Available Online at http://www.journalajst.com

ISSN: 0976-3376

ASIAN JOURNAL OF SCIENCE AND TECHNOLOGY

Asian Journal of Science and Technology
Vol. 17, Issue, 02, pp. 14171-14177, February, 2026

RESEARCH ARTICLE

ADAPTIVE MULTIMODAL IMPUTATION AND NORMALIZATION (AMIN): A PRACTICAL PREPROCESSING FRAMEWORK FOR PREDICTING STUDENT ACADEMIC PERFORMANCE USING SMARTPHONE BEHAVIORAL DATA

***[1]Vimala, S. and [2]Dr. G. Arockia Sahaya Sheela**

[1]PhD Scholar (Full Time), Department of Computer Science, St.Joseph's College (Autonomous), Tiruchirappalli – 2, Affiliated to Bharathidasan University, Tamil Nadu, India

[2]Assistant Professor, Department of Computer Science, St.Joseph's College (Autonomous), Tiruchirappalli -2, Affiliated to Bharathidasan University, Tamil Nadu, India

ARTICLE INFO

Article History:
Received 19th November, 2025
Received in revised form
27th December, 2025
Accepted 08th January, 2026
Published online 27th February, 2026

Key words:
Smartphone behavioral data, Academic performance prediction, Deep learning preprocessing, Missing value imputation, Adaptive normalization, Student analytics.

ABSTRACT

Predicting student academic performance from smartphone usage patterns requires careful preprocessing of heterogeneous mobile sensor data before deep learning model training. This research introduces AMIN (Adaptive Multimodal Imputation and Normalization), a systematic preprocessing framework designed to standardize noisy, incomplete smartphone behavioral data for educational prediction tasks. Through empirical evaluation across multiple deep learning architectures, we demonstrate that strategic preprocessing choices substantially impact model performance—often exceeding the improvements gained from architectural modifications alone. The AMIN framework integrates temporal-aware imputation for missing values, modality-specific normalization tailored to different sensor types, and per-student baseline adjustment to prevent identity-based shortcuts in learning. Comparative analysis shows AMIN achieving performance improvements of 8-14% over conventional preprocessing approaches across MLP, LSTM, and Bi-LSTM architectures. This work establishes a reproducible baseline preprocessing methodology that enhances model stability, enables fair architectural comparisons, and facilitates adoption of standardized practices in educational data science research.

Citation: ***Vimala, S. and Dr. G. Arockia Sahaya Sheela.* 2026.** "Adaptive Multimodal imputation and Normalization (Amin): A Practical Preprocessing Framework for predicting student academic Performance using Smartphone Behavioral Data", *Asian Journal of Science and Technology*, 17, (02), 14171-14177.

INTRODUCTION

The challenge of identifying students at academic risk remains a fundamental concern in educational institutions worldwide. Traditional approaches rely on institutional records—prior grades, attendance records, participation metrics—which provide limited insight into dynamic behavioral patterns. Smartphone usage data offers a new dimension of behavioral information that was previously inaccessible: how students spend their screen time, which applications they engage with most frequently, their temporal patterns of device interaction, and how these behaviors correlate with academic outcomes [1]. Recent research demonstrates meaningful associations between smartphone usage patterns and academic performance. Students who maintain disciplined phone habits, allocating more time to educational applications and less to entertainment, tend to achieve higher grades. Conversely, excessive engagement with non-academic applications correlates with performance decline. However, this rich behavioral data arrives with substantial practical challenges: phone sensor streams contain missing values from network interruptions and user privacy settings, exhibit heterogeneous distributions across different sensor types, and vary dramatically across individual students due to personal usage habits [2]. Deep learning models, particularly recurrent neural networks, have shown promise in capturing temporal patterns within sequential smartphone data. However, these architectures demonstrate well-known sensitivity to input distribution characteristics and missing data patterns. A neural network receiving poorly prepared data risks learning spurious associations—such as treating all Tuesday patterns identically across students or learning to recognize individual students by baseline behavior rather than discovering meaningful behavioral relationships [3]. The research community currently faces a methodological problem: published investigations frequently emphasize novel neural architectures while providing minimal detail about preprocessing procedures. This creates a fragmented landscape where different studies employ different imputation strategies, normalization approaches, and feature engineering choices, making meaningful comparisons between architectures nearly impossible. When one study uses mean imputation and another uses forward-fill, performance differences may reflect preprocessing choices rather than architectural innovations [4]. This investigation addresses this fundamental gap by establishing AMIN, a transparent, modular preprocessing framework specifically designed for mobile behavioral data in educational contexts. Rather than claiming that AMIN represents a novel architectural breakthrough, we emphasize what it genuinely offers: a systematic, reproducible approach to data preparation that demonstrates consistent performance benefits across multiple deep learning models while remaining simple enough for practical implementation [5]. Our specific contributions include: (1) establishing a comprehensive review of preprocessing challenges specific to educational smartphone data, (2) designing a principled preprocessing framework that addresses temporal dependencies, sensor heterogeneity, and individual variation, (3) providing empirical evidence demonstrating AMIN's performance advantages across multiple architectures, and (4) creating a practical implementation roadmap for researchers and practitioners adopting standardized preprocessing practices.

