Action Rules Mining for Student Persistence at a Public HBCU

Willie Moore and Li-Shiang Tsay
North Carolina A&T State University, Greensboro, USA

Wdmoore@aggies.ncat.edu
ltsay@ncat.edu

Abstract: Predictive student-success analytics identify students at risk but often do not specify candidate actions. This study applies Rule-Based Action Rules Mining (RB-ARM) with checkpoint-based temporal feature engineering to student persistence and term academic standing at a public Historically Black College and University. The integrated dataset contains 39,413 student-term records combining student information, advising appointments and check-ins, and online tutoring logs. Engagement counts were encoded within three academic-calendar windows. Rules were discovered on an oversampled training partition using minimum support and confidence thresholds. Interestingness was assessed with observational uplift and action rule lift on original-training, temporal validation, and hold-out test partitions. Three persistence action rules identified candidate changes from no late-term (week 9 through end of term) advising appointments to exactly one or two or more appointments among first-year undergraduates. For the rule proposing two or more appointments, test-partition persistence rates were 69.8% in the pre-action group and 88.3% in the post-action group, yielding observational uplift of 18.5 percentage points and an action rule lift of 1.26. These comparisons describe associations between student-term groups rather than causal effects or within-student changes. No qualifying action rules were discovered for term academic standing, online tutoring, or check-in engagement under the configured mining requirements. The study contributes an interpretable framework linking candidate support-service actions to specific student contexts and within-term windows. It also documents constraints on rule formation involving recorded engagement density and compatible classification-rule pairs. The framework is transferable in principle, while the institution-specific candidate actions require prospective evaluation before operational use.

Keywords: Action rules mining, Educational data mining, HBCU, Prescriptive analytics, Student persistence, Temporal feature engineering

1. Introduction

Historically Black Colleges and Universities (HBCUs) occupy a critical position in the United States higher education ecosystem, yet their key persistence and graduation indicators consistently trail comparable non-HBCU institutions on publicly reported rates (Gordon et al. 2021). This gap carries both equity and fiscal consequences as performance-based funding models increase attention to measurable student success outcomes (Miller 2022). HBCUs therefore require decision-support tools that move beyond descriptive dashboards and predictive risk scores toward prescriptive analytics capable of recommending specific, feasible, time-bounded interventions.

Predictive models such as random forests and gradient-boosting machines can identify which students are at risk, but they do not articulate what should be done differently (Dol and Jawandhiya 2023; Barbeiro et al. 2024). This “prediction-to-action gap” is particularly acute at HBCUs, where limited advising ratios make untargeted outreach inefficient and where each misallocated intervention represents a real opportunity cost (Richards and Awokoya 2012). Action Rules Mining (ARM) builds on the action-rule concept introduced by Raś and Wieczorkowska (2000), with a comprehensive treatment of its foundations and algorithms provided by Dardzińska (2013). ARM addresses this gap by discovering rules of the form: given a fixed stable context, transitioning a flexible attribute from state α to state β is associated with transitioning the outcome from undesirable to desirable. Each rule is annotated with support, confidence, and lift, providing transparent metrics for institutional decision-making.

Despite applications in healthcare (Bagavathi and Tzacheva 2019), software quality assurance (Oueslati et al. 2024), and customer retention (Duan and Raś 2022), ARM has seen minimal application in education and virtually none in HBCU contexts. This paper addresses that gap with an RB-ARM pipeline with checkpoint-based temporal feature engineering for student success analytics at a public HBCU. The pipeline is informed by Astin's Input–Environment–Outcome (I-E-O) framework (Astin 1993), using student background and prior academic characteristics as input measures, support-service engagement as a measured dimension of the college environment, and term-to-term persistence and term academic standing as outcomes.

The contributions of this paper are:

1. A documented, reproducible RB-ARM pipeline integrating multiple institutional data sources, with checkpoint-based CP-window feature engineering, profile-specific mining thresholds, and re-evaluation across multiple data partitions.
2. Empirical evidence that RB-ARM can discover persistence-oriented action rules in a public-HBCU context.
3. Documentation of boundary conditions under which ARM produces null results in educational data, specifically participation density floors and asymmetric conditional structures.
4. A pipeline transferable in principle to institutions with student-linked, time-stamped support-interaction records aligned to academic term calendars.

