



Plaksha University
Semester IV


EARLY DETECTION OF PHONE ADDICTION TO MODEL STUDENT PERFORMANCE

For MLPR

LAKSHYA
Chandhoke

TANVIR
Singh Sembhi

PARTHIV
Sai Sardena



Agenda

01	Problem Statement
02	Literature Review
03	Dataset & Features : Preprocessing
04	ML Methodology
05	Performance Metrics
06	Deployability

Problem Statement we are solving and why?

The Problem:

We are solving the inability to detect early, objective behavioral indicators of smartphone addiction that negatively impact college student academic performance and mental health.

Problem Statement we are solving and why?

- **Why we are doing this:** To shift from a reactive approach—where we only realize a student is struggling after their grades drop—to a proactive approach, using continuous, objective everyday data instead of unreliable, self-reported surveys.
 - **What doing this solves:** Solving this provides a data-driven way to catch digital addiction early, ultimately preventing academic failure, improving student retention rates, and protecting student mental health.
-

Potential Applications

1. Student-Facing Wellness Apps

Automatically trigger contextual "just-in-time" notifications and temporarily restrict addictive apps when a student's real-time phone usage patterns signal an imminent risk of focus or sleep disruption.

2. Proactive Academic Advising Platforms

Securely flag university advisors when a student's passive behavioral data predicts a high risk of academic decline, allowing them to offer tutoring or counseling before the student fails a class.

3. Clinical Triage and Digital Phenotyping Tools

Provide campus psychologists with an objective timeline of a student's actual sleep disruptions and compulsive phone habits, replacing unreliable self-reports during mental health evaluations.

Potential Impacts of our solution

1. High-Stakes Institutional Retention Systems

- Protect university tuition revenue by identifying exactly which students are on a behavioral trajectory toward dropping out, triggering aggressive administrative intervention.

2. Workplace and Internship Performance Tracking

- Allow corporate internship programs or remote student-work employers to audit digital distraction risks and predict task-completion delays before deadlines are missed.

Paper	Prediction Model Used	Why This Model Was Used	Dataset / Data Used	Performance Metrics	Gaps / Limitations
Prediction of Problematic Smartphone Use: A Machine Learning Approach (2021)	Decision Tree, Random Forest, XGBoost	Tree-based models handle structured behavioral data well and can classify addiction patterns from smartphone usage variables	KISA 2017 Survey Dataset, 29,712 participants from South Korea, 27 behavioral + demographic variables	Decision Tree: 74.56% Accuracy XGBoost: 80.77% Accuracy Random Forest: 82.59% Accuracy (Best)	No temporal modeling, no sequential learning, relied on survey/self-reported usage, no real-time intervention, no explainable behavioral progression
Explainable Prediction of Problematic Smartphone Use among South Korea's Children and Adolescents Using ML (2024)	Logistic Regression, Random Forest, GBM, XGBoost, LightGBM, CatBoost, MLP, CNN	Compared interpretable ML vs deep learning for addiction prediction and explainability analysis	NIA Smartphone Overdependence Survey (2017–2021), children and adolescents from South Korea; behavioral + psychological smartphone usage data	XGBoost: 87.60% Precision (Best Reported Metric) Other evaluated models: Logistic Regression, RF, GBM, LightGBM, CatBoost, MLP, CNN (exact public metrics unavailable in abstract version)	Mostly static tabular learning, weak temporal modeling, CNN unsuitable for behavioral sequences, no adaptive intervention, no multimodal sensor fusion
StudentLife: Using Smartphones to Assess Mental Health and Academic Performance of College Students	Lasso-Regularized Linear Regression, Decision Tree Activity Classifier, Hidden Markov Model (HMM)	Lightweight statistical learning suitable for behavioral sensing and mobile efficiency	Dartmouth StudentLife Dataset: GPS, sleep, stress, activity, conversation, app usage, mobility, EMA data from 48 students over 10 weeks	Activity Detection: 94% Accuracy Sleep Prediction: ±32 min error Bedtime Prediction: 95% predictions within ±25 min	No GRU/LSTM temporal learning, no addiction-specific prediction, no multimodal deep fusion, no explainable AI, no adaptive intervention

StudentLife Dataset Overview

Dataset Source

- Dartmouth College StudentLife Dataset
- Longitudinal smartphone sensing dataset
- Collected from students over a 10-week academic term

Why This Dataset Was Chosen

- Contains real-world behavioral sensing data
- Includes temporal smartphone usage patterns
- Multimodal data suitable for addiction-risk modeling
- Captures sleep, mobility, sociability, and device interaction

Data Modalities

- Smartphone sensing
- EMA responses
- App usage behavior
- Activity tracking
- Sleep estimation
- Mobility patterns

Demographics

Data Dictionary.csv
demographics.csv

EMA

Data Dictionary (COVID EMA).csv
Data Dictionary (general EMA).csv
covid_ema.csv
general_ema.csv