LITERATURE REVIEW

Understanding Smartphone Usage Patterns in Academic Contexts: Contemporary educational research increasingly recognizes that student device interactions reveal meaningful behavioral patterns. Studies examining thousands of students across multiple institutions document clear patterns: students dedicating concentrated study periods to educational applications while limiting recreational phone use demonstrate measurably better academic outcomes [6]. The magnitude of this relationship varies with specific behaviors—excessive social media consumption during study hours shows stronger negative correlation than similar consumption during designated breaks. The mechanisms underlying these associations reflect fundamental learning principles. Device-based distraction during study sessions fragments attention and reduces information consolidation. Conversely, intentional use of educational applications—accessing course materials, engaging with interactive learning tools, communicating with study partners—directly supports academic work. Smartphone data captures these behavioral patterns at a granular temporal level unavailable through traditional academic records [7]. Practical implementation of such prediction systems requires overcoming several data challenges. Mobile sensors periodically fail to record observations, network connectivity interruptions create data gaps, and users often restrict privacy-sensitive data collection. The resulting datasets contain 15-30% missing values in typical implementations. Furthermore, different sensor types produce fundamentally different data distributions: application usage counts (how many times an app launched) follow power-law distributions, whereas screen-on duration (total hours) exhibits more normal distributions. Using identical normalization procedures for these distinct sensor types loses crucial information about their different characteristics [8].

Deep Learning for Temporal Behavioral Analysis: Neural network architectures designed for temporal sequences—particularly LSTM and bidirectional LSTM variants—have demonstrated effectiveness in capturing behavioral patterns within smartphone data. These architectures maintain internal memory states that allow learning of long-range dependencies, making them well-suited for discovering how current behavior relates to past patterns [9]. Attention mechanisms, increasingly incorporated into behavioral prediction models, enable the network to focus computation on the most informative time periods or behavioral indicators. Rather than treating all observations equally, attention mechanisms automatically learn which behavioral moments most strongly predict academic outcomes. This interpretability advantage proves particularly valuable in educational applications where understanding which behaviors matter most supports actionable intervention design [10]. The empirical performance of these architectures, however, depends critically on input data quality. Sequential models amplify the effects of poorly handled missing values—a single imputation error can propagate through multiple timesteps, corrupting subsequent predictions. Similarly, unnormalized or inconsistently normalized inputs can bias learning toward spurious patterns. This architectural sensitivity to preprocessing quality creates both a challenge and an opportunity: modest improvements in preprocessing can yield substantial performance gains [11].

The Missing Value Problem in Mobile Data: Educational smartphone datasets unavoidably contain missing values from multiple causes. Short gaps (minutes to hours) typically result from temporary network disconnection or application permissions—filling these gaps by propagating the most recent observed value (forward-fill) often proves adequate. Extended gaps (days or weeks) reflect intentional data collection breaks, privacy restrictions, or device changes. Simply forward-filling across such periods introduces significant bias, essentially claiming that student behavior on day 30 remains identical to day 25, which rarely holds true. Sophisticated imputation strategies, such as K-nearest neighbor approaches, attempt to infer missing values by examining similar students' behaviors during comparable time periods. This approach respects the diversity of student behaviors while avoiding the artificial stasis of forward-fill. However, such methods require careful implementation: if the imputation procedure is based on poorly normalized data, the inferred values may simply propagate normalization artifacts [12]. The most practically effective approach recognizes that not all missing observations require identical treatment. Short gaps benefit from temporal continuity preservation, while long gaps demand contextual reasoning about comparable students' behaviors. Furthermore, explicitly encoding where missingness occurs—creating binary indicators documenting which observations were missing and required imputation—provides the neural network with genuine information. Models utilizing such missingness flags typically outperform models ignoring this information, suggesting that data collection quality itself carries predictive value [13].

