Power System Protection and Control
ISSN:1674-3415

ADAPTIVE EVIDENCE-DRIVEN DATA INTELLIGENCE: A SELF-VALIDATING FRAMEWORK FOR RELIABLE AI DECISION-MAKING UNDER CONTINUOUS DATA DRIFT

Chaitanya Nagabhyru

270 Sherman avenue, Apartment B3, New Haven, Connecticut, USA, 06511

Received: 17/08/2026 **Revised**: 21/09/2026 **Accepted**: 05/10/2026

***ABSTRACT*:**

Machine learning systems are trained on a snapshot of the world but deployed into a world that keeps changing. When the data that a model sees in production departs from its training distribution, predictions can degrade silently, and decisions that rely on them become unreliable. This paper proposes an adaptive, evidence-driven and self-validating framework for maintaining reliable AI decision-making under continuous data drift. The framework combines a multimodal representation layer, a statistical evidence layer that quantifies drift and uncertainty, a validation gate that converts evidence into trust states, an adaptive response layer that recalibrates, retrains or defers to human judgement, and an audit ledger that records why each action was taken. Using three illustrative figures, we show how multimodal contrastive alignment feeds anomaly detection, compliance validation and risk scoring; how the statistical significance of a drift test depends jointly on shift magnitude and sample size; and how the same synthetic shift can improve performance at one hospital while degrading it at another, with large class-level variation hidden behind stable aggregate metrics. We derive design rules for monitoring, response policy and governance, and we discuss limitations and research directions.

***Keywords**: Data Drift, Self-Validating AI, Evidence-Driven Decision-Making, Model Monitoring, Multimodal Learning, Contrastive Learning, Calibration, AI Governance.*

INTRODUCTION

Machine learning models are usually evaluated on held-out data that resembles the data on which they were trained. In deployment, this assumption fails. Customer behaviour changes, sensors age, clinical protocols and imaging equipment differ between hospitals, regulations alter what is recorded, and adversaries adapt. The resulting departure of production data from the training distribution is known as dataset shift or data drift [1], [2]. Its consequences are serious because a drifting model continues to produce confident predictions even as their accuracy falls, and ground-truth labels often arrive late or not at all, so the degradation can remain invisible [3].
The problem is acute in high-stakes domains. Finlayson and colleagues have described how dataset shift affects clinical AI, where changes in patient populations, practice patterns and imaging hardware can reduce the reliability of models that performed well at validation [4]. Similar concerns arise in finance, audit and compliance, where models support decisions about fraud, risk and regulatory conformity. In these settings, an organisation needs more than a prediction. It needs evidence that the prediction can still be trusted and a record of what was done when it could not.

Existing practice tends to treat drift detection, retraining and governance as separate activities. Drift detectors raise alarms but do not say how much they should change a decision; retraining pipelines run on schedules unrelated to evidence; and governance documents are written once and rarely updated. The result is a gap between what a model is believed to do and what it is doing now.

We argue for an alternative: a self-validating framework in which every automated decision is accompanied by current evidence about the reliability of the data and the model, and in which the system adapts its own behaviour in response to that evidence. We call this adaptive, evidence-driven data intelligence. The key ideas are that evidence should be quantitative and comparable across modalities, that it should be converted into explicit trust states that govern what the system is allowed to do, and that responses and their justifications should be logged for audit.

DOI: 10.46121/pspc.54.4.03

The contributions of this paper are as follows. First, we formalise the problem of reliable decision-making under continuous drift and define self-validation. Second, we present a five-layer framework and map it to a multimodal audit-integration network. Third, we analyse how the strength of drift evidence depends on shift magnitude and sample size, and we draw design rules for monitoring windows. Fourth, we examine multi-site healthcare drift results to show why aggregate monitoring is insufficient and why site-level and class-level evidence is required. Fifth, we provide a response policy, governance guidance and an implementation checklist.

BACKGROUND AND RELATED WORK

A. Types of Dataset Shift

Dataset shift is conventionally decomposed according to which part of the joint distribution $P(X, Y)$ of inputs X and labels Y changes [1]. Table I summarises the main types. In practice, several occur at once and cannot be cleanly separated, especially when labels are unavailable.

TABLE I. TYPES OF DATASET SHIFT

Type	What changes	Example	Detection signal
Covariate shift	Input distribution P(X); P(Y|X) unchanged	New scanner model changes image statistics	Input or embedding drift tests
Prior (label) shift	Class prevalence P(Y)	Seasonal change in disease frequency	Shift in predicted class distribution
Concept drift	Relationship P(Y|X)	Fraud tactics evolve; new coding practice	Performance decline; delayed labels
Data-quality drift	Missingness, formats, schema	Upstream system alters a field	Quality checks and schema validation
Multimodal drift	Relations between modalities	Report templates change while images do not	Cross-modal alignment statistics

B. Detecting Drift

Detection methods fall into three groups. Statistical two-sample tests compare a reference window with a current window of features or model outputs. Examples include the Kolmogorov-Smirnov test for univariate features, kernel-based tests such as maximum mean discrepancy [5], and population stability indices commonly used in credit scoring. Sequential detectors such as the Drift Detection Method and ADWIN monitor streams of errors or values and adapt their window sizes [6], [7]. Learning-based detectors train a classifier to distinguish reference from current data, with its accuracy indicating the degree of shift. Rabanser and colleagues compared combinations of dimensionality reduction and statistical testing and concluded that no single detector dominates, and that detection on low-dimensional representations of model outputs is often effective [3]. Reviews of concept-drift learning provide broader surveys of detection and adaptation strategies [8], [9].

C. Uncertainty, Calibration and Reliability

A model that knows when it does not know is more useful than one that is merely accurate. Calibration measures whether predicted probabilities match observed frequencies, and the Brier score is a proper scoring rule that combines calibration and discrimination [10]. Modern neural networks are often poorly calibrated [11], and calibration degrades further under shift, as shown by Ovadia and colleagues in a large benchmark of predictive uncertainty [12]. Conformal prediction provides distribution-free prediction sets with coverage guarantees under exchangeability, although its guarantees weaken under drift unless it is adapted [13].

D. Representation Learning Across Modalities

Enterprise data is rarely a single table. Audit and compliance processes combine structured records, free text and images or scanned documents. Contrastive learning aligns representations of related items and separates unrelated ones. The supervised contrastive loss of Khosla and colleagues uses class labels to define positives and has been shown to improve representation quality and robustness [14]. These properties are attractive for drift monitoring because a well-structured latent space makes distances and densities more meaningful as drift indicators.

