

A Machine Learning-Based Student Dropout Risk Prediction and Early Warning System

Kritanat Chungnoy¹, Tanatorn Tanantong^{2*}, Naphat Jakkraphacharakul¹, Nongnapat Sunontanam¹, Supada Boonyavattanavijit¹, Thanyaluck Tungthananithirath¹, and Surasit Uypatchawong³

¹College of Interdisciplinary Studies, Thammasat University, Pathum Thani, 12121, Thailand

²Thammasat University Research Unit in Data Innovation and Artificial Intelligence, Department of Computer Science, Faculty of Science and Technology, Thammasat University, Pathum Thani, 12121, Thailand

³Department of Computer, Faculty of Science and Technology, Sakon Nakhon Rajabhat University, 4700, Thailand

* Corresponding author. E-mail address: tanatorn@sci.tu.ac.th

Received: 4 March 2026; Revised: 27 June 2026; Accepted: 4 September 2026; Available online: 1 October 2026

Abstract

Student dropout rates in higher education have become a significant problem for universities worldwide, impacting both student academic achievement and institutional stability. To address this problem, this research proposes the Student Dropout Risk Prediction and Early Warning System (SDRP-EWS), a machine learning framework designed to predict the risk of eventual student dropout. Utilizing multi-semester academic performance data from 20,919 undergraduate students at Sakon Nakhon Rajabhat University, Thailand (2014–2024), the study developed logistic regression, random forest, and XGBoost models. To support early and continuous monitoring, these models were evaluated at three distinct prediction horizons: following the first semester, after the second semester, and after three or more semesters. Cohen's d effect size analysis compared ROC-AUC distributions across engineered feature sets, confirming the practical advantage of the selected model configurations. The results show that XGBoost consistently outperformed other models across all prediction stages, achieving F1-scores for the dropout class ranging from 0.65 in the early stages up to 0.93 in later semesters with comprehensive academic histories. This study's primary contribution is integrating the top-performing multi-semester model into a functional, web-based prototype dashboard tailored for Thai higher education. By providing academic advisors with intuitive risk scores and key contributing factors, the system demonstrates strong potential as an evidence-based decision-support tool to facilitate proactive, targeted institutional interventions.

Keywords: Higher Education Analytics, Early warning system, Machine Learning, XGBoost, Dropout Prediction

Introduction

Student attrition in higher education has become a major concern for universities worldwide. It affects not only students' learning experiences and future career opportunities but also the quality, reputation, and financial stability of higher education institutions (Tinto, 2012; Marbouti et al., 2016). In the context of rapid technological and societal transformation, higher education also plays an important role in developing the people skills individuals need to adapt to emerging economic and digital environments (Tanantong et al., 2024; Tantiathimongkhon et al., 2026). Many students struggle to adapt to university-level studies due to increased academic pressure, low motivation, and social or economic challenges, which may ultimately lead to dropout (Tinto, 2012; Markson et al., 2025). Current student support processes, such as academic counseling and periodic progress reviews, are often reactive, time-consuming, and slow, so students may receive support only when they are already at high risk (Marbouti et al., 2016). Although universities possess a wealth of data, including grades, registration records, and learning management system logs, they often underutilize this information to provide systematic student support (Bañeres et al., 2023). In recent years, learning analytics and

machine learning techniques have been increasingly applied to predict student dropout risk and support early warning systems (Marbouti et al., 2016; Bañeres et al., 2023; Córdova-Esparza et al., 2025; Dasi & Kanakala, 2022; Cho et al., 2023).

Previous research has employed classification models such as logistic regression, decision trees, random forests, and gradient boosting methods to estimate dropout risk. Some studies also incorporate behavioral data from online platforms, including login frequency, assignment submission patterns, and forum participation. These studies show that tree-based ensemble methods such as random forests, gradient boosting, and XGBoost effectively capture nonlinear relationships and interaction effects in educational datasets (Alhardi & ALAN, 2024; Bañeres et al., 2023; Tian et al., 2023; Carballo-Mendivil et al., 2025). In addition, predictive performance improves when models incorporate well-designed features, longitudinal academic records, and appropriate strategies for handling imbalanced datasets (Agrusti et al., 2020; Tian et al., 2023; Fernández et al., 2018; He & Ma, 2013; Albugami et al., 2025). However, a critical comparison of these existing approaches reveals a significant pedagogical and technical dichotomy. While models heavily reliant on behavioral logs achieve high localized accuracy within virtual learning environments, they exhibit low portability to traditional institutional frameworks that lack pervasive digital tracking. Conversely, studies utilizing purely administrative records often analyze student data statically, failing to contrast how predictive performance and risk indicators shift dynamically across sequential enrollment terms. This lack of comparative synthesis leaves a critical research gap regarding how multi-semester administrative data can be systematically structured across different algorithmic paradigms to maximize early-intervention utility.

However, existing research still has several limitations. Many prediction systems estimate dropout risk only at the end of a semester, limiting universities' ability to identify at-risk students early enough for timely intervention (Marbouti et al., 2016). Some studies rely heavily on behavioral data from learning management systems, which may not be consistently available in all institutions. In addition, some research emphasizes overall accuracy as the main evaluation metric, even though recall, F1-score, and AUC provide more meaningful insights for imbalanced classification problems. Furthermore, relatively few studies describe the complete workflow from data preparation and feature engineering to implementation in institutional advisory systems, particularly in non-Western contexts such as Thailand (Bañeres et al., 2023; Tian et al., 2023).

To address these limitations, this study develops a machine learning-based early warning system using multi-semester academic performance data from Sakon Nakhon Rajabhat University, Thailand. The study performs detailed data preparation and feature engineering and compares three machine learning algorithms—logistic regression, random forest, and XGBoost—across three prediction stages (Term 1, Term 2, and Term 3 onwards). Model performance was evaluated using recall, F1-score, ROC-AUC, and PR-AUC. In addition, a prototype web-based dashboard was developed to provide descriptive risk scores and support academic advisors in monitoring student risk levels.

To enhance methodological rigor and guide our empirical evaluation, this investigation explicitly addresses the following research questions How do the predictive capabilities (specifically Recall and F1-score) of baseline linear models and tree-based ensemble models evolve when evaluated dynamically across successive multi-semester horizons?

To what extent does custom, architecture-specific feature engineering mitigate the mathematical classification challenges imposed by severe class imbalance within institutional academic records?

