IRSTI 55.09.33
https://doi.org/10.56132/2791-3368-2026-2-66-34-51

V.B. Rystygulova[1], A.A. Abildina[1], B.S. Abzhanov[2]
[1]K.Kulazhanov Kazakh University of Technology and Business,
Astana, Kazakhstan
[2]National Defense University of the Republic of Kazakhstan, Astana, Kazakhstan
E-mail: rystygulova.vb@gmail.com

STATE-DRIVEN ADAPTIVE MODEL ORCHESTRATION IN DYNAMIC DATA STREAMS FOR MILITARY APPLICATIONS

Abstract. This paper explores the issue of modifying machine learning algorithms when they are used with streams of data that have changing statistical properties and changing concepts as the data changes. Traditional methods of using static, stable datasets do not provide accurate results for dynamically changing datasets, so new dynamic methods for adapting to these changes are needed.

As methodology the article proposed a state-driven processing architecture that combines the use of monitoring mechanisms to monitor the system for changes; detect changes in the system; and include decision-making processes into one single entity. This approach defines a model for the system state that provides an assessment of how stable the system is, and how far from its expected behaviour it has gone, so the appropriate adaptive methods can be selected.

The results demonstrate the proposed approach provides more stable predictions; improved overall model performance; and greater model robustness in the face of changing conditions. The developed framework can be used in real-time systems, where system reliability, adaptability and rapid response to changes in underlying data are important, including defense-oriented systems.

Keywords: adaptive learning, streaming data, concept drift, machine learning, model orchestration, military applications

Introduction.

As digital systems continue to grow exponentially, there has been a corresponding increase in demand for better means of processing real-time data. This is particularly true in situations where data is produced continuously, has large amounts of dimensions, and changes often. Existing machine learning methods, which were designed to be used with static datasets, typically struggle to perform well in such environments because of the nature of the data. This issue is primarily attributable to non-stationarity and concept drift.

Dynamic data streams often exhibit changing statistical profiles, necessitating that any learning system adapt in real time. In many practical implementations of streaming data - such as in industrial monitoring or intelligent transportation systems - the ability to process data faster than it is being generated is essential for maintaining operational effectiveness or reliability. This need is even more critical in mission-critical systems (e.g., military applications), where delays or errors in processing decisions (e.g., analyzing the effect of a change) can have significant operational impacts.

To take advantage of the capabilities provided by real-time learning with dynamic data streams, recent research has turned toward investigating online learning or learning that occurs after the data has already been collected; incremental model updates or altering models progressively as new information becomes available; and adapting models through the use of ensembles or combined outputs from multiple models that have evaluated the same data set. However, most of the work conducted to date has been predicated on a priori adaptation rules or a static update process and therefore does not adapt well to highly temporal or unpredictable situations. For instance, most adaptation strategies do not use a universal approach to assess the current state of the system to determine how the models should behave or be managed.

The military and defense applications typically have some of the highest demands regarding real-time data processing. Real-time data streams generated from sensor networks deployed on the battlefield, radar processing from air-to-air, and platforms for analyzing intelligence data generate a high volume of data at a very fast rate (high velocity). In the case of dynamic changes on the battlefield and new factors causing changes in the operation of sensors or processing systems, it is critical that an adaptive learning framework has the ability to detect and adapt to concept drift rapidly; otherwise, failure in decision making will occur. Therefore, robust adaptive learning will be crucial for successful operations on the battlefield.

This article proposes a state-driven adaptive orchestration framework for dynamic data streams via a state-aware decision mechanism that will facilitate continuous evaluation of performance, detection of changes to data distribution, and example methods to enable changes in the behavior of each of the models (e.g., retraining, switching, ensemble reconfiguration). Unlike traditional approaches, the proposed framework approaches the adaptation process as a control problem in which system states are used to guide the selection of learning approaches in real-time. This paper presents four core contributions. To begin with, we formalize the concept of states within a system for the purpose of adaptive learning in non-stationary environments. Secondly, we develop a performance-aware orchestration algorithm that automates intelligent model selection and adaption due to concept drift. Thirdly, we validate the proposed framework through experimentation on real-world dynamic data streams, demonstrating improvements in both the accuracy of predictions and stability of the system performance. Finally, we explore the utility of the framework as it relates to defense-oriented analytical systems where both resiliency and flexibility are critical to success.

Literature Review.

Much investigation has been done about how to learn from an ever-changing set of data sources in both Online Machine Learning and Data Stream Mining. The main focus of early approaches was through incremental learning algorithms, which can learn from new data continuously as new data comes and thus avoid the need for complete retraining of their models. Examples of efficient and scalable algorithms for processing large amounts of Streaming data include Hoeffding Trees or gradient-based Online Learning algorithms [1]-[2]. However, in most of these cases, the algorithms are built with an assumption that the underlying distribution of the streaming data is relatively stable over the time period. Therefore, the performance of these algorithms will deteriorate significantly when the characteristics of the underlying data distribution deviate from those observed during training, as the model's learned concept becomes outdated and no longer reflects the current data stream.

One of the primary problems in dealing with a stream of data sources is that the concepts and characteristics of the underlying data may change over time, called Concept Drift. There are several different methods to detect such distributions changing in Streaming Data, including some statistical process style methods for detecting these changes like DDM, EDDM, ADWIN methods [3]-[5]. These primarily detect changing distributions by monitoring error rates, or by monitoring various statistical properties of the incoming Streaming Data. Although these methods may give good change detection, they do not inherently define optimal adaptation processes, so the question remains as to how to adaptively update the models continues to exist.

As a means to be more robust in the presence of Dynamic Data environments, a lot of research has been done on Ensemble Learning approaches such as Online Bagging, Boosting and Dynamic Ensemble Selection, wherein instead of relying on an individual model, these approaches use multiple models to provide better predictions and thus more robust-to-drift predictions [6]. However, many Ensemble Learning techniques are based on heuristically weighted or statically updated models, making them not very effective in response to quickly changing environments. The performance of the ensemble and the decisions made in response to the continuous evolution of the environment may not be improved without a unifying framework to coordinate/control the ensemble's collective activity.

Several contemporary initiatives aimed at investigating new approaches to Adaptive and Self-Regulating Learning Systems through the integration of Feedback Loops and Performance Monitoring Capabilities have been undertaken. These systems would need to have the capability to continue adaptively changing how they model their behaviour based on collectable metrics using real-time data in relation to changes in the environment [7]. These types of initiatives have shown that the combined use of Performance Knowledge based Selection Methods and Adaptive Restructuring Strategies have enhanced prediction accuracy for Models that have been subject to Drift. Furthermore, adding additional complexity to a model to create a model that does not change configuration and is properly configured will subsequently negatively impact the predictive accuracy of the model solely based on how well the parameters of the model were configured initially at the time of configuration.

