



Adaptive Machine Learning Models for Real-Time Decision Support: A Survey

Research Article

<https://stem.techspherejournal.com>

Article Info

Revised Date: 30th October 2025

Accepted Date: 4th November 2025

Published Date: 10th November 2025

Keywords

Adaptive Machine Learning

Real-Time Decision Support

Online Learning

Continual Learning

Dynamic Environments

Author Details

Onwujei Augustine Ikechukwu^{1*}, Obaze Caleb Akachukwu²

¹ Department of Software Engineering, Dennis Osadebay University, Asaba, Delta State.

² Department of Computer Science, Dennis Osadebay University, Asaba, Delta State.

*Corresponding author's email: onwujei.ikechukwu@dou.edu.ng

DOI: <https://doi.org/10.5281/zenodo.17568180>

This is an open-access article under the [CC BY-SA](https://creativecommons.org/licenses/by-sa/4.0/) license



ABSTRACT

This survey provides a comprehensive overview of adaptive machine learning models for real-time decision support systems. In dynamic environments where data distributions can change over time, traditional static models often become obsolete, leading to degraded performance. Adaptive models, which can learn and evolve from continuous data streams, offer a promising solution. This paper categorizes and discusses various adaptive machine learning techniques, including online learning, continual learning, and reinforcement learning. We explore their application across diverse domains such as finance, healthcare, and autonomous systems. Furthermore, we identify and analyze the key challenges in the deployment and maintenance of these models, including concept drift, catastrophic forgetting, and the need for model interpretability. Finally, we highlight emerging trends and future research directions that are poised to shape the future of adaptive real-time decision support.

1 Introduction

1.1 The Need for Real-Time Decision Support

In today's rapidly evolving digital landscape, the ability to make instantaneous, data-driven decisions is no longer a competitive advantage but a necessity for organizational success [1]. Industries ranging from finance and healthcare to e-commerce and manufacturing are increasingly dependent on real-time data to navigate hyper-competitive and volatile market environments [2]. This demand is fuelled by several factors, including the digital transformation of business processes, rising consumer expectations for immediate services, and the pressures of global competition. [1]. Real-time decision-making, which involves making choices based on current, streaming data, allows organizations to identify emerging trends, respond swiftly to market changes, and optimize their operations for enhanced efficiency [3].

Data-driven decision-making provides a systematic and objective approach that can lead to increased profitability and productivity [4]. By leveraging real-time data, businesses can gain a deeper understanding of their customers, personalize experiences, and ultimately improve satisfaction and loyalty [5]. In critical sectors like healthcare, immediate access to patient data can significantly influence treatment decisions and improve outcomes, while in finance, split-second trading decisions are made based on real-time market data [6]. The ability to process and act on information as it is generated is transformative, enabling organizations to be more agile, proactive, and resilient in the face of constant change [7].



1.2 The Limitations of Static Machine Learning Models

Traditional machine learning models, often referred to as static or batch models, are trained on a fixed, historical dataset [8]. This approach involves a well-defined workflow of data collection, model training, evaluation, and deployment. Once trained, these models operate with a static understanding of the world, making predictions based on the patterns and relationships learned from the initial dataset [9]. While effective in stable environments where data distributions do not change significantly over time, these models have inherent limitations when faced with the dynamism of real-world scenarios [10].

A primary drawback of static models is their inability to adapt to new data in real-time. In environments where data is constantly evolving, such as in financial markets or e-commerce, static models can quickly become outdated, leading to a degradation in prediction accuracy [11]. To incorporate new information, these models must be retrained from scratch, a process that can be computationally expensive, time-consuming, and resource-intensive, especially with large datasets [12]. This periodic retraining creates latency, meaning that the insights generated by the model may no longer be relevant by the time they are available [13]. Furthermore, static models are vulnerable to "concept drift," a phenomenon where the statistical properties of the data change over time, rendering the learned patterns obsolete [14]. This lack of flexibility and adaptability makes static models ill-suited for applications that require immediate responses to changing conditions [15].

1.3 The Emergence of Adaptive Machine Learning

In response to the limitations of static models, the field of adaptive machine learning has emerged as a powerful paradigm for building intelligent systems capable of learning and evolving in dynamic environments [16]. Unlike their static counterparts, adaptive machine learning models are designed to continuously learn from a stream of incoming data, updating their internal parameters and decision-making logic in real-time [16]. This ability to adapt to new information allows these models to remain relevant and accurate even as the underlying data distributions change [17]. Adaptive machine learning encompasses a variety of techniques that enable models to learn from experience and adjust their behavior accordingly [18]. This approach is particularly well-suited for real-world applications where data is generated continuously and unpredictably [19]. By processing data sequentially as it arrives, adaptive models can provide immediate insights and predictions, overcoming the latency issues associated with batch learning. The core principle of adaptive machine learning is to create systems that are more robust, efficient, and agile [18]. As new data is ingested, these models can identify and respond to emerging patterns, a process that is crucial for tasks such as real-time fraud detection, personalized recommendations, and autonomous navigation [20]. The more data an adaptive model is exposed to, the "smarter" and more accurate it becomes over time, effectively learning from its past interactions to improve future performance.

