From Engagement Trajectories to Targeted Support: An Intervention-Oriented Framework for Reducing Learning Gaps in MOOCs

Stefania Zourlidou[1], Rekha Agali Virupaksha[2]

[1, 2] University of Koblenz, Institute for Web Science and Technologies (WeST), Germany

Abstract

Massive Open Online Courses (MOOCs) expand access to learning, but instructors and platforms often struggle to translate noisy and sparse engagement traces into timely support that is both pedagogically meaningful and feasible to deliver at scale. This paper proposes a three-stage, workload-aware intervention framework that bridges early-warning evidence in learner traces and course-level decision making. Stage 1 derives interpretable signals from engagement trajectories and assessment participation (e.g., inactivity accumulation, engagement variability, and assessment-to-content imbalance). Stage 2 maps these signals into risk archetypes using transparent, educator-auditable rules that resolve conflicting evidence and prioritize actionable explanations (e.g., early disengagers, deadline-driven cramming, irregular persistence, and assessment-dominant strategies). Stage 3 operationalizes archetypes under capacity constraints by selecting intervention policies that explicitly balance anticipated instructional value against available instructional resources, enabling graduated support (universal low-intensity supports plus budgeted, targeted higher-intensity supports when risk persists). We use the public HarvardX–MITx Person–Course Academic Year 2013 dataset as a motivating context and describe a filtered working sample of 338,223 person-course records from that release to illustrate the scale and structural sparsity of MOOC learner data, motivating budgeted, stage-specific interventions rather than blanket “at-risk” labeling. The contribution is a practitioner-oriented blueprint for designing interventions that remain faithful to evidence in learner traces while acknowledging the operational realities of large-scale online education.

Keywords: *MOOCs, learning gaps, engagement trajectories, targeted support, capacity-aware interventions*

1. Introduction

MOOCs have enabled wide participation in higher education and professional upskilling, but high dropout rates and persistent learning gaps remain common. In this paper, we use learning gaps broadly to mean emerging differences in learners’ ability to sustain participation, access core course activities, and progress toward successful completion. Foundational MOOC studies documented substantial heterogeneity in how learners engage with course content and assessments, including early disengagement and irregular participation patterns [1,2]. More recent work has emphasized that early-course behavior can be predictive of later success or dropout risk [3,4]. In parallel, the learning analytics community has produced a rich set of methods for profiling learners and predicting outcomes using platform traces [5,6]. Despite these advances, the practical adoption barrier often lies in the step after prediction: translating analytics into timely, feasible, and pedagogically meaningful actions. Instructors and course teams need to decide (i) which engagement signals should trigger a response, (ii) what kind of support is appropriate for a given behavioral pattern, and (iii) how to prioritize interventions under limited capacity. To address this gap, we reframe early warning as an intervention design problem and contribute an intervention-oriented framework that converts interpretable engagement-trajectory signals into actionable support policies. The framework builds on our prior empirical work, which identified these signals through hybrid-weighted clustering and anomaly analysis on the same MOOC release [18]. This paper does not evaluate predictive accuracy or intervention effects. Instead, it presents a rule-based framework for translating learner-trace evidence into feasible support decisions.
The paper makes three contributions: (1) a compact set of trajectory signals intended for intervention triggering (Table 1), (2) an interpretable set of risk archetypes and a corresponding intervention playbook (Table 3), and (3) a workload-aware evaluation perspective that accounts for the number of learners targeted each week (Table 4). Figure 1 summarises the overall pipeline. The intended audience is course teams and platform operators who already have access to learning analytics dashboards but lack a structured way of converting their outputs into bounded, repeatable instructional actions. We therefore prioritise interpretability, low implementation cost, and the ability for an instructor to inspect

why a particular learner was flagged. Throughout the paper, the HarvardX-MITx Person-Course dataset serves as descriptive context for the scale and sparsity challenges of MOOC support. It is not used as an evaluation benchmark, because our contribution is a design framework, not a predictive model.

