Open Access

Journal of Artificial Intelligence & Robotics

Article type: Review Article

Applying machine learning to student diversity in high school classrooms: A K-means clustering approach to learner profiling

Eric CK Cheng*

The Education University of Hong Kong, Hong Kong.

***Corresponding Author: Eric CK Cheng**

The Education University of Hong Kong, Hong Kong.

Email: eckcheng@eduhk.hk

Received: Sep 26, 2026

Accepted: Oct 19, 2026

Published Online: Oct 26, 2026

Website: www.joaiar.org

License: © **Cheng ECK** (2026). *This Article is distributed under the terms of Creative Commons Attribution 4.0 International License.*

Volume 3 [2026] Issue 2

Abstract

High school teachers face classes of twenty-five to forty students differing in learning speed, prerequisite knowledge and preference, yet can author only three to five instructional variants per unit. This paper proposes K-means clustering to assign students in each unit to data-derived learner profiles mapped to pre-authored packages governing tiered tasks, pacing, intervention allocation and modality choice. The method suits the problem because learner type is unlabelled, cluster assignments translate directly into teacher action, and computational demands match school capability. A pipeline integrating platform logs, diagnostics and engagement signals is specified, alongside evaluation combining cluster validity, stability, normalised learning gain, cost-sensitive equity metrics and teacher agreement.

Keywords: K-means clustering; Learner profiling; Differentiated instruction; Educational data mining; Adaptive learning.

Introduction

High school classrooms are structurally heterogeneous units in which a single teacher delivers a fixed curriculum, on a fixed timetable, to somewhere between twenty-five and forty adolescents who differ along at least three consequential dimensions. Learning speed varies in the number of practice attempts and the elapsed time each student requires to reach mastery of a standard. Background knowledge varies in the prerequisite skills students carry forward from earlier grades, feeder schools and out-of-school experiences. Learning preferences vary in preferred modality, topic interest, and engagement and affective patterns [1,2]. Conventional instruction delivers identical content at uniform pacing and therefore simultaneously under-serves struggling learners, who need additional scaffolding, and advanced learners, who need extension [3,4]. The documented consequence is persistent achievement gaps, disengagement and suboptimal outcomes across the whole class [5,6].

The operational difficulty in schools is not a lack of belief in differentiation but a lack of tractable differentiation. While it is infeasible for a teacher to author thirty-two individual learning paths per unit, authoring three to five instructional variants per unit and assigning students to them defensibly is achievable. This is precisely the decision that a machine learning model can inform.

The decision driven by the model output is a placement decision made at the start of each teaching unit, a cycle of roughly two to four weeks, in which the model assigns every student to a data-derived learner profile. Each profile is mapped by the department to a concrete, pre-authored instructional package, and the assignment drives four school actions. The first is flexible grouping and tiered tasks, meaning which of the differentiated task sets- core, scaffolded or extension- a student receives at the start of the unit. The second is pacing and mastery gating, meaning whether a student proceeds on the default schedule or is routed through additional prerequisite practice before the summative task [7]. The third is intervention allocation, meaning which students are placed in limited-capacity resources such as after-school tutoring, learning support periods or peer tutoring pairings. The fourth is resource and modality selection, meaning whether the student receives video first, text first or practice first materials, and which reinforcement schedule is applied [8,9].

Citation: Cheng ECK. Applying machine learning to student diversity in high school classrooms: A K-means clustering approach to learner profiling. J Artif Intell Robot. 2026; 3(2): 1043.

Profiles are recomputed each unit, so assignment remains provisional and reversible rather than hardening into a permanent stream. This is an important design requirement given the well-documented risk that algorithmic categorisation entrenches existing inequities [10].

Method application and suitability

The chosen technique is K-means clustering, an unsupervised method applied here to multidimensional student profiling. K-means partitions students into k groups by minimising within-group distance across an engineered feature vector spanning speed, prior knowledge and preference indicators. It is selected for four reasons.

First, the problem is definitionally unsupervised. There is no ground truth label for learner type. Learning speed, background knowledge and preference are latent constructs that must be discovered from behaviour rather than predicted against a known target. Zhang [10] applies K-means specifically for student profiling on multiple simultaneous characteristics in order to generate personalised pathways, and notes its scalability advantage, namely that systems can apply group-level adaptations while personalising within clusters. This matches the school's real constraint of a small number of authorable instructional variants.