2. Related Work

2.1 Student Success Analytics

Educational Data Mining (EDM) has produced substantial work on predicting student outcomes from student information system (SIS) data, learning management system (LMS) telemetry, and co-curricular engagement records (Dol and Jawandhiya 2023; Alyahyan and Düştegör 2020; Rahman et al. 2021). Decision trees, random forests, gradient-boosting machines, and neural networks have forecast course failure, stop-out, persistence, and graduation (Barbeiro et al. 2024; Alwarthan et al. 2022; Vaarma and Li 2024). However, these models terminate at risk-score generation, leaving practitioners without specific intervention recommendations (Miller 2022; Santos and Henriques 2023).

The analytics maturity continuum, from descriptive and diagnostic to predictive and prescriptive (Govindarajan 2021; Lepenioti et al. 2020; Wolniak and Grebski 2023), frames the gap this paper addresses. ARM operates at the prescriptive level by embedding a recommended attribute transformation within every discovered rule.

2.2 Action Rules Mining

Action rules were introduced by Raś and Wieczorkowska (2000). Subsequent developments include tree-structured rule-based discovery systems (Tsay and Raś 2005; Raś et al. 2009), methods incorporating meta-actions (Raś and Dardzińska 2009; Tzacheva and Raś 2010), and scalable distributed implementations (Bagavathi and Tzacheva 2019). Tzacheva et al. (2022) provide a survey of actionable pattern discovery methods, including approaches for large datasets and distributed computing. The ActionRulesDiscovery library (Sýkora and Kliegr 2020) provides an open-source Python implementation of rule-based action rule mining. In this approach, an action rule pairs two classification rules sharing a stable context but differing in the values of one or more flexible attributes, with one rule predicting the undesirable outcome and the other predicting the desirable outcome. The resulting rule identifies candidate changes for a specified subgroup. Support and confidence summarize empirical rule strength, while observational uplift quantifies differences in outcome rates between groups matching the pre-action and post-action conditions.

Educational applications of ARM remain limited; one example is Liu et al. (2012), who proposed an action-rule discovery approach incorporating pruning through micro-actions and demonstrated its application in education. The present study applies RB-ARM in an HBCU context, combining checkpoint-based temporal feature engineering, profile-specific threshold design, and multi-partition validation.

2.3 ARM at HBCUs

To our knowledge, no prior study has applied ARM to student success data at an HBCU. The HBCU context introduces specific considerations, including smaller enrollment, high proportions of Pell-eligible and first-generation students (Gordon et al. 2021; Richards and Awokoya 2012), and lower recorded participation density in some support-service streams. These characteristics constrain the flexible attribute space and what ARM can discover, making boundary-condition documentation particularly important here.

3. Methodology

3.1 Research Design and Theoretical Framework

This study follows the Cross-Industry Standard Process for Data Mining (CRISP-DM) methodology (Chapman et al. 2000; Shearer 2000) and is grounded in Astin's I-E-O framework (Astin 1993). Stable attributes (demographics, academic preparedness) map to Inputs; flexible attributes (support-service engagement counts) map to Environment; and persistence and term GPA map to Outcomes. ARM discovers ΔF transformations (changes in

Environment) associated with improved Outcomes while conditioning on fixed Inputs. This paper presents a condensed journal version of the RB-ARM with checkpoint-based temporal feature engineering framework and evaluation reported in (Author 2026).

3.2 Data Sources and Study Population

De-identified longitudinal student-term records were assembled from three institutional sources at a public HBCU: (1) the student information system (SIS), providing enrollment, demographic, academic performance, and persistence data; (2) a student success management system (SSS), providing advising appointment and front-desk check-in records; and (3) an online tutoring platform (E-Tutoring), providing tutoring session logs. Records were linked at the student–term level using anonymized identifiers and institutional term codes.