E. Operational and Regulatory Context

Hidden technical debt in machine learning systems, including monitoring gaps and unstable data dependencies, has been recognised as a cause of production failures [15], and rubrics for ML test readiness explicitly include

DOI: 10.46121/pspc.54.4.03

data and prediction monitoring [16]. Regulatory frameworks increasingly require continuous oversight: the NIST AI Risk Management Framework emphasises measurement and management of risk across the AI lifecycle [17], and the European Union Artificial Intelligence Act requires post-market monitoring and logging for high-risk systems [18]. A framework for reliable decision-making must therefore generate auditable evidence as a first-class output.

PROBLEM FORMULATION

Consider a deployed model f trained on data from a distribution $P_0(X, Y)$. At time t the model receives inputs from $P_t(X, Y)$, which may differ from P_0. A decision d_t is taken using the model's output, and its quality depends on the unobserved loss of f under P_t. Because labels are delayed or absent, the loss cannot be measured directly. The system therefore relies on proxy evidence E_t, which may include drift statistics on inputs and outputs, estimated uncertainty, data-quality indicators and, when available, delayed performance feedback.

We define a system as self-validating if it satisfies three conditions. First, it computes E_t continuously and attaches a summary of it to each decision or batch of decisions. Second, it maps E_t to a trust state that constrains its own autonomy, for example by allowing automatic action, requiring review or blocking action. Third, it records the evidence, the state and the action in a tamper-evident log so that decisions can be reconstructed later.

Four requirements follow. Timeliness: evidence must be fresh enough to act upon before harm accumulates. Calibrated sensitivity: the system should raise alarms when shift matters, not whenever any shift exists, since benign drift is common and alarm fatigue is costly. Granularity: evidence must be available at the level at which harm occurs, such as site, subgroup or class, and not only in aggregate. Accountability: every adaptation must be traceable to evidence and approved under policy.

THE ADAPTIVE EVIDENCE-DRIVEN FRAMEWORK

The framework has five layers, summarised in Table II. Data flow from the representation layer to the evidence layer, then to the validation gate, which determines the response layer's behaviour; the audit ledger records every transition.

TABLE II. LAYERS OF THE FRAMEWORK

Layer	Purpose	Typical outputs
1 Representation	Encode multimodal data into a shared latent space	Embeddings; anomaly, compliance and risk signals
2 Evidence	Quantify drift, uncertainty, quality and performance	Statistics with effect sizes and confidence intervals
3 Validation gate	Map evidence to trust states	Green, amber or red state per site and class
4 Adaptive response	Act in proportion to evidence	Recalibration, retraining, rollback, deferral
5 Audit ledger	Record evidence, decisions and approvals	Append-only, tamper-evident log

A. Representation Layer: Multimodal Alignment

Fig. 1 shows a multimodal network that we use as the representation and intelligence layer. Data are acquired and pre-processed, then passed through modality-specific encoders. A cross-modal alignment network maps the encoder outputs into a shared latent space, and training uses a supervised contrastive objective. The audit intelligence module sits downstream of the latent space and produces three outputs: anomaly detection, compliance validation and risk scoring.

DOI: 10.46121/pspc.54.4.03

Power System Protection and Control
ISSN:1674-3415

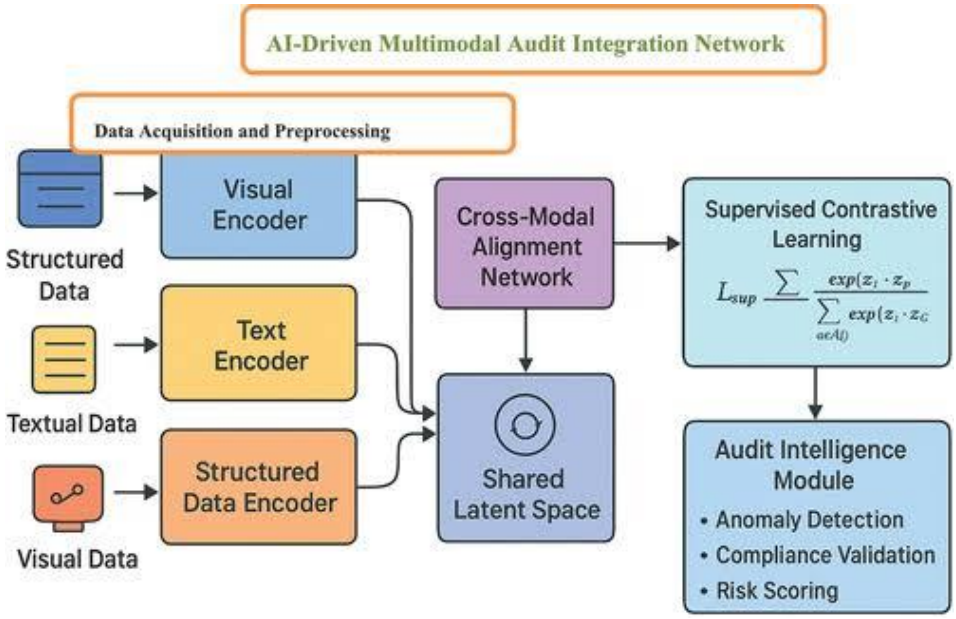

Fig. 1. AI-driven multimodal audit integration network: modality encoders, cross-modal alignment into a shared latent space, supervised contrastive learning and an audit intelligence module.

The supervised contrastive loss in the figure follows the form

$$L_{sup} = \Sigma_{i \in I} (-1 / |P(i)|) \Sigma_{p \in P(i)} \log [\exp(z_i \cdot z_p / tau) / \Sigma_{a \in A(i)} \exp(z_i \cdot z_a / tau)]$$

where z_i is the normalised embedding of sample i, $P(i)$ is the set of other samples sharing its label, $A(i)$ is the set of all other samples in the batch, and tau is a temperature parameter [14]. The loss pulls together embeddings of samples with the same label, regardless of modality, and pushes apart those with different labels. For our purposes the shared latent space has a second function. Because it places structured records, text and images in a common geometry, distributional distance and density measures computed there become comparable across modalities, and a single drift statistic can monitor evidence from heterogeneous sources. The figure's three modality encoders are labelled for visual, text and structured data respectively; in a deployment the encoder for each modality is chosen to suit the data, such as a vision transformer for scans, a language model for documents and a tabular network for ledgers.

The audit intelligence module links the representation layer to evidence. Anomaly detection flags samples that lie in low-density regions of the latent space or are poorly explained by the model. Compliance validation checks whether records satisfy rules and whether predictions are consistent with policy. Risk scoring combines these signals into a score that supports prioritisation. Each of these outputs is itself monitored, since a change in the rate of anomalies or in the distribution of risk scores is a drift signal.

