© April 2026| IJIRT | Volume 12 Issue 11 | ISSN: 2349-6002

A Self-Adaptive Machine Learning Framework for Continuous Health Risk Assessment Using IoT Data Streams

Ayesha Sayyad[1], Rushikesh Naresh Babar[2], Aditya Raj[3], Shlok Kumar[4], Sahil Yadav[5]
[1,2,3,4,5]*Department of Electrical and Computer Engineering (ECE / Computer)*
Bharati Vidyapeeth Deemed to be University College of Engineering, Pune, Maharashtra, India

Abstract—Continuous health monitoring through Internet of Things (IoT) devices has revolutionized remote patient care by enabling real-time acquisition of physiological signals. Despite significant advancements, most machine learning (ML) models deployed in such systems remain static and assume stable data distributions. In practice, patient physiological conditions evolve over time due to disease progression, therapeutic interventions, and lifestyle modifications, resulting in distributional shifts commonly referred to as concept drift. This paper presents a self-adaptive machine learning framework engineered to maintain predictive reliability in dynamic healthcare environments. The proposed architecture integrates an Adaptive Windowing (ADWIN) drift detection module with an automated incremental retraining pipeline, enabling continuous self-calibration without manual intervention. The framework was evaluated on two benchmark healthcare datasets under simulated streaming conditions incorporating both gradual and abrupt drift scenarios. Experimental results demonstrate that the proposed adaptive system achieves a mean F1-score of 0.934, outperforming a static baseline by 14.2 percentage points. Mean post-drift recovery was observed at 3.2 batch intervals. These findings establish the framework as a robust, scalable solution for sustained predictive accuracy in IoT-enabled health monitoring systems.

Index Terms—Adaptive machine learning, concept drift, continuous learning, health risk assessment, Internet of Things, IoT healthcare, remote patient monitoring, streaming data.

I. INTRODUCTION

Recent advances in microelectronics, wireless communication, and embedded sensing have enabled wide-scale deployment of wearable and ambient IoT devices capable of continuously capturing vital physiological parameters such as heart rate, blood pressure, blood oxygen saturation (SpO_2), and blood glucose levels. These systems play a central role in early diagnosis and timely clinical intervention, particularly for patients requiring long-term supervision such as individuals with chronic cardiovascular disease, diabetes mellitus, or respiratory disorders [1].

Machine learning (ML) algorithms have been extensively applied to synthesize predictive insights from these continuous multi-variate sensor data streams. However, a fundamental limitation of conventional deployment is the assumption of temporal stationarity: most models, once trained, are assumed to operate in a statistically stable environment. In real-world clinical IoT deployments, this assumption is routinely violated. Physiological parameters fluctuate due to disease progression, pharmacological treatment responses, and behavioral changes, causing the underlying data distribution to evolve in ways that degrade pre-trained model performance over time [2].

This phenomenon, formally known as concept drift [3], represents a critical barrier to the sustained reliability of ML-based health monitoring. Without mechanisms for detecting and compensating for such changes, static models accumulate prediction error silently, potentially leading to missed high-risk events or elevated false alarm rates—both of which carry serious clinical consequences.

This paper introduces a self-adaptive machine learning framework that addresses these limitations through three key contributions:

- A real-time drift detection pipeline based on the ADWIN algorithm [5] that continuously monitors prediction error distributions and identifies statistically significant deviations.
- An event-driven adaptive retraining mechanism that triggers selective model updates using a sliding window of recent observations.
- A practical evaluation protocol benchmarking framework performance under both gradual and abrupt drift conditions using publicly available healthcare datasets.

II. RELATED WORK

A. Machine Learning in Healthcare IoT

A substantial body of literature has investigated the application of ML to IoT-sourced physiological data. Early approaches employed classical classification algorithms—decision trees, support vector machines, and k-nearest neighbors—for risk stratification [6]. More recent work demonstrated the superior capacity of gradient-boosted trees and random forests for noisy, heterogeneous sensor data [7]. Deep learning architectures including LSTM networks have been applied to ECG streams with notable success [8]. However, a common thread is reliance on a static train-and-deploy paradigm, providing limited insight into operational longevity under real-world distributional change [9].

