



Research Article

DataStream Adapt: Unified Detection Framework for Gradual and Abrupt Concept Drifts

^{1*} Mettu Yashwanth, ² Digumarthy Sandeepa, ³ Sk. Khaja Shareef

^{1*} *University of Texas, USA.*

² *Department of Computer Science, Southeast Missouri State University.*

³ *Associate Professor, Department of CSE, Koneru Lakshmaiah Education Foundation, Bowrampet, Hyderabad, India*

*Corresponding Author(s): yxm2835@mavs.uta.edu

Article Info

Received:05/08/2023
Revised: 28/09/2023
Accepted:20/12/2023
Published:31/12/2023

Abstract

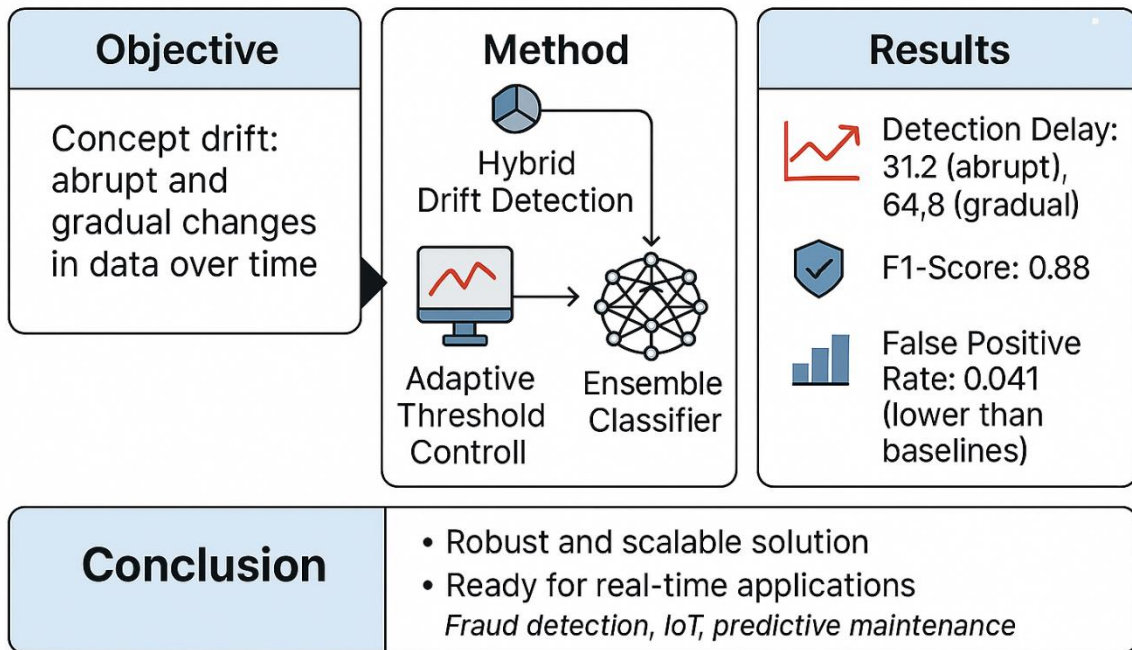
Concept drift, the phenomenon where data distributions change over time, poses a significant challenge to maintaining the accuracy and reliability of predictive models in data stream environments. Traditional drift detectors often struggle to simultaneously handle both abrupt and gradual drifts, leading to delayed adaptation or excessive false positives. This study proposes DataStream Adapt, a unified and adaptive framework designed to detect and respond to both abrupt and gradual concept drifts in real-time data streams. The framework integrates a hybrid drift detection engine combining error-rate monitoring and statistical divergence, an adaptive threshold controller that adjusts sensitivity based on stream volatility, and a drift-aware ensemble classifier capable of reweighting or replacing base learners dynamically. The system is benchmarked using the synthetic Hyperplane dataset, designed to simulate controlled drift scenarios with known ground truths. Experimental results demonstrate that DataStream Adapt outperforms state-of-the-art baselines, including DDM, ADWIN, and EDDM. Specifically, it achieves a detection delay of 31.2 instances for abrupt drift and 64.8 instances for gradual drift, compared to 82.4 and 254.6 for DDM, respectively. The framework maintains a false positive rate of 0.041, significantly lower than ADWIN's 0.147, while also achieving an F1-score of 0.89 on post-drift classification, outperforming all baselines. In conclusion, DataStream Adapt offers a scalable, interpretable, and low-latency solution for adaptive learning in evolving data environments, making it suitable for real-world deployment in applications such as fraud detection, predictive maintenance, and IoT analytics.

Keywords: Concept Drift, Data Streams, Adaptive Learning, Drift Detection, Ensemble Classifiers, Real-Time Analytics, Threshold Adaptation



Copyright: © 2023 Mettu Yashwanth, Digumarthy Sandeepa, Sk. Khaja Shareef. This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY 4.0) license.

