

## Adaptive Recommendation Strategies for MOOCs: Linking Learner Motivation, Behavior, and Completion Outcomes

Sandeepi, Research Scholar, Department of Computer Science, NIILM University, Kaithal, Haryana  
Dr. Sandeep Chahal, Professor, Department of Computer Science, NIILM University, Kaithal, Haryana

### Abstract

The dropout problem has scarcely changed in MOOCs for over a decade. Coursera, edX, and Udemy claim completion rates between 5% and 15%, which is disappointing given how many individuals start these courses with intent. This research contends that the fundamental reason for this gap between enrollment and completion is that these platforms' recommendation engines don't understand the learner. A research of 9,556 MOOC learners found that five motivational traits affect participation, course rating, and dropout probability. We also develop and test a Q-Learning-based adaptive recommendation scheme that dramatically outperforms baselines. Precision@5 improved 76.5%, anticipated completion rate increased from 44.7% to 61.3%, and dropout rate decreased from 44.2% to 28.7%. Our 'limited engagement bias' finding—dormant learners give courses the highest average ratings ( $M = 47.53$ )—is a major issue for platforms that feed those ratings back into recommendation engines without checking who provided them.

**Keywords:** MOOC, personalized recommendation, Q-Learning, learner motivation, dropout prediction, user profiling, reinforcement learning, engagement behavior, completion rate, limited engagement bias.

### 1. Introduction

The completion rate problem is hardly a mystery, as most MOOC researchers know but rarely admit. We approximately understand dropouts. They enrolled in something that wasn't right for them, was too long, pitched at the incorrect level, or wasn't related to their goals. The recommendation system picked something based on what other people enrolled in or that matched a few phrases in their profile, but it was inaccurate. That's the fundamental issue. In the past decade, the global MOOC market has risen rapidly. Nearly 220 million learners have registered on major platforms by 2021, with over 19,000 courses in Python programming, machine learning, business strategy, and public health (Shah, 2021). However, most of these courses remain unfinished. Jordan (2015) found a median completion rate of 12.6% across 221 MOOCs, with several courses below 10%. Recent data from individual platforms reveals this figure has not increased despite years of platform design and content quality investment. The literature has mostly explained why MOOCs attract 'casual browsers'—curious students who don't intend to finish. Some learners may experience this, but it oversimplifies and leads platform designers and academics to view dropout as an expected characteristic rather than a failing. This study examines the data, revealing a more nuanced picture. Dropout is concentrated among underserved students. Predictable. With the correct recommendation infrastructure, it's possible.

The paper offers empirical findings from a study of 9,556 MOOC learners and introduces a Q-Learning-based adaptive recommendation framework to fill customization system deficiencies. The paper is organized: Section 2 briefly reviews relevant prior work; Section 3 describes the data, methodology, and framework design; Section 4 presents empirical results across four research questions; Section 5 discusses implications and theoretical contributions; and Section 6 concludes with limitations and future work.

### 2. Background and Related Work

#### 2.1 The Dropout Problem: What the Literature Actually Shows

Based on learners' course content interaction patterns, **Kizilcec et al. (2013)** proposed a typology of MOOC learner behavior: finishing, auditing, disengaging, and sampling. This investigation was crucial since it proved that 'dropout' is multiple things that happen for different causes. A learner who watches only lecture videos and never takes quizzes is different

from one who finishes the first three weeks then stops.

**Yang et al. (2013)** found that dropout signals are detected within one to two weeks of a learner's course trajectory. Early intervention was made possible by this finding, but the field has been slow to apply it to recommendation system design. Most MOOC platforms lack Week 1 engagement indicator early warning mechanisms.

Content-based filtering (matching courses to learner interests based on topic tags and descriptions), collaborative filtering (recommending what similar learners also took), and hybrid combinations of the two have dominated MOOC recommendation. Static, well-structured recommendation settings suit these methods. They perform poorly in MOOCs for two reasons. First, learner profiles are frequently brief: platforms collect demographic data upon registration and maybe an interest area. Second, as learners gain knowledge, change professional priorities, or discover they don't require what they thought they did, their needs change. Static registration profiles cannot follow these changes.