1.4 Scope and Contributions of this Survey

This survey provides a comprehensive overview of the state-of-the-art in adaptive machine learning models for real-time decision support. Our goal is to offer a structured and insightful examination of the key concepts, methodologies, applications, and challenges in this rapidly advancing field. The primary contributions of this paper are threefold:

1. **A Unified Taxonomy:** We present a clear and organized taxonomy of adaptive machine learning models, categorizing them into distinct approaches such as online learning, continual learning, and reinforcement learning. This classification helps to delineate the unique characteristics, strengths, and weaknesses of each approach.
2. **Cross-Disciplinary Application Insights:** We explore the practical application of these models across a diverse range of domains, including finance, healthcare, autonomous systems, and e-commerce. By highlighting real-



world use cases, we demonstrate the transformative impact of adaptive machine learning on decision-making processes in various industries.

3. **Identification of Key Challenges and Future Directions:** We identify and discuss the significant challenges that researchers and practitioners face in the development and deployment of adaptive models. These include issues of concept drift, catastrophic forgetting, data quality, model interpretability, and system-level complexities. Furthermore, we outline promising future research directions that are poised to shape the next generation of adaptive real-time decision support systems.

Through this comprehensive survey, we aim to provide a valuable resource for both newcomers to the field and experienced researchers, fostering a deeper understanding of the principles and practices of adaptive machine learning and inspiring further innovation in this critical area of artificial intelligence.

2 Background: From Batch to Real-Time Learning

2.1 Batch Learning Paradigm

The traditional approach to training machine learning models is known as batch learning, or offline learning.[1] In this paradigm, the model is trained on the entirety of a collected dataset at once.[2] The process is static; the system learns from a fixed, complete set of data before it is deployed.[3] Once the training phase is complete, the model's parameters are frozen, and it begins to make predictions in a production environment without further learning.[3] To incorporate new data or adapt to changes in the underlying data patterns, the model must be taken offline and retrained from scratch on an updated dataset that includes both old and new data.[2] This method is computationally intensive but can be effective for stable environments where data distributions do not change frequently.[3][4].

2.2 Streaming Data and its Characteristics

In contrast to the static datasets used in batch learning, many modern applications generate data continuously, creating what is known as a data stream. Streaming data is not a finite collection but an unbounded, continuous flow of data points ordered in sequence.[5] This type of data is defined by several key characteristics, often referred to as the "Vs" of Big Data.[6] These include:

- a. **Volume:** Streaming data often involves massive quantities of information generated at a high scale.[7]
- b. **Velocity:** The data is generated and arrives at a very high speed, requiring immediate processing to derive timely insights.[6][7]
- c. **Variety:** The data can come in various formats, ranging from structured numerical data to unstructured text, images, and videos.[5][7]

The continuous and high-velocity nature of streaming data presents significant challenges for traditional batch processing methods, which are not designed to handle a constant influx of information.[8]

2.3 The Concept of Online Learning

Online learning, also known as incremental learning, emerged as a direct response to the limitations of the batch paradigm when dealing with streaming data.[3][4] Unlike batch learning, which processes the entire dataset at once, online learning updates the model sequentially, instance by instance or in small mini-batches, as new data arrives.[1][3] This incremental updating allows the model to learn on the fly and adapt to changes in the data over time without the need for complete retraining.[2][3][9] The model continuously evolves and refines its knowledge as it is exposed to more data, making it well-suited for dynamic environments where patterns can change rapidly.[10] This approach is more



resource-efficient in terms of computation and memory, as it does not require storing and reprocessing the entire historical dataset for each update.[1].

2.4 Real-Time Decision Support Systems (RT-DSS)

A Real-Time Decision Support System (RT-DSS) is an information system designed to support organizational decision-making by providing immediate, data-driven insights.[11][12] These systems are crucial in environments where decisions must be made quickly in response to rapidly changing conditions.[12] An RT-DSS integrates and analyzes streaming data from various sources to offer actionable intelligence, often through interactive dashboards, alerts, and visualizations.[11][13]

The critical role of RT-DSS is evident across numerous sectors. In finance, they power algorithmic trading and real-time fraud detection systems.[13] In healthcare, clinical decision support systems analyze patient data streams to assist physicians with diagnoses and treatment plans.[14] Logistics and supply chain management rely on RT-DSS to optimize routes and manage inventory based on real-time information.[15] By leveraging continuous data analysis, these systems empower organizations to move beyond reactive problem-solving and adopt a proactive approach, enhancing efficiency, mitigating risks, and capitalizing on opportunities as they arise.[11]

3 A Taxonomy of Adaptive Machine Learning Models

Adaptive machine learning models can be broadly categorized into several distinct paradigms, each with its own methodologies and ideal use cases. This section provides a taxonomy of these models, focusing on online learning, continual learning, and reinforcement learning as key approaches for building systems that can learn and adapt in real time.