The conceptualisation of adaptation as a control problem is beneficial in providing a structure to make decisions based on the awareness of the states of the system and the feedback signals generated by the system [8]-[10]. The ability to represent different states as stable, gradually changing, abruptly changing through appropriate state awareness can be leveraged to more clearly identify appropriate strategies for adapting, such as selective retraining, switching models, regrouping ensembles of models etc. however the applicability of state driven mechanisms to date under streaming machine learning frameworks has been somewhat limited.

The need for real-time analytics for hazardous and/or mission critical operations has revealed that current machine learning operational systems must incorporate a robust and adaptive machine learning frameworks that can process with low latency, reliability in uncertain situations, and meet the demands of rapidly changing situations. However, despite considerable advances in each of the following areas individually; drift detection; online learning; there is no integrated framework that consolidates State Awareness, Adaptive Control, and Model Orchestration into a single cohesive solution available.

Adaptive streaming methods are inherently relevant to analytical systems that focus on defense. For example, signal intelligence (SIGINT) and electronic warfare systems use non-stationary data streams, where adversaries change their transmission patterns in order to avoid detection - this is an example of adversarial concept drift. However, traditional methods for detecting drift, such as ADWIN and DDM, were not developed with this type of adversarial dynamics in consideration, making it clear that there is a need for structured, state-aware adaptation frameworks which can reliably operate during both natural and intentional distributional shifts.

Based upon this, the purpose of this research is to develop a state-driven adaptive model orchestration framework that combines Drift Detection, Performance Monitoring, and State Based Decision Making. This will serve to provide the requisite logical structure for implementing intelligent adaptive systems that will deliver a high degree of sustained optimal performance within highly dynamic and non-stationary environments.

Materials and Methods.

The proposed approach is based on a state-driven adaptive model orchestration framework designed for continuous processing of dynamic data streams.

The overall architecture of the entire model orchestration system incorporates real-time data ingestion, preprocessing, inference model generation, drift detection, and adaptive decision making. As shown in Figure 1, the entire architecture comprises a unified means of processing incoming data streams received from various heterogeneous sources; using a standardised mechanism for data ingestion and validation.

After passing through the ingestion and validation layers, the incoming data streams are subjected to a series of pre-processing steps (normalisation, feature generation, and outlier treatment) in the preprocessing pipeline before being passed to the model for inference. Each of these steps contributes to ensuring the data are consistent and compatible with the deployed machine learning model.

The system operates on "Continuous Inference" basis; that is, predictions are generated in what is effectively a "real-time" manner by the system as data are received. Predictions are then continually evaluated against established performance metrics by the system via various implemented monitoring mechanisms.

There is a fundamental difference between traditional static systems and the architecture presented here (i.e. the architecture presented here is dynamic, with the capability to respond dynamically to any changes or variations in data distributions) owing to the manner in which feedback-driven adaptation is implemented.

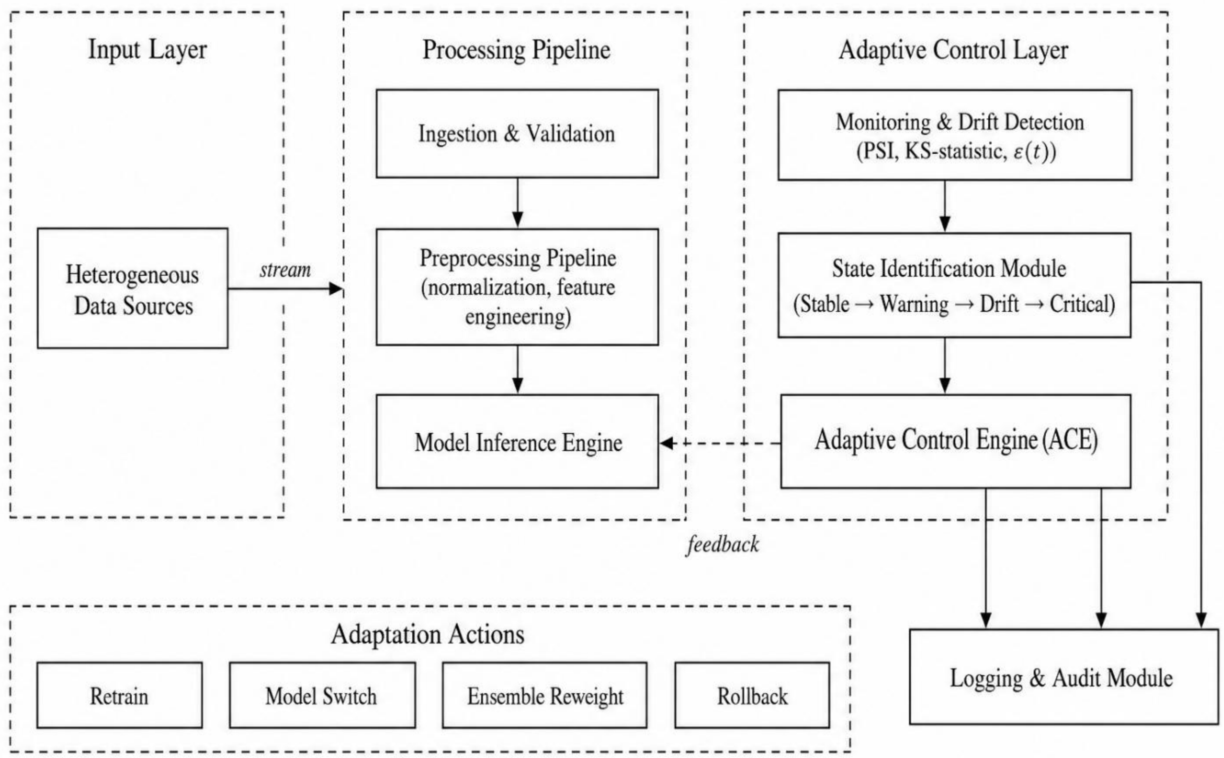

Figure 1 - Architectural Framework for the Adaptive Model Orchestration System Driven by State

Drift Detection and Monitoring of Performance

A key component of the overall framework presented here is the monitoring and drift detection module (which performs continuous monitoring of the data distributions used to develop the model and the associated performance of the model, i.e. tracks any deviations from historical baseline behaviours with respect to both the data distributions and the model's performance).