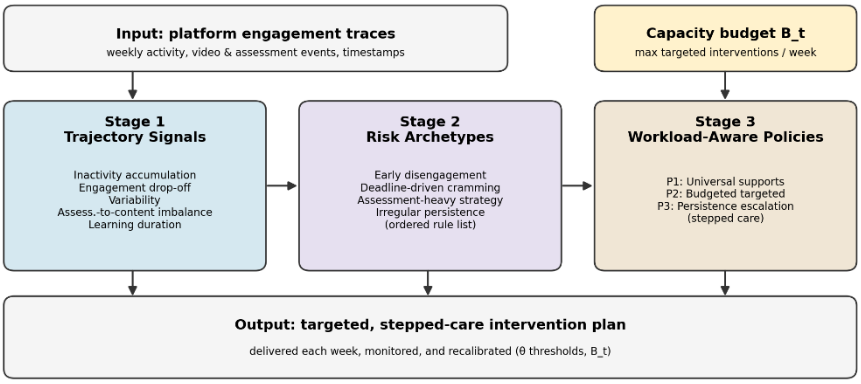

Fig. 1. Three-stage intervention-oriented framework. Engagement traces feed Stage 1 signal extraction, Stage 2 maps signals to interpretable risk archetypes, Stage 3 selects graduated-support policies under a weekly capacity budget B_t.

2. Related Work

Early MOOC research highlighted engagement diversity and the need to interpret learner behavior beyond completion rates [1,2]. Work on engagement subpopulations introduced typologies that distinguish disengaged, auditing, and completing behaviors [2], while dropout prediction studies established that early-week traces can provide useful signals for risk detection [4,8]. Beyond dropout, performance prediction and certification prediction have been explored using prior knowledge, activity traces, temporal features, and interpretable feature sets [3,9]. However, there is a well-known gap between analytics outputs and actionable instructional decisions. Surveys of educational data mining and learning analytics emphasize the importance of aligning models with educational objectives and interpretability requirements [5,6]. Recent systematic reviews also summarize the multi-factor nature of MOOC engagement and dropout, pointing to the importance of contextual and motivational explanations [10]. Motivations and challenges reported by learners and instructors further suggest that interventions should support pacing and reduce friction during critical attrition windows [11].
This paper builds on that line of work by organizing engagement-trajectory evidence into a framework that is explicitly optimized for intervention decisions. Our emphasis is on (i) interpretability of signals and (ii) operational feasibility of responses at MOOC scale. It also builds directly on our prior empirical work on MOOC learning-gap detection [18]. That work examined hybrid-weighted K-Means, OPTICS, and Local Outlier Factor on the HarvardX-MITx Person-Course release, identified engagement signals, such as quiz-to-video ratio, inactive weeks, engagement variability, and engagement slope, that separate learners into clusters with distinct underperformance rates, and analysed clustering failure modes that produce technically valid but outcome-flat partitions. The present paper extends that empirical line of work by translating the previously identified engagement signals into an educator-auditable, workload-aware intervention framework. It serves as the design-and-deployment counterpart to the earlier learner-segmentation analysis.

3. An Intervention-Oriented Framework

The framework proceeds in three stages: (1) compute interpretable trajectory signals, (2) map signal combinations to intervention-ready risk archetypes, and (3) select interventions using workload-aware policies. The framework is designed so that each stage can be implemented by course teams without specialized machine learning infrastructure.

3.1 Stage 1: Trajectory Signal Extraction

Instead of relying solely on aggregate activity counts, we focus on trajectory signals that capture engagement shape over time. Temporal analyses of MOOC study patterns have shown that the timing and regularity of activity can differentiate meaningful behaviors even when total engagement volume is similar [12]. Table 1 lists the proposed signal set.

Signal	Operational definition (example)	Intervention rationale
Inactivity accumulation	Count of course weeks with no recorded events up to week t	Directly indicates disengagement windows, commonly predictive early in a course [4,8].
Engagement drop-off	Ratio of mean engagement in recent weeks vs. initial weeks (e.g., weeks 1-2)	Detects loss of momentum after initial onboarding, supports timely re-engagement prompts [2,11].
Engagement variability	Dispersion of weekly engagement (e.g., standard deviation over weeks 1..t)	High variability can indicate unstable study routines or deadline-driven activity, requiring pacing scaffolds [12,13].
Assessment-to-content imbalance	Smoothed quiz-to-video ratio or related proxy of assessment-first strategies	Flags potentially shallow strategies and test-taking behavior, suggests adaptive review or concept-check follow-up [9,14].
Learning duration	Time span between first and last recorded activity (within the course window)	Supports distinguishing early dropouts from late disengagers, informs when interventions may still help [1,10].

