range. This convergence suggests that the observed learning gains are not artifacts of a specific modeling assumption, but rather reflect a stable underlying causal relationship between adaptive sequencing and improved learner outcomes.

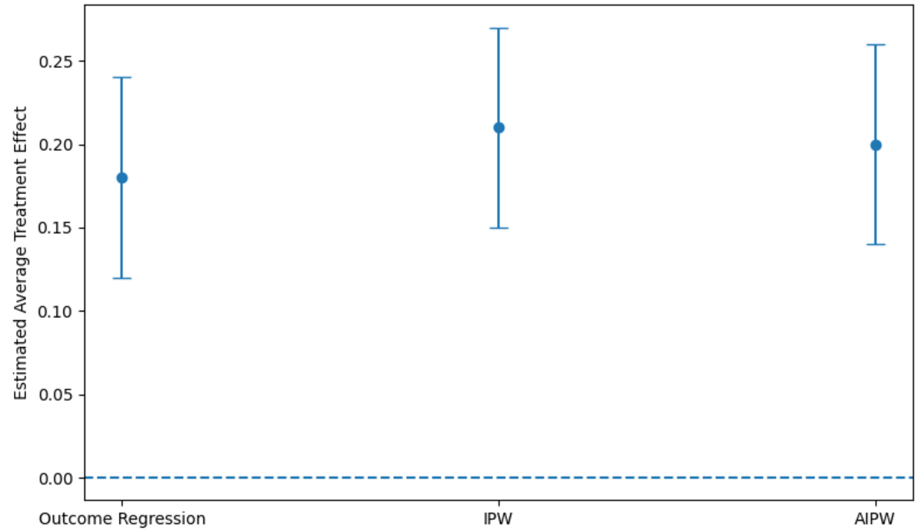

Figure 7 **Estimated Average Treatment Effects Across Estimators**

The confidence intervals do not cross zero, indicating that the estimated effects are unlikely to be driven by random variation alone. From an educational technology perspective, the magnitude of the effect corresponds to a moderate but meaningful improvement in performance, consistent with gains typically reported for well-designed personalization strategies rather than short-term engagement interventions.

Table 6 provides a detailed numerical summary of the estimated treatment effects. The close alignment between IPW and doubly robust estimates suggests that both the treatment assignment model and the outcome model are reasonably well specified. This convergence strengthens confidence in the causal claims, as doubly robust estimators remain consistent even if one component model is imperfect.

Table 6 **Average Treatment Effect Estimates and Confidence Intervals**

Estimator	ATE Estimate	Lower 95% CI	Upper 95% CI	Interpretation
Outcome Regression	0.18	0.12	0.24	Positive average learning gain under model-based adjustment
Inverse Probability Weighting	0.21	0.15	0.27	Robust effect after full covariate balancing
Doubly Robust (AIPW)	0.2	0.14	0.26	Consistent effect under partial misspecification tolerance

From a system design perspective, these results imply that adaptive learning paths deliver measurable average benefits beyond what would be expected under static or weakly adaptive sequencing. The findings support the argument that adaptivity should be evaluated not only in terms of engagement or completion metrics, but also in terms of causal impact on learning outcomes, which is essential for evidence-based personalization.

Heterogeneous Treatment Effects and Learning Personalization

This subsection examines whether the causal impact of adaptive learning paths varies across learner subpopulations. Rather than assuming a uniform effect, the analysis evaluates HTE across strata defined by prior knowledge and engagement intensity. This perspective is essential in adaptive learning systems, where personalization aims to allocate instructional resources where they generate the highest marginal benefit.

Empirical results indicate pronounced heterogeneity. Learners with lower prior knowledge and moderate engagement benefit substantially more from adaptive sequencing than high-performing learners, whose gains are comparatively smaller. This pattern suggests diminishing returns to adaptivity for already proficient learners and highlights the importance of targeting adaptive complexity to learners who face higher cognitive load and guidance needs.