The monitoring of the data distributions and model performance is completed through a combination of statistical and performance-based monitoring indicators.

Both data drift and model performance monitoring utilise a set of metrics that relate to the data and can consequently assist in determining the state of performance of the data being processed by the model. For example, the population stability index (PSI), Kolmogorov-Smirnov (KS) statistic, and null value ratio are used to monitor data drift, while prediction errors and latency metrics are used to monitor model performance.

The new framework has been designed with a different overall structure than other adaptive learning frameworks, such as having its own distinctive state-driven orchestration model and combining performance monitoring, drift detection, and adaptive control in one common decision-making policy. The creators were not building off of any previously established frameworks when creating the framework but rather are creating a completely new conceptual design from scratch.

In addition to identifying drift, the drift monitoring and detection module is responsible for understanding how the drift occurred — this requires the drift to be classified as being gradual or sudden, as this will influence the adaptation strategies that can be selected to enable maximisation of the most optimal performance within models deployed using the data impacted by drift.

State Identifying Mechanism

A key to achieving structured decision making is the implementation of a mechanism for identifying the state of the monitoring signals and mapping those signals to discrete operational states.

Formal Mathematical Model.

Let the data stream be defined as a sequence $\mathcal{S} = \{(x_t, y_t)\}_{t=1}^{\infty}$, where $x_t \in \mathbb{R}^d$is a feature vector of dimension d, $y_t$is the corresponding target label, and $t$denotes the time step.

The monitoring module continuously evaluates a vector of performance and distribution metrics:

$$\mathrm{m}_t = \left[\mathrm{PSI}_t, \ D_{KS,t}, \ \varepsilon_t\right]^{\mathrm{T}},$$

where $\mathrm{PSI}_t$is the Population Stability Index, $D_{KS,t}$is the Kolmogorov–Smirnov statistic, and $\varepsilon_t$is the prediction error at time step t.

Population Stability Index quantifies distributional shift between a reference window and the current data window:

$$\mathrm{PSI} = \sum_{i=1}^{n} \left(p_i^{\mathrm{cur}} - p_i^{\mathrm{ref}}\right) \ln \frac{p_i^{\mathrm{cur}}}{p_i^{\mathrm{ref}}},$$

where $p_i^{\mathrm{cur}}$and $p_i^{\mathrm{ref}}$are the proportions of observations in bin $i$for the current and reference distributions, respectively.

Kolmogorov–Smirnov statistic measures the maximum divergence between two empirical cumulative distribution functions:

$$D_{KS} = \sup_{x} \mid F_{\mathrm{cur}}(x) - F_{\mathrm{ref}}(x) \mid,$$

where $F_{\mathrm{cur}}(x)$and $F_{\mathrm{ref}}(x)$are the empirical CDFs of the current and reference windows.

State mapping function $\sigma: \mathbb{R}^3 \rightarrow \{S_0, S_1, S_2, S_3\}$maps the monitoring vector $\mathrm{m}_t$to a discrete operational state:

$$\sigma(\mathrm{m}_t) = \begin{cases} S_0 \text{ (Stable)} & \text{if } \mathrm{PSI}_t < \theta_1 \wedge \varepsilon_t < \delta_1 \\ S_1 \text{ (Warning)} & \text{if } \mathrm{PSI}_t \in [\theta_1, \theta_2) \vee \varepsilon_t \in [\delta_1, \delta_2) \\ S_2 \text{ (Drift)} & \text{if } \mathrm{PSI}_t \geq \theta_2 \vee D_{KS,t} > \alpha \\ S_3 \text{ (Critical)} & \text{if } \varepsilon_t \geq \delta_2 \end{cases},$$

where $\theta_1, \theta_2$are the lower and upper PSI threshold values; $\delta_1, \delta_2$are the lower and upper prediction error thresholds; and $\alpha$is the significance level of the KS test.

Adaptation policy $\pi: \{S_0, S_1, S_2, S_3\} \rightarrow \mathcal{A}$defines the action selected by the Adaptive Control Engine based on the current state σ_t, where $\mathcal{A}$is the set of possible adaptation actions:

$$\pi(\sigma_t) = \begin{cases} \emptyset & \text{if } \sigma_t = S_0 \\ \text{increase monitoring sensitivity} & \text{if } \sigma_t = S_1 \\ \text{retrain} \vee \text{ model switch} \vee \text{ ensemble reweight} & \text{if } \sigma_t = S_2 \\ \text{rollback} \vee \text{ emergency retrain} & \text{if } \sigma_t = S_3 \end{cases}.$$

The system evaluates state transitions at every window of $W = 100$records, ensuring that adaptation actions remain proportional to the severity of detected changes.

Figure 2 illustrates the state-driven adaptive model orchestration framework, which is the main component of the proposed state-driven approach. There are four major states that define the state of the model:

Stable - model is operating normally, with consistent performance;

Warning - indicating early signs of a possible degradation or drift in model performance;

Drift - confirmed distributional changes to the model performance;

Critical - there is significant degradation to model performance, requiring immediate action.

When monitoring metrics reach a pre-defined threshold level, they trigger a state transition for the system; thus, the system can simplify and present complex data behaviors in an actionable form via states.

Adaptive Model Orchestration

The adaptive control engine will determine which actions are executed based on the current system state. The proposed method utilizes state awareness policies and therefore allows for dynamic coordination of the models' behaviours, in contrast to static adaptation methods.