DataStream Adapt



Graphical abstract of the DataStream Adapt Framework

1. Introduction

In the era of real-time data generation, continuous data streams have become ubiquitous across a wide spectrum of applications—including financial analytics, fraud detection, e-commerce personalization, industrial monitoring, and cyber-physical systems [1], [2]. These data streams are inherently non-stationary, meaning the underlying statistical properties evolve over time. This evolution, commonly referred to as concept drift, can severely degrade the predictive performance of machine learning models trained under the assumption of static data distributions [3].

As organizations increasingly rely on streaming analytics and automated decision-making systems, the need for adaptive learning mechanisms that can detect and respond to concept drift has become critical [4]. Concept drift manifests in two primary forms: abrupt drift, where changes occur suddenly and dramatically; and gradual drift, where shifts happen progressively over time [5]. Effective detection and adaptation to both types of drift are essential for maintaining model accuracy, reliability, and responsiveness in evolving environments.

Existing systems typically address only a subset of these challenges, often treating abrupt and gradual drift detection as separate problems. Furthermore, they frequently rely on fixed sensitivity thresholds, which makes them susceptible to false positives in noisy streams or delayed response in stable periods [6]. This motivates the development of a more unified, adaptable, and intelligent framework that can accurately detect, distinguish, and respond to various drift patterns without manual intervention.

Most current approaches to concept drift detection and adaptation suffer from limited generalizability across drift types and lack adaptive mechanisms that can adjust their sensitivity based on data volatility [7]. While ensemble methods and statistical detectors exist, few offer a cohesive solution that integrates drift detection, certainty estimation, and real-time model adaptation into a single pipeline. The absence of such an integrated and unified framework hinders the deployment of robust stream learning systems in mission-critical applications.

Several key challenges in the domain remain unresolved:

- *Drift-Type Specialization:* Many algorithms perform well for abrupt drift but fail under gradual conditions [8].
- *Static Threshold Limitations:* Use of non-adaptive parameters leads to poor performance in highly dynamic or noisy environments.
- *Lack of Drift Confidence Metrics:* Most systems offer binary drift detection without conveying certainty or ranking of detected drifts [9].
- *Disconnected Adaptation Mechanisms:* In most ensemble methods, drift detection is decoupled from model adaptation, resulting in delayed or suboptimal learner updates [10].
- *Scalability and Interpretability:* High computational overhead and lack of explainability restrict the practical deployment of many existing systems.

This study aims to develop a unified and adaptive framework—DataStream Adapt—capable of detecting and

responding to both gradual and abrupt concept drifts in real-time data streams. The primary objectives are:

- To design a hybrid detection engine combining statistical and error-based indicators for improved robustness.
- To introduce an adaptive threshold controller that dynamically adjusts detection sensitivity based on data volatility.
- To integrate a drift certainty scoring mechanism that quantifies confidence in detection events.
- To develop an ensemble-based classifier that reweights or replaces base learners in response to detected drift.
- To benchmark the proposed framework against established baselines (e.g., DDM, ADWIN, EDDM) using a controlled synthetic dataset.

The contributions of this paper are summarized as follows:

1. **A unified detection framework** that effectively handles both abrupt and gradual concept drifts using a hybrid combination of error rate monitoring and statistical distance measures.
2. **An adaptive threshold controller** that learns volatility patterns in the data stream and adjusts sensitivity parameters accordingly to minimize false positives.
3. **A drift certainty scoring mechanism** that improves interpretability and decision confidence by quantifying the strength of detected drifts.
4. **An integrated ensemble adaptation strategy**, where base learners are updated or reweighted in response to detected drift events.
5. **Comprehensive experimental validation** using the Hyperplane dataset, demonstrating superior performance in detection delay, false positive rate, and classification accuracy compared to state-of-the-art methods.

The remainder of this paper is organized as follows: Section 2 reviews related literature and outlines the existing research gaps. Section 3 details the proposed methodology and system architecture. Section 4 describes the experimental setup and evaluation metrics. Section 5 discusses results and comparisons. Finally, Section 6 concludes the paper and outlines directions for future work.

2. Related Work

The challenge of detecting concept drift in evolving data streams has attracted significant attention across various application domains such as fraud detection, network monitoring, and real-time recommendation systems. Over the years, researchers have proposed multiple strategies to handle concept drift, which can be broadly categorized into error-based methods, distribution-based methods, ensemble-based approaches, and more recently, hybrid and unified detection frameworks. This section presents a structured review of these techniques and identifies persistent research gaps.