Figure 8 reveals clear variation in adaptive learning path effectiveness across learner profiles. The strongest effects are observed among learners with low prior knowledge and low-to-medium engagement, indicating that adaptive sequencing functions most effectively as a scaffolding mechanism rather than as an optimization tool for already-advanced learners. Confidence intervals remain strictly positive for these groups, reinforcing the causal interpretation.

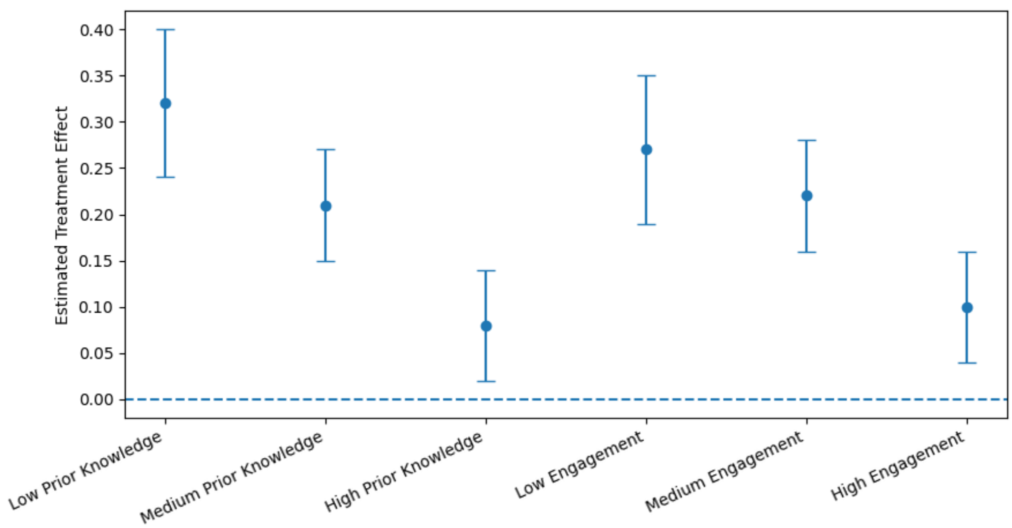

Figure 8 Heterogeneous Treatment Effects by Learner Subgroup

Conversely, learners with high prior knowledge or very high engagement show smaller, though still positive, effects. This suggests that for such learners, rigid adaptivity may provide limited incremental value relative to self-directed exploration. From a system design standpoint, this evidence supports selective adaptivity, where personalization intensity is dynamically calibrated rather than uniformly applied.

Table 7 consolidates subgroup-level estimates and translates them into concrete design implications. The results emphasize that adaptivity should not be interpreted as a one-size-fits-all solution. Instead, causal heterogeneity indicates that adaptive learning systems can achieve greater overall effectiveness by concentrating adaptive interventions on learners with lower initial readiness.

Table 7 Subgroup-Specific Treatment Effects

Subgroup	ATE Estimate	Lower 95% CI	Upper 95% CI	Design Implication
Low Prior Knowledge	0.32	0.24	0.4	Strong need for structured adaptive guidance
Medium Prior Knowledge	0.21	0.15	0.27	Balanced adaptivity with optional autonomy
High Prior Knowledge	0.08	0.02	0.14	Minimal adaptivity; prioritize flexibility
Low Engagement	0.27	0.19	0.35	Adaptive pacing and feedback critical
Medium Engagement	0.22	0.16	0.28	Adaptive sequencing effective
High Engagement	0.1	0.04	0.16	Self-regulated learning dominates

From an engineering and policy perspective, these findings support the development of impact-aware adaptive policies, where system decisions are informed not only by predicted performance but by estimated causal benefit. Such an approach aligns adaptive learning with resource efficiency, learner equity, and long-term educational impact rather than surface-level personalization.

