Electronics and Information Technologies, 2026, 34, 65–82
http://publications.lnu.edu.ua/collections/index.php/electronics/index

UDC: 004.41

CLOUD-AGNOSTIC INTELLIGENT MONITORING ARCHITECTURE FOR CLIMATE ANOMALIES IN EARLY WARNING SYSTEMS

Vasyl Lyashkevych *, ***Yurii Marchuk***, ***Petro Kulyk***, ***Andrii Patynko***, ***Taras Kerod***
Department of System Design
Ivan Franko National University of Lviv,
50 Drahomanova Str., Lviv, 79005, Ukraine

Lyashkevych, V., Marchuk, Y., Kulyk, P., Patynko, A., & Kerod, T. (2026). Cloud-agnostic Intelligent Monitoring Architecture for Climate Anomalies in Early Warning Systems. *Electronics and Information Technologies*, 34, 65–82. https://doi.org/10.30970/eli.34.4

ABSTRACT

Background. Climate-related hazards require early warning systems to combine heterogeneous environmental observations, detect anomalous spatiotemporal patterns, and disseminate warning signals to stakeholders. Existing solutions separate software architecture, anomaly-detection models, and decision logic to transform deviations into early signals.

Materials and Methods. The study develops a cloud-agnostic intelligent monitoring architecture for early warning systems, supported by ML-based anomaly identification and decision support. A hierarchical anomaly model is introduced for climate zones, areas, and regions. It evaluates four types of evidence: deviation from normal state, consistency with neighboring regions, consistency with parent area, and the influence of contextual factors. The method combines region-specific detector training, multi-component anomaly scoring, persistence-based warning confirmation, and CEP-based escalation. A synthetic spatiotemporal climate graph was used to compare the Robust Z-score statistical method, the Isolation Forest anomaly-detection algorithm, the One-Class Support Vector Machine method, and the proposed HCOIM-EWS method, where HCOIM-EWS denotes Hierarchical Context-Orchestrated Intelligent Monitoring for Early Warning Systems.

Results and Discussion. The proposed architecture integrates ML-based detector selection, region-specific training, hierarchical anomaly scoring, CEP-based warning logic, and notification distribution. In simulation, HCOIM-EWS achieved the lowest false-alarm count while maintaining comparable event-level recall. Compared with Robust Z-score, Isolation Forest, and One-Class SVM, the proposed method reduced false alarms by 55.0%, 22.6%, and 49.7%. The method is suitable for warning-oriented intelligent monitoring, where operational reliability and false-alarm reduction are as important as point-level sensitivity.

Conclusion. The proposed method extends anomaly detection in early warning systems from isolated local time-series analysis to hierarchical, context-aware, and event-driven intelligent monitoring. This confirms the technical feasibility of the architecture and justifies its use as a reusable template for climate-risk monitoring platforms.

Keywords: anomaly detection, early warning systems, intelligent monitoring, cloud-agnostic architecture, complex event processing, machine learning.

INTRODUCTION

Early warning systems (EWSs) are key instruments of disaster risk reduction because they support timely monitoring, forecasting, communication, and coordinated response to hazardous events [1].

Anomaly detection identifies patterns that deviate from expected behavior, including outliers, novelties, exceptions, or potentially dangerous states [2, 3]. In environmental and meteorological settings, anomaly detection becomes a basis for identifying deviations that may precede floods, storms, droughts, heatwaves, or other natural hazards [3]. Recent ML-based studies also demonstrate the growing role of artificial intelligence in modeling extreme weather events, quality control of meteorological observations, and heavy-rainfall observation analysis [4-6].

This paper uses weather singularities as a motivating class of climate-related deviations and studies their identification through hierarchical anomaly monitoring. Regardless of terminology, the core problem is how to build a scalable system that ingests heterogeneous climate signals, synchronizes them in time and space, trains and deploys multiple anomaly detectors, and transforms anomaly evidence into actionable warning signals.

Recent studies show that EWSs are moving from isolated forecasting tools to modular, digital, and intelligent architectures [7], with an emphasis on interoperability, resilience, governance, and integration of analytical and distributive workflows [8, 9]. Their practical implementation increasingly relies on scalable big data processing, automated machine learning, and climate zone segmentation to support distributed model training and region-specific analysis [10–12]. At the same time, anomaly detection in environmental monitoring is evolving from purely local time series analysis to spatiotemporal and context-dependent formulations, where a local anomaly is interpreted in relation to neighboring regions and broader climate regimes [13–15].

In this study, intelligent monitoring is understood as a continuous process that combines data acquisition, anomaly identification, contextual interpretation, and warning-oriented decision support. Such a perspective is consistent with recent work on intelligent monitoring, optimization of analytical pipelines by machine learning and genetic algorithms, and scenario-based environmental modeling [16–18].

This work addresses the lack of an integrated cloud-agnostic intelligent monitoring architecture that combines ML-based anomaly identification, warning-oriented decision support, hierarchical anomaly modeling, and EWS reference architecture. The objective is to develop and evaluate a cloud-agnostic intelligent monitoring architecture supported by the author-defined HCOIM-EWS method, which connects ML-based anomaly detection, warning-oriented decision logic, and reusable deployment principles for early warning systems. In this study, HCOIM-EWS denotes Hierarchical Context-Orchestrated Intelligent Monitoring for Early Warning Systems and is not used as a standardized or officially established IT-industry term. The study demonstrates the feasibility of the method through simulation on a synthetic hierarchical climate graph. The main contributions are: (i) a formal hierarchical anomaly model for climate-anomaly monitoring; (ii) the author-defined HCOIM-EWS method for transforming anomaly candidates into warning-level evidence; (iii) a simulation-based comparison with three baseline anomaly-identification methods; and (iv) a cloud-agnostic reference architecture for early warning systems.

MATERIALS AND METHODS

The proposed solution assumes a hierarchically organized monitored space composed of climate zones, areas, and regions. For each region, the system maintains synchronized multivariate observations, local and neighboring context, and higher-level climatic aggregates. The event-driven cloud-agnostic architecture integrates distributed

APIs, Kafka, Big Data storage and preprocessing, AutoML training, and CEP-based warning fusion.

Problem formulation and decision-support objective

The problem addressed in this study is the detection of anomalous deviations associated with weather singularities that may serve as early indicators of natural hazards. Unlike conventional forecasting systems focused on isolated regional observations, the proposed method assumes that weather conditions in a given region may be influenced by patterns observed in neighboring regions and broader climatic domains. Accordingly, anomaly detection should not be restricted to a single local time series but should incorporate both vertical scaling across broader and narrower climatic levels and horizontal orchestration across adjacent regions.