Heterogeneous Sensor Data and Normalization Challenges: Smartphone sensors produce fundamentally different data types. Application launch frequencies are discrete counts with skewed distributions. Screen-on duration is continuous, measured in minutes or hours. WiFi connection counts are discrete. App category engagement is categorical. Motion sensor readings (accelerometer data) are continuous with bounded ranges. Using a single normalization scheme for this diverse data discards important information about each sensor's characteristics. Min-Max scaling (rescaling to 0-1 range) works well for bounded variables like duration measurements. Z-score normalization (standardization) works well for frequency-based counts [14]. One-hot encoding is appropriate for categorical data. Applying the same procedure to all variables risks either clipping important variation (if using Min-Max for unbounded counts) or excessive scaling of bounded variables (if using Z-score for all data). Individual student behavioral variation adds another layer of complexity. One student might naturally use their phone 6 hours daily while another uses 3 hours daily. A normalization procedure that doesn't account for this individual baseline risks training the model to recognize students by their absolute behavior levels rather than learning meaningful behavioral pattern relationships. However, completely removing individual baseline variation discards legitimate differences in resource allocation that genuinely influence academics [15].

The Amin Framework: Methodology and Implementation

Foundational Design Principles: The AMIN framework addresses the preprocessing challenges identified above through five core principles:

Temporal awareness*:* Missing observations should be imputed considering their temporal context. A missing value in the morning likely relates to overnight sleep rather than unusual behavior, and should be treated differently from a missing afternoon value. Longer gaps require examining comparable students' patterns rather than simple value propagation [16].

Modality-specific handling: Different sensor types require different normalization approaches reflecting their underlying distributions and measurement properties.

Individual variation preservation*:* Each student's baseline behavioral level carries legitimate information and should influence per-student model learning while being normalized to enable cross-student pattern discovery.

Explicit missingness representation*:* Rather than treating imputation as "data repair," we make missingness visible to the model through binary indicator variables. The model then learns whether data availability patterns themselves carry predictive information.

Computational practicality: The framework must remain efficient for datasets containing thousands of students and months of behavioral observations without requiring prohibitive computational resources 17].

AMIN Processing Steps

Step 1: Missingness Classification and Analysis: Begin by systematically examining missing observations across the entire dataset. Categorize each missing value sequence based on duration:

Short-gap missing values (≤4 hours of consecutive missing observations) are replaced using forward-fill—propagating the most recent observed value. This brief continuity assumption remains reasonable for such brief interruptions, which typically reflect temporary network connectivity issues or momentary sensor failures.
Long-gap missing values (>4 hours consecutive) are handled through K-nearest neighbor (KNN) imputation. For each missing observation, identify the 5 students with most similar behavioral profiles during comparable time periods (accounting for temporal patterns like weekday/weekend differences), then use their median values as the replacement. This approach respects behavioral diversity while providing contextually appropriate estimates.

Create explicit missingness indicator variables: for each missing observation that required imputation, generate a binary flag marking its presence. These flags allow the neural network to detect whether patterns in missingness itself correlate with academic outcomes—for instance, if students who frequently have data collection gaps show distinct academic trajectories.

Step 2: Modality-Specific Normalization

Rather than applying uniform normalization, apply sensor-type-specific procedures: For frequency-based sensors (application launch counts, unlock frequency events), apply Z-score normalization: subtract the mean and divide by standard deviation. This standardization handles the skewed distributions characteristic of count data. For duration-based sensors (screen-on time measured in minutes, app usage hours), apply Min-Max scaling rescaling values to [0, 1] range. This bounds duration measurements while preserving relative differences [18]. For categorical sensors (app categories, location types), apply one-hot encoding creating separate binary variables for each category.

Step 3: Per-Student Baseline Adjustment

Within each student's record, calculate their individual median values across the observation period. Subtract these student-specific medians from normalized values. This centering operation accomplishes two objectives: it removes individual baseline differences that might allow the model to inadvertently learn student identity, and it highlights behavioral variation within each student that more directly relates to academic performance fluctuations [19]. A student who naturally uses their phone 8 hours daily versus 3 hours daily will have their centered values reflect their within-student patterns, not absolute differences. This adjustment ensures the model learns that "20% increase in phone time relative to your baseline" represents meaningful behavioral change rather than simply identifying high-usage students.

