



OLR-WAA: Adaptive and Drift-Resilient Online Regression with Dynamic Weighted Averaging

Mohammad Abu-Shaira¹ · Weishi Shi¹

Received: 17 February 2025 / Revised: 25 June 2025 / Accepted: 1 August 2025 / Published online: 28 November 2025
 © The Author(s) 2025

Abstract

Real-world datasets frequently exhibit evolving data distributions, reflecting temporal variations and underlying shifts. Overlooking this phenomenon, known as concept drift, can substantially degrade the predictive performance of the model. Furthermore, the presence of hyperparameters in online models exacerbates this issue, as these parameters are typically fixed and lack the flexibility to dynamically adjust to evolving data. This paper introduces “OLR-WAA: An Adaptive and Drift-Resilient Online Regression with Dynamic Weighted Average”, a hyperparameter-free model designed to tackle the challenges of non-stationary data streams and enable effective, continuous adaptation. The objective is to strike a balance between model stability and adaptability. OLR-WAA incrementally updates its base model by integrating incoming data streams, utilizing an exponentially weighted moving average. It further introduces a unique optimization mechanism that dynamically detects concept drift, quantifies its magnitude, and adjusts the model based on real-time data characteristics. Rigorous evaluations show that it matches batch regression performance in static settings and consistently outperforms or rivals state-of-the-art online models, confirming its effectiveness. Concept drift datasets reveal a performance gap that OLR-WAA effectively bridges, setting it apart from other online models. In addition, the model effectively handles confidence-based scenarios through a conservative update strategy that prioritizes stable, high-confidence data points. Notably, OLR-WAA converges rapidly, consistently yielding higher R^2 values compared to other online models.

Keywords Online learning · Online regression · Adaptive learning · Non-stationary data streams · Concept drift · Hyperparameters optimization

Abbreviations

OLR-WA	Online Regression with Weighted Average	ORR	Online Ridge Regression
OLR-WAA	Online Regression with Weighted Average Adaptive	OLR	Online Lasso Regression
SGD	Stochastic Gradient Descent	PA	Passive-Aggressive Algorithm
MBGD	Mini-Batch Gradient Descent	PA-I, PA-II, PA-III	Variants of Passive-Aggressive Learning
LMS	Least Mean Squares (Widrow-Hoff Algorithm)	SAM	Self-Adjusting Memory
RLS	Recursive Least Squares	MSE	Mean Squared Error
ORR	Online Ridge Regression	RMSE	Root Mean Squared Error
		MAE	Mean Absolute Error
		R^2	Coefficient of Determination
		CCPP	Combined Cycle Power Plant Dataset
		KCHSD	King County House Sales Dataset
		MCPD	Medical Cost Personal Dataset
		1KC	1000 Companies Profit Dataset
		KPI	Key Performance Indicator
		KPI-Win	KPI Window
		KWS	KPI Window Size
		LB	Lower Bound

✉ Mohammad Abu-Shaira
 shairaabu-shaira@my.unt.edu

Weishi Shi
 weishi.shi@unt.edu

¹ Department of Computer Science and Engineering,
 University of North Texas, Denton, Texas, USA

UB	Upper Bound
α	Weighting factor for model averaging
δ	Scaling factor for KPI window size
λ	Regularization parameter
ϵ	Insensitivity parameter for Passive-Aggressive Algorithm
τ (Threshold)	Dynamically determined based on standard deviation and a scaling factor
σ (Sigma)	Standard Deviation
Z+	Natural Number ≥ 1
DM	Drift Magnitude
MBKPI	Mini-Batch KPI
SM	Scale Map
t	Time step or iteration index
η (Learning Rate)	Step size for gradient-based optimization methods
i.i.d. (Independent and Identically Distributed)	Assumption that samples are drawn from the same distribution and are mutually independent
P (Covariance Matrix)	Matrix used in Recursive Least Squares (RLS) to estimate parameter uncertainty
I	Identity Matrix
CUR-MB	Current Mini-Batch
μ KPI	Mean of the KPI-Window (KPI-Win)
EWMA	Exponentially Weighted Moving Average
K (Mini-Batch Size)	The number of samples used in each mini-batch for model updates
U (User-Defined Value)	A manually specified value used as a threshold or configuration parameter
E (Epochs)	The number of full passes through the training dataset

1 Introduction

Traditional machine learning methods operate predominantly in a batch processing paradigm, where the complete dataset is assumed to be available for all computations. This framework enables iterative analysis, including multiple interactions with the data and rigorous cross-validation through subset exclusion. Moreover, it benefits from the absence of temporal constraints, allowing extensive training time, and assumes a stationary data distribution, typically modeled as independent and identically distributed (*i.i.d.*). However, these assumptions [1] present an oversimplified

and often unrealistic representation of real-world data, making batch-based approaches less suitable for dynamic and evolving datasets.

The assumptions underlying batch learning significantly constrain its practical applicability, particularly in dynamic environments. For instance, a machine learning model trained to predict stock prices relies on historical market data, yet the inherent nonstationarity of financial markets causes predictive accuracy to degrade over time [2]. Economic fluctuations such as recessions and inflation, along with structural market changes such as new companies go public and existing ones file for bankruptcy, continuously reshape stock price behaviors. As a result, the initial training data become increasingly outdated, rendering the model less representative of current market conditions. Without regular updates and adaptation, its predictions lose reliability, highlighting the limitations of batch learning in evolving data environments.

Modern streaming environments, characterized by continuous data arrival and the demand for timely predictions, render the batch paradigm impractical. Users cannot afford to wait for complete model retraining with each new data increment, necessitating a shift toward online learning. Unlike batch processing, online learning operates with only a subset of the data at any given time, provides timely responses, and adapts to evolving data distributions [3]. Consider, for instance, a smart city traffic management system. Real-time predictions of traffic flow and congestion are crucial. Given the rapid fluctuations in traffic conditions due to accidents, road closures, or unforeseen events, continuous model updates with real-time traffic data, potentially exhibiting both temporal and spatial correlations, are essential. This enables the system to provide drivers with accurate, up-to-the-minute information, empowering them to make informed routing decisions and thereby mitigating congestion in real time. The adoption of online models brings forth a transformative paradigm in predictive modeling, offering agility, efficiency, and adaptability in the face of ever-changing data landscapes [4].

However, online learning encounters several challenges, including effectively managing *concept drift* [5, 6], a prevalent issue in real-world streaming applications where data characteristics evolve unpredictably over time. These shifts can arise due to changing user behaviors, evolving market trends, seasonal variations, or unforeseen external factors, requiring models to continuously adapt while maintaining stability and predictive accuracy. Additionally, the presence of hyperparameters in online learning models presents a significant challenge, more specifically, with the existence of concept drift. Consequently, expecting users to manually recalibrate these parameters each time concept drift occurs is impractical [7].

Among online learning approaches, *Linear Regression* and its online counterpart play a crucial role in data analysis and predictive modeling [8]. Linear regression models the relationship between a dependent variable y and independent variables X as $y = X\beta + \epsilon$, where β represents the regression coefficients and ϵ is the error term, assumed to follow a normal distribution. Its online counterpart incrementally updates the model as new data arrives, making it well-suited for streaming environments. This adaptability is essential in fields such as finance for stock price forecasting, healthcare for predicting patient outcomes, marketing for demand forecasting, and manufacturing for quality control and process optimization.

Existing online regression methods face several limitations that hinder their effectiveness in dynamic, real-world environments. Many rely on fixed hyperparameters (such as learning rate or regularization), making them highly sensitive and ill-suited for non-stationary data with concept drift, since manually adjusting these parameters at each occurrence of drift is neither feasible nor practical. In addition to this lack of adaptivity, models like OSGD and OMDG, although classified as online methods due to their ability to process one data point or mini-batch at a time, retain previously seen data. This leads to increasing memory usage over time and contradicts the core principle of memory efficiency expected in true online learning. Another common limitation is the lack of a controlled decaying or weighting mechanism, which is essential for maintaining the balance between adaptability and stability. This is a fundamental objective in online learning and is particularly more important in confidence-based scenarios [9, 10], where preserving trust in expert-verified points requires assigning them higher weights to ensure model stability. Moreover, most models do not include built-in drift detection and adaptation mechanisms, leaving them unprepared to respond effectively to evolving data distributions.

This paper introduces OLR-WAA, a novel and versatile online linear regression model specifically designed to handle dynamic data streams in real-time environments. Key contributions include: (1) a novel strategy for dynamically redefining the model in response to emerging data patterns; (2) a proactive, in-memory drift detection and adaptation mechanism that employs variable thresholding and leverages key performance indicators (KPIs) without assuming any specific data distribution, making it suitable for high-dimensional and large-scale data streams; (3) the use of an Exponentially Weighted Moving Average with an automatically inferred decay factor, enabling adaptive response to drift while maintaining a balance between stability and adaptability; and (4) competitive performance on par with batch models in stationary regression scenarios, while effectively managing concept drift in dynamic environments,

demonstrating the robustness and adaptability of OLR-WAA in real-time learning.

In conclusion, OLR-WA (Online Linear Regression with Weighted Average) is a novel online regression model designed for adaptive learning in streaming environments. It incrementally updates its base model using an exponentially weighted moving average to ensure smooth adaptation. The enhanced variant, OLR-WAA, where the final 'A' stands for Adaptive, extends this approach by featuring a novel mechanism for real-time drift detection and automatic hyperparameter inference, enabling robust performance under evolving data distributions. Extensive evaluations across diverse benchmarks demonstrate that our method matches batch regression in static settings and consistently outperforms leading online models in dynamic scenarios. Moreover, OLR-WAA effectively bridges performance gaps caused by concept drift, further distinguishing it from existing approaches. The source code and datasets are publicly available at our GitHub repository [11].

2 Related Work

Linear regression [12, 13] is one of the most widely used and fundamental algorithms in statistical and machine learning. It models the relationship between one or more independent variables and a dependent variable, serving as a mathematical framework for predictive analysis. Additionally, in certain cases, it can be utilized to infer causal relationships. This section first provides a brief overview of batch regression before delving into a detailed exploration of online linear regression and its various models.

Regression can be categorized into simple and multiple regression. Simple linear regression examines the relationship between two continuous variables, where x serves as the predictor (independent variable) and y as the response (dependent variable) [13]. It is mathematically represented as $y = w_0 + w_1x + \epsilon$, where w_0 and w_1 are model parameters and ϵ represents the error term.

Multivariate linear regression (MLR) extends this concept to multiple predictor variables, modeling the response as $y = w_0 + w_1x_1 + \dots + w_mx_m + \epsilon$, where w_i (for $i = 0, 1, \dots, m$) denote the coefficients associated with each predictor. In this formulation, y represents the observed dependent variable, while $w_0 + w_1x_1 + \dots + w_mx_m$ corresponds to the predicted value, and ϵ captures the residual error. The coefficients are typically estimated using the Moore-Penrose pseudoinverse, a widely used method for computing the least-squares solution [13], as formalized in Eq. 1, where $\hat{W}$ denotes the estimated parameter values and N represents the number of data samples.

$$\hat{W} = (X^T X)^{-1} X^T y \quad (1)$$

$$\text{where } \mathbf{W} = \begin{bmatrix} w_0 \\ w_1 \\ \vdots \\ w_M \end{bmatrix}, \mathbf{X} = \begin{bmatrix} 1 & x_{11} & x_{12} & \dots & x_{1M} \\ 1 & x_{21} & x_{22} & \dots & x_{2M} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_{n1} & x_{n2} & \dots & x_{nM} \end{bmatrix},$$

$$y = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{bmatrix}$$

The Mean Squared Error (MSE) objective in linear regression is convex due to its quadratic form, ensuring a unique global minimum. This allows the problem to be efficiently solved using optimization techniques like gradient descent, improving both stability and accuracy [14]. A more general principle for parameter estimation is Maximum Likelihood Estimation (MLE), a method that estimates the parameters of a statistical model by fitting a probability distribution to the data in a way that maximizes the likelihood function [15]. The least squares approach used to determine model parameters in linear regression is a specific case of MLE [16].

Beyond the regression methods discussed earlier, several alternative approaches cater to varying data characteristics and modeling requirements. Among these, Support Vector Regression (SVR) [17] has proven to be a powerful technique for real-valued function estimation, designed to handle noise and uncertainty through Vapnik's ϵ -insensitive loss. It introduces a margin of tolerance where small errors are ignored, while larger deviations are penalized, enabling balanced treatment of over- and underestimations. SVR's regularized objective promotes generalization by controlling model complexity and penalizing deviations beyond the margin [18], making it effective across diverse, noisy datasets.

Statistical learning theory, online convex optimization (OCO), and game theory form the theoretical backbone of online learning techniques [19]. Convexity plays a central role in designing efficient algorithms for streaming data, as many online learning problems can be framed as OCO tasks. In this setting, an algorithm selects a decision from a convex set at each time step and incurs a loss defined by a convex function. The performance is evaluated through cumulative regret, which measures the gap between the algorithm's total loss and that of the best fixed decision in hindsight. Minimizing this regret is fundamental to ensuring the algorithm remains competitive over time.

In general, *Online Learning* models adhere to the framework outlined in Algorithm 1, which is designed to iteratively update model parameters while minimizing the cumulative loss incurred over time.

```

1: Initialize: Set model parameters  $w_1 = 0$ 
2: for each round  $t = 1, \dots, T$  do
3:   Receive training instance  $x_t$ 
4:   Predict label  $\hat{y}_t \in \mathbb{R}$ 
5:   Observe true label  $y_t \in \mathbb{R}$ 
6:   Compute loss  $l(\hat{y}_t, y_t, x_t)$ 
7:   Update model parameters  $w_t$ 
8: end for
    
```

Algorithm 1 General Online Learning Framework [20]

This section provides a comprehensive analysis of prominent online linear regression algorithms, examining their strengths, limitations, and applicability. Through a rigorous comparison, we explore their underlying principles, performance characteristics, and real-world suitability, offering insights into advancements, challenges, and future research directions in the field.