Formal hierarchical anomaly model

Let R denote the set of monitored regions, A – set of climate areas, and Z – set of climate zones. For each region $r \in R$ and time moment t, the monitoring system observes a synchronized multivariate environmental vector $x_r(t) \in R^m$. Let $N(r)$ denote the set of horizontally neighboring regions and $P(r)$ denote the parent climatic area or zone associated with region r.

The normal behavior of a region is represented by the expected state vector $\widehat{x_r}(t)$, estimated from historical synchronized observations, temporal context, and higher-level climatic context. We define a regional anomaly as a deviation that simultaneously affects at least one of four dimensions: local deviation, horizontal coherence, vertical coherence, or contextual stress, which represents deviation of external environmental drivers from their expected baseline:

$$D_{loc}(r,t) = \left|\left|x_r(t) - \hat{x}_r(t)\right|\right|_{\Sigma_r^{-1}}, \tag{1}$$

where Σ_r is the covariance matrix of normal regional behavior,

$$D_{hor}(r,t) = \frac{1}{|N(r)|}\sum_{n=1}^{N(r)} w_{rn}\left|\left|x_r(t) - x_n(t)\right|\right|_2, \tag{2}$$

where w_{rn} is the strength of spatial linkage between neighboring regions,

$$D_{ver}(r,t) = \left|\left|x_r(t) - g_{P(r)}(t)\right|\right|_2, \tag{3}$$

where $g_{P(r)}(t)$ is the aggregated climatic state of the parent area or zone,

$$D_{ctx}(r,t) = \left|\left|c_r(t) - \hat{c_r}(t)\right|\right|_2, \tag{4}$$

where $c_r(t)$ is a vector of contextual drivers and $\hat{c}_r$ is its expected baseline,

$$A(r,t) = \alpha \cdot D_{loc}(r,t) + \beta \cdot D_{hor}(r,t) + \gamma \cdot D_{ver}(r,t) + \delta \cdot D_{ctx}(r,t), \tag{5}$$

where $\alpha + \beta + \gamma + \delta = 1$, and $\alpha, \beta, \gamma, \delta \geq 0$. A point is marked as a local anomaly candidate if $A(r,t) > \tau_r$; a warning is generated only after persistence and cross-region confirmation rules are satisfied within a sliding time window H.

The score is an operational monitoring indicator for warning-oriented decision logic, not a replacement for physical meteorological forecasting models.

Proposed HCOIM-EWS anomaly identification method

To operationalize the formal anomaly model, the author-defined method Hierarchical Context-Orchestrated Intelligent Monitoring for Early Warning Systems (HCOIM-EWS) is proposed. In this paper, HCOIM-EWS is used as the author's designation for a hierarchical context-orchestrated intelligent monitoring method for early warning systems, not as a standardized, generally accepted, or official IT-industry term. It integrates machine learning-based anomaly identification, hierarchical contextual interpretation, and warning decision rules. It considers:

- multi-source ingestion and synchronization of environmental data streams through APIs, Kafka topics, and storage services;
- construction of region-level, area-level, and zone-level feature spaces with explicit horizontal and vertical linking;
- construction and periodic retraining of region-specific anomaly detectors using synchronized multivariate feature vectors, baseline estimation, unsupervised anomaly scoring, and warning-oriented threshold calibration;
- online inference for each region with computation D_{loc}, D_{hor}, D_{ver}, and D_{ctx};
- aggregation of local anomaly events in the CEP layer to identify warning templates based on persistence, neighborhood support, and hierarchical escalation;
- dissemination of warning signals to the alert subsystem, dashboards, and external stakeholders.

Thus, HCOIM-EWS transforms ML-based anomaly candidates into operational warning events through hierarchical context and decision logic.

Machine learning model construction

The machine learning component was constructed as a region-specific anomaly identification pipeline. The objective was to adapt anomaly scoring to each region and interpret local results in spatial and hierarchical contexts.

For each monitored region r, the input data were represented as a synchronized multivariate environmental vector:

$$x_r(t) = \left[x_{r,1}(t),\ x_{r,2}(t),\ \ldots,\ x_{r,m}(t)\right], \tag{6}$$

where $x_r(t)$ describes synchronized environmental and contextual variables.

In the simulation, each region had 720 synchronized observations. These observations were used to estimate the baseline behavior of each region and to train or calibrate anomaly detectors.

The procedure included the stages summarized in **Table 1**. The ML component consists of a region-specific anomaly identification layer and a warning-oriented decision layer. The first layer detects candidate deviations, while the second layer determines whether these deviations are persistent, spatially supported, and sufficiently meaningful to generate an early warning. To clarify how the machine learning component was constructed, **Table 1** summarizes the model-building procedure used in the proposed method.

Table 1 shows that HCOIM-EWS extends point-level anomaly detection with horizontal, vertical, contextual, and CEP-based confirmation mechanisms.

Table 1. **Machine learning model-building procedure**

Stage	Input	Operation	Output
Data synchronization	Heterogeneous environmental observations	Time alignment, grouping by region, normalization	Synchronized regional feature vectors
Baseline estimation	Historical regional observations	Estimation of expected regional behavior	Region-specific normal behavior model
Baseline detector construction	Normalized feature vectors	Robust Z-score method, Isolation Forest algorithm, and One-Class SVM method	Comparative anomaly scores
Author-defined HCOIM-EWS score construction	Regional, neighboring, parent-level, and contextual features	Computation of D_{loc}, D_{hor}, D_{ver}, D_{ctx}	Aggregated anomaly score $A(r,t)$
Thresholding	Anomaly scores	Quantile-based threshold selection	Local anomaly candidates
Warning decision	Local anomaly candidates	CEP rules, persistence check, spatial and hierarchical confirmation	Warning-level anomaly events

Cloud-agnostic EWS workflow

At the application level, the proposed early warning system is organized as a pipeline that transforms heterogeneous environmental data into warning-oriented information products. The workflow begins with data acquisition and normalization, continues with regional forecasting and anomaly detection, and ends with complex event processing and signal dissemination to end users and response services.

Simulation design

Anomalies that appear weak or noisy at one spatial scale may become interpretable when neighboring regions and parent climatic structures are considered together. Vertical scaling moves from smaller territorial units toward larger climatic aggregations and back, allowing the monitoring system to assess whether local instability is consistent with broader climatic states.

Climate regimes motivate the hierarchical grouping of regions. The Köppen–Geiger classification supports this conceptual partitioning and justifies region-specific models.

Horizontal and vertical orchestration evaluate the target region along with neighboring regions for anomaly interpretation. This is important when the local signal needs to be amplified, attenuated, or contextualized using spatial coherence (**Fig. 1**).