The integrated dataset includes 39,413 student–term records representing 11,231 unique undergraduates across eight terms from Fall 2022 through Spring 2025. Summer terms were retained in the integrated dataset but excluded from the mining and evaluation partitions due to compressed calendars and reduced CP-window exposure.

These engagement streams were selected because they produced structured event-level records suitable for temporal aggregation. Other engagement types, including advisor case notes and early-alert resolutions, were excluded because they were qualitative or not consistently translatable into countable flexible attributes.

3.3 Temporal Feature Engineering: Cut-Point Windows

Support-service engagement variables were temporally encoded using within-term cut-point (CP) windows:

- CP1: weeks 1–3 (early term)
- CP2: weeks 4–8 (mid-term)
- CP3: weeks 9–end of term (late term; week numbers > 20 rolled into CP3)

Week numbers were computed from official term start dates. For each student–term record, SSS appointment, SSS check-in, and E-Tutoring counts were aggregated within each CP window, producing nine CP-windowed count variables. Before mining, counts were discretized into {0, 1, 2plus} categories to address right-skewed distributions.

The temporal dimension of this study is the checkpoint-based feature construction, where engagement is aggregated into academic-calendar-aligned windows (CP1, CP2, CP3). The rule-generation algorithm itself does not model cross-term sequences or transitions; genuine temporal action rule mining is reserved for future work.

3.4 Variable Roles

Variables were assigned to three roles:

- Stable: Demographics, academic preparedness indicators, class level, cumulative GPA bin, enrollment status, and two dynamically constructed course-signature features (18 attributes when both course-signature features are present; 17 when one is absent due to term composition).
- Flexible: Nine CP-windowed support-service engagement counts plus online_courses_ratio_bin (10 total). Three academic-outcome variables (term credits passed ratio, term credits earned, failed courses count) were excluded as outcomes rather than prescribable interventions.
- Decision: Examined twice, once for persistence (persisted_flag: returned next term or graduated) and once for term academic standing (term GPA ≥ 2.0 vs. < 2.0).

For persistence mining, SSS Apt Any CP3 is the discretized count of SSS advising appointments of any type during the CP3 window. SSS check-in variables (SSS Chk Any CP1, CP2, CP3) were mined as separate flexible attributes within the same profile and produced no rules. An attribute is considered prescribable if an institutional actor (e.g., an advisor or tutoring coordinator) can influence its value through a specific action within the term.

3.5 Temporal Data Partitioning

The dataset was partitioned using a primary temporal split:

- Training: Fall 2022, Spring 2023, Fall 2023 (15,688 student–term records; base persistence rate 0.846)
- Validation: Spring 2024 (5,150 records; base rate 0.818)
- Test: Fall 2024 (5,751 records; base rate 0.883)

Spring 2025 remains as a terminal labeling term (confirming Fall 2024 persistence outcomes) but is not used for mining or evaluation, since subsequent-term return cannot be observed.

3.6 Class Imbalance Handling

The training partition exhibited substantial class imbalance (84.6% persistence rate). Term-aware random oversampling was applied exclusively to the training partition, duplicating real minority-class records (capped at 5 additional copies per original record) within each term. The resulting oversampled training set contained 24,295 rows with a persistence rate of 54.6%. Random oversampling by duplication was selected over synthetic methods such as SMOTE because the ARM framework operates on fully discretized categorical data where interpolation is undefined, and because preserving real attribute combinations is essential for rules that must correspond to observed student patterns (He and Garcia 2009).

For the GPA outcome, a separate oversampled training set was created using the same procedure (24,506 rows; 50.8% good-standing rate). Oversampling was never applied to validation or test partitions.

3.7 Mining Configuration

Three profiles were configured to examine advising and check-in engagement, online tutoring engagement, and course modality. Minimum support thresholds were set separately for each profile to accommodate differences in the frequency of observed attribute values. The lowest threshold was assigned to the online tutoring profile because recorded engagement was sparse. Each profile was mined across five staged configurations, described below. Minimum support thresholds apply to classification-rule antecedents (stable conditions and flexible states). Reported action rule support uses left support, sup(r1), including the undesirable outcome.