Table 1. Intervention-oriented engagement trajectory signals.

3.2 Stage 2: Risk Archetypes

Stage 2 translates the signal set from Stage 1 into a small number of intervention-ready risk archetypes. The aim is not to discover clusters, but to produce labels that are (i) interpretable, (ii) stable enough for operational use, and (iii) directly connected to a plausible instructional response. Each archetype specifies: (1) a signal pattern (expressed as simple conditions on Table 1 signals), (2) a likely barrier or learning-gap mechanism, and (3) a recommended first-line action from the intervention playbook (Table 3). Archetypes are therefore designed as interpretable support categories rather than statistical groupings.

Rule template. Let S(t) be the vector of signals from Table 1 computed up to week t for a learner. An archetype a is defined by a Boolean rule f_a(S(t)) that returns TRUE when the observed early trajectory matches the archetype pattern. If no rule fires, the learner remains unassigned and receives only universal low-intensity supports (Stage 3, Policy P1).

Conflict resolution and decision order. Multiple rules may fire simultaneously (e.g., a learner can be both highly variable and also show a strong drop-off). To ensure a single actionable label, we apply an ordered decision list that prioritizes conditions that suggest imminent disengagement and require earlier intervention. A practical default order is: (1) Early Disengagement, (2) Deadline-Driven Cramming, (3) Assessment-Dominant Strategy, and (4) Irregular Persistence.

Example ordered rules (course-calibrated thresholds θ are set from early-course distributions): (1) Early Disengagement if inactivity accumulation ≥ θ_inactive OR drop-off ratio ≤ θ_drop, (2) Deadline-Driven Cramming if engagement variability ≥ θ_var AND a late surge indicator is TRUE, (3) Assessment-Dominant Strategy if assessment-to-content imbalance ≥ θ_imbalance (assessment-dominant), (4) Irregular Persistence if engagement variability ≥ θ_var but without sustained decline, otherwise Unassigned.

These archetypes are intentionally coarse. Their purpose is to provide educators with an auditable explanation and a clear first step, not a definitive diagnosis. In deployment, thresholds and rule order can be tuned to course context, and archetype definitions can be validated by monitoring whether the corresponding interventions reduce subsequent risk signals.

3.3 Stage 3: Workload-Aware Intervention Policies

Stage 3 turns archetype labels into workload-aware intervention decisions. The key inputs are (i) the set of learners flagged at week t, (ii) the intensity and timing of available interventions, and (iii) an explicit capacity budget B_t, defined as the maximum number of targeted interventions that can be delivered in that week without overwhelming learners or staff. Budgeting is central in large-scale courses because even low-cost actions can scale to thousands of messages, and excessive targeting can reduce trust and increase opt-outs. We therefore recommend a graduated-support approach that starts with universal, low-intensity supports and escalates to higher-intensity, targeted supports only when risk persists and budget allows.

Policy families. We recommend three deployable policy types: (P1) Universal supports (e.g., pacing reminders and 'start-here' prompts) delivered to all learners, (P2) Budgeted targeted supports delivered only to the top-ranked flagged learners up to B_t (e.g., structured catch-up pathways, adaptive review prompts, or tailored remediation links matched to the archetype), and (P3) Persistence-based escalation, where learners flagged for consecutive weeks move to a higher-intensity targeted tier while one-time flags receive only P1 or light P2.

Prioritization within a budget can be implemented with a simple priority score derived from Table 1 signals (e.g., combining inactivity, drop-off ratio, and variability) and selecting the highest-scoring learners until the weekly budget is exhausted. A lightweight monitoring loop (delivery counts, opt-outs, and subsequent engagement shifts) supports recalibration of θ thresholds and B_t over time.