Fig. 1 shows how the target region is compared with adjacent regions. If the target region deviates while neighboring regions remain stable, the system may treat the signal as a local anomaly candidate. If the deviation is also supported by neighboring regions, the anomaly evidence becomes stronger and may contribute to warning-level escalation.

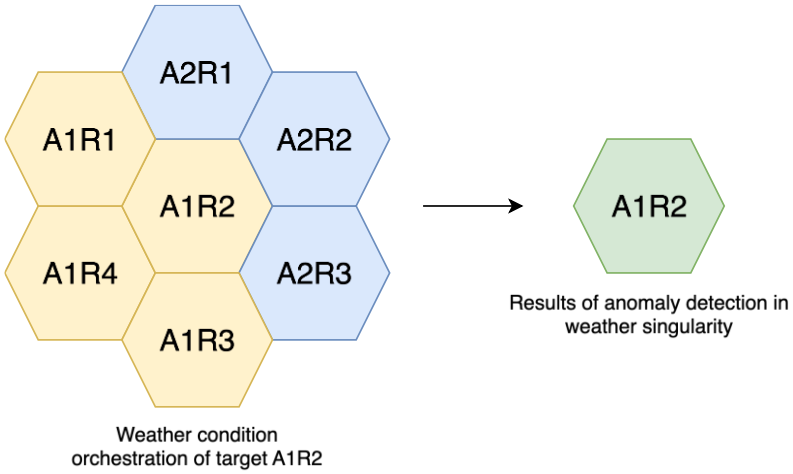

Fig. 1. Horizontal orchestration of a target region and neighboring regions for spatial coherence assessment. Source: developed by the authors

To test the feasibility of the anomaly model and the proposed monitoring method, a synthetic spatiotemporal climate graph was generated for 36 regions organized into 3 zones × 3 areas × 4 regions. Each regional time series contained 720 synchronized observations and combined seasonal structure, local trend, global climatic oscillation, zone-specific variation, area-specific variation, and random noise.

In this study, the synthetic spatiotemporal climate graph is understood as an artificially generated hierarchical model of climate zones, areas, and regions. Each region is represented as a graph node with a synchronized environmental time series, while links between nodes describe spatial neighborhood relations and membership in higher climatic levels. This representation allows controlled generation of normal states, injected anomalies, and ground-truth anomaly episodes for comparative evaluation of anomaly-identification methods.

Four anomaly types were injected into the synthetic environment:

- local spikes;
- regional drifts affecting all regions inside an area;
- hierarchy breaks where one area diverged from its zone-level regime;
- latent-driver anomalies affecting multiple semantically related regions.

In total, 33 anomaly episodes were injected and used as ground truth.

Four approaches were compared: the Robust Z-score statistical method, the Isolation Forest anomaly-detection algorithm, the One-Class SVM method, and the proposed author-defined HCOIM-EWS intelligent monitoring method.

All diagrams, workflows, algorithms, and architecture figures presented in this paper were redrawn and unified by the authors. Figures based on external concepts are explicitly marked as “developed by the authors based on [source]”, while fully original figures are marked as “developed by the authors”.

RESULTS AND DISCUSSION

Anomaly model simulation

The simulation inspected whether the model distinguishes isolated local deviations from spatially and hierarchically meaningful warning patterns. The synthetic climate graph contained 36 monitored regions arranged into 3 zones, 9 areas, and 36 regions. For each region, a synchronized multivariate time series of 720 observations was generated. Each observation included seasonal variation, local noise, regional trend, zone-specific climatic

oscillation, area-level variation, and contextual drivers. **Table 2** summarizes the main configuration parameters of the synthetic simulation environment.

Table 2. **Configuration of the synthetic simulation environment**

Parameter	Value used in simulation	Purpose
Number of climatic zones	3	Representation of broad climatic regimes
Number of climatic areas	9	Intermediate aggregation level
Number of regions	36	Region-specific anomaly detection
Regions per area	4	Local spatial neighborhood construction
Observations per region	720	Synchronized time-series simulation
Total regional observations	25,920	Complete synthetic monitoring dataset
Number of input dimensions	6	Multivariate environmental and contextual representation
Injected anomaly episodes	33	Ground truth for event-level validation
Compared methods	4	Robust Z-score method, Isolation Forest algorithm, One-Class SVM method, HCOIM-EWS method
Evaluation unit	Anomaly episode	Warning-level rather than point-level validation

The generated observations were used to compute four partial anomaly components: local deviation, horizontal deviation, vertical deviation, and contextual stress. The local deviation component reacted to abrupt changes in a single regional signal. The horizontal component evaluated whether the target region remained coherent with neighboring regions. The vertical component measured whether the region was consistent with its parent climatic area or zone. The contextual component represented abnormal pressure caused by external or latent drivers. The final anomaly score combined these components using weighted aggregation and a region-specific threshold estimated from baseline observations.

The injected anomalies were designed to activate different components of the proposed anomaly model. **Table 3** shows the relationship between anomaly type, spatial scope, expected dominant anomaly component, and warning-level interpretation.

The design supports component-level inspection because different anomaly classes should activate different evidence patterns.

The simulation confirmed that the four-component anomaly model produces more interpretable warning evidence than a purely local detector. Local spikes were mainly reflected by the local deviation component. Regional drifts increased both local and horizontal components. Hierarchy breaks were most clearly reflected by vertical incoherence. Latent-driver anomalies increased contextual stress and produced spatially distributed but semantically coherent deviations. Therefore, the model can explain not only that an anomaly occurred, but also which type of spatial, hierarchical, or contextual inconsistency contributed to the warning signal.

Table 3. **Injected anomaly types and expected response of the proposed anomaly model**

Anomaly type	Spatial scope	Expected dominant component	Warning-level interpretation
Local spike	Single region	Local deviation, D_{loc}	Short, abrupt regional disturbance
Regional drift	All regions inside one area	Local deviation and horizontal coherence, D_{loc} + D_{hor}	Gradual area-level instability
Hierarchy break	One area diverges from its zone	Vertical coherence, D_{ver}	Inconsistency between the area and the parent climatic regime
Latent-driver anomaly	Several semantically related regions	Contextual stress and horizontal coherence, D_{ctx} + D_{hor}	Distributed anomaly caused by hidden contextual pressure

The most important observation is that not every detected point anomaly should be escalated into an early warning. In the proposed model, an event is escalated only when the anomaly score remains persistent within a sliding window and is supported by neighboring regions, parent climatic structures, or contextual drivers. This makes the model suitable for early warning systems, where excessive sensitivity to isolated outliers may produce many false alarms and reduce trust in the warning service.

The model was inspected at the component level to check whether each component contributes distinct anomaly evidence.