### 2.2 Motivation and Self-Determination Theory

Self-Determination Theory (SDT), created by Deci and Ryan (1985, 2000), is the most widely accepted theory for why some MOOC students persist. SDT contrasts intrinsic motivation—engaging in an activity out of true interest or enjoyment—from extrinsic motivation—engaging due to external pressure, requirements, or rewards. According to the notion, intrinsically driven learners will engage more, persist more, and be happier with their learning. This prediction holds well in MOOCs. Hew and Cheung (2014) showed that intrinsic motivation best predicted MOOC completion in their engagement study review. Professional learners with clear intrinsic goals had much higher completion rates than employer-mandated learners, according to Milligan and Littlejohn (2014). The literature has only lately translated these motivational characteristics into recommendation system design operational aspects. This study's main contribution is that translation.

### 2.3 Reinforcement Learning for Recommendation

Deep Q-networks' effectiveness in game-playing and sequential decision-making has spurred interest in reinforcement learning (RL) recommendation methods (Mnih et al., 2015). RL is appealing in recommendation contexts because it naturally models the sequential and interactive user-platform relationship: each recommendation is a system decision, which observes the user's response, updates its model, and improves the next decision (**Zhao et al., 2018**).

In a comprehensive survey of deep reinforcement learning in recommender systems, **Chen et al. (2019)** found that RL approaches outperform static collaborative filtering and content-based approaches in dynamic user preferences and long-term engagement optimization. Dynamic preferences and long-term participation match the MOOC setting, making RL a perfect choice for MOOC recommendation. Few research have empirically validated RL for MOOC recommendation. That gap is addressed in this study.

## 3. Data, Methodology, and Framework Design

### 3.1 Data Sources and Pre-processing

A study used two complementary data sources. A structured survey was given to Coursera, edX, Udemy, and FutureLearn MOOC learners to collect data on motivational orientation, content format preferences, course duration preferences, prior educational background, working status, and self-reported engagement patterns. Secondary data included Kaggle and UCI Machine Learning Repository datasets and learner interaction logs from publicly accessible MOOC platform pages. The final dataset has 9,556 valid learner records after Python (pandas, scikit-learn) preprocessing to remove missing values, duplicates, and outliers. SPSS 26 was used for statistical analysis.

Active learners engaged with at least 60% of available content weekly (N = 5,495, 57.5%), Intermittent learners engaged with 20–59% (N = 2,568, 26.9%), and Dormant learners engaged

with less than 20%. The categories were created from platform log data and evaluated against survey self-reported involvement levels, demonstrating good concordance ( $\kappa = 0.71$ ).

### 3.2 Motivational Profiling

Using a five-item survey, learners were asked to choose their main MOOC enrollment reason from five categories and elaborate. Career Progression (skills for professional advancement), Knowledge Acquisition (learning for its own sake), Certification/CV Enhancement (employment credentials), Domain Exploration (curiosity-driven sampling of new fields), and External Requirement were the five categories. These categories are based on SDT (Deci & Ryan, 2000) and MOOC motivation typologies (Hew & Cheung, 2014).

### 3.3 The Q-Learning Recommendation Framework

The Q-Learning approach treats MOOC recommendation as sequential decision-making. The present learner profile—a vector of demographic, behavioral, motivational, and real-time engagement dimensions—dictates the system's state. Action space is course recommendations. Course completion with a high rating ( $\geq 4$  stars) earns +10, moderate rating (+6 or +2), sustained engagement without completion (+4), early dropout (Week 1–2), late dropout (after Week 3), passive browsing (0), and re-enrollment in a recommended course (+3).