Table 1: Mining Profiles

Profile	Flexible Attributes	Support	Confidence	Stages
Advising Appointments and Check-ins	Six attributes: advising appointment counts and front-desk check-in counts, each measured across three checkpoint windows	4%	70%	5
Online Tutoring Engagement	Three attributes: online tutoring session counts across three checkpoint windows	2%	70%	5
Online Course Enrollment	One attribute: categorized proportion of enrolled courses taken online	5%	70%	5

Each profile was mined under five configurations with increasing limits on rule complexity, expressed as (maximum stable attributes, maximum flexible attributes): (1,1), (1,2), (2,2), (3,2), and (3,3). Rules with identical stable conditions, flexible transitions, and outcomes were counted once across configurations.

3.8 Action Rule Metric Definitions

RB-ARM pairs classification rules predicting undesirable (r1) and desirable (r2) outcomes under shared stable conditions (Dardzińska 2013; Sýkora and Kliegr 2020):

- Left support sup(r1): count matching r1's conditions and undesirable outcome.
- Right support sup(r2): count matching r2's conditions and desirable outcome.
- Action rule support: sup(r1); confidence: conf(r1) × conf(r2).
- conf(r1) = P(undesirable | context, pre-action state); conf(r2) = P(desirable | context, post-action state).

For persistence, r1 predicts non-persistence, so P_pre = 1 − conf(r1); P_post = conf(r2).

- Observational uplift: 100 × (P_post − P_pre) percentage points.
- Action rule lift: P_post / P_pre (P_pre > 0).

R1 (Table 3): conf(r1) = 0.787 implies P_pre = 0.213; with P_post = 0.801, uplift is 58.8 pp and lift 3.76. These are observational group comparisons, not causal effects or within-student changes.

3.9 Proof-of-Concept: Association Rule Screening

Before ARM, a preliminary association-rule analysis explored persistence patterns involving support-service engagement within stable contexts. Apriori was applied to the support-user subset across all terms (21,460 records), using 5% minimum support and 60% minimum confidence. The analysis generated 75 classification-style association rules of the form context + CP-windowed support engagement → persisted_flag = 1. Several discovered rules included late-term (CP3) engagement. These exploratory findings informed the subsequent ARM analysis, which examined engagement transitions within shared stable contexts using CP-windowed attributes.

3.10 Cross-Partition Evaluation Protocol

After discovery on the oversampled training partition, each rule was re-evaluated without re-mining on three natural-distribution datasets: the original unoversampled training partition, the temporal validation partition, and the temporal hold-out test partition. No rule discovery or parameter tuning was performed on the validation or test partitions.

4. Results

4.1 Persistence Rules

The persistence mining run produced three unique action rules after deduplication. All three were generated by the SSS_arm profile and involve SSS advising appointments in the CP3 window for first-year undergraduate students. The E-Tutoring and OTHR profiles produced zero rules.

Table 2: Discovered Persistence Action Rules

#	Narrative action rule
R1	For first-year undergraduates, the candidate action is to move from no recorded CP3 advising appointments to exactly one appointment, targeting a change from non-persistence to persistence.
R2	For first-year undergraduates, the candidate action is to move from no recorded CP3 advising appointments to two or more appointments, targeting a change from non-persistence to persistence.
R3	For newly admitted first-year undergraduates, the candidate action is to move from no recorded CP3 advising appointments to two or more appointments, targeting a change from non-persistence to persistence.

Note: CP3 covers weeks 9 through the end of the term. Counts include all recorded SSS advising appointment types. The proposed changes reflect observational associations between student-term groups, not demonstrated causal effects.

4.2 Development Metrics

Table 3: Persistence Rules: Oversampled Development Metrics

#	Sup(r1): Left support (AR support)	Sup(r2): Right support	P_pre	P_post	Uplift	Lift	Conf(r1)
R1	2,072	875	0.213	0.801	58.8pp	3.76	78.7%
R2	2,072	974	0.213	0.883	67.0pp	4.14	78.7%
R3	1,747	878	0.207	0.882	67.5pp	4.26	79.3%

On oversampled data, approximately 21% of first-year students with no CP3 advising appointments persisted (P_pre), versus 80–88% of those with CP3 appointments (P_post). These metrics overstate the true effect because oversampling inflates the non-persister population, deflating P_pre and inflating uplift.