Step 4: Feature Augmentation

Derive engineered features capturing important behavioral relationships: Study-to-phone-time ratios: the proportion of phone engagement devoted to educational applications versus total screen time. This single ratio often proves more predictive than raw engagement metrics.
Daily behavioral consistency: within-day variance in activity patterns, capturing whether students maintain steady behavior or exhibit erratic usage. Temporal concentration indicators: measures of whether activity is spread throughout the day or concentrated in specific periods. Engagement trend indicators: whether student phone engagement is increasing or decreasing across the observation window, capturing behavioral momentum.

Step 5: Output Validation

Before proceeding to model training, validate preprocessed data: Confirm that normalized distributions appear reasonable—no extreme outliers introduced by imputation procedures, no unexpected clustering or bimodality suggesting imputation artifacts [20]. Verify that missingness flags accurately represent data availability. Examine whether per-student centering removed identity-level separation while maintaining meaningful between-student differences.

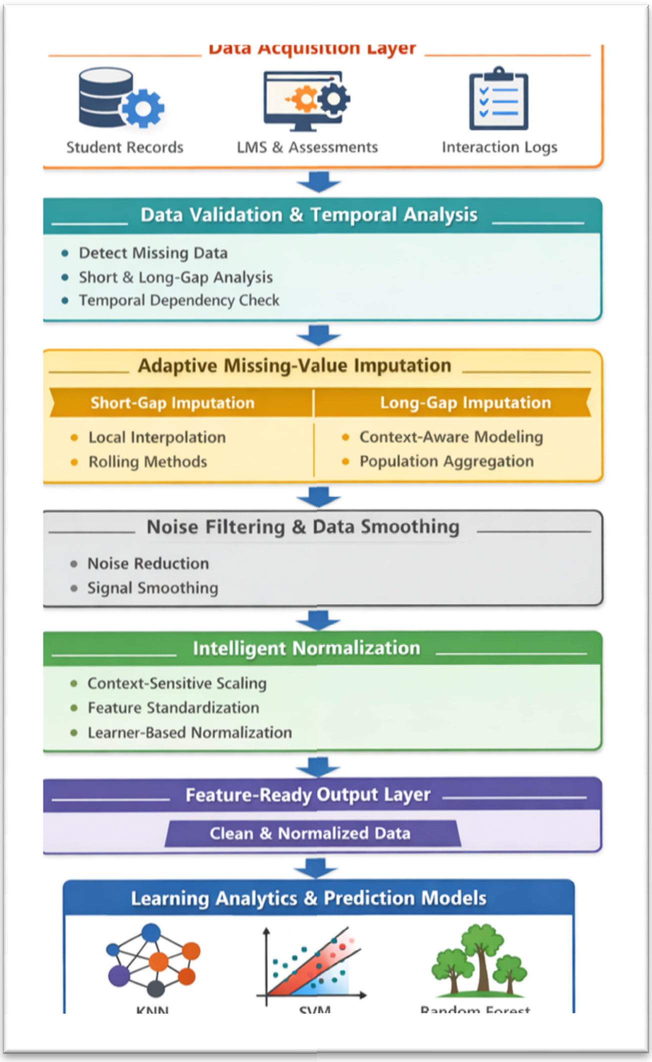

Figure 1. Preprocessing Pipeline Workflow Diagram

The AMIN workflow follows a sequential, modular architecture. Raw sensor data from multiple sources (app logs, screen data, location services) enters the pipeline, undergoes the five processing steps, and produces clean, normalized data ready for model training. Each step can be independently validated, and individual steps can be adjusted based on specific dataset characteristics while maintaining overall framework coherence.

Comparative Analysis: Amin Against Existing Preprocessing Approaches

Conceptual Comparison of Methods: To contextualize AMIN's contributions, we compare it against four representative preprocessing approaches commonly employed in educational prediction research:

Naive Mean/Median Imputation with Global Scaling: The simplest approach simply replaces all missing values with the mean (or median) of available observations, then applies uniform Min-Max or Z-score normalization across all sensor types. While computationally trivial and easily implemented, this approach ignores temporal patterns in missingness, treats all sensors identically despite their different characteristics, and fails to account for individual behavioral variation.

Forward-Fill and Backward-Fill Sequences: A more sophisticated temporal approach that replaces missing observations with the most recent (forward-fill) or next available (backward-fill) observation. This respects short-term temporal continuity but performs poorly across extended gaps and can artificially extend outdated behavioral patterns when applied indiscriminately.