This study makes three main contributions. First, it proposes a multi-semester prediction framework (Term 1, Term 2, and Term 3 onwards) to support earlier identification of at-risk students throughout their academic trajectory. Second, it introduces time-based feature generation with effect-size analysis to improve the interpretability of predictive indicators. Third, it presents a functional dashboard tailored to Thai higher education, illustrating a complete workflow from data preparation to practical implementation.

Materials and Methods

Fig. 1 presents the overall framework of the proposed student dropout prediction system. The process starts with academic and student profile data, followed by data preparation and feature engineering. The data were then segmented by academic terms and split into training and testing sets. Multiple machine learning models were trained and evaluated, and statistical comparison was performed to select the final model. The selected model outputs students' dropout risk as a percentage and corresponding risk levels for early warning purposes.




Figure 1 System architecture of the Student Dropout Risk Prediction and Early Warning System (SDRP-EWS)

Data and study context

This study used institutional data from Sakon Nakhon Rajabhat University, a large public university in Thailand that offers undergraduate programs across six faculties: education, science and technology, management science, humanities and social sciences, industrial technology, and agricultural technology. The dataset initially contained approximately 2.1 million records covering all undergraduate students enrolled between academic years 2014 and 2024, resulting in 31,090 unique students, which were extracted from the university's student information and academic records. Sakon Nakhon Rajabhat University officially authorized and provided these data. For each student, the available variables include term-by-term course grades, cumulative GPA, faculty, gender, upper-secondary GPA, and final study status (graduated or dropout).

Student data were derived from university status codes. Graduates were labeled as "Graduated," while cases of expulsion or withdrawal were categorized as "Dropout". We removed current students and partial data to ensure the labels reflect actual outcomes. After filtering out 964,698 blank entries (46% of the total), skipping summer semesters (2.8%), and removing anomalies such as a Semester 1 GPA of 0.00 with no further information, we were left with 20,919 student datasets for analysis (85.25% graduated, 14.75% withdrawn). The resulting dataset represents each student's academic performance from initial enrollment until graduation or dropout, and provides the basis for constructing early-term and late-term prediction scenarios.

Data exploration

Raw data was stored in two main tables: the course-level grade table and the student-level table. The course-level table had one row per student-course-semester, so the same student ID appeared multiple times across different courses and semesters. In contrast, the student-level table had one row per student with demographic and status information, leading to inconsistencies between the two data sources. Initial investigations also revealed several rows in the grade table that were missing all fields, indicating fictitious records, which were subsequently removed. Furthermore, a small number of students had unusually long study periods and irregular semester structures.

In field terms, both semester and academic year were included in a single string. There was also a small summer term entry, representing less than 3% of the total. Statistics on terms per student show that studying more than eight regular semesters is rare and usually associated with irregular pathways. Therefore, we thought special handling was needed during pre-processing.

Data preparation

Data preparation aimed to clean the raw data table, level out the details, and create a suitable semester-level GPA panel for students to use in modeling. First, rows in the grade table that had no data in any columns were removed, resulting in the loss of 964,698 records that contained no data. Then, the original semester strings were split into two fields: the numerical semester index (`term_num`, with values 1–3 for the first, second, and summer semesters) and the academic year (`term_year`). Because summer semester records accounted for only about 2.8% of all rows and represented a different academic style, we excluded records with `term_num` = 3 from the analysis at this stage. Similarly, enrollment records beyond the 8th regular semester were truncated. These subsets contained very few data points compared to the standard academic timeline (Terms 1 to 8); therefore, we treated them as statistical noise that could distort the core longitudinal patterns and reduce the model's generalizability and predictive stability.

Next, we rolled up course-level records into semester GPAs for each student. For each student in each semester, the Grade Point Average (GPA) is derived from the weighted average of course grades. The data is then arranged chronologically using the academic year and semester number. Appropriate labels are assigned to each student, such as Semester 1, Semester 2, etc., to track progress accurately. The data is converted from a long format (multiple rows per student) to a wide format, one row per student, with the GPA column spanning multiple semesters. Dropped semesters past 8 to match normal 4–5 year undergrad timelines and cut out noise from those rare marathon study cases.

Finally, we merged the cleaned semester GPA table with the student-level table using student IDs, resulting in a single large dataset containing demographic information, multi-semester GPA histories, and final status. We filtered out anomalies, such as a first semester GPA of 0.00 with no further data, leaving 20,919 students and 14 key variables ready for model building.

As shown in Table 1, the data remaining after the cleaning and merging steps consists of a total of 14 variables, which are grouped into three categories: Student Basic Information (3 variables: `STUDENT_ID`, `FAC_NAME`, `GENDER_NAME`), Educational Status (1 variable: `STD_STATUS_GROUP_NAME`), and

Table 1 Input Variables and Engineered Features in SDRP-EWS

Feature Name	Description	Note
STUDENT_ID	Student Identification Number	
TERM{n}	Grade Point Average in each semester	The variable <i>n</i> is equal to 1–8
FAC_NAME	Student's Faculty Name	
GENDER_NAME	Student's Gender	
GPA	Cumulative Grade Point Average	
STD_STATUS_GROUP_NAME	Current Educational Status	"Dropout", "Graduated"
OLD_GPA_M6	Cumulative GPA after completing Grade 12	

Academic Performance (10 variables: TERM{n} [for n=1–8 semesters], GPA, and OLD_GPA_M6)

Feature engineering

To boost predictive power and improve interpretability, we built features from multi-semester GPA trends and other academic data (Mustofa et al., 2025; Liu et al., 2025). In addition to raw term-level GPAs, several aggregate features were constructed to capture overall performance and variation, including the average GPA across completed terms, the minimum and maximum GPA achieved, the GPA range, and the standard deviation of term-level GPA. Difference-based features, such as the change in GPA from the first term to the latest term and the improvement or decline relative to upper-secondary GPA, were also computed to describe performance over time.

Further, binary indicators were introduced to represent key academic events and risk thresholds, for example, the number of failed courses and withdrawn or incomplete grades, flags for having at least one or multiple failures, and threshold-based indicators for low cumulative performance (e.g., GPA below 2.5, 2.0, or 1.75) or low first-term GPA. Interactive features were constructed, such as a product of term GPAs and a composite risk score incorporating failed courses, very low GPAs, and declining trends into a single index, to help both linear and tree models to capture complex patterns in the data. In total, the feature generation process generated 38 acquired features in addition to the original variables, and for each prediction stage (first term, second term, subsequent terms), only features observed up to that time were included in the relevant model.