3.4 Deployment Considerations

The framework is intentionally designed to be implemented on top of an existing learning analytics dashboard, not as a standalone system. A minimal deployment requires three artefacts: (i) a weekly batch job that computes the Table 1 signals from platform logs, (ii) a small rule file expressing the ordered Stage 2 archetype list and the current threshold values θ, and (iii) a delivery layer that can dispatch templated messages or surface in-course banners under the Stage 3 policies P1-P3. None of these components requires model training, GPU infrastructure, or feature engineering beyond the aggregates that most MOOC platforms already expose.

In practice we recommend that course teams begin with a pilot deployment without learner-facing messages, in which signals and archetype labels are computed and logged each week. After two to three weeks of pilot data, instructors can inspect the distribution of triggered archetypes, sanity-check a small random sample of flagged learner trajectories, and tune θ and B_t before activating P1 supports. This staged roll-out is also the natural time to confirm that the rule order in Stage 2 reflects the course's pedagogical priorities, since the same signal pattern can plausibly justify more than one first-line action depending on the discipline.

4. Data Context and Implementation Requirements

We use the public HarvardX-MITx Person-Course Academic Year 2013 dataset, a de-identified dataset of learner-course records from HarvardX and MITx offerings on edX, covering Fall 2012 to Summer 2013 across 16 courses [15]. Prior analyses using this de-identified release report 641,138 person-course observations from 476,532 unique learners after aggregation [17]. For the descriptive scale-and-sparsity context used in this paper, we draw on a 338,223-record analytic subset derived from this release in our prior work [18], restricted to five selected MOOCs (CS50x, CB22x, ER22x, PH207x, and PH278x) and obtained after the study's preprocessing pipeline (exclusion of internally inconsistent records flagged in the release, zero-imputation of structural non-engagement, and removal of records not meeting the analytic inclusion criteria). Within that subset, temporal trajectory analyses focused on a single 21-week course to avoid cross-course differences in calendar length, pacing, assessment cadence, and instructional design. Because the public file contains aggregate counts and does not provide per-event timestamps, weekly activity was approximated by distributing each learner's total events across their active weeks, and computationally expensive analyses, such as DBSCAN with DTW distances, were applied to a subsample [18]. The subset is used here only as descriptive context for the scale and sparsity challenges of MOOC support design, not as a benchmark for evaluating predictive performance or intervention effects.

Initial inspection of the working sample highlights substantial structural sparsity: in our prior analysis of this subset, 90.16% of learners exhibited no recorded video interactions and 62.04% showed no activity in key course events such as problem sets and chapter interactions [18]. This matters operationally because trace-based support rules can only act on learners for whom some behavioural evidence is available.

Characteristic	Value	Notes
Dataset	HarvardX–MITx Person–Course Academic Year 2013	Public de-identified person–course release
Public release	641,138	Person-course observations across 476,532 learners, Fall 2012-Summer 2013
Working sample	338,223	Filtered subset used for the descriptive budgeting context in this paper
Video interactions	Substantial sparsity in video-event traces (90.16% of learners with no video interactions in prior analysis [18])	Illustrates uneven content-view traces in large MOOC data
Assessed components	Substantial sparsity in assessed-component traces (62.04% of learners with no activity in key course events [18])	Early assessed-event traces may be absent for many learners

Table 2. Dataset context and descriptive indicators used to motivate workload-aware intervention budgeting.

The signals in Table 1 require timestamped or week-level platform logs in deployment (e.g., weekly activity counts, video and assessment events, and timestamps of first/last activity). The public person–course release is therefore used here to contextualize scale and structural sparsity, not to estimate all week-level trajectory features directly. These descriptive indicators are not used as performance results as they provide operational motivation for Stage 3 budgeting by quantifying how many learners may lack actionable trace evidence early in a course.

5. Intervention Playbook

Table 3 presents a practical playbook that maps archetypes to scalable interventions. The interventions are designed to be platform-feasible and to support learners during known attrition windows, particularly the early weeks of a MOOC [4,11].