2.1 Online Linear Regression using Stochastic Gradient Descent (SGD)

Stochastic Gradient Descent (SGD), or Online SGD [21], is a fundamental optimization method for large-scale and online learning [22]. It updates model parameters using a single data point per iteration, making it highly scalable with per-step complexity $O(M)$, where M is the number of parameters [23]. SGD is commonly used to minimize Mean Squared Error (MSE) in regression tasks [16], using the negative gradient direction for weight updates [24]. Its stochastic nature introduces gradient noise that acts as implicit regularization, aiding generalization in streaming contexts [25].

Despite its efficiency [26, 27], SGD requires access to previously seen data, leading to memory overhead. It is sensitive to learning rate settings: large values can cause divergence, while small ones may slow convergence [26, 28]. SGD is also vulnerable to outliers [29] and its implicit retention of past patterns makes it less suitable for non-stationary environments.

To improve convergence, adaptive methods like AdaGrad [30], RMSProp [31], and Adam [32] adjust learning rates based on historical gradients. SGDM introduces a momentum term $\beta \in [0, 1)$ to smooth updates and improve stability, especially in non-convex settings [33]. Additionally, mini-batch gradient descent offers a middle ground by reducing variance while maintaining scalability [34].

2.2 Online Regression using Mini-Batch Gradient Descent (MBGD)

Mini-Batch Gradient Descent (MBGD) strikes a balance between Batch Gradient Descent (BGD) and Stochastic Gradient Descent (SGD), offering both computational efficiency and reduced gradient variance in online regression tasks [35, 36]. It updates weights using small batches of size

K , allowing faster convergence and improved stability, especially in streaming and evolving data contexts. By tuning the batch size, practitioners control the trade-off between gradient noise and generalization [37, 38], with smaller batches increasing variance and larger ones reducing stochasticity [27, 39]. The overall time complexity is $O(I \cdot K \cdot M)$ for I iterations and M features, with $O(K \cdot M)$ per iteration [40].

MBGD stabilizes learning in noisy, dynamic environments [27] and benefits from adaptive optimizers like Adam and RMSprop, which adjust learning rates based on past gradients [31, 32]. Its inherent noise also acts as a regularizer, mitigating overfitting and promoting better generalization in online settings [41].

To scale MBGD for large and decentralized data, asynchronous and parallel variants like Hogwild! and Federated Learning have been introduced [42, 43]. Momentum-based techniques and variance-reduction methods such as SVRG and SAGA further improve convergence and robustness under concept drift [28, 44]. Nonetheless, MBGD still requires careful tuning of fixed batch sizes and static optimizer settings under evolving data distributions.

2.3 Widrow-Hoff (LMS)

The Widrow-Hoff algorithm, also known as the Least Mean Squares (LMS) algorithm or Adaline, was introduced by Widrow and Hoff in 1960 [45]. It is a foundational technique in online learning, adaptive filtering, and neural networks, particularly for linear regression and classification [46]. LMS employs a stochastic gradient-based approach that iteratively updates weights using only the current data point, making it efficient for streaming settings [1].

LMS is computationally simple, avoiding matrix inversions required by more complex methods like Recursive Least Squares (RLS) [47]. With a per-iteration complexity of $O(P)$, where P is the number of parameters, it is highly suitable for real-time applications [1].

To overcome LMS's sensitivity to fixed learning rates, Normalized LMS (NLMS) introduces a time-varying step size proportional to the inverse energy of the input [1, 48]. More recent extensions integrate LMS with neural networks and reinforcement learning for nonlinear and autonomous system adaptation [49].

Despite its simplicity, LMS minimizes only the current error without incorporating historical data, limiting its ability to capture long-term patterns [1]. It also requires prior knowledge of input statistics for optimal learning rate selection, and may converge slowly compared to adaptive methods [50]. NLMS can further introduce instability in highly dynamic environments due to potential overcorrection [51].

2.4 Recursive Least Squares (RLS)

Recursive Least Squares (RLS) is a classic adaptive filtering algorithm widely used in system identification and adaptive control [1, 52]. Unlike stochastic methods like LMS, RLS employs a deterministic update mechanism to minimize cumulative squared error over all seen data. Its use of the matrix inversion lemma allows rapid convergence, especially under correlated input signals, outperforming LMS in both accuracy and speed [53].

RLS's key advantage lies in its fast adaptation and lack of a learning rate [1, 53]. The algorithm recursively updates a covariance matrix and parameter vector, ensuring convergence to the optimal solution equivalent to batch least squares [54]. A forgetting factor λ controls how past data influences updates, enabling adaptability to nonstationary environments [54]. Its memory of all prior observations makes it effective in stable, low-noise conditions and scenarios with highly correlated inputs [55].

However, this performance comes at a cost, as RLS has quadratic time complexity $O(P^2)$, making it less practical for high-dimensional or resource-constrained environments [53]. It is also prone to numerical instability due to recursive matrix operations, particularly in ill-conditioned or finite-precision settings [56]. Issues such as round-off errors and finite word-length effects can degrade reliability.

To mitigate these challenges, variants like Fast RLS (FRLS) [57] and Sliding-Window RLS [58] have been proposed, reducing complexity while maintaining accuracy. Additional stabilization methods such as regularization, QR decomposition, and square-root filtering have been developed to improve numerical robustness in practical applications [59].

2.5 Online Ridge Regression (ORR)

Ridge regression introduces L2 regularization to mitigate overfitting and multicollinearity by penalizing large coefficients, thereby promoting model stability and reducing variance [60]. Unlike ordinary least squares, it balances mean squared error minimization with coefficient shrinkage, making it effective for high-dimensional data [15]. Online Ridge Regression (ORR) extends this principle to streaming settings, updating parameters iteratively and incorporating regularization directly into the stochastic gradient updates [61].

ORR is well-suited for large-scale applications such as adaptive control, finance, and recommendation systems, offering linear time complexity per iteration [40]. It ensures real-time learning while maintaining regularization, which prevents overfitting and stabilizes updates over time.

However, ORR's effectiveness depends heavily on selecting an appropriate regularization parameter λ . A

poorly chosen λ can lead to underfitting or overfitting [62]. To address this, adaptive techniques like Bayesian optimization [63] and regularization path tracking [64] have been proposed to tune λ dynamically.

Recent advancements have improved ORR's adaptability and robustness in nonstationary environments. Incremental kernel ridge regression with predictive sampling has shown near-optimal performance in online contexts [65]. Other works introduced novel ridge estimators for high-dimensional regression under multicollinearity [66], and alternative algorithms such as forward-based stochastic methods have demonstrated better regularization resilience and theoretical guarantees [67].

2.6 Online Lasso Regression (OLR)

Lasso regression, or Least Absolute Shrinkage and Selection Operator, introduces L1 regularization to penalize large coefficients and enforce sparsity in linear models, enabling simultaneous regression and feature selection [68, 69]. The regularization parameter λ controls the trade-off between accuracy and model simplicity [15].

Online Lasso Regression (OLR) extends this framework to streaming data by incrementally updating model parameters with each new observation [61]. It is particularly suitable for real-time applications, combining the adaptability of online learning with Lasso's inherent feature selection, which efficiently manages complexity in high-dimensional environments.

OLR relies on stochastic gradient descent, which enables scalability for large and streaming datasets. However, its performance is sensitive to the regularization parameter. Poor tuning of λ may lead to underfitting or overfitting [15]. To address this challenge, adaptive methods such as data-driven regularization tuning with theoretical guarantees have been proposed [70].

Recent innovations include the Online Linearized Lasso by Yang et al. [71], a memory-efficient approach that improves convergence rates by constraining the solution path. For cases where Lasso struggles, especially when predictors outnumber observations, the elastic net offers a more robust alternative by combining L1 and L2 regularization [72]. These developments enhance OLR's reliability across domains like real-time analytics, finance, and recommendation systems.

2.7 Online Passive-Aggressive (PA)

The Online Passive-Aggressive (PA) algorithm [73] is an efficient margin-based online learning method closely related to Support Vector Machines. It updates model parameters only when a prediction error exceeds a predefined threshold,

remaining passive when the loss is zero and aggressive otherwise. This selective update mechanism enables fast and focused learning for real-time regression tasks.

PA offers low time complexity ($O(M)$), quick convergence, and robustness to noise by emphasizing examples near the decision boundary [73]. Its updates are driven by an ϵ -insensitive hinge loss, which ignores small errors and focuses updates on significant deviations. The algorithm has three main variants (PA-I, PA-II, PA-III), differing in how they compute the update magnitude based on regularization and aggressiveness parameters, allowing flexibility in learning dynamics [73].

However, PA's performance is sensitive to the aggressiveness parameter C . Small values may slow convergence, while large values can cause overfitting due to overly aggressive updates [73]. The fixed threshold also makes it less adaptable in dynamic or noisy environments, limiting its responsiveness to evolving data patterns.

Recent enhancements address these limitations. Adaptive PA variants incorporate side information and dynamic threshold selection to improve convergence and robustness in non-stationary environments [74]. Other advancements apply adaptive regularization techniques to stabilize learning in high-dimensional settings and improve generalization under drift [75].

In summary, while significant progress has been made in online linear regression, key limitations remain. These can be summarized as follows:

1. Most online learning algorithms rely on hyperparameters that are impractical to tune manually, especially in the presence of concept drift.
2. Many approaches lack a decay or weighting mechanism, which is crucial for balancing adaptability to new data with stability over time.
3. Existing methods are not drift-aware; they do not detect, quantify, or respond to distributional changes, nor do they adaptively optimize hyperparameters or update strategies accordingly.

These limitations highlight the need for a more flexible and adaptive solution. This motivates the development of OLR-WAA, a drift-aware online regression framework that dynamically adjusts hyperparameters and model behavior to maintain an ideal balance between adaptiveness and stability in evolving data streams.

3 OLR-WA Method

This section presents a comprehensive overview of the “OLR-WA: Online Regression with Weighted Average in Multivariate Data Streams,” highlighting its core principles and distinctive features that establish it as a robust and efficient approach for online regression. It begins by outlining the problem settings, defining the online regression framework and the challenges posed by evolving data streams. Then, it explores the first variant, OLR-WA, detailing its parameter update mechanism and foundational characteristics. Following this, the discussion extends to OLR-WAA, showcasing its advancements over OLR-WA through the development of a novel memory-based automatic hyperparameter optimization mechanism, designed to enhance the model’s adaptability in the presence of concept drift.

3.1 Problem Settings

Consider $\mathcal{X} \subseteq \mathbb{R}^d$ as the input feature space and $\mathcal{Y} \subseteq \mathbb{R}$ as the continuous output space. At each time step t , a new data instance (x_t, y_t) is received, where $x_t \in \mathcal{X}$ represents the feature vector, and $y_t \in \mathcal{Y}$ is the corresponding real-valued target. The data stream is assumed to be potentially infinite, requiring models to process and learn from data sequentially with limited memory resources [76].

The goal of the online regression model is to iteratively update the predictive function $f_t : \mathcal{X} \rightarrow \mathcal{Y}$ in order to minimize the cumulative regression error over time. In typical online learning scenarios, data arrives incrementally, allowing the model to adapt continuously as new observations are introduced. However, real-world data streams often exhibit non-stationarity, where the underlying data distribution changes over time, which is known as *concept drift*.

Formally, let $\{(x_1, y_1), (x_2, y_2), \dots\}$ be a sequence of data points generated from an unknown joint probability distribution $p(x, y)$. Concept drift occurs when this joint distribution shifts between time steps, which can be mathematically expressed as below, where t_n and t_{n+1} denote consecutive time points, reflecting the dynamic nature of the data.

$$\exists x, y : p_{t_n}(x, y) \neq p_{t_{n+1}}(x, y)$$

Concept drift can occur in various forms, such as sudden (abrupt), gradual, or incremental changes, each presenting distinct challenges to maintaining model performance [5, 76]. Effectively managing concept drift is essential for preserving the performance and reliability of online regression models. This work aims to develop an adaptive online regression framework designed to accurately detect and

respond to dynamic changes in data distributions, ensuring sustained predictive accuracy in evolving environments.

3.2 Online Learning Parameter Update

The OLR-WA framework facilitates incremental learning from continuous data streams by maintaining two distinct hyperplanes: the base hyperplane, $f_{\text{base}}(x)$, representing the cumulative knowledge acquired from historical data, and the incremental hyperplane, $f_{\text{inc}}(x)$, capturing insights derived from the current mini-batch at time t , with x denoting the feature vector. The associated norm vectors, V_{base} and V_{inc} , are combined through the Exponentially Weighted Moving Average (EWMA) method. A dynamic smoothing factor $\alpha \in (0, 1]$ regulates this combination, assigning a weight of α to V_{inc} to prioritize recent information, while $(1 - \alpha)$ preserves the influence of prior knowledge encapsulated in the base model. The aggregated norm vector, V_{Avg} , undergoes continuous updates, utilizing the intersection point P_{int} of the base and incremental hyperplanes to refine the base model W_{base} . This iterative process aims to incrementally align W_{base} with W_{inc} as more data becomes available, thereby enhancing the model’s adaptability and learning efficiency. The learning objective of our model is defined as follows:

$$w_{t+1} = \arg \min_w \left\{ \frac{1}{2} \|w - w_t\|^2 + \alpha \mathcal{L}(y_{\text{inc}}, f(w; X_{\text{inc}})) + (1 - \alpha) \mathcal{L}(y_{\text{base}}, f(w; X_{\text{base}})) \right\} \quad (2)$$

The objective function effectively balances *stability* and *adaptability* in an online regression setting. The first term, $\frac{1}{2} \|w - w_t\|^2$, acts as a regularization component, promoting smooth transitions between model updates by penalizing large deviations from the previous weight vector w_t . The second term, $\alpha \mathcal{L}(y_{\text{inc}}, f(w; X_{\text{inc}}))$, represents the mean squared error (MSE) loss calculated over the incremental mini-batch. Here, $\alpha \in (0, 1]$ serves as a weighting factor that emphasizes recent data, thereby enhancing the model’s responsiveness to new patterns or shifts in the data distribution. Conversely, the term $(1 - \alpha) \mathcal{L}(y_{\text{base}}, f(w; X_{\text{base}}))$ corresponds to the MSE loss associated with the base (historical) data. This term ensures that the model retains knowledge from previously encountered data, with $(1 - \alpha)$ controlling the degree of emphasis on past observations. Together, these components enable the model to adapt efficiently to new information while maintaining robustness against concept drift.