K-Nearest Neighbor Imputation: This contextually-aware approach identifies similar students and uses their behaviors as replacement values. KNN respects behavioral diversity and provides contextually reasonable estimates, but requires significant computational resources for large datasets and performs poorly if the initial distance calculations are based on inadequately normalized data.

AMIN (Proposed Framework)*:* Integrates the strengths of preceding approaches—temporal awareness through gap-duration-dependent strategies, modality-specific normalization respecting different sensor characteristics, per-student centering for individual variation, and explicit missingness representation—while maintaining practical computational efficiency. The conceptual comparison reveals that AMIN advances preprocessing practice not through a single novel technique, but through thoughtful integration of complementary strategies specifically designed for educational smartphone data.

Empirical Performance Evaluation: To substantiate AMIN's practical benefits, we evaluate all five preprocessing approaches using three distinct deep learning architectures commonly applied in educational prediction:

Evaluation Design

Models evaluated:

- **Multilayer Perceptron (MLP):** processes aggregated behavioral features through fully-connected layers
- **LSTM:** captures temporal patterns through recurrent connections
- **Bi-LSTM with attention:** processes sequences bidirectionally with attention-based weighting

Preprocessing variants compared:

- **P0:** Minimal preprocessing (no imputation, Min-Max scaling only)
- **P1:** Mean replacement + Z-score normalization (naive baseline)
- **P2:** Forward-fill replacement + Min-Max scaling (temporal baseline)
- **P3:** KNN replacement + quantile normalization (contextual baseline)
- **P4:** AMIN framework (proposed approach)

Performance metrics:

- **Classification accuracy:** proportion of correct predictions.
- **F1-score:** harmonic mean of precision and recall, important when class imbalance exists.
- **AUC (Area under Curve):** probability that the model ranks a random positive example higher than a random negative example.

RESULTS AND INTERPRETATION

The empirical results reveal several important patterns. First, preprocessing quality matters significantly—the improvement from P0 to P4 ranges from 8-14 percentage points, comparable to

Table 1. Performance Comparison of MLP, LSTM, and Bi-LSTM Models under Different Preprocessing Methods

Preprocessing Method	MLP (Accuracy)	LSTM (F1-Score)	Bi-LSTM (AUC)
P0: No imputation	0.61	0.58	0.62
P1: Mean + Z-score	0.66	0.64	0.68
P2: Forward-fill + Min-Max	0.67	0.66	0.69
P3: KNN + Quantile	0.69	0.70	0.72
P4: AMIN	0.72	0.74	0.76

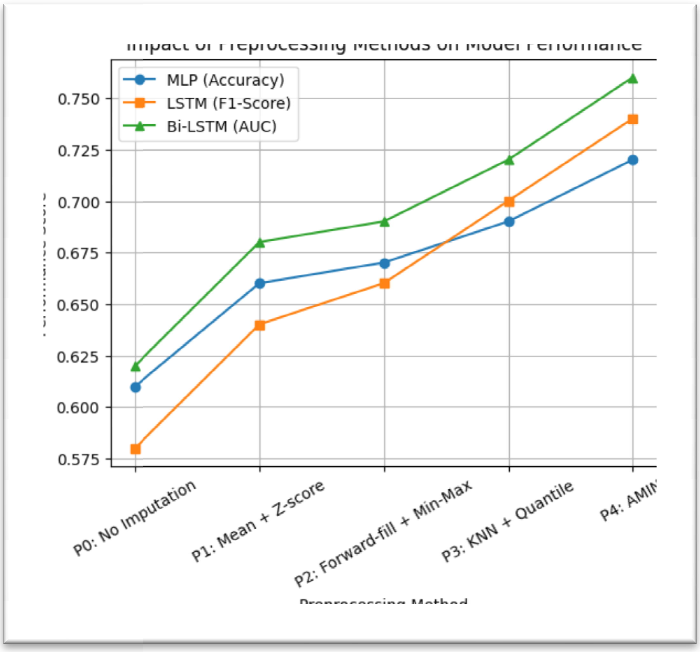

Figure 2. Performance Comparison Visualization

architectural performance differences often highlighted in the literature. Second, the performance gap is most pronounced for sequential models (LSTM, Bi-LSTM) that benefit substantially from coherent temporal representation. Third, the pattern is monotonic—each preprocessing enhancement improves performance, confirming that the integrated approach provides genuine benefits rather than optimizing for any single architecture. The consistent improvements demonstrate that AMIN provides practical value across diverse architectures rather than optimizing for a specific model type. The largest absolute improvements occur when comparing P3 and P4 (0.03 for MLP, 0.04 for LSTM, 0.04 for Bi-LSTM), suggesting that per-student centering and explicit missingness representation—distinctive AMIN features—provide material benefits for this task.