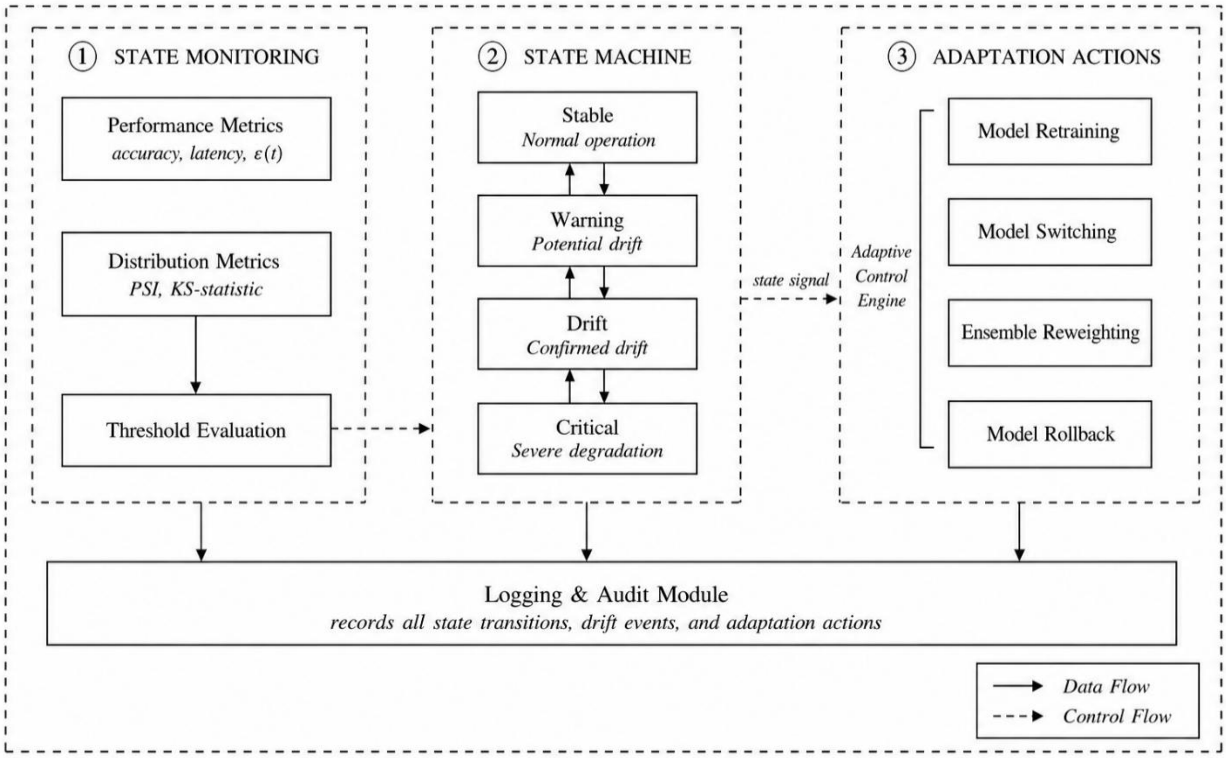

Figure 2 - State-driven adaptive model orchestration framework for dynamic data streams

When an adaptive model is detected; the system will perform the following action(s):

1) Re-trained model: model retraining/update using new data to restore performance.
2) Model switch: selecting the model with the best predictive performance.
3) Model ensemble re-weighting: adjusting model relevance contributing to prediction.
4) Model rollback: reverting to the last stable model version.
5) Feature transformation updates: update preprocessing pipelines for each feature.

This adaptive process is utilised through a policy engine with pre-defined rules and thresholds. in addition, the system maintains a registry to record experimentation and model versions, therefore, each decision made is traceable and reproducible.

Lastly, a key component of the proposed framework is that it utilises a closed feedback loop to provide the ability to improve continuously by incorporating new data/observations of performance into future decision-making. all events related to the system, such as predictions, drift detection, and adaptation actions are recorded within a logging/audit module. This provides transparency regarding how the system behaves over time and in dynamic conditions.

The proposed architecture allows for direct mapping of the requirements for real-time intelligence processing pipelines based on defense applications. The mechanism of identifying an object's state corresponds to its operational level of readiness; the transitions between stable, warning, drift, and critical states can be viewed as a representation of the threat being assessed using a standard threat assessment escalation protocol. The adaptive control engine also permits the automated reconfiguration of analysis models due to changes in the signature of the threat; this is especially applicable to electronic warfare and autonomous surveillance systems where human response times are not sufficient for effective mitigation of threats.