Table 2–6 covers all 38 attributes we developed from raw grade data. These attributes transform baseline scores into clear risk indicators that can detect complex academic patterns. In short, they enable the model to detect the risk of dropping out of school early and guide students who are likely to experience difficulties.

Table 2 Interaction Feature Engineering

Feature Name	Description	Note
avg_gpa_squared	Squared cumulative grade point average (GPA).	
COUNT_F_squared	Squared number of failed courses.	
avg_gpa_x_COUNT_F	Product of cumulative GPA and the number of failed courses.	These features are constructed by squaring or multiplying existing variables to help both linear and tree-based
avg_gpa_x_trend	Product of cumulative GPA and the GPA change trend.	
TERM1_x_TERM2	Product of first-term GPA and second-term GPA.	
TERM2_x_TERM3	Product of second-term GPA and third-term GPA.	

Table 3 Academic Performance Feature engineering

Feature Name	Description	Note
avg_gpa_up_to_now	Cumulative grade point average (GPAX) computed from all available	Summarize the overall level and variability of academic performance based on term-level GPA.
min_gpa_up_to_now	Lowest term GPA achieved among all completed terms.	
max_gpa_up_to_now	Highest term GPA achieved among all completed terms.	
gpa_range	Difference between the highest and lowest term GPA.	
gpa_std	Standard deviation of term-level GPA across completed terms.	
latest_available_gpa	Total number of terms for which GPA information is available.	
num_terms_with_data	GPA in the most recent term with available data.	

Table 4 Trend and Difference Feature engineering

Feature Name	Description	Note
gpa_change_from_start	Change in GPA relative to the first term.	These features capture trends in academic performance relative to the first term or to upper-secondary GPA, and detect patterns of decline or recovery over time.
improvement_from_hs	Improvement in performance compared with upper-secondary GPA.	
declining_trend	Indicator for a declining GPA trend (1 = yes, 0 = no).	
improving_trend	Indicator for an improving GPA trend (1 = yes, 0 = no).	
decline_last_term	Indicator that the GPA in the latest term is lower than in the immediately preceding term (1 = yes, 0 = no).	
consecutive_decline_2	Indicator that GPA has decreased for two consecutive terms (1 = yes, 0 = no).	
long_decline_3terms	Indicator that GPA has decreased for three consecutive terms (1 = yes, 0 = no).	
overall_gpa_stability	Measure of GPA stability.	
has_recovered	Indicator of performance recovery (1 = yes, 0 = no).	

Table 5 Event and risk thresholds Feature engineering

Feature Name	Description	Note
COUNT_F	Cumulative number of courses in which the student has received a grade of	These features represent risk events such as repeated failures, course withdrawals or incomplete grades, and GPA falling below specific thresholds in different terms, as well as grouping students into performance categories.
has_F	Indicator for having received at least one F (1 = yes, 0 = no).	
multiple_F	Indicator for having received at least two Fs (1 = yes, 0 = no).	
excessive_F	Indicator for having received at least three Fs (1 = yes, 0 = no).	
has_WIU	Indicator for having at least one withdrawal or incomplete grade (W/I/U) (1 = yes, 0 = no).	
low_gpa	Indicator that the cumulative GPA (avg_gpa) is below 2.5 (1 = yes, 0 =	
very_low_gpa	Indicator that the cumulative GPA (avg_gpa) is below 2.0 (1 = yes, 0 = no).	
critical_gpa	Indicator that the cumulative GPA (avg_gpa) is below 1.75 (1 = yes, 0 =	
early_warning	Indicator that first-term GPA is below 2.0 (1 = yes, 0 = no).	
term{x}_low	Indicator that GPA in term {x} is below 2.5 (1 = yes, 0 = no).	
term{x}_excellent	Indicator that GPA in term {x} is at least 3.5 (1 = yes, 0 = no).	
performance_category	Categorical grouping of overall performance levels.	

Table 6 Composite risk and context Feature engineering

Feature Name	Description	Note
risk_score	A composite risk score computed from the number of failed courses, very low GPA, and declining GPA trends.	These features aggregate multiple risk factors into a single index and
current_term	The current term being predicted (a feature that informs the model which point in the study timeline is under consideration).	indicate the current term at which the prediction is made.

Model development and evaluation

The predictive task was formulated as a binary classification problem that distinguished between “Graduated” and “Dropout” students. Three algorithms were evaluated: logistic regression as a baseline interpretable linear classifier, Random Forest as a tree-based ensemble model, and XGBoost as a gradient-boosting framework optimized for tabular data (Burkov, 2019; Theobald, 2021). For each algorithm, separate models were trained for three prediction stages, namely after the first term, after the second term, and after three or more terms. At each stage, only the features that would realistically be available at that point in a student’s academic journey were used.

The dataset was randomly split at the student level into 80% training data and 20% testing data to avoid information leakage between terms of the same student. To prevent data leakage and eliminate mathematical bias during training, we omitted unique student identifiers (STUDENT_ID) and the target outcome labels (STD_STATUS_GROUP_NAME) from the predictor input space. Furthermore, the feature deployment framework was strictly segmented into stage-specific and architecture-specific subsets, ensuring models only utilized data realistically available at each decision point:

- **XGBoost Feature Set:** Term 1 incorporates baseline performance attributes (avg_gpa_up_to_now, min_gpa_up_to_now, max_gpa_up_to_now), high-school transition (improvement_from_hs), early risks (has_F, multiple_F, low_gpa, early_warning), and current_term. Term 2 adds longitudinal shifts (gpa_change_from_start, gpa_std_up_to_now, decline_last_term). Term 3 onwards introduces cumulative trajectory metrics (consecutive_decline_2).

- **Logistic Regression Additional Feature Set:** Beyond the core XGBoost attributes, Term 1 includes polynomial transformations (avg_gpa_squared, count_F_squared, avg_gpa_x_count_F). Term 2 incorporates interaction dynamics (declining_trend, avg_gpa_x_trend, term1_x_term2). Term 3 onwards integrates missingness indicators and terminal statuses (num_term_with_data, latest_available_gpa, improving_term, term2_x_term3).

- **Random Forest Additional Feature Set:** Beyond the core XGBoost attributes, Term 1 utilizes broader statistical spans (gpa_range, gpa_std, excessive_F, very_low_gpa, critical_gpa, term_low/excellent, performance_category, risk_score). Term 2 includes simplified behavioral directional flags (declining_trend, improving_trend). Term 3 onwards incorporates comprehensive volatility metrics (num_term_with_data, latest_available_gpa, improving_term, long_decline_3terms, overall_gpa_stability, has_recovered).