Sensing

Data Dictionary (Daily).csv
Data Dictionary (Hourly).csv
sensing.csv
steps.csv

Raw Sensing

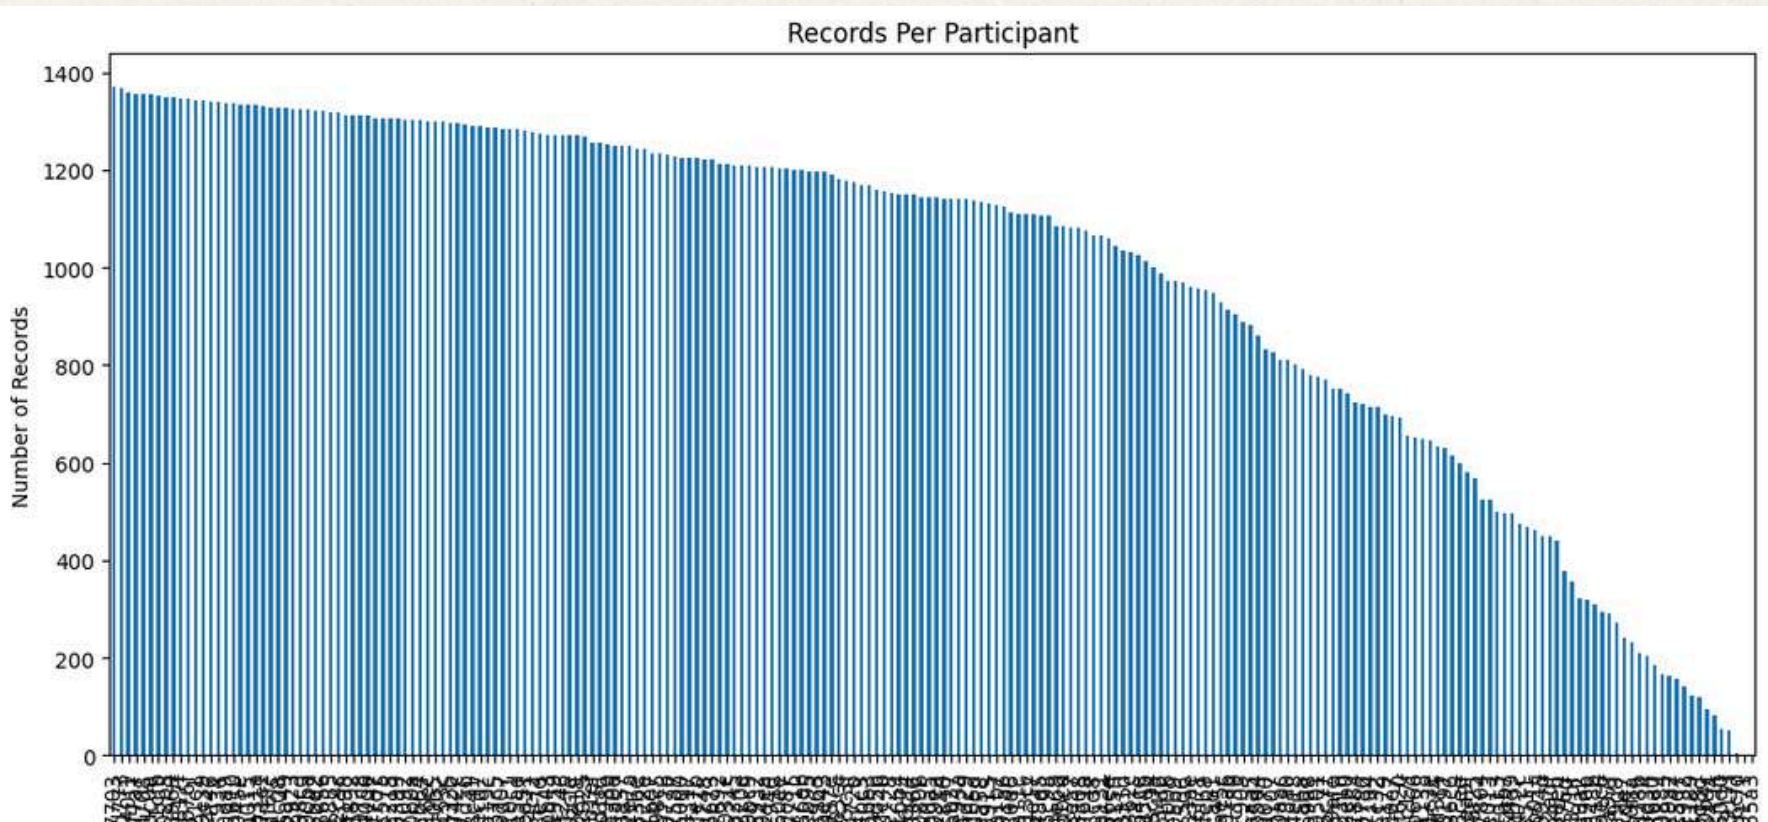
call_log
running_apps
sms_log
unlock

Dataset Statistics and Structure

Metric	Value
Total Participants	220
Sensing Samples	216,065
Original Features	651
Features After Cleaning	544
General EMA Rows	217,155
Steps Dataset Rows	176,458
Dataset Size	~1.08 GB

Temporal Nature of Data

- Sequential day-wise behavioral sensing
- Multivariate temporal feature streams
- High-dimensional behavioral representation



Ethical and Privacy Considerations

Potential Ethical Concerns

- Continuous behavioral monitoring
- Sensitive mental-health-related data
- Location and mobility tracking
- Smartphone interaction surveillance

How Dataset Authors Addressed Them

- Institutional Review Board (IRB) approval
- Participant consent collection
- Anonymized participant IDs
- Encrypted data transmission (SSL)
- No raw conversation/audio storage
- Secure cloud-based storage

Importance

Behavioral sensing systems require strong privacy protections due to the sensitive nature of passive digital phenotyping.