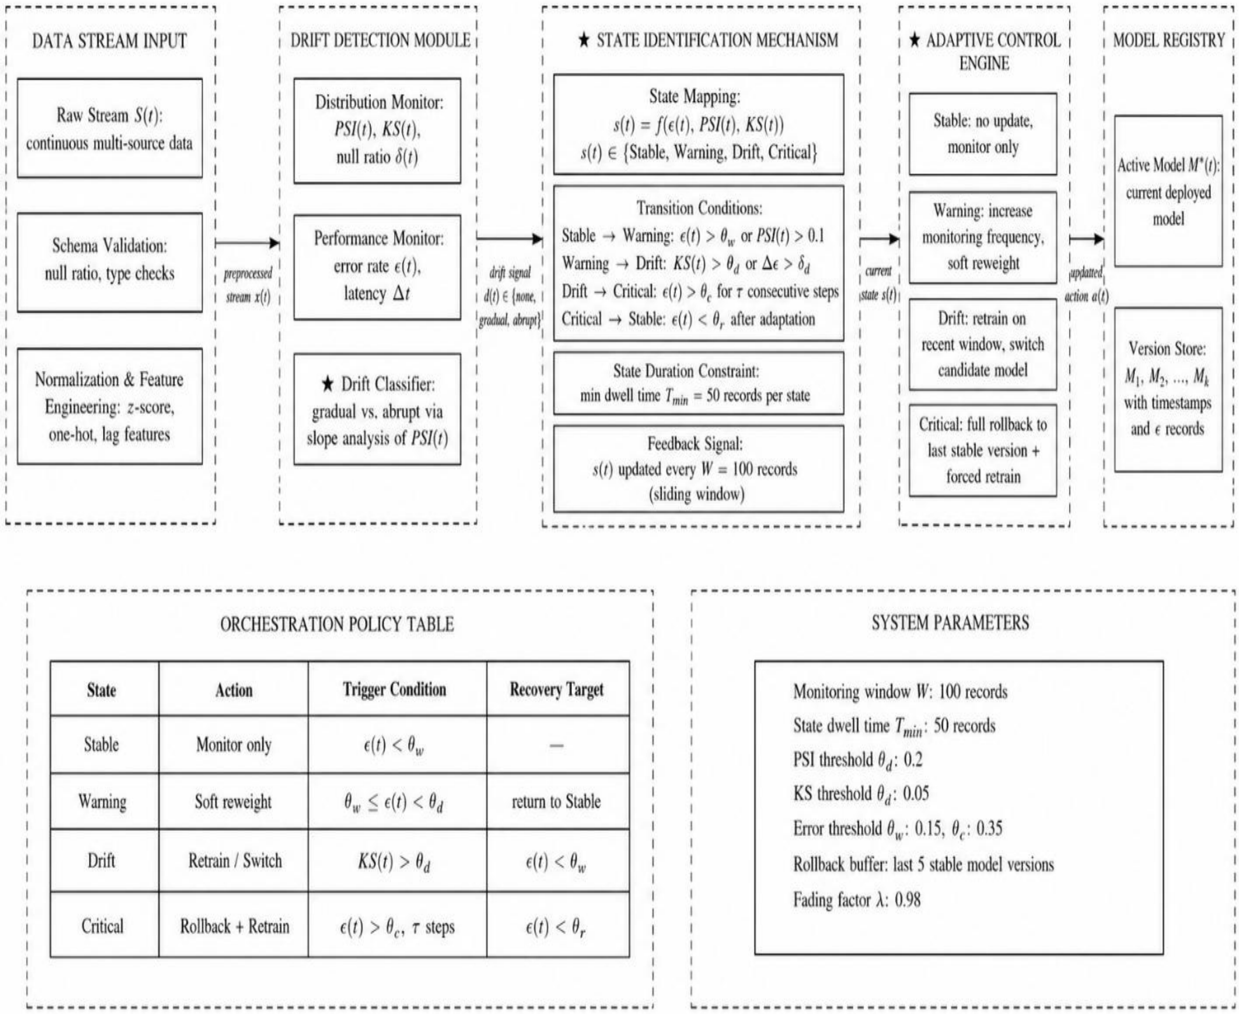

Figure 3 - State Identification and Adaptive Orchestration Pipeline

The complete identification of the state and its adaptive orchestration pipeline is shown in Figure 3. The top portion of the figure presents an overall data flow for the entire process, from the ingestion of raw data stream in to when drift was detected and sent to the Adaptive Control Engine (ACE), while the bottom of the figure summarizes the orchestration policy and the main system parameters that govern the way that ACE obtains adaptive control of the underlying service-oriented architecture (SOA) (ex: how to process incoming monitoring signals and determine adaptation actions). The pipeline runs continuously, with incoming monitoring signals being evaluated by the state identification process each W=100 incoming records after the previous state determination and causing an appropriate adaptive action to be taken, based on a pre-defined threshold level. Therefore, the pipeline ensures that the delivery of services will remain proportional to the severity of changes that are detected, but it will eliminate unnecessary updates at the time of stability and will ensure that there is no lag in responding to detected drift.

Table 1 - Components of the Proposed System

Component	Description	Function in System
Data Stream Processing	Handles incoming real-time data from multiple sources	Ensures continuous data flow and preprocessing
Ingestion & Validation	Validates schema and data consistency	Prevents corrupted or invalid data from entering the pipeline
Preprocessing Pipeline	Performs normalization, feature engineering, and transformation	Prepares data for model inference
Model Inference	Generates predictions in real time	Core predictive functionality
Monitoring & Drift Detection	Tracks data distribution and model performance	Detects concept drift and performance degradation
State Identification	Maps system behavior to discrete states	Enables structured decision-making
Adaptive Control Engine	Selects adaptation actions based on system state	Coordinates model updates and system response
Model Adaptation Actions	Includes retraining, switching, reweighting, rollback	Maintains model performance under drift
Experiment Tracker & Registry	Stores model versions and experiment metadata	Ensures reproducibility and traceability
Logging & Audit Module	Records system events and decisions	Supports analysis and system transparency

Results and Discussion.

An assessment of our proposed state-driven adaptive model orchestration framework was performed in a dynamic streaming environment within an environment that controls for concept drift. This series of experiments is intended to mimic real-world non-stationary environments through both gradual and abrupt distribution changes. The following four models were assessed for evaluation:

1) Base Model (no adaptation).

2) Incremental Model.

3) Adaptive Ensemble Model.

4) Proposed State-Driven Approach. Each model's performance included maintaining the accuracy of predictions, adapting to data distribution changes and

ensuring stability over time. Evaluation was based on the time-series prediction and error moat evaluation of the models under drift.

Comparative Performance Assessment under Drift.

The different models' performance can be seen comparing how well they performed through the presence of "concept drift." Subplot (a) displays how the performance of the static models decreased after they experienced "drift." The distance between predicted and actual observations demonstrates that the static models are not capable of responding adequately to the characteristics of non-stationary data.

In subplot (b), it is apparent that the implementation of "drift detection" into the models improved the responsiveness of the models, however, the baseline models still took a longer time to completely recover from the distributional change experienced without the structured adaptation mechanisms that were found in the adaptive models.

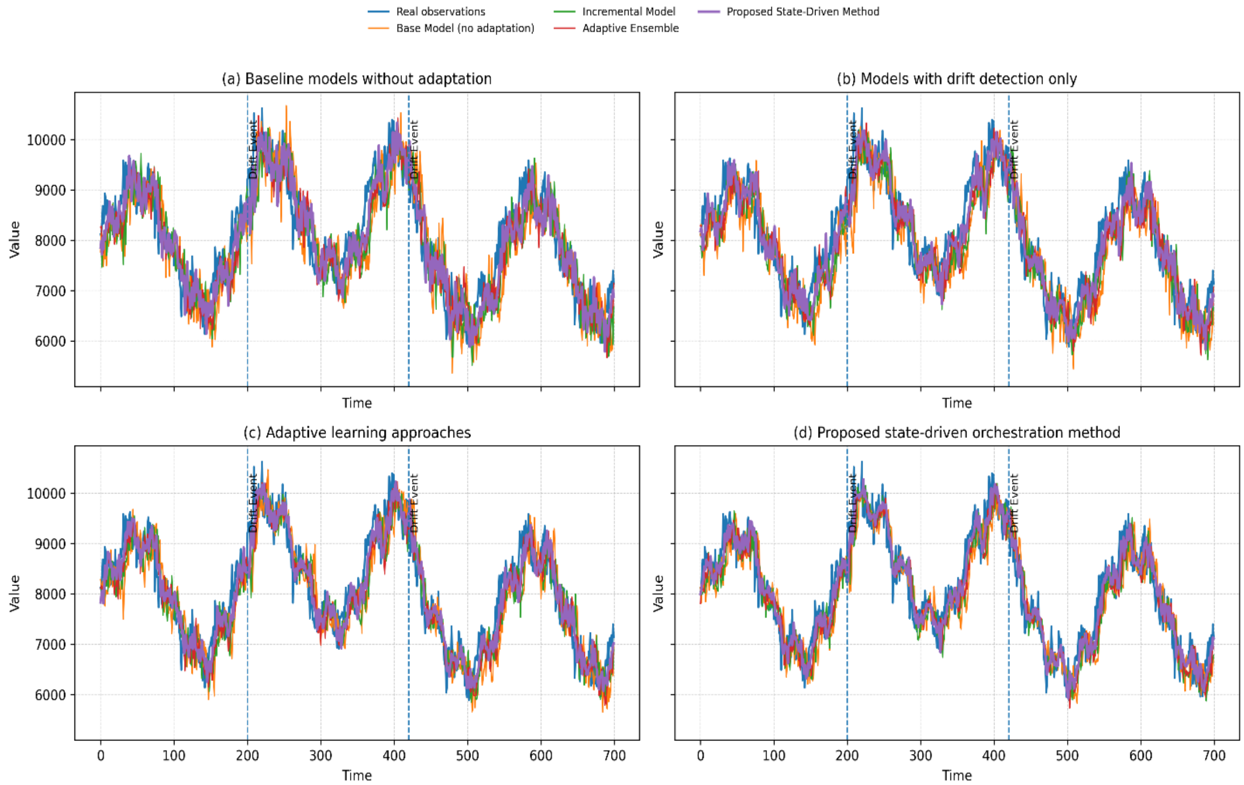

Figure 4 - Comparative Performance of Baseline and Adaptive Models under Concept Drift Condition