3.1 Online Learning Models

Online learning models process data sequentially, updating the model's parameters with each new instance or small batch of instances. This incremental approach allows them to adapt to evolving data streams without the need to be retrained from scratch.

Stochastic Gradient Descent (SGD) is a fundamental optimization algorithm that is particularly well-suited for online learning. Unlike traditional batch gradient descent, which computes the gradient of the loss function over the entire dataset, SGD updates the model's parameters based on the gradient of the loss for a single, randomly selected training example at each iteration.[1][2] This makes SGD computationally efficient and allows for rapid model updates, which is essential for handling high-velocity data streams.[3]

Several variants of SGD have been developed to improve its convergence and stability. These include Mini-batch Gradient Descent, which updates the parameters based on a small subset of the training data, reducing the variance of the updates and leading to more stable convergence.[1] Other popular variants include Momentum, which helps accelerate SGD in the relevant direction and dampens oscillations, and adaptive learning rate methods like Adagrad and Adam, which adjust the learning rate for each parameter individually.

The Perceptron is one of the earliest and simplest online learning algorithms for binary classification. It processes training examples one by one and updates its weight vector only when it makes a misclassification. This error-driven update rule makes it a computationally efficient and straightforward approach for linear classification tasks in a streaming setting.



Passive-Aggressive (PA) algorithms are a family of online learning algorithms that are particularly effective for large-scale learning.[4] The name "Passive-Aggressive" reflects their update strategy: if a new data point is correctly classified, the model remains "passive" and does not change.[4] However, if a misclassification occurs, the model becomes "aggressive" and updates its parameters to correct the mistake, while also aiming to keep the change to the model as small as possible.[4] Like the Perceptron, PA algorithms do not require a learning rate but often include a regularization parameter to control the trade-off between correcting mistakes and maintaining the stability of the model.[4][5]

Hoeffding Trees, also known as Very Fast Decision Trees (VFDT), are a type of incremental decision tree designed specifically for learning from high-speed data streams.[6][7] Traditional decision trees require multiple passes over the entire dataset to determine the best split at each node.[8] In contrast, Hoeffding Trees use a statistical result called the Hoeffding bound to make split decisions with a certain level of confidence after observing a relatively small number of examples.[6][7]

This allows the tree to be built incrementally, one data point at a time, without storing the entire dataset in memory.[9] Hoeffding Trees have theoretical guarantees that, given enough data, they will converge to the same tree that would have been built by a batch decision tree learner.[8] Various extensions of the Hoeffding Tree have been developed to handle concept drift and other challenges in streaming data, such as the Hoeffding Adaptive Tree.[10]

3.2 Continual Learning (Lifelong Learning)

Continual learning, or lifelong learning, aims to develop models that can learn a sequence of tasks without forgetting previously learned knowledge. This is a significant challenge in machine learning, as neural networks are prone to a phenomenon known as catastrophic forgetting.

Catastrophic forgetting, also known as catastrophic interference, occurs when a neural network trained on a new task loses the knowledge it acquired from previous tasks.[11][12][13] This happens because the network's weights are adjusted to minimize the error for the new task, which can overwrite the representations that were important for the old tasks.[12] The core challenge in continual learning is to balance the stability of previously learned knowledge with the plasticity required to learn new information.[11]

Regularization-based approaches to continual learning add a penalty term to the loss function that discourages changes to the weights that are important for previous tasks. A prominent example is Elastic Weight Consolidation (EWC).[14][15] EWC calculates the importance of each weight for a previously learned task and then penalizes changes to the more important weights when learning a new task.[16] This is analogous to synaptic consolidation in the brain, where important connections are stabilized to protect them from being overwritten.[15] EWC effectively creates "elastic" connections, allowing the network to learn new tasks while retaining knowledge of old ones.[17]

Replay-based methods address catastrophic forgetting by storing a subset of data from previous tasks in a memory buffer.[13][18] When the model is trained on a new task, it is also trained on a small sample of the stored data from past tasks.[19] This "experience replay" helps to refresh the model's memory of the old tasks and prevents it from overfitting to the new task.[20] While effective, these methods can be memory-intensive if a large amount of past data needs to be stored.[20]

Architecture-based approaches dynamically modify the model's architecture to accommodate new tasks. Progressive Neural Networks (PNNs) are a key example of this approach. When a new task is introduced, a PNN adds a new "column" (a new neural network) to the existing architecture. The new column can access the features learned by the



previous columns, allowing for knowledge transfer from old tasks to the new one.[23] By freezing the weights of the old columns, PNNs are immune to catastrophic forgetting by design. However, the size of the model grows with each new task, which can be a limitation in terms of scalability.

3.3 Reinforcement Learning (RL) for Adaptive Decision-Making

Reinforcement learning is a paradigm where an agent learns to make a sequence of decisions in an environment to maximize a cumulative reward signal. RL is inherently adaptive, as the agent continuously learns and refines its policy through interaction with the environment.