To ensure reproducibility, we systematically handled missing data and categorical variables. Rows with widespread missing values or insufficient longitudinal records (such as students with only first-term data and no definitive administrative status) were dropped, and uninformative columns with high sparsity were removed. For

unelapsed future semesters, missing term-level GPAs were imputed using the cohort's global mean for that specific term. Categorical variables were converted into standard numerical representations using One-Hot Encoding. The feature selection procedure relied on explicit architectural compatibility; rather than a unified filter, features were tailored to match each classifier's mathematical strengths and limitations, resulting in different feature counts across the final models.

To ensure a rigorous and mathematically fair comparison, hyperparameter tuning was conducted uniformly across all models using RandomizedSearchCV configured with identical optimization parameters: 30 iterations ($n_iter = 30$), a 3-fold stratified cross-validation strategy ($cv = 3$), and the Area Under the Receiver Operating Characteristic curve as the primary evaluation objective ($scoring = roc_auc$). Class imbalance within the institutional database was uniformly counteracted via algorithmic penalty tuning: the `scale_pos_weight` parameter was initialized for the XGBoost architecture, while equivalent cost-sensitive adjustments via `class_weight` were implemented for both Logistic Regression and Random Forest models. Decision thresholds were then optimized independently on the validation set by evaluating values from 0.10 to 0.85 in steps of 0.05 to maximize the F1-score for the minority class. Because the data are imbalanced, evaluation focused on recall, F1-score, ROC-AUC, and PR-AUC as key metrics to assess how effectively each model identified at-risk students while controlling false positives at each prediction stage. The model configuration that provided the best balance across these metrics was selected for further analysis and later integration into the early-warning system prototype.

Table 7, the number of features used in each testing model, shows a clear increasing trend across the Terms. Random Forest consistently used the most features at every stage (26, 33, and 64 features), followed by Logistic Regression (20, 28, and 45 features), while XGBoost used the fewest features overall (16, 21, and 34 features) in sequential order.

Table 7 Number of Features Used in Each Testing Model

Model	Term_1 (features)	Term_2 (features)	Term_3+ (features)
Random Forest	26	33	64
Logistic Regression	20	28	45
XGboost	16	21	34
Baseline	6	7	13

Results

This experiment presents the results of three models: logistic regression, random forest, and XGBoost, used to predict the risk of student dropout in three periods: after the first semester, after the second semester, and after three semesters or more. The models were evaluated using Recall, F1-score, and ROC-AUC on a separate test dataset. The best-performing model then underwent further statistical analysis and was integrated into an early warning system.

Descriptive results of the dataset

The final modeling dataset consists of 20,919 students from the academic years 2014 to 2024. After data collection, cleaning, and anomaly removal, students were divided into two groups: Graduated and Dropout, with a 1:4.75 ratio. This imbalance led this research to use Recall and F1-Score as the main indicators to evaluate the model's performance.

Model performance across terms

Following the experimental setup described above, Tables 8–11 show the performance of Logistic Regression, Random Forest, and XGBoost in three prediction periods, evaluated from three feature sets designed for XGBoost, Logistic Regression, and Random Forest, respectively. The experiments showed that XGBoost performed better than Logistic Regression and Random Forest. These results show that the Gradient Boosting Tree architecture can better capture complex relationships and repeated failure patterns in academic features than linear and bagging models. Notably, even when XGBoost is applied to feature sets designed for other algorithms, the model still performs better. Furthermore, models trained on three or more terms and on all three feature sets yielded XGBoost F1 scores and ROC-AUC values close to 1.0, demonstrating its high accuracy in identifying students at risk of dropping out.

Table 8 The results of the experiments using the baseline feature set

Model	Term	Features	Recall	F1	ROC-AUC	PR_AUC
Logistic Regression	1	6	0.661	0.611	0.841	0.664
	2	7	0.674	0.606	0.849	0.686
	3	13	0.836	0.757	0.933	0.852
Random Forest	1	6	0.660	0.615	0.859	0.695
	2	7	0.653	0.623	0.867	0.721
	3	13	0.822	0.794	0.947	0.888
XGboost	1	6	0.708	0.593	0.864	0.708
	2	7	0.713	0.612	0.870	0.729
	3	13	0.856	0.761	0.947	0.891

Table 9 The results of the experiments using the feature set specifically engineered for XGboost

Model	Term	Features	Recall	F1	ROC-AUC	PR_AUC
Logistic Regression	1	16	0.584	0.635	0.843	0.663
	2	21	0.606	0.637	0.852	0.691
	3	34	0.830	0.883	0.985	0.952
Random Forest	1	16	0.592	0.621	0.858	0.694
	2	21	0.555	0.636	0.858	0.713
	3	34	0.883	0.911	0.990	0.970
XGboost	1	16	0.621	0.640	0.866	0.709
	2	21	0.601	0.657	0.870	0.732
	3	34	0.891	0.926	0.992	0.973

Table 10 The results of the experiments using the feature set specifically engineered for logistic regression

Model	Term	Features	Recall	F1	ROC-AUC	PR_AUC
Logistic Regression	1	20	0.603	0.633	0.844	0.669
	2	28	0.598	0.641	0.848	0.706
	3	45	0.865	0.877	0.985	0.953
Random Forest	1	20	0.529	0.617	0.858	0.693
	2	28	0.603	0.635	0.858	0.715
	3	45	0.892	0.923	0.991	0.967
XGboost	1	20	0.614	0.642	0.866	0.710
	2	28	0.588	0.652	0.870	0.732
	3	45	0.892	0.923	0.991	0.970

Table 11 The results of the experiments using the feature set specifically engineered for Random Forest

Model	Term	Features	Recall	F1	ROC-AUC	PR_AUC
Logistic Regression	1	26	0.581	0.633	0.842	0.653
	2	33	0.610	0.639	0.851	0.693
	3	64	0.822	0.875	0.984	0.951
Random Forest	1	26	0.597	0.621	0.858	0.694
	2	33	0.647	0.636	0.859	0.713
	3	64	0.869	0.915	0.989	0.969
XGboost	1	26	0.621	0.633	0.842	0.708
	2	33	0.647	0.657	0.870	0.732
	3	64	0.881	0.923	0.991	0.974

To provide a comprehensive view of the early warning system's practical utility, Table 12 presents a single consolidated summary of the final selected model configuration at each prediction stage. This comprehensive view integrates the optimized decision thresholds, exact confusion matrix counts, and detailed per-class classification metrics to facilitate direct comparison across the academic timeline.