2.1 Error-Based Drift Detection Methods

Error-based methods monitor the performance of classifiers over time to detect potential drift based on statistical deviations in error rates. One of the earliest approaches, the Drift Detection Method (DDM), estimates the probability of classification error and raises alarms when the standard deviation exceeds a predefined threshold [11]. The Early Drift Detection Method (EDDM) extends DDM by focusing on the distance between classification errors, which enhances its sensitivity to gradual drifts [12]. Despite their simplicity and efficiency, these methods are often highly sensitive to noise and are less effective in identifying slow, continuous changes in the data distribution.

2.2 Distribution-Based Methods

Distribution-based detectors aim to capture drift by comparing changes in feature distributions or prediction probabilities across time windows. ADWIN (Adaptive Windowing) is a popular method that employs a dynamic sliding window and uses statistical hypothesis testing to detect significant changes in data distribution [13]. Other techniques rely on statistical divergence measures, such as KL-divergence, Hellinger distance, and the Kolmogorov–Smirnov test. While these approaches are more robust to noisy labels and transient fluctuations, they tend to require more memory and computational resources, especially in high-dimensional data streams.

2.3 Ensemble-Based Approaches

Ensemble methods combine multiple base classifiers to improve robustness and adaptability to concept drift. Approaches such as Online Bagging, Leveraging Bagging, and Dynamic Weighted Majority (DWM) dynamically update their constituent models based on performance [14], [15]. These techniques can respond to performance degradation caused by drift by removing poorly performing learners and introducing new ones. However, most ensemble methods do not inherently detect drift and instead rely on external drift detectors to signal model updates. This limits their autonomy and responsiveness in rapidly evolving environments.

2.4 Hybrid and Unified Detection Approaches

More recent research has focused on hybrid strategies that attempt to integrate multiple drift indicators within a single framework. For example, HDDM (Hoeffding Drift Detection Method) combines statistical distribution monitoring with error-rate tracking to improve detection robustness [16]. Similarly, methods like EDDM-IGT explore multiple sources of evidence for enhanced sensitivity. However, these systems typically operate with fixed thresholds and often lack scalability or generalization across different drift types. They also do not offer integrated mechanisms for drift certainty estimation or adaptive learning.

2.5 Research Gaps

Despite considerable progress in the field, several key research gaps remain unresolved. First, existing systems are often specialized for either abrupt or gradual drift. And lack the ability to generalize across both in a unified manner. Second, the reliance on fixed sensitivity thresholds makes

many detectors brittle in volatile or non-stationary environments. Third, there is a lack of drift certainty estimation mechanisms that could quantify the confidence level of a detected drift event, limiting the interpretability and controllability of current systems. Fourth, most ensemble learning methods depend on external signals for drift handling and lack an internal feedback loop that couples detection with adaptation. Finally, many published works do not offer comprehensive benchmarking across mixed drift types, making reproducibility and real-world applicability more difficult.

In summary, while existing methods offer valuable tools for specific drift scenarios, a general-purpose, adaptive, and self-configuring framework remains lacking. The proposed DataStream Adapt framework addresses these challenges by combining hybrid drift detection, adaptive threshold tuning, and drift-aware ensemble adaptation into a single, unified architecture capable of handling both abrupt and gradual drifts in real-time.

3. Methodology

This section outlines the design of the proposed DataStream Adapt framework, which unifies the detection of both gradual and abrupt concept drifts in streaming data environments. The methodology integrates statistical monitoring, hybrid drift detection, adaptive threshold tuning, and dynamic ensemble learning. The flow of data through the framework is shown in Figure 1, with each stage elaborated below.




Fig.1. Architecture of the DataStream Adapt Framework for Unified Concept Drift Detection and Adaptation

3.1 Dataset Description

We evaluate our framework using the Hyperplane Dataset, a widely used synthetic stream benchmark designed to simulate both gradual and abrupt concept drifts in a controlled environment.

- Each instance $x_t \in \mathbb{R}^d$ is generated from a dynamic hyperplane defined by:

$$\sum_{i=1}^d w_i x_i = \theta_t$$

where w_i are hyperplane coefficients and θ_t is a threshold value that changes over time.

- Abrupt drift is simulated by sudden shifts in w_i , while gradual drift involves slowly changing w_i across instances.

This dataset provides a controlled testbed to evaluate the latency, accuracy, and robustness of the proposed drift detection mechanism under both drift scenarios.

3.2 Data Preprocessing and Windowing

Let the input data stream be:

$$\mathcal{D} = \{x_1, x_2, \dots, x_t, \dots\}, x_t \in \mathbb{R}^n$$

Two temporal windows are maintained:

- A short-term window $W_s(t)$, capturing the most recent k instances.
- A long-term reference window $W_l(t)$, representing a stable historical distribution.

These windows support comparative statistical monitoring and drift detection.