The Q-table was initialized with uniform values and updated using the Q-Learning update rule:  $Q(s,a) \leftarrow Q(s,a) + \alpha[r + \gamma \max_{a'} Q(s',a') - Q(s,a)]$ , where  $\alpha = 0.01$  is the learning rate,  $\gamma = 0.9$  is the discount factor, and  $\varepsilon = 0.1$  sets the exploration–exploitation tradeoff in the  $\varepsilon$ -gre Train the framework on 80% of the dataset and test it on 20%. The Interactive Course Recommendation Framework (ICRF), a collaborative filtering technique, and Automatic Information Retrieval (AIR), a content-based matching engine, were used to compare performance.

## 4. Results

### 4.1 Motivational Profile Distribution and Engagement Outcomes

Table 1 displays the sample's motivational traits and engagement and completion outcomes. The striking pattern matches SDT expectations. Knowledge-motivated students finish 63.7 percent and dropout 22.1 percent least. Externally driven learners have the worst outcomes, with a 31.9% completion rate and 54.7% dropout rate—more than twice the highest-performing group.

**Table 1. Motivational Profile Distribution and Associated Engagement Outcomes (N = 9,556)**

Motivation Profile	N (%)	Completion Rate (%)	Dropout Rate (%)	Mean Course Rating	Mean Engagement Score
Career Progression	2,847 (29.8%)	58.4	28.2	4.1 / 5.0	74.3
Knowledge Acquisition	2,391 (25.0%)	63.7	22.1	4.4 / 5.0	81.2
Certification / CV	1,911 (20.0%)	47.3	38.4	3.8 / 5.0	61.4
Domain Exploration	1,339 (14.0%)	61.2	24.6	4.3 / 5.0	77.8
External Requirement	1,068 (11.2%)	31.9	54.7	3.2 / 5.0	48.2
Total / Weighted Avg.	9,556 (100%)	52.8	32.1	4.0 / 5.0	71.3

The gap between intrinsically motivated learners (Knowledge Acquisition + Domain Exploration completion rate ~62%) and extrinsically motivated learners (Certification + External completion rate ~40%) is significant. It demonstrates that motivational orientation is

one of the most significant elements a recommendation system should capture at onboarding, not as a curiosity but as a main driver of its suggestion approach.

#### 4.2 Content Format and Duration Preferences by Learner Type

Table 2 shows working status-based format and duration preferences. The course time split is most important: 47.3% of working professionals choose 1–4 week micro-courses, whereas 54.2% of students prefer 4–12 week courses. This disparity is so significant that the recommendation engine needs duration-based filtering. A platform that suggests a 12-week course to a working professional may set them up to drop out, not because they lack dedication but because their time architecture is different.

**Table 2. Content Format and Course Duration Preferences by Working Status**

Preference Dimension	Students (%)	Professionals (%)	Overall (%)
Video Tutorials (preferred format)	52.4	44.1	48.7
Interactive / Lab Content	21.3	31.8	26.0
Text / PDF Materials	14.7	12.9	13.9
Audio / Podcast Style	6.8	7.4	7.1
Micro-course (1–4 weeks)	18.6	47.3	31.3
Standard MOOC (4–12 weeks)	54.2	38.9	47.3
Long-term Module (12+ weeks)	27.2	13.8	21.4

Professional learners choose interactive lab content over video lectures, compared to 21.3% of students. Adult learning literature shows that professionals prefer practice-based, application-focused learning over expository education. A recommendation system that ignores format choices will always suggest courses that fit a learner's topic but not their preferred method.

#### 4.3 The Rating Paradox: Limited Engagement Bias

A surprising conclusion in this study is the link between participation category and course rating behavior. See Figure 1 for the visual pattern and Table 3 for descriptive information. Contrary to expectations, dormant learners rate courses highly.