In dynamic environments where the rules, rewards, or transitions can change over time, traditional RL algorithms that assume a stationary environment may fail. Adaptive RL techniques are designed to enable agents to adjust their strategies in response to these changes. This involves a continuous feedback loop where the agent explores the environment, exploits its current knowledge to make decisions, and updates its policy based on the feedback it receives. This makes RL well-suited for real-time decision-making in complex and unpredictable settings.

Two of the most fundamental approaches in RL are Q-learning and policy gradient methods.

- a. Q-learning is a value-based method where the agent learns a Q-function that estimates the expected future reward of taking a particular action in a given state. It is an off-policy algorithm, meaning it can learn from past experiences stored in a replay buffer, which can improve sample efficiency.
- b. Policy Gradient methods directly learn a policy that maps states to actions. Instead of learning a value function, these methods adjust the policy's parameters in the direction that increases the expected reward. Policy gradient methods are often on-policy, meaning they learn from the actions taken by the current policy. They are particularly well-suited for continuous action spaces.

Meta-reinforcement learning (meta-RL) is an advanced area of RL that aims to train agents that can "learn to learn." The goal of meta-RL is to enable an agent to quickly adapt to new, unseen tasks with only a small amount of experience. This is achieved by training the agent on a distribution of related tasks. During this meta-training phase, the agent learns a general learning algorithm that it can then apply to solve new tasks at test time rapidly. Meta-RL is a promising approach for developing truly adaptive agents that can generalize their knowledge and skills to novel situations.

4 Applications of Adaptive Models in Real-Time Decision Support

Adaptive machine learning models are being deployed across a wide range of industries to power real-time decision support systems. Their ability to learn from streaming data and adapt to changing conditions makes them invaluable in dynamic environments. This section explores some of the key applications of these models in finance, healthcare, autonomous systems, and other domains.

4.1 Finance

The financial sector is characterized by high-velocity data and the need for split-second decision-making. Adaptive models are well-suited to this environment, where market conditions can change in an instant.

In algorithmic trading, adaptive models are used to predict short-term price movements and execute trades at optimal times. These models can continuously learn from real-time market data, including price feeds, order books, and news sentiment, to adjust their trading strategies. Reinforcement learning, in particular, is used to develop agents that can learn to make profitable trading decisions in complex and non-stationary market environments. Furthermore, adaptive



models are employed for market impact analysis, predicting how a large trade will affect the price of an asset and allowing traders to minimize adverse price movements.

Online learning is extensively used in real-time fraud detection systems to identify and prevent fraudulent transactions as they occur.[1][2] As fraudsters constantly change their tactics, static, rule-based systems quickly become outdated. Adaptive models can learn from a continuous stream of transaction data, identifying new and evolving patterns of fraudulent behavior. By flagging suspicious activities in real-time, these systems can block fraudulent transactions before they are completed, significantly reducing financial losses for both institutions and their customers.

Adaptive machine learning models are transforming credit scoring and risk management by providing a more dynamic and accurate assessment of creditworthiness.[3][4][5][6] Unlike traditional static credit scoring models, adaptive models can continuously learn from new data, such as a borrower's recent transaction history and repayment behavior.[5][7][8] This allows for a more up-to-date evaluation of risk, which is particularly beneficial for individuals with limited credit histories.[9] By incorporating real-time data, lenders can make more informed decisions, adjust credit limits dynamically, and proactively manage the risk of default.[5][10]

4.2 Healthcare

In healthcare, real-time data from electronic health records, medical devices, and wearables offers new opportunities for improving patient care and clinical outcomes. Adaptive models are crucial for turning this data into actionable insights.

Adaptive models are being integrated into Clinical Decision Support Systems (CDSS) to enhance diagnostic accuracy and personalize treatment recommendations.[11][12][13] These systems can analyze real-time patient data, including vital signs, lab results, and clinical notes, to provide physicians with timely and relevant information. For instance, an adaptive model could learn from the outcomes of different treatments for patients with similar profiles to recommend the most effective course of action for a new patient.

Streaming data from wearable sensors and bedside monitors can be used to continuously track a patient's health status.[14] Adaptive models can analyze these data streams in real time to detect subtle changes that may indicate a deterioration in the patient's condition.[14] These early warning systems can alert clinicians to potential adverse events, such as sepsis or cardiac arrest, allowing for timely intervention and improved patient outcomes.

Adaptive models play a key role in the field of personalized medicine, where treatment plans are tailored to the individual characteristics of each patient. These models can continuously learn from a patient's genomic data, lifestyle information, and response to treatment to adapt their therapeutic regimen over time. This is particularly important in areas like oncology, where a tumor's genetic makeup can evolve, requiring adjustments to the treatment strategy.[15][16]

4.3 Autonomous Systems

Autonomous systems, such as self-driving cars and robots, operate in highly dynamic and unpredictable environments. Adaptive models are essential for enabling these systems to perceive their surroundings, make decisions, and act safely and effectively.