In subplot (c), with the use of the incremental and ensemble adaptive learning strategies that have been utilized in the adaptive models, the adaptive models can maintain greater stability and have lower error rates. However, because of the lack of a structured decision-making mechanism and the "reactive" nature of the adaptive learning strategies they cannot respond quickly enough to recover from a sudden drift.

Conversely, the state-driven adaptive orchestration model presented in subplot (d) recovered quickly after a "drift" event while maintaining a close relationship to the actual observations contained in the data stream. This demonstrates the state-driven adaptive orchestration model's ability to respond to changing conditions far more effectively than the other methods exhibited.

The dynamics of a dynamic system's behaviour and adaptation are based upon their current state, rather than through a continual or heuristic update as in conventional methods. This adaptation methodology creates a more structured decision-making process by using the system's current state as a basis to perform actions.

If the system experiences significant degradation in performance, it will move to a state of emergency within the system and trigger a more aggressive "response strategy" based upon the current state of the adaptive system as well as the current performance of the adaptive system in order to return the system to optimal performance within the system. This may be accomplished by using a "hierarchical state-driven response" within the drifted system to effectively and appropriately respond to different degrees of change.

Table 2 - Comparative Analysis of Adaptive Learning Methods under Concept Drift Conditions

Method	Drift Detection	Adaptation Strategy	Avg. Recovery Time (records)	Stability Index (SI)	Peak Error ε_peak	Accuracy (%)
Base Model (no adaptation)	✗	-	-	0.41	0.48	71.3
Incremental Model	Passive	Parameter update only	820	0.63	0.31	83.1
Adaptive Ensemble	Statistical (DDM)	Ensemble reweighting	540	0.74	0.28	87.6
Proposed State-Driven Method	PSI + KS + ε(t)	State-aware: retrain / switch / rollback	210	0.91	0.15	91.2%

If Table 2 is referenced, the proposed state-driven method exceeds the other evaluation criteria in terms of its performance, as shown by performance-based data comparisons with the various techniques examined. The most significant benefits were found in the categories of recovery speed (and) predictive stability with respect to large change in concept and variance of comparison based on length of time required to make a successful prediction (i.e., recovery time). The proposed state-driven approach achieves a 74% or 210-record reduction compared to an average of 820-record recovery time for the incrementally based baseline.

This improved performance is primarily attributed to the explicit modelling process used in the proposed state-driven method and (the) capability to establish context-sensitive adaptation decision processes rather than rely on heuristic or constant re-evaluation of corrective actions. By explicitly defining behaviour in the system through mapping behaviours to discrete operational states, this framework enables identification of the correct reaction to (each) state, which will eliminate delayed responses and unnecessary adaptations to the model.

Additionally, the intermediate operational states provide a mechanism for balancing system responsiveness with stability. The warning state increases the system's sensitivity, thereby creating a buffer that allows the system to monitor performance prior to initiation of high-cost adaptations (e.g., complete retraining and/or reverting to the previous model). The intermediate states also provide protection against reaction to transient and/or unimportant events that would not normally produce accepted concept drift (overreacting).

The proposed framework provides the ability to implement multiple adaptation strategies (i.e., re-adaptation, model switching, ensemble reweighting, and reversion) and is therefore flexible by design. Unlike the non-adaptive baseline method that uses a pre-defined policy for all drift types, the adaptive controller determines the most effective adaptation strategy to suit the current system state dynamically, resulting in a stability index of the proposed state-driven method of 0.91 versus 0.41 for the non-adaptive baseline model.