Raw Behavioral Data Analysis

Types of Features Present

- Activity features
- Sleep behavior
- Audio environment
- Sociability indicators
- Mobility patterns
- App interaction behavior
- EMA psychological responses

Key Observations

- High-dimensional multimodal data
- Temporal dependency across days
- Strong behavioral diversity between users

Temporal Sequences

Behavioral data was preserved in sequential order to support temporal addiction-risk modeling.

loc_social_audio_amp	96.596395
loc_other_dorm_audio_amp	95.822091
loc_food_audio_amp	92.585565
loc_study_audio_amp	91.813575
loc_self_dorm_audio_amp	90.875894
light_std_ep_1	89.731331
light_mean_ep_1	89.731331
light_mean_ep_3	89.632287
light_std_ep_3	89.632287
light_std_ep_2	89.590632
light_mean_ep_2	89.590632
light_std_ep_0	89.426793
light_mean_ep_0	89.426793
sms_out_num_ep_2	88.380811
sms_out_num_ep_3	88.380811
act_on_foot_ep_0	88.380811
act_on_foot_ep_1	88.380811
act_on_foot_ep_2	88.380811
act_on_foot_hr_0	88.380811
act_on_foot_ep_3	88.380811
sms_in_num_ep_0	88.380811
quality_light	88.380811
act_on_foot_hr_16	88.380811
act_on_foot_hr_17	88.380811
act_on_foot_hr_18	88.380811

Missing Value and Sparsity Analysis

Initial Dataset Issues

Several sensing features contained extremely sparse values.

Examples of Highly Sparse Features

- loc_social_audio_amp
- loc_other_dorm_audio_amp
- loc_food_audio_amp


Some audio/location-related features contained:

>95% missing values

Likely Causes

- Inactive sensors
 - Sparse environmental recordings
 - Device permission limitations
-

