An AI-Driven Early Warning System for Wellness Risk Detection Using Wearable and Self-Reported Data

Sellappan Palaniappan[1], Rajasvaran Logeswaran[2*] and Anitha Velayutham[3]

[1,2,3]Faculty of Computing and Digital Technology, HELP University, Kuala Lumpur, Malaysia
*Loges@ieee.org

Received: 20 December 2025 Revised: 22 January 2026 Accepted: 27 January 2026

Abstract

Wearable sensors and digital self-reporting tools enable continuous monitoring of individual wellness. This is particularly useful for remote monitoring of persons living alone, an increasingly common scenario in the modern world. However, transforming heterogeneous, high-frequency data into actionable insights remains a major challenge. This study proposes an artificial intelligence (AI)-driven early warning system for wellness risk detection, integrating wearable-derived physiological signals with self-reported behavioral and psychological indicators. Trained using supervised machine learning models, the proposed system identifies individuals at elevated wellness risk and provides probabilistic early warnings for preventive intervention. A simulation-based experimental framework was in this work to evaluate the model performance, interpretability and operational feasibility. Logistic Regression and Random Forest classifiers were compared using standard evaluation metrics, including accuracy, precision, recall, F1-score and ROC-AUC. The obtained results demonstrate strong predictive performance, with ROC-AUC values exceeding 0.92. The results highlight fatigue, stress and sleep duration as key contributors to early warning signals. Such AI-driven systems are feasible for wellness monitoring of the individual as well as the larger population, maintaining self-empowerment and reduced hospital observations stays, while emphasizing personal health management, interpretability, ethical considerations and real-world deployability.

Keywords- Artificial intelligence, Early warning systems, Wellness analytics, Wearable data, Machine learning, Preventive healthcare

1. Introduction

Healthcare systems worldwide are increasingly shifting from reactive and treatment-centered, to a preventive and wellness-oriented approach. This transition is driven by rising chronic disease prevalence, aging populations and escalating healthcare costs. Wellness risk detection, i.e. identifying early signs of physical or psychological deterioration before clinical manifestation, plays a critical role in enabling timely intervention and personalized care. Advances in artificial intelligence (AI), wearable sensors, as well as digital platforms, now enable operationalization of such early warning systems at scale.

Wearable devices are currently being used to capture physiological signals, such as heart rate, physical activity and sleep patterns. Often connected to mobile applications, they allow real-time self-reporting of physical health. Additionally, with the power of AI-enabled platforms, information on psychological and subjective well-being may also be obtained via health surveys, quizzes and even analysis of social media posts. When combined, these data streams provide a rich profile of an individual's wellness. However, the complexity, heterogeneity and volume of such data pose significant analytical challenges. AI-driven technology can overcome the challenges faced by traditional statistical approaches that often struggle to capture the nonlinear interactions and temporal patterns inherent in wellness data.

Early warning systems (EWS) that have long been used for disaster management, epidemiological surveillance and financial risk detection, have recently gain in terms of improved prediction accuracy, lead time, and adaptability through the incorporation of AI technology [1][2]. In healthcare, AI-enabled EWS have been successfully applied to patient deterioration detection, infectious disease surveillance and clinical decision support [3]-[5]. However, there has been exploration of wellness-focused EWS that operate outside acute clinical settings.

This study addresses this gap by proposing an AI-driven EWS designed specifically for wellness risk detection using wearable and self-reported data. Unlike disease-centric models, the proposed system emphasizes early, subclinical indicators of risk, and prioritizes interpretability and usability for preventive care. The objectives of this study are threefold: (1) to design a scalable AI-based framework for wellness risk detection, (2) to evaluate its predictive performance using simulated realistic data, and (3) to identify key wellness indicators that drive early warning signals.

The remainder of this paper is organized as follows. Section 2 reviews related work and methodological foundations for this study. Section 3 presents the proposed methodological framework, with the simulation study design given in Section 4 and results reported in Section 5. Section 6 discusses the implications, limitations and ethical considerations of the study, with the conclusions and recommendations for future direction given in Section 7. Some pseudocode of parts of the implementation are provided in Appendix A as a guide to more indepth exploration..

2. Related Work and Background

2.1 Early Warning Systems (EWS) and Artificial Intelligence (AI)

EWS are designed to detect emerging risks at an early stage, enabling timely preventive or mitigating actions before adverse outcomes materialize. Traditionally, EWS have relied on rule-based thresholds, expert-defined heuristics, or classical statistical models. While effective in well-structured and stable environments, these approaches often struggle in complex, dynamic settings characterized by noisy, high-dimensional, and nonlinear data.