The local deviation component was sensitive to abrupt single-region changes and therefore provided fast initial detection. However, this component alone was not sufficient for warning-level decisions because local noise, sensor artifacts, or short-term weather fluctuations may also produce high local anomaly scores. The horizontal coherence component reduced this limitation by comparing the target region with neighboring regions. If the deviation was not supported by spatial context, the probability of escalation decreased. This behavior is important for suppressing spatially inconsistent noise. An ablation-style inspection showed the expected behavior of the model (**Table 4**).

Table 4. **Component-level inspection of the proposed anomaly model**

Model variant	Included components	Expected strength	Expected limitation
Local-only model	D_{loc}	Fast reaction to abrupt regional deviations	High sensitivity to noise and isolated outliers
Local + horizontal model	D_{loc}+ D_{hor}	Suppression of spatially inconsistent false alarms	Weaker interpretation of hierarchy-level deviations
Local + horizontal + vertical model	D_{loc}+ D_{hor} + D_{ver}	Detection of inconsistencies between regions, areas, and zones	Limited explanation of hidden contextual drivers
Full HCOIM-EWS model	D_{loc} + D_{hor}+ D_{ver}+ D_{ctx}+ CEP rules	Balanced warning-level detection with reduced false alarms	Requires calibration of weights, thresholds, and CEP rules

The vertical coherence component evaluated whether regional behavior was consistent with the parent's climatic area or zone. This component was especially useful for hierarchy-break anomalies, where a local or area-level pattern diverged from the broader climatic regime. The contextual stress component captured abnormal influence from external or latent drivers. It was useful in cases where several regions changed in a coordinated way, but the change was not fully explained by local or neighboring time-series behavior.

The inspection confirms that each additional component contributes to a distinct type of monitoring evidence. The local component provides initial sensitivity, the horizontal component checks spatial support, the vertical component evaluates hierarchical consistency, and the contextual component captures external or latent pressure. CEP rules then transform the aggregated anomaly score into warning-level evidence.

Thus, the weighted anomaly score is a warning-oriented evidence function rather than a universal meteorological risk index.

Anomaly detection method implementation

The HCOIM-EWS method was implemented as a sequence of data-processing, model-training, inference, and warning-aggregation operations aligned with the proposed cloud-agnostic architecture (**Fig. 2**). The implementation assumes that each monitored region has its own synchronized data stream, historical baseline, neighboring regions, and parent climatic area or zone.

At the ingestion stage, heterogeneous environmental inputs are collected through distributed APIs and transformed into unified time-indexed observations. Missing or delayed values are handled through time alignment, caching of the last valid values, and normalization by input type. The resulting observations are written to the storage layer and also published as events for online inference.

At the feature-construction stage, the system builds three connected feature spaces: region-level features, neighborhood-level features, and parent-level climatic aggregates. Region-level features are used to estimate local deviation. Neighborhood-level features are used to estimate horizontal coherence. Parent-level aggregates are used to estimate vertical coherence. Contextual drivers are represented as an additional feature vector and compared with their expected baseline.

Fig. 2 shows that HCOIM-EWS separates local anomaly candidate generation from the final warning generation. This separation prevents local anomaly candidates from being treated as warnings without persistence and spatial support.

At the training stage, region-specific models are trained or calibrated on synchronized baseline observations. The Robust Z-score statistical method estimates robust statistical deviation, the Isolation Forest anomaly-detection algorithm learns isolation-based abnormality, the One-Class Support Vector Machine method estimates the boundary of normal behavior, and the proposed HCOIM-EWS intelligent monitoring method constructs a multi-component anomaly score from local, horizontal, vertical, and contextual evidence. In a production environment, this stage can be implemented through AutoML services that periodically retrain region-specific models and update thresholds.

At the online inference stage, each new observation is processed as follows:

- receive synchronized observation for region r at time t;
- compute local deviation D_{loc} using the regional baseline;
- compute horizontal deviation D_{hor} using neighboring regions $N(r)$;
- compute vertical deviation D_{ver} using the parent climatic aggregate $P(r)$;
- compute contextual stress D_{ctx} using contextual drivers;
- calculate the aggregated anomaly score $A(r,t)$;
- compare $A(r,t)$ with the region-specific threshold τ_r;
- publish a local anomaly event if the threshold is exceeded;