Table 12 per-class precision/recall/F1-score metric for best model

Term	Selected Model Configuration	Decision Threshold	Target Class	Predicted Graduate	Predicted Dropout	Precision	Recall	F1-Score	Support
1	Random Forest (RF-engineered)	0.70	Graduate	3,250	207	0.92	0.94	0.93	3,457
			Dropout	269	458	0.68	0.63	0.65	727
2	XGBoost (XG-engineered)	0.65	Graduate	3,215	242	0.93	0.93	0.93	3,457
			Dropout	262	465	0.67	0.64	0.66	727
3+	Logistic Regression (LR-engineered)	0.60	Graduate	3,422	35	0.98	0.99	0.99	3,457
			Dropout	58	669	0.94	0.92	0.93	727

Decision threshold values for risk identification were determined using an estimated range from 0.1 to 0.85, in increments of 0.05, based on the validation dataset. We selected the value with the highest F1-score to effectively identify at-risk students. This approach balances detecting actual student dropouts with reducing false positives. With a dropout rate of approximately 14.75%, the obtained threshold values were: Term 1: 0.70 (high to limit false positives in the early stages of high uncertainty), Term 2: 0.65, and Term 3: 0.60 (lower to utilize more complete data to

Effect size comparison across feature sets

To optimize XGBoost performance across prediction periods and measure differences relative to other designs, we performed an effect size analysis. This involved pairwise comparisons of ROC-AUC values between different models at each term: after term 1, after term 2, and after three or more terms. Four XGBoost models were considered: the baseline model using only basic, unengineered features, and three other models trained with features specifically designed for Logistic Regression (LR Engineering), Random Forest (RF Engineering), and XGBoost (XG Engineering), respectively. Repeated Stratified Cross-Validation was used to obtain the distribution of ROC-AUC values for each term of the four models. Cohen's d was then calculated for all models.

Six and eighteen pairwise comparisons were performed overall. Cohen's *d* measures the difference in the mean ROC-AUC values between models using the standard deviation. A higher absolute value indicates a larger practical difference, while a value close to zero indicates a statistically significant difference with limited practical benefit. Each pairwise comparison was then performed. The analysis reported the mean ROC-AUC for each model, the corresponding Cohen's *d*, the qualitative effect size (insignificant, small, medium, or large), and the *p*-value from the paired *t*-test to identify the "winner" model with the higher mean ROC-AUC. This analysis allowed the research to identify which XGBoost setting yielded optimal performance on a term basis and to determine whether the observed differences relative to the baseline and other feature sets were practically significant. In this research, Cohen's *d* was calculated from the distribution of ROC-AUC values obtained from Repeated Stratified Cross-Validation, enabling measurement of practical performance differences between XGBoost settings in standardized units.

Tables 13–15 present the pairwise effect size analysis conducted to identify the most suitable XGBoost configuration at each prediction stage. In term 1, the XGBoost using RF engineering features exhibited the highest average ROC-AUC and outperformed the baseline and other features. Although Cohen's *d* was insignificant to slightly positive, it consistently showed a small but sustained advantage. In term 2, XGBoost using XG engineering features showed the highest performance, with most effect sizes positive compared to the baseline and RF and LR features, indicating a moderate and stable improvement. From term 3 onward, XGBoost with LR engineering features showed the highest average ROC-AUC, with a slightly positive but consistent Cohen's *d* compared with

other configurations, supporting its selection when sufficient cumulative learning data is available.

Table 13 XGboost pairwise effect size comparisons term1

Comparison Model		Performance Mean		Effect Size		
Model 1	Model 2	Mean Model 1	Mean Model 2	Cohen's <i>d</i>	Magnitude	<i>p</i> -value
Baseline	LR engineering	0.8690	0.8698	-0.0924	Negligible	1.3×10^{-4}
Baseline	RF engineering	0.8690	0.8698	-0.0921	Negligible	7.2×10^{-5}
Baseline	XG engineering	0.8690	0.8697	-0.0762	Negligible	9.9×10^{-5}
RF engineering	XG engineering	0.8698	0.8697	+0.0161	Negligible	0.1683
XG engineering	LR engineering	0.8697	0.8698	-0.0158	Negligible	0.1529
RF engineering	LR engineering	0.8698	0.8698	+0.0005	Negligible	0.9689

Table 14 XGboost pairwise effect size comparisons term2

Comparison Model		Performance Mean		Effect Size		
Model 1	Model 2	Mean Model 1	Mean Model 2	Cohen's <i>d</i>	Magnitude	<i>p</i> -value
Baseline	LR engineering	0.8750	0.8745	+0.0535	Negligible	0.0615
Baseline	RF engineering	0.8750	0.8748	+0.0221	Negligible	0.2885
Baseline	XG engineering	0.8750	0.8750	-0.0012	Negligible	0.9538
RF engineering	XG engineering	0.8748	0.8750	-0.0225	Negligible	0.0646
XG engineering	LR engineering	0.8750	0.8745	+0.0553	Negligible	0.0120
RF engineering	LR engineering	0.8748	0.8745	+0.0327	Negligible	0.0964

Table 15 XGboost pairwise effect size comparisons term3+

Comparison Model		Performance Mean		Effect Size		
Model 1	Model 2	Mean Model 1	Mean Model 2	Cohen's d	Magnitude	p-value
Baseline	LR engineering	0.9908	0.9913	-0.3959	Small	1.4×10^{-9}
Baseline	RF engineering	0.9908	0.9911	-0.2463	Small	3.4×10^{-7}
Baseline	XG engineering	0.9908	0.9912	-0.2848	Small	4.2×10^{-5}
RF engineering	XG engineering	0.9911	0.9912	-0.0357	Negligible	0.0122
XG engineering	LR engineering	0.9912	0.9913	-0.1064	Negligible	0.0270
RF engineering	LR engineering	0.9911	0.9913	-0.1416	Negligible	0.2541

This result supports the selection of such a model as the primary choice when more complete cumulative academic performance data is available. Considering all prediction intervals as a whole, most Cohen's d values were insignificant to slightly insignificant, indicating that statistically significant differences in ROC-AUC values in many cases offer limited practical benefit. However, the consistent direction of positive effect size across all comparisons reinforces the evidence for identifying a model configuration that provides statistically reliable and practically meaningful improvements for academic counseling.