OLR-WA utilizes the Exponentially Weighted Moving Average (EWMA) [77] to integrate the base and incremental

models. At the start of each iteration, the weighting factor (α) is either user-defined or set to its default value ($\alpha = 0.5$). In OLR-WAA, however, (α) is dynamically adjusted to optimize adaptability and robustness in response to concept drift, rendering OLR-WAA a hyperparameter-free model. This adaptive weighting strategy is elaborated in Sect. 3.4. The weighted average formula in Eq. 3 maintains consistency across dimensions, integrating the norm vectors of the base and incremental hyperplanes using the dynamically computed (α) and its complement ($1 - \alpha$):

$$V_{\text{Avg}} = \alpha \cdot V_{\text{inc}} + (1 - \alpha) \cdot V_{\text{base}} \quad (3)$$

In the realm of N-dimensional geometry, the definition of a new hyper-plane necessitates the utilization of a norm vector and a point [78]. Let $\mathbf{n} = \langle n_1, n_2, \dots, n_N \rangle$ denote the norm vector, representing the directional characteristics of the desired hyper-plane, and let $P = (X_{1p}, X_{2p}, \dots, X_{Np})$ be a point lying within the hyper-plane. To establish the equation of the new hyper-plane, consider an arbitrary point $Q = (x_1, x_2, \dots, x_N)$ on the hyper-plane. The vector connecting P and Q , denoted as $\overrightarrow{PQ} = \langle x_1 - X_{1p}, x_2 - X_{2p}, \dots, x_N - X_{Np} \rangle$, lies within the hyper-plane and is orthogonal to the norm vector $\mathbf{n}$. Hence, the orthogonality condition can be expressed as $\mathbf{n} \cdot \overrightarrow{PQ} = n_1(x_1 - X_{1p}) + n_2(x_2 - X_{2p}) + \dots + n_N(x_N - X_{Np}) = 0$. This methodology enables the precise definition of a new hyper-plane in N-dimensional space based on its norm vector and a known point lying within the hyper-plane.

Figure 1 illustrates the process of defining the updated hyperplane within a two-dimensional space, which is then adopted as the new base model for the next iteration. This example reflects an adversarial scenario where the incoming

data significantly diverges from the existing base model. The base hyperplane, shown in dark orange, represents the knowledge accumulated up to time ($t - 1$), while the current incoming batch at time (t) is highlighted in deep pink. The corresponding norm vectors for both hyperplanes are computed and visually represented using matching color schemes.

Assuming $\alpha = 0.5$, the weighted average vector, depicted in purple, is calculated using Eq. 3. The new hyperplane, labeled as line l_1 , is established based on the intersection of the base and incremental boundaries, guided by the weighted average norm vector. Additionally, the figure highlights two more hyperplanes: l_2 , corresponding to a higher α value, which accelerates adaptation by giving more importance to new data, and l_3 , representing a confidence-based strategy that prioritizes stability by focusing on data points with higher reliability. These boundaries showcase OLR-WA's flexibility in addressing diverse scenarios, including adversarial environments where a higher α promotes rapid adaptation to significantly divergent data, and confidence-driven contexts where conservative updates ensure model stability. A practical example of the confidence-based approach is sentiment analysis, where data verified by domain experts or trusted sources can be assigned greater weights, reflecting their higher confidence and relevance.

The process of obtaining the point of intersection varies depending on the dimensionality of the space. For instance, in a two-dimensional setting, the intersection of two lines is found by solving a system of two equations with two variables, resulting in a single exact point of intersection. However, in three dimensions, determining the intersection of two planes involves solving a system of two equations with three variables, leading to an infinite number of solutions that represent the points forming the intersecting line.




Fig. 1 OLR-WA one step online learning

In higher dimensions, the intersection of N-Dimensional hyper-planes, is an (N-1)-dimensional hyper-plane which will be generated by solving two equations of N variables. This can be done using methods like Gaussian elimination, or using software packages. The solution will be a point $(x_1, x_2, \dots, x_n)$ that satisfies both equations, representing the intersection of the two hyper-planes. Note that there can be multiple solutions, depending on the hyper-planes' orientation and position in the N-dimensional space. However, any solution is an accurate, to simply demonstrate this in a 2-dimensional space, a directional vector along with any point in the line of the intersection suffices to define the new hyperplane.

The system relies on the presence of an intersection point. In the absence of one, two scenarios may occur, as described by Layton et al. [79]. First, in the “Coincide” case, no update is required as the base and incremental models are identical. Second, in the “Parallel” case, the system calculates a weighted midpoint and uses it as the intersection point. Figure 2 illustrates the concept of the weighted midpoint. In (a), with $\alpha = 0.5$, the weighted midpoint is positioned exactly at the center. In (b), when $\alpha = 0.8$, the weighted midpoint shifts towards the incremental model, along with the corresponding norm vector, collectively defining the new hyperplane.

```

1: Input: Data stream  $\{\mathbf{X}, \mathbf{y}\}$ ,  $\alpha = (0, 1]$ , (default  $\alpha = 0.5$ )
2:  $\mathbf{W}_{\text{base}} \leftarrow \arg \min_{\mathbf{w}} \mathcal{L}_{\text{MSE}}(\mathbf{X}_{\text{base}}, \mathbf{y}_{\text{base}})$   $\triangleright$  (default  $\mathbf{X}_{\text{base}}^\dagger \mathbf{y}_{\text{base}}$ ),
    $\mathbf{X}_{\text{base}}, \mathbf{y}_{\text{base}}$  is the first mini-batch
3: for  $t \leftarrow 1$  to  $T$  do
4:    $\mathbf{W}_{\text{inc}} \leftarrow \arg \min_{\mathbf{w}} \mathcal{L}_{\text{MSE}}(\mathbf{X}_{\text{inc}}, \mathbf{y}_{\text{inc}})$ 
5:   if  $\frac{\mathbf{W}_{\text{base}, i}}{\mathbf{W}_{\text{inc}, i}} = \text{constant} \quad \forall i$  then  $\triangleright$  Check if the hyperplanes coincide
6:     continue
7:   end if
8:    $\mathbf{P}_{\text{int}} \leftarrow \text{FindIntersectPoint}(\mathbf{W}_{\text{base}}, \mathbf{W}_{\text{inc}})$ 
9:    $\mathbf{V}_{\text{base}} \leftarrow \text{ComputeNormVector}(\mathbf{W}_{\text{base}})$ 
10:   $\mathbf{V}_{\text{inc}} \leftarrow \text{ComputeNormVector}(\mathbf{W}_{\text{inc}})$ 
11:   $\mathbf{V}_{\text{base}} \leftarrow \frac{\mathbf{V}_{\text{base}}}{\|\mathbf{V}_{\text{base}}\|_2}$   $\mathbf{V}_{\text{inc}} \leftarrow \frac{\mathbf{V}_{\text{inc}}}{\|\mathbf{V}_{\text{inc}}\|_2}$ 
12:   $\mathbf{V}_{\text{avg}} \leftarrow (1 - \alpha)\mathbf{V}_{\text{base}} + \alpha\mathbf{V}_{\text{inc}}$ 
13:   $\mathbf{W}_{\text{base}} \leftarrow \text{DefineHyperplane}(\mathbf{V}_{\text{avg}}, \mathbf{P}_{\text{int}})$ 
14: end for
15: return  $\mathbf{W}_{\text{base}}$ 

```

Algorithm 2 OLR-WA




Fig. 2 Weighted middle point




Fig. 3 KPIs Window

3.3 Normalizing the Norm Vectors

The norm vector encapsulates both magnitude and direction, with its magnitude varying depending on the coefficients of the corresponding hyperplane. Consequently, *normalizing* the norm vectors is a critical step to ensure consistency and comparability. Consider the scenario where $\alpha = 0.5$, and the weighted average operation simplifies to a summation, represented by the placement of norm vectors at the origin. Without normalization, directly utilizing the norm vectors in their original form would yield a significantly different average norm vector, thereby substantially impacting the influence of the weighting factor α .

Furthermore, normalizing the vectors is a critical step in controlling potential overfitting, a phenomenon where the model fits the training data too closely, capturing noise and thereby reducing its ability to generalize to unseen data. By normalizing the norm vectors, OLR-WA effectively mitigates this risk, ensuring that the model prioritizes learning the underlying patterns rather than memorizing the training data.

3.4 Memory-Based Automatic Hyperparameters Optimization

The presence of hyperparameters in online learning presents significant challenges, primarily due to the evolving nature of data distributions over time, a phenomenon known as concept drift [80]. In such scenarios, it is impractical for users to continuously update hyperparameters with each occurrence of concept drift. Online learning models are inherently designed to adapt to dynamic data streams; however, their ability to do so varies considerably. Some models experience substantial performance degradation at the point of drift but rapidly adjust to the new concept, whereas others require a prolonged period to effectively accommodate the distributional shift and achieve adaptation.

OLR-WAA introduces an in-memory, built-in hyperparameter optimization technique designed for automatic drift detection and adaptation in dynamic data streams. This method leverages memory resources and key performance indicators (KPIs) without assuming any specific underlying data distribution, making it well-suited for managing high-dimensional datasets and efficiently processing large-scale data. The framework is integrated with the Self-Adjusting Memory (SAM) architecture, enhancing

real-time processing capabilities and adaptability. Additionally, it incorporates variable thresholding techniques, further reinforcing its effectiveness in dynamically evolving environments.

OLR-WAA's hyperparameter optimization technique is designed to achieve four key objectives. First, it facilitates drift detection by identifying shifts in the underlying data streams. Second, it quantifies the magnitude of these drifts, offering valuable insights into their significance. Third, it employs a hyperparameter adjustment mechanism that leverages fine-tuned parameters to dynamically adapt the model to immediate fluctuations. Lastly, it ensures model calibration using the optimized hyperparameters to maintain performance consistency. The optimization process is detailed in Algorithm 3, beginning at line 15.

OLR-WAA systematically captures Key Performance Indicators (KPIs) from incoming data streams, which can be processed as mini-batches or individual data points, denoted as (X_{inc}, y_{inc}) . The KPIs for each processed mini-batch are stored in a compact, sequentially structured repository referred to as the KPI-Window (KPI-Win). These KPIs may include metrics such as Accuracy, Precision, Recall, and F1-score, among others. The selection of KPIs is flexible and depends on the specific problem domain. For instance, in *regression* tasks, commonly used performance metrics include Mean Absolute Error (MAE), Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and R-squared (R^2).

The size of the KPI-Window (KWS) is a bounded variable, determined by the formulation presented in Eq. 4. In online learning settings, where the number of incoming data samples is theoretically infinite (∞), the parameter N represents the number of the most recent data points retained within the rolling window, while K denotes the mini-batch size. The scaling factor (δ) is selected based on specific application requirements, with an empirically chosen value of $\delta = 0.05$. The KWS is constrained within predefined lower (LB) and upper (UB) bounds; any values falling below LB are set to LB, whereas values exceeding UB are capped at UB to ensure stability and efficient resource utilization.

The bounds LB and UB are introduced to prevent the KPI-Window from becoming either too narrow or too wide. A very narrow window with insufficient observations may result in noisy KPI estimates, reducing the reliability of performance monitoring. On the other hand, an excessively large window may include outdated observations, which

can lead to inaccurate predictions at times when adaptability is crucial in online learning. Additionally, a large window may delay the model's response to concept drift and increase computational overhead. These bounds are therefore essential to maintain a balance between stability and responsiveness. Empirically, we found that setting the lower bound to 11 and the upper bound to 31 achieves this balance

effectively across various datasets with different drift conditions. These values were selected through grid-search validation to ensure both reliable adaptation and efficient memory use.

$$LB \leq KWS = \left(\frac{N}{K}\right) \times \delta \leq UB \quad (4)$$

```

1: Input: Data stream  $\{\mathbf{X}, \mathbf{y}\}$ ,  $\alpha = (0, 1]$  (default  $\alpha = 0.5$ ),  $z \approx [1.5, 2.5]$ ,
2:    $kpi \subseteq \{R^2, \text{Loss, etc.}\}$ 
3:  $\mathbf{W}_{\text{base}} \leftarrow \arg \min_{\mathbf{w}} \mathcal{L}_{\text{MSE}}(\mathbf{X}_{\text{base}}, \mathbf{y}_{\text{base}})$   $\triangleright$  (default  $\mathbf{X}_{\text{base}}^\dagger \mathbf{y}_{\text{base}}$ ,
    $\mathbf{X}_{\text{base}}, \mathbf{y}_{\text{base}}$  is the first mini-batch)
4: for  $t \leftarrow 1$  to  $T$  do
5:    $\mathbf{W}_{\text{inc}} \leftarrow \arg \min_{\mathbf{w}} \mathcal{L}_{\text{MSE}}(\mathbf{X}_{\text{inc}}, \mathbf{y}_{\text{inc}})$ 
6:   if  $\frac{\mathbf{W}_{\text{base}, i}}{\mathbf{W}_{\text{inc}, i}} = \text{constant} \quad \forall i$  then  $\triangleright$  Check if the hyperplanes coincide
7:     continue
8:   end if
9:    $\mathbf{P}_{\text{int}} \leftarrow \text{FindIntersectPoint}(\mathbf{W}_{\text{base}}, \mathbf{W}_{\text{inc}})$ 
10:   $\mathbf{V}_{\text{base}} \leftarrow \text{ComputeNormVector}(\mathbf{W}_{\text{base}})$ 
11:   $\mathbf{V}_{\text{inc}} \leftarrow \text{ComputeNormVector}(\mathbf{W}_{\text{inc}})$ 
12:   $\mathbf{V}_{\text{base}} \leftarrow \frac{\mathbf{V}_{\text{base}}}{\|\mathbf{V}_{\text{base}}\|_2}$   $\mathbf{V}_{\text{inc}} \leftarrow \frac{\mathbf{V}_{\text{inc}}}{\|\mathbf{V}_{\text{inc}}\|_2}$ 
13:   $\mathbf{V}_{\text{avg}} \leftarrow (1 - \alpha)\mathbf{V}_{\text{base}} + \alpha\mathbf{V}_{\text{inc}}$ 
14:   $\mathbf{W}_{\text{base}} \leftarrow \text{DefineHyperplane}(\mathbf{V}_{\text{avg}}, \mathbf{P}_{\text{int}})$ 
15:  /* Start hyperparameter optimization once the
16:   KPI window reaches its threshold size. */
17:   $\text{MB}_{\text{kpi}} \leftarrow \text{MbKPI}(\mathbf{X}_{\text{inc}}, \mathbf{y}_{\text{inc}}, \mathbf{W}_{\text{base}}, kpi)$ 
18:   $\text{AddItem}(L, \text{MB}_{\text{kpi}})$ 
19:   $\tau, \mu, \sigma, \text{low}, \text{high}, \text{DM} \leftarrow \text{MeasureKPIs}(L, z)$ 
20:   $\text{drift} \leftarrow \text{DetectDrift}(\mu, L(-1), \tau)$ 
21:  if  $\text{drift}$  then
22:     $\text{RemoveItem}(L, -1)$ 
23:     $\text{SM} \leftarrow \text{DefineScaleMap}(\mu, \text{low}, \text{high})$   $\triangleright$  SM is the scale map
24:     $\alpha' \leftarrow \text{TuneHyperparams}(\text{SM}, \text{DM})$ 
25:     $\mathbf{V}_{\text{avg}} \leftarrow (1 - \alpha')\mathbf{V}_{\text{base}} + \alpha'\mathbf{V}_{\text{inc}}$ 
26:     $\mathbf{W}_{\text{base}} \leftarrow \text{UpdateHyperplane}(\mathbf{V}_{\text{avg}}, \mathbf{P}_{\text{int}})$ 
27:  end if
28: end for
29: return  $\mathbf{W}_{\text{base}}$ 