B. Evidence Layer

The evidence layer computes four families of indicators. **Input drift** is measured by two-sample tests or distances between reference and current windows, on raw features, on latent embeddings and on model outputs. **Uncertainty and calibration** are measured by predictive entropy, ensemble disagreement, the Brier score and reliability diagrams. **Data quality** captures missingness, schema violations, range violations and duplicate rates. **Delayed performance** is computed on labelled samples as they become available, with confidence intervals that reflect the small sample. Each indicator is accompanied by its effect size and its sampling uncertainty, because a statistic without its uncertainty cannot support a decision.

C. Validation Gate

The validation gate converts evidence into one of three trust states. In the **green** state, evidence is consistent with the reference and automated action is permitted. In the **amber** state, evidence of material drift or elevated uncertainty exists, and the system continues but increases sampling for labelling, narrows the actions it may take automatically and flags decisions for review. In the **red** state, evidence indicates that reliability cannot be assured, and the system falls back to a safe default, such as human decision, a conservative rule or a previous validated model. Thresholds are set by the cost of errors in each domain and are reviewed periodically.

DOI: 10.46121/pspc.54.4.03

Power System Protection and Control
ISSN:1674-3415

EVIDENCE STRENGTH: SHIFT MAGNITUDE AND SAMPLE SIZE

A common practice is to run a statistical test on each monitoring window and raise an alarm when the p-value falls below a threshold such as 0.05. The reliability of this practice depends on the interaction between how large the shift is and how many samples are in the window. Fig. 2 illustrates this interaction. It plots drift-test p-values against shift magnitude, expressed as a percentage, for window sizes from 500 to 4,000 samples, with the significance threshold shown as a horizontal line and shaded bands indicating variability across repetitions.

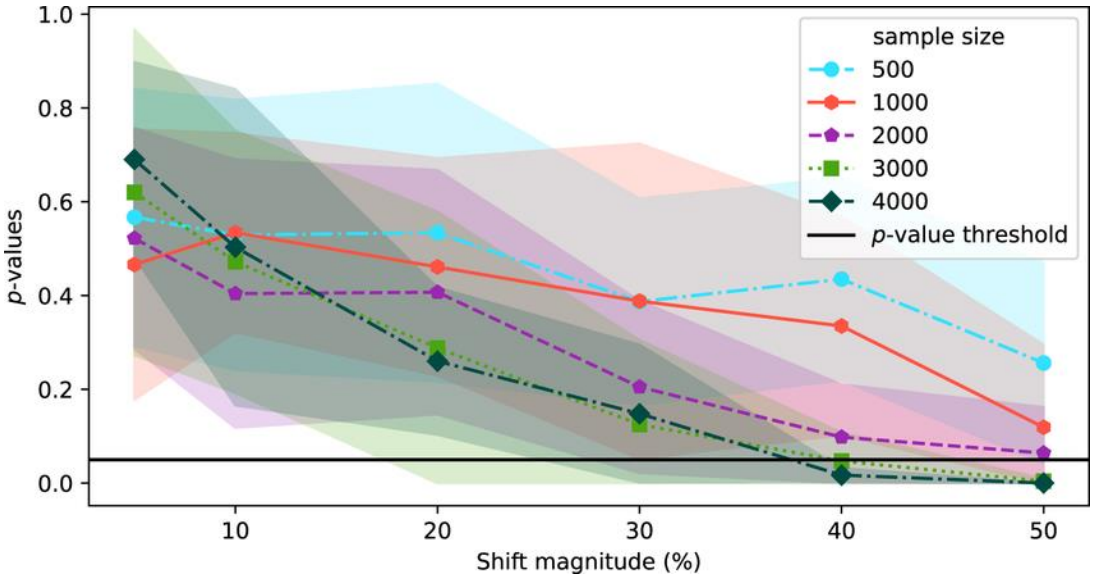

Fig. 2. Drift-test p-values versus shift magnitude for window sizes of 500 to 4,000 samples, with the significance threshold shown.

A. What the Figure Shows

- **Small shifts are not detected.** At a shift of about five percent, p-values for every window size lie well above the threshold, roughly between 0.45 and 0.7. A small shift is statistically indistinguishable from sampling noise at these sample sizes.
- **P-values fall as shift increases, and fall faster for larger windows.** At a shift of thirty percent the 3,000 and 4,000-sample windows have p-values near 0.13 to 0.15, whereas the 500 and 1,000-sample windows remain near 0.39. By forty percent the largest windows approach the threshold, and by fifty percent they are effectively zero.
- **Small windows may never detect a moderate shift.** The 500-sample window still has a p-value above 0.25 at the largest shift shown, and the 1,000-sample window, while declining, remains above the threshold. The curve for 500 samples is also non-monotonic, indicating that noise can dominate the trend.
- **Variability is large.** The shaded bands are wide and overlap across window sizes, especially at small shifts. A single window's p-value is therefore a noisy measurement, and individual alarms or non-alarms should not be treated as certain.

B. Design Implications

Three design rules follow. First, window size is a design parameter with a direct effect on detection power, and it should be chosen from the smallest shift that matters to the decision, not from convenience. Where traffic is low, windows must be longer, which increases detection delay; this timeliness trade-off should be stated explicitly. Second, p-values should be accompanied by effect sizes. A large window can produce a very small p-value for a trivial shift, causing false alarms, while a small window can fail to flag a consequential shift. Third, decisions should be based on persistence across windows and on aggregated evidence, for example through sequential tests or by requiring that a statistic remain beyond its threshold for several consecutive windows, rather than on a single test. When many features or subgroups are tested simultaneously, multiple-testing corrections or false-discovery-rate control are needed so that the number of alarms is not inflated by chance.

The population stability index is a widely used effect-size measure for binned features, computed as

DOI: 10.46121/pspc.54.4.03

Power System Protection and Control
ISSN:1674-3415

$$PSI = \Sigma_k \left(a_k - e_k \right) \times \ln \left(a_k / e_k \right)$$

where a_k and e_k are the actual and expected proportions in bin k. Industry conventions treat values below about 0.1 as stable and values above about 0.25 as a significant shift, though such thresholds are heuristics and should be calibrated to the application. Combining an effect-size indicator with a significance test provides a more robust evidence summary than either alone.

DRIFT ACROSS SITES: A HEALTHCARE ILLUSTRATION

Medical imaging provides a demanding test of drift monitoring because scanners, protocols and patient populations differ across institutions. Fig. 3 shows the change in model performance under increasing synthetic shift at two hospital sites. The left panels give the change relative to the unshifted baseline in AUROC, precision, recall, F1 and Brier score loss; the right panels give the change in F1 for each of ten findings, such as atelectasis, cardiomegaly, consolidation, edema, effusion, fracture, lung lesion, pneumonia and pneumothorax.