4.3 Cross-Partition Validation

Table 4: Rule R2 Cross-Partition Validation

Partition	N	P_pre	P_post	Uplift	Lift	Sup(r1)	Sup(r2)
Oversampled Train	24,295	0.213	0.883	67.0pp	4.14	2,072	974
Original Train	15,688	0.555	0.971	41.6pp	1.75	449	974
Validation	5,150	0.722	0.774	5.2pp	1.07	54	236
Test	5,751	0.698	0.883	18.5pp	1.26	70	226

Four key patterns emerge:

Directional consistency. P_post exceeds P_pre and lift exceeds 1.0 on all four partitions, including both temporally out-of-sample hold-out sets.

Expected attenuation. Uplift decreased from 67.0pp (oversampled) to 41.6pp (original training) and was lower on both hold-out partitions: 5.2pp on validation and 18.5pp on test. This attenuation is driven primarily by changes in P_pre, which is artificially deflated by oversampling (0.213) relative to its natural value (0.555).

Post-action persistence rates. For two or more CP3 advising appointments, P_post ranged from 0.774 to 0.971 across partitions.

Small-cell limitation. Action rule support (left support) is modest on the hold-out partitions (54 validation, 70 test records). However, these are not the denominators for P_pre; the implied pre-action denominators are approximately 194 (validation) and 232 (test). The corresponding 95% Wilson intervals for P_pre are approximately 0.655–0.780 (validation) and 0.636–0.754 (test). Validation and test uplift estimates should therefore be interpreted with caution, with uncertainty estimated directly for the difference P_post − P_pre rather than inferred from left support alone.

4.4 Term GPA Outcome

The GPA mining run produced zero action rules across all profiles and stages. No stable subgroup produced both a left-side and right-side classification rule meeting the 70% confidence threshold.

Table 5: P(GPA ≥ 2.0 | SSS APT Any CP3 Engagement) by Class Level

Class Level	P(good | engaged)	n(engaged)	P(good | not engaged)	P(bad | not engaged)	Left rule ≥ 70%?	Right rule ≥ 70%?
First Year	0.469	3,152	0.181	0.819	✓	✗
Second Year	0.546	2,201	0.338	0.662	✗	✗
Third Year	0.612	2,371	0.444	0.556	✗	✗
Fourth Year	0.684	2,446	0.629	0.371	✗	✗

Note: This table presents an illustrative diagnostic slice by class level. ARM mining evaluated the full multivariate stable-context space, not class level alone. "Left rule ≥ 70%?" indicates whether the undesirable-outcome classification rule (P(bad | not engaged)) meets the 70% confidence threshold. "Right rule ≥ 70%?" indicates whether the desirable-outcome classification rule (P(good | engaged)) meets the 70% confidence threshold. Both sides must meet the threshold within the same stable context for an action rule to form.

The diagnostic reveals an asymmetric conditional structure. Among first-year students, the left-side threshold is met (P(bad|not engaged) = 0.819 > 0.70), but the right side fails (P(good|engaged) = 0.469 < 0.70). Even with

engagement, fewer than half of first-year students achieve good academic standing. Among fourth-year students, the right side approaches but does not reach the threshold (0.684), while the left side is far below (0.371). No subgroup simultaneously satisfies both requirements.

Within individual subgroups, the difference in good-standing probability between engaged and non-engaged students can be substantial, up to 29 percentage points for first-year students, but these do not manifest as ARM-feasible rule pairs because both classification rules must simultaneously exceed the confidence threshold within the same stable context.

4.5 Null Findings for E-Tutoring and Check-Ins

No action rules involving E-Tutoring sessions or SSS check-ins were discovered for either outcome. Within the selected E-Tutoring data stream, CP1 had only 0.54% nonzero records, CP2 had 3.73%, and CP3 had 5.48% in the oversampled training data. Sparse engagement may limit antecedent support within stable contexts under the tutoring profile's 2% minimum support threshold (Table 1). The null finding indicates a boundary condition for this specific data stream at its observed density within the current analytic scope, not a conclusion about the effectiveness or availability of tutoring support.