```

Algorithm 3 OLR-WAA




Fig. 4 Limits - Gauss Dist. of KPI-Win on $KPI = R^2$




Fig. 5 Scale Map

In an illustrative scenario featuring 31 entries in KPI-Win, where each entry encompasses a set of KPIs related to the consumed mini-batch, the initial 30 entries are allocated for historical data. The final entry is designated for the current mini-batch, denoted as CUR-MB. The CUR-MB entry is transient and dynamically updates in response to concurrent computations, while the remaining entries represent our *Baseline Statistics*, serving as a point of reference for comparing performance changes. Figure 3 illustrates the

data structure utilized within OLR-WAA, where each entry represents a bag comprising one or more KPIs.

OLR-WAA structures historical KPIs within a normal distribution to facilitate the calculation of a dynamic threshold τ . This dynamic approach circumvents the limitations of fixed thresholds, which may not be universally applicable across different KPIs or data streams due to their diverse characteristics and evolving patterns. The threshold τ is determined using Equation 5, where σ represents the standard deviation of the distribution, and z is a scaling factor that typically corresponds to the number of standard deviations from the mean. The value of z is user-defined, with a commonly accepted range between 1.5 and 2.5. A lower z value, typically below 1.5, can be selected to identify subtle variations, and subsequently leading to the inference of a higher smoothing factor α , thereby ensuring a greater level of adaptation. Conversely, a higher z value adopts a more conservative approach, allowing for less sensitivity to minor fluctuations.

The proposed mechanism determines the optimal hyperparameter adjustments based on the detected drift magnitude, which is calculated using Eq. 8. In this context, μ_{KPI} represents the mean of the KPI-Window (KPI-Win) for the specified KPI, while $\text{CUR-MB}_{\text{KPI}}$ denotes the KPI value of the current data feed under evaluation. Furthermore, the mechanism establishes two critical limits, referred to as ‘low’ and ‘high,’ which are derived from the computed variable threshold using Eqs. 6 and 7, and as illustrated in Fig. 4. Data feeds that fall within these thresholds, between the mean and one of the boundaries, are regarded as minor deviations, indicative of incremental drift. Conversely, data feeds that exceed these thresholds signal significant deviations, identifying abrupt drift. Figure 5 presents an illustrative example of a dynamically constructed scale map, which is updated iteratively. Given that the range of the smoothing factor α is defined as $(0, 1]$, and the drift magnitude is quantified, the appropriate value of α can be inferred accordingly.

$$\text{Threshold } (\tau) = z \times \sigma \quad (5)$$

$$\text{low} = \mu_{\text{KPI}} - \tau \quad (6)$$

$$\text{high} = \mu_{\text{KPI}} + \tau \quad (7)$$

$$\text{DM} = \mu_{\text{KPI}} - \text{CUR-MB}_{\text{KPI}} \quad (8)$$

In certain scenarios, the nature of the learning task can be anticipated in advance, allowing for the selection of an appropriate adaptation strategy. For instance, if the task is *time-based*, the model is expected to continuously adapt to evolving patterns over time. Alternatively, if the task is *confidence-based*, the existing model is already well-established

and highly reliable, thereby carrying significant weight in decision-making. In such cases, employing a fixed smoothing factor, whether a high value (e.g., 1) or a low value (e.g., 0.1), may initially appear effective. A high smoothing factor places greater emphasis on recent data, enabling rapid adaptation to changes, whereas a low smoothing factor prioritizes historical knowledge, ensuring stability over time.

However, using a fixed smoothing factor presents two critical challenges. First, a high smoothing factor can result in *overfitting*, where the model becomes excessively sensitive to recent changes, potentially capturing noise instead of meaningful trends. Furthermore, it increases the risk of *catastrophic forgetting*, where the model fails to retain valuable past knowledge, leading to suboptimal long-term performance. Second, a low smoothing factor will result in very low adaptability in responding to rapid changes in the data stream, as it places greater emphasis on historical data while minimizing the influence of recent observations. This can lead to delayed adjustments in dynamic environments, reducing the model’s ability to effectively capture and react to sudden shifts or concept drifts. In response, the proposed OLR-WAA balances *adaptiveness* (responding to new changes) with *robustness* (preserving learned knowledge), ensuring responsiveness to evolving data while maintaining stability.

3.5 Theoretical Analysis of Convergence

3.5.1 Model Description

The OLR-WA model incrementally learns from streaming data by maintaining two hyperplanes: the base hyperplane $f_{\text{base}}(x)$, which encapsulates the accumulated knowledge from past observations, and the incremental hyperplane, $f_{\text{inc}}(x)$, which reflects the knowledge derived from the current mini-batch (instant feed) at time t , where x denotes the feature vector. The corresponding norm vectors of these boundaries, denoted as V_{base} and V_{inc} , respectively, are integrated using the Exponentially Weighted Moving Average (EWMA) technique. The dynamic weighting factor $\alpha \in (0, 1]$ serves as a smoothing parameter, assigning a weight of (α) to the incremental norm vector, emphasizing the contribution of the most recent data, while the remaining weight $(1 - \alpha)$ accounts for the historical knowledge retained in the base model. The combined norm vector V_{Avg} is iteratively updated, leveraging the intersection point P_{int} to redefine the base norm vector, with the objective of progressively aligning with V_{inc} as the number of observations increases.

3.5.2 Assumptions

1. **Linearity:** The relationship between the features and the target variable follows a linear model.
2. **Concept Stability [81]:** The underlying data distribution is subject to temporal evolution and may experience various forms of drift. However, it is assumed that each incoming mini-batch represents a coherent and stable concept, meaning that the data within the mini-batch pertains to a single, consistent concept rather than being fragmented across multiple evolving concepts. This assumption ensures that the model can effectively adapt to gradual changes without ambiguity.
3. **Boundedness of Data:** The input features and target variable are confined within a finite range, ensuring the model's stability and preventing unrealistic predictions beyond the defined input domain.
4. **Continuity:** Hyperplanes are continuous functions.

3.5.3 Convergence Criterion

Our objective is to demonstrate the asymptotic convergence of the OLR-WA model, specifically focusing on the intersection point and the weighted average norm vector, as the number of observations approaches infinity.

Step 1: Convergence of Incremental Norm Vector V_{inc} :

Under the assumption of concept stability [81], the norm vector V_{inc} of the new incoming mini-batches (instant feeds) is expected to adapt and approximate optimal values over time as $t \rightarrow \infty$, facilitating improved classification performance for any given instant feed at time t .

Step 2: EWMA-Based Integration and Computation of the Weighted Average Norm Vector:

Upon receiving the regression hyperplanes from both the base and incremental models at every time point t , where the base model encapsulates the past accumulated knowledge from all previous observations and the incremental model reflects the knowledge derived from the current mini-batch, the OLR-WA framework computes their corresponding norm vectors, denoted as V_{base} and V_{inc} . To effectively capture the evolving hyperplanes, OLR-WA combines these norm vectors to derive a weighted average norm vector, V_{Avg} , at each iteration t . This vector encapsulates both the direction and magnitude of the evolving classification model.

The integration process employs the Exponentially Weighted Moving Average (EWMA) approach, where the weighting factor α determines the relative importance of the most recent incremental mini-batches (instant feeds), while $1 - \alpha$ accounts for previously accumulated knowledge. The

weighting factor α serves as an adaptation parameter, allowing the model to balance the responsiveness to new information with the retention of previous learning.

The weighted average norm vector, V_{Avg} , is computed as follows:

$$V_{\text{Avg}} = \alpha \cdot V_{\text{inc}} + (1 - \alpha) \cdot V_{\text{base}}$$

Step 3: Convergence of OLR-WA Model

Convergence of V_{Avg} to V_{inc}

Let $\Delta_t = \|V_{\text{Avg}} - V_{\text{inc}}\|$ represent the norm of the difference between the weighted average norm vector and the incremental norm vector associated with the evolving hyperplanes. We aim to prove that $\Delta_t \rightarrow 0^+$ as $t \rightarrow \infty$.

Theorem 1 *Let t denote the number of iterations. Let α denote the smoothing parameter. As $t \rightarrow \infty$, the EWMA converges to the true mean or central tendency of the observed process.*

3.5.4 Proof Sketch

1. Recall that in the OLR-WA model, the weighted average norm vector V_{Avg} is obtained by combining the norm vectors of the base and incremental regression hyperplanes using Exponentially Weighted Moving Averages (EWMA) with dynamic weighting factor α .
2. Therefore, $V_{\text{Avg}} = \alpha \cdot V_{\text{inc}} + (1 - \alpha) \cdot V_{\text{base}}$.
3. The weighted average norm vector is computed iteratively using the EWMA formula to ensure an adaptive balance between recent and historical information.

By systematically incorporating information from all past observations, an iterative application of the EWMA formulation results in the following expansion:

$$\bar{X}_t = (1 - \alpha)\bar{X}_{t-1} + \alpha X_t \quad (9)$$

$$\bar{X}_t = (1 - \alpha)\bar{X}_{t-1} + \alpha X_t$$

$$\bar{X}_t = (1 - \alpha) [(1 - \alpha)\bar{X}_{t-2} + \alpha X_{t-1}] + \alpha X_t$$

$$\bar{X}_t = (1 - \alpha)^2 \bar{X}_{t-2} + \alpha(1 - \alpha)X_{t-1} + \alpha X_t$$

$$\bar{X}_t = (1 - \alpha)^2 [(1 - \alpha)\bar{X}_{t-3} + \alpha X_{t-2}] + \alpha(1 - \alpha)X_{t-1} + \alpha X_t$$

$$\bar{X}_t = (1 - \alpha)^3 \bar{X}_{t-3} + \alpha(1 - \alpha)^2 X_{t-2} + \alpha(1 - \alpha)X_{t-1} + \alpha X_t$$

$$\bar{X}_t = (1 - \alpha)^3 \bar{X}_{t-3} + \alpha(1 - \alpha)^2 X_{t-2} + \alpha(1 - \alpha) X_{t-1} + \alpha X_t$$

$$V_{\text{Avg}} = (1 - \alpha)^t V_{\text{base}} + \alpha \sum_{k=0}^{t-1} (1 - \alpha)^k V_{\text{inc}, t-k}$$

Through exponential smoothing, the assigned weights decrease exponentially as observations recede further into the past. The smoothing parameter α controls the rate of decay, with higher values favoring recent observations and lower values retaining historical influence over a longer period.

Applying the same principles to the OLR-WA model, it follows that as V_{Avg} approaches V_{inc} , the hyperplane defined by the intersection point converges, ensuring that $\Delta_t \rightarrow 0^+$.

3.6 Time and Memory Complexity

3.6.1 Time Complexity

OLR-WA utilizes Moore-Penrose pseudo-inverse once in the computation of the base hyperplane, and (T) times for the computation of the incremental hyperplanes, using the Moore-Penrose pseudo-inverse is a particularly common choice as it is polynomial time. Given X to be a N by M matrix, where N is the number of samples and M the number of features. The matrix multiplications each require $O(NM^2)$, while multiplying a matrix by a vector is $O(NM)$. Computing the inverse requires $O(M^3)$ in order to compute the LU or (Cholesky) factorization. Asymptotically, $O(NM^2)$ dominates $O(NM)$, so this calculation is ignored. Since the normal equation is used, the assumption that $N > M$ holds, otherwise the matrix $X^T X$ would be singular (and hence non-invertible), which means that $O(NM^2)$ asymptotically dominates $O(M^3)$. Therefore, the total complexity for the pseudo inverse is $O(NM^2)$.

OLR-WA processes data incrementally, enabling efficient computation by handling smaller data batches. Unlike the batch version, the online OLR-WA model achieves a lower time complexity, estimated at approximately $O(KM^2)$ per iteration, where K represents the number of samples in the mini-batch. This incremental approach not only reduces computational overhead but also ensures faster and more scalable processing, particularly in real-time or large-scale data environments.

This analysis presents the complexity in a “static context”; however, it is important to consider the practical implications. In practice, the batch model would require reprocessing all existing and new data points if encountering new data, resulting in a significantly higher time complexity compared to online models. This is primarily because the batch model’s computational burden increases

with the growing size of data points (N) when new data is introduced, while, in online, K remains fixed as a small mini-batch throughout.