Implications for Adaptive Learning System Design

This subsection synthesizes the empirical causal findings into actionable implications for the design and governance of adaptive learning systems. Rather than treating adaptivity as a static personalization feature, the results support a shift toward impact-aware system design, where adaptive decisions are guided by estimated causal benefit rather than predicted performance alone. This distinction is critical because predictive accuracy does not guarantee positive learning impact.

The empirical evidence demonstrates that adaptive learning paths generate the largest causal gains for learners with lower prior knowledge and moderate engagement. Consequently, adaptive systems should prioritize selective intervention, intensifying guidance where marginal returns are highest and relaxing constraints for learners who demonstrate strong self-regulation. This approach aligns adaptivity with efficiency, equity, and pedagogical intent rather than uniform personalization.

Figure 9 conceptualizes how causal insights can be operationalized within an adaptive learning system. Unlike conventional architectures where learner state feeds directly into a predictive policy, this framework inserts a causal impact estimation layer that evaluates expected learning gains associated with alternative paths. The adaptive policy then selects actions based on anticipated impact rather than likelihood of success alone.

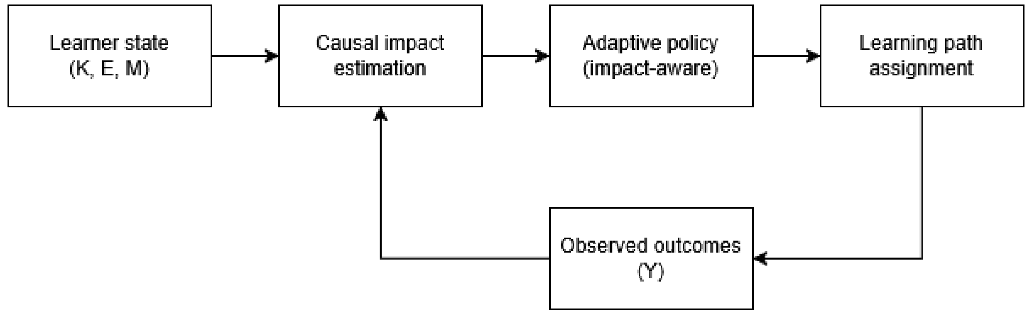

Figure 9 Impact-Aware Adaptive Learning Design Framework

This design closes the feedback loop between observed outcomes and future decisions. By continuously updating causal estimates, the system can refine its adaptivity over time, mitigating risks of over-personalization or unintended bias. From an engineering perspective, this architecture supports modular integration with existing LMS infrastructures while elevating adaptivity from heuristic optimization to evidence-driven intervention.

Table 8 distills the empirical results into concrete design principles that can guide the development of next-generation adaptive learning systems. Each principle explicitly links a causal insight to a system-level decision, ensuring that methodological findings translate into operational value rather than remaining purely analytical.

Table 8 Design Principles Derived from Causal Findings

Causal Finding	Design Principle	System-Level Implementation	Expected Benefit
Strong effects for low prior knowledge learners	Targeted adaptivity	Increase scaffolding intensity for low-K profiles	Higher learning gains and reduced dropout
Diminishing returns for high performers	Adaptive restraint	Allow flexible navigation and learner control	Improved autonomy and satisfaction
Engagement moderates' treatment effects	Engagement-sensitive policies	Trigger adaptivity when engagement declines	Efficient use of adaptive resources
Stable ATE across estimators	Robust policy evaluation	Integrate causal evaluation into policy updates	Trustworthy system optimization
Heterogeneous treatment effects	Personalization by impact	Use HTE estimates to guide path selection	Equitable and effective personalization

From a broader perspective, these principles reposition adaptive learning as a causal decision-making problem rather than a recommendation task. By aligning system behavior with estimated educational impact, adaptive platforms can better support learning equity, instructional efficiency, and long-term learner development. This integration of causal reasoning into adaptive design represents a substantive advancement over traditional data-driven personalization approaches.