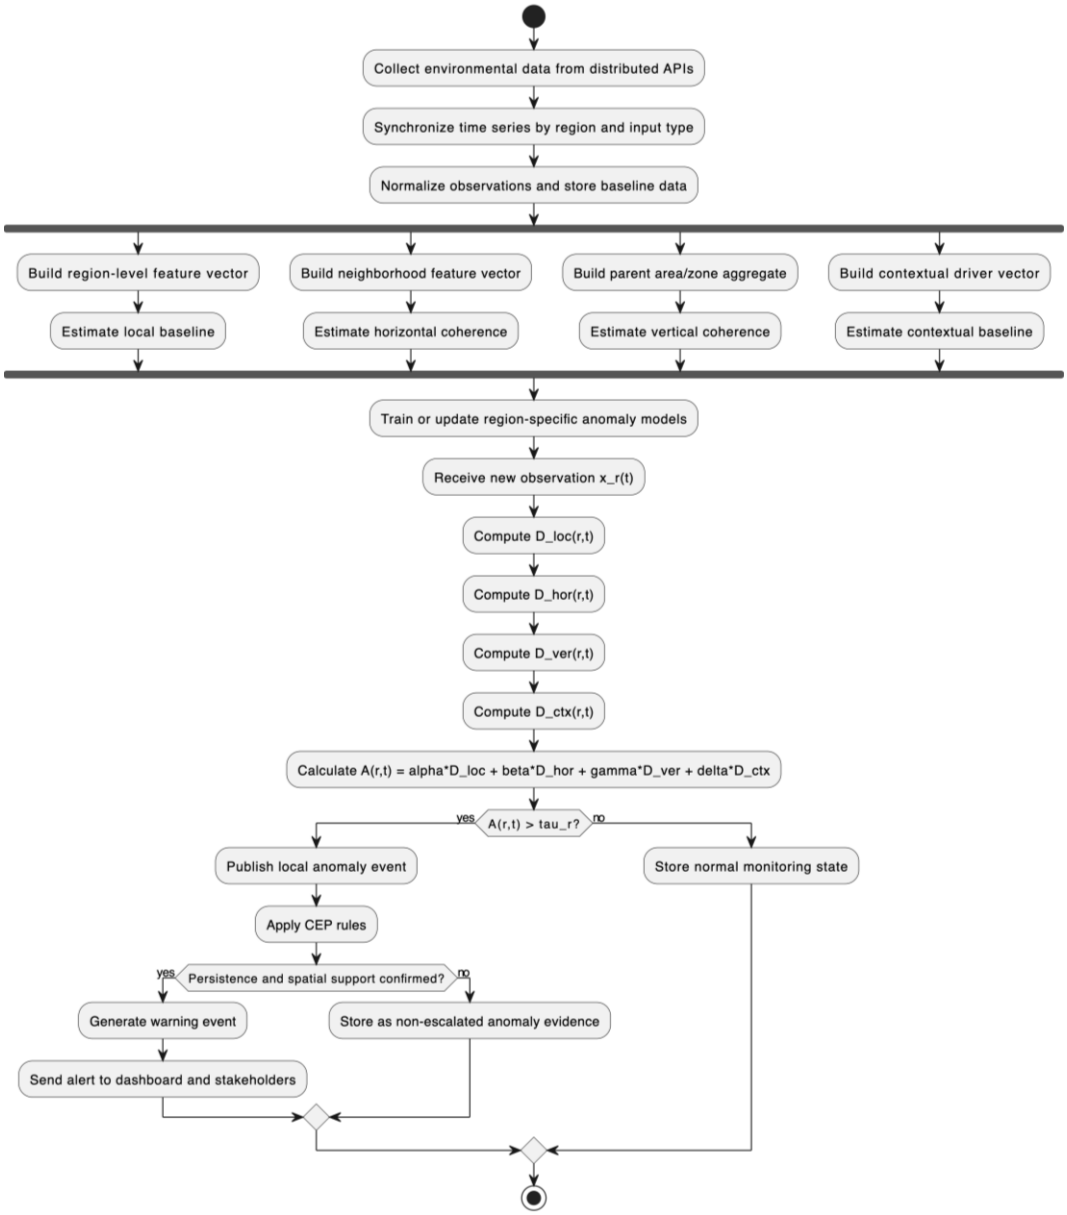

Fig. 2. Implementation flow of the HCOIM-EWS method from data ingestion to warning generation. Source: developed by the authors

- apply CEP rules for persistence, neighborhood support, and hierarchical escalation;
- generate a warning event only when the escalation conditions are satisfied.

The implementation separates ML-based anomaly candidate generation from warning-level decision-making. This separation is essential for early warning systems because a point anomaly is only a signal candidate, whereas a warning must be persistent, spatially supported, and operationally meaningful. As a result, HCOIM-EWS acts not only as an anomaly detector but also as an intelligent monitoring method that transforms local deviations into structured warning evidence.

Validation protocol

The validation protocol was designed to evaluate the proposed method from the viewpoint of warning management rather than only point-level anomaly classification. This distinction is important because an early warning system should not simply maximize the number of detected abnormal points. It should detect meaningful anomaly episodes while minimizing false alarms and maintaining acceptable detection delays. **Fig. 3** summarizes the validation procedure: synthetic graph generation, normal data simulation, anomaly injection, method comparison, and event-level evaluation.

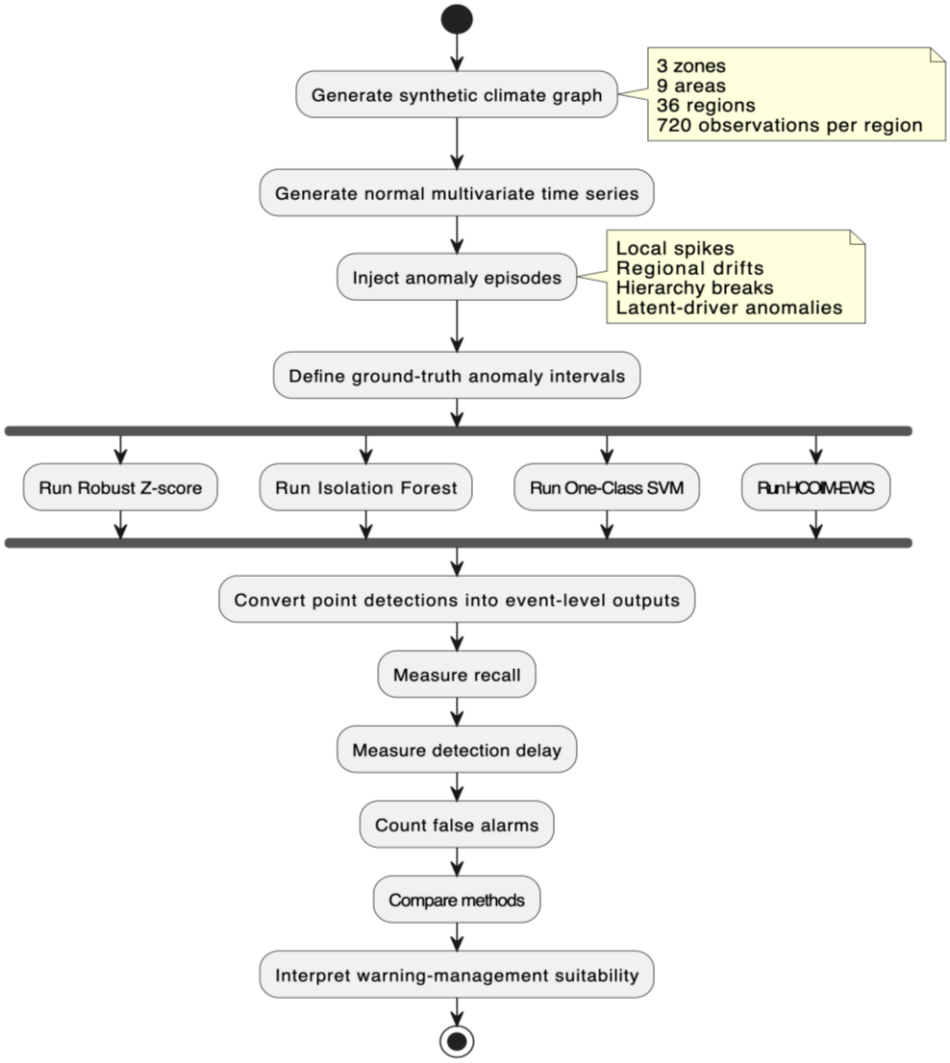

Fig. 3. Validation protocol for event-level evaluation of the HCOIM-EWS method. Source: developed by the authors

The main evaluation unit was an anomaly episode rather than an isolated anomalous point. An injected episode was counted as detected if at least one warning was generated within the allowed time window. Detection delay was calculated as the difference between the start of the injected episode and the first corresponding detection. False alarms were counted when the method generated warning-level events outside the ground-truth anomaly intervals.

HCOIM-EWS achieved the lowest false-alarm count while preserving a recall level comparable to the best baseline methods. This confirms that the proposed method is suitable for warning-oriented monitoring, where the operational cost of excessive false positives is high.