OLR-WA method comprises several procedural stages, the computational complexities of which are disregarded, as it is dominated by the pseudo-inverse computation $O(KM^2)$. These stages are succinctly outlined as follows:

1. **Hyper-plane Coincidence Check:** The initial step involves confirming the alignment of two hyper-planes. This verification process, entailing iterations solely over specified norm vectors, demonstrates a time complexity of $O(M)$, where M signifies the dimensionality of the space.
2. **Linear Equation System Solution:** The resolution of a linear equation system can be achieved through various methodologies, such as Gaussian elimination. In the context of this method, which necessitates solving a system with two equations, the highest order time complexity is $O(N^3)$ [82], where N is the number of equations, since OLR-WA consists of only two equations, solving a system of two linear equations will be constant.
3. **Weighted Midpoint Determination (Parallel Planes):** For scenarios involving parallel planes, the computation of the weighted midpoint involves fundamental mathematical operations, consequently resulting in a constant time complexity.
4. **Defining a New Hyper-plane:** The process of defining a new hyper-plane entails M iterations over the corresponding norm vector, thereby giving rise to a time complexity of $O(M)$, where M is the number of dimensions.

OLR-WAA integrates the aforementioned procedures without adaptive optimization mechanism comprising four core components: (1) drift detection, (2) drift magnitude quantification, (3) hyperparameter optimization, and (4) hyperparameter inference for adaptation. These components are executed sequentially for each incoming mini-batch and are succinctly outlined as follows:

1. **KPI Computation:** The evaluation of performance metrics such as R^2 and MSE typically necessitates computing predictions of the form $\hat{y}_i = \mathbf{x}_i^T \mathbf{w}$, for $i = 1, \dots, K$, which involves a total computational complexity of $O(KM)$.
2. **KPI-Win Insertion:** The computed KPIs are appended to the KPI-Win with a constant time complexity of $O(1)$.
3. **Statistical Measures:** This step involves computing the mean μ and standard deviation σ of the KPI-Win of length $|L|$, which requires $O(|L|)$ time. Additionally, the

calculation of the threshold, as well as the lower and upper limits, involves standard operations with constant time complexity $O(1)$.

4. **Drift Magnitude Quantification:** The deviation between the current KPI and the historical mean is calculated with a constant time complexity of $O(1)$.
5. **Scale Map:** The interval between the KPI-Win mean and the lower or upper bound is partitioned into a fixed number of equal-width regions. Each region is then mapped to a corresponding subinterval within the full hyperparameter range, enabling scaled adaptation based on the degree of deviation. This procedure involves constant-time arithmetic and indexing, yielding a time complexity of $O(1)$.
6. **Hyperparameter Inference:** Based on the position of the current KPI within the scale map, this step identifies and returns the corresponding hyperparameter value. The operation has a time complexity of $O(|L|)$.
7. **Model Adjustment via Inferred Hyperparameter:** This step leverages all short-lived information from the current mini-batch, re-executes the EWMA update using the newly inferred hyperparameter and define the updated hyperplane, with an overall time complexity bounded by $O(M)$.

In summary, the optimization mechanism in OLR-WAA has a time complexity bounded by $O(KM)$ and does not alter the overall complexity of OLR-WA, which remains dominated by the pseudo-inverse operation with a complexity of $O(NM^2)$.

3.6.2 Memory Complexity

The memory footprint of OLR-WAA is primarily determined by the model parameters and the KPI-Win buffer used for performance monitoring and drift analysis. The model maintains a weight vector $w \in \mathbb{R}^M$, where M is the number of input features, requiring $O(M)$ space. In addition, the KPI-Win buffer stores the most recent $|L|$ performance entries, where each entry contains s scalar KPIs. Here, s denotes

the number of evaluation metrics tracked per mini-batch, such as the mean squared error (MSE) and the coefficient of determination (R^2), which corresponds to $s = 2$ in typical configurations. This results in a total memory complexity of $O(M + |L| \cdot s)$, which remains minimal and scales well in streaming environments, making the approach suitable for deployment in high-throughput streaming environments.

3.7 Evaluation Metric

The coefficient of determination (R^2) is preferred in regression scenarios because it offers bounded and interpretable values, making it more informative than other common metrics such as MSE, MAE, RMSE, and MAPE [83]. Specifically, R^2 values range from negative infinity to 1, with values close to 1 indicating a strong predictive model. Unlike MSE or RMSE, which are unbounded and sensitive to the scale of the target variable, R^2 provides a normalized measure of fit that reflects the proportion of variance in the dependent variable explained by the independent variables. This scale-invariance allows for easier comparison of model performance across datasets and domains [83]. Moreover, R^2 generates high scores only when the majority of ground truth elements are predicted accurately, reflecting both the explained variance and the distribution of the data. As the most widely used measure of model fit in linear regression [8], R^2 is often the default metric in statistical software and machine learning libraries. Its interpretability, statistical grounding [84], and widespread adoption across scientific disciplines justify its recommendation as a standard evaluation criterion for regression models.

However, in one of our experiments, which was targeted to compute the average rank of models across multiple datasets in preparation for potential statistical significance testing, we selected the Mean Squared Error (MSE) as the evaluation metric. In addition to being the default cost function in regression [15], MSE was chosen due to its compatibility with parametric statistical tests that rely on assumptions of normally distributed errors, as commonly assumed in classical regression settings [8]. Unlike the coefficient of

Table 1 OLR-WA Sensitivity to Hyperparameters on Normal Regression Summary of Performance Measures using R^2 on the Last Iteration

Datasets	Algorithms				
	Equal Weights $\alpha=0.5$	Favor New Data $\alpha=0.9$	Radically Favor New data $\alpha=0.995$	Favor Existing Data $\alpha=0.1$	Radically Favor Existing Data $\alpha=0.005$
DS1 [11]	0.97423	0.97154	0.97064	0.97612	0.97590
DS2 [11]	0.98260	0.98027	0.97947	0.98402	0.98389
DS3 [11]	0.98204	0.98018	0.97952	0.98193	0.97792
DS4 [11]	0.92482	0.91634	0.91326	0.92911	0.92144
MCPD [85]	0.73125	0.69246	0.67260	0.74188	0.72409
1KC [86]	0.90773	0.90545	0.90449	0.90646	0.90474
KCHSD [87]	0.54048	0.45209	0.41781	0.57392	0.57453
CCPP [88]	0.92202	0.90869	0.90279	0.92785	0.92841

Table 2 OLR-WA Sensitivity to Hyperparameters on Adversarial Scenarios DS5, DS6 for Time-Based, and DS7, DS8 for Confidence-Based Summary of Performance Measures using R^2 on the Last Iteration

Datasets	Algorithms				
	Equal Weights $\alpha=0.5$	Favor New Data $\alpha=0.9$	Radically Favor New data $\alpha=0.995$	Favor Existing Data $\alpha=0.1$	Radically Favor Exist- ing Data $\alpha=0.005$
DS5 [11]	0.98601	0.98513	0.98484	N/A	N/A
DS6 [11]	0.60221	0.93739	0.93535	N/A	N/A
DS7 [11]	N/A	N/A	N/A	0.07826	0.97815
DS8 [11]	N/A	N/A	N/A	0.58033	0.93191

determination (R^2), which is bounded in $(-\infty, 1]$ and relative to the variance of the response variable, MSE serves as an absolute measure of predictive error. It consistently penalizes large deviations and preserves cardinal information about model performance [83]. These properties make MSE especially suitable for reliably ranking models across heterogeneous datasets, particularly when differences in performance are subtle or when dataset scales vary.

4 Hyperparameters Sensitivity Analysis

In this section we study the sensitivity of OLR-WA to its hyperparameter (α) across three distinct usage types of OLR-WA: Normal Online Linear Regression, Time-Based Adversarial Scenarios, and finally Confidence-Based Adversarial Scenarios.

4.1 Normal Linear Regression Scenarios

Table 1 presents the performance results for each normal linear regression dataset, across diverse hyperparameter configurations. The findings indicate that OLR-WA demonstrates minimal sensitivity to the hyperparameter (α) within normal linear regression scenarios.

4.2 Adversarial Scenarios

Table 2 presents the performance results for DS5, and DS6 for Time based adversarial Scenarios, and DS7, and DS8 for Confidence-Based Adversarial Scenarios, cross diverse hyperparameter configurations. The findings indicate that OLR-WA demonstrates high sensitivity to the hyperparameter (α) within adversarial scenarios.

However, the conducted ablation study evaluates results on the test data after processing all data points, which is important but may not capture the model's sensitivity to internal fluctuations within each encountered mini-batch. To address this, the experimental section includes a dedicated part for concept drift experiments during data streams.

5 Experiments

This section presents experiments on synthetic and real public datasets, employing a rigorous methodology for robustness. To reduce randomness, results are averaged over 5 random seeds and evaluated using 5-fold cross-validation. Each reported R^2 value is based on 25 experiments with different data splits, ensuring stability and consistency.

In our experimentation, specific guidelines were diligently adhered to, guaranteeing an impartial evaluation of all models and preventing any form of bias.

1. Initialization - During the initialization phase, all models' weights were set to an array of zeros, establishing a uniform and consistent starting point for each model.
2. Feature Engineering - In the process of feature engineering, caution was exercised to limit any potential undue impact resulting from feature manipulation. Consequently, minimal feature engineering techniques, such as normalization and one-hot encoding, were employed to ensure the preservation of data integrity and maintain a balanced evaluation of the models.
3. Hyperparameters Tuning - Hyperparameters were meticulously fine-tuned for each model and experiment, taking into account factors such as data set size and the number of features. The optimal outcomes for each model were carefully identified and subsequently reported.
4. Mini-Batch Size (K) - For models utilizing mini-batches, such as MBGD and OLR-WA, the mini-batch size (K) was determined using the formula provided in Eq. 10. The mini-batch size was set to either four times the number of dimensions in the dataset or a user-defined value, whichever was larger. In our experiments, the user-defined value was deliberately kept low to prioritize the formula-based size. This approach was adopted to strike a balance between efficient computation and model performance during our experimentation.
5. Epochs Selection:
 - (a) For SGD, ORR, OLR online models, the appropriate hyperparameter was determined using the formula $N \times \mathbb{Z}^+$, where $\mathbb{Z}^+$ is a Natural Number > 1 .

Table 3 Summary of the Main Properties of the Datasets Considered. Used for 1st Experiment: Performance Analysis on Normal Regression Scenarios

Dataset	Type	Datapoints	Dimensions	Noise	Training (5-Folds)	Testing (5-Folds)	Measurement Focus
DS1 [11]	Synthetic	1000	3	10	800	200	Normal Regression
DS2 [11]	Synthetic	10,000	20	20	8000	2000	Normal Regression
DS3 [11]	Synthetic	10,000	200	25	8000	2000	Normal Regression
DS4 [11]	Synthetic	50,000	500	50	40,000	10,000	Normal Regression
MCPD [85]	Real	1338	7	N/A	1070	267	Normal Regression
1KC [86]	Real	1000	5	N/A	800	200	Normal Regression
KCHSD [87]	Real	21,613	21	N/A	17,290	4322	Normal Regression
CCPP [88]	Real	9568	5	N/A	7654	1914	Normal Regression
DS5 [11]	Synthetic	5000	20	20	4500	500	Time-Based
DS6 [11]	Synthetic	10,000	200	40	9000	1000	Time-Based
DS7 [11]	Synthetic	5000	20	20	4500	500	Confidence-Based
DS8 [11]	Synthetic	10,000	200	40	9000	1000	Confidence-Based
DS9 [11]	Synthetic	1000	2	20	800	200	Convergence
DS10 [11]	Synthetic	1000	2	40	800	200	Convergence

(b) For mini-batch models, the formula $\frac{N}{K} \times \mathbb{Z}^+$ was utilized. Here, N refers to the number of training samples, M represents the data-set dimensions, and $\mathbb{Z}^+$ is a natural number. $\mathbb{Z}^+$ was systematically increased until no further improvement in the coefficient of determination (R^2) was observed.

$$K \text{ (Mini-Batch Size)} = \max(U, (M \times \mathbb{Z}^+)) \quad (10)$$

5.1 Performance Analysis on Stationary Linear Regression Scenarios

A comparative analysis was conducted to evaluate the performance of different online regression models, employing synthetic and real-world publicly available data-sets. This experiment, utilizing OLR-WA, was conducted under stationary regression settings where no concept drift is present. The evaluation utilized 5-fold cross-validation and seed averaging, with the R^2 metric measured at the end of the last iteration. To provide a concise overview of our experimental setup, Table 3 presents key properties of the specific data-sets employed in this study.

Performance Evaluation Table 4 presents the performance results for each data-set, along with the hyperparameters utilized for each model. The inclusion of these hyperparameters enhances transparency and facilitates a comprehensive analysis of the models' performance. Below are some observed insights about the performance evaluation for this experiment:

1. The batch model, depicted in the first column, serves as a benchmark for evaluating performance and sets the standard for all online models.

2. OLR-WA performance is remarkably significant and very close to the batch model, the difference is very slight and almost in the 3'd digit after the decimal point across all data-sets, except the 1KC data-set, where the batch model performance is 0.93615 while OLR-WA is 0.90773.
3. RLS performance significantly deteriorates on the more challenging DS4 data-set, where it achieves an R^2 value of approximately 0.62667. Additionally, it demonstrates a low R^2 score of 0.66197 on the CCPP data-set.
4. PA has demonstrated its proficiency in capturing underlying patterns and relationships effectively, as evidenced by competitive R^2 values across various datasets. However, its performance on the 1KC dataset appears to be somewhat limited, with an R^2 value of approximately 0.78565, in contrast to the batch model's 0.93615. Similarly, on the CCPP dataset, PA achieved an R^2 value of 0.66327, notably lower than the batch model's 0.92855.
5. The Widrow-Hoff algorithm exhibits good performance in most of the data-sets, although it experienced some degradation with an R^2 value of 0.60772 in the 1KC data-set, whereas the batch model R^2 value is 0.93615.

In conclusion, the application of OLR-WA with its default hyperparameter setting of $\alpha = 0.5$ consistently delivered top-tier performance across all evaluated data-sets.