**Table 3. Descriptive Statistics: Course Ratings by Participation Category**

Participation Category	N	Mean Rating	Std. Deviation	Min	Max
Active Learners	5,495	43.17	15.023	10.00	80.00
Intermittent Learners	2,568	40.57	13.061	12.00	75.00
Dormant Learners	1,493	47.53	11.712	18.00	78.00
Total	9,556	43.12	14.207	10.00	80.00

A one-way ANOVA showed a significant difference:  $F(3, 9235) = 27.635, p < .001$ . Using Tukey HSD post-hoc testing, all three pairwise comparisons were significant ( $p < .001$ ), with the biggest difference between Dormant and Intermittent learners (mean difference = 6.96 rating points). Limited engagement bias is our interpretation. Dormant learners have barely touched a course's content. Their grade includes first-impression criteria including the instructor's charisma in the promotional video, the platform's visual design, and the course title's appeal, not the learning experience. Active learners have done the homework, attended all the lectures, probably struggled, and made a more nuanced judgment. Their ratings average lower but provide more course quality diagnostics. Any recommendation engine that uses aggregate course ratings as a primary quality signal without controlling for rater engagement is biased toward courses that make a good first impression rather than courses that provide real

educational value. A participation-weighted rating approach would create a far more accurate quality signal by weighting ratings from high-engagement learners.

#### 4.4 Dropout Predictors: Early Signals and Course Type

Early engagement signal analysis showed dropout risk is recognized before the point of no return. Table 4 displays the mean Week 1 engagement markers for eventual completers and non-completers, along with their Cohen's d effect sizes.

**Table 4. Early Engagement Indicators by Completion Outcome with Effect Sizes**

Week 1 Indicator	Completers (Mean)	Non-Completers (Mean)	Cohen's d	Interpretation
% of Week-1 Content Consumed	74.3%	28.6%	1.94	Very Large
Logins in First 7 Days	8.4	2.1	1.82	Very Large
Days to First Quiz Attempt	3.2 days	9.7 days	1.61	Large
Assignment Submissions (Week 1)	1.7	0.3	1.67	Large
Forum Posts in First Week	2.8	0.4	1.43	Large

All five indicators have large-to-very large effects. The best predictor of dropping out is Week 1 content intake ( $d = 1.94$ ): students who consume less than 30% are at high risk by Day 7. A system monitoring these signals can identify at-risk learners within the first week and intervene with content re-recommendations, encouraging messages, or format alternatives before dropout. A Chi-Square test found a significant correlation between course type (free vs. paid) and participation category ( $\chi^2(3) = 76.346, p < .001$ ). Interestingly, active learners (81.1%) were mostly in free courses, whereas dormant learners were split evenly between free (51.8%) and paid (48.2%). This undermines the idea that financial investment predicts commitment. Based on SDT, free courses self-select for intrinsically driven learners, whereas paid courses attract a larger community of extrinsically motivated students.

#### 4.5 Q-Learning Framework Performance

Table 5 presents the performance comparison across the three recommendation models. The Q-Learning framework shows consistent and substantial improvements over both baselines on all seven metrics.

**Table 5. Performance Comparison: AIR, ICRF, and Q-Learning Models**

Metric	AIR Model	ICRF Model	Q-Learning	Q-L vs. ICRF (% Gain)
Precision@5	0.318	0.421	0.743	+76.5%
Recall@10	0.294	0.387	0.681	+75.9%
F1 Score	0.305	0.402	0.709	+76.4%
NDCG@5	0.341	0.448	0.762	+70.1%
Mean Avg. Precision	0.278	0.374	0.694	+85.6%
Predicted Completion Rate	38.2%	44.7%	61.3%	+37.1%
Predicted Dropout Rate	51.4%	44.2%	28.7%	-35.1 pp

Better precision (76.5%), mean average precision (85.6%), and anticipated dropout rate (44.2% to 28.7%) are significant. These improvements would transform user experience at scale. Q-Learning achieves useful performance levels quickly — Precision@5 of 0.54 within the first 6–10 learner interactions — so it does not need an implausibly large amount of interaction data to provide useful personalization.

The suggestion quality was also driven by profile dimensionality. Precision@5 was 0.314 for a demographic-only profile (age, gender, education, field of study). Adding format, duration, and motivation data improved this to 0.527 (+67.8%). Dynamic real-time engagement signals raised it to 0.743 (+40.8% over Config. B, +136.6% over A). The monotonic improvement at each layer shows that each new dimension of the learner profile captures something that was missing, improving recommendation quality by adding something meaningful.