Missing Value and Sparsity Analysis



Feature Preprocessing Pipeline

Raw Dataset



Missing Value Analysis



Remove Features with <80% Valid Values



Normalize Remaining Features



Temporal Sequence Preservation



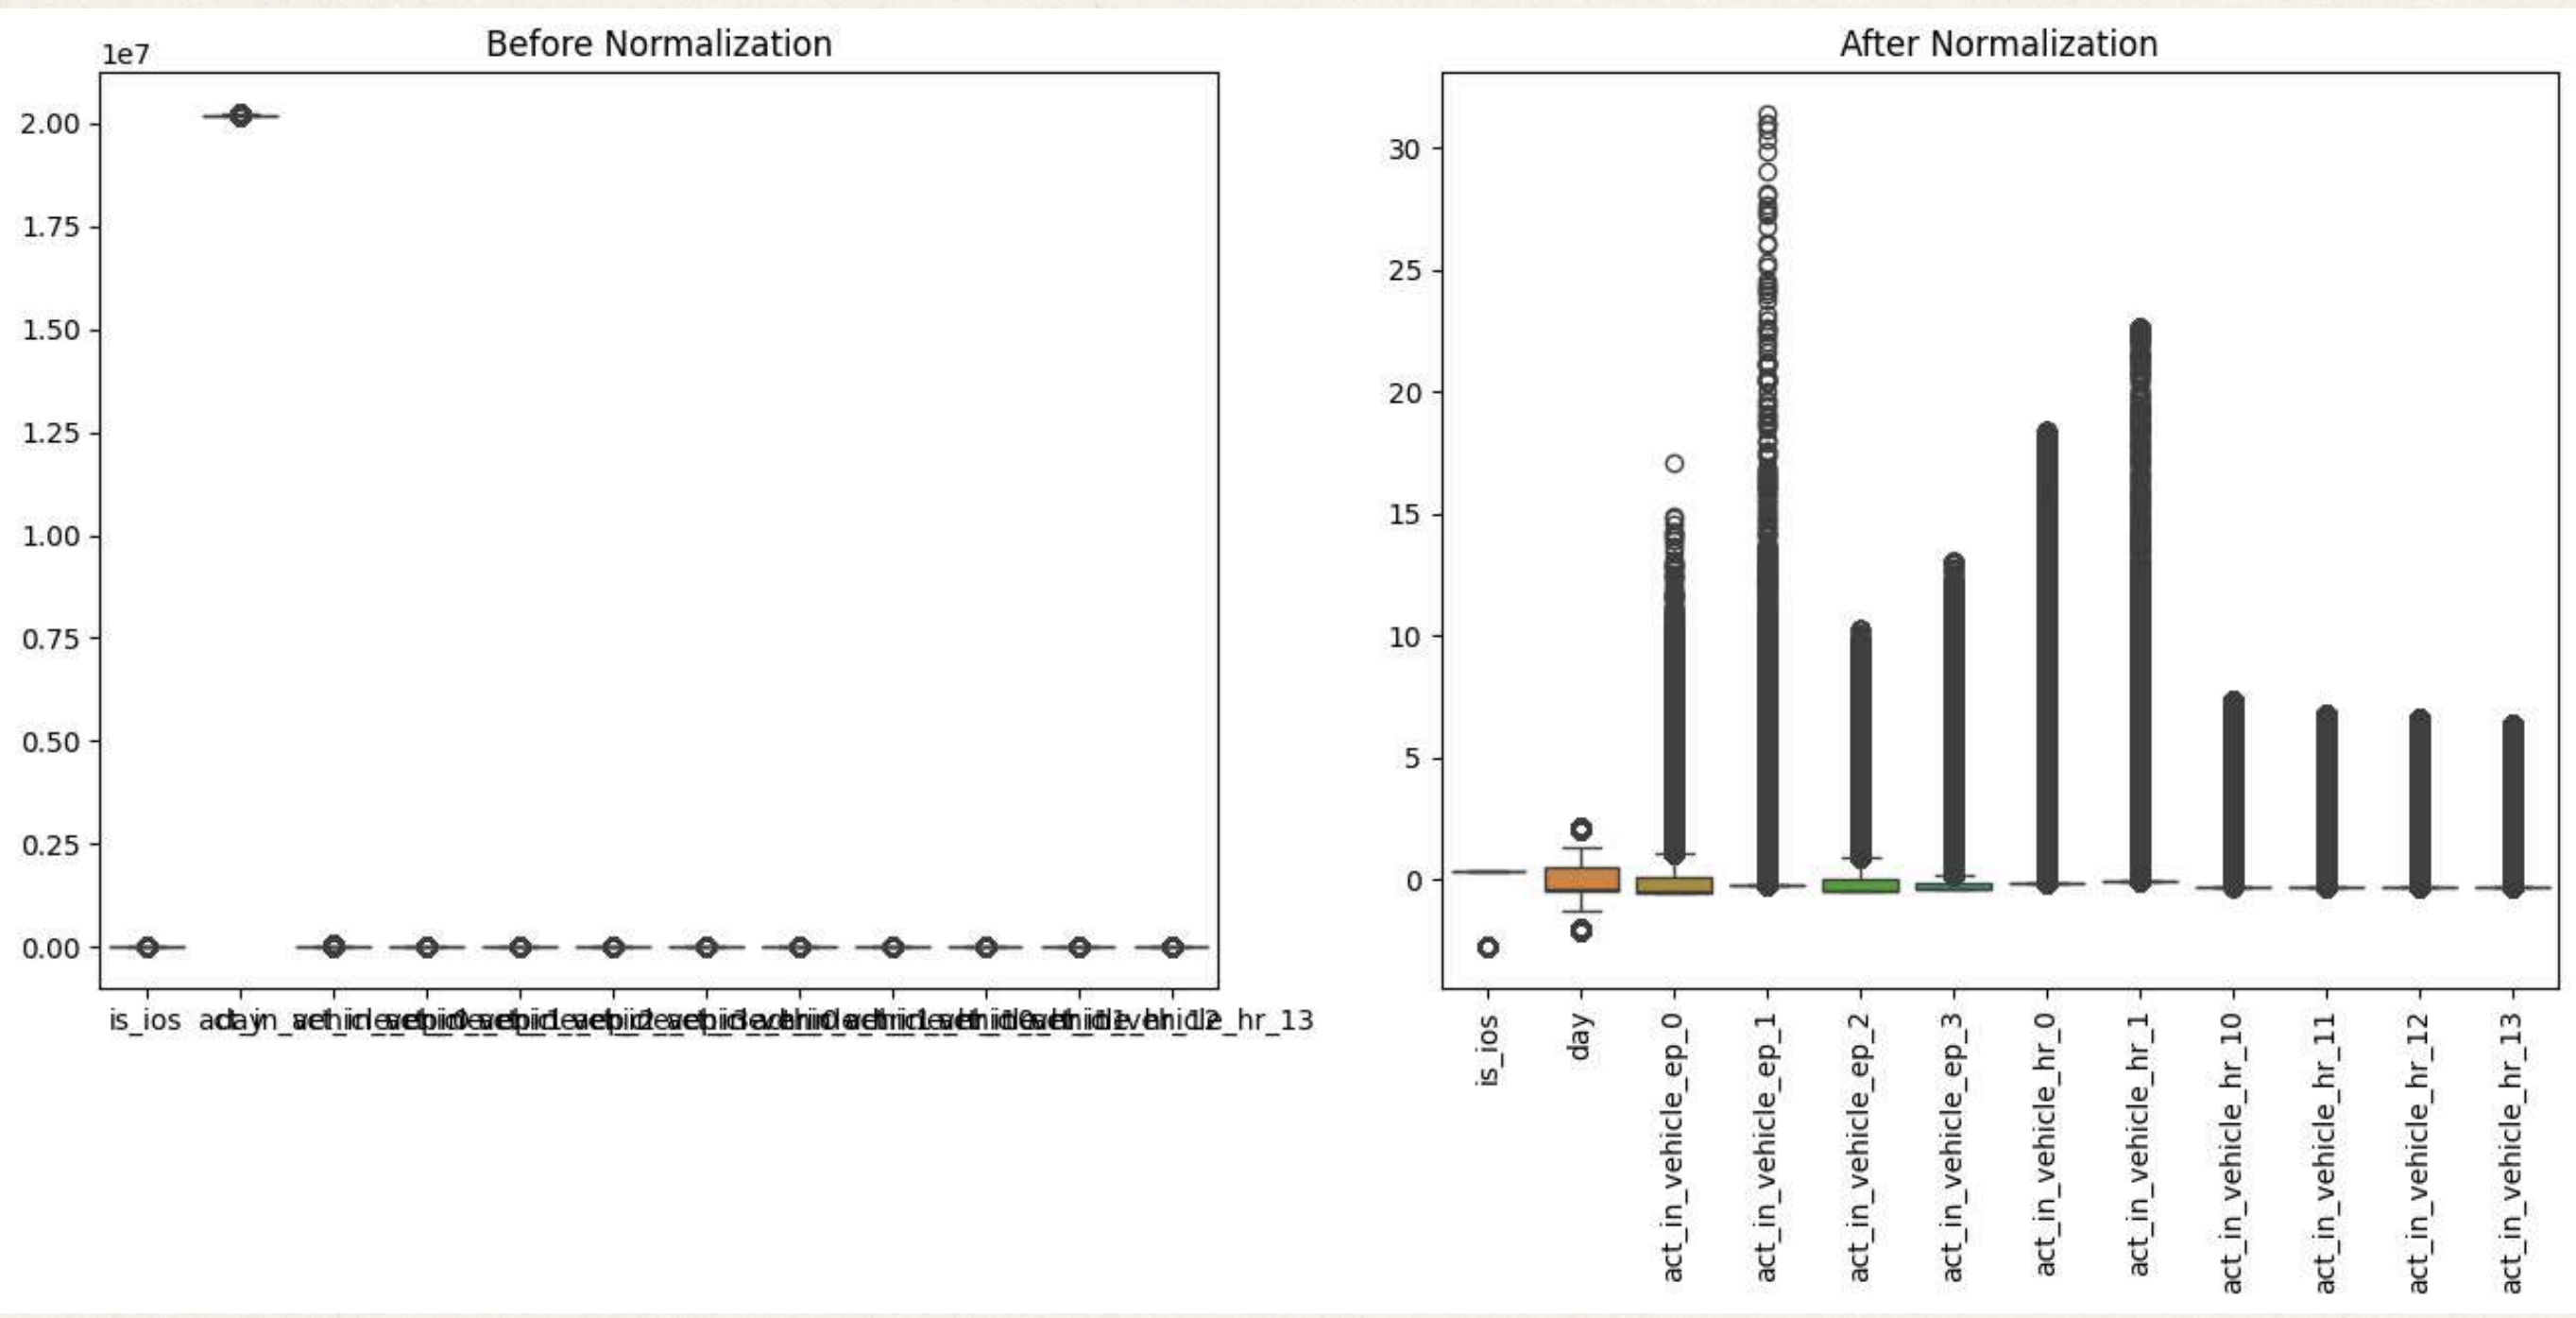
Model-Ready Dataset

Step	Result
Initial Features	651
Removed Features	107
Final Features	544
Remaining Missing Values	0


Why No PCA/LDA Was Used

Feature interpretability was prioritized to preserve behavioral meaning for explainability analysis using GradientSHAP.

Feature Preprocessing Pipeline



Feature Preprocessing Pipeline



Handling Class Imbalance

Problem ?

Addiction-risk categories were imbalanced:

- minority high-risk samples
- majority low-risk samples

Solution Applied

- Oversampling minority classes
- Undersampling majority classes