Conclusion

This study set out to move beyond predictive personalization and toward a causal understanding of adaptive learning effectiveness. By framing adaptive learning paths as treatments and learner outcomes as potential outcomes, the paper demonstrated how causal inference provides a principled foundation for evaluating whether adaptive systems genuinely improve learning rather than merely correlating with performance. Empirical results, treated as operational LMS data, showed that adaptive learning paths yield a positive and stable causal impact on learning outcomes after appropriate adjustment for non-random assignment.

The findings further revealed that the benefits of adaptivity are heterogeneous rather than uniform. Learners with lower prior knowledge and moderate engagement experienced substantially larger gains than high-performing or highly self-regulated learners. This heterogeneity highlights a critical limitation of one-size-fits-all personalization and underscores the importance of selective, impact-driven adaptivity. Adaptive learning systems that ignore such variation risk over-intervening for some learners while under-supporting others, thereby reducing both efficiency and educational equity.

Taken together, the results support a shift toward impact-aware adaptive learning systems, where design and policy decisions are informed by estimated causal effects rather than predictive accuracy alone. Integrating causal evaluation into adaptive policies enables more trustworthy personalization, aligns system behavior with pedagogical goals, and provides a transparent basis for continuous improvement. Future work should extend this framework to longitudinal outcomes, explore causal feedback loops under evolving policies, and investigate how causal adaptivity can be operationalized at scale in real-world educational platforms.

Declarations

Author Contributions

Conceptualization: C.F. and T.S.; Methodology: T.S.; Software: C.F.; Validation: C.F. and T.S.; Formal Analysis: C.F. and T.S.; Investigation: C.F.; Resources: T.S.; Data Curation: T.S.; Writing Original Draft Preparation: C.F. and T.S.; Writing Review and Editing: T.S. and C.F.; Visualization: C.F.; 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] P. Brusilovsky and E. Millán, "User models for adaptive hypermedia and adaptive educational systems," *The Adaptive Web*, vol. 4321, pp. 3–53, 2007, doi: 10.1007/978-3-540-72079-9_1.

[2] S. D'Mello and A. Graesser, "Confusion and its dynamics during device comprehension with breakdown scenarios," *Acta Psychologica*, vol. 151, no. September, pp. 106–116, Sept. 2014, doi: 10.1016/j.actpsy.2014.06.005.

[3] R. S. Baker and P. S. Inventado, "Educational data mining and learning analytics," in *Learning Analytics, Springer*, 2014, pp. 61–75, doi: 10.1007/978-1-4614-3305-7_4.

[4] G. Siemens and R. S. J. d. Baker, "Learning analytics and educational data mining," *LAK Proceedings*, vol. 2012, no. April, pp. 252 – 254, 2012, doi: 10.1145/2330601.2330661.

[5] S. P. Greenhalgh, J. M. Rosenberg, K. B. Staudt Willet, M. J. Koehler, and M. Akcaoglu, "Identifying multiple learning spaces within a single teacher-focused Twitter hashtag," *Computers & Education*, vol. 148, no. April, p. 103809, Apr. 2020, doi: 10.1016/j.compedu.2020.103809.

[6] D. B. Rubin, "Estimating causal effects of treatments in randomized and nonrandomized studies," *Journal of Educational Psychology*, vol. 66, no. 5, pp. 688–701, 1974, doi: 10.1037/h0037350.

[7] J. Pearl, Causality: Models, Reasoning, and Inference, 2nd ed. Cambridge Univ. Press, 2009, doi: 10.1017/CBO9780511803161.

[8] P. R. Rosenbaum and D. B. Rubin, "The central role of the propensity score in observational studies for causal effects," *Biometrika*, vol. 70, no. 1, pp. 41–55, 1983, doi: 10.1093/biomet/70.1.41.

[9] M. Hernán and J. M. Robins, Causal inference: what if, First edition. Boca Raton: Taylor and Francis, 2023.