Feature importance and risk patterns

Feature importance for the XGBoost models was evaluated after hyperparameter tuning with RandomizedSearchCV, using the best-performing configuration at each prediction stage compatible with the Scikit-learn API. Feature importance values were extracted via the model's standard attribute, using gain as the importance metric, which reflects the average improvement in the objective function from splits involving each feature. The resulting scores were combined with feature names to construct a ranked table for the first-term model (Term1), the second-term model (Term2), and the three-or-more-terms model (Term 3+), sorted from the most to the least important variables and used to interpret the role of individual predictors as well as to check for potential data-leakage risks among highly influential features.

Table 16 shows that combining observed academic performance across multiple semesters with behavioral risk indicators such as the frequency of F grades and cumulative GPA improves the model's ability to consistently identify students at risk of dropping out. In the early semester model, failing-related variables including multiple_F, COUNT_F, excessive_F, and has_F, along with performance categories, faculty coding, and cumulative GPA (avg_gpa_up_to_now and min_gpa_up_to_now), were the most prominent predictors. In the third-semester model and beyond, low-semester GPA indicators (term4_low and term3_low), missing-GPA indicators (e.g., TERM5_missing), and combinations such as risk_score and repeated-failing variables (multiple_F and COUNT_F) became the most influential features. These results suggest that both semester academic status and cumulative GPA patterns play crucial roles in detecting dropout risk. Students with repeated failures or a GPA in the critical range show persistent academic difficulty, while declining GPA patterns may signal increased risk. Although the current GPA remains within the institution's requirements, monitoring dropout risk should not rely solely on a single point-in-time GPA cutoff; it should also incorporate cumulative academic performance over several semesters to support timely, proactive academic intervention.

Table 16 Feature Importance Score Rank

RANK	ImportantScore Term1	Score	ImportantScore Term2	Score	ImportantScore Term3+	Score
1	multiple_F	0.674	multiple_F	0.443	term4_low	0.388
2	COUNT_F	0.123	COUNT_F	0.235	TERM7	0.208
3	has_F	0.067	excessive_F	0.068	TERM8	0.091
4	excessive_F	0.056	has_F	0.060	TERM4	0.033
5	performance_category	0.016	min_gpa_up_to_now	0.022	excessive_F	0.024
6	FAC_ENCODED	0.012	avg_gpa_up_to_now	0.018	multiple_F	0.021
7	avg_gpa_up_to_now	0.008	FAC_ENCODED	0.017	risk_score	0.020
8	COUNT_WIU	0.008	TERM2	0.017	term3_low	0.018
9	TERM1	0.008	COUNT_WIU	0.012	TERM5_missing	0.017
10	OLD_GPA_M6	0.007	OLD_GPA_M6	0.011	COUNT_F	0.015

In addition, the effect size analysis further supports these findings by quantifying the magnitude of differences between dropout and graduated students. Features with medium-to-large Cohen's *d* values, such as cumulative GPA, repeated course failures, and composite risk indicators, also rank among the most influential predictors identified by the XGBoost model based on gain. This alignment indicates that variables showing strong group separation in statistical analysis are effectively used by the machine learning models, reinforcing both the validity of the feature engineering process and the interpretability of the proposed dropout prediction framework.

These top features translate directly to actionable insights for academic advisors. For example, Term 1 (multipleF, COUNTF): Flag students with ≥ 2 F grades for immediate tutoring/remediation to prevent early attrition. Term 2 (multipleF, excessiveF): Monitor consecutive failures; schedule counseling if ≥ 3 Fs signal deepening academic distress. Term 3+ (term4low, riskscore): Intervene on low mid-term GPA (< 2.5) or high composite risk via probation plans or faculty advising.

Prototype early warning system

The selected models were deployed to enable fully automated dropout-risk prediction in practice. The best-performing XGBoost model, chosen from all three prediction stages (after the first term, after the second term, and after three or more terms) and systematically evaluated feature sets, was adopted as the core prediction engine. A lightweight Web API was then implemented using the FastAPI framework to handle data input and return prediction results in real time. (Sushmitha & Shivamurthy, 2023) The API receives user inputs such as faculty, gender, the number of F grades, and term-wise GPA values, preprocesses them into the required format (including categorical encoding), and forwards the processed features to the XGBoost model to estimate the probability of dropout.

To ensure consistent operation across environments, Docker containerized the FastAPI Web API, the persisted XGBoost model, and all required dependencies. This containerization enables reliable deployment across different servers and allows the prediction service to integrate seamlessly into a web application for academic staff. As a result, advisors can obtain timely risk estimates and use the model outputs to support data-driven decision-making and personalized counseling for students

Fig. 2 shows that the student dropout risk assessment system consists of three main components: (1) Student Profile Input, which collects general information and academic performance indicators such as cumulative GPA, failed courses, and GPA history for each semester; (2) Risk Assessment Output, which displays the predicted dropout probability, the risk level, and key indicators that explain the assessment results; and (3) Future Risk

Projection, which allows users to simulate the impact of the projected GPA in the next semester on the dropout probability, to support student assistance planning.

1 Student Profile

GENERAL INFORMATION

Faculty: Faculty of Science and Technology

Gender: Male GPAX (Cumulative): 2.55

ACADEMIC PERFORMANCE

Number of F Grades: 2

Term GPA History (Enter available data)

Year	Term 1	Term 2
Year 1	2.5	2.6
Year 2		
Year 3		
Year 4		
Year 5		

Calculate Risk >

2 ANALYSIS REPORT Risk Assessment Date: Dec 13, 2025 ID: 094248

DROPOUT PROBABILITY

63.8%

High Risk Level

Current GPAX: 2.55 Failed Subjects (F): 2 Data Coverage: 2 Terms

Strategic Recommendation

Critical Academic Deficit: Multiple course failures detected. Immediate intervention by an academic advisor is required to restructure the study plan.

3 Future Projection Simulation

Analyze impact of next term's performance.

Expected Next Term GPA: 2.5

Simulate

Projection Results

Current Risk: 63.8%
Future Risk: 12.5%
Change: + Decreased by 51.4%

RECOMMENDATION

Achieving a GPA of 2.50 in the upcoming term is projected to reduce the dropout probability by 51.4%.

Figure 2 Dropout Risk Prediction Process in SDRP-EWS

Fig.3 shows a webpage designed for academic advisors, allowing them to download standard student information forms for their supervised students. These forms record relevant student data, which can then be uploaded to the system for automatic processing and dropout-risk analysis. This helps advisors effectively identify at-risk students and plan appropriate academic support.