The integration of AI has significantly expanded the capabilities of modern EWS. Machine learning (ML), deep learning and data-driven analytics, allow systems to learn latent patterns, temporal dependencies and nonlinear interactions that are difficult to capture using conventional methods [6]. Recent systematic reviews and conceptual frameworks highlight the success of AI-enabled EWS in forecasting accuracy, adaptability and real-time responsiveness [7]. These studies underscore the ability of AI models to integrate heterogeneous data sources (e.g. sensor streams, behavioral indicators and contextual information) into unified predictive frameworks.

Past literature also consistently identify model interpretability, data quality, robustness and ethical deployment as key challenges [8]. Addressing these is especially critical in human-centered applications, where trust, transparency and accountability are essential for adoption

2.2 AI in Healthcare and Wellness Monitoring

AI applications in healthcare have expanded rapidly over the past decade, with significant advances in predictive modeling, risk stratification and clinical decision support systems. AI-driven EWS are increasingly used to detect early signs of patient deterioration, adverse events and disease progression, by analyzing electronic health records (EHRs), physiological signals and wearable sensor data [3] [9].

In wellness monitoring, AI research has mainly focused on detecting pre-diagnostic risk states rather than overt disease outcomes. Studies in this area commonly integrate behavioral, physiological and subjective indicators (e.g. physical activity, sleep patterns, heart rate metrics, perceived stress and self-reported fatigue) to assess overall wellness risk [10], [11]. This complex multidimensional data is required to cater for the complex and holistic nature of wellness, which cannot be adequately captured by clinical biomarkers alone.

Wearable-based analytics are better able to identify early deviations, enabling detection of stress, fatigue and sleep disturbances, which may be indicators to more serious health conditions. However, wearable data alone often lacks contextual information related to psychological or social factors. Self-reported measures, despite their subjective nature, provide

 © IJETCEAI | www.ijetceai.com

valuable insights into perceived well-being and symptom burden. The integration of wearable-derived and self-reported data has been shown to enhance predictive performance, though it introduces additional methodological challenges related to feature engineering, data fusion and interpretability [12].

2.3 Supervised Learning for AI-Driven Early Warning Systems

Supervised machine learning models are widely used in early warning applications due to their flexibility and ability to incorporate labeled outcome data. Models such as Logistic Regression and Random Forests remain popular in healthcare and wellness analytics, because they provide a good balance of predictive performance, computational efficiency and interpretability [13]. Logistic Regression provides transparent probabilistic outputs, while ensemble methods like Random Forests are well-suited for handling heterogeneous inputs and capturing nonlinear interactions among features [14].

Model evaluation for EWS includes performance metrics such as overall accuracy, sensitivity (recall), false-negative rates and receiver operating characteristic area under the curve (ROC–AUC) [15]. The combination of metrics overcomes weaknesses associated with single metrics, such as prediction errors (e.g. false negatives). In the wellness context, high sensitivity is often prioritized over high specificity, to ensure that at-risk individuals are identified early, even at the expense of increased false positives.

Furthermore, transparent reporting of model behavior, including feature importance and uncertainty, is essential [16]. Explainability and interpretability [17]-[19] are increasingly recognized as prerequisites for ethical and responsible AI deployment [20]-[22], particularly for systems that influence health-related decisions. This is also seen in other sensitive decision making roles (e.g. a banker needs to be able to explain to the customer exactly why the customer is not eligible for a loan).

3. Methodology

3.1 System Overview

The proposed AI-driven EWS is designed to support proactive wellness risk detection by integrating heterogeneous data streams and transforming them into actionable risk signals. The framework consists of four interconnected components: (1) data acquisition, (2) feature processing, (3) predictive modelling, and (4) risk communication. These components form a modular pipeline that can be adapted to different wellness monitoring contexts and data availability constraints. Figure 1 shows a schematic overview of the proposed system architecture, illustrating the interaction between these components.

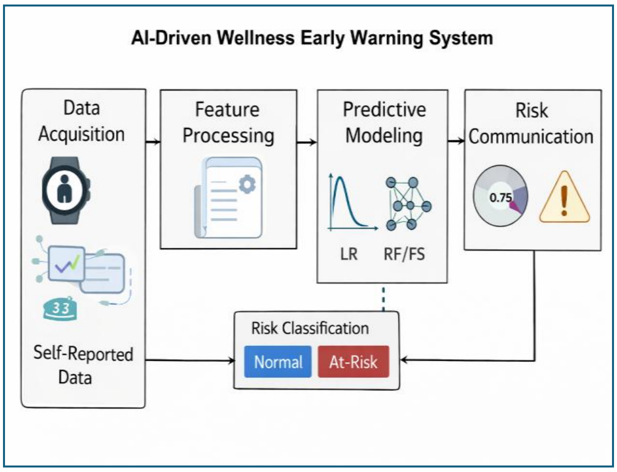

Figure 1. System architecture of the AI-driven wellness EWS

Data acquisition involves the collection of wearable-derived physiological measurements, and self-reported wellness indicators. Wearable devices continuously capture objective signals, such as physical activity, sleep duration and resting heart rate, while self-reported inputs provide subjective assessments of stress and fatigue. These complementary data sources enable both continuous monitoring and contextual interpretation of individual wellness states.