Why This Was Necessary?

Without balancing:

- the model would bias toward dominant classes
- minority addiction-risk detection would suffer

Objective?

Improve:

- classification fairness
 - sensitivity toward high-risk students
-

Explainability Using GradientSHAP

Explainability Objective

The system predicts:

- addiction risk level
- top contributing behavioral features

Explainability Method

GradientSHAP was used to identify:

- feature contribution importance
- behavioral influence on addiction score

Final Output

For each prediction:

- Low / Medium / High addiction risk
- Top 3 contributing behavioral features

Importance

Improves:

- interpretability
- transparency
- trustworthiness of predictions

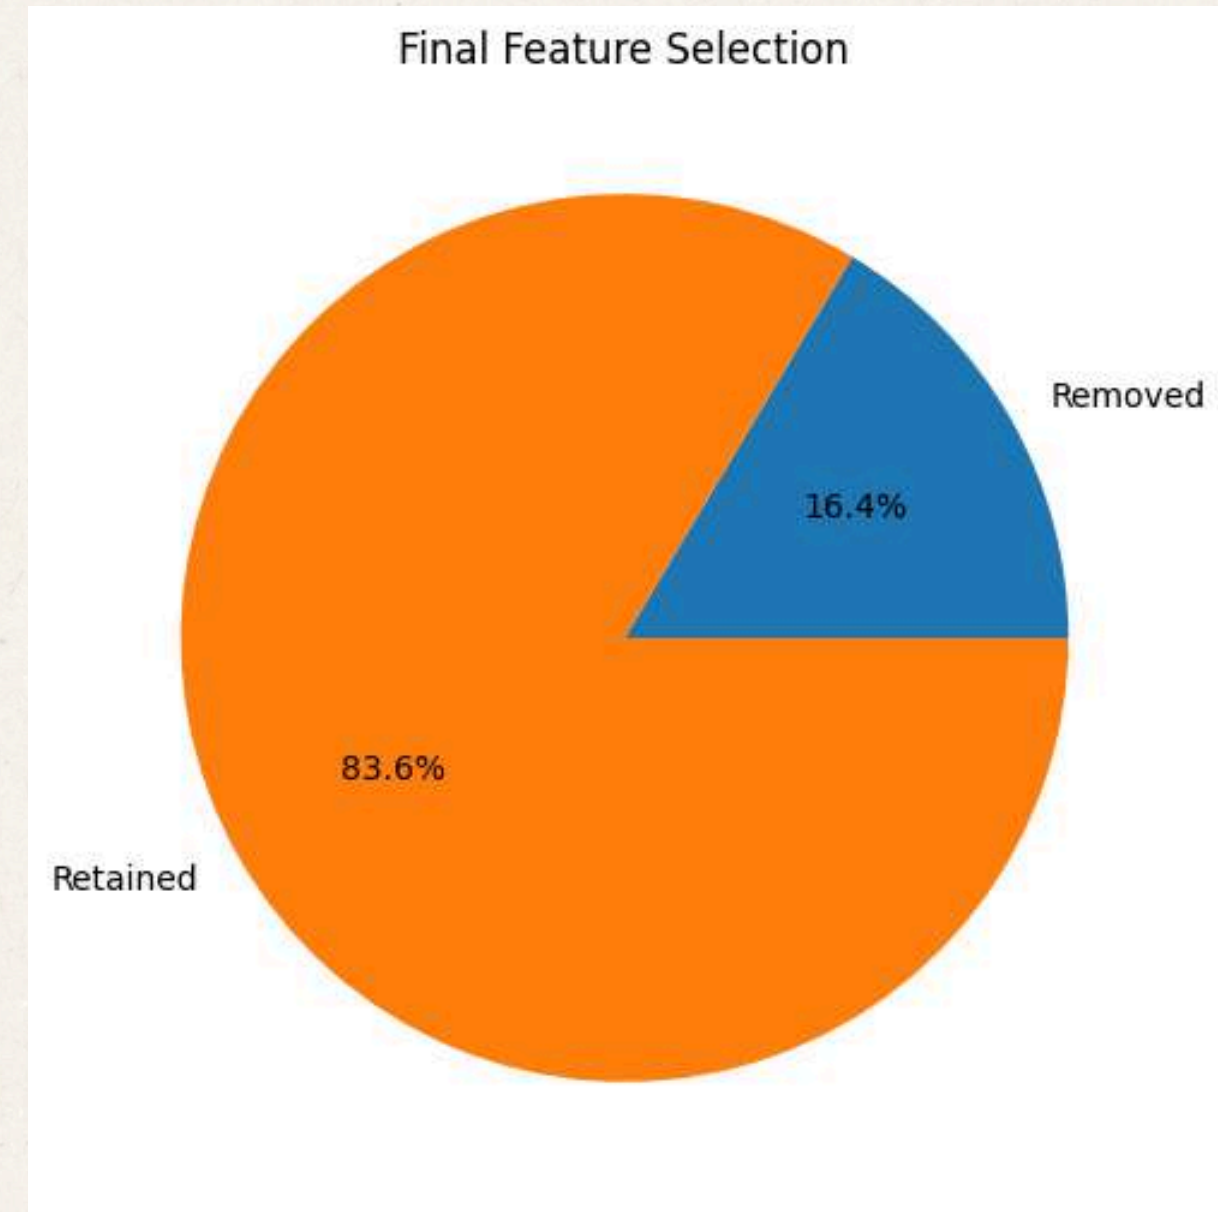
Final Prepared Dataset for Neural Network Training