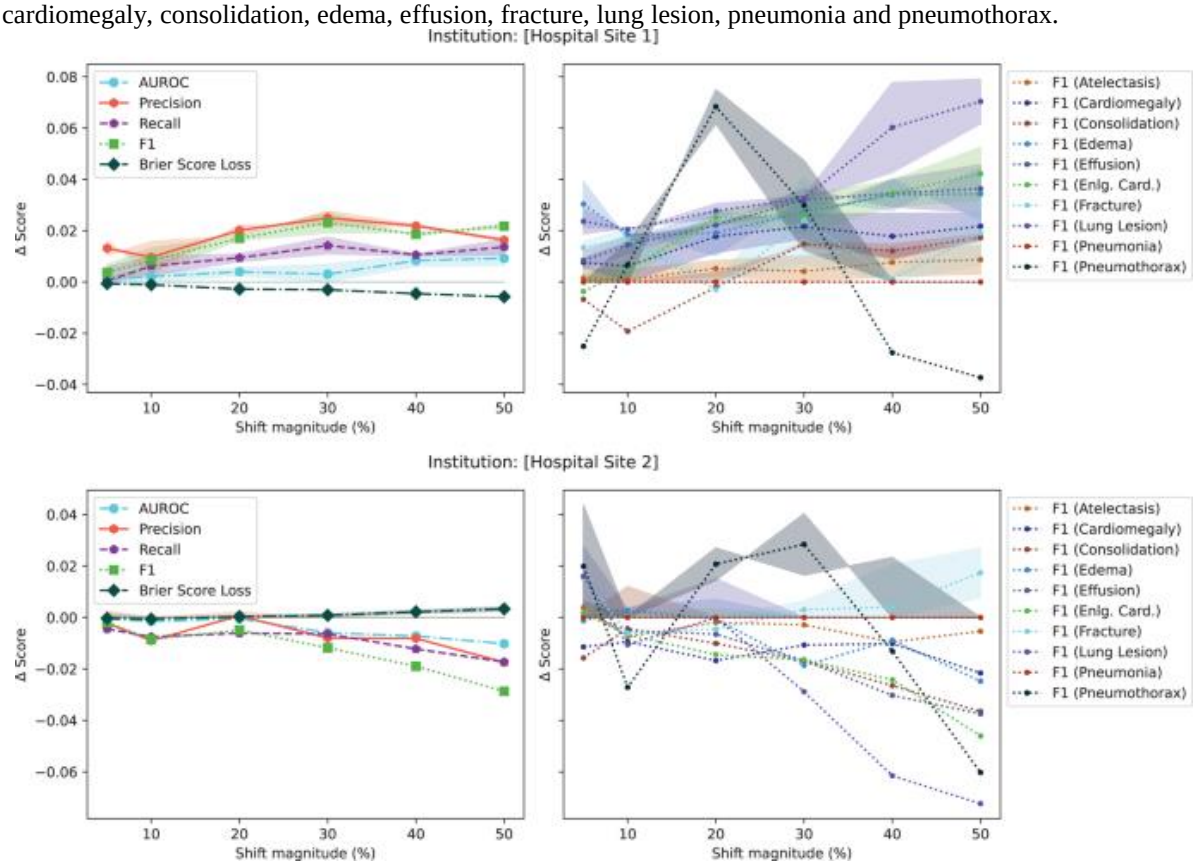

Fig. 3. Change in aggregate metrics (left) and per-class F1 (right) under increasing shift at two hospital sites.

A. Aggregate Metrics Can Move in Opposite Directions

At the first site, aggregate changes are small and mostly positive. Precision, recall and F1 rise by roughly 0.01 to 0.025 for moderate shifts, AUROC is close to unchanged, and the Brier score loss decreases slightly, which indicates marginally better probabilistic accuracy. At the second site, the same range of shift magnitudes produces the opposite pattern. Precision, recall, F1 and AUROC all decline, with F1 falling by about 0.03 at the largest shift, and the Brier score loss increases slightly. The same intervention therefore improves some measures at one institution and harms them at another. A framework that monitors a single pooled metric, or that assumes a shift has a consistent direction, would misjudge the effect on at least one site.

B. Aggregates Hide Class-Level Volatility

The class-level panels reveal considerably more variation than the aggregates. At the first site the per-class changes in F1 range from about minus 0.04 to plus 0.07. One finding, pneumothorax, rises sharply at a shift of around twenty percent, then falls below baseline at the largest shifts, while another class climbs steadily. At the second site most classes decline as shift increases, with some reaching about minus 0.07, whereas fracture rises slightly at high shift and pneumothorax swings between positive and strongly negative values. The shaded bands,

DOI: 10.46121/pspc.54.4.03

which indicate uncertainty, are widest for the rarer findings, so changes for these classes are both large and imprecisely estimated.

C. Lessons for Evidence Design

- **Monitor at the level of harm.** Evidence should be computed per site and per class or subgroup, with aggregates as summaries and not as substitutes.
- **Report uncertainty with every change.** For rare classes, label scarcity makes estimates noisy; the validation gate should treat wide intervals as a reason for caution, not as reassurance.
- **Do not assume monotonic degradation.** Some metrics improve temporarily or in one direction, so alarms should be triggered by deviation from the baseline in either direction when the decision depends on stability.
- **Prefer calibration-aware evidence.** Brier score and calibration changes can move independently of discrimination measures such as AUROC, so both families should be tracked.
- **Local adaptation matters.** Because sites respond differently, a single global update may help one site and harm another; adaptation should be validated per site before release.

ADAPTIVE RESPONSE POLICY

The response layer turns evidence states into actions. Table III gives a representative policy. Responses are ordered by cost and risk: lightweight actions are applied first, and heavier ones only when evidence persists.

TABLE III. EXAMPLE RESPONSE POLICY

Evidence state	Typical condition	Response
Green	No material drift; calibration stable	Continue automated operation; routine sampling
Amber: input drift only	Significant, persistent drift with stable labelled performance	Increase labelling; recalibrate; narrow auto-actions; alert owner
Amber: calibration loss	Rising Brier score or reliability error	Recalibrate; raise review rate for uncertain cases
Red: performance decline	Labelled metrics below threshold with confidence	Fall back to validated model or human; start retraining
Red: data failure	Schema or quality violation	Block affected inputs; escalate to data steward

A. Recalibration and Re-weighting

When drift is moderate and discrimination is stable, probability recalibration, for example by temperature scaling or isotonic regression on recent labelled data, can restore calibration at low cost. Importance weighting can adjust for covariate shift when density ratios can be estimated reliably. These measures are fast and reversible but do not help when the relationship between inputs and outcomes has changed.