Student ID	Name	Prediction	Probability	Level	Action
STD00069	Harper Lopez	Dropout	74.8%	High	View
STD00089	Joshua Harris	Dropout	67.3%	High	View
STD00063	Samuel Wright	Dropout	64.7%	High	View
STD00019	Joshua Harris	Dropout	62.7%	High	View
STD00087	Christopher Lewis	Dropout	59.4%	Medium	View
STD00054	Emma Wilson	Dropout	58.5%	Medium	View
STD00006	Matthew Thomas	Dropout	58.3%	Medium	View
STD00099	Christopher Lewis	Dropout	54.9%	Medium	View
STD00064	Andrew Martinez	Dropout	54.3%	Medium	View
STD00081	David Walker	Dropout	53.7%	Medium	View

Prev Page 1 / 10 Next

Figure 3 Web-Based Advisor Dashboard of SDRP-EWS

Fig. 4 shows the screens for counseling professors, which display the predicted risk of absenteeism for each student in care, including key information such as student ID, first name, last name, and risk probability values in both percentages and risk levels (high, medium, and low). In addition, the dashboard includes a feature to

view individual risk details, allowing advisory faculty to examine the factors affecting each student's future learning status in detail.

Import Student Data Download CSV Template

Upload a CSV/XLSX file to analyze student dropout risks in batch.

Choose File no file selected

Analyze

Student ID	Name	Prediction	Probability	Level	Action
------------	------	------------	-------------	-------	--------

Figure 4 webpage displaying prediction results for academic advisors

The Fig. 5 illustrates the detailed information associated with the specific factors contributing to dropout risk for individual students. By presenting the underlying components that inform each prediction, the system enables academic advisors to assess students' academic situations accurately. This detailed view supports the development of targeted, evidence-based intervention strategies that address the key factors underlying potential student attrition.

Student Analysis Report ×

Student: Harper Lopez (STD00069)
Faculty: N/A

Dropout Probability
74.8%
(High)

Key Contributing Factors:

Total F Grades: 5 Subjects

Declining GPA Trend: Significantly Declining GPA Trend

Figure 5 Additional details about individual student risk

Discussion

The findings indicate that incorporating longitudinal academic performance data improves the accuracy of predicting student dropout risk. Across the three prediction phases, the XGBoost model consistently outperformed logistic regression and random forest in terms of recall, F1-score, ROC-AUC, and PR-AUC. However, it is important to contextualize this superiority; our Cohen's d effect-size analysis revealed that the practical differences across various XGBoost feature-set configurations were mostly negligible to small. This suggests that while gradient boosting provides statistical improvements over baseline models, the practical magnitude of these differences when altering feature sets is relatively modest. These findings suggest that tree-based boosting methods are particularly effective for capturing non-linear relationships within educational datasets, especially when modeling patterns related to academic performance and repeated course attempts.

Compared with studies that conduct prediction only at the end of a semester (Marbouti et al., 2016), the staged prediction approach enables earlier and more accurate identification of at-risk students. Notably, recall

increased from 62% in the first semester to 92% by the third semester, indicating that prediction accuracy improves as additional academic history becomes available. It must be explicitly acknowledged that the near-perfect performance observed in the later prediction stage (Term 3+) reflects the availability of a much more complete academic history. Therefore, predictions at this stage function more as later-stage detection rather than true "early warning" compared to the Term 1 and Term 2 models. Furthermore, while this prototype advisory tool shows potential for early identification, practical claims must be moderated; this study has not yet provided empirical intervention evidence that institutional actions based on this tool effectively reduce dropout rates in practice.

In contrast to the approach presented by Bañeres et al. (2023), which relies on behavioral data from learning management systems, this study's model uses only academic performance data already available within university administrative systems. This characteristic makes the framework potentially practical for public universities in Thailand, where financial and technological resources may be limited. The analysis also reveals patterns in the academic trajectories of students who eventually drop out. Key indicators include repeated course failures, consistently low-grade point averages, declining academic performance across semesters, and elevated risk scores generated by the model. These findings suggest that student dropout is often linked to cumulative academic difficulties rather than isolated setbacks.

Several limitations should be acknowledged. The model was developed using data from Sakon Nakhon Rajabhat University, which may limit the generalizability of the findings to other higher education institutions with different academic structures, grading policies, or definitions of student dropout. Changes in educational policy between 2014 and 2024 may also influence the characteristics of more recent student cohorts. Methodologically, using a random 80/20 train-test split over this 10-year period is a limitation; a temporal validation design (e.g., training on past cohorts and testing on future cohorts) would likely provide a more realistic estimate of future deployment performance. In addition, this study did not evaluate demographic differences or incorporate contextual factors such as financial conditions, mental health, or family background, which may also affect dropout risk. Finally, researchers must rigorously address ethical implications before full-scale deployment in a real-world setting, particularly regarding student privacy, informed consent for data use, and the potential for algorithmic bias. Future research may extend this work by integrating additional data sources such as learning management system activity, attendance records, and student engagement indicators. Another direction would be to link the prediction model with institutional intervention strategies, including academic advising, tutoring programs, or financial support, and evaluate whether these interventions reduce dropout risk over time through causal validation. Cross-institutional validation and comparison with more complex models, such as neural networks, would further strengthen the approach's generalizability and practical impact.

Conclusion and Suggestions

This study developed a machine learning-based framework for predicting student dropout risk using longitudinal academic records from a higher education institution. The Student Dropout Risk Prediction and Early Warning System (SDRP-EWS) demonstrates a potential for staged identification of students at risk of discontinuing their studies by analyzing academic performance patterns across multiple semesters. Among the evaluated models, XGBoost achieved the strongest predictive performance across all prediction stages. The

findings highlight the importance of academic indicators such as cumulative GPA, course failures, and repeated course attempts in identifying dropout risk. The staged prediction approach also has practical value: early-term predictions provide initial warning signals, while later predictions become more accurate as additional academic data become available. However, the exceptionally high performance in the later prediction stage reflects the availability of a more complete academic history; therefore, these late-stage metrics should not be interpreted the same way as very early-stage predictions. The prototype early warning dashboard shows how predictive models can help academic advisors identify at-risk students and plan targeted institutional support. Although this study is based on data from a single institution, the proposed framework provides a practical foundation for data-driven student retention strategies. Future research may extend beyond predictive modeling to conduct causal validation and empirical intervention evaluations to systematically measure actual reductions in student attrition. Additionally, future work should incorporate fairness analysis to detect and mitigate potential algorithmic bias, integrate behavioral and socio-economic factors, and validate the model across multiple universities to enhance its generalizability and real-world applicability.