Simulation-based comparison of four methods

The compared methods were evaluated from the perspective of warning-level anomaly detection. Therefore, the evaluation focused not only on whether anomalous points were detected, but also on whether complete anomaly episodes were identified with acceptable delay and a limited number of false alarms. **Table 5** summarizes the threshold quantile used for each method, event-level recall, average detection delay, and the number of false alarms generated during the simulation.

Table 5. **Simulation results for the compared methods**

Compared approach	Threshold quantile	Event-level recall	Detection delay	False alarms
Robust Z-score	0.960	0.641	2.88	160
Isolation Forest	0.975	0.641	6.20	93
One-Class SVM	0.985	0.667	2.65	143
HCOIM-EWS	0.975	0.641	3.16	72

HCOIM-EWS achieved the lowest false-alarm count while preserving recall comparable to Robust Z-score and Isolation Forest. One-Class SVM had slightly higher recall but much more false alarms. Thus, the proposed method provides a more balanced warning-management profile, prioritizing stable and context-supported warning generation rather than sensitivity to isolated anomalous points. This behavior is especially important for early warning systems, where excessive false positives may reduce user trust and overload downstream response services.

The compared ML models were constructed on the same synchronized feature space, but they differed in the way anomaly evidence was interpreted. The Robust Z-score statistical method evaluated deviations from a robust statistical baseline. The Isolation Forest anomaly-detection algorithm identified observations that were easier to isolate in the feature space. The One-Class SVM method estimated the boundary of normal behavior. The proposed author-defined HCOIM-EWS intelligent monitoring method extended this point-level logic by adding spatial, hierarchical, and contextual interpretation.

To make the comparison more explicit, **Table 6** presents the relative false-alarm reduction achieved by HCOIM-EWS in comparison with the baseline approaches.

Table 6. **Relative improvement of HCOIM-EWS by false-alarm reduction**

Baseline approach	False alarms of baseline	False alarms of HCOIM-EWS	Absolute reduction	Relative reduction
Robust Z-score	160	72	88	55.0%
Isolation Forest	93	72	21	22.6%
One-Class SVM	143	72	71	49.7%

Thus, HCOIM-EWS mainly improves operational reliability by reducing false warnings. **Table 7** summarizes the operational interpretation of these results.

The method is therefore best interpreted as a balanced warning-management method, not as the most sensitive point detector.

Cloud-agnostic intelligent monitoring reference architecture for EWS

The proposed cloud-agnostic architecture treats anomaly identification and warning generation in EWSs as a sequence of distributed events and analytics operations rather than as a monolithic forecasting task. The monitored climate space is partitioned into zones, areas, and regions. Vertical and horizontal orchestration of those partitions allows the monitoring system to interpret local instability in a broader climatic context and to scale to hundreds of region-specific models.

Table 7. **Operational interpretation of model-evaluation results**

Evaluation aspect	Best-performing method	Interpretation
Lowest false-alarm count	HCOIM-EWS	Best warning-management reliability
Highest event-level recall	One-Class SVM	Highest sensitivity, but with high false-alarm burden
Shortest detection delay	One-Class SVM	Fastest detection in the simulation
Lowest delay among methods with low false alarms	HCOIM-EWS	Balanced compromise between delay and false-alarm reduction
Best operational trade-off	HCOIM-EWS	Most suitable for warning-oriented anomaly identification

The reference architecture consists of five logical layers: data acquisition, event and storage infrastructure, ML training and inference, warning decision support, and alert dissemination. In this interpretation, the architecture is not only a reference software design but also a practical implementation environment for intelligent monitoring: Kafka provides asynchronous event exchange, Spark-based processing supports scalable feature preparation, AutoML manages model training and retraining, and CEP converts anomaly evidence into warning-level decisions [10, 11].

Region-specific anomaly detectors publish local anomaly events that are aggregated by the context orchestration and CEP layers. Confirmed warning events are disseminated to dashboards, notification services, emergency services, and user-facing applications, reflecting the role of communication and response capacity in effective EWSs [19, 20].

Because environmental inputs differ in update periods, quality, and temporal resolution, the architecture performs time alignment, value reconciliation, caching, grouping, and normalization before training and inference.

The training contour includes historical data storage, feature preparation, model training, model registry, threshold calibration, and deployment to online inference services. Cloud-agnostic deployment is achieved by using provider-independent components such as containerized services, Kubernetes orchestration, Spark-compatible processing, object or distributed storage, and event streaming.

Fig. 4 presents the proposed reference architecture as an implementation template for HCOIM-EWS. The architecture is cloud-agnostic, meaning that its components can be deployed in public, private, or hybrid cloud environments without dependence on the proprietary services of a particular cloud provider.

Fig. 4 shows how regional anomaly identification, hierarchical context orchestration, CEP-based warning confirmation, and alert dissemination are connected in one cloud-agnostic monitoring architecture.

The study has several limitations. The validation was performed on synthetic data, which allows controlled inspection of anomaly types but does not fully reproduce the complexity of real climate and sensor data. The thresholds and weights of the anomaly score were selected for simulation purposes and require calibration on historical observations. In addition, the proposed architecture was evaluated conceptually and through simulation rather than through full industrial deployment. These limitations define the next stage of validation, which should include calibration on real historical climate observations, testing on streaming sensor data, sensitivity analysis of weights and thresholds, and evaluation in a pilot deployment scenario.