Archetype	Signal pattern (example)	Likely barrier	Scalable intervention
Early disengagement	Inactivity accumulation rises by week 2-3, engagement drop-off < 1	Onboarding friction, low self-efficacy, unclear expectations	Week-2 re-engagement nudge, 'start-here' guide, short diagnostic quiz with feedback, optional catch-up path
Deadline-driven cramming	High engagement variability with late spikes, drop-off after deadlines	Time management issues, reliance on last-minute study	Pacing scaffold: weekly plan and reminders, chunked deadlines, 'two-step' practice (preview then quiz)
Assessment-dominant strategy	High assessment-to-content imbalance (quiz-to-video ratio high)	Superficial engagement, test-taking without conceptual grounding	Adaptive review prompts after each quiz, short concept-check quizzes, concept-check items with explanations

Archetype	Signal pattern (example)	Likely barrier	Scalable intervention
Irregular persistence	Moderate engagement but high variability without sustained decline	Competing commitments, unstable routines	Routine-building nudges, flexible pacing options, goal-setting prompts, social participation suggestions

Table 3. Risk archetypes and a scalable intervention playbook.

6. Workload-Aware Targeting and Evaluation

A common failure mode in practice is to deploy an early-warning rule that flags an impractically large fraction of learners, leading to interventions that are too diffuse to be meaningful. We propose evaluating intervention policies along three dimensions (Table 4): (i) lead time, (ii) targeting quality, and (iii) intervention workload.

Evaluation dimension	What to measure	Why to measure
Lead time	Earliest week where a stable signal triggers (e.g., week 2 vs. week 5)	Earlier targeting enables low-cost supports before disengagement hardens [4,8].
Targeting quality	Association between triggers and later underperformance or disengagement	Ensures interventions focus on learners likely to benefit, aligning analytics with educational objectives [6,9].
Intervention workload	Proportion and count of learners triggered per week, intervention type cost	Supports feasible deployment, avoids overwhelming staff or learners, enables capacity planning.

Table 4. Workload-aware evaluation dimensions for intervention policies.

Instructors can implement workload-aware targeting using simple budget rules (e.g., intervene on the top 10% highest-risk eligible learners each week) or graduated-support strategies (e.g., universal low-cost nudges plus targeted higher-cost supports). To make the budget intuition concrete, consider an illustrative course scenario with 30,000 enrollees. If 10% of enrollees (3,000 learners) generate sufficient early-week traces to be eligible for trajectory-based flagging, then a P2 budget of 5% of this eligible subset corresponds to 150 targeted contacts in a given week.
As a further illustration, suppose θ_inactive is calibrated so that the Early Disengagement rule fires on the top quartile of this eligible subset. Roughly 750 learners would then be candidates for P2 supports in week 2. Setting B_t = 150 would therefore force an explicit prioritisation step, with the remaining candidates routed to P1 universal messaging until the following week, when persistence-based escalation is re-evaluated. This kind of simple calculation is intended as a planning illustration, not an empirical performance result, and shows how analytics outputs can be translated into a delivery plan that an instructional team can execute.

7. Discussion

The framework reframes learner analytics outputs as intervention inputs. The approach supports educator trust and practical deployment by emphasizing trajectory shape signals and interpretable archetypes. The archetypes align with established observations of MOOC engagement patterns, including early disengagers and irregular participants [2,12]. Compared with prior work, our contribution is prescriptive and not descriptive. Earlier MOOC typologies [2] characterise how learners behave but stop short of mapping behaviours to instructional responses, while dropout-prediction models [4,8] estimate risk but rarely specify what to do once a learner is flagged. Our framework occupies the intermediate layer: it converts interpretable signals into a small set of support categories, each carrying a default action, constrained by an explicit workload budget so that the prescription is both pedagogically reasonable and operationally deliverable.
A deployment trial should track the three dimensions in Table 4 across at least one full course offering, examining whether the framework yields (i) earlier P1 universal supports than baseline instructor-driven interventions, (ii) improved or stable week-over-week engagement among learners receiving escalated

P2/P3 support relative to a matched comparison group, and (iii) a weekly intervention workload within the configured B_t budget. Degradation on any dimension should trigger threshold and rule-order recalibration rather than abandonment of the framework.
From a ‘future of education’ perspective, the primary value lies in scalable personalization: capacity-aware evaluation helps institutions plan support offerings across large cohorts and the same signal-archetype-policy pipeline can be transferred to university LMS courses, bootcamps, and corporate training where instructional capacity similarly constrains support delivery. A further implication concerns the role of the instructor. By exposing thresholds, rule order, and budgets as explicit operational parameters, the framework deliberately keeps a human in the loop: instructors can override archetype assignments, change priority order, or suspend Stage 3 escalation around exam weeks. We see this transparency as a precondition for responsible adoption in publicly funded courses, where learners and regulators may legitimately ask why a particular support was offered.