5. Discussion

5.1 Practical Implications

The discovered persistence rules identify a specific, operationalizable pattern: first-year undergraduates who engage with at least one advising appointment in the CP3 window are observationally associated with substantially higher persistence rates. For R2, involving two or more appointments, the hold-out test partition shows the same directional association (uplift 18.5 percentage points, lift 1.26).

This finding is consistent with retention literature identifying the first year as the highest-risk period for departure (Tinto 1993; Pascarella and Terenzini 2005) and research emphasizing sustained advisor–student contact at HBCUs (Richards and Awokoya 2012). The present study extends this literature by identifying a specific temporal and appointment-count pattern within a prescriptive rule framework.

The concentration of rules in CP3 (weeks 9 through end of term) warrants careful interpretation. CP3 advising may serve as a re-engagement touchpoint for at-risk students, or may mark already-engaged students more likely to persist. The absence of CP1 and CP2 rules may reflect participation density constraints rather than an absence of early-term effects. The observational design cannot distinguish these mechanisms.

5.2 Methodological Implications

The four-level evaluation framework demonstrates that oversampling inflates development-set metrics systematically and predictably. ARM studies commonly report metrics computed on the same dataset used for mining (e.g., Raś and Wieczorkowska 2000; Dardzińska 2013; Sýkora and Kliegr 2020), conflating discovery-phase artifacts with population-level estimates. The multi-partition approach separates these concerns and provides transparent evidence of attenuation.

The null findings for tutoring and check-ins, alongside the asymmetric conditional structure for GPA, illustrate constraints on ARM applications. Sparse recorded engagement may limit antecedent support within stable contexts, while action-rule formation also requires compatible classification rules meeting the confidence thresholds.

5.3 The GPA Null Finding

The class-level diagnostic in Table 5 illustrates a classification-rule confidence constraint. Among first-year students, non-engagement met the left-side threshold, but the pooled engaged-group good-standing rate (0.469) fell below the right-side threshold. Whether prior academic performance explains this pattern should be evaluated in future work.

5.4 Data Environment Constraints

A central finding is that ARM's prescriptive power is constrained by the density and breadth of the flexible attribute space. Sufficient recorded engagement density within the selected flexible attributes is a prerequisite for rule formation. The null findings for E-Tutoring and check-in variables reflect the density of recorded engagement within the selected streams during the analytic window, not a characterization of institutional support capacity. This constraint is specific to the data streams and time period analyzed.

5.5 Limitations

The most significant limitation is the observational, retrospective design. Students who attend advising appointments may differ systematically from those who do not, so the observed association may partly reflect self-selection rather than a causal effect.

The findings are based on data from a single public HBCU. The pipeline is portable, but the specific rules are context-dependent.

Hold-out action rule support (left support) is small (54–70); statistical precision depends on the full matching-group denominators. All discovered rules involve CP3 (late-term). Earlier-window rules would be more actionable, but their absence may reflect sparse recorded engagement or other mining constraints.

The GPA outcome was operationalized as binary (≥ 2.0 vs. < 2.0), which collapses meaningful within-range variation.

6. Conclusion

This paper presented a Rule-Based Action Rules Mining pipeline with checkpoint-based temporal feature engineering for prescriptive student success analytics at a public HBCU. The pipeline integrates multiple institutional data sources, applies CP-window temporal feature engineering, uses profile-specific mining thresholds, and validates discovered rules across temporally out-of-sample partitions.

For persistence, three action rules were discovered and re-evaluated across original-training, validation, and test partitions, all involving SSS advising appointments in the CP3 window for first-year undergraduates. Rule R2 exhibited directionally consistent positive associations across all four evaluation partitions, with a hold-out test uplift of 18.5 percentage points and a lift of 1.26.

For term academic standing, no rules were discovered. The diagnostic analysis revealed an asymmetric conditional structure that prevents both classification rules from simultaneously meeting the confidence threshold within any stable subgroup.