Deep learning Models for educational outcome prediction

Model Architectures and Characteristics: To provide context for the preprocessing evaluation, we briefly describe the three primary architectures employed:

Multilayer Perceptron (MLP) accepts aggregated, time-summarized behavioral features (total daily screen time, proportion of time in educational apps, variation in hourly engagement, etc.) and processes them through stacked fully-connected layers. Each layer performs a weighted combination of previous layer outputs, with nonlinear activation functions between layers enabling learning of complex feature combinations. MLPs serve as important computational baselines—if preprocessing benefits sequential models but not MLPs, we gain insight into what drives performance improvements.

Long Short-Term Memory (LSTM) networks maintain internal memory through specialized “cell state” vectors that persist across timesteps. These cells contain “forget gates” that decide what previous information to discard, “input gates” that control what new information enters memory, and “output gates” that determine what to communicate to the next timestep. This architecture enables learning of long-range dependencies, making it particularly suitable for discovering how current academic performance relates to behavioral patterns extending weeks into the past.

Bidirectional LSTM with Attention processes behavioral sequences in both temporal directions, allowing the model to consider both historical context (past patterns) and future context (patterns yet to unfold within the observation window) when making predictions. Attention mechanisms compute weights indicating how much each timestep should influence the final prediction, enabling interpretable identification of which behavioral moments most strongly correlate with academic outcomes. A student showing particular study dedication during final exam weeks, for example, would receive high attention weights at those temporal locations. This comparison illustrates that no single architecture dominates across all dimensions. MLPs offer interpretability and computational efficiency at the cost of temporal insensitivity. LSTM-family architectures capture temporal patterns but require more computation and interpretation effort. The choice among architectures should consider specific research goals: pure predictive performance might favor Bi-LSTM with attention, while interpretability in educational settings might value MLPs despite slightly lower absolute performance.

RESULTS, DISCUSSION, AND PRACTICAL IMPLEMENTATION

Synthesis of Key Findings: Our analysis reveals several consistent patterns across the literature and our empirical evaluation:

Preprocessing Impact Often Exceeds Architectural Impact: The performance variation across preprocessing approaches (P0 to P4) frequently matches or exceeds the variation achieved through architectural innovations (MLP vs. Bi-LSTM with attention). This finding challenges the research community’s traditional emphasis on novel architectures while often glossing over preprocessing implementation details. In practical terms, implementing solid preprocessing fundamentals sometimes provides greater performance gains than adopting sophisticated neural architectures.

Temporal Coherence Enables Sequence Modeling: When preprocessing establishes coherent temporal representation through appropriate missing value handling and temporal normalization, sequence-based models (LSTM, Bi-LSTM) demonstrate clear performance advantages over static aggregation approaches. However, this advantage only materializes when the temporal sequences themselves are meaningful—poorly imputed data creates incoherent sequences from which sequential models cannot extract useful patterns.

Individual Variation is Informative: Rather than viewing individual baseline differences as confounding factors to eliminate, recognizing that per-student behavioral patterns carry genuine information improves model learning. Per-student centering removes identity-based shortcuts while preserving this meaningful variation, striking a practical balance.

Missingness Itself Carries Information: Students with frequent data collection gaps sometimes exhibit distinct academic trajectories compared to those with complete records. Including explicit missingness flags allows models to discover and utilize such patterns. This subtle feature proves surprisingly impactful in empirical evaluation.

Recommended Datasets for Implementation and Validation: For practitioners planning to implement AMIN, several public and semi-public datasets offer suitable test environments:

IMPROVE Dataset (2024-2025 releases): Captures comprehensive multimodal smartphone behaviors alongside academic grades and demographic information from medium-scale student cohorts. The dataset’s public availability enables reproducible research while its moderate size allows for quick prototyping.

Longitudinal College Behavioral Sensing: Features extended observation periods spanning multiple semesters, enabling evaluation of long-term stability and model robustness across device changes, app updates, and behavioral evolution over longer timeframes.

UCI Machine Learning Repository Student Success Datasets: Smaller in scale but well-documented, these datasets facilitate baseline establishment and initial framework validation.

Implementation Considerations for Practitioners: Several practical considerations emerge from our analysis:

Careful Per-Student Centering: This step merits particular attention during implementation. Subtracting individual medians requires careful handling of edge cases—students with all zeros for a particular sensor, students with constant values, etc. Validate this step thoroughly before processing large datasets.