8. Limitations and Ethical Considerations

Several limitations should be considered before deployment. First, engagement traces are imperfect proxies for learning and may miss off-platform study, peer discussion, or use of external materials. Archetype labels should therefore be treated as conversation starters with learners, not as diagnoses. Second, signal thresholds θ are course-dependent and assume that early-course distributions are stable enough to calibrate against. In courses with very small early enrolment or highly skewed cohorts, calibration may be unreliable, and instructor review is essential before rule activation. Third, the framework itself has not yet been validated experimentally. We propose archetypes and a playbook on the basis of prior literature and dataset descriptive evidence, but we do not show that the resulting interventions improve learning outcomes. A field trial measuring learning gains, retention, and learner experience under different intervention budgets is the natural next step. Fourth, interventions can have unintended consequences. Excessive nudging may produce notification fatigue, opt-outs, or stigma effects, particularly when learners infer that they have been singled out as ‘at risk’. Graduated-support escalation and explicit learner-facing explanations are intended to mitigate these effects but do not eliminate them. Fifth, the public HarvardX-MITx Person-Course release reflects MOOC behaviour from the 2012-2013 academic year. Modern platforms (current Coursera, edX, FutureLearn) differ in structure, mobile-first interaction, and learner demographics, so transferability of specific thresholds should not be assumed without recalibration. Sixth, privacy and transparency are essential when learner data drive personalised messaging. Deployments should follow responsible data practices [16], including clear opt-out mechanisms and learner-facing explanations of why support was offered. Relatedly, although prior empirical work [18] included demographic variables (e.g. age, gender, education) for explanatory analysis, the intervention rules proposed here should rely on behavioural engagement signals rather than protected or sensitive learner attributes unless their use is explicitly justified, audited, and ethically approved.

9. Conclusion and Future Work

This paper presented an intervention-oriented framework that maps interpretable engagement-trajectory signals into actionable, scalable support strategies for MOOCs. By combining a compact signal set (Table 1), intervention-ready archetypes (Table 3), and workload-aware evaluation (Table 4), the framework offers a practical bridge between learning analytics and course design decisions. The framework is intended as a design and implementation aid, not as an intervention-effectiveness study, and its next step is empirical validation through field deployment.
Several extensions follow naturally. The most important is empirical validation through a controlled or quasi-experimental field study in which Stage 3 policies are activated in one course offering and compared against an instructor-driven baseline along the dimensions in Table 4. A second extension is to broaden the rule library beyond the four archetypes described here for project-based, group-assessed, or self-paced certification tracks where engagement rhythms differ from cohort-based MOOCs. A third is to integrate the framework with adaptive content systems so Stage 3 actions can include lightweight content adjustments such as alternate pathways or scaffolded retrieval-practice items. On the methodological side, future work should investigate principled ways of setting and updating the thresholds θ and the weekly budget B_t, for example by treating threshold calibration as a constrained optimisation problem where the objective is week-level intervention precision under a hard cap on instructional cost. Rule order may also need to shift across the course timeline, and learner-

facing transparency mechanisms should be designed and tested as part of the intervention bundle from the start, instead of being added only after deployment.