5.2 Performance Analysis on Adversarial Scenarios

This experiment aims to understand the model's behavior under varying settings of the hyperparameter α in adversarial data environments. A high α value is used in Time-Based scenarios to enable rapid adaptation to recent patterns, while a low α is used in Confidence-Based scenarios to emphasize

Table 4 OLR-WA Performance Analysis on Normal Regression Scenarios Summary of Performance Measures using R^2 on the Last Iteration t_s : time in seconds, E: epochs, $\mathbb{Z}^+$: Natural Number ≥ 1

Datasets	Algorithms									
	Batch Regression (Pseudo-Inverse)	SGD $\eta = 0.01$ [DS3, DS3 =0.01, KCHSD 0.0001], $E=N \times$ ($\mathbb{Z}^+ = 2$) [IKC $\mathbb{Z}^+ = 3$, CCPP $\mathbb{Z}^+ = 5$]	MBG $\eta = 0.01$ [DS4=0.001], $K=M \times (\mathbb{Z}^+ = 5)$ $E = \frac{N}{K} \times (\mathbb{Z}^+ = 10)$ [DS1 $\mathbb{Z}^+ = 5$, DS3 $\mathbb{Z}^+ = 100$, KCHSD $\mathbb{Z}^+ = 20$, IKC, CCPP $\mathbb{Z}^+ = 40$]	LMS $\eta = 0.01$ [IKC=0.001, KCHSD=0.0001]	ORR $\eta = 0.01$ [DS3, DS4, KCHSD, MCPD=0.001], $E=N \times (\mathbb{Z}^+ = 2)$ [IKC $\mathbb{Z}^+ = 3$, KCHSD, CCPP $\mathbb{Z}^+ = 5$] $\lambda = 0.1$ [IKC, [MCPD, KCHSD, CCPP=0.001]]	OLR $\eta = 0.01$ [DS3, DS4, MCPD, KCHSD= 0.001], $E=N \times (\mathbb{Z}^+ = 2)$ [IKC $\mathbb{Z}^+ = 3$, KCHSD, CCPP $\mathbb{Z}^+ = 5$], $\lambda = 0.1$ [MCPD, CCPP=0.01], [IKC, KCHSD=0.001], [KCHSD=0.0001]]	RLS $\lambda = .99$ $\delta = 0.01$	PA $C = 0.01$ $\epsilon = 0.01$, [DS1, DS2, DS3 $C = 0.1$, $\epsilon = 0.1$]	OLR-WA $\alpha = .5$ [KCHSD $\alpha = .1$, $\frac{1}{B}K = N \times (\mathbb{Z}^+ = 0.1)$ $K=M \times (\mathbb{Z}^+ = 5)$]	
DS1 [11]	$r^2: 0.97637$ $t_s: 0.00121s$	$r^2: 0.97561$ $t_s: 0.01391s$	$r^2: 0.97580$ $t_s: 0.00434s$	$r^2: 0.97616$ $t_s: 0.00669s$	$r^2: 0.96760$ $t_s: 0.01958s$	$r^2: 0.97565$ $t_s: 0.02026s$	$r^2: 0.97624$ $t_s: 0.01073s$	$r^2: 0.97378$ $t_s: 0.00866s$	$r^2: 0.97423$ $t_s: 0.01185s$	
DS2 [11]	$r^2: 0.98418$ $t_s: 0.01431s$	$r^2: 0.98231$ $t_s: 0.14151s$	$r^2: 0.98384$ $t_s: 0.01453s$	$r^2: 0.98243$ $t_s: 0.06802s$	$r^2: 0.96987$ $t_s: 0.19516s$	$r^2: 0.97959$ $t_s: 0.20398s$	$r^2: 0.98268$ $t_s: 0.12351s$	$r^2: 0.97642$ $t_s: 0.05957s$	$r^2: 0.98260$ $t_s: 0.03571s$	
DS3 [11]	$r^2: 0.98299$ $t_s: 0.08145s$	$r^2: 0.98116$ $t_s: 0.15012s$	$r^2: 0.98290$ $t_s: 0.24452s$	$r^2: 0.98161$ $t_s: 0.07244s$	$r^2: 0.96901$ $t_s: 0.18893s$	$r^2: 0.97933$ $t_s: 0.20211s$	$r^2: 0.96001$ $t_s: 5.14062s$	$r^2: 0.96705$ $t_s: 0.09319s$	$r^2: 0.98204$ $t_s: 0.39293s$	
DS4 [11]	$r^2: 0.92973$ $t_s: 49.89187s$	$r^2: 0.90663$ $t_s: 5.95597s$	$r^2: 0.92964$ $t_s: 156.83463s$	$r^2: 0.90715$ $t_s: 5.02476s$	$r^2: 0.85535$ $t_s: 1.92488s$	$r^2: 0.86412$ $t_s: 3.93231s$	$r^2: 0.62667$ $t_s: 1907.88762s$	$r^2: 0.87255$ $t_s: 3.80154s$	$r^2: 0.92464$ $t_s: 8.08490s$	
MCPD [85]	$r^2: 0.74321$ $t_s: 0.00332s$	$r^2: 0.73092$ $t_s: 0.02735s$	$r^2: 0.73007$ $t_s: 0.00905s$	$r^2: 0.73637$ $t_s: 0.01364s$	$r^2: 0.74070$ $t_s: 0.02973s$	$r^2: 0.74217$ $t_s: 0.02825s$	$r^2: 0.73830$ $t_s: 0.02185s$	$r^2: 0.72893$ $t_s: 0.00650s$	$r^2: 0.74080$ $t_s: 0.00358s$	
IKC [86]	$r^2: 0.93615$ $t_s: 0.00111s$	$r^2: 0.92930$ $t_s: 0.01471s$	$r^2: 0.91195$ $t_s: 0.01702s$	$r^2: 0.60772$ $t_s: 0.00470s$	$r^2: 0.90294$ $t_s: 0.02493s$	$r^2: 0.90322$ $t_s: 0.02582s$	$r^2: 0.85503$ $t_s: 0.00807s$	$r^2: 0.78565$ $t_s: 0.00479s$	$r^2: 0.90773$ $t_s: 0.00606s$	
KCHSD [87]	$r^2: 0.57859$ $t_s: 0.02213s$	$r^2: 0.57277$ $t_s: 0.31353s$	$r^2: 0.57431$ $t_s: 0.17558s$	$r^2: 0.56344$ $t_s: 0.11620s$	$r^2: 0.56432$ $t_s: 0.88060s$	$r^2: 0.57040$ $t_s: 0.95015s$	$r^2: 0.48088$ $t_s: 0.18357s$	$r^2: 0.53533$ $t_s: 0.12020s$	$r^2: 0.57395$ $t_s: 0.08367s$	
CCPP [88]	$r^2: 0.92855$ $t_s: 0.00807s$	$r^2: 0.91838$ $t_s: 0.22912s$	$r^2: 0.91989$ $t_s: 0.29177s$	$r^2: 0.89785$ $t_s: 0.04680s$	$r^2: 0.92550$ $t_s: 0.39269s$	$r^2: 0.92611$ $t_s: 0.58023s$	$r^2: 0.66197$ $t_s: 0.07742s$	$r^2: 0.66327$ $t_s: 0.04727s$	$r^2: 0.92202$ $t_s: 0.05153s$	

Table 5 OLR-WA Performance Analysis on Adversarial Scenarios Summary of Performance Measures using R^2 on the Last Iteration N/A: Minus R^2 , E =Epochs, Z^+ : Natural Number ≥ 1

Datasets	Algorithms	SGD η = [DS5, DS7 = 0.01, DS6, DS8=0.001], $E=N \times (Z^+ = 2)$	MBGD η = 0.01, $K=M \times (Z^+ = 5)$, $E=\frac{N}{K} \times (Z^+ = 2)$	LMS η = [DS5, DS7 = 0.01, DS6, DS8=0.001], $E=N \times (Z^+ = 2)$	ORR η = [DS5, DS7 = 0.01, DS6, DS8 = 0.001], $E=N \times (Z^+ = 2)$, $\lambda = [DS5, DS7 = 0.01]$	OLR η = [DS5, DS7 = 0.01, DS6, DS8 = 0.001], $E=N \times (Z^+ = 2)$, $\lambda = [DS5, DS7 = 0.01, DS6, DS8 = 0.01]$	RLS λ = [DS5, DS6 = 99, DS7, DS8 = 99], $\lambda = 0.01$	PA C = 0.1, $\epsilon = 0.1$, $[DS6C = .01, \epsilon = .01]$	OLR-WA α = [DS5, DS6 = 95, DS7, DS8 = 0.005], $BK = N \times 0.1$, $K=M \times (Z^+ = 5)$
DS5 [11]	N/A	N/A	N/A	0.98528	N/A	N/A	0.98546	0.97812	0.98498
DS6 [11]	N/A	N/A	N/A	0.93810	N/A	N/A	0.87134	0.91145	0.93634
DS7 [11]	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	0.97815
DS8 [11]	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	0.93191

stability and trust in the existing model. This setup employs OLR-WA, which involves manual tuning of α , and highlights the need for dynamic adjustment, a feature introduced in OLR-WAA.

As summarized in Table 3, datasets DS5 through DS8 were used for the experiment. DS5 and DS6 correspond to time-based scenarios, where the model is expected to prioritize recent observations following a distributional shift. DS7 and DS8 represent confidence-based scenarios, evaluating the model's ability to maintain performance by relying on historically consistent patterns.

Performance Evaluation Table 5 presents the performance results for each data-set, along with the hyperparameters utilized for each model. Based on the experimental findings of DS5 and DS6, specifically designed for Time-Based scenarios, it becomes evident that the model's performance improves when it can effectively accommodate new data. Notably, OLR-WA, Widrow-Hoff, PA, and RLS algorithms have demonstrated remarkable outcomes in terms of their ability to adapt to dynamic changes in the data. In contrast, the performance results of other online models, namely SGD, MBGD, Online Ridge, Online Lasso, and even the standard batch model, were notably unsatisfactory within this specific scenario.

The experimental results obtained from DS7 and DS8, which are tailored for Confidence-Based scenarios, highlight the significance of the model's capability to effectively incorporate previous data for improved performance. Particularly, the OLR-WA algorithm has exhibited impressive outcomes by consistently adapting to changes in data patterns while maintaining a conservative approach. Conversely, the performance of all other online models proved to be notably inadequate within this particular scenario.

The accomplishments of the Confidence-Based OLR-WA algorithm hold significant relevance across various scenarios, such as sentiment analysis tasks, where certain labeled data points have undergone verification by experts or trusted sources. In such cases, assigning higher weights to these points based on their elevated confidence levels can prove advantageous. A real case study of the applicability of linear regression on sentiment analysis is conducted by the authors in [89]. Another example, similarly consider the example of an Amazon store with an extensive product inventory and a daily influx of hundreds of new products. In this context, employing a Confidence-Based approach, which inherently favors the larger existing product pool, or products that have been selected for trustworthiness, can potentially enhance the model's prediction performance.

As a key highlight of our visual experiment, Fig. 6a and b illustrate the behavior of OLR-WA in Time-Based and Confidence-Based scenarios. Figure 6a depicts a Time-Based adversarial setting, where OLR-WA, with its

Fig. 6 Comparison of two adaptation strategies





Fig. 7 Convergence analysis models cost by number of data points considered on DS10 for multiple seeds, 5-Fold cross validation

hyperparameter set to $\alpha = 0.95$, successfully adapted to a newly introduced negatively correlated adversarial model. In contrast, Fig. 6b presents a confidence-based scenario, in which $\alpha = 0.05$, fine-tuning enabled OLR-WA to preserve confidence in existing data, minimizing the impact of new data on its performance. These findings are significant because they provide valuable insight into automatic weight selection and adaptive hyperparameter tuning, particularly in the presence of concept drift or adversarial conditions.

5.3 Convergence Analysis

This experiment provides a comprehensive comparative analysis of the convergence behavior of online regression models, with a particular focus on OLR-WA. The evaluation examines OLR-WA's stability and efficiency relative to other models, employing the R^2 metric to assess predictive performance and Mean Squared Error (MSE) to analyze convergence trends, both monitored iteratively. To ensure

robustness and generalizability, the analysis leverages 10 datasets with diverse characteristics, capturing a wide range of data distributions and complexities to thoroughly evaluate convergence dynamics.

Initially, reference datasets DS9 and DS10 [11], characterized by reduced dimensionality, are employed to constrain information, thereby challenging model performance while enabling detailed convergence monitoring with smaller data batches. Subsequently, the models were assessed on a broader set of datasets with varying dimensionality and complexity to evaluate their generalization capabilities and convergence behavior under diverse data conditions.