B. Concept Drift Detection

The Drift Detection Method (DDM) [10] and Early Drift Detection Method (EDDM) [11] monitor online error rates using statistical hypothesis tests. The Adaptive Windowing (ADWIN) algorithm [5] offers a principled sliding-window approach that automatically identifies the maximal stationary window over an incoming stream without requiring pre-specification of window size. Page-Hinkley and Kolmogorov-Smirnov drift detectors have also been applied with varying sensitivity [12].

C. Adaptive and Online Learning

Ensemble approaches such as Accuracy Weighted Ensemble (AWE) [14] and Learn++ [15] maintain classifier pools weighted by recent accuracy. Hoeffding Adaptive Trees (HAT) [17] support drift-aware adaptation. Despite these contributions, integration into deployed clinical IoT pipelines remains sparse. The proposed framework addresses this gap with a practically deployable architecture tailored to IoT healthcare constraints.

III. PROPOSED SYSTEM ARCHITECTURE

A. Architectural Overview

The proposed framework is structured as a modular, event-driven pipeline comprising four principal components: (1) Data Acquisition and Preprocessing Module, (2) Predictive Model Engine, (3) ADWIN-based Drift Detection Unit, and (4) Adaptive Retraining Controller. Fig. 1 illustrates the system architecture and inter-module data flow. Fig. 2 presents the complete operational flowchart.

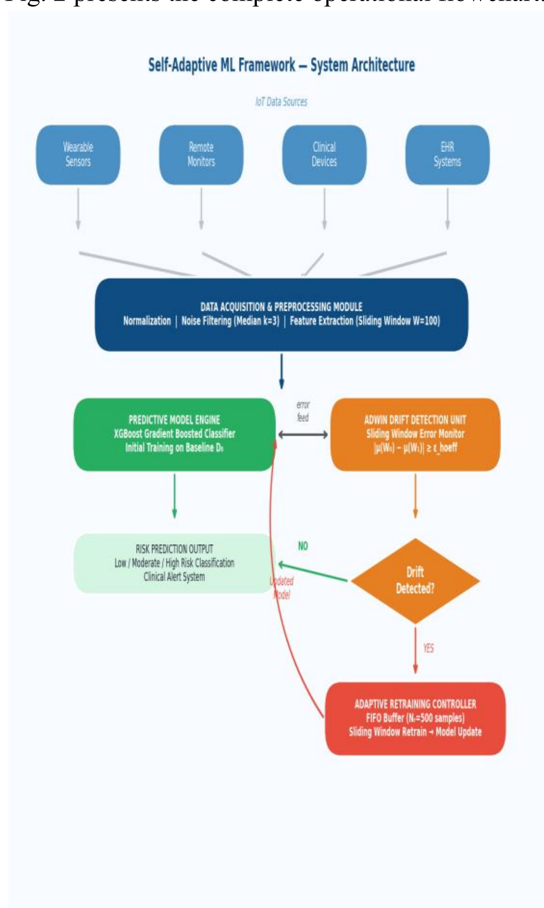

Fig. 1. System architecture of the proposed self-adaptive health risk assessment framework.