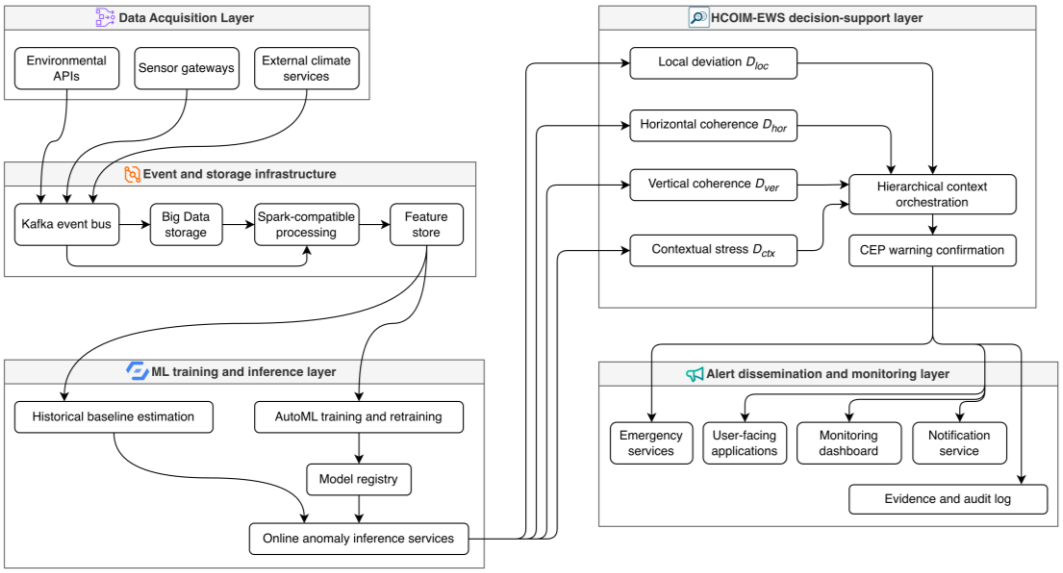

Fig. 4. Cloud-agnostic HCOIM-EWS reference architecture connecting environmental data acquisition, event streaming, ML training and inference, hierarchical context orchestration, CEP-based warning confirmation, and alert dissemination. Source: developed by the authors

CONCLUSION

This paper proposed a cloud-agnostic intelligent monitoring architecture for climate anomalies in early warning systems, supported by an integrated ML-based anomaly-identification and decision-support method. The architecture combines a formal hierarchical anomaly model, region-specific anomaly detector training, warning-oriented decision logic, the author-defined HCOIM-EWS method, and a reusable cloud-agnostic reference architecture.

The proposed anomaly model integrates local deviation, horizontal coherence, vertical coherence, and contextual stress. These components allow the system to distinguish isolated local deviations from spatially and hierarchically meaningful warning patterns. The HCOIM-EWS method operationalizes this model by transforming local anomaly candidates into warning-level evidence using persistence, neighborhood support, hierarchical consistency, and CEP-based escalation rules.

Simulation on a synthetic climate graph with 36 regions and 33 injected anomaly episodes showed that the author-defined HCOIM-EWS intelligent monitoring method achieved the lowest false-alarm count at a comparable event-recall level relative to the

Robust Z-score statistical method, the Isolation Forest anomaly-detection algorithm, and the One-Class SVM method. The proposed method reduced false alarms by 55.0% compared with the Robust Z-score statistical method, 22.6% compared with the Isolation Forest anomaly-detection algorithm, and 49.7% compared with the One-Class SVM method, while preserving comparable event-level recall.

The obtained results confirm that the proposed method is suitable for warning-oriented anomaly identification, where false-alarm reduction, contextual interpretation, and decision-support reliability are as important as point-level anomaly sensitivity. Future work should include validation on real historical and streaming climate data, calibration of thresholds and CEP rules, integration of data- and model-drift management, and extension of XAI and MLOps mechanisms for industrial-scale early warning systems.

ACKNOWLEDGMENTS AND FUNDING SOURCES

The authors received no financial support for the research, writing, and publication of this article.

COMPLIANCE WITH ETHICAL STANDARDS

The authors declare that the research was conducted in the absence of any conflict of interest.

AUTHOR CONTRIBUTIONS

Conceptualization, [V.L., Y.M., P.K., A.P., T.K.]; methodology, [V.L., Y.M., T.K.]; validation, [V.L., Y.M., P.K.]; formal analysis, [V.L., Y.M., A.P.]; investigation, [Y.M., P.K., A.P.]; resources, [P.K., T.K.]; data curation, [A.P.]; writing – original draft preparation, [V.L., Y.M., T.K.]; writing – review and editing, [V.L., Y.M., T.K.]; visualization, [T.K., A.P.]; supervision, [V.L.]; project administration, [V.L.].

All authors have read and agreed to the published version of the manuscript.

REFERENCES

[1] Rokhideh, M., Fearnley, C., & Budimir, M. (2025). Multi-hazard early warning systems in the Sendai Framework for disaster risk reduction: Achievements, gaps, and future directions. *International Journal of Disaster Risk Science, 16*, 103–116. https://doi.org/10.1007/s13753-025-00622-9

[2] Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys, 41*(3), Article 15. https://doi.org/10.1145/1541880.1541882

[3] Sgueglia, A., Di Sorbo, A., Visaggio, C. A., & Canfora, G. (2022). A systematic literature review of IoT time series anomaly detection solutions. *Future Generation Computer Systems, 134*, 170–186. https://doi.org/10.1016/j.future.2022.04.005

[4] Camps-Valls, G., Fernández-Torres, M.-Á., Cohrs, K.-H., Höhl, A., Castelletti, A., Pacal, A., Robin, C., Martinuzzi, F., Papoutsis, I., Prapas, I., Pérez-Aracil, J., Weigel, K., Gonzalez-Calabuig, M., Reichstein, M., Rabel, M., Giuliani, M., Mahecha, M. D., Popescu, O.-I., Pellicer-Valero, O. J., ... Williams, T. (2025). Artificial intelligence for modeling and understanding extreme weather and climate events. *Nature Communications, 16*, Article 1919. https://doi.org/10.1038/s41467-025-56573-8

[5] Spohn, T. K., Walsh, E., Horan, K., O'Donoghue, J., Charnecki, T., Haslam, M., & Gallagher, S. (2026). A machine learning approach using autoencoders to perform quality control on meteorological data. *Environmental Data Science, 5*, Article e1. https://doi.org/10.1017/eds.2026.10030