Figure 7 illustrates the convergence behavior of various algorithms based on MSE using DS10 [11]. Most algorithms reach optimal performance around epoch 150. However, OLR-WA exhibits a distinct convergence pattern, starting with a significantly lower cost and continuously sustaining this diminished cost throughout. This early and consistent cost advantage establishes one of our key contributions in

Table 6 Convergence Analysis Models Performance by Number of Data Points Considered on DS9 for Multiple Seeds, 5 Folds Cross Validation Hyperparameters are the same as Fig. 7

Model	Number of Training Data Points															
	10	20	30	40	50	60	70	80	90	100	110	120	130	140	150	Final
OLR-WA	0.86265	0.87026	0.87771	0.88329	0.88861	0.89288	0.89591	0.89875	0.90098	0.90161	0.90433	0.90615	0.90810	0.90986	0.91084	0.91293
SGD	0.32402	0.50374	0.62607	0.72745	0.78288	0.82651	0.85521	0.87468	0.88738	0.89584	0.90132	0.90371	0.90599	0.90806	0.90799	0.91233
MBGD	0.32226	0.51114	0.64383	0.74555	0.79964	0.83936	0.85881	0.87445	0.88487	0.89264	0.89909	0.90194	0.90535	0.90592	0.90829	0.91177
LMS	0.25216	0.48899	0.66102	0.72896	0.78882	0.82736	0.84889	0.86801	0.88491	0.89585	0.90026	0.90347	0.90625	0.90919	0.91053	0.91348
ORR	0.28438	0.44982	0.59813	0.69093	0.74606	0.78599	0.81481	0.84090	0.85622	0.86723	0.88113	0.88760	0.89168	0.89530	0.89552	0.90296
OLR	0.27546	0.51746	0.64119	0.72773	0.77723	0.81353	0.84392	0.86510	0.88134	0.89082	0.89761	0.90173	0.90476	0.90693	0.90769	0.91299
RLS	0.13441	0.30085	0.43636	0.49871	0.57590	0.63265	0.67494	0.71707	0.75557	0.78608	0.80571	0.82775	0.83981	0.85533	0.87163	0.91513
PA	0.76154	0.88878	0.89449	0.90396	0.86813	0.89807	0.90726	0.91309	0.90740	0.90570	0.90363	0.90456	0.90786	0.90941	0.90577	0.90626

Table 7 Convergence Analysis Summary of Performance Measures Using MSE on the First 10 Iterations E : Epochs, K : Mini-batch Size, Z^+ : Natural Number ≥ 1

Datasets	Algorithms																
	SGD	MBGD	LMS	ORR	OLR	RLS	PA	OLR-WA									
	$\eta = 0.01[\text{DS3, DS4}$ $=0.001,$ $E = \frac{N}{K} \times (\mathbb{Z}^+ = 5)$ [DS3, DS4 $\mathbb{Z}^+ = 100, \text{KCHSD}\mathbb{Z}^+ = 20,$ [CCPP $\mathbb{Z}^+ = 5]$	$\eta = 0.01[\text{KC} = 0.1],$ $E = \frac{N}{K} \times (\mathbb{Z}^+ = 5)$ [DS3, DS4 $\mathbb{Z}^+ = 100, \text{KCHSD}\mathbb{Z}^+ = 20,$ [KC, CCPP $\mathbb{Z}^+ = 40]$	$\eta = 0.01[\text{DS3, DS4} = 0.001,$ $\eta = 0.01[\text{DS3, DS4, KCHSD,}$ $=0.001], E = N \times (\mathbb{Z}^+ = 2)$ [$\text{KC}\mathbb{Z}^+ = 3, \text{KCHSD, CCPP}$ $\mathbb{Z}^+ = 5], \lambda = 0.1[\text{KC} = 0.01,$ MCPD, KCHSD, C CCPP = 0.01]	$\eta = 0.01[\text{DS3, DS4, KCHSD,}$ $=0.001], E = N \times (\mathbb{Z}^+ = 2)$ [$\text{KC}\mathbb{Z}^+ = 3, \text{KCHSD, CCPP}$ $\mathbb{Z}^+ = 5], \lambda = 0.1[\text{KC} = 0.01,$ MCPD, KCHSD, C CCPP = 0.01]	$\eta = 0.01[\text{DS3, DS4, KCHSD,}$ $=0.001],$ $E = N \times (\mathbb{Z}^+ = 2)$ [$\text{KC}\mathbb{Z}^+ = 3, \text{KCHSD, CCPP}$ $\lambda = 0.1[\text{KC,}$ CCPP = 0.01], [KCHSD = 1e-4]	$\lambda = 0.01[\text{DS3, DS4,}$ $[\text{KC}\delta = 0.1]$ $\alpha = 0.001],$ $C = 0.01$ $\lambda = 0.1[\text{KC,}$ CCPP = 0.01], [KCHSD = 1e-4]	$\lambda = 0.01, \epsilon = 0.1, [\text{KCHSD}$ DS6, DS7, DS8 $C = 0.01$ $\epsilon = 0.01]$	$\alpha = .5$ [KCHSD $\alpha = .9$ $C = 0.01$ $\epsilon = 0.01]$									
Name	K	MSE	rank	MSE	rank	MSE	rank	MSE	rank	MSE	rank	MSE	rank				
DS1 [11]	10	1865.20581	4	5976.67276	8	1778.39859	3	1961.97400	6	1933.74953	5	3080.20632	7	541.27625	2	455.90216	1
DS2 [11]	30	2430.97513	5	23760.09556	8	2390.53622	4	2825.94930	6	2122.09717	3	4715.74581	7	1564.91181	2	841.70463	1
DS3 [11]	300	3282.66236	6	31822.14283	8	3275.53508	5	3887.32562	7	3218.66053	4	1593.63189	2	2349.97800	3	1165.10646	1
DS4 [11]	1,000	5550.43115	5	31435.32987	8	5547.42274	4	5658.62652	6	5429.84169	3	12569.60925	7	5407.04210	2	3565.32212	1
MCPD [85]	25	0.33973	8	0.85781	8	0.31242	1	0.33924	3	0.33995	5	0.36409	7	0.36793	7	0.31687	2
KCH [86]	39	0.00362	6	0.00709	8	0.00138	2	0.00358	4	0.00359	5	0.00151	3	0.00464	7	0.00097	1
KCHSD [87]	30	0.88983	7	0.69907	6	0.89519	8	0.59394	5	0.57859	4	0.49323	1	0.50757	3	0.50742	2
CCPP [88]	20	0.05089	4	0.18292	7	0.23790	8	0.05065	3	0.05058	2	0.09088	6	0.07098	5	0.00407	1
Average		1,641.32	5.125	11,624.50	7.625	1,624.16744	4.375	1,791.86	5	1,588.17	3.875	2,745.02	4.875	1,233.02	3.875	753.61	1.250




Fig. 8 Concept drift types [76]

Finally, another thorough convergence experiment was conducted. Table 7 presents the performance of the models by evaluating the MSE over the initial ten mini-batches. We utilized the same number of mini-batches (K) at each iteration (t) to ensure consistency and comparability across different models. The experiment evaluates the average rank of each model's performance across various datasets, assigning a rank of 1 to the top-performing model. Among the models assessed, OLR-WA model achieves the lowest average rank of 1.25. This finding underscores the significance

Table 8 Dataset Properties. Drift Locations: Abrupt (A) marks sudden drift onset. Incremental (I) drift recurs every K points. Gradual (G) shows drift as a sequence of concepts C_1 and C_2 . E.D.: Euclidean distance between coefficients in space. C-Start: Start Concept. C-End: End Concept. E.D. (Consec.): E.D. between each two consecutive concepts in sequence

Dataset	Drift Type	Data-points	Dimensions	Noise	Drift Locations	E.D. (C-Start to C-End)	E.D. (Consec.)
DS11 [11]	A	10k	10	20	5k	327.23	327.23
DS12 [11]	I	10k	10	20	every 1k	1205.10	[14.8, 18.5, 23.9, 31.8, 44.6, 66.9, 111.5, 223.1, 669.5]
DS13 [11]	G	10k	10	20	[2.5k-c1, 1k-c2, 1k-c1, 2k-c2, 1k-c1, 2.5k-c2]	2049.14	[1158.3, 1205.1, 1885.4, 1885.4, 2167.7]

this work, showing that OLR-WA not only adapts quickly but also maintains superior performance over time. As shown in Table 6, OLR-WA achieved the highest R^2 score of 0.86265 within the first 10 data points and sustained high R^2 values from the 90th data point onward. PA also demonstrated rapid convergence, stabilizing at a high R^2 by the 70th data point, while the majority of models converged between the 120th and 150th data points.

of further investigating the statistical implications of online regression models.

5.4 Performance Analysis on Concept Drift Datasets.

In this section, four top-performing models are evaluated across diverse types of concept drift, as illustrated in Fig. 8, using the concept drift datasets summarized in Table 8.




Fig. 9 Models performance analysis on abrupt drift dataset DS11




Fig. 10 OLR-WAA performance analysis on incremental drift dataset DS12




Fig. 11 OLR-WAA performance analysis on gradual drift dataset DS13

This experiment employs OLR-WAA, the specialized variant designed to handle concept drift through the automatic inference of the hyperparameter α .

5.4.1 Performance Analysis on Abrupt Drift datasets

Performance Evaluation: Figure 9 illustrates the performance of OLR-WAA, PA, RLS, and LMS on abrupt drift datasets, measured using the R^2 metric. The yellow region indicates the onset of a new drifted concept starting at data point 5000. OLR-WAA demonstrates superior robustness to abrupt drift, while PA, RLS, and LMS exhibit varying degrees of sensitivity, with RLS being the most affected.

It is important to note that RLS includes a hyperparameter $\lambda \in (0, 1]$, known as the forgetting factor, which determines the relative weighting of recent versus past observations. A value of $\lambda = 1$ assigns equal importance to all data points, while $\lambda < 1$ applies exponential forgetting, emphasizing more recent data. In our online learning experimental setup, λ is fixed at 0.99 by default to enable mild forgetting while avoiding any risk of data snooping. However, since RLS lacks a mechanism to dynamically adjust this parameter in response to drift, it may be suboptimal in non-stationary environments. This limitation is addressed in OLR-WAA, which incorporates an internal mechanism to automatically adjust hyperparameters upon each detected concept drift.

5.4.2 Performance Analysis on Incremental Drift datasets

Performance Evaluation: Figure 10 illustrates the performance of the models under incremental drift, where the yellow lines indicate the onset of new concepts in an incremental manner. The models exhibit reduced sensitivity to incremental drift compared to abrupt drift, with OLR-WAA demonstrating the best overall performance. Additionally, LMS shows a significant degradation in performance beyond data point 8000, which may be attributed to the limitations of a fixed learning rate in such settings.

5.4.3 Performance Analysis on Gradual Drift datasets

Performance Evaluation: Figure 11 illustrates the performance of the models on the gradual drift dataset DS13. The green line indicates the onset of concept 2, while the yellow line marks the reintroduction of concept 1. Overall, all models exhibit significant performance degradation, primarily due to fixed hyperparameters that lack adaptability in the presence of concept drift. Notably, OLR-WAA demonstrates superior resilience, emerging as a top-performing model in handling gradual concept drift.

6 Conclusion and Future Work

In streaming environments, batch learning is often unsuitable due to its reliance on full data access, ample training time, and the assumption of a stationary *i.i.d.* distribution. For instance, in the stock market, models trained on outdated historical data may fail to capture rapidly evolving patterns. Online learning addresses these challenges by enabling real-time updates from sequential data and handling distribution shifts. However, it faces limitations such as difficulty adapting to concept drift and the impracticality of manually tuning hyperparameters during continuous change.

This paper presents OLR-WAA: Adaptive and Drift-Resilient Online Regression with Dynamic Weighted Averaging, a hyperparameter-free model designed for robust, real-time regression in evolving data streams. OLR-WAA incrementally updates its base model by integrating streaming data using an exponentially weighted moving average, facilitating smooth adaptation to evolving patterns. To address concept drift, OLR-WAA introduces a novel optimization mechanism that proactively detects and responds to drift using real-time, in-memory processing. This mechanism incorporates holistic drift handling by combining drift detection, severity quantification, adaptive hyperparameter adjustment. It is designed for large-scale, high-dimensional datasets by relying on performance-driven KPIs without making distributional assumptions. Additionally, the

incorporation of dynamic thresholding that does not assume any prior data distribution enhances the effectiveness of OLR-WAA in adapting to varying data patterns, managing high-dimensional inputs, and supporting diverse performance KPIs.

We conducted four types of experiments: stationary, adversarial, convergence, and concept drift, across a diverse set of datasets. The empirical evaluation establishes OLR-WAA and its variant OLR-WAA as a top-tier model for handling a diverse range of scenarios. It ranked consistently top in stationary tasks and excelled in both time-based and uniquely in confidence-based adversarial scenarios. In convergence, it was the only model to achieve low cost and continuously sustain that diminished cost throughout. In contrast to other online models, which exhibit fluctuating performance on different concept drift datasets, OLR-WAA consistently demonstrated superior robustness across all types of concept drift.

Future efforts will concentrate on expanding the OLR-WAA approach to a broader spectrum of machine learning tasks, aiming to boost its versatility and applicability across various domains. Concurrently, we plan to implement OLR-WAA as an online web service for drift-aware, real-time classification, thereby embracing the principles of Machine Learning as a Service (MLaaS).

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s41019-025-00312-y>.

Acknowledgements The authors have no acknowledgements to declare.

Author contributions Mohammad Abu-Shaira (M.A.) led the conceptualization and development of the research project, ensuring a strong theoretical foundation and alignment with emerging trends in online regression. He conducted an extensive literature review, synthesizing relevant studies to identify research gaps and establish a comparative framework. He was responsible for designing and executing experiments, including data collection, preprocessing, and model implementation. M.A. performed in-depth data analysis, interpreting results, and drawing key insights to validate the study's hypotheses. Additionally, he took the lead in drafting the manuscript, structuring the narrative, and developing coherent arguments. He created figures, tables, and visualizations to effectively present the findings. M.A. also played a crucial role in revising and refining the manuscript, integrating feedback, and collaborating with co-authors to enhance clarity, coherence, and impact. Weishi Shi (W.S.) provided mentorship and strategic guidance, supporting the professional development of the first author throughout the research process. He contributed to shaping the research direction, ensuring methodological soundness and theoretical rigor. W.S. critically reviewed and refined manuscript drafts, enhancing clarity, coherence, and scholarly quality. He offered expert insights on data analysis techniques, helping to interpret results and strengthen the research's empirical foundation. Additionally, he facilitated alignment between the conducted studies and established machine learning theories, proposing innovative ideas to extend the research scope and broaden its impact.

Funding No funding was received for this research.

Data Availability The project, including the implementation and synthetic datasets, is available in the following GitHub repository:

GitHub Repository: (<https://github.com/anonymous273800/OLR-WA>) Additionally, publicly available datasets used in this study can be accessed through the following sources. These references are also cited within the manuscript: 1. Medical Cost Personal Dataset: (<https://www.kaggle.com/datasets/mirichoi0218/insurance>) 2. 1000 Companies Profit Dataset: (<https://www.kaggle.com/datasets/rupakroy/1000-companies-profit>) 3. Combined Cycle Power Plant Dataset: (<https://www.kaggle.com/datasets/gova26/airpressure>) 4. King County House Sale Dataset: (<https://www.kaggle.com/datasets/minasameh55/kingcountyr-houses-aa>).