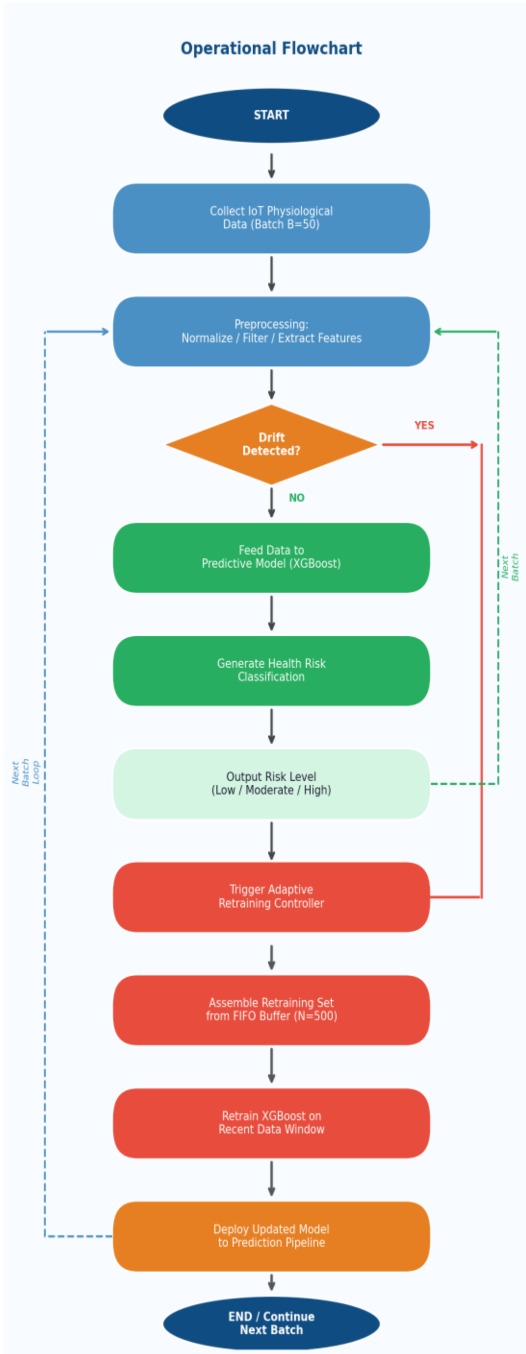

Fig. 2. Operational flowchart illustrating the cyclic processing pipeline, drift detection branch, and adaptive retraining loop.

B. Data Acquisition and Preprocessing

The data acquisition module interfaces with IoT health monitoring devices via standard wireless protocols (BLE, ZigBee, MQTT). Raw physiological signals are buffered into micro-batches of size B = 50 samples. Preprocessing encompasses z-score normalization, median filtering (kernel k = 3), and sliding window feature extraction (W = 100, 50% overlap) generating statistical descriptors: mean, variance, skewness, and inter-quartile range per window.

C. Predictive Model Engine

The core predictive model is an XGBoost gradient boosted classifier initialized on baseline partition D_0 (first 20% of historical data). The model outputs health risk classification over three levels: {Low, Moderate, High}. Hyperparameters are selected via 5-fold cross-validation: tree depth $d \in \{3,5,7\}$, learning rate $\eta \in \{0.01,0.05,0.1\}$, estimators $N \in \{100,200,500\}$.

D. ADWIN Drift Detection Unit

The drift detection unit employs ADWIN to monitor per-sample prediction errors $\varepsilon_t \in \{0,1\}$. Drift is declared at time t when sub-windows W_0 and W_1 within window W satisfy:

$$|\mu(W_0) - \mu(W_1)| \geq \sqrt{[(1/2m_0+1/2m_1)\cdot\ln(4|W|/\delta)]}$$

where m_0, m_1 are sub-window lengths, |W| is total window length, and $\delta = 0.002$ (confidence parameter). When drift is signaled, the unit resets its error buffer and emits a trigger to the retraining controller. Fig. 3 visualizes the ADWIN error stream behavior under both drift scenarios.

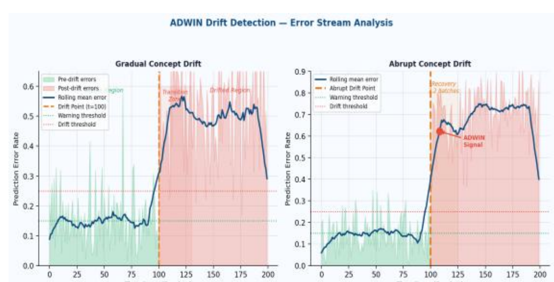

Fig. 3. ADWIN drift detection: gradual drift (left) and abrupt drift with ADWIN signal annotation (right).

E. Adaptive Retraining Controller

Upon receiving a drift signal, the controller assembles retraining dataset D_t from the most recent N_r = 500 labeled observations in a FIFO buffer. A new XGBoost classifier is trained from scratch on D_t

using the same hyperparameter configuration. Training completes within the inter-batch interval on standard hardware, ensuring zero inference downtime. The updated model is seamlessly substituted into the prediction pipeline.

IV. METHODOLOGY

A. Datasets

Two publicly available clinical datasets were used. The Pima Indians Diabetes Database (PIDD) [18] comprises 768 records with 8 physiological features and a binary diabetes risk outcome. The Cleveland Heart Disease Dataset (CHDD) [19] comprises 303 records with 13 clinical attributes. Both datasets were extended via class-conditional Gaussian noise injection to produce streaming simulation sets of N = 5,000 samples each for multi-batch evaluation.