Second, the method produces an output that maps directly onto teacher action. Unlike a continuous risk score, a cluster assignment is a discrete, interpretable group that a teacher can operationalise as a table, a task set or a tutoring roster. Cluster centroids can be reported as plain language profiles, such as slower to mastery, strong prerequisites and high video engagement, which supports the interpretability and teacher trust requirements that deep architectures struggle to meet [11,12].

Third, the approach is empirically precedented in adaptive education. Real-time analytics frameworks that categorise students and dynamically adjust learning materials are the operational pattern reported by Navyamol [8]. Beulah et al. [4] group learning objects against learner style models in order to drive personalised content selection, and Essa et al. [2] show that behaviour-based clustering of interaction data is preferable to static self-report instruments for identifying learning styles. Clustering is therefore the standard entry point for profile-driven personalisation.

Fourth, the method is proportionate to school capability. K-means is computationally cheap, runs on a laptop or a modest cloud instance, requires no labelled outcome data, and can be retrained overnight. This contrasts with Dynamic Key Value Memory Networks, Proximal Policy Optimisation, or transformer-based knowledge tracing, all of which require large interaction corpora and sustained engineering support [2,11]. Prior work using large-scale knowledge tracing depended on datasets of the order of 352,000 records and 17,000 students, a volume that a single high school will never hold.

Several practical constraints qualify the choice. The first concerns the selection of k and cluster validity, since K-means requires k to be specified in advance and assumes roughly spherical, similarly sized clusters. Educational data are rarely so well behaved, and the elbow and silhouette criteria may disagree, so k is constrained to between three and five for pedagogical usability and stability is validated before deployment. The second concerns feature scaling and dimensionality sensitivity, because distance-based clustering is dominated by high-variance features and unstandardised scales, such as minutes on task set against proportion correct, which will distort the partition. Standardisation and dimensionality reduction are therefore prerequisites rather than refinements. The third concerns instability and data latency, since cluster membership is volatile early in a term when interaction data are sparse, a cold start problem. Assignments should not be made before a minimum activity threshold is met, and profiles should be refreshed on the unit cycle rather than daily, to avoid churn that teachers cannot act upon.

The remaining constraints are ethical and organisational. Clusters trained on historical data can encode socioeconomic, language or disability related bias and thereby widen rather than close gaps [10], and they may also become de facto ability labels that lower teacher expectations. Mitigation requires excluding demographic attributes as clustering inputs, auditing cluster composition for disparate representation, and presenting profiles as instructional recommendations rather than student attributes. Interaction, affective and performance data are sensitive and subject to FERPA- and GDPR type constraints, requiring consent, data minimisation and access control. Finally, benefits do not materialise without professional development, since many educators lack preparation in data-driven instruction [5], and clustering optimises measurable constructs while underrepresenting creativity, collaboration and critical thinking. It must therefore augment rather than replace teacher judgement.

Data pipeline and feature engineering

Six raw sources are required. Learning management system and online practice platform logs supply item-level responses with timestamps, attempts, hint requests and resource clicks, and are high in volume, high in noise and available in near real time. The student information system contributes prior year grades, attendance, course enrolment and learning support flags, which are low in volume and high in reliability but updated only termly. Skill-tagged diagnostic pre-tests are administered at the start of each unit. Formative and summative assessments provide gradebook scores carrying rubric or standard tags. Preference and engagement signals include resource type, dwell time, modality choice, session timing and a short self-report modality survey. Optional affective proxies include idle time, rapid guessing and repeated abandonment, and their inclusion is justified by Hatta et al. [11], who found emotion level and quiz score to be the dominant features in personalisation decisions.

Records are joined on a pseudonymised student key. Rapid guess responses with implausibly short latencies and idle inflated durations are winsorised, and duplicate submissions and test accounts are removed. Students below a minimum activity threshold, for example fewer than twenty graded interactions in the window, are held out of clustering and handled by teacher judgement, which avoids unstable cold-start assignment. Data are then aggregated into a student-by-unit matrix, one row per student per teaching unit, so that the data structure matches the decision cadence. All features are z-standardised and principal component analysis is applied to collapse correlated behavioural features before clustering,

mitigating the distance distortion problem noted above. Sampling is stratified by grade and subject so that the largest cohort does not dominate clusters.