## 5. Discussion

### 5.1 What the Motivation Data Actually Tells Us

If you take SDT seriously, it shouldn't be surprise that 54.7% of learners who are motivated by external factors leave out, which is almost twice the rate of learners who are motivated by knowledge acquisition. But the effects on platform design are more extreme than what most platforms are ready to accept. This means that motivation should be a top priority in the enrollment funnel and the recommendation system, not an afterthought. If a platform knows that a learner signed up because their company required it, it should respond differently—by providing more scaffolding, more visibility into progress, and more social accountability features—than if it knows that the learner signed up because they are really interested in the subject. This isn't about punishing students who are driven by things outside of school or offering them worse service. It is about recognizing that individuals require various forms of help to sustain involvement. There is a lot of research on motivation in schools that shows that extrinsic motivation can turn into intrinsic motivation over time if the learning environment is set up to help with this, such as by giving students more freedom, making them feel competent, and connecting them to a learning community (Deci et al., 1999). A recommendation system that understands a learner's initial motivation profile can actively work to make these conditions happen.

### 5.2 Why Limited Engagement Bias Matters More Than People Think

The limited engagement bias result has effects that go beyond MOOC platforms. Any recommendation system that uses ratings, like those for movies, restaurants, or online courses, has the danger that the ratings will come mostly from people who didn't utilize the product much since they don't care as much about how accurate their judgment is. In the case of MOOCs, the pattern is made worse by the fact that inactive learners often miss the parts of the course that are truly valuable, such as the hard problem sets, the later lectures that require a lot of thought, and the peer discussion threads where real learning happens. The five-star rating is for the welcome video and the page that gives an overview of the course. The fix is easy to understand, but it needs to be done carefully: weight ratings based on how deep the involvement is. It should mean more to get a rating from a student who finished 80% of the course than from one who only finished 10%. An empirical question that platforms could adjust with held-out validation data is how much more. A basic binary weighing system—full weight for learners who finished at least half of the content and half weight for those who finished less—would greatly lessen the bias we saw.

### 5.3 The Early Warning Opportunity

The most practically significant finding in this study is the magnitude of the Week 1 effect sizes. A Cohen's *d* of 1.94 for Week 1 content consumption indicates a huge influence. In medical research, a Cohen's *d* above 0.8 is seen as "large" and often important for clinical practice. An effect size of almost 2.0 suggests that the distributions of people who finished and those who didn't are almost completely different. You can very accurately guess if a learner

will finish the course if you know how much content they consumed in Week 1. It is sad that platforms have access to this information in real time, as it is stored in their server logs, but most of them do nothing beneficial with it. If a student has only finished 8% of the Week 1 content by Day 7, they are almost certainly going to drop out. The platform could send them a personalized suggestion for a different course, a version of the current content that is easier to understand, or just a simple message saying, "We noticed you haven't had much time this week. Here is a shorter version of Module 1." Once the monitoring system is set up, these interventions don't cost much to put into place. The problem is not with the technology; it's with the structure. Someone needs to decide that preventing dropouts is a priority that needs to be worked on.

#### 5.4 Limitations

There are some problems with this study that should be honestly pointed out. First, the performance numbers for the Q-Learning framework originate from a validation dataset that was not used in a live randomized controlled deployment. The 61.3% expected completion rate is based on a model, and the actual results of the deployment could be different, either higher or lower. We have not yet reached the gold standard of a genuine RCT, where learners are randomly assigned to the Q-Learning system or the current platform recommendation.

Second, the dataset was put together mostly from English-language MOOC platforms, even if it was big. More and more people from all around the world are taking MOOCs, especially in India, China, Brazil, and Sub-Saharan Africa. It remains uncertain whether the motivational profiles, format preferences, and engagement patterns we recorded are applicable to these populations. There is substantial evidence to suggest that cultural factors—such as attitudes towards online learning, interactions with credentialing systems, and available study time—differ in ways that influence both the characteristics of learner profiles and the most effective recommendation tactics for various populations.