B. Retraining and Fine-Tuning

When labelled evidence indicates a performance decline, or when drift is large and persistent, the model should be retrained or fine-tuned on recent data. To avoid forgetting stable patterns, training data should combine recent and historic samples, and candidate models should be validated against both. The decision to promote a candidate should rely on a pre-registered validation protocol, including per-site and per-class checks, not on aggregate accuracy alone.

C. Safe Deployment of Updates

New models should be released in shadow mode, where they score live traffic without affecting decisions, and then through canary deployment to a small fraction of traffic with automatic rollback if evidence worsens. Versioning of data, code, features and models allows any decision to be reproduced. Every update, including the evidence that justified it, is written to the audit ledger.

D. Deferral and Human Oversight

In the red state, or whenever uncertainty exceeds a threshold, the system should abstain and route the case to a human. Selective prediction, in which the model answers only when its confidence is adequate, trades coverage

DOI: 10.46121/pspc.54.4.03

for reliability and provides a continuous mechanism for deferral. Human reviewers also supply the labels needed to refresh evidence, so deferral and learning reinforce one another. Reviewer workload must be modelled so that thresholds do not overwhelm capacity.

GOVERNANCE AND THE AUDIT LEDGER

The audit ledger is a structured, append-only record that stores for each monitoring window the data snapshot identifiers, the evidence indicators with uncertainty, the resulting trust state, the response taken, the approver, and the model and data versions in use. Cryptographic hashing or write-once storage makes tampering detectable. This record serves several purposes. It enables incident investigation by reconstructing why a given decision was made. It provides the documentation of monitoring and logging that frameworks such as the NIST AI Risk Management Framework and the EU AI Act expect [17], [18]. It supports model documentation practices, which can be extended to include live reliability summaries beside static descriptions of intended use [19].

Roles must be explicit. A model owner is accountable for performance and approves updates; a data steward is accountable for input quality; a risk function sets thresholds and reviews exceptions; and an independent validator checks candidate models. Thresholds for trust states, escalation timelines, and the criteria for decommissioning a model should be documented and reviewed after significant incidents. Evidence about subgroups also supports fairness oversight, since drift that affects one population more than another is an equity risk as well as a performance risk, as the site-level variation in Fig. 3 illustrates.

IMPLEMENTATION GUIDANCE

A practical implementation can be built incrementally on standard MLOps tooling. Table IV lists the core indicators, the typical computation and the intended use. The following sequence has proved sensible in similar systems.

TABLE IV. CORE MONITORING INDICATORS

Indicator	Computation	Use
Feature and embedding drift	Two-sample tests or distances with effect size, per window	Early warning of input shift
Output distribution	Change in predicted class or score distribution	Detect prior shift and model behaviour changes
Calibration	Brier score, reliability error, calibration curves	Probability reliability for decisions
Per-site and per-class performance	Metrics with confidence intervals on labelled samples	Locate harm; detect opposing site effects
Data quality	Missingness, schema and range violations	Prevent silent upstream failures
Operational	Latency, review rate, alarm precision	Tune thresholds and workload

- **Establish the reference.** Freeze reference windows and baseline metrics at the time of validation, including per-site and per-subgroup values with confidence intervals.
- **Instrument the pipeline.** Log inputs, embeddings, outputs and uncertainty for every prediction, subject to privacy requirements, and collect labels with known delay.
- **Select windows and thresholds deliberately.** Use the sample-size analysis in Section V to select window length, and define effect-size thresholds with domain experts.
- **Start with alerting, then add gating.** Run the evidence layer in observation mode, tune thresholds against historical incidents, and only then allow it to constrain automated actions.
- **Rehearse failure.** Test rollback, fallback and deferral paths using simulated drift, and run periodic drills.
- **Review regularly.** Evaluate alarm precision, detection delay and reviewer workload, and adjust the framework as the data and the business change.

LIMITATIONS

The framework has limitations. Evidence from proxies cannot replace ground truth, and drift in inputs does not always imply performance loss; conversely, concept drift may degrade performance without any visible change in the inputs. The illustrations in Figs. 2 and 3 summarise drift experiments under synthetic shifts and are intended

DOI: 10.46121/pspc.54.4.03

to show qualitative behaviour; details such as the test statistic, the shift generation and the dataset should be specified when reproducing them, and conclusions should be validated on the organisation's own data. The multimodal network in Fig. 1 is described at the architectural level, and the benefit of contrastive alignment for drift monitoring requires empirical evaluation. Finally, adaptive updating introduces its own risks, including feedback loops in which the model's actions alter the data it later sees, and adversaries who may attempt to induce drift deliberately, so adaptation must be conservative and monitored.

FUTURE RESEARCH DIRECTIONS

Several directions follow. Unlabelled performance estimation, which predicts accuracy from model outputs and drift statistics, could strengthen proxy evidence. Conformal methods that remain valid under drift would provide prediction sets with reliable coverage. Causal and shift-stable modelling could reduce sensitivity to site-specific changes [20]. Federated and site-specific adaptation could allow institutions to benefit from shared learning without exchanging data. Standardised benchmarks and shared reporting of drift evidence across multimodal settings would improve comparability. Finally, human-factors studies on how reviewers use trust states and explanations would help calibrate the interaction between automation and oversight.

CONCLUSION

Reliable AI decision-making cannot be established once and assumed thereafter. Data keep changing, and the evidence that a model remains trustworthy has to be produced continuously. This paper proposed a self-validating framework in which a multimodal representation layer supplies comparable signals, an evidence layer quantifies drift, uncertainty and quality, a validation gate converts evidence into trust states, an adaptive response layer acts proportionately, and an audit ledger records the reasoning.

The illustrative analyses yielded concrete design rules. Drift-test evidence depends on both shift magnitude and sample size, so windows must be sized for the smallest consequential shift and p-values read together with effect sizes and persistence. The same shift can help one site and harm another, and aggregate metrics conceal large class-level swings, so evidence must be computed at the level of site and class with uncertainty reported alongside. Response policies should favour cheap, reversible actions first, validate candidates per site, release updates safely and defer to humans when evidence is weak.

Organisations that adopt this approach gain more than robustness. They gain the ability to explain, at any moment, why an automated decision was or was not trusted, which is the foundation of accountable AI in high-stakes settings.