Feature processing transforms raw sensor streams and questionnaire responses into structured, analysis-ready feature vectors. Temporal aggregation is applied to summarize high-frequency wearable data over fixed observation windows (e.g., daily or weekly averages), reducing noise and improving model stability. The resulting feature vectors serve as standardized inputs to supervised machine learning models.

Predictive modeling constitutes the core analytical component of the system. Supervised classifiers are trained to estimate the probability that an individual belongs to an elevated wellness risk category. These probabilistic outputs are subsequently translated into EWS signals through predefined decision thresholds.

Finally, risk communication focuses on delivering interpretable outputs, such as risk probabilities and binary warning flags, to end users or downstream systems. This component is essential for enabling timely intervention and supports integration with digital health dashboards or mobile wellness applications.

3.2 Input Features

The EWS incorporates a multi-modal feature set encompassing demographic, behavioral, physiological, and self-reported indicators of wellness. By integrating these diverse features into a unified representation, the proposed framework captures the multidimensional nature of wellness and aligns with contemporary approaches in AI-driven health analytics. Examples of the different levels of information that are analysed and utilized in the proposed framework based on the mentioned features are as follows:

- Demographic information, such as age, provides contextual baseline characteristics that may influence wellness trajectories.
- Behavioral features, including average daily step count, capture patterns of physical activity that are strongly associated with overall health and lifestyle behaviors.
- Physiological features derived from wearable sensors include sleep duration and resting heart rate. Sleep duration serves as a proxy for recovery and circadian regulation, while resting heart rate reflects cardiovascular and autonomic nervous system functioning. Deviations in these measures have been linked to stress, fatigue, and emerging health risks.
- Self-reported measures, specifically stress score and perceived fatigue, provide subjective insights that are not directly observable through wearable sensors. Although inherently subjective, these measures capture psychological and emotional dimensions of wellness that are critical for early risk detection. Prior studies have shown that combining subjective and objective indicators improves predictive performance and robustness in wellness and health monitoring systems [10] [12].

3.3 Predictive Models

Two supervised learning models are evaluated within the proposed framework, namely: (1) Logistic Regression, and (2) Random Forest [23]. Logistic Regression is employed as a baseline model due to its simplicity, interpretability, and ability to produce well-calibrated probabilistic outputs. Its linear decision boundary provides a transparent reference for assessing model behavior and feature influence.

The Random Forest model, an ensemble-based approach, is used to capture nonlinear relationships and complex interactions among input features. By aggregating predictions from multiple decision trees, Random Forest reduces variance and improve generalization performance, particularly in heterogeneous and noisy datasets. Additionally, Random Forest provides intrinsic feature importance measures, supporting model interpretability and explanatory analysis.

The primary estimand of interest in this study is the individual-level probability of wellness risk, defined as the posterior probability that an individual belongs to the at-risk class, given their observed features. This probabilistic formulation allows flexible thresholding for early warning generation and supports risk stratification rather than binary classification alone.

Model training, validation, and inference are conducted using a reproducible Python-based workflow. Detailed implementation steps, including data simulation, feature preparation, model fitting, evaluation, and visualization procedures, are provided in Appendix A.

4. Simulation Study Design

4.1 Data Generation

Due to ethical, privacy, and regulatory constraints associated with real-world wellness and health data, this study adopts a simulation-based experimental design. Synthetic datasets are generated to approximate realistic distributions observed in wearable-derived physiological signals and self-reported wellness indicators, while avoiding the use of identifiable personal information.

The simulated dataset represents a heterogeneous population with varying demographic, behavioral, and physiological characteristics. The features employed include age, average daily step count, sleep duration, resting heart rate, perceived stress score, and self-reported fatigue. These variables are sampled from distributions informed by existing literature on wearable analytics and wellness monitoring to reflect plausible ranges and inter-individual variability.

Binary class labels are assigned to indicate normal or at-risk wellness states. The at-risk class represents individuals exhibiting combinations of elevated stress, reduced physical activity, abnormal sleep patterns, or increased fatigue. This formulation reflects the objective of early warning systems, which aim to identify emerging risk patterns before overt clinical deterioration occurs.