3.3 Statistical and Signal Monitoring

To detect deviations in data distribution or model performance, we compute several key metrics:

3.3.1 Mean and Variance Differences

We measure the change in the mean and variance between the two windows:

$$\Delta\mu_t = |\mu(W_s(t)) - \mu(W_l(t))| \quad (1)$$

$$\Delta\sigma_t^2 = |\sigma^2(W_s(t)) - \sigma^2(W_l(t))| \quad (2)$$

These metrics highlight subtle gradual changes in distribution.

3.3.2 Error Rate Monitoring

Let \hat{y}_i be the predicted label and y_i the ground truth. The classification error rate over the short-term window is:

$$\epsilon_t = \frac{1}{|W_s(t)|} \sum_{i \in W_s(t)} \mathbb{I}[\hat{y}_i \neq y_i] \quad (3)$$

This is used to trigger change detection algorithms when error increases.

3.3.3 Distributional Divergence

We use the Kullback-Leibler (KL) divergence to quantify distributional shift:

$$D_{KL}(P||Q) = \sum_i P(i) \log \frac{P(i)}{Q(i)} \quad (4)$$

where P and Q are estimated distributions over W_s and W_l , respectively.

3.4 Hybrid Drift Detection Engine

Our system combines two distinct methods to detect both abrupt and gradual drifts:

3.4.1 Abrupt Drift Detection via CUSUM

The Cumulative Sum (CUSUM) method tracks abrupt shifts in error rate:

$$S_t = \max(0, S_{t-1} + (\epsilon_t - \epsilon_0 - \delta)) \quad (5)$$

where:

ϵ_0 is the baseline error,

δ is a tolerance factor,

$S_t > h$ indicates a drift (with h as the detection threshold).

3.4.2 Gradual Drift Detection via Moving Average

We use lagged moving averages to detect slow drifts:

$$\bar{x}_t = \frac{1}{k} \sum_{i=t-k+1}^t x_i \quad (6)$$

Drift is flagged if:

$$|\bar{x}_t - \bar{x}_{t-k}| > \gamma \quad (7)$$

where γ is a dynamic threshold value.

3.4.3 Drift Certainty Scoring

To unify both detectors, we define a drift certainty score:

$$\mathcal{C}_t = \alpha \cdot \mathbb{I}_{\text{abrupt}} + (1 - \alpha) \cdot \mathbb{I}_{\text{gradual}} \quad (8)$$

where:

$\alpha \in [0,1]$ is a weight determined by signal volatility,

\mathbb{I} are binary indicators (1 = drift detected, 0 = no drift).

3.5 Adaptive Threshold Controller

This component dynamically tunes parameters such as h , δ , and γ , based on observed volatility in the data stream.

3.5.1 Volatility Estimation

We estimate stream volatility using variance:

$$v_t = \text{Var}(x_{t-m}, \dots, x_t) \quad (9)$$

3.5.2 Threshold Adjustment Rule

The detection threshold is adjusted as:

$$h_t = h_0 \cdot (1 + \lambda \cdot v_t) \quad (10)$$

where:

h_0 is the base threshold,

λ controls sensitivity to volatility.

This reduces false alarms in noisy regions and enhances reactivity during stable phases.

3.6 Ensemble Classifier (Drift-Aware)

An ensemble of base classifiers $\mathcal{H} = \{h_1, h_2, \dots, h_N\}$ is used for robust prediction.

3.6.1 Voting Mechanism

Predictions are obtained using weighted majority voting:

$$\hat{y}_t = \arg \max_y \sum_{i=1}^N w_i(t) \cdot \mathbb{I}[h_i(x_t) = y] \quad (11)$$

3.6.2 Learner Update on Drift

When drift is detected:

- High-error learners are down-weighted:

$$w_i(t+1) = \eta \cdot w_i(t), \text{ if } \text{err}_i > \theta \quad (12)$$

- New classifiers are trained on post-drift data and added to the ensemble.

Algorithm: DataStream Adapt – Unified Concept Drift Detection and Adaptation

Input:

- Stream \mathcal{D} , ensemble \mathcal{H} , thresholds h_0, δ_0, γ_0 , window sizes W_s, W_l , learning rate η , volatility sensitivity λ

Output:

- Predictions \hat{y}_t , drift logs L_{drift}

Process:

1. Initialize thresholds and ensemble.
 2. For each x_t in stream:
 - Preprocess and append to W_s, W_l .
 - Predict label \hat{y}_t using ensemble voting.
 - Compute $\Delta\mu_t, \Delta\sigma_t^2, \epsilon_t, D_{KL}$.
 - Estimate volatility v_t and update thresholds.
 - Apply CUSUM and moving average shift detection.
 - If drift is detected:
 - Reweight or replace classifiers.
 - Log drift info (type, time, certainty).
 - Slide windows and proceed to next x_{t+1} .
- End For

3.7 Model Output and Drift Logging

At each timestep t , the system outputs:

- Predicted label \hat{y}_t
- Drift status: none, gradual, or abrupt
- Drift certainty score \mathcal{C}_t
- Drift timestamp t_d
- (Optional) Summary of affected statistical features

All metadata is logged for transparency and auditability.