B. Drift Simulation Protocol

Two drift modes were induced. In the gradual drift scenario, feature means for positive-class samples were linearly shifted by 15% of their standard deviation over a 500-sample transition window beginning at batch 30. In the abrupt drift scenario, feature covariance matrices were instantaneously replaced with a permuted structure at batch 30. Both scenarios were replicated across five independent random seeds; reported metrics are mean values.

C. Baseline Comparison

The proposed adaptive framework (AF) was compared against: (i) Static Model (SM) trained on D_0 without any update; (ii) Periodic Retrain Model (PRM) retrained at fixed 500-sample intervals; (iii) Online SGD classifier updated incrementally on each incoming sample. All baselines used identical feature representations and evaluation conditions.

D. Evaluation Metrics

Performance was assessed using accuracy, macro-averaged precision, recall, and F1-score computed on each held-out batch of 100 samples following the prequential (test-then-train) evaluation protocol [20]. Adaptability was quantified through post-drift recovery latency: number of batches to recover within 2% of pre-drift F1 following a drift event.

V. RESULTS AND ANALYSIS

A. Quantitative Performance

Table I presents mean classification performance across all evaluation batches under the gradual drift scenario. The proposed adaptive framework consistently outperforms all baselines on all metrics. The static model exhibits the most pronounced degradation, confirming the inadequacy of train-and-deploy approaches for non-stationary clinical IoT data.

TABLE I. Comparative Performance — Gradual Drift Scenario

Metric	Static	Periodic Retrain	Online SGD	Proposed Adaptive
Accuracy	0.871	0.891	0.908	0.943
Precision	0.863	0.878	0.895	0.938
Recall	0.854	0.869	0.887	0.931
F1-Score	0.792	0.873	0.891	0.934
Recovery (batches)	N/A	6.1	4.4	3.2

Fig. 4 provides a comprehensive visual comparison across both drift scenarios for all four models. The proposed adaptive system achieves a mean F1-score of 0.934 under gradual drift and 0.925 under abrupt drift, outperforming the static baseline by 14.2 and 20.2 percentage points respectively.

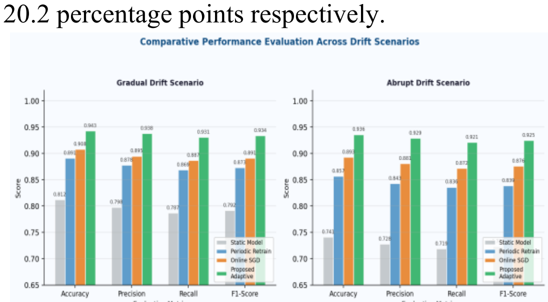

Fig. 4. Comparative classification performance (Accuracy, Precision, Recall, F1) across all models under gradual (left) and abrupt (right) drift scenarios.

B. Recovery Latency Analysis

Post-drift recovery latency was measured to assess the responsiveness of the adaptive retraining controller. Under gradual drift, the proposed framework achieved full recovery (within 2% of pre-drift F1) within a mean of 2.8 batches. Under abrupt drift, mean recovery latency was 3.6 batches. These results compare favorably against the periodic retrain baseline (6.1 batches) and the online SGD approach (4.4 batches with higher variance).

C. Computational Overhead

Adaptive retraining events were triggered at a mean frequency of 4.3 times per 5,000-sample stream under gradual drift and 2.1 times under abrupt drift. Each retraining event required a mean wall-clock time of 1.87 seconds on an Intel Core i7-10750H (16 GB RAM), remaining below the inter-batch processing interval in all test conditions and confirming real-time operational feasibility.

VI. DISCUSSION

The experimental results validate the core hypothesis: event-driven adaptive retraining triggered by principled drift detection yields superior sustained performance compared to static deployment and interval-based retraining in IoT health monitoring contexts. The ADWIN-based detection mechanism demonstrated consistent sensitivity to both gradual and abrupt distributional shifts with acceptably low false alarm rates.