By using a controlled simulation environment, the study enables systematic evaluation of model behavior, interpretability, and robustness under known data-generating assumptions. This allows for better analysis of the factors under examination, with minimal influence of noisy data.

4.2 Experimental Setup

The simulated dataset is partitioned into training and testing subsets using stratified sampling to preserve class proportions across splits. This ensures balanced model evaluation and reduces bias associated with class imbalance, which is common in early warning and risk detection scenarios.

The two supervised learning models (Logistic Regression and Random Forest) are trained on the same feature set for comparative evaluation. Hyperparameters are selected using standard defaults to emphasize methodological transparency and reproducibility. Model training and inference are implemented using Python-based machine learning libraries, as detailed in Appendix A.

Model performance is evaluated using multiple complementary metrics to capture both overall predictive accuracy and sensitivity to at-risk cases. In early warning contexts, minimizing false negatives is critical, as missed detections may delay intervention. Accordingly, recall and ROC-AUC are emphasized alongside accuracy and precision. Table 1 provides a summary of the evaluation metrics employed in the study and their interpretive roles.

Table 1. Model Performance Metrics

Metric	Description
Accuracy	Overall correctness of classification
Precision	Proportion of predicted at-risk cases that are true positives
Recall	Sensitivity to at-risk cases
F1-score	Harmonic mean of precision and recall
ROC–AUC	Discriminative ability across classification thresholds

4.3 Output Generation

The AI-driven EWS produces two primary outputs for each individual: (1) a probabilistic wellness risk score, and a corresponding (2) binary early warning flag. The risk score represents the estimated probability of belonging to the at-risk class, while the binary flag is generated by applying a predefined decision threshold selected to balance sensitivity and specificity.

In addition to numerical outputs, the system generates a set of visualization artifacts to support model evaluation and interpretability. These include wellness risk label distributions, confusion matrices summarizing classification outcomes, receiver operating characteristic (ROC) curves illustrating discrimination performance, and feature importance plots highlighting the relative contribution of input variables to model predictions. These outputs form the basis of the empirical evaluation in the next Section.

5. Results and Discussion

5.1 Distribution of Wellness Risk Labels

Figure 2 shows the distribution of wellness risk labels in the simulated dataset. The dataset exhibits a moderate class imbalance, with a larger proportion of individuals classified as normal compared to at-risk. This imbalance reflects realistic wellness monitoring scenarios, where adverse outcomes typically occur less frequently than normal states.

Such class distributions pose challenges for predictive modeling, particularly in early warning contexts where minority-class detection is critical. The presence of imbalance further motivates the use of evaluation metrics beyond accuracy, such as recall and ROC–AUC, to ensure adequate sensitivity to at-risk individuals.

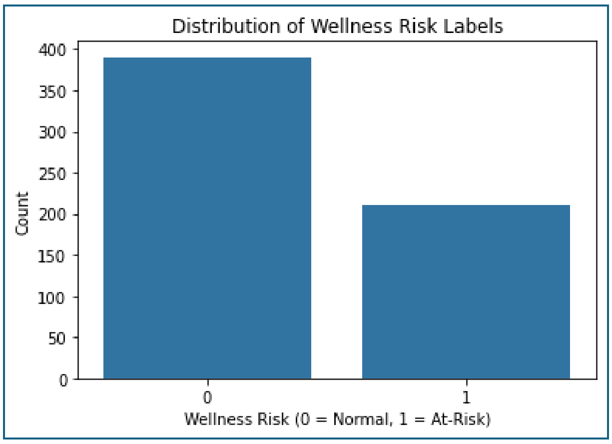

Figure 2. Distribution of wellness risk labels (0 = Normal, 1 = At-risk)

5.2 Model Performance Comparison

Table 2 summarizes the predictive performance of Logistic Regression and Random Forest models. Both models demonstrate strong classification capability across multiple evaluation metrics.

Table 2. Model Performance Comparison

Model	Accuracy	Precision	Recall	F1-score	ROC–AUC
Logistic Regression	**0.856**	0.761	**0.857**	**0.806**	**0.945**
Random Forest	0.839	**0.793**	0.730	0.760	0.929

Logistic Regression achieves the highest ROC–AUC (0.945), indicating superior discriminative ability across decision thresholds. Its high recall value suggests strong sensitivity to at-risk cases, which is particularly desirable in early warning applications. The Random Forest model, while slightly lower in overall ROC–AUC, demonstrates higher precision, indicating fewer false positives when predicting wellness risk.