REFERENCES

1. **Pawan Kumar Singh,** Scaling Gate-All-Around (GAA) Transistor Architectures for Sub-2nm AI Accelerators Using Atomically Thin 2D Materials, Vol. 33 No. 08 (2024): JOCAAA, Journal Name: Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), https://eudoxuspress.com/index.php/pub/article/view/5713
2. **Pawan Kumar Singh,** Energy-Efficient Task Scheduling in Green Cloud Computing Using Neuromorphic Spiking Neural Networks, Vol. 33 No. 08 (2024): JOCAAA, Journal Name: Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), https://eudoxuspress.com/index.php/pub/article/view/5712
3. **Rangavinay Teja Royal Jagga**, 2026, DOI: 10.1109/GCWCN66157.2025.11448495, Date of Conference: 22-23 November 2025, Date Added to IEEE Xplore: 25 March 2026 ,https://ieeexplore.ieee.org/document/11448495
4. **Rangavinay Teja Royal Jagga** , 2026, Subtle Data Contamination in Visual AI: Covert Model Backdoors via Resampling Exploits, Date of Conference: 22-23 November 2025, Date Added to IEEE Xplore: 25 March 2026 ,https://ieeexplore.ieee.org/document/11448514
5. **Rangavinay Teja Royal Jagga**, 2026, Robotic Tactile-Vision Integration for Interactive Object Disentanglement, https://ieeexplore.ieee.org/document/11448337

DOI: 10.46121/pspc.54.4.03

6. **Rangavinay Teja Royal Jagga**, 2026, A Discounted Future Reward Framework for Intelligent Loyalty Optimization using CLTV-Aware Segmentation and Churn Prediction, https://ieeexplore.ieee.org/document/11497560
7. **Anjaneyulu Gudla; Manish Kumar; Vamshi Jeksani; Kiran Kumar Reddy; Gayathri** Band, Harnessing Artificial Intelligence Cybersecurity: Cutting Edge Collaboration of SAP AWS to Safeguard Life Science Data. 2026 https://doi.org/10.1109/aiei69164.2026.11497833
8. **G. Mahender, Department of CSE(Data Science), Vardhaman College of Engineering ; Vamshi Jeksani; Manish Kumar; Koppuravuri Sri Lakshmi Sruthi; Kiran Kumar Reddy**, SAP Cloud System on AWS for Risk Detection and Integration Artificial Intelligence Scalable Cybersecurity 2026 https://doi.org/10.1109/aiei69164.2026.11497435
9. **Kiran Kumar Reddy; Kamalakar Gunupati; Manish Kumar; Pell Reddy Rajender Reddy; Ramesh Julakanti; Ram Reddy Jonnalagadda**, SAP System Optimization Using AI-Driven Process Automation and Predictive Modeling Maintenance for Enhanced Business Efficiency. 2025 https://doi.org/10.1109/computingcon64838.2025.11376613
10. **Amilineni, M.V. 2025**. Audit-Safe Agentic RAG Framework for Regulated Enterprise Systems: A Trustworthy AI Architecture for SAP, Finance, Healthcare, and Government Workflows. INTERNATIONAL JOURNAL OF ADVANCES IN SIGNAL AND IMAGE SCIENCES. (Jan. 2025), 34–46. DOI: https://doi.org/10.29284/7fjzk883
11. **Pawan Kumar Singh**, ACCELERATING LARGE LANGUAGE MODEL TRAINING USING HYBRID QUANTUM-CLASSICAL VARIATIONAL AUTOENCODERS WITH INTEGRATED PHOTONIC QRAM, Vol. 54 No. 3 (2026): July-September 2026, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/501, DOI: https://doi.org/10.46121/pspc.54.3.45
12. **Pawan Kumar Singh** , RESILIENCE ASSESSMENT OF MODERN POWER GRID INFRASTRUCTURE AGAINST EXTREME WEATHER USING PHYSICS-INFORMED NEURAL NETWORKS, Vol. 54 No. 1 (2026): January-March 2026, Power System Protection and Control, ISSN-1674-3**415,**https://pspac.info/index.php/dlbh/article/view/502, DOI: https://doi.org/10.46121/pspc.54.1.64
13. **Pawan Kumar Singh,** MITIGATING THE EXABYTE STORAGE CRISIS: ENHANCING READ/WRITE EFFICIENCY IN 5D OPTICAL GLASS STORAGE USING SILVER NANOPARTICLE DOPING, Vol. 38 No. 11s (2025), International Journal of Applied Mathematics (IJAM) , pISSN 1311-1728, eISSN 1314-8060 ,https://ijamjournal.org/ijam/publication/index.php/ijam/article/view/2153
14. **Mr. Pawan Kumar Singh (Author), Miss. Pari Nidhi Singh (Author**) ,Intelligent Bot, ASIN : B0DBGCQY8X,Publisher : Namya Press, Publication date : November 25, 2022,Language : English, Print length : 105 pages,ISBN-10 : 9390124867,ISBN-13 : 978-9390124862,Item Weight : 7.5 ounces, Dimensions : 6.24 x 0.43 x 9.24 inches,https://www.amazon.com/Intelligent-Bot-Pawan-Kumar-Singh/dp/B0DBGCQY8X/ref=sr_1_5?crid=2VAXN5BRN5SMS&dib=eyJ2IjoiMSJ9.v3OCQvHBpy20RIjfMWh0qbOl-RLOlUJIC8Xa0Pjeqilvo2sbZ_suN3LbSCw_OA_ZD_4lgjRLMGlBwZQQ1ehZEfDpt_MJiOUfVuqqzMav0I5_mb6pN5DQog0U7SA826n1DQQ3QAtWwopHy4lC5IfhZXttKj9faiyATwhdtDATSSV1-r7FmzM2VQyIA0jxcfq6E6hp9ECfH1fby0li9kUfGpLxMq4GktlzUIJO5ggrhsQ.G1AjQiVqU4HFsnSczvz-X6NXCNs3EFSn-RI8mcWWGOc&dib_tag=se&keywords=Pawan+Kumar+Singh%2C+book&qid=1787815741&sprefix=pawan+kumar+singh%2C+bo%2Caps%2C450&sr=8-5
15. **Mr. Pawan Kumar Singh (Author),** STORAGE CRISES DUE TO RISE OF AI: A Comprehensive Guide to Managing the Storage Challenges in the Age of Artificial Intelligence, ASIN : B0HFK12ZCT, Publisher : EB1 Profile, Accessibility : Learn more, Publication date : August 17, 2026,Edition : 1st Language : English, File size : 1.7 MB, Screen Reader : Supported, Enhanced typesetting : Enabled X-Ray : Not Enabled, Word Wise : Not Enabled ,Print length : 261 pages ,ISBN-13 : 979-8951479112,Page Flip : Enabled, https://www.amazon.com/STORAGE-CRISES-DUE-RISE-Comprehensive-

DOI: 10.46121/pspc.54.4.03

ebook/dp/B0HFK12ZCT/ref=sr_1_1?crid=9MF6C4WDVJR5&dib=eyJ2IjoiMSJ9.DQ8fesBch7ELmRnilDt-YA.NBSCHLmetvnbhkxwIpOoj3s0IVsCV8y6-ZCfCP6Ul1c&dib_tag=se&keywords=Pawan+Kumar+Singh%2C+book%2C+eb1+Profile&qid=1787816040&sprefix=pawan+kumar+singh%2C+book%2C+eb1+profile%2Caps%2C420&sr=8-1