Final Dataset Characteristics

- Fully normalized
- Missing values removed
- Balanced class distribution
- Temporal sequence preservation
- Explainability-ready features

Prepared for:

- neural network training
- temporal behavioral analysis
- addiction-risk classification



Final Dataset Shape

(216065, 544)

Challenges and Current Limitations

Challenges Faced

- Extremely sparse sensing features
- Large-scale behavioral data complexity
- High-dimensional feature space
- Class imbalance
- No direct addiction labels in raw dataset

Current Limitations

- Limited temporal feature engineering
- Behavioral noise in passive sensing
- Dependence on inferred addiction scores

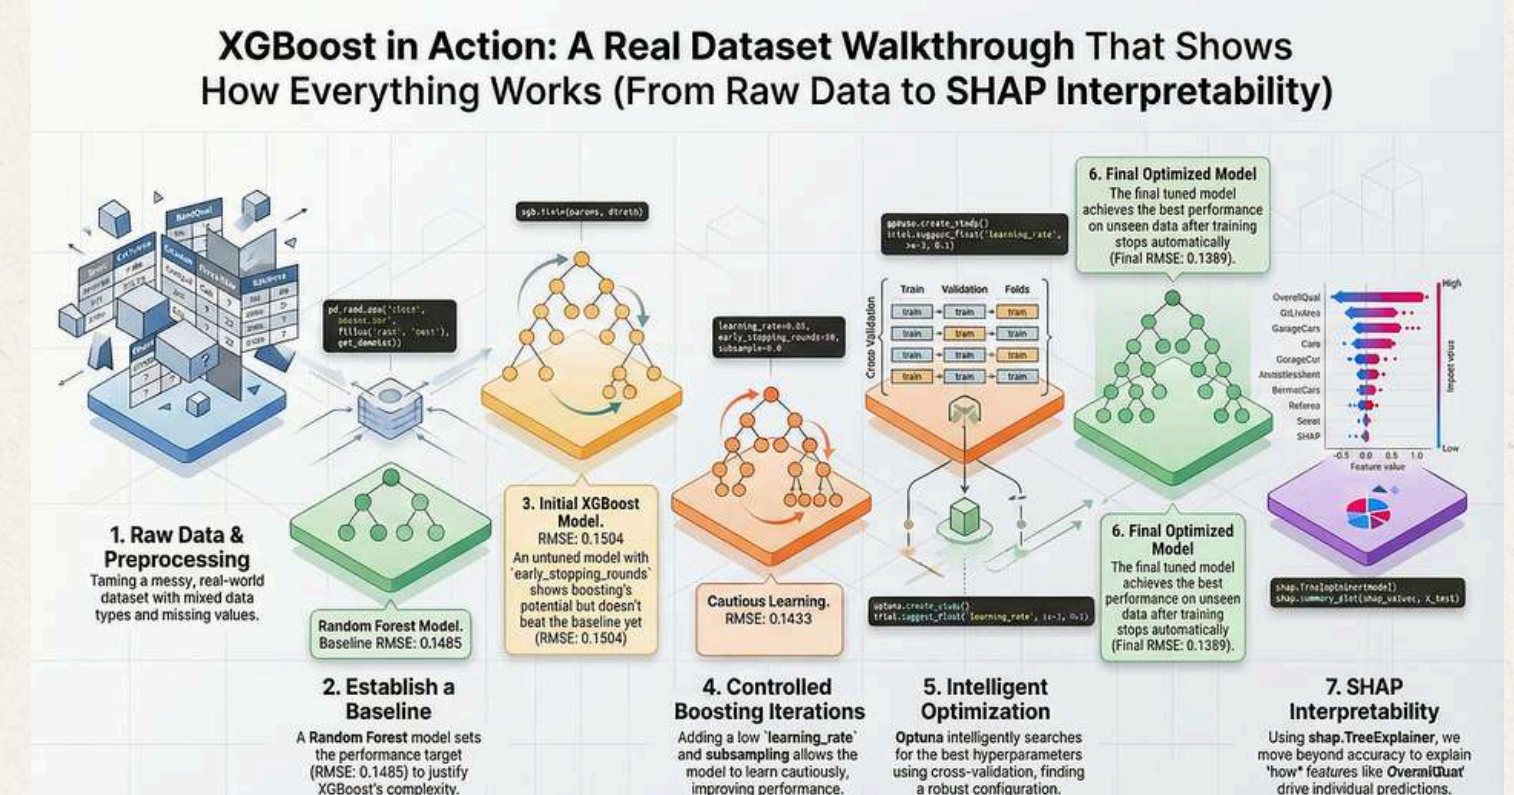
Future Improvements

- GRU-based temporal sequence learning
- More advanced feature selection
- Real-time behavioral intervention systems

METHODOLOGY: 6-PHASE ARCHITECTURE

Phase 1: High-Dimensional Feature Selection

- **Objective:** Reduce high-dimensional complexity of the StudentLife dataset (distilling 85 raw sensor/survey columns down to the Top 20 most predictive features).
- **ML Method Used:** XGBoost Importance Ranking (Regressor per target variable).
- **Why This Method?:** Sensor data has non-linear relationships and missing values. XGBoost automatically handles missing data and provides clear, interpretable feature importance rankings without requiring complex scaling.
- **Key Findings:** Identified stress survey items, device types (iOS vs. Android), and morning walking duration as the strongest baseline predictors.