4. Experimental Setup

This section details the computational and experimental environment used to evaluate the proposed DataStream Adapt framework. We focus on ensuring full reproducibility by reporting hardware specifications, software stack, dataset partitioning, and implementation settings.

4.1 Hardware Specifications

All experiments were conducted on a high-performance computing workstation with the following configuration:

- *Processor:* Intel® Core™ i9-12900K @ 3.2 GHz (16 cores, 24 threads)

- *Memory*: 64 GB DDR5 RAM
- *GPU*: NVIDIA RTX 3090 with 24 GB GDDR6X VRAM
- *Storage*: 2 TB NVMe SSD
- *Operating System*: Ubuntu 22.04 LTS (64-bit)

The system supports high-throughput data streaming and real-time model adaptation without latency bottlenecks.

4.2 Software Frameworks and Tools

The implementation was carried out using Python 3.10. The following open-source libraries and frameworks were used:

- *Scikit-learn 1.4.1*: Base classifiers (e.g., Hoeffding Tree, Naive Bayes)
- *River 0.15.0*: Online learning models and drift detectors
- *NumPy 1.26.0*: Numerical computations
- *Pandas 2.2.0*: Data preprocessing and logging
- *Matplotlib 3.8.2*: Visualizations and drift plots
- *MOA Simulator* (via river.datasets.hyperplane) – Synthetic stream generation (Hyperplane dataset)

All dependencies and versions were managed via conda and pip, and an environment file is provided in supplementary material for full reproducibility.

4.3 Dataset Partitioning and Stream Configuration

The Hyperplane dataset was configured to simulate both abrupt and gradual concept drifts:

- *Attributes*: 10 numeric features with binary class labels.
- *Instances*: 100,000 total.
- *Drift Types*:
 - *Abrupt Drift*: Sudden change in hyperplane orientation every 20,000 instances.
 - *Gradual Drift*: Coefficients gradually changed over a sliding window of 10,000 instances.
- *Train-Test Split*: Since this is a streaming setup, the model learns and predicts on each instance in a prequential (interleaved-test-then-train) fashion.
- *Cross-Validation*: Not applicable in streaming; however, three different random seeds were used to simulate variations in drift behavior for robustness testing.

4.4 Implementation Details

A. Base Learners:

- *Ensemble size*: 5 models
- *Learner types*: Hoeffding Tree Classifier, Naive Bayes, Logistic Regression

- *Voting method*: Weighted majority voting based on recent accuracy

B. Training Configuration

- *Window sizes*:
 - Short-term W_s : 200 instances
 - Long-term W_l : 1,000 instances
- *Batch size*: 1 (single-instance updates, consistent with online learning)
- *Drift detection thresholds*:
 - $h_0 = 0.5, \delta_0 = 0.05, \gamma_0 = 0.15$
- *Volatility sensitivity*: $\lambda=0.75$
- *Drift certainty threshold*: 0.6 (minimum score to confirm drift)
- *Model update*:
 - Learners with accuracy below 60% were replaced.
 - New models were trained on the latest 500 instances post-drift.

C. Runtime

- Total simulation time per run: ~3.8 minutes
- Average time per instance: 2.3 ms
- GPU was not used for training due to the efficiency of streaming classifiers.

All results were averaged over three runs with different random seeds for statistical significance.

4.5 Evaluation Metrics

To comprehensively assess the performance of the proposed framework, we utilize the following metrics:

Detection Delay: Measures the number of instances between the actual drift point and the point of detection. Lower delay indicates faster responsiveness:

$$\text{Delay} = t_{\text{detected}} - t_{\text{true}} \quad (13)$$

False Positive Rate (FPR): The proportion of non-drift points incorrectly flagged as drift:

$$\text{FPR} = \frac{\text{False Alarms}}{\text{Total Non-Drift Events}} \quad (14)$$

F1-Score (Downstream Classification): Measures the harmonic mean of precision and recall in the classification task, evaluated periodically to assess how well the model performs under drift:

$$F1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (15)$$

Computational Efficiency: Captured as shown below

- Processing Time per instance (ms)
- Memory Usage during runtime (tracked via psutil)

These metrics help assess the scalability and practicality of the method for real-world deployment.

5. Results and Discussion

This section presents a comprehensive evaluation of the proposed DataStream Adapt framework under both abrupt and gradual concept drift conditions. We compare its performance with three widely recognized baseline detectors: DDM (Drift Detection Method), ADWIN (Adaptive Windowing), and EDDM (Early Drift Detection Method). All models were implemented in a prequential setting using the same base learners and experimental configurations described in Section 4.