16. **Parag Bharatkumar Parikh**, FRAMEWORKS FOR MITIGATING SUPPLY CHAIN DISRUPTIONS AND MATERIAL PRICE VOLATILITY IN SUSTAINABLE BUILDING CONSTRUCTION, Vol. 53 No. 2 (2025): April-June 2025 , , Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/447, DOI: https://doi.org/10.46121/pspc.53.2.34
17. **Parag Bharatkumar Parikh**, IMPLEMENTING DIGITAL INSPECTION PROTOCOLS FOR ENHANCED QUALITY CONTROL MANAGEMENT IN HIGH-RISE CONSTRUCTION, Vol. 54 No. 3 (2026): July-September 2026, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/439, DOI: https://doi.org/10.46121/pspc.54.3.12
18. **Parag Bharatkumar Parikh** , EVALUATING THE FEASIBILITY AND EFFICIENCY OF 3D-PRINTED CONCRETE TECHNOLOGY IN SUSTAINABLE MODULAR CONSTRUCTION, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/issue/view/30, DOI: https://doi.org/10.46121/pspc.54.1.61
19. **Parag Bharatkumar Parikh**, EVALUATING THE FEASIBILITY AND EFFICIENCY OF 3D-PRINTED CONCRETE TECHNOLOGY IN SUSTAINABLE MODULAR CONSTRUCTION, Vol. 54 No. 1 (2026): January-March 2026, System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/440, DOI: https://doi.org/10.46121/pspc.54.1.61
20. **Ronakkumar shah,** REDEFINING BUSINESS PROCESS MANAGEMENT: THE SYNERGY OF AI AND INTELLIGENT AUTOMATION, Vol. 54 No. 3 (2026): July-September 2026, System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/448,
21. DOI: https://doi.org/10.46121/pspc.54.3.14
22. **Ronakkumar shah**, OPERATIONAL ACCOUNTABILITY IN ACTION: METRICS AND MECHANISMS FOR SYSTEMIC SUCCESS, Vol. 53 No. 4 (2025): October-December 2025, System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/452,
23. DOI: https://doi.org/10.46121/pspc.53.4.40
24. **Ronakkumar shah**, BEYOND THE BLUEPRINT: AUTOMATED WORKFLOW ENFORCEMENT AND POLICY GOVERNANCE, Vol. 53 No. 2 (2025): April-June 2025 , System Protection and Control, ISSN-1674-3415 ,https://pspac.info/index.php/dlbh/article/view/451, DOI: https://doi.org/10.46121/pspc.53.2.35
25. **Tushita Agarwal; Swaminathan Sethuraman,** Mapping Customer Behavior: Visual Analysis of Online Retail Interactions, Date Added to IEEE Xplore: 17 February 2026,
26. DOI: https://doi.org/10.1109/IC2NC67409.2025.11376261
27. **Tushita Agarwal,** Efficient Data Summarization and Comparison via Grid-Based Statistical Proxies for Large-Scale Datasets**, Date Added to IEEE Xplore: 03 December 2025,**
28. DOI: https://doi.org/10.1109/ICBATS66542.2025.11258564
29. **Dip Bharatbhai Patel,** Building scalable business intelligence systems in the cloud: A technical approach, International Journal of Science and Research Archive, 2021, 02(01), 212-215., Received on 03 February 2021; revised on 17 April 2021; accepted on 19 April 2021,DOI: https://doi.org/10.30574/ijsra.2021.2.1.0047
30. **Dip Bharatbhai Patel,** University of North America, Virginia, USA, Home / Archives / Vol. 5 No. 02 (2025): Volume 05 Issue 02 / Articles, https://www.academicpublishers.org/journals/index.php/ijdsml/article/view/5958
31. **Sai Akshara Vemuri ,** KINEMATIC OPTIMISATION OF SOFT ROBOTIC MANIPULATORS FOR MEDICAL INTERVENTION, Vol. 54 No. 3 (2026): July-September 2026, System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/450
32. **Sai Akshara Vemuri**, https://www.periodicos.ulbra.org/index.php/acta/article/view/528

33. **Aziz, A., Chandusha, K., Naga Kiran, B., Prasad, M. L. M., Golla, K., Rajeswari, S., & Katta**, S. K. R. (2026). Exploring Machine Learning Applications for Genomic Data Analysis in Personalized Medicine. International Journal of Drug Delivery Technology, 16(38s), 1166-1173. https://doi.org/10.25258/ijddt.16.38s.127,https://www.periodicos.ulbra.org/index.php/acta/article/view/687
34. **Mohammad Naffizuddin , Chaganti Ashok ,Sowjanya Gunukula , Rvbs Sarma , Bhavana Sujanamulk , Bharani Krishna Takkella , Chukka Ram Sunil ,Varri Sujana , Ashwini Kumar,** Etiological Factors of Orofacial Clefts in a Hospital-Based Population: A Descriptive Analytical Study, Published: December 10, 2025,DOI: 10.7759/cureus.98868, DOI : https://doi.org/10.7759/cureus.98868
35. **Manoj Kumar Kagitha**, Automated Security Compliance in Ci/Cd: A Natural Language Processing (NLP) Approach to Detecting Vulnerabilities in Infrastructure-As-Code (Terraform/Bicep), Publication year 2021, Volume 2021, Issue 12, https://computerfraudsecurity.com/index.php/journal/article/view/970
36. **Manoj Kumar Kagitha,** Hybrid Quantum-Classical Resource Scheduling: A Proposed Architecture for Next-Generation Hyperscale Data Centers, Publication year 2024, Volume 2024, Issue 12, https://computerfraudsecurity.com/index.php/journal/article/view/969
37. **Manoj Kumar Kagitha,** Self-Healing Cloud Infrastructure: A Hybrid Reinforcement Learning Framework for Predicting Microservice Failures in Azure Kubernetes Service (AKS**),**Publication year 2023, Volume 2023, Issue 12,https://computerfraudsecurity.com/index.php/journal/article/view/971
38. **Manoj Kumar Kagitha**, GREENOPS IN THE CLOUD: AN AI-DRIVEN TELEMETRY ANALYSIS MODEL FOR REDUCING CARBON FOOTPRINT IN HIGH-AVAILABILITY CLUSTERS, Vol. 28 No. 6 (2020): JOCAAA, https://www.eudoxuspress.com/index.php/pub/article/view/5054
39. Lakshmi Rahul Reddy Mareddy. (2026). Security and Performance Assessment of Encryption Techniques for Cloud Platforms. International Journal of Communication Networks and Information Security (IJCNIS), 18(1), 68–80. Retrieved from https://ijcnis.org/index.php/ijcnis/article/view/8735
40. Lakshmi RahulReddy Mareddy, Context-Aware Learning Approaches for Improving Prediction Accuracy in Dynamic Systems, International Journal of Computational Mathematical Ideas (IJCMI), Vol.17 No.3(2025): IJCMI Vol17 NO3 202 5 https://www.ijcmi.com/index.php/journal/article/view/15
41. Lakshmi Rahul Reddy , Mareddy. (2026). Revisiting Autonomy in Artificial Intelligence through Agentic System Design. Journal of Computational Analysis and Applications (JoCAAA), 35(1), 707–713. https://eudoxuspress.com/index.php/pub/article/view/4780
42. **Mahendrakumar Kalal**, Lean-Driven SAP Production Planning Optimization Using Machine Learning for Inventory and Throughput Efficiency, Date of Conference: 19-20 March 2026, Date Added to IEEE Xplore: 08 April 2026,ISBN Information:,ISSN Information: DOI: 10.1109/ISDFS69419.2026.11459029, https://ieeexplore.ieee.org/abstract/document/11459029
43. **Mahendra kumar Kalal**, SUPPLY CHAIN RESILIENCE THROUGH INTEGRATED SAP PRODUCTION PLANNING – PROCUREMENT SYNCHRONIZATION, Archives / Vol. 54 No. 1 (2026): January-March 2026 / Articles, https://pspac.info/index.php/dlbh/article/view/259
44. **Dr Manoj Goel (December 2018**). THE ATTITUDE OF DIFFERENT STATES STUDENTS TOWARDS PHYSICAL EDUCATION; STUDYING IN PHYSICAL EDUCATION COLLEGES OF MAHARSHTRA STATE, https://ijeponline.com/index.php/journal/article/view/965/847