REFERENCES

[1] A. Anderson, D. Huttenlocher, J. Kleinberg, and J. Leskovec, "Engaging with massive online courses," in Proceedings of the 23rd International Conference on World Wide Web (WWW '14), ACM, 2014, pp. 687–698.
[2] R. F. Kizilcec, C. Piech, and E. Schneider, "Deconstructing disengagement: analyzing learner subpopulations in massive open online courses," in Proceedings of the 3rd International Conference on Learning Analytics and Knowledge (LAK '13), ACM, 2013, pp. 170–179.
[3] G. Kennedy, C. Coffrin, P. de Barba, and L. Corrin, "Predicting success: how learners' prior knowledge, skills and activities predict MOOC performance," in Proceedings of the 5th International Conference on Learning Analytics and Knowledge (LAK '15), ACM, 2015, pp. 136–140.
[4] M. Kloft, F. Stiehler, Z. Zheng, and N. Pinkwart, "Predicting MOOC dropout over weeks using machine learning methods," in Proceedings of the EMNLP 2014 Workshop on Analysis of Large Scale Social Interaction in MOOCs, 2014, pp. 60–65.
[5] C. Romero and S. Ventura, "Educational data mining: a review of the state of the art," IEEE Transactions on Systems, Man, and Cybernetics, Part C, vol. 40, no. 6, pp. 601–618, 2010.
[6] C. Romero and S. Ventura, "Educational data mining and learning analytics: an updated survey," Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, vol. 10, no. 3, e1355, 2020.
[7] O. Zawacki-Richter, A. Bozkurt, U. Alturki, and A. Aldraiweesh, "What research says about MOOCs - an explorative content analysis," International Review of Research in Open and Distributed Learning, vol. 19, no. 1, pp. 242–259, 2018.
[8] J. Whitehill, K. Mohan, D. Seaton, Y. Rosen, and D. Tingley, "MOOC dropout prediction: how to measure accuracy?," in Proceedings of the 4th ACM Conference on Learning @ Scale, 2017, pp. 161–164.
[9] G. Kostopoulos, T. Panagiotakopoulos, S. Kotsiantis, C. Pierrakeas, and A. Kameas, "Interpretable models for early prediction of certification in MOOCs: a case study on a MOOC for smart city professionals," IEEE Access, vol. 9, pp. 165881–165891, 2021.
[10] H. Huang, L. Jew, and D. Qi, "Take a MOOC and then drop: a systematic review of MOOC engagement pattern and dropout factor," Heliyon, vol. 9, no. 4, e15220, 2023.
[11] K. F. Hew and W. S. Cheung, "Students' and instructors' use of massive open online courses (MOOCs): motivations and challenges," Educational Research Review, vol. 12, pp. 45–58, 2014.
[12] M. Shirvani Boroujeni and P. Dillenbourg, "Discovery and temporal analysis of MOOC study patterns," Journal of Learning Analytics, vol. 6, no. 1, pp. 16–33, 2019.
[13] J. Qiu, J. Tang, T. X. Liu, J. Gong, C. Zhang, Q. Zhang, and Y. Xue, "Modeling and predicting learning behavior in MOOCs," in Proceedings of the 9th ACM International Conference on Web Search and Data Mining (WSDM '16), 2016, pp. 93–102.
[14] X. Zhu, Y. Ye, L. Zhao, and C. Shen, "MOOC behavior analysis and academic performance prediction based on entropy," Sensors, vol. 21, no. 19, 6629, 2021.
[15] MITx and HarvardX, "HarvardX-MITx Person-Course Academic Year 2013 De-Identified dataset, version 3.0," Harvard Dataverse, 2014. doi:10.7910/DVN/26147.
[16] J. P. Daries, J. Reich, J. Waldo, E. M. Young, J. Whittinghill, A. D. Ho, D. T. Seaton, and I. Chuang, "Privacy, anonymity, and big data in the social sciences," Communications of the ACM, vol. 57, no. 9, pp. 56–63, 2014.
[17] S. Jiang, K. Schenke, J. S. Eccles, D. Xu, and M. Warschauer, "Cross-national comparison of gender differences in the enrollment in and completion of science, technology, engineering, and mathematics Massive Open Online Courses," PLoS ONE, vol. 13, no. 9, e0202463, 2018. doi: 10.1371/journal.pone.0202463.
[18] R. A. Virupaksha, S. Zourlidou, and F. Hopfgartner, "Extending clustering and causal analysis techniques for identifying and explaining learning gaps in MOOCs," in Proc. International Conference on Interactive Mobile Communication, Technologies and Learning (IMCL 2025), 2025 (to appear).