The framework is transferable in principle to institutions with student-linked, time-stamped support-interaction data aligned to academic term calendars. The specific rules are local; the method is general. Future work includes expanding the flexible attribute space to incorporate additional engagement data streams, longitudinal multi-term rule mining, and prospective validation of discovered rules through controlled pilot studies.

Ethics Declaration: This article is derived in part from the author's doctoral dissertation (Author 2026). This retrospective study of de-identified student records was approved by the Institutional Review Board of the author's enrolled university.

AI Declaration: The author used generative AI tools for editorial assistance in revising wording, organization, and clarity. All technical content, calculations, interpretations, and references were reviewed and verified by the author, who assumes full responsibility for the manuscript.

References

Alwarthan, S., Aslam, N. and Khan, I. U. (2022) "An explainable model for identifying at-risk student at higher education", IEEE Access, Vol 10, pp 107649–107662. doi:10.1109/ACCESS.2022.3211070.

Alyahyan, E. and Düştegör, D. (2020) "Predicting academic success in higher education: Literature review and best practices", International Journal of Educational Technology in Higher Education, Vol 17, No. 1, Art. 3.

Astin, A. W. (1993) What Matters in College? Four Critical Years Revisited, Jossey-Bass, San Francisco, CA.

Author (2026) [Doctoral dissertation], details withheld for blind review.

Bagavathi, A. and Tzacheva, A. A. (2019) "Scalable action mining for recommendations to reduce hospital readmission", Proceedings of the IEEE 20th International Conference on Information Reuse and Integration (IRI), pp 159–166.

Barbeiro, L., Gomes, A., Correia, F. B. and Bernardino, J. (2024) "A review of educational data mining trends", Procedia Computer Science, Vol 237, pp 88–95. doi:10.1016/j.procs.2024.05.083.

Chapman, P., Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., Shearer, C. and Wirth, R. (2000) CRISP-DM 1.0: Step-by-Step Data Mining Guide, CRISP-DM Consortium..

Dardzińska, A. (2013) Action Rules Mining, Studies in Computational Intelligence, Vol 468, Springer, Berlin, Heidelberg. doi:10.1007/978-3-642-35650-6

Dol, S. M. and Jawandhiya, P. M. (2023) "A review of data mining in education sector", Journal of Engineering Education Transformations, Vol 36, No. S2, pp 13–23. doi:10.16920/jeet/2023/v36is2/23003.

Duan, Y. and Raś, Z. W. (2022) "Developing customer attrition management system: Discovering action rules for making recommendations to retain customers", Applied Intelligence, Vol 53, No. 9, pp 10485–10501.

Gordon, E. K., Hawley, Z. B., Kobler, R. C. and Rork, J. C. (2021) "The paradox of HBCU graduation rates", Research in Higher Education, Vol 62, No. 3, pp 332–358. doi:10.1007/s11162-020-09598-5.

Govindarajan, M. (2021) "Educational data mining techniques and applications", in Advancing the Power of Learning Analytics and Big Data in Education, IGI Global, Hershey, PA, pp 234–251.

He, H. and Garcia, E. A. (2009) "Learning from imbalanced data", IEEE Transactions on Knowledge and Data Engineering, Vol 21, No. 9, pp 1263–1284. doi:10.1109/TKDE.2008.239.

Lepenioti, K., Bousdekis, A., Apostolou, D. and Mentzas, G. (2020) "Prescriptive analytics: Literature review and research challenges", International Journal of Information Management, Vol 50, pp 57–70. doi:10.1016/j.ijinfomgt.2019.04.003.

Liu, Y., Zhao, X. and Zhang, Y. (2012) "Action rules mining triggered by micro-actions and its application in education", in Sambath, S. and Zhu, E. (eds.) Frontiers in Computer Education, Advances in Intelligent and Soft Computing, Vol 133, Springer, Berlin, Germany, pp 57–64. doi:10.1007/978-3-642-27552-4_10.

Miller, T. E. (2022) "A case study of an institution continuing to advance student success, the University of South Florida", Strategic Enrollment Management Quarterly, Vol 10, No. 2, pp 49–58.