Three engineered feature blocks mirror the three diversity dimensions. The speed block comprises attempts to first correct response per standard, time to mastery measured as median minutes to reach three consecutive correct responses, a pace index expressing student time to mastery as a ratio of the class median, the slope of a fitted learning curve across the unit, hint request rate, and retention decay measured as the accuracy drop from unit end to a delayed check, which reflects the forgetting factors used in hybrid knowledge tracing models. The background knowledge block comprises the diagnostic pre-test skill vector, a prerequisite mastery ratio expressing mastered prerequisites as a proportion of those required by the unit's concept graph, the prior year grade in the subject, and a Bayesian Knowledge Tracing mastery probability per standard used as an input feature to clustering rather than as a standalone model [7]. The preference block comprises modality engagement shares for video, text and interactive practice as proportions of total resource time, optional task selection pattern by topic, a time of day and session length signature, persistence after error rate, self-reported modality preference, and a frustration proxy combining abandonment and rapid guessing [12,2].

Deriving the prerequisite ratio from a departmental concept graph is a deliberate design choice. Graph-encoded prerequisite structure has been shown to improve the pedagogical validity of sequencing decisions [4], and embedding it in the feature set means that clusters respect curricular structure rather than raw behaviour alone [13].

Model evaluation

Evaluation must address both cluster quality and instructional value, because a statistically clean partition is worthless if it fails to improve learning or if it distributes harm unevenly. Six families of metrics are therefore applied.

Internal cluster validity is assessed primarily by the silhouette coefficient, with a target above 0.35 for solutions of three to five clusters, corroborated by the Davies-Bouldin and Calinski-Harabasz indices, and with the elbow criterion on within-cluster sum of squares used to select k. These indices quantify whether the profiles are genuinely separated or merely arbitrary cuts through a continuum.

The Adjusted Rand Index measures stability between clusterings derived from bootstrap resamples and between clusterings from consecutive units. Low stability, conventionally an index below 0.6, would mean that teachers were regrouping students on noise, so stability functions as a deployment gate rather than an afterthought.

Predictive and construct validity are established through between-cluster differences in end-of-unit assessment scores and in time to mastery, reported as effect sizes using eta squared or Cohen's d, which confirm that the profiles capture instructionally meaningful variation. Cluster assignment should also remain a significant predictor of outcome over and above a prior grade baseline.

Instructional impact provides the decision-level metric. Since the clusters exist to drive differentiated instruction, the outcome measure is the normalised learning gain, calculated as the post-test score minus the pre-test score divided by the maximum score minus the pre-test score, compared between profile-driven differentiated classes and business-as-usual classes. This is the metric used to benchmark adaptive tutoring systems, where gains of 0.72 have been reported against 0.58 for rule-based and 0.45 for static instruction [14]. Complementary operational metrics include practice efficiency, meaning attainment per minute of practice, following the reported thirty-nine per cent efficiency improvement and thirty-three per cent duration reduction for concept-aware sequencing, together with the mastery attainment rate per standard.

Equity metrics are treated as cost-sensitive, because the dominant cost in this setting is asymmetric. Misassigning a capable student to a low expectation profile is more damaging than over-provisioning support. Evaluation therefore includes variance and gap reduction, measured as change in the standard deviation of attainment and in the gap between the lowest and highest quartile, since closing gaps is the stated objective [13]. It also includes representation parity, meaning the distribution of socioeconomic, English as an additional language and learning support students across clusters, monitored for disparate concentration [10]. Finally, it includes recall on at-risk identification, meaning the proportion of students who ultimately fail the unit who had been placed in a support-eligible profile, prioritised over precision because unnecessary support is cheaper than missed support.

Human validity is checked using Cohen's kappa between model profile assignment and independent teacher classification of the same students. Moderate agreement, in the region of 0.4 to 0.6, is the desired result, being high enough to establish credibility yet low enough to indicate that the model surfaces information teachers miss. Teacher-reported actionability is logged as a qualitative deployment criterion, consistent with the argument that augmenting rather than replacing teacher decision-making is the appropriate design goal for school-based artificial intelligence [5,12]. Taken together, these metrics align the model with the two objectives that matter to a school, which are maximising learning gain per unit of scarce instructional time and ensuring that personalisation narrows rather than widens the diversity-driven achievement gap.