These results suggest a trade-off between sensitivity and specificity across model types, highlighting the importance of model selection based on deployment context and risk tolerance.

5.3 Confusion Matrix and Classification Analysis

Figure 3 provides further analysis of classification behavior via the confusion matrix for the Random Forest model. The matrix shows a relatively low number of false negatives, indicating that the model successfully identifies a substantial proportion of at-risk individuals.

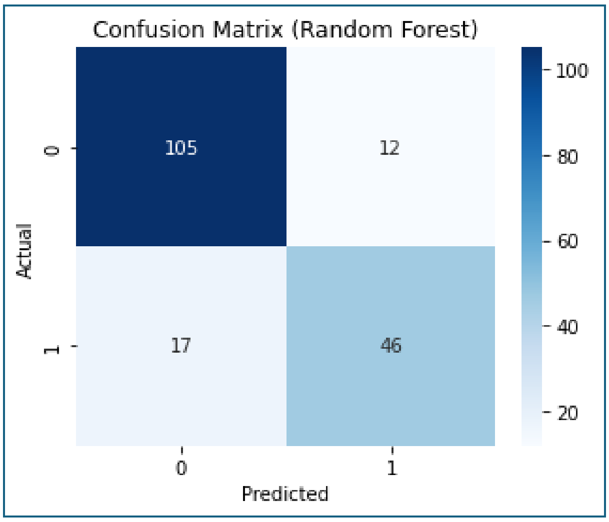

Figure 3. Confusion matrix for Random Forest classifier

False positives are present but remain within acceptable bounds for a wellness EWS, where precautionary alerts are often preferable to missed detections. This aligns with the system's preventive objective, prioritizing early identification over strict specificity.

5.4 ROC Curve Analysis

Figure 4 shows the receiver operating characteristic (ROC) curve for the Random Forest model.

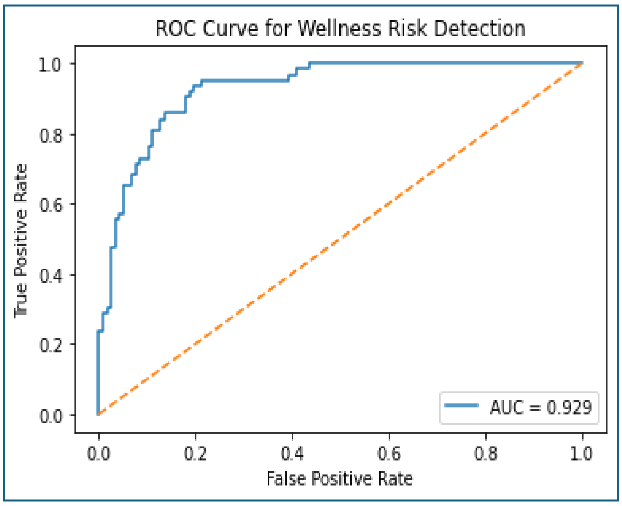

Figure 4. ROC curve for wellness risk detection

The curve demonstrates strong separation between true positive and false positive rates across thresholds, with an AUC value of 0.929. The ROC analysis confirms that the model maintains robust discriminative performance even under varying decision thresholds. This characteristic is particularly valuable for adaptive EWS, where alert thresholds may be tuned based on operational constraints or individual risk profiles.

5.5 Feature Importance Analysis

Table 3 summarizes the feature importance scores derived from the Random Forest model and is visually represented in Figure 5. Self-reported fatigue emerges as the most influential predictor, followed by the stress score and sleep duration.

Table 3. Feature Importance (Random Forest)

Feature	Importance
Self-reported fatigue	0.308
Stress score	0.230
Sleep hours	0.154
Average steps	0.154
Resting heart rate	0.106
Age	0.047

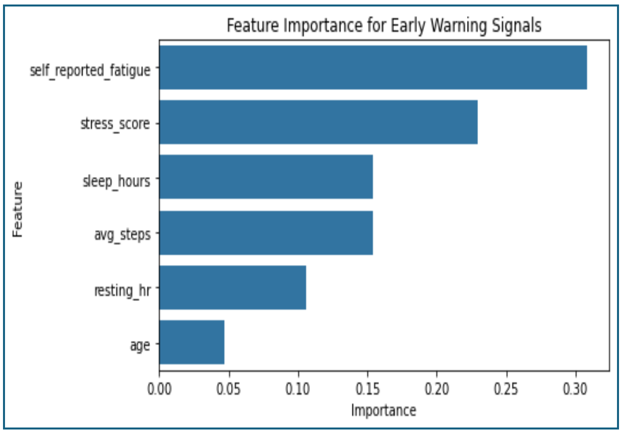

Figure 5. Feature importance for early warning signals

Behavioral and subjective measures collectively account for a substantial proportion of predictive importance, underscoring the value of integrating self-reported data with wearable-derived metrics. Demographic variables such as age contribute comparatively less, suggesting that short-term behavioral and physiological signals are more informative for early wellness risk detection.