Acknowledgements

This work was supported by Thammasat University Research Unit in Data Innovation and Artificial Intelligence.

Author Contributions

Kritanat Chungnoy: Conceptualization, Methodology, Modeling and analysis, Writing original draft preparation, Writing –review and editing

Tanatorn Tanantong: Conceptualization, Methodology, Data curation, Data analysis, Validation, Writing – review and editing, Corresponding author

Naphat Jakkracharakul: Data curation, Data Preprocessing, Feature engineering, Statistical analysis, Data interpretation, Technical validation, Writing –review and editing.

Nongnapat Sunontanam: Conceptualization, Methodology, Data acquisition, Data validation, Writing – review and editing

Supada Boonyavattanavijit: System implementation, Model integration, Performance evaluation, Visualization, Writing – original draft.

Thanyaluck Tungthanathirath: Conceptualization, Machine learning model development, Software development, Writing –review and editing

Surasit Uypatchawong: Conceptualization, Data curation, Data analysis, Writing –review and editing

Conflict of Interests

All authors declare that they have no conflicts of interest.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Declaration of Generative AI and AI-assisted Technologies

During the preparation of this manuscript, the authors used ChatGPT and Claude for code writing in model training and web interface design, along with basic grammar checks throughout. After using these tools, the authors carefully reviewed and edited everything and take full responsibility for the final published version.

References

- Agrusti, F., Mezzini, M., & Bonavolontà, G. (2020). Deep learning approach for predicting university dropout: A case study at Roma Tre University. *Journal of E-Learning and Knowledge Society*, 16(1), 44–54. <https://doi.org/10.20368/1971-8829/1135192>
- Albugami, S., Albugami, H., & Wali, A. M. (2024). From Data to Decision: Machine Learning and Explainable AI in Student Dropout Prediction. *Journal of eLearning and Higher Education*, 2024, Article 246301, <https://doi.org/10.5171/2024.246301>
- Alhardi, A., & ALAN, S. (2024). Predicting student dropout in higher education using machine learning techniques: a predictive model using XGBoost algorithm. *International Conference on Engineering Technologies (ICENTE'24)* (p. 31). Turkey, Konya.
- Bañeres, D., Rodríguez-González, M. E., Guerrero-Roldán, A. E., & Cortadas, P. (2023). An early warning system to identify and intervene online dropout learners. *Computers & Education*, 202, 104840. <https://doi.org/10.1016/j.compedu.2023.104840>
- Burkov, A. (2019). *The hundred-page machine learning book*. Andriy Burkov.
- Carballo-Mendivil, B., Arellano-González, A., Ríos-Vázquez, N. J., & Lizardi-Duarte, M. D. P. (2025). Predicting student dropout from day one: XGBoost-based early warning system using pre-enrollment data. *Applied Sciences*, 15(16), 9202. <https://doi.org/10.3390/app15169202>
- Cho, C. H., Yu, Y. W., & Kim, H. G. (2023). A study on dropout prediction for university students using machine learning. *Applied Sciences*, 13(21), 12004. <https://doi.org/10.3390/app132112004>
- Córdova-Esparza, D. M., Terven, J., Romero-González, J. A., & Córdova, K. E. (2025). Predicting and preventing school dropout with business intelligence: Insights from a systematic review. *Information*, 16(4), 326. <https://doi.org/10.3390/info16040326>
- Dasi, H., & Kanakala, S. (2022). Student dropout prediction using machine learning techniques. *International Journal on Informatics Visualization*, 6(3), 678–683. <https://doi.org/10.30630/joiv.6.3.2276>
- Fernández, A., García, S., Galar, M., Prati, R. C., Krawczyk, B., & Herrera, F. (2018). *Learning from imbalanced data sets*. Springer. <https://doi.org/10.1007/978-3-319-98074-4>
- He, H., & Ma, Y. (2013). *Imbalanced learning: Foundations, algorithms, and applications*. Wiley-IEEE Press. <https://doi.org/10.1002/9781118646106>

- Liu, Z., Zhou, X., & Liu, Y. (2025). Student dropout prediction using ensemble learning with SHAP-based explainable AI analysis. *Journal of Social Systems and Policy Analysis*, 2(3), 111–132. <https://doi.org/10.62762/JSSPA.2025.321501>
- Marbouti, F., Diefes-Dux, H. A., & Madhavan, K. (2016). Models for early prediction of at-risk students in a course using standards-based grading. *Computers & Education*, 103, 1–15. <https://doi.org/10.1016/j.compedu.2016.09.006>
- Markson, R., Thiago, R., Tiago, R., Rafael, T., Vinicius, R., Emanuel, M., Roberto, M., & Cristian, C. (2025). Student dropout prediction through machine learning optimization: Insights from Moodle log data. *Scientific Reports*, 15, 93918. <https://doi.org/10.1038/s41598-025-93918-1>
- Mustofa, S., Emon, Y. R., Mamun, S. B., Akhy, S. A., & Ahad, M. T. (2025). A novel AI-driven model for student dropout risk analysis with explainable AI insights. *Computers and Education: Artificial Intelligence*, 8(1), 100352. <https://doi.org/10.1016/j.caeai.2024.100352>
- Sushmitha, K., & Shivamurthy, R. C. (2023). A real-time web application using machine learning for predicting university dropouts. *International Journal of Scientific Development and Research*, 8(8), 63–67.
- Tanantong, T., Moolngearn, P., Kraiwanit, T., Limna, P., & Rafiyya, A. (2024) People Skills in the 21st Century: A Perspective on the Smart City in an Emerging Economy. *Human Behavior and Emerging Technologies*, 2024, 5211958. <https://doi.org/10.1155/2024/5211958>
- Tantiathimongkhon, T., Ketcham, M., & Rattanasiriwongwut, M. (2026). The strategic-behavioral analysis: Factors influencing AI-Powered learning platform adoption in Thai higher education. *Journal of Information Technology Education: Research*, 25, 12. <https://doi.org/10.28945/5694>
- Theobald, O. (2021). *Machine learning for absolute beginners: A plain English introduction* (3rd ed.). Scatterplot Press.
- Tian, X., Nunes, B. P., Liu, Y., & Manrique, R. (2023). Predicting student engagement using sequential ensemble model. *IEEE Transactions on Learning Technologies*, 17, 939–950. <https://doi.org/10.1109/TLT.2023.334286>
- Tinto, V. (2012). *Completing College: Rethinking Institutional Action*. University of Chicago Press. <https://doi.org/10.7208/chicago/9780226804545.001.0001>