Third, we only look at completion and course rating as our outcome measures. We don't know if taking a MOOC course really led to learning new skills, getting a better job, or any other real-world result that would make the time and money spent worthwhile. This is a general deficit in MOOC research that has to be filled with long-term studies that also collect follow-up data.

#### 6. Conclusion

This research aimed to demonstrate that MOOC dropout is not inevitable; rather, it is a foreseeable, systematic failure resulting from a specific and rectifiable cause: recommendation systems that lack sufficient understanding of their users to effectively meet their needs. The data continually backs up this claim. Motivation is very important and can be measured. There are actual and systemic differences between students and professionals when it comes to their preferences for format and duration. Platforms already have indications that can tell them if a student is likely to drop out within seven days of enrolling. A Q-Learning-based recommendation framework that takes all of this seriously—one that creates rich, dynamic learner profiles and continuously modifies its suggestion strategy based on what it sees happening—works better than other methods by margins that are big enough to matter at scale. This work makes three important theoretical contributions: (1) the idea of limited engagement bias in course rating systems, which applies to any recommendation context where user engagement depth varies, not just MOOCs; (2) the use of Self-Determination Theory motivational constructs as active variables in a reinforcement learning recommendation architecture; and (3) the empirical demonstration of the engagement–relevance–completion virtuous cycle, where better recommendations lead to more engagement, which leads to better profile data, which leads to better recommendations. The most significant practical recommendation from this study does not pertain to algorithms. It's all about what matters most. Platforms need to see dropout as an issue that can be fixed through engineering, put money

into early warning systems that respond to Week 1 signs, and build the user profile infrastructure that makes personalization truly useful. There is already technology that can achieve this. We now think that the proof that it works is rather robust. What is left is the will of the organization to put it into action.

### References

1. Chen, M., Beutel, A., Covington, P., Jain, S., Belletti, F., & Chi, E. (2019). Top-K off-policy correction for a REINFORCE recommender system. *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, 456–464.
2. Deci, E. L., & Ryan, R. M. (1985). *Intrinsic motivation and self-determination in human behavior*. Springer.
3. Deci, E. L., & Ryan, R. M. (2000). The “what” and “why” of goal pursuits: Human needs and the self-determination of behavior. *Psychological Inquiry*, 11(4), 227–268.
4. Deci, E. L., Koestner, R., & Ryan, R. M. (1999). A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. *Psychological Bulletin*, 125(6), 627–668.
5. Hew, K. F., & Cheung, W. S. (2014). Students’ and instructors’ use of massive open online courses (MOOCs): Motivations and challenges. *Educational Research Review*, 12, 45–58.
6. Jordan, K. (2015). Massive open online course completion rates revisited: Assessment, length, and attrition. *The International Review of Research in Open and Distributed Learning*, 16(3), 341–358.
7. Kizilcec, R. F., Piech, C., & Schneider, E. (2013). Deconstructing disengagement: Analyzing learner subpopulations in massive open online courses. *Proceedings of the Third International Conference on Learning Analytics and Knowledge*, 170–179.
8. Milligan, C., & Littlejohn, A. (2014). Supporting professional learning in a massive open online course. *The International Review of Research in Open and Distributed Learning*, 15(5), 197–213.
9. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A., Veness, J., Bellemare, M., Graves, A., Riedmiller, M., Fidjeland, A., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533.
11. Shah, D. (2021). By the numbers: MOOCs in 2021. *Class Central*.
12. Yang, D., Sinha, T., Adamson, D., & Rose, C. P. (2013). Turn on, tune in, drop out: Anticipating student dropouts in massive open online courses. *Proceedings of the NIPS Data-Driven Education Workshop*, 1–9.
13. Zhao, X., Zhang, L., Xia, L., Tang, J., & Yin, D. (2018). Deep reinforcement learning for list-wise recommendations. *Proceedings of the 12th ACM Conference on Recommender Systems*, 95–103.