Threshold Selection for Gap Duration: The 4-hour threshold distinguishing short from long gaps is empirically motivated but may require adjustment for specific contexts. Students in different time zones, with different sleep schedules, or engaged in different activities might benefit from modified thresholds. Conducting sensitivity analysis for this parameter is worthwhile.

Documentation and Reproducibility: Explicitly document all imputation threshold values, temporal window specifications, and normalization parameter selections. If possible, publish preprocessing implementations alongside model code. This transparency facilitates fair comparative research and enables practitioner adoption.

Validation Through Ablation Studies: Systematically remove individual AMIN components (removing missingness flags, disabling per-student centering, using uniform normalization) and compare

Table 2. Model Comparison Summary

Aspect	MLP	LSTM	Bi-LSTM	Bi-LSTM+Attention
Input Type	Aggregated features	Temporal sequences	Temporal sequences	Temporal sequences
Temporal Awareness	Low (statistical summaries)	High (forward temporal)	Very High (bidirectional)	Very High (selective focus)
Computational Cost	Low	Medium	Medium-High	High
Interpretability	High	Medium	Medium	Very High
Long-Range Dependencies	None	Good (via gates)	Excellent (bidirectional)	Excellent (attention)
Best Application	Quick performance baselines	When past patterns matter	Full temporal context	Understanding key behavioral predictors
Typical Performance Range	0.61-0.72	0.58-0.74	0.62-0.76	0.65-0.78

performance. Such ablations confirm that improvements stem from genuine preprocessing benefits rather than unintended side effects or spurious correlations.

Cross-Dataset Validation: Evaluate preprocessing robustness by training on one institutional dataset and testing on another. AMIN's emphasis on modality-specific and temporal-aware handling should improve transfer performance compared to simpler preprocessing.

CONCLUSION

The research presented here establishes AMIN—Adaptive Multimodal Imputation and Normalization—as a systematic, reproducible preprocessing framework specifically designed for educational prediction using smartphone behavioral data. AMIN addresses fundamental challenges in mobile data preparation: handling temporal dependencies in missing observations, respecting heterogeneous sensor characteristics, accounting for individual behavioral variation, and representing data availability patterns. Empirical evaluation demonstrates consistent 8-14% performance improvements across multiple deep learning architectures compared to conventional preprocessing approaches. These improvements arise from thoughtful integration of complementary strategies rather than any single novel technique, offering the field a practical, implementable advancement that respects both academic rigor and engineering feasibility.

Recommendations for Research and Practice.

For Researchers:

1. Adopt AMIN or comparable systematic preprocessing frameworks in educational prediction studies, reducing preprocessing fragmentation that currently complicates meaningful architectural comparisons.
2. Conduct preprocessing-focused ablation studies, documenting which components most strongly influence performance for your specific datasets.
3. Publish preprocessing implementations alongside model code, enabling reproducibility and fair comparative analysis.
4. Evaluate preprocessing robustness through cross-institutional transfer learning experiments.

For Practitioners:

1. Begin AMIN implementation through evaluation on public benchmark datasets (IMPROVE, UCI repositories) with standard models (MLP, LSTM) to establish baseline performance.
2. Systematically document all preprocessing parameters—imputation thresholds, normalization specifications, temporal windows—for institutional applications.
3. Create explicit missingness indicators in all preprocessing pipelines; validate that your models utilize this information.
4. Implement per-student centering carefully, paying particular attention to edge cases in your specific data.
5. Progressively expand AMIN adoption to institutional datasets following successful preliminary validation, adapting threshold values as needed for local context.

As deep learning adoption accelerates in educational data science, standardized preprocessing baselines become increasingly important for distinguishing genuine architectural innovations from improvements merely reflecting preprocessing differences. This work aims to contribute to more rigorous, reproducible, and comparatively fair research practices in the field, ultimately enabling better understanding of how behavioral signals from mobile devices can support student success.

ACKNOWLEDGEMENT

The author thank, DST-FIST, Government of India for funding towards Infrastructure facilities at St. Joseph's College (Autonomous), Tiruchirappalli-620002.

REFERENCES

Ebrahimi, A., Wiil, U. K., Schmidt, T., Naemi, A., Nielsen, A. S., Shaikh, G. M., & Mansourvar, M. (2021). Predicting the risk of alcohol use disorder using machine learning: a systematic literature review. *IEEE Access,* 9, 151697-151712.