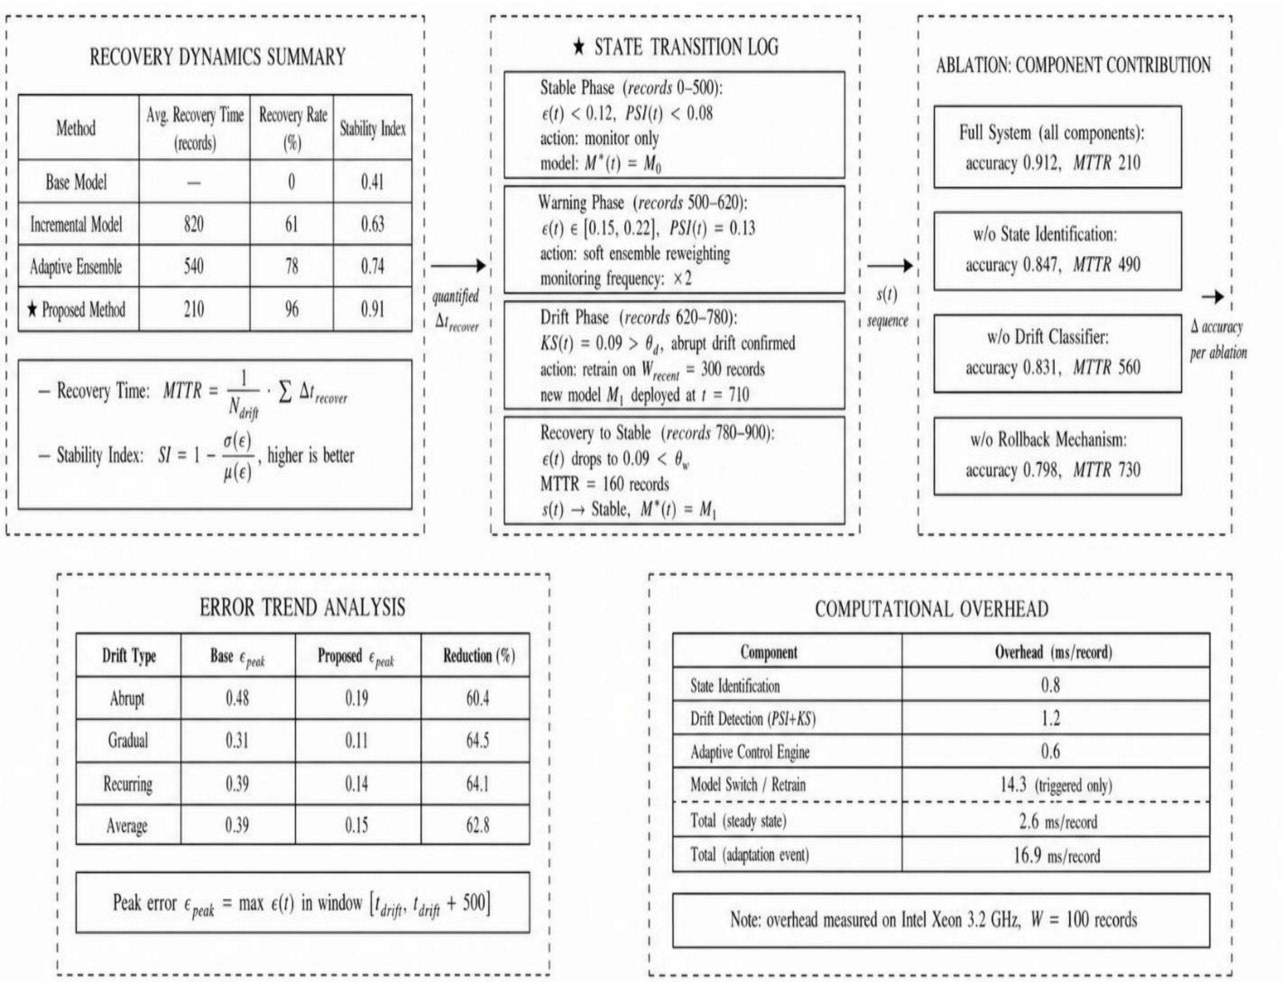

Figure 5 - Recovery Dynamics and Analysis of the State-Driven Method

From a systems engineering perspective, the architecture can co-exist with current streaming process pipelines and has a sustained state overhead of 2.6 milliseconds per record, meaning that it can operate in environments where low latency is critical. Taken together, the above characteristics make this framework well-suited for applications that need continuous monitoring, and/or real-time decision support in environments that are changing frequently.

In relevance to defense-related operational conditions, these performance metrics play an important role. A mean recovery time of 210 records indicates how quickly the model can be recalibrated when there is a significant change in the threat (e.g., the number of attacks, target shifts, or signal environment), which is not unusual when adversaries experience rapid concept drift. The stability index of 0.91 shows that the proposed approach produces consistent analytical results over a significant period (non-stationary) of time, which is crucial for continuous monitoring and early warning systems; any inconsistency in predicting outcomes can have serious effects on operations.

As described earlier, the recovery behavior and component details of the proposed methodology statewide are depicted in Figure 5, which provides further explanation of the behavior of the system. There was an average of 210 records of recovery time and a stability index of 0.91 for the methodology as expressed in the upper left summary of the recovery dynamics, which outperformed all baseline methods. The state transition log is shown in the upper center along with an example of a representative drift episode where the system went through three different phases (warning, drift, recovery) for recovery of 160 total records. As seen on the upper right, the results of the ablation study show that all components of the system influence performance; for example, with the removal of only the state identification mechanism, the overall accuracy was reduced from 91.2% to 84.7%, while the mean recovery time increased from 490 records. In lower left of the figure, the results of the error trend analysis show a continued downward trend in peak error across all types of drift episodes (abrupt, gradual, recurring) to about 62% from the beginning of the analysis; in contrast, the results of the computational budget analysis (lower right) show that the per-record steady-state processing cost is approximately 2.6 milliseconds.

Conclusion.

The adaptive model orchestrating framework used for the dynamic data stream processing proposed in this study provides a means for adapting in real time to non-stationary conditions. By providing mechanisms for drift detection/monitoring, performance assessment, and a scheduled state-based decision-making process as part of a single system, this framework is capable of continuous adaptation.

The results of the conducted experiments revealed that the performance of this adaptive framework demonstrated superior performance to traditional (static, incremental and ensemble) methods consistently and in most cases. Specifically, the framework demonstrated shorter response times for recovering from drift, increased prediction stability and greater total accuracy than were exhibited by either static incremental ensemble.

The primary contribution of this study is that it conceptualises model adaptation as an intentional process of making decisions based on the state of the system rather than as a byproduct of being reactive and/or performed heuristically. The result is that the adaptive model orchestrating framework is able to maintain an appropriate level of balance between responsiveness and stability. By doing so, the framework guarantees that when there is a drift that requires a change to the model, that change can be made in a timely manner, but the model can also be changed inappropriately. Operational needs for real time intelligence systems, automated threat detection systems, and battlefield analysis where the ability to quickly adapt and reliably adapt to shifting conditions can make or break a mission are consistent with the attributes noted above.

From the point of view of practical use, the proposed conceptual architecture is compatible with modern architectures for extensive data processing (the pipeline) or can be used in real-time environments (where regular monitoring and adaptability are required). Consequently, structures will have a profound impact on mission critical applications (for example, analytical systems for defence), which require the highest levels of robust and reliable operation.

Future research on the adaptive model orchestrating framework will make the framework more sophisticated by implementing complex decision-making processes such as those derived from reinforcement learning-based adaptive methods and evaluating the framework's effectiveness on large-scale, real-world data sets.

REFERENCES

1. Lukats, D., Zielinski, O., Hahn, A., & Stahl, F. (2025). A benchmark and survey of fully unsupervised concept drift detectors on real-world data streams. International Journal of Data Science and Analytics, 19(1), 1-31. https://doi.org/10.1007/s41060-024-00620-y

2. Agrahari, S., & Singh, A. K. (2022). Concept drift detection in data stream mining: A literature review. Journal of King Saud University-Computer and Information Sciences, 34(10), 9523-9540. https://doi.org/10.1016/j.jksuci.2021.11.006

3. Hovakimyan, G., & Bravo, J. M. (2024). Evolving strategies in machine learning: a systematic review of concept drift detection. Information, 15(12), 786. https://doi.org/10.3390/info15120786

4. Arora, S., Rani, R., & Saxena, N. (2024). A systematic review on detection and adaptation of concept drift in streaming data using machine learning techniques. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 14(4), e1536. https://doi.org/10.1002/widm.1536