6. Conclusion

The findings of this study demonstrate the feasibility and effectiveness of a proposed AI-driven EWS for wellness risk detection, that integrates wearable-derived and self-reported data. Both evaluated models exhibit strong discriminative performance, as evidenced by high ROC–AUC values, indicating their ability to distinguish between normal and at-risk wellness states across a range of decision thresholds. These results support the growing body of evidence that machine learning approaches can meaningfully enhance early risk detection in preventive health contexts.

The feature importance analysis provides valuable interpretability, revealing that self-reported fatigue and perceived stress are the most influential predictors of wellness risk. This observation aligns with established wellness and behavioral health research, which consistently identifies fatigue and stress as early indicators of declining well-being [24]. The prominence of subjective self-reports underscores their complementary role alongside physiological and behavioral signals, offering contextual insights that may not be fully captured by sensor data alone. Rather than diminishing objectivity, the integration of self-reported measures enhances the system's holistic understanding of wellness dynamics.

Importantly, the proposed system is designed for early warning rather than clinical diagnosis. Its outputs, i.e. probabilistic risk scores and binary warning flags, are intended to support awareness, self-regulation, and timely preventive action, rather than to replace professional medical judgment. This positioning is critical for responsible deployment, particularly in wellness and population health settings, where low-cost, scalable monitoring solutions are needed.

Despite these promising results, several limitations warrant consideration. First, the use of simulated data, while necessary due to privacy and ethical constraints, may not fully capture the complexity, noise, and behavioral variability present in real-world wearable datasets. Second, the current framework relies on static feature aggregation, which may obscure temporal patterns and short-term fluctuations that are informative for early risk detection. Incorporating longitudinal and sequential modeling approaches, such as recurrent neural networks or temporal probabilistic models, represents a natural direction for future research.

Additionally, practical deployment raises important ethical and governance considerations. Issues related to data privacy, informed consent, algorithmic bias, and transparency must be carefully addressed to ensure trust and equitable use. Providing explainable outputs, clear communication of uncertainty, and user control over data sharing will be essential for real-world adoption. Future work should also explore adaptive alert thresholds tailored to individual risk profiles and contextual factors, further enhancing the system's responsiveness and usability.

References

[1] T. Tiggeloven, S. Pfeiffer, A. Matanó, M. van den Homberg, L. Thalheimer, M. Reichstein, and S. Torresan, "The role of artificial intelligence for early warning systems: Status, applicability, guardrails, and ways forward," iScience, vol. 28, no. 11, Art. no. 113689, 2025, doi: 10.1016/j.isci.2025.113689.

[2] M. Cate, "AI-driven early warning systems: Design and implementation," Apr. 2022.

[3] I. Villanueva-Miranda, G. Xiao, and Y. Xie, "Artificial intelligence in early warning systems for infectious disease surveillance: A systematic review," Frontiers in Public Health, vol. 13, Art. no. 1609615, 2025, doi: 10.3389/fpubh.2025.1609615.

[4] H. Arsalan, "Enhancing predictive healthcare using AI-driven early warning systems," in Proc. AAAI Conf. Artificial Intelligence, vol. 39, no. 28, 2025, doi: 10.1609/aaai.v39i28.35326.

[5] R. Lamsal and T. V. V. Kumar, "Artificial intelligence and early warning systems," in AI and Robotics in Disaster Studies, Singapore: Springer, 2020, pp. 25–44, doi: 10.1007/978-981-15-4291-6_2.

[6] L. Karunakaran, "AI-driven disaster prediction and early warning systems: A systematic literature review," International Journal of Research in Computing, vol. 4, no. 2, pp. 44–55, 2025.

[7] S. Huang, "AI-driven early warning systems for supply chain risk detection: A machine learning approach," Academic Journal of Computing & Information Science, vol. 8, no. 9, pp. 92–107, 2025, doi: 10.25236/AJCIS.2025.080913.

[8] M. Wever, M. Shah, and N. O'Leary, "Designing early warning systems for detecting systemic risk: A case study and discussion," Futures, vol. 136, Art. no. 102882, 2022, doi: 10.1016/j.futures.2021.102882.

[9] O. H. Embarak and S. Hawarna, "Automated AI-driven system for early detection of at-risk students," Procedia Computer Science, vol. 231, pp. 151–160, 2024, doi: 10.1016/j.procs.2023.12.187.