Reinforcement learning is a core technology in the development of autonomous vehicles and robots.[17][18] These systems use a variety of sensors, including cameras, LiDAR, and radar, to perceive their environment. Adaptive models can process this sensory data in real time to understand the current situation and predict the actions of other agents,



such as pedestrians and other vehicles. This allows the autonomous system to make safe and intelligent decisions, such as when to change lanes or how to navigate a cluttered environment.

In smart cities, a vast network of IoT sensors generates a continuous stream of data on everything from traffic flow and energy consumption to air quality and public safety.[19][20] Adaptive machine learning models are used to analyze this data in real time to optimize city services.[15] For example, intelligent traffic management systems can use real-time data to adjust traffic signal timing and reroute vehicles to reduce congestion. Similarly, adaptive models can be used for dynamic resource allocation in IoT networks and for environmental monitoring.[16][17]

4.4 Other Domains

The applications of adaptive models extend to numerous other domains where real-time decision-making is critical.

Online learning is widely used in e-commerce to personalize the shopping experience in real time.[29] Recommendation systems can continuously update their suggestions based on a user's browsing history, clicks, and purchases. This allows for highly relevant and timely recommendations that can significantly increase user engagement and sales.[1]

In cybersecurity, the threat landscape is constantly evolving, with new types of attacks emerging all the time.[37] Traditional signature-based detection systems are often ineffective against these novel threats. Adaptive machine learning models can be used to build anomaly detection systems that learn the normal patterns of network traffic and user behavior.[12][13][18] When a deviation from these normal patterns is detected, the system can raise an alert, allowing for a rapid response to potential security breaches.

5 Challenges and Open Research Questions

The deployment of adaptive machine learning models in real-time decision support systems, while powerful, is not without its significant challenges. This section outlines the primary obstacles and open research questions related to concept drift, data, model, and system-level issues.

5.1 Concept Drift

Concept drift is a fundamental challenge in streaming data environments, where the statistical properties of the data change over time, leading to a degradation in model performance.[1][2]

Concept drift can manifest in several ways, each requiring a different approach for detection and adaptation. The main types of concept drift are:

- a. **Sudden or Abrupt Drift:** This occurs when the data distribution changes unexpectedly and rapidly. An example of this would be a sudden shift in consumer purchasing behavior due to a global event like a pandemic.[2][3]
- b. **Gradual Drift:** This involves a slow and continuous change in the data distribution over a more extended period. An example is the gradual evolution of customer preferences in a recommendation system.[2][4]
- c. **Incremental Drift:** This is a subtype of gradual drift where changes happen in small, sequential steps.[5]
- d. **Recurring Drift:** This type of drift involves cyclical or seasonal changes where a previously seen concept may reappear. A common example is the seasonal demand for certain products in e-commerce.[2][5]

Effectively managing concept drift involves two key stages: detection and adaptation.[6]



- a. **Drift Detection:** This involves monitoring the model's performance or the statistical properties of the incoming data to identify when a drift is occurring. Common methods include statistical process control, monitoring the error rate of the model, and using window-based techniques to compare recent data with older data.[1][7]
- b. **Drift Adaptation:** Once a drift is detected, the model must adapt to the new data distribution. Adaptation strategies can be broadly categorized as:
 - **Online Learning:** The model continuously updates with new data, which naturally allows it to adapt to gradual changes.
 - **Window-based approaches:** The model is trained on a sliding window of recent data, effectively forgetting older, potentially irrelevant information.
 - **Ensemble methods:** An ensemble of models is maintained, and their predictions are combined. When a drift is detected, underperforming models can be replaced or retrained.

A significant challenge is the trade-off between sensitivity to drift and the rate of false alarms, where the system incorrectly identifies a drift.[8].

5.2 Data-Related Challenges

The nature of streaming data introduces several challenges related to data quality and labeling.

Streaming data is often noisy, incomplete, and can come in various formats.[9] Unlike batch processing, where data can be thoroughly cleaned and preprocessed offline, in a streaming context, these tasks must be performed in real-time with limited computational resources.[10] Issues such as missing values, outliers, and data inconsistencies need to be handled on the fly to prevent them from negatively impacting the model's performance.[11][12] Efficient and scalable data validation and cleansing pipelines are crucial for maintaining data quality in real-time systems.[11][13]

Many adaptive machine learning models, especially in supervised learning settings, require labeled data to learn and update. However, obtaining accurate labels in real-time can be a significant bottleneck.[14] Manual labeling is often too slow and expensive for high-velocity data streams. While there are semi-supervised and active learning techniques that aim to reduce the labeling effort, the challenge of providing timely and accurate feedback to the model remains an active area of research. Real-time annotation tools and chaining labeling jobs are emerging as potential solutions.[15][16] The difficulty in acquiring timely labels also complicates the evaluation of model performance, as the ground truth may be delayed.[17].

5.3 Model-Related Challenges

Beyond the data itself, there are inherent challenges in developing and evaluating adaptive models for real-time applications.