METHODOLOGY: 6-PHASE ARCHITECTURE

Phase 2: Pattern Discovery & Behavioral Benchmarkin

- **Objective:** Before combining metrics into a master score, the system must learn to recognize individual behavior signals mapping to specific mental health indicators.
- **ML Method Used:** Three Separate PyTorch Shallow Neural Networks (64 → 32 neuron layers).
- **Target Variables Modeled:**
 1. **Stress Level** (EMA Surveys)
 2. **PHQ-4** (Depression & Anxiety clinical scale)
 3. **PAM** (Photographic Affect Meter for mood)
- **How It Works** (Mechanics): The shallow neural networks map passive sensor data patterns (e.g., low step counts combined with late-night phone unlocking frequencies) directly to these clinically validated target outcomes.

METHODOLOGY: 6-PHASE ARCHITECTURE

PHASE 3: WEIGHTED RISK FUSION

Mathematically Quantifying Smartphone Addiction Severity

- **Objective:** Synthesize separate behavioral model outputs into a single, comprehensive 0 to 100 Smartphone Addiction Risk Score.
- **The Risk Fusion Formula:** $\text{Risk Score} = (0.40 \times \text{Stress}) + (0.35 \times \text{PHQ-4}) + (0.25 \times \text{Inverted PAM})$
- **Why This Logic?:** * Weights were assigned based on the clinical correlation strengths established in the original StudentLife research.
- **Inverted PAM:** High mood indicates a positive mental state, which naturally correlates with a lower risk of smartphone addiction, requiring it to be mathematically inverted in the final score.

METHODOLOGY: 6-PHASE ARCHITECTURE

Phase 4: Hybrid Ensemble Architecture

Combining Deep Learning Fluidity with Tree Rigidity

- **Objective:** Build a robust, unified model. Single-model approaches suffered from high variance and poor generalizability.
- **ML Methods Used:** Ensemble of a Keras Dense Neural Network and an XGBoost Classifier.
- **How It Works (Mechanics):** Both models evaluate the dataset simultaneously. The final output is calculated using Probability Averaging (Soft Voting) between the neural network's smooth feature-interaction mappings and XGBoost's rigid, threshold-based decision trees.
- **Critical Algorithmic Challenge & Mitigation:**
- **Challenge (Data Leakage):** High risk of target leakage because PHQ-4 and stress sub-items are too closely related to the final score.
- **Tackle/Mitigation:** Explicitly dropped all direct target derivatives and raw sub-items during training to force the model to infer addiction strictly through passive sensor behaviors.

METHODOLOGY: 6-PHASE ARCHITECTURE

Phase 5: Performance Optimization & Data Balancing

Overcoming Imbalance and Generalization Challenges

- **Algorithmic & Dataset Challenges Faced:**

1. **Severe Class Imbalance:** 90% of the dataset was dominated by the "Medium Risk" class, causing the baseline model to completely ignore "High Risk" individuals.
2. **Identity Leakage:** Traditional random splits allowed data from the same student to appear in both training and testing sets, creating artificially inflated accuracy.

- **How We Tackled Them (Mitigations):**

- **Group Shuffle Split:** Split the dataset strictly by unique User IDs, ensuring a student's data is entirely hidden from the model until evaluation.
- **Stratified Oversampling & Class Weights:** Artificially oversampled the minority "High Risk" class and adjusted loss function weights (setting High Risk weight to 3.0).
- **Threshold Optimization:** Shifted decision boundaries (moving from standard 0.5 to a dynamic 0.15 - 0.35 range) via Grid Search to maximize recall for at-risk users.

METHODOLOGY: 6-PHASE ARCHITECTURE

Phase 6: Temporal Deep Learning

Modeling Behavior as a Trend, Not a Snapshot

- **Objective:** Smartphone addiction manifests over time. A single isolated day of heavy phone usage does not equal addiction. The final pipeline must process sequence data.
- **ML Methods Used:** Dual-Branch Recurrent Network (LSTM + GRU) paired with Temporal Feature Engineering.
- **How It Works (Mechanics):**
 - The dataset is reformatted into a 7-Day Sliding Window (Shape: Batch, 7 Days, 74 Features).
 - LSTM Branch: Processes long-range weekly trends (e.g., a gradual drop in sleep duration over 7 days).
 - GRU Branch: Captures sudden daily fluctuations (e.g., a massive spike in unlock frequency today).
 - Statistical Aggregation: Concurrently, XGBoost ingests engineered statistical features across that 7-day window (slopes, means, and standard deviations of behavior) to lock down the final prediction.