Oueslati, K., Laberge, G., Lamothe, M. and Khomh, F. (2024) "Mining action rules for defect reduction planning", Proceedings of the ACM on Software Engineering, Vol 1, No. FSE, Art. 102, pp 2309–2331. doi:10.1145/3660809.

Pascarella, E. T. and Terenzini, P. T. (2005) How College Affects Students: A Third Decade of Research, Vol 2, Jossey-Bass, San Francisco, CA.

Rahman, M., Ahmed, M. I. and Hossain, M. S. (2021) "Analysis of student's achievement through educational data mining", Proceedings of the 2021 International Conference on Information Systems and Advanced Technologies (ICISAT), pp 1–5. doi:10.1109/ICISAT54145.2021.9678406.

Raś, Z. W. and Dardzińska, A. (2009) "Action rules discovery based on tree classifiers and meta-actions", in Rauch, J., Raś, Z. W., Berka, P. and Elomaa, T. (eds.) Foundations of Intelligent Systems: ISMIS 2009, Lecture Notes in Artificial Intelligence, Vol 5722, Springer, Berlin, Heidelberg, pp 66–75. doi:10.1007/978-3-642-04125-9_10.

Raś, Z. W. and Wieczorkowska, A. (2000) "Action-rules: How to increase profit of a company", in Principles of Data Mining and Knowledge Discovery: PKDD 2000, Lecture Notes in Computer Science, Vol 1910, Springer, pp 587–592.

Raś, Z. W., Tsay, L.-S. and Dardzińska, A. (2009) "Tree-based algorithms for action rules discovery", in Mining Complex Data, Studies in Computational Intelligence, Vol 165, Springer, pp 153–163. doi:10.1007/978-3-540-88067-7_9.

Richards, D. A. R. and Awokoya, J. T. (2012) Understanding HBCU Retention and Completion, Frederick D. Patterson Research Institute, UNCF, Fairfax, VA.

Santos, R. M. and Henriques, R. (2023) "Accurate, timely, and portable: Course-agnostic early prediction of student performance from LMS logs", Computers and Education: Artificial Intelligence, Vol 5, Art. 100175. doi:10.1016/j.caeai.2023.100175.

Shearer, C. (2000) "The CRISP-DM model: The new blueprint for data mining", Journal of Data Warehousing, Vol 5, No. 4, pp 13–22.

Sýkora, L. and Kliegr, T. (2020) "Action rules: Counterfactual explanations in Python", RuleML+RR 2020 Challenge, CEUR Workshop Proceedings, Vol 2644, Aachen, Germany, pp 28–41.

Tinto, V. (1993) Leaving College: Rethinking the Causes and Cures of Student Attrition, 2nd ed., University of Chicago Press, Chicago, IL.

Tsay, L.-S. and Raś, Z. W. (2005) "Action rules discovery: system DEAR2, method and experiments", Journal of Experimental & Theoretical Artificial Intelligence, Vol 17, No. 1–2, pp 119–128. doi:10.1080/09528130512331315855.

Tzacheva, A. A. and Raś, Z. W. (2010) "Association action rules and action paths triggered by meta-actions", Proceedings of the IEEE International Conference on Granular Computing (GrC), pp 772–776. doi:10.1109/GRC.2010.64.

Tzacheva, A. A., Chatterjee, S. and Raś, Z. (2022) "Cloud mining actionable pattern discovery in big data: A survey", Transactions on Machine Learning and Artificial Intelligence, Vol 10, No. 3, pp 42–65. doi:10.14738/tmlai.103.12397.

Vaarma, M. and Li, H. (2024) "Predicting student dropouts with machine learning: An empirical study in Finnish higher education", Technology in Society, Vol 76, Art. 102474. doi:10.1016/j.techsoc.2024.102474.

Wolniak, R. and Grebski, W. (2023) "The concept of diagnostic analytics", Scientific Papers of Silesian University of Technology: Organization and Management Series, No. 175, pp 651–669. doi:10.29119/1641-3466.2023.175.41.