Adaptive models must be able to process high-velocity data streams with minimal latency.[18] This requires algorithms that are computationally efficient and can scale to handle large volumes of data.[19] As models become more complex, their computational and memory requirements can become a significant barrier to real-time performance.[20] There is an ongoing need for research into scalable architectures and optimization techniques that can balance model complexity with the demands of real-time processing.

Evaluating the performance of an adaptive model in a streaming context is more complex than in a batch setting. Traditional evaluation methods like k-fold cross-validation are not applicable to data streams. Instead, techniques like prequential or interleaved-test-then-train evaluation are used, where each new data point is first used to test the model and then to train it. A significant challenge arises when there is a delay in receiving the ground truth labels, making it



difficult to assess the model's performance in real-time.[17] This necessitates the development of new evaluation metrics and methodologies specifically designed for streaming data.

As adaptive models are increasingly deployed in high-stakes domains like healthcare and finance, the ability to understand and interpret their decisions becomes crucial. Many powerful machine learning models, such as deep neural networks, are often considered "black boxes," making it difficult to understand their internal workings. The field of Explainable AI (XAI) aims to develop techniques that can provide insights into why a model made a particular prediction. For adaptive models that are continuously changing, the challenge of providing real-time, understandable explanations is even greater. Ensuring that adaptive models are not only accurate but also transparent and fair is a critical area of ongoing research.

5.4 System-Level Challenges

The successful deployment and maintenance of adaptive models also involve significant system-level challenges.

Deploying and managing machine learning models that are continuously learning in a production environment is a complex task. The field of MLOps (Machine Learning Operations) has emerged to address these challenges by applying DevOps principles to the machine learning lifecycle. For adaptive models, this includes creating reproducible pipelines for data processing, model training, and deployment. Continuous monitoring of model performance and data distributions is essential to detect issues like concept drift and model degradation. Automating the entire machine learning lifecycle, from data ingestion to model retraining and deployment, is a key goal of MLOps for adaptive systems.

Real-time decision support systems often process sensitive personal and proprietary data, raising significant security and privacy concerns. The continuous flow of data in these systems creates a larger attack surface for potential security breaches. Ensuring data privacy is paramount, especially with regulations like GDPR. Techniques such as federated learning, where models are trained on decentralized data without the data ever leaving the local device, and differential privacy, which adds noise to the data to protect individual privacy, are promising areas of research for building secure and privacy-preserving adaptive systems.

5.5 Explainable AI (XAI) for DRL in Robotics

As DRL-powered robots are deployed in safety-critical applications, such as autonomous driving and medical robotics, the ability to understand and trust their decision-making processes becomes paramount. Explainable AI (XAI) aims to develop techniques for making the "black box" nature of deep neural networks more transparent and interpretable.

In the context of DRL for robotics, XAI can provide insights into:

- a. Why a robot chose a particular action: By visualizing the parts of the input state that were most influential in the robot's decision, we can gain a better understanding of its reasoning process.
- b. What a robot has learned: XAI can help to analyze the internal representations of the DRL agent and identify the features of the environment that it deems important.
- c. When a robot is likely to fail: By understanding the limitations of the learned policy, we can identify situations where the robot's behavior may be unreliable.

The development of XAI methods for DRL is crucial for the certification and public acceptance of autonomous robotic systems, ensuring that they are not only capable but also trustworthy and accountable.

6 Future Trends and Directions

The field of adaptive machine learning is continuously evolving, driven by new research, increasing computational power, and the growing demand for more intelligent and autonomous systems. This section explores several key trends and future directions that are poised to shape the future of real-time decision support.



6.1 Hybrid and Ensemble Models

A significant trend in adaptive machine learning is the development of hybrid and ensemble models that combine different techniques to leverage their respective strengths. Ensemble methods, which integrate multiple models to improve predictive performance, are being adapted for real-time systems.[1][2] For instance, an ensemble could combine the outputs of a Local Outlier Factor, a One-Class Support Vector Machine, and an Autoencoder through a weighted average to enhance anomaly detection.[3] This approach can lead to more robust and accurate predictions than any single model could achieve on its own. Hybrid models that merge different machine learning paradigms, such as combining online learning with reinforcement learning, are also gaining traction.[1][4] These composite models are better equipped to handle the complex and multifaceted challenges of real-world data streams, such as dealing with imbalanced datasets and various types of concept drift.[1][4].