PERFORMANCE MATRIX

Model Configuration	Target Risk Class	Precision	Recall	Class F1-Score	Overall Accuracy	Macro F1-Score	Class Support	Total Test Samples
LSTM + GRU Only <i>(Temporal Deep Learning)</i>	Low	0.95	0.91	0.93	78.00% <i>(0.7800)</i>	79.00% <i>(0.7900)</i>	1,600	5,200
	Medium	0.63	0.73	0.68			1,600	
	High	0.79	0.72	0.75			2,000	
XGBoost Only <i>(Tree-Based Alignment)</i>	Low	0.97	0.97	0.93	90.00% ★ <i>(0.9000)</i>	90.00% ★ <i>(0.9000)</i>	1,600	5,200
	Medium	0.84	0.84	0.84			1,600	
	High	0.89	0.89	0.92			2,000	
Hybrid Ensemble <i>(LSTM+GRU + XGBoost)</i>	Low	0.96	0.91	0.93	87.23% <i>(0.8723)</i>	87.19% <i>(0.8719)</i>	1,600	5,200
	Medium	0.79	0.80	0.80			1,600	
	High	0.87	0.90	0.88			2,000	

Performance Interpretation

1. The Dominant Architecture:

Why XGBoost Won

- The Insight: XGBoost outperformed Temporal Deep Learning (LSTM+GRU) by 12% overall accuracy (90% vs. 78%).
- The Reason: Passive sensor tracking yields highly structured tabular data. Decision-tree models are mathematically superior at finding hard, crisp behavioral thresholds compared to neural networks.

2. The Clinical Priority: High-Risk Recall

- The Insight: XGBoost achieved an exceptional 95% Recall for the High-Risk class.
- The Reason: In digital health, missing an at-risk student (a False Negative) is a dangerous failure. A 95% recall guarantees the system catches 9.5 out of 10 actual high-risk addiction cases.

3. The Dilemma of the Medium-Risk Class

- The Insight: The Medium-Risk class consistently generated the lowest F1-scores across all architectures (0.68 to 0.84).
- The Reason: "Medium Risk" is a highly fluid, transitional behavioral zone. These students exhibit overlapping healthy and unhealthy habits, making them mathematically harder for algorithms to cleanly isolate.

4. Why the Ensemble Didn't Beat XGBoost

- The Insight: The Hybrid Ensemble dropped top-end accuracy slightly to 87.23%.
- The Reason: The Ensemble relied on Soft Probability Voting. The weaker confidence metrics of the deep learning branch (78% accuracy) dragged down XGBoost's sharp, decisive metrics.

-

Performance Interpretation

1. The Dominant Architecture:

Why XGBoost Won

- The Insight: XGBoost outperformed Temporal Deep Learning (LSTM+GRU) by 12% overall accuracy (90% vs. 78%).
- The Reason: Passive sensor tracking yields highly structured tabular data. Decision-tree models are mathematically superior at finding hard, crisp behavioral thresholds compared to neural networks.

2. The Clinical Priority: High-Risk Recall

- The Insight: XGBoost achieved an exceptional 95% Recall for the High-Risk class.
- The Reason: In digital health, missing an at-risk student (a False Negative) is a dangerous failure. A 95% recall guarantees the system catches 9.5 out of 10 actual high-risk addiction cases.

3. The Dilemma of the Medium-Risk Class

- The Insight: The Medium-Risk class consistently generated the lowest F1-scores across all architectures (0.68 to 0.84).
- The Reason: "Medium Risk" is a highly fluid, transitional behavioral zone. These students exhibit overlapping healthy and unhealthy habits, making them mathematically harder for algorithms to cleanly isolate.

4. Why the Ensemble Didn't Beat XGBoost

- The Insight: The Hybrid Ensemble dropped top-end accuracy slightly to 87.23%.
- The Reason: The Ensemble relied on Soft Probability Voting. The weaker confidence metrics of the deep learning branch (78% accuracy) dragged down XGBoost's sharp, decisive metrics.

-

Deployment at Plaksha University

Feasibility and Student-Centric Benefits

- **Core Feasibility:** Yes. Deployment directly aligns with Plaksha's emphasis on data science, tech innovation, and student wellness.

The Opportunity (Why We Should Do It)

- **Self-Awareness Hub:** Provides an objective, data-driven mirror for students to catch unhealthy habits early.
- **Proactive Well-being:** Shifts campus mental health support from reactive counseling to automated, early-stage behavioral insights.

Deployment Strategy (How to Do It)

- **Opt-In Pilot Cohort:** Launch as a voluntary, student-led research trial within a specific residential community on campus.
- **Privacy-First Edge Computing:** Process sensitive sensor logs locally on the student's phone, uploading only anonymous risk scores to ensure trust.

DEPLOYMENT: REAL-WORLD SCALING CHALLENGES

1. High-Dimensional Data Bottleneck

- The Challenge: Continuously gathering and syncing 85 distinct passive sensor features across thousands of devices causes severe network congestion.
- The Reality: This extreme feature density drains student phone batteries rapidly and consumes excessive network bandwidth.

2. Algorithmic Complexity at Scale

- The Challenge: Processing complex 7-day sliding window matrix operations through parallel LSTM and GRU branches demands heavy continuous memory.
- The Reality: While localized testing on 5,200 entries is seamless, scaling the live inference engine to track hundreds of active students will heavily strain campus server infrastructure.

3. Ethical & Permission Friction

- The Challenge: Continuous background tracking requires deep, invasive system-level application permissions on iOS and Android.
- The Reality: Overcoming user hesitation and establishing absolute institutional trust regarding personal data monitoring is a major non-technical barrier.

References

1.

Lee, S. H., Kim, S. M., & Cho, H. (2021). Prediction of problematic smartphone use: A machine learning approach. International Journal of Environmental Research and Public Health, 18(12), 6458.

2.

Kim, J. H., Lee, H. J., & Park, S. H. (2024). Explainable prediction of problematic smartphone use among South Korea's children and adolescents using machine learning. International Journal of Medical Informatics, 188, 105474.

3.

Wang, R., Chen, F., Chen, Z., Li, T., Harari, G., Tignor, S., Zhou, X., Ben-Zeev, D., & Campbell, A. T. (2014). StudentLife: Assessing mental health, academic performance and behavioral trends of college students using smartphones. Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp), 3–14.

Thank you