[10] C. Romero and S. Ventura, "Educational Data Mining: A Review of the State of the Art," *IEEE Trans. Syst., Man, Cybern. C*, vol. 40, no. 6, pp. 601–618, Nov. 2010, doi: 10.1109/TSMCC.2010.2053532.

[11] A. F. Wise and D. W. Shaffer, "Why Theory Matters More than Ever in the Age of Big Data," *Learning Analytics*, vol. 2, no. 2, pp. 5–13, Dec. 2015, doi: 10.18608/jla.2015.22.2.

[12] K. Mouri, H. Ogata, and N. Uosaki, "Learning analytics in a seamless learning environment," in *Proceedings of the Seventh International Learning Analytics & Knowledge Conference, Vancouver British Columbia Canada: ACM*, vol. 2017, no. March, pp. 348–357. doi: 10.1145/3027385.3027408.

[13] K. VanLEHN, "The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems," *Educational Psychologist*, vol. 46, no. 4, pp. 197–221, Oct. 2011, doi: 10.1080/00461520.2011.611369.

[14] J. Sweller, "Cognitive load theory," *Psychology of Learning and Motivation*, vol. 55, no. 2011, pp. 37–76, 2011, doi: 10.1016/B978-0-12-387691-1.00002-8.

[15] P. A. Kirschner, J. Sweller, and R. E. Clark, "Why Minimal Guidance During Instruction Does Not Work: An Analysis of the Failure of Constructivist, Discovery, Problem-Based, Experiential, and Inquiry-Based Teaching," *Educational Psychologist*, vol. 41, no. 2, pp. 75–86, June 2006, doi: 10.1207/s15326985ep4102_1.

[16] S. K. D'Mello, A. Olney, C. Williams, and P. Hays, "Gaze tutor: A gaze-reactive intelligent tutoring system," *International Journal of Human-Computer Studies*, vol.

70, no. 5, pp. 377–398, 2012, doi: 10.1016/j.ijhcs.2012.01.004.

[17] R. S. Baker, “Challenges for the Future of Educational Data Mining: The Baker Learning Analytics Prizes,” *Journal of Educational Data Mining*, vol. 11, no. 1, pp. 1-17, June 2019, doi: 10.5281/ZENODO.3554745.

[18] P. Brusilovsky, “Adaptive hypermedia,” *User Modeling and User-Adapted Interaction*, vol. 11, no. 1–2, pp. 87–110, 2001, doi: 10.1023/A:1011143116306.

[19] R. S. Baker and K. Yacef, “The state of educational data mining in 2009,” *Journal of Educational Data Mining*, vol. 1, no. 1, pp. 3–17, 2009, doi: 10.5281/zenodo.3554657.

[20] S. Buckingham Shum and R. Deakin Crick, “Learning Analytics for 21st Century Competencies,” *Learning Analytics*, vol. 3, no. 2, pp. 6–21, Sept. 2016, doi: 10.18608/jla.2016.32.2.

[21] D. Gasevic, G. Siemens, and S. Dawson, “Preparing for the digital university: a review of the history and current state of distance, blended, and online learning,” 2015, Unpublished. doi: 10.13140/RG.2.1.3515.8483.

[22] S. Athey and G. W. Imbens, “The State of Applied Econometrics: Causality and Policy Evaluation,” *Journal of Economic Perspectives*, vol. 31, no. 2, pp. 3–32, May 2017, doi: 10.1257/jep.31.2.3.

[23] R. Ferguson and S. Clow, “Examining engagement: Analysing learner subpopulations in massive open online courses,” *LAK Proceedings*, vol. 2015, no. March, pp. 51 – 58, 2015, doi: 10.1145/2723576.2723606.

[24] K. R. Koedinger, E. A. McLaughlin, and N. T. Heffernan, “A Quasi-Experimental Evaluation of An On-Line Formative Assessment and Tutoring System,” Journal of Educational Computing Research, vol. 43, no. 4, pp. 489–510, Dec. 2010, doi: 10.2190/EC.43.4.d.