6.2 Federated and Edge Learning

Federated learning and edge learning are emerging as critical paradigms for building adaptive systems that are both privacy-preserving and efficient.[5] Instead of sending raw data to a centralized server for processing, these approaches enable model training directly on decentralized edge devices, such as smartphones or IoT sensors.[6][7] In federated learning, only the model updates, not the raw data, are sent to a central server for aggregation, which significantly enhances data privacy and security.[7][8] This is particularly important for sensitive applications in healthcare and finance.[6] By processing data at the edge, these methods also reduce latency and bandwidth consumption, making them ideal for real-time applications like autonomous vehicles and industrial IoT.[5][8] Adaptive federated learning is an active area of research, focusing on developing algorithms that can dynamically adjust the learning process in resource-constrained edge environments.[8][9]

6.3 Automated Machine Learning (AutoML) for Adaptive Systems

Automated Machine Learning (AutoML) aims to automate the end-to-end process of applying machine learning, from data preprocessing and feature engineering to model selection and hyperparameter optimization.[10][11] While AutoML has traditionally focused on static, batch learning scenarios, there is a growing interest in developing AutoML systems for online and adaptive learning.[12][13][14] An adaptive Online AutoML (OAML) system could continuously search for the optimal pipeline configuration, including the best preprocessing algorithms and online learning models, as the data distribution evolves.[12][13][14] This would automate the process of adapting to concept drift and ensure that the model remains performant over time.[12][14] By democratizing access to advanced machine learning and streamlining the development of adaptive systems, AutoML has the potential to accelerate the adoption of real-time decision support across various industries.[10]

6.4 Agentic AI and Autonomous Decision-Making

The rise of Agentic AI represents a significant leap towards more autonomous and intelligent systems. Agentic AI systems are designed to perceive their environment, make complex decisions, and act to achieve specific goals with minimal human intervention.[15][16] Unlike traditional AI models that often require predefined rules or human input, agentic systems can operate independently, learn from their interactions, and adapt their strategies in real-time.[17][18] This is particularly relevant for applications in finance, where AI agents can autonomously execute trading strategies, and in healthcare, where they can assist in managing complex patient data and adapting treatment plans.[18] Agentic AI is poised to revolutionize business operations by enabling autonomous decision-making and enhancing operational efficiencies.[17][18]

6.5 Quantum Machine Learning

While still in its early stages, Quantum Machine Learning (QML) holds the long-term potential to revolutionize real-time decision-making by harnessing the power of quantum computing.[19] Quantum computers, with their ability to



perform certain computations exponentially faster than classical computers, could significantly accelerate the training of complex machine learning models.[20]. QML algorithms could enhance optimization, enabling more accurate and timely predictions, which would be transformative for fields like drug discovery and financial modeling.[20][22] In the context of real-time systems, quantum-enhanced models could process vast and complex datasets at unprecedented speeds, enabling rapid analysis and decision-making.[20] Although significant hardware and algorithmic challenges remain, the ongoing research into hybrid quantum-classical models suggests a promising future for QML in solving problems that are currently intractable for even the most powerful supercomputers [20].

7 Conclusion

This survey has provided a comprehensive examination of adaptive machine learning models and their critical role in the development of real-time decision support systems. As the pace of data generation accelerates and the need for instantaneous, intelligent decision-making becomes paramount across industries, the limitations of traditional static models have become increasingly apparent. Adaptive models, with their inherent ability to learn and evolve from continuous data streams, represent a fundamental shift towards more dynamic, resilient, and effective artificial intelligence.

7.1 Summary of Key Findings

Our exploration began by establishing the necessity of real-time decision support in today's data-rich world and highlighting the shortcomings of the batch learning paradigm in dynamic environments. We then presented a taxonomy of adaptive machine learning models, detailing three principal approaches:

- a. **Online Learning:** Models like Stochastic Gradient Descent and Hoeffding Trees that update incrementally with each new data point, making them highly efficient for high-velocity streams.
- b. **Continual Learning:** Techniques designed to learn a sequence of tasks without suffering from catastrophic forgetting, enabling models to accumulate knowledge over a lifetime.
- c. **Reinforcement Learning:** A paradigm where agents learn optimal decision-making policies through continuous interaction with an evolving environment.

We demonstrated the widespread and transformative impact of these models through a diverse range of applications, from algorithmic trading and real-time fraud detection in finance to clinical decision support and patient monitoring in healthcare, and navigation and control in autonomous systems. However, the path to deploying these systems is fraught with significant challenges. We identified critical hurdles including concept drift, where data distributions change over time; data-related issues like real-time quality control and labeling; model-related complexities such as scalability, evaluation in a streaming context, and the crucial need for interpretability; and system-level challenges encompassing MLOps for continuously learning models and pressing security and privacy concerns.

7.2 Concluding Remarks on the Future of Adaptive Real-Time Decision Support

The trajectory of adaptive machine learning points towards a future where intelligent systems are not merely tools for analysis but are active, autonomous partners in the decision-making process. Emerging trends such as hybrid and ensemble models, federated and edge learning, and the automation of machine learning through AutoML are set to make these systems more powerful, private, and accessible. The rise of Agentic AI further signals a move towards systems that can independently reason, act, and adapt in complex, real-world scenarios.

The transformative potential of adaptive real-time decision support is immense, promising to unlock new efficiencies, drive innovation, and solve some of the most pressing challenges in science, industry, and society. However, realizing this full potential is contingent on our ability to successfully navigate the challenges outlined in this survey. Future



research must focus not only on developing more sophisticated algorithms but also on creating robust frameworks for deploying, monitoring, and explaining these dynamic models. As we move forward, ensuring that these powerful adaptive systems are secure, private, interpretable, and aligned with human values will be the cornerstone of building a future where real-time, data-driven intelligence can be harnessed responsibly for the greater good.