Declarations

Conflict of interest The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Fontenla-Romero, Ó., Guijarro-Berdiñas, B., Martínez-Rego, D., Pérez-Sánchez, B., Peteiro-Barral D (2013) Online machine learning. In: Efficiency and Scalability Methods for Computational Intellect, pp. 27–54. IGI Global,
- Box, G.E., Jenkins, G.M., Reinsel, G.C., Ljung GM (2015) Time Series Analysis: Forecasting and Control. John Wiley & Sons, Hoboken, New Jersey . Published simultaneously in Canada
- Celik B, Singh P, Vanschoren J (2023) Online automl: an adaptive automl framework for online learning. *Mach Learn* 112(6):1897–1921
- Shalev-Shwartz S et al (2012) Online learning and online convex optimization. *Foundations and Trends® in Machine Learning* 4(2):107–194
- Webb GI, Hyde R, Cao H, Nguyen HL, Petitjean F (2016) Characterizing concept drift. *Data Min Knowl Discov* 30(4):964–994
- Ditzler G, Polikar R (2012) Incremental learning of concept drift from streaming imbalanced data. *IEEE Trans Knowl Data Eng* 25(10):2283–2301
- Barbaro B (2018) Tuning hyperparameters for online learning. PhD thesis, case western reserve university
- Montgomery, D.C., Peck, E.A., Vining GG (2021) Introduction to Linear Regression Analysis. John Wiley & Sons
- Crammer, K., Dredze, M., Pereira F (2008) Exact convex confidence-weighted learning. *Adv Neural Inf Process Syst* 21
- Prasad S, Bruce LM (2008) Decision fusion with confidence-based weight assignment for hyperspectral target recognition. *IEEE Trans Geosci Remote Sens* 46(5):1448–1456
- Abu-Shaira M (2025) Online regression with weighted adaptive average. GitHub repository. <https://github.com/mabushaera/OLR-WAA>
- Yaser S. Abu-Mostafa, H-T L (2012) Malik magdon-ismail: learning from data. AMLBook. <https://amlbook.com/>
- Maulud D, Abdulazeez AM (2020) A review on linear regression comprehensive in machine learning. *J Appl Sci Technol Trends* 1(4):140–147
- Boyd S, Vandenberghe L (2004) Convex Optimization. Cambridge University Press, Cambridge
- Hastie, T., Tibshirani, R., Friedman, J.H., Friedman JH (2009) The elements of statistical learning: data mining, inference, and prediction vol. 2. Springer
- Bishop C (2006) Pattern recognition and machine learning. Springer google schola 2:5–43
- Awad, M., Khanna, R., Awad, M., Khanna R (2015) Support vector regression. *Efficient learning machines: Theories, concepts, and applications for engineers and system designers*, 67–80
- Smola AJ, Schölkopf B (2004) A tutorial on support vector regression. *Stat Comput* 14:199–222
- Hoi SC, Sahoo D, Lu J, Zhao P (2021) Online learning: a comprehensive survey. *Neurocomputing* 459:249–289
- Smola, A., Vishwanathan S VN (2008) Introduction to machine learning. cambridge university press,
- Ding, Q., Hsieh, C.-J., Sharpnack J (2021) An efficient algorithm for generalized linear bandit: Online stochastic gradient descent and thompson sampling. In: *Int Conf Artif Intell Stat*, pp. 1585–1593. PMLR
- Bottou L (2010) Large-scale machine learning with stochastic gradient descent. In: *Proceedings of COMPSTAT'2010: 19th International Conference on Computational Statistics Paris France, August 22-27, 2010 Keynote, Invited and Contributed Papers*, pp. 177–186. Springer
- Jin B, Lu X (2018) On the regularizing property of stochastic gradient descent. *Inverse Probl* 35(1):015004
- Bottou L et al (1991) Stochastic gradient learning in neural networks. *Proceedings of Neuro-Nimes* 91(8):12
- Hardt, M., Recht, B., Singer Y (2016) Train faster, generalize better: Stability of stochastic gradient descent. In: *International Conference on Machine Learning*, pp. 1225–1234. PMLR
- Goodfellow, I., Bengio, Y., Courville A (2016) Deep Learning. MIT press
- Ruder S (2016) An overview of gradient descent optimization algorithms. arXiv preprint [arXiv:1609.04747](https://arxiv.org/abs/1609.04747)
- Johnson, R., Zhang T (2013) Accelerating stochastic gradient descent using predictive variance reduction. *Advances in neural information processing systems* 26
- Shah, V., Wu, X., Sanghavi S (2020) Choosing the sample with lowest loss makes sgd robust. In: *Int Conf Artif Intell Stat* pp. 2120–2130. PMLR
- Duchi, J., Hazan, E., Singer Y (2011) Adaptive subgradient methods for online learning and stochastic optimization. *J mach learn res* 12(7)
- Tieleman T (2012) Lecture 6.5-rmsprop: divide the gradient by a running average of its recent magnitude. *Coursera Neural network mach learn* 4(2):26
- Kingma DP (2014) Adam: A method for stochastic optimization. arXiv preprint. [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)
- Liu Y, Gao Y, Yin W (2020) An improved analysis of stochastic gradient descent with momentum. *Adv Neural Inf Process Syst* 33:18261–18271
- Danner, G., Jelasity M (2015) Fully distributed privacy preserving mini-batch gradient descent learning. In: *Distributed Applications and Interoperable Systems: 15th IFIP WG 6.1 International Conference, DAIS 2015, Held as Part of the 10th International Federated Conference on Distributed Computing Techniques*,

- DisCoTec 2015, Grenoble, France, June 2-4, 2015, Proceedings 15, pp. 30–44. Springer
35. Dekel, O., Gilad-Bachrach, R., Shamir, O., Xiao L (2012) Optimal distributed online prediction using mini-batches. *Journal of Machine Learning Research* 13(1)
 36. Bottou L, Curtis FE, Nocedal J (2018) Optimization methods for large-scale machine learning. *SIAM Rev* 60(2):223–311
 37. Li, M., Zhang, T., Chen, Y., Smola AJ (2014) Efficient mini-batch training for stochastic optimization. In: *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 661–670
 38. Avrithis Y (2021) Iterative label cleaning for semi-supervised learning. PhD thesis, National and Kapodistrian University of Athens
 39. Montavon, G., Orr, G., Müller KR (2012) *Neural Networks: Tricks of the Trade* vol. 7700. Springer
 40. J, N., Bhandari N (2021) Computational complexity of gradient descent algorithm. *TechRxiv* <https://doi.org/10.36227/techrxiv.14544000.v1>. Preprint
 41. Hoffer, E., Hubara, I., Soudry D (2017) Train longer, generalize better: closing the generalization gap in large batch training of neural networks. *Advances in neural information processing systems* 30
 42. Recht, B., Re, C., Wright, S., Niu F (2011) Hogwild!: A lock-free approach to parallelizing stochastic gradient descent. *Advances in neural information processing systems* 24
 43. McMahan, B., Moore, E., Ramage, D., Hampson, S., Arcas BA (2017) Communication-efficient learning of deep networks from decentralized data. In: *Artificial Intelligence and Statistics*, pp. 1273–1282. PMLR
 44. Defazio, A., Bach, F., Lacoste-Julien S (2014) Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. *Advances in neural information processing systems* 27
 45. Widrow, B., Hoff M E *et al* (1960) Adaptive switching circuits. In: *IRE WESCON Convention Record*, vol. 4, pp. 96–104. New York
 46. Haykin S (2002) *Adaptive filter theory*. Prentice Hall google scholar 2:333–346
 47. Hsia T (1983) Convergence analysis of lms and nlms adaptive algorithms. *ICASSP'83. IEEE International Conference on Acoustics Speech and Signal Processing* 8:667–670 (IEEE)
 48. Slock DT (1993) On the convergence behavior of the LMS and the normalized LMS algorithms. *IEEE Trans Signal Process* 41(9):2811–2825
 49. Liu W, Principe JC, Haykin S (2011) *Kernel Adaptive Filtering: a Comprehensive Introduction*. John Wiley & Sons, Hoboken, NJ
 50. Kwong RH, Johnston EW (1992) A variable step size lms algorithm. *IEEE Trans Signal Process* 40(7):1633–1642
 51. Ghauri, S.A., Sohail MF (2013) System identification using lms, nlms and rls. In: *2013 IEEE Student Conference on Research and Development*, pp. 65–69. IEEE
 52. Islam SAU, Bernstein DS (2019) Recursive least squares for real-time implementation [lecture notes]. *IEEE Control Syst Mag* 39(3):82–85
 53. Haykin S (2013) *Adaptive Filter Theory: International Edition*, 5/E. Pearson
 54. Bittanti, S., Bolzern, P., Campi, M., Coletti E (1988) Deterministic convergence analysis of rls estimators with different forgetting factors. In: *Proceedings of the 27th IEEE Conference on Decision and Control*, pp. 1530–1531. IEEE
 55. Li, Z., Li C (2011) Lms and rls algorithms comparative study in system identification. In: *2011 International Conference on Multimedia Technology*, pp. 5428–5430. IEEE
 56. Douglas SC (2000) Numerically-robust $O(n/\sup 2)$ rls algorithms using least-squares prewhitening. In: *2000 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No. 00CH37100)*, 1, 412–415. IEEE
 57. Xiao Y, Ma L, Ward RK (2007) Fast rls fourier analyzers capable of accommodating frequency mismatch. *Signal Process* 87(9):2197–2212
 58. Djigan VI (2006) Multichannel parallelizable sliding window RLS and fast RLS algorithms with linear constraints. *Signal Process* 86(4):776–791
 59. Liu K, Hsieh S-F, Yao K, Chiù C-T *et al* (1991) Dynamic-range, stability, and fault-tolerant capability of finite-precision RLS systolic array based on givens rotations. *IEEE Trans Circuits Syst* 38(6):625–636
 60. Witten, D., James G (2013) *An Introduction to Statistical Learning with Applications in R*. Springer publication
 61. Mohri, M., Rostamizadeh, A., Talwalkar A (2018) *Foundations of Machine Learning*. MIT press
 62. Stephenson W, Frangella Z, Udell M, Broderick T (2021) Can we globally optimize cross-validation loss? Quasiconvexity in ridge regression. *Adv Neural Inf Process Syst* 34:24352–24364
 63. Snoek, J., Larochelle, H., Adams RP (2012) Practical bayesian optimization of machine learning algorithms. *Advances in neural information processing systems* 25
 64. Huang, C., Yao Y (2018) A unified dynamic approach to sparse model selection. In: *International Conference on Artificial Intelligence and Statistics*, pp. 2047–2055. PMLR
 65. Xu, S., Zhang, X., Liao S (2019) New online kernel ridge regression via incremental predictive sampling. In: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pp. 791–800
 66. Akhtar N, Alharthi MF (2024) A comparative study of the performance of new ridge estimators for multicollinearity: insights from simulation and real data application. *AIP Adv.* <https://doi.org/10.1063/5.0236631>
 67. Ouhamma R, Maillard O-A, Perchet V (2021) Stochastic online linear regression: the forward algorithm to replace ridge. *Adv Neural Inf Process Syst* 34:24430–24441
 68. Ransam J, Cook J (2018) Lasso regression. *Br J Surg* 105(10):1348–1348
 69. Tibshirani R (1996) Regression shrinkage and selection via the lasso. *J R Stat Soc Ser B Stat Methodol* 58(1):267–288
 70. Chichignoud M, Lederer J, Wainwright MJ (2016) A practical scheme and fast algorithm to tune the lasso with optimality guarantees. *J Mach Learn Res* 17(1):8162–8181
 71. Yang, S., Yan, Y., Zhu, X., Sun Q (2023) Online linearized lasso. In: *International Conference on Artificial Intelligence and Statistics*, pp. 7594–7610. PMLR
 72. Zou H, Hastie T (2005) Regularization and variable selection via the elastic net. *J R Stat Soc Ser B Stat Methodol* 67(2):301–320
 73. Crammer, K., Dekel, O., Keshet, J., Shalev-Shwartz, S., Singer Y (2006) Online passive aggressive algorithms. *Online passive aggressive algorithms*
 74. Shi, R., Ying, J., Palomar DP (2024) Adaptive passive-aggressive framework for online regression with side information. In: *The Thirty-eighth Annual Conf Neural Inf Process Syst*
 75. Wu, S., Lu, F., Raff, E., Holt J (2024) Stabilizing linear passive-aggressive online learning with weighted reservoir sampling. In: *The Thirty-eighth Annual Conf Neural Inf Process Syst*
 76. Gama J, Žliobaitė I, Bifet A, Pechenizkiy M, Bouchachia A (2014) A survey on concept drift adaptation. *ACM Comput Surv* 46(4):1–37
 77. Shumway RH, Stoffer DS, Stoffer DS (2000) *Time Series Analysis and Its Applications*, vol 3. Springer, New York
 78. Stewart, J., Clegg, D.K., Watson S (2020) *Calculus: Early Transcendentals*. Ceng Learn
 79. Layton W, Sussman M (2014) *Numerical linear algebra*. University of Pittsburgh, Pittsburgh, pp 28–39

80. Lu J, Liu A, Dong F, Gu F, Gama J, Zhang G (2018) Learning under concept drift: a review. *IEEE Trans Knowl Data Eng* 31(12):2346–2363
81. Buzmakov A, Kuznetsov SO, Napoli A (2014) Is concept stability a measure for pattern selection? *Procedia Comput Sci* 31:918–927
82. Jeannerod C-P, Pernet C, Storjohann A (2013) Rank-profile revealing gaussian elimination and the cup matrix decomposition. *J Symbol Comput* 56:46–68
83. Chicco D, Warrens MJ, Jurman G (2021) The coefficient of determination r-squared is more informative than smape, mae, mape, mse and rmse in regression analysis evaluation. *PeerJ Comput Sci* 7:623
84. Nagelkerke NJD (1991) A note on a general definition of the coefficient of determination. *Biometrika* 78(3):691–692
85. Medical Cost Personal Datasets. online - kaggle. <https://www.kaggle.com/datasets/mirichoi0218/insurance>
86. 1000 Companies Profit. online - kaggle. <https://www.kaggle.com/datasets/rupakroy/1000-companies-profit>
87. King County House Sale Data. online - kaggle. <https://www.kaggle.com/datasets/minasameh55/king-country-houses-aa>
88. Combined Cycle Power Plant Datasets. online - kaggle. <https://www.kaggle.com/datasets/gova26/airpressure>
89. Cakra, Y.E., Trisedya, B.D.: Stock price prediction using linear regression based on sentiment analysis. In: 2015 International Conference on Advanced Computer Science and Information Systems (ICACSIS), pp. 147–154 (2015). IEEE