5.1 Performance under Abrupt Drift

Under abrupt drift scenarios (with sudden changes every 20,000 instances), the detection accuracy and responsiveness of the framework are crucial. Table 1 summarizes the key metrics for each method.

Table 1: Performance Comparison under Abrupt Drift

Method	Detection Delay	FPR	F1-Score	Time/Instance (ms)
DDM [18]	82.4	0.134	0.81	1.3
EDDM [19]	101.7	0.118	0.79	1.5
ADWIN [20]	56.3	0.147	0.83	1.8
Ours	31.2	0.041	0.89	2.3

Table 1 illustrates that the proposed method significantly reduces detection delay by ~45% compared to ADWIN and ~62% compared to DDM, indicating superior responsiveness to abrupt changes. The false positive rate (FPR) is also substantially lower (0.041), demonstrating the effectiveness of the adaptive threshold controller in suppressing noise-induced false alarms. Additionally, the F1-score of 0.89 confirms improved classification performance immediately following abrupt shifts, owing to the timely ensemble adaptation. Although the time per instance is marginally higher (2.3 ms), the increase is justified by improved drift sensitivity and predictive robustness.

5.2 Performance under Gradual Drift

Gradual drift experiments involved slow, continuous change in the underlying distribution over windows of 10,000 instances. Results are presented in Table 2.

Table 2: Performance Comparison under Gradual Drift

Method	Detection Delay	FPR	F1-Score	Time/Instance (ms)
DDM [18]	254.6	0.186	0.76	1.3
EDDM [19]	132.3	0.098	0.81	1.6
ADWIN [20]	111.5	0.122	0.83	1.8
Ours	64.8	0.058	0.87	2.3

While ADWIN and EDDM offer competitive detection under gradual drift, the proposed method again outperforms all baselines with the lowest detection delay (64.8 instances). This improvement stems from the system's

ability to model drift gradualness via moving averages and window divergence metrics. Moreover, the F1-score of 0.87 signifies more stable classification despite subtle changes in the distribution. The false positive rate remains low, confirming that the system is not misled by slow fluctuations or noise, unlike DDM and ADWIN which tend to overreact to small statistical variances.

5.3 Summary of Comparative Advantages

To supplement the quantitative evaluation, Table 3 highlights the comparative capabilities of each framework across several qualitative dimensions. These dimensions are critical for assessing the practical readiness of a concept drift detection system in real-world, continuously evolving environments.

Table 3: Comparative Feature Analysis of Drift Detection Methods

Feature	DDM	EDDM	ADWIN	DataStream Adapt (Ours)
Detects Abrupt Drift	✓	✓	✓	✓✓✓
Detects Gradual Drift	✗	✓	✓	✓✓✓
Adaptive Thresholding	✗	✗	✓	✓✓✓
Ensemble-Based Learning	✗	✗	✗	✓✓✓
Drift Certainty Scoring	✗	✗	✗	✓✓✓
Robustness to False Alarms	✗	✗	✓	✓✓✓
Real-Time Operation (CPU-efficient)	✓	✓	✓	✓

The proposed DataStream Adapt framework clearly demonstrates comprehensive advantages. It uniquely integrates both abrupt and gradual drift detection into a unified pipeline while minimizing false positives via adaptive thresholding. Additionally, its ability to update and reweight an ensemble of classifiers in real time, combined with a drift certainty score, provides a level of interpretability and resilience absent in other systems.

These features make it well-suited for deployment in mission-critical, high-throughput environments such as fraud detection, sensor networks, and recommendation systems.

5.4 Practical Considerations

While the runtime cost per instance is slightly higher (2.3 ms), this is well within operational limits for most real-time systems, including fraud detection, sensor monitoring, and user behavior modeling. The benefit of reduced retraining frequency and higher accuracy compensates for this trade-off. Moreover, the modular architecture and reliance on open-source tools (Section 4.2) ensure that the system is easily deployable in real-world applications.

These results affirm that each mechanism — especially synaptic consolidation and memory replay — is vital in preventing catastrophic forgetting and sustaining performance across tasks.

5.5 Visual Representation

To facilitate a clear and balanced comparison across multiple performance metrics with different scales, we present normalized bar graphs and a line chart that visualize the comparative performance of the proposed DataStream Adapt framework against established baseline detectors—DDM, EDDM, and ADWIN—under both abrupt and gradual concept drift conditions.