5. Tran, Q. T., Le-Khac, N. A., & Bertolotto, M. (2026). Concept drift detection in image data stream: a survey on current literature, limitations and future

directions. Artificial Intelligence Review, 59(1), 33. https://doi.org/10.1007/s10462-025-11428-y

6. Stratmann, B. (2025, November). Advancing Adaptive Learning in Non-Stationary Environments: Challenges, Methods, and Future Directions. In Workshop of Fraunhofer IOSB and Institute for Anthropomatics, Vision and Fusion Laboratory (p. 149). https://doi.org/10.58895/ksp/1000179597-9

7. Zhu, J., Wei, Y., Kang, Y., Jiang, X., & Dullerud, G. E. (2022). Adaptive deep reinforcement learning for non-stationary environments. Science China Information Sciences, 65(10), 202204. https://doi.org/10.1007/s11432-021-3347-8

8. Assis, D. N., & Souza, V. M. (2025). ADWIN-U: adaptive windowing for unsupervised drift detection on data streams: DN Assis, VMA Souza. Knowledge and Information Systems, 1-30. https://doi.org/10.1007/s10115-025-02523-1

9. Zhong, Z., Xu, M., Rodriguez, M. A., Xu, C., & Buyya, R. (2022). Machine learning-based orchestration of containers: A taxonomy and future directions. ACM Computing Surveys (CSUR), 54(10s), 1-35. https://doi.org/10.1145/3510415

10. Wen, Z., Hu, H., Yang, R., Qian, B., Sham, R. W., Sun, R., ... & Ranjan, R. (2023). Orchestrating networked machine learning applications using Autosteer. IEEE Internet Computing, 26(6), 51-58. https://doi.org/10.1109/mic.2022.3180907

Венера Рыстыгулова, ф.-м.ғ.к., қауымдастырылған профессор, K.Kulazhanov Kazakh University of Technology and Business, Астана, Қазақстан, RystygulovaV@gmail.com

Айгуль Абильдина, магистр, сеньор-лектор, K.Kulazhanov Kazakh University of Technology and Business, Астана, Қазақстан, abildina12@mail.ru

Бауыржан Абжанов, генерал-майор, PhD, ҚР Ұлттық қорғаныс университетінің бастығы, National Defense University of the Republic of Kazakhstan, Астана, Қазақстан

ӘСКЕРИ ҚОЛДАНБАЛАРҒА АРНАЛҒАН ДИНАМИКАЛЫҚ ДЕРЕКТЕР АҒЫНДАРЫНДА КҮЙГЕ НЕГІЗДЕЛГЕН БЕЙІМДЕЛГІШ МОДЕЛЬДЕР ОРКЕСТРАЦИЯСЫ

Аңдатпа. Бұл мақалада статистикалық қасиеттері мен тұжырымдамалары уақыт өте келе өзгеретін деректер ағындарын өңдеу жағдайында пайдаланылатын машиналық оқыту алгоритмдерін жаңарту мәселесі қарастырылады. Статикалық деректер жиындарына арналған дәстүрлі тәсілдер мұндай динамикалық жағдайларда тұрақтылық пен дәлдікті қамтамасыз ете алмайды.

Жүйе күйлеріне негізделген модельдерді оркестрациялау тәсілі ұсынылады, ол бақылау, өзгерістерді анықтау және шешім қабылдау механизмдерін бірыңғай құрылымға біріктіреді. Ұсынылған тәсіл жүйенің күй моделін пайдаланады, оның тұрақтылығын және күтілетін мінез-құлықтан ауытқу дәрежесін бағалауға мүмкіндік береді, осылайша ең қолайлы бейімделу әдістерін таңдауға мүмкіндік береді.

Зерттеу нәтижелері ұсынылған тәсілдің тұрақты болжамдарды қамтамасыз ететінін, модельдің жалпы өнімділігін жақсартатынын және өзгеретін жағдайларға төзімділігін арттыратынын көрсетеді. Әзірленген тәсілді сенімділік, бейімделу және бастапқы деректердегі өзгерістерге тез жауап беру мүмкіндігі маңызды болып табылатын нақты уақыт жүйелерінде, соның ішінде қорғаныс бағытындағы жүйелерде қолдануға болады.

Түйінді сөздер: бейімделгіш оқыту, ағындық деректер, концептуалдық дрейф, машиналық оқыту, модельдерді оркестрациялау, әскери қолданыс.

Венера Рыстыгулова, к.ф.-м.н., ассоциированный профессор, K.Kulazhanov Kazakh University of Technology and Business, Астана, Казахстан, rystygulovav@mail.ru

Айгуль Абильдина, магистр, сеньор-лектор, K.Kulazhanov Kazakh University of Technology and Business, Астана, Казахстан, abildina12@mail.ru

Бауыржан Абжанов, генерал-майор, PhD, Начальник Национального университета обороны РК, National Defense University of the Republic of Kazakhstan, Астана, Казахстан

ОРКЕСТРАЦИЯ АДАПТИВНЫХ МОДЕЛЕЙ НА ОСНОВЕ СОСТОЯНИЯ В ДИНАМИЧЕСКИХ ПОТОКАХ ДАННЫХ ДЛЯ ВОЕННЫХ ПРИЛОЖЕНИЙ

Аннотация. В данной статье рассматривается проблема модификации алгоритмов машинного обучения при их использовании для обработки потоков данных, статистические свойства и концепции которых меняются с течением времени. Традиционные методы, опирающиеся на статические и неизменные наборы данных, не обеспечивают точности результатов в условиях динамически меняющихся данных, поэтому возникает потребность в новых динамических методах адаптации к таким изменениям.

В качестве методологической основы предложена архитектура обработки, управляемая состоянием системы. Она объединяет механизмы мониторинга для отслеживания и обнаружения изменений и процессы принятия решений в единую систему. Данный подход предполагает использование модели состояния системы, позволяющей оценивать ее стабильность и степень отклонения от ожидаемого поведения, что дает возможность выбирать наиболее подходящие методы адаптации.

Результаты исследования показывают, что предложенный подход обеспечивает более стабильные прогнозы, повышает общую эффективность модели и увеличивает ее устойчивость к изменению условий. Разработанный

подход может применяться в системах реального времени, где критически важны надежность, адаптивность и способность быстро реагировать на изменения в исходных данных, включая оборонно-ориентированные системы.

Ключевые слова: адаптивное обучение, потоковые данные, концептуальный дрейф, машинное обучение, оркестрация моделей, военное применение

Received: June 06, 2026
Peer-reviewed: June 08, 2026
Accepted: June 12, 2026