[10] L. Poorani, A. Malik, J. Rani, M. Varsha, and R. Ravi, "AI for disaster response: Enhancing early warning systems and recovery efforts," International Journal of Engineering Technology and Sciences, vol. 11, no. 9, pp. 74–83, 2024.

[11] N. S. Nandini, "AI-based early warning systems for climate change and extreme weather events," International Journal of Scientific Research & Engineering Trends, vol. 11, no. 2, pp. 1–8, 2025.

[12] P. Charan, M. M. Siddiqui, V. Yadav, C. Pathak, Y. Narayan, and Z. H. Khan, "AI-enhanced early warning systems for natural disaster detection and mitigation using wireless sensor networks," in Proc. 2nd Int. Conf. Computational and Characterization Techniques in Engineering & Sciences (IC3TES), Lucknow, India, 2024, pp. 1–6, doi: 10.1109/IC3TES62412.2024.10877562.

[13] S. Pfeiffer, M. Dewi, K. Dall, M. Reichstein, M. Usamah, S. E. Werners, H. Shekhar, and M. Hagenlocher, "Strengthening equitable, impact-based early warning through artificial intelligence: Four key perspectives," PreventionWeb, Aug. 2024, https://www.preventionweb.net/drr-community-voices/strengthening-equitable-impact-based-early-warning-through-artificial.

[14] G. Xiao, I. Villanueva-Miranda, and Y. Xie, "Machine learning and deep learning methods for early outbreak detection," Health Data Science, vol. 3, no. 1, pp. 1–14, 2024.

[15] A. Rajkomar, J. Dean, and I. Kohane, "Machine learning in medicine," New England Journal of Medicine, vol. 380, no. 14, pp. 1347–1358, 2019.

[16] E. J. Topol, "High-performance medicine: The convergence of human and artificial intelligence," Nature Medicine, vol. 25, pp. 44–56, 2019.

[17] F. Doshi-Velez and B. Kim, "Towards a rigorous science of interpretable machine learning," arXiv:1702.08608, 2017.

[18] J. L. Hernández-Orallo et al., "AI evaluation: Past, present and future," Artificial Intelligence, vol. 279, Art. no. 103215, 2020.

[19] S. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Proc. NIPS, 2017, pp. 4765–4774.

[21] World Health Organization, Ethics and Governance of Artificial Intelligence for Health, Geneva, Switzerland, 2021.

[22] European Commission, Ethics Guidelines for Trustworthy AI, Brussels, Belgium, 2019.

[23] S. K. Saha, “A Comparative Analysis of Logistic Regression and Random Forest for Individual Fairness in Machine Learning”, International Journal of Advanced Engineering Research and Science (IJAERS), vol. 12, no. 5, pp. 33-37, doi: 10.22161/ijaers.125.5.

[24] S. Attrey, S. Joshi, and V. Bhatt, “Mindful path to psychological wellbeing: a comprehensive review”, Journal of Behavioral and Cognitive Therapy, vol. 35, no. 4, p. 100537, doi: 10.1016/j.jbct.2025.100537.