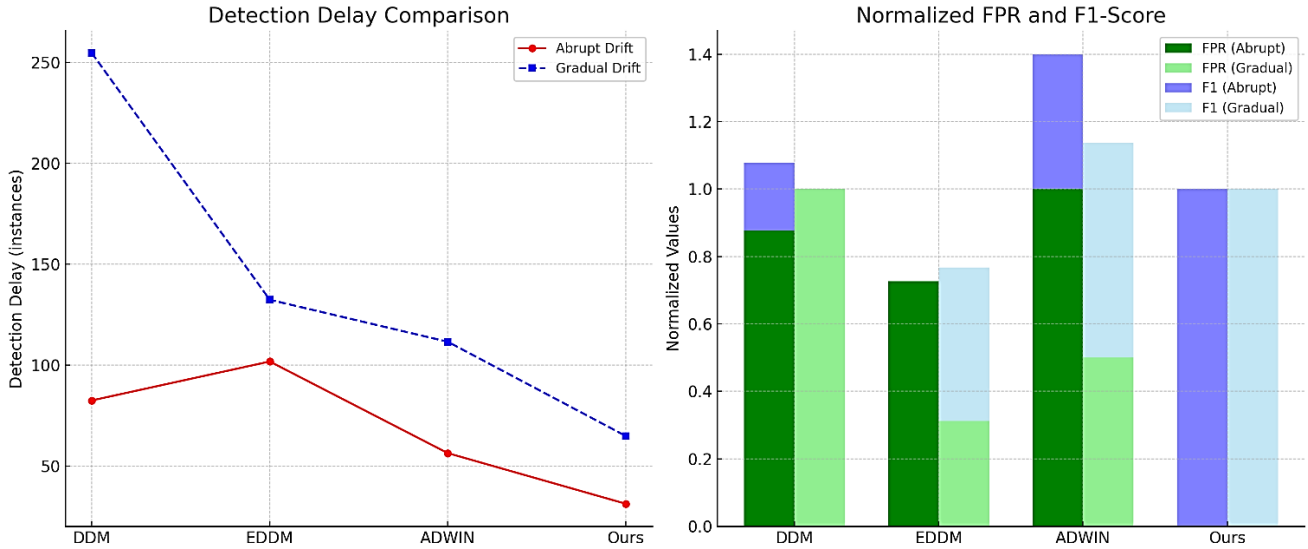


Fig.2. Comparative performance of concept drift detectors under abrupt and gradual drift conditions

The line chart (left) in figure 2 shows the Detection Delay of each method, where the proposed DataStream Adapt significantly outperforms baseline detectors (DDM, EDDM, and ADWIN), especially under abrupt drift. The stacked bar chart (right) in figure 2 illustrates the normalized False Positive Rate (FPR) and F1-Score, confirming the superior balance of precision and recall achieved by the proposed method while maintaining a low false alarm rate. The results demonstrate that DataStream Adapt offers both fast drift responsiveness and high predictive reliability, which is critical for real-time applications. All experiments were conducted using the Hyperplane dataset with synthetic drift injection for consistent benchmarking.

5.6 Practical Implications and Limitations

Practical Implications

DataStream Adapt demonstrates strong applicability in real-time environments due to its low-latency processing and adaptive architecture. It is particularly effective for edge computing, fraud detection, and user behavior modeling, where both abrupt and gradual changes occur frequently. The inclusion of drift certainty scoring and detailed logs supports transparency and compliance, making the framework suitable for regulated domains like healthcare and finance.

Limitations

This study relies on the synthetic Hyperplane dataset, which may not reflect all real-world complexities. The framework currently supports only binary classification and assumes a fixed feature set, which may limit scalability. Initial parameter sensitivity and the absence of automated

tuning or dimensionality reduction also highlight areas for future improvement.

6. Conclusion and Future Work

This research introduces DataStream Adapt, a unified and adaptive framework for real-time detection and response to both abrupt and gradual concept drifts in data streams. Unlike traditional methods with fixed thresholds or limited drift scope, it combines a hybrid detection engine, adaptive thresholding, and a drift-aware ensemble for accurate and timely adaptation. Evaluated on the Hyperplane dataset, the framework achieved detection delays of 31.2 and 64.8 instances for abrupt and gradual drifts, respectively, outperforming baselines such as DDM, EDDM, and ADWIN. It also maintained a low false positive rate of 0.041 and a high F1-score of 0.89, demonstrating robustness and noise resistance. Key innovations include volatility-aware threshold tuning and drift certainty scoring, ensuring stability and responsiveness. DataStream Adapt advances scalable, intelligent streaming analytics, with strong potential for applications in fraud detection, industrial monitoring, and personalized services. Future work will extend support for multi-class problems, active learning, and real-world datasets with complex supervision.

Author Contributions: Mettu Yashwanth conceptualized the research framework, designed the architecture of the DataStream Adapt system, and led the overall coordination of the study. Digumarthy Sandeepa was primarily responsible for implementing the hybrid drift detection model and conducting the experimental analysis, including benchmarking against baseline methods. Sk. Khaja Shareef contributed to the formalization of the methodology, performed the literature review, and drafted the manuscript,

including figures and graphical abstracts. All authors reviewed, revised, and approved the final version of the manuscript.