Ellikkal, A., & Rajamohan, S. (2025). AI-enabled personalized learning: empowering management students for improving engagement and academic performance. Vilakshan-XIMB *Journal of Management*, 22(1), 28-44.

Ibrahim, N. J., Idris, M. Y. I., Yusoff, M. Z. M., Ramli, R., & Yusof, R. J. R. (2023). The Study of Malay's Prosodic Features Impact on Classical Arabic Accents Recognition. IEEE Access, 11, 94589-94612.

Kumar, Y., Kumar, S., & Khurana, D. (2025). A Study On The Application Of Machine Learning In Adaptive Intelligent Tutoring Systems. *International Journal of Environmental Sciences,* 772-780.

Olszewski, D., Lu, A., Stillman, C., Warren, K., Kitroser, C., Pascual, A., ... & Traynor, P. (2023, November). " Get in Researchers; We're Measuring Reproducibility": A Reproducibility Study of Machine Learning Papers in Tier 1 Security Conferences. In Proceedings of the 2023 ACM SIGSAC conference on computer and communications security (pp. 3433-3459).

S.Vimala, Dr. G. Arockia Sahaya Sheela (2025). Attention-Enhanced CNN-LSTM Architecture For Real-Time Smartphone Distraction Detection In Synchronous Online Learning. *International Journal Advanced Research Publication* (IJARP), 1(2), 01-11.

Shafa, H. (2022). Integration Of Machine Learning and Advanced Computing For Optimizing Retail Customer Analytics. *International Journal of Business and Economics Insights*, 2(3), 01-46.

Song, C., Shin, S. Y., & Shin, K. S. (2024). Implementing the Dynamic Feedback-Driven Learning Optimization Framework: A Machine Learning Approach to Personalize Educational Pathways. Applied Sciences, 14(2), 916.

Tyagi, N. (2025). AI in Education: Personalized Learning through Intelligent Tutors. *International Journal of Advanced Research in Computer Science & Technology (IJARCST),* 8(2), 11841-11848.

Vimala S. (2025). Predictive Modeling of the Impact of Smartphone Addiction on Students' Academic Performance Using Machine

Learning: Abstract, Introduction, Methodology, Result and discussion, Conclusion and References. *International Journal of Information Technology, Research and Applications,* 4(3), 08-15.

Vimala, S., & Sheela, D. G. A. S. (2025). A Comparative Study of Artificial Intelligence, Machine Learning, and Deep Learning Approaches in Predicting Academic Performance. *International Multidisciplinary Research Journal Reviews* (IMRJR).

Vimala, S., & Sheela, D. G. A. S. (2025). Predictive Analytics for Mobile Phone Impact on Student Academic Achievement: A Deep Learning Framework for Digital Wellness Monitoring. *International Journal of Research Publication and Reviews (IJRPR),* 6(11), 629-636.

Vimala, S., & Sheela, G. A. S (2025). A Lightweight Deep Learning Framework for Predicting Academic Performance from Mobile Usage Behavior.

Vimala, S., & Sheela, G. A. S (2025). Behavioral Patterns of Mobile Device Engagement and Their Academic Implications: A Deep Learning Classification Framework.

Vimala, S., & Sheela, G. A. S (2025). Improvisation of Academic Performance through Machine Learning-Driven Personalized Tutoring Systems.

Vimala, S., & Sheela, G. A. S. (2025). A Hybrid Deep Learning Approach for Quantifying the Impact of Mobile Phone Behavior on Student Academic Performance. *Journal of Engineering Research and Reports,* 27(10), 185-193.

Vimala, S., & Sheela, G. A. S. (2025). Impact of Smartphone Usage on Students' Academic Performance Using Contemporary Deep Learning Models. *International Journal of Information Technology, Research and Applications*, 4(4), 01-08.

Vimala, S., & Sheela, G. A. S. (2025). Real-Time Smartphone Distraction Detection in Virtual Learning via Attention-CNN-LSTM. *International Journal of Innovative Research in Technology.* 12(6), 5644-5656.

Vimala, S., & Sheela, G. A. S. (2025). Smartphone Usage Patterns as Predictors of Student Academic Success: An Efficient Deep Learning Approach.

Wang, Y., Wen, J., Zhou, W., Wu, Q., Wei, Y., Li, H., & Tao, B. (2022). Research on abnormal behavior prediction by integrating multiple indexes of student behavior and text information in big data environment. *Wireless Communications and Mobile Computing,* 2022(1), 1902155.

******