DOI: 10.46121/pspc.54.4.03

45. **Ilakiya Ulaganathan, Tagore Engg. College, Anna University, Chennai, India, ; FNU Antara; Santhosh Chitraju Gopal Varma; Zokir Mamadiyarov; Ziyadulla Turayev, Published** in: 2025 International Conference on Sustainability, Innovation & Technology (ICSIT), Date of Conference: 22-23 August 2025, Date Added to IEEE Xplore: 19 December 2025, ISBN Information:
https://ieeexplore.ieee.org/document/11294748/references#references
46. **Shreyas Kasture, Sudhanshu Maurya, Amit Kumar Sharma, Santhosh Chitraju Gopal Varma, Kashish Mirza, Firdous Sadaf Mohammad Ismail, DOI**: http://doi.org/10.11591/ijai.v15.i2.pp1746-1761, IAES International Journal of Artificial Intelligence (IJ-AI), ISSN/e-ISSN 2089-4872/2252-8938
https://ijai.iaescore.com/index.php/IJAI/article/view/26932
47. **Shiv Shankar Dwivedi, Mr. Santosh Kumar Sharma,** Future-Proofing Data Security with Quantum-Resistant Cryptography, Vol. 11 No. 3 (2025): IJCESEN,
DOI: https://doi.org/10.22399/ijcesen.3447, https://ijcesen.com/index.php/ijcesen/article/view/3447
48. **Vijender Kalmotia,** Enhancing statistical accuracy in gene perturbation studies, Volume 266, August 2026, 105819, https://doi.org/10.1016/j.biosystems.2026.105819,
https://www.sciencedirect.com/science/article/pii/S0303264726001292
49. **Naveen Undrathi,** Design and implementation of generative AI-driven data transformation pipelines using spring boot for scalable financial applications, eISSN: 2581-9615 || CODEN: WJARAI || Impact Factor 8.2 || CrossRef DOI, September 2026 (Volume 31, Issue 3), DOI url: https://doi.org/10.30574/wjarr.2026.30.3.1629,
50. **Naveen Undrathi,** Agentic generative AI-enabled microservices framework for intelligent financial data processing in java-based enterprise systems, September 2026 (Volume 31, Issue 3), eISSN: 2581-9615 || CODEN: WJARAI || Impact Factor 8.2 ||
https://doi.org/10.30574/wjarr.2026.30.3.1628
51. **Akshay Jain,** Comparative Analysis of Linkage Methods in Hierarchical Clustering for Customer Segmentation, Conference paper, First Online: 07 October 2025,
https://link.springer.com/chapter/10.1007/978-981-96-7289-9_44
52. **Ismail Hossain,** Department of Computer Science, George Mason University, Fairfax, VA, USA, Mridul Banik, Department of Computer Science, Colorado State University, Fort Collins, CO, USA, https://ieeexplore.ieee.org/abstract/document/11446534
DOI: 10.1109/FMLDS67896.2025.00106
53. **Ismail Hossain, Mridul Banik , Fahmid Al Farid , Jia Uddin , Hezerul bin Abdul Karim** , Advancing Large Language Models for Low-Resource Languages: A Systematic Review of Pretraining, Adaptation, and Ethical Challenges, 10.32604/cmes.2026.075507, Date: 2026
https://www.sciencedirect.com/org/science/article/pii/S152614922600322X
54. **Ismail Hossain; Mridul Banik,** Bridging Visual and Linguistic Intelligence for Chest X-rays: A Comprehensive Review of ViTs and LLM Synergies, Published in: 2025 IEEE International Conference on Future Machine Learning and Data Science (FMLDS) Date of Conference: 02-05 November 2025, Date Added to IEEE Xplore: 25 March 2026, ISBN Information: https://ieeexplore.ieee.org/abstract/document/11446534
55. **Rifat Sarker Aoyon1, Ismail Hossain2, M. Abdullah-Al-Wadud3, Jia Uddin4,** A Secured and Continuously Developing Methodology for Breast Cancer Image Segmentation via U-Net Based Architecture and Distributed Data Training,
https://www.sciopen.com/article/10.32604/cmes.2025.060917
56. **M. B. Vibha, M. Chandrika, Samitha Khaiyum & P. Rakshitha Kiran,** Predicting sedentary behavior in adults using stacked LSTM modeling, published: 27 November 2024
Volume 16, pages 346–355 (2025), https://link.springer.com/article/10.1007/s13198-024-02622-2
57. **Bellary Chiterki Anil, Jayasimha Sondekoppa Rajkumar, T. L. Divya, Samitha khaiyum**

Power System Protection and Control
ISSN:1674-3415

Rakshitha Kiran P., Balakrishnan Ramadoss, A Radiomics-based Framework for Liver Cancer Analysis using Explainable Artificial Intelligence (XAI) Methods, Vol. 15 No. 3 (2025): June, 2025, https://etasr.com/index.php/ETASR/article/view/10377

DOI: 10.46121/pspc.54.4.03