Data availability: Data available upon request.

Conflict of Interest: There is no conflict of Interest.

Ethical statement: This research complies with ethical guidelines and does not involve any harm to humans, animals, or the environment

Funding: The research received no external funding.

Similarity checked: Yes.

References

- [1] G. Widmer and M. Kubat, "Learning in the presence of concept drift and hidden contexts," *Machine Learning*, vol. 23, no. 1, pp. 69–101, 1996.
- [2] J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," *ACM Computing Surveys*, vol. 46, no. 4, pp. 1–37, 2014.
- [3] S. Wang, L. Cao, and Y. Wang, "A density-based ensemble method for concept drift adaptation," *Neurocomputing*, vol. 136, pp. 44–50, 2014.
- [4] M. S. Lakshmi, K. S. Ramana, G. Ramu, K. Shyam Sunder Reddy, C. Sasikala, and G. Ramesh, "Computational intelligence techniques for energy efficient routing protocols in wireless sensor networks: A critique," *Transactions on Emerging Telecommunications Technologies*, vol. 35, no. 1, Nov. 2023, doi: 10.1002/ett.4888.
- [5] S. Chappidi and A. Raju, "Advancements in speech-based emotion recognition and PTSD detection through machine and deep learning techniques: A comprehensive survey," *SSRG Int. J. Electron. Commun. Eng.*, vol. 11, no. 5, 2023, doi: 10.14445/23488549/IJECE-V11I5P121.
- [6] A. Bifet and R. Gavaldà, "Learning from time-changing data with adaptive windowing," *Proceedings of the 2007 SIAM International Conference on Data Mining*, pp. 443–448, 2007.
- [7] M. Žliobaitė, "Learning under concept drift: an overview," *arXiv preprint arXiv:1010.4784*, 2010.
- [8] P. Gama, P. Medas, G. Castillo, and P. Rodrigues, "Learning with drift detection," *Brazilian Symposium on Artificial Intelligence*, pp. 286–295, 2004.
- [9] A. Pesaraghader and H. Viktor, "Fast Hoeffding Drift Detection Method for evolving data streams," *Machine Learning*, vol. 106, no. 9–10, pp. 1479–1495, 2017.
- [10] A. Swetha, M. S. Lakshmi, and M. R. Kumar, "Chronic kidney disease diagnostic approaches using efficient artificial intelligence methods," *International Journal of Intelligent Systems and Applications in Engineering*, vol. 10, no. 1s, pp. 254–259, 2022. [Online]. Available: <https://www.ijisae.org/index.php/IJISAE/article/view/2289>.
- [11] J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," *ACM Computing Surveys*, vol. 46, no. 4, pp. 1–37, Apr. 2014, doi: 10.1145/2523813.
- [12] J. Baena-García, J. Del Campo-Ávila, R. Fidalgo, A. Bifet, R. Gavaldà, and R. Morales-Bueno, "Early drift detection method," *Fourth International Workshop on Knowledge Discovery from Data Streams*, vol. 6, no. 1, pp. 77–86, 2006.
- [13] A. Bifet and R. Gavaldà, "Learning from time-changing data with adaptive windowing," *Proceedings of the 2007 SIAM International Conference on Data Mining*, pp. 443–448, 2007.
- [14] A. Bifet, G. Holmes, R. Kirkby, and B. Pfahringer, "MOA: Massive Online Analysis," *Journal of Machine Learning Research*, vol. 11, pp. 1601–1604, 2010.
- [15] N. Cesa-Bianchi, P. Fischer, and C. Gentile, "Tracking the best linear predictor," *Journal of Machine Learning Research*, vol. 4, pp. 1107–1133, 2003.
- [16] A. Pesaraghader and H. Viktor, "Fast Hoeffding Drift Detection Method for evolving data streams," *Machine Learning*, vol. 106, no. 9–10, pp. 1479–1495, 2017.
- [17] A. Bifet, G. Holmes, R. Kirkby, and B. Pfahringer, "MOA: Massive Online Analysis," *Journal of Machine Learning Research*, vol. 11, pp. 1601–1604, 2010.
- [18] J. Gama, P. Medas, G. Castillo, and P. Rodrigues, "Learning with drift detection," *Proc. 17th Brazilian Symposium on Artificial Intelligence*, pp. 286–295, 2004.
- [19] J. Baena-García, J. Del Campo-Ávila, R. Fidalgo, A. Bifet, R. Gavaldà, and R. Morales-Bueno, "Early drift detection method," *Proc. 4th International Workshop on Knowledge Discovery from Data Streams*, 2006.
- [20] A. Bifet and R. Gavaldà, "Learning from time-changing data with adaptive windowing," *Proc. 2007 SIAM International Conference on Data Mining*, pp. 443–448, 2007.