Appendix: Implementation Details

This appendix provides the full Python implementation of the AI-driven EWS for wellness risk detection described in this study. The code covers data simulation, feature preparation, supervised model training and evaluation, generation of early warning outputs, and visualization of key results. Standard scientific computing and machine learning libraries are used to ensure reproducibility and accessibility. The implementation is intended to support transparency, facilitate replication of the reported findings, and enable future extensions using real-world wearable and self-reported wellness data.

```
# ------------------------------
# 0. Imports
# ------------------------------
import numpy as np
import pandas as pd
import matplotlib.pyplot as plt
import seaborn as sns

from sklearn.model_selection import train_test_split
from sklearn.preprocessing import StandardScaler
from sklearn.metrics import (
    accuracy_score,
    precision_score,
    recall_score,
    f1_score,
    roc_auc_score,
    confusion_matrix,
    roc_curve
)
from sklearn.linear_model import LogisticRegression
from sklearn.ensemble import RandomForestClassifier

np.random.seed(42)

# ------------------------------
# 1. Simulate Wellness Dataset
# ------------------------------
n_users = 600

data = pd.DataFrame({
    "age": np.random.randint(18, 65, n_users),
    "avg_steps": np.random.normal(7000, 2000, n_users),
    "sleep_hours": np.random.normal(6.5, 1.2, n_users),
    "resting_hr": np.random.normal(72, 8, n_users),
    "stress_score": np.random.uniform(1, 10, n_users),
    "self_reported_fatigue": np.random.uniform(1, 10, n_users)
})

# Latent wellness risk score
risk_latent = (
    0.04 * (75 - data["avg_steps"] / 100) +
    0.6 * (7 - data["sleep_hours"]) +
    0.05 * (data["resting_hr"] - 70) +
    0.4 * data["stress_score"] +
    0.5 * data["self_reported_fatigue"] +
    np.random.normal(0, 1, n_users)
)

# Binary early warning label (1 = at-risk)
data["wellness_risk"] = (risk_latent > np.percentile(risk_latent,
65)).astype(int)

# ------------------------------
# 2. Train-Test Split & Scaling
# ------------------------------
X = data.drop(columns=["wellness_risk"])
y = data["wellness_risk"]

X_train, X_test, y_train, y_test = train_test_split(
    X, y, test_size=0.3, stratify=y, random_state=42
)

scaler = StandardScaler()
X_train_scaled = scaler.fit_transform(X_train)
X_test_scaled = scaler.transform(X_test)

# ------------------------------
# 3. Models
# ------------------------------
log_model = LogisticRegression(max_iter=1000)
rf_model = RandomForestClassifier(
    n_estimators=200,
    max_depth=6,
    random_state=42
)

log_model.fit(X_train_scaled, y_train)
rf_model.fit(X_train, y_train)

# ------------------------------
# 4. Predictions & Metrics
# ------------------------------
def evaluate(model, X_eval, y_eval, prob=True):
    y_pred = model.predict(X_eval)
    y_prob = model.predict_proba(X_eval)[:, 1] if prob else None

    return {
        "Accuracy": accuracy_score(y_eval, y_pred),
        "Precision": precision_score(y_eval, y_pred),
        "Recall": recall_score(y_eval, y_pred),
        "F1-score": f1_score(y_eval, y_pred),
        "ROC-AUC": roc_auc_score(y_eval, y_prob) if prob else np.nan
    }

log_metrics = evaluate(log_model, X_test_scaled, y_test)
rf_metrics = evaluate(rf_model, X_test, y_test)

metrics_table = pd.DataFrame([log_metrics, rf_metrics],
                             index=["Logistic Regression", "Random Forest"])

print("\n=== Table 1: Model Performance Metrics ===")
print(metrics_table.round(3))

# ------------------------------
```

```
# 5. Confusion Matrix (Best Model)
# ------------------------------
best_model = rf_model
y_pred_best = best_model.predict(X_test)

cm = confusion_matrix(y_test, y_pred_best)

# ------------------------------
# 6. Feature Importance
# ------------------------------
importance = pd.DataFrame({
    "Feature": X.columns,
    "Importance": rf_model.feature_importances_
}).sort_values(by="Importance", ascending=False)

print("\n=== Table 2: Feature Importance (Random Forest) ===")
print(importance)

# ------------------------------
# 7. FIGURES
# ------------------------------

# Figure 1: Class Distribution
plt.figure(figsize=(6,4))
sns.countplot(x=data["wellness_risk"])
plt.title("Distribution of Wellness Risk Labels")
plt.xlabel("Wellness Risk (0 = Normal, 1 = At-Risk)")
plt.ylabel("Count")
plt.show()

# Figure 2: Confusion Matrix
plt.figure(figsize=(5,4))
sns.heatmap(cm, annot=True, fmt="d", cmap="Blues")
plt.title("Confusion Matrix (Random Forest)")
plt.xlabel("Predicted")
plt.ylabel("Actual")
plt.show()

# Figure 3: ROC Curve
y_prob = best_model.predict_proba(X_test)[:, 1]
fpr, tpr, _ = roc_curve(y_test, y_prob)

plt.figure(figsize=(6,4))
plt.plot(fpr, tpr, label=f"AUC = {roc_auc_score(y_test, y_prob):.3f}")
plt.plot([0,1], [0,1], linestyle="--")
plt.title("ROC Curve for Wellness Risk Detection")
plt.xlabel("False Positive Rate")
plt.ylabel("True Positive Rate")
plt.legend()
plt.show()

# Figure 4: Feature Importance
plt.figure(figsize=(7,4))
sns.barplot(data=importance, x="Importance", y="Feature")
plt.title("Feature Importance for Early Warning Signals")
plt.show()

# ------------------------------
# 8. Early Warning Example Output
# ------------------------------
sample_users = X_test.iloc[:5]
sample_probs = best_model.predict_proba(sample_users)[:, 1]

early_warning_table = sample_users.copy()
early_warning_table["Risk_Probability"] = sample_probs
early_warning_table["Early_Warning"] = (sample_probs >
0.7).astype(int)

print("\n=== Table 3: Sample Early Warning Outputs ===")
print(early_warning_table.round(2))
```