Conclusion

Student diversity in high school classrooms is fundamentally a problem of allocation under scarcity, where the scarce resources are instructional time, authoring effort and intervention capacity. The contribution of machine learning here is not to determine how to teach but to make defensible the prior question of which students receive which of a small number of variants a department has already designed. K-means clustering suits that workflow because its output is structurally identical to the decision being made, since a partition into three to five groups is what a teacher can act upon, a centroid described in plain language is what a teacher can trust, and a model that trains overnight without labelled outcomes is what a school can operate. Deep knowledge tracing and policy learning approaches offer finer adaptation but presuppose data volumes and engineering support a single secondary school cannot absorb.

The surrounding design choices matter more than the algorithm. Aggregating to a student-by-unit matrix aligns model resolution with the instructional cycle and prevents group churn that would destroy teacher confidence. Standardisation and dimensionality reduction stop the partition being determined

by whichever feature carries the largest variance. A minimum activity threshold keeps cold-start students with teacher judgement. A prerequisite mastery ratio drawn from a concept graph gives profiles curricular rather than merely behavioural meaning. Recomputing every unit keeps assignment provisional, which is the strongest safeguard against clustering hardening into permanent streaming.

Evaluation is correspondingly two-tiered. Validity and stability indices establish only that profiles are real and reproducible, which is necessary but insufficient. The decision-level question is whether profile-driven differentiation raises normalised learning gain while reducing attainment dispersion and quartile gaps. Because misassigning a capable student to a low expectation profile costs more than redundant support, recall on at-risk identification is prioritised over precision, and cluster composition is audited continuously for disparate representation. Moderate rather than high agreement with teacher classification is the target, since a model reproducing existing intuitions adds nothing while one diverging entirely forfeits credibility.

Three limitations bound the claim. Clustering captures only the measurable, leaving creativity, collaboration and critical reasoning to teacher judgement, so the system augments rather than replaces professional decision-making. Excluding demographic attributes suppresses direct discrimination without eliminating proxy effects, making composition audits a permanent operating requirement. Benefits also depend on professional development and on consent, minimisation and access control obligations that are organisational rather than technical. Future work should examine profile trajectories longitudinally, compare Gaussian mixture and hierarchical alternatives that relax the spherical cluster assumption, and test whether the profile-to-package mapping can be refined from outcomes without reintroducing prohibitive data demands. The broader conclusion holds regardless. A simple, transparent and frequently recomputed clustering model, embedded in a pipeline respecting curricular structure and an evaluation regime treating equity and stability as deployment gates, is a realistic mechanism for converting recognition of student diversity into instruction that responds to it.

References

1. Yadav S, Maurya M, Chaudhary R. Machine learning in education: personalised learning and educational technology. Pharma Innovation. 2019.
2. Personalised adaptive learning technologies based on machine learning techniques to identify learning styles: a systematic literature review. IEEE Access. 2023; 11: 48392-48409.
3. Khazanchi R, Khazanchi P. Artificial intelligence in education: a closer look into intelligent tutoring systems. 2021: 256-277.
4. Beulah C, Latha C, Kirubakaran E. Personalised learning path delivery in web-based educational systems using a graph theory-based approach. 2013.
5. Alawneh YJJ, Sleema H, Salman FN, Alshammat MF, Oteer RS, Alrashidi NKN. Adaptive learning systems: revolutionising higher education through AI-driven curricula. 2024; 1: 1-5.
6. Jose BC, Kumar M, UdayaBanu T, Nagalakshmi M. Assessing the effectiveness of adaptive learning systems in K-12 education. 2024; 01.
7. Hou WS, Li Y. Intelligent algorithms to enhance education in the United States. 2025.
8. Navyamol KT. Personalised learning pathways using machine learning in educational portals. 2025.
9. Satyashree, Manjunath C. NeuroAI: an AI-powered personalised learning tutor for students with diverse learning challenges. 2026.
10. Zhang Q. Harnessing artificial intelligence for personalised learning pathways: a framework for adaptive education management systems. Applied and Computational Engineering. 2024; 82: 167-172.
11. Hatta M, et al. Actor-critic reinforcement learning for personalised STEM learning path optimisation. Journal of Information Systems and Informatics. 2025; 7: 2780-2802.
12. Janardhana DR, Arif H, U HB, Rukmini T, Hanumanthappa H, Gaikwad V. Personalised learning systems using AI enhancing adaptive education. International Journal of Environmental Sciences. 2025: 668-676.
13. Xu Q. Design of an artificial intelligence-based personalised learning system. 2025: 424-429.
14. Guzman ST, Cruz-Mercado MCD. Developing an intelligent tutoring system for personalised skill development using reinforcement learning. 2025.