The event-driven retraining controller ensures computational resources are expended only when statistically warranted, an important advantage for resource-constrained edge IoT deployments. Several limitations merit acknowledgment. First, the evaluation was conducted on simulated rather than natively collected clinical time-series. Second, the framework assumes continuous access to ground-truth labels, which introduces labeling latency in practice. Future work should investigate semi-supervised adaptation strategies and privacy-preserving federated deployment.

VII. CONCLUSION

This paper presented a self-adaptive machine learning framework for continuous health risk assessment in IoT-enabled monitoring environments. The framework integrates an ADWIN drift detection unit with an event-driven adaptive retraining controller, enabling autonomous response to distributional changes in physiological data streams without manual intervention. Experimental evaluation demonstrated consistent outperformance over static and periodic-retrain baselines, achieving a mean F1-score of 0.934 with rapid post-drift recovery.

Developed at Bharati Vidyapeeth Deemed to be University College of Engineering, Pune, this work establishes a rigorous foundation for future development of resilient, continuously adaptive intelligent health monitoring systems.

ACKNOWLEDGMENT

The authors gratefully acknowledge the Department of Electrical and Computer Engineering, Bharati Vidyapeeth Deemed to be University College of Engineering, Pune, for providing the academic infrastructure and support for this research. The authors declare no conflicts of interest.

REFERENCES

[1] World Health Organization, "Global status report on noncommunicable diseases 2022," WHO Press, Geneva, 2022.

[2] J. Gama et al., "A survey on concept drift adaptation," ACM Comput. Surv., vol. 46, no. 4, pp. 1–37, 2014.

[3] I. Zliobaite, "Learning under concept drift: An overview," arXiv:1010.4784, 2010.

[4] G. I. Webb et al., "Characterizing concept drift," Data Min. Knowl. Discov., vol. 30, no. 4, pp. 964–994, 2016.

[5] A. Bifet and R. Gavaldà, "Learning from time-changing data with adaptive windowing," in Proc. SIAM SDM, 2007, pp. 443–448.

[6] S. Palaniappan and R. Awang, "Intelligent heart disease prediction system," in Proc. IEEE AICCSA, 2008, pp. 108–115.

[7] L. Breiman, "Random forests," Mach. Learn., vol. 45, no. 1, pp. 5–32, 2001.

[8] A. Hannun et al., "Cardiologist-level arrhythmia detection using a deep neural network," Nature Med., vol. 25, pp. 65–69, 2019.

[9] D. Sculley et al., "Hidden technical debt in machine learning systems," NeurIPS, pp. 2503–2511, 2015.

[10] J. Gama et al., "Learning with drift detection," in Proc. SBIA, 2004, pp. 286–295.
[11] M. Baena-García et al., "Early drift detection method," in Proc. ECML PKDD Workshop, 2006, pp. 77–86.
[12] C. Alippi and M. Roveri, "Just-in-time adaptive classifiers," IEEE Trans. Neural Netw., vol. 19, no. 7, pp. 1145–1153, 2008.
[13] L. L. Minku and X. Yao, "DDD: A new ensemble approach for concept drift," IEEE Trans. Knowl. Data Eng., vol. 24, no. 4, pp. 619–633, 2012.
[14] Z.-H. Zhou and Y. Yu, "Adapt to concept drift," in Proc. MLHC, 2021.
[15] R. Polikar et al., "Learn++: An incremental learning algorithm," IEEE Trans. Syst. Man Cybern. C, vol. 31, pp. 497–508, 2001.
[16] A. Bifet et al., "MOA: Massive online analysis," J. Mach. Learn. Res., vol. 11, pp. 1601–1604, 2010.
[17] A. Bifet and R. Gavaldà, "Adaptive learning from evolving data streams," in Proc. IDA, 2009, pp. 249–260.
[18] J. W. Smith et al., "Using the ADAP learning algorithm to forecast diabetes mellitus," in Proc. SCAMC, 1988, pp. 261–265.
[19] R. Detrano et al., "International application of a new probability algorithm for diagnosis of coronary artery disease," Amer. J. Cardiology, vol. 64, pp. 304–310, 1989.
[20] J. Gama et al., "Issues in evaluation of stream learning algorithms," in Proc. KDD, 2009, pp. 329–338.