[6] Sun, H., Liu, Y., Zhang, Y., et al. (2025). A machine learning-based quality control algorithm for heavy rainfall observations. *Remote Sensing, 17*(24), Article 3976. https://doi.org/10.3390/rs17243976
[7] Islam, M. M., Hasan, M., Mia, M. S., Masud, A. A., & Islam, A. R. M. T. (2025). Early warning systems in climate risk management: Roles and implementations in eradicating barriers and overcoming challenges. *Natural Hazards Research, 5*(3), 523–538. https://doi.org/10.1016/j.nhres.2025.01.007
[8] Chovanec, D., Kollár, B., Halúsková, B., Kubás, J., Pawęska, M., & Ristvej, J. (2025). A component-based approach to early warning systems: A theoretical model. *Applied Sciences, 15*(6), Article 3218. https://doi.org/10.3390/app15063218
[9] Girotto, C. D., Piadeh, F., Bkhtiari, V., Behzadian, K., Chen, A. S., Campos, L. C., & Zolgharni, M. (2024). A critical review of digital technology innovations for early warning of water-related disease outbreaks associated with climatic hazards. *International Journal of Disaster Risk Reduction, 100*, Article 104151. https://doi.org/10.1016/j.ijdrr.2023.104151
[10] Zaharia, M., Xin, R. S., Wendell, P., Das, T., Armbrust, M., Dave, A., Meng, X., Rosen, J., Venkataraman, S., Franklin, M. J., Ghodsi, A., Gonzalez, J., Shenker, S., & Stoica, I. (2016). Apache Spark: A unified engine for big data processing. *Communications of the ACM, 59*(11), 56–65. https://doi.org/10.1145/2934664
[11] Hutter, F., Kotthoff, L., & Vanschoren, J. (Eds.). (2019). *Automated machine learning: Methods, systems, challenges*. Springer. https://doi.org/10.1007/978-3-030-05318-5
[12] Beck, H. E., McVicar, T. R., Vergopolan, N., Berg, A., & Wood, E. F. (2023). High-resolution (1 km) Köppen–Geiger maps for 1901–2099 based on constrained CMIP6 projections. *Scientific Data, 10*, Article 724. https://doi.org/10.1038/s41597-023-02549-6
[13] García, J. A., Acero, F. J., Martínez-Pizarro, M., & Lara, M. (2024). A Bayesian hierarchical spatio-temporal model for summer extreme temperatures in Spain. *Stochastic Environmental Research and Risk Assessment, 38*(9), 3393–3410. https://doi.org/10.1007/s00477-024-02754-8
[14] Dereszynski, E. W., & Dietterich, T. G. (2011). Spatiotemporal models for data-anomaly detection in dynamic environmental monitoring campaigns. *ACM Transactions on Sensor Networks, 8*(1), Article 3. https://doi.org/10.1145/1993042.1993045
[15] Bâra, A., Văduva, A. G., & Oprea, S.-V. (2024). Anomaly detection in weather phenomena: News and numerical data-driven insights into climate change in Romania's historical regions. *International Journal of Computational Intelligence Systems, 17*, Article 134. https://doi.org/10.1007/s44196-024-00536-2
[16] Lyashkevych, V. Y. (2026). Intelligent monitoring as an information technology for context-aware decision-making strategy selection. *Ukrainian Journal of Information Technology, 8*(1), 39–50. https://doi.org/10.23939/ujit2026.01.039
[17] Lyashkevych, V. Y. (2025). Sustainable optimization of consolidated data processing algorithms based on machine learning and genetic algorithms. *Electronics and Information Technologies, 32*, 39–54. https://doi.org/10.30970/eli.32.3
[18] Gura, V., Ostrovska, O., & Lyashkevych, V. (2025). Scenario modelling and impact estimation of a local pollutant on the environment. *Applied Computer Systems, 30*(1), 195–201. https://doi.org/10.2478/acss-2025-0021

[19] Coughlan de Perez, E., Berse, K. B., Depante, L. A. C., Easton-Calabria, E., Evidente, E. P. R., Ezike, T., Heinrich, D., Jack, C., Lagmay, A. M. F. A., Lendelvo, S., Marunye, J., Maxwell, D. G., Murshed, S. B., Orach, C. G., Pinto, M., Poole, L. B., Rathod, K., Shampa, & Van Sant, C. (2022). Learning from the past in moving to the future: Invest in communication and response to weather early warnings to reduce death and damage. *Climate Risk Management, 38*, Article 100461. https://doi.org/10.1016/j.crm.2022.100461

[20] Sadiq, A.-A., Dougherty, R. B., Tyler, J., & Entress, R. (2023). Public alert and warning system literature review in the USA: Identifying research gaps and lessons for practice. Natural Hazards, 117(2), 1711–1744. https://doi.org/10.1007/s11069-023-05926-x

ХМАРНО-НЕЗАЛЕЖНА АРХІТЕКТУРА ІНТЕЛЕКТУАЛЬНОГО МОНІТОРИНГУ КЛІМАТИЧНИХ АНОМАЛІЙ У СИСТЕМАХ РАННЬОГО ПОПЕРЕДЖЕННЯ

Василь Ляшкевич*, Юрій Марчук, Петро Кулик, Андрій Патинко, Тарас Керод
Львівський національний університет імені Івана Франка, вул. Драгоманова, 50, Львів, 79005, Україна

АНОТАЦІЯ

Вступ. Кліматичні небезпеки вимагають систем раннього попередження, здатних поєднувати гетерогенні спостереження за станом довкілля, виявляти аномальні просторово-часові закономірності та передавати попереджувальні сигнали зацікавленим сторонам. Наявні рішення часто розділяють програмну архітектуру, моделі виявлення аномалій і логіку прийняття рішень, яка перетворює відхилення на ранні сигнали.

Матеріали та методи. У дослідженні розроблено хмарно-незалежну архітектуру інтелектуального моніторингу для систем раннього попередження, що підтримується ідентифікацією аномалій і прийняттям рішень на основі машинного навчання. Вона оцінює чотири типи ознак: відхилення від нормального стану, узгодженість із сусідніми регіонами, узгодженість із батьківською областю та вплив контекстних чинників. Синтетичний просторово-часовий кліматичний граф використано для порівняння статистичного методу Robust Z-score, алгоритму виявлення аномалій Isolation Forest, методу One-Class Support Vector Machine та запропонованого методу HCOIM-EWS, де HCOIM-EWS означає Hierarchical Context-Orchestrated Intelligent Monitoring for Early Warning Systems.

Результати. Запропонований метод інтегрує вибір детекторів на основі машинного навчання, регіонально-специфічне навчання, ієрархічну оцінку аномалій, логіку прийняття рішень щодо попереджень на основі системи опрацювання комплексних подій та хмарно-незалежну еталонну архітектуру моніторингу. У моделюванні HCOIM-EWS досяг найнижчої кількості хибних тривог за порівнюваного рівня повноти виявлення подій. Порівняно з Robust Z-score, Isolation Forest і One-Class SVM запропонований метод зменшив кількість хибних тривог на 55,0 %, 22,6 % і 49,7 % відповідно. Метод придатний для інтелектуального моніторингу, орієнтованого на попередження, де операційна надійність і зменшення хибних тривог є такими ж важливими, як і чутливість до точкових аномалій.

Висновки. Запропонований метод розширює виявлення аномалій у системах раннього попередження від аналізу ізольованих локальних часових рядів до ієрархічного, контекстно-залежного та подієво-орієнтованого інтелектуального моніторингу. Це підтверджує технічну доцільність архітектури та обґрунтовує її використання як багаторазового шаблону для платформ моніторингу кліматичних ризиків.

Ключові слова: виявлення аномалій, системи раннього попередження, інтелектуальний моніторинг, хмарно-незалежна архітектура, обробка складних подій, машинне навчання

Received / Одержано	Revised / Доопрацьовано	Accepted / Прийнято	Published / Опубліковано
02 June, 2026	24 June, 2026	25 June, 2026	29 June, 2026