References

- [1] R. Majdzadeh, "Big data revolution: transforming business landscapes through data-driven decision making," *Soc. Sci. Spectr.*, vol. 3, no. 1, 2024.
- [2] J. Nwoke, "Harnessing predictive analytics, machine learning, and scenario modeling to enhance enterprise-wide strategic decision-making," *Int. J. Comput. Appl. Technol. Res.* <https://doi.org/10.7753/IJCATR1404>, vol. 1010, 2025.
- [3] O. H. Olayinka, "Big data integration and real-time analytics for enhancing operational efficiency and market responsiveness," *Int J Sci Res Arch*, vol. 4, no. 1, pp. 280–296, 2021.
- [4] A. Bousdekis, K. Lepenioti, D. Apostolou, and G. Mentzas, "A review of data-driven decision-making methods for industry 4.0 maintenance applications," *Electronics*, vol. 10, no. 7, p. 828, 2021.
- [5] N. Rane, "Enhancing customer loyalty through Artificial Intelligence (AI), Internet of Things (IoT), and Big Data technologies: improving customer satisfaction, engagement, relationship, and experience," *Internet Things (IoT), Big Data Technol. Improv. Cust. Satisf. Engag. Relationship, Exp. (October 13, 2023)*, 2023.
- [6] G. Jumaev and M. Jonqobilov, "Real Time Decision Making and Information Extraction Algorithms Models and Applications," in *Proceedings of the 8th International Conference on Future Networks & Distributed Systems*, 2024, pp. 638–641.
- [7] M. A. S. Goraya, M. Z. Yaqub, M. A. Khan, M. S. Akram, and H. Alofaysan, "Transforming performance: how agility, response, resilience and support shape success in digital strategies," *Inf. Technol. People*, 2024.
- [8] S. F. Ahmed *et al.*, "Deep learning modelling techniques: current progress, applications, advantages, and challenges," *Artif. Intell. Rev.*, vol. 56, no. 11, pp. 13521–13617, 2023.
- [9] M. Krenn *et al.*, "Forecasting the future of artificial intelligence with machine learning-based link prediction in an exponentially growing knowledge network," *Nat. Mach. Intell.*, vol. 5, no. 11, pp. 1326–1335, 2023.
- [10] F. Liu and D. Panagiotakos, "Real-world data: a brief review of the methods, applications, challenges and opportunities," *BMC Med. Res. Methodol.*, vol. 22, no. 1, p. 287, 2022.
- [11] E. Dritsas and M. Trigka, "Machine Learning in e-Commerce: Trends, Applications, and Future Challenges," *IEEE Access*, 2025.
- [12] J. Yang, S. Wu, D. Zhang, J. Dai, F. Li, and G. Chen, "Rethinking Learned Cost Models: Why Start from Scratch?," *Proc. ACM Manag. Data*, vol. 1, no. 4, pp. 1–27, 2023.
- [13] R. R. Dos Santos, E. K. Viegas, A. O. Santin, and V. V. Cogo, "Reinforcement learning for intrusion detection: More model longness and fewer updates," *IEEE Trans. Netw. Serv. Manag.*, vol. 20, no. 2, pp. 2040–2055, 2022.
- [14] D. Pelosi, D. Cacciagrano, and M. Piangerelli, "Explainability and Interpretability in Concept and Data Drift: A Systematic Literature Review," *Algorithms*, vol. 18, no. 7, p. 443, 2025.
- [15] A. S. Adepoju, "Adaptive Program Management Strategies for AI-Based Cyber Defense Deployments in Critical Infrastructure and Enterprise Digital Transformation Initiatives".
- [16] A. S. Shethiya, "Adaptive Learning Machines: A Framework for Dynamic and Real-Time ML Applications," *Ann. Appl. Sci.*, vol. 5, no. 1, 2024.
- [17] S. Patchipala, "Tackling data and model drift in AI: Strategies for maintaining accuracy during ML model inference," *Int. J. Sci. Res. Arch.*, vol. 10, no. 2, pp. 1198–1209, 2023.
- [18] O. Gheibi, D. Weyns, and F. Quin, "Applying machine learning in self-adaptive systems: A systematic literature review," *ACM Trans. Auton. Adapt. Syst.*, vol. 15, no. 3, pp. 1–37, 2021.
- [19] V. Velayutham, S. Kumar, A. Kumar, S. Raha, and G. C. Saha, "Analysis of deep learning in real-world applications: Challenges and progress," *Tujin Jishu/Journal Propuls. Technol.*, vol. 44, no. 2, p. 2023, 2023.
- [20] S. A. Farooqi, A. Memon, S. Zamir, K. Malik, W. Batool, and H. Zahid, "Navigating AI in the real world: Transformations, regulations, and challenges," *Policy Res. J.*, vol. 2, no. 4, pp. 1083–1099, 2024.