

# Temporal Pattern Recognition and Unsupervised Anomaly Detection for Early Warning of Disease Progression in Longitudinal Health Records

Sunday Oladimeji Adegoke  
Honeywell Flour Mills  
Lagos, Nigeria

**Abstract:** Longitudinal health records provide a rich temporal view of patient trajectories, enabling early identification of disease progression before overt clinical deterioration occurs. As healthcare systems increasingly rely on continuous data streams from electronic health records, wearable sensors, and remote monitoring platforms, robust temporal pattern recognition methods have become critical for proactive clinical decision-making. However, early warning systems in healthcare must operate under conditions of sparse labels, heterogeneous data sources, and evolving disease dynamics, making unsupervised anomaly detection particularly well suited to this domain. This work examines temporal pattern recognition and unsupervised anomaly detection frameworks for early detection of disease progression in longitudinal health records. From a broad perspective, the study situates anomaly detection within predictive population health, chronic disease management, and preventive care, emphasizing its role in identifying subtle deviations from personalized baselines rather than population-wide thresholds. The discussion then narrows to temporal modeling techniques, including sequence-aware statistical models, representation learning approaches, recurrent and attention-based architectures, and time-series clustering methods that capture latent disease dynamics without reliance on labeled progression events. Key methodological challenges are addressed, including handling irregular sampling, missing data, patient heterogeneity, and non-stationary clinical processes. The abstract further highlights evaluation strategies that prioritize early detection sensitivity, temporal consistency, and clinical relevance, rather than retrospective classification accuracy alone. Particular attention is given to minimizing false alarms and supporting clinician interpretability to ensure practical adoption in real-world settings. The study concludes by identifying future research directions, including integration of multimodal longitudinal data, adaptive personalization of anomaly thresholds, and alignment with explainability and safety requirements for clinical deployment. Temporal unsupervised anomaly detection is positioned as a foundational capability for scalable, early-warning disease surveillance and individualized care pathways.

**Keywords:** Temporal pattern recognition; Unsupervised anomaly detection; Longitudinal health records; Disease progression modeling; Early warning systems; Healthcare time-series analysis

## 1. INTRODUCTION

### 1.1 Background and Motivation

The rapid expansion of longitudinal health data has fundamentally reshaped clinical observation and monitoring [1]. Electronic health records now capture repeated laboratory tests, vital signs, medication histories, and clinical assessments over extended periods, while wearable devices and remote monitoring technologies generate continuous streams of physiological data outside traditional care settings [2]. Together, these developments have shifted clinical practice from episodic measurement toward persistent observation.

This transition has intensified the clinical need for early warning systems capable of detecting disease progression before overt deterioration occurs. Many acute and chronic conditions evolve gradually, with subtle deviations preceding clinically obvious events [3]. Early identification of such changes enables timely intervention, improved outcomes, and more efficient use of healthcare resources. However, translating dense longitudinal data into actionable insight remains challenging.

Conventional rule-based thresholds are poorly suited to this task. Fixed limits fail to account for patient-specific baselines, temporal trends, and multivariate interactions, often triggering late or excessive alerts [4]. Supervised machine learning approaches, while powerful, depend on labeled outcomes that are scarce, inconsistently defined, or retrospectively assigned in real-world datasets [5]. Moreover, supervised models tend

to focus on prediction of discrete events rather than continuous progression.

These limitations motivate a shift toward temporal pattern recognition using unsupervised machine learning. By learning normative trajectories directly from longitudinal data, unsupervised models can identify deviations without relying on predefined labels [6]. When coupled with explicit temporal modeling, such approaches offer a promising pathway for early, patient-specific detection of disease progression grounded in observed clinical behavior rather than rigid rules.

### 1.2 Problem Definition and Research Gap

Disease progression can be conceptualized as a latent process that unfolds over time rather than a sequence of isolated events. Physiological systems adapt, compensate, and deteriorate gradually, producing trajectories that encode clinically meaningful information [3]. Detecting progression therefore requires models that capture temporal structure, persistence, and directionality.

A fundamental challenge is the absence of reliable labels for progression events in routine clinical data. Many deteriorations are undocumented until late stages, and diagnostic timestamps often lag behind physiological change [7]. This makes supervised learning ill-suited for early detection, as the ground truth is incomplete, delayed, or ambiguous.

Existing anomaly detection methods only partially address this challenge. Many approaches operate on static snapshots, ignore temporal dependencies, or rely on population-level

norms that obscure individual variability [8]. Even when time is considered, models often lack mechanisms to represent evolving baselines or sustained deviation.

The resulting gap lies in the lack of patient-specific, time-aware unsupervised methods that explicitly model longitudinal patterns. There is a need for frameworks that learn normal temporal behavior from individual records and identify progression as deviation from learned trajectories rather than from population averages or fixed thresholds [9].

### 1.3 Research Objectives and Contributions

This work aims to address the identified gap by developing a temporal machine learning framework for early anomaly detection in longitudinal health records. The primary objective is to model disease progression as deviation from learned patient-specific temporal patterns rather than as classification of predefined events [1].

The framework employs unsupervised learning to capture normative trajectories directly from longitudinal data, avoiding dependence on sparse or delayed labels [6]. Temporal structure is explicitly encoded to distinguish transient variability from sustained change, enabling early identification of clinically relevant deviations.

A further objective is rigorous quantitative comparison with established clinical standards, including rule-based thresholds and composite early warning scores. Performance is evaluated using safety-relevant metrics such as detection latency and false-alarm burden to ensure clinical relevance [4].

Crucially, the framework emphasizes interpretability and safety. Model components are designed to produce transparent anomaly scores linked to observable features and temporal patterns, supporting clinician trust and auditability [7]. By aligning computational outputs with clinical reasoning, this work contributes an approach that is not only technically robust but also suitable for safety-critical decision support.

### 1.4 Manuscript Organization

The remainder of this manuscript is organized as follows. Section 2 reviews conceptual foundations and regulatory considerations for safety-critical anomaly detection. Section 3 describes the longitudinal data sources, temporal structure, and safety-preserving preprocessing strategies. Section 4 presents the interpretable representation and temporal feature encoding framework. Section 5 details the explainable anomaly detection mechanism and training strategy. Section 6 introduces anomaly scoring and thresholding methods. Section 7 reports quantitative and qualitative evaluation results. Section 8 discusses clinical implications and limitations, and Section 9 concludes with future research directions.

## 2. RELATED WORK AND THEORETICAL FOUNDATIONS

### 2.1 Longitudinal Data Modeling in Healthcare

Longitudinal data modeling has become central to modern healthcare analytics as electronic health records increasingly capture repeated observations over time [7]. Laboratory tests, vital signs, medication records, and clinical assessments form multivariate time-series that reflect patient evolution rather than isolated states. Modeling such data enables analysis of disease dynamics, treatment response, and early deterioration.

Disease trajectories differ fundamentally from cross-sectional snapshots. While snapshots provide momentary assessment,

trajectories encode progression, recovery, and compensation mechanisms that unfold over time [10]. Clinicians routinely rely on trends rather than single measurements when making decisions, highlighting the importance of temporal modeling.

However, longitudinal healthcare data present significant challenges. Sampling is irregular, driven by clinical events rather than fixed schedules, complicating the application of standard time-series methods [8]. Missingness is pervasive and often informative, reflecting clinical judgment rather than data loss. Additionally, measurements may vary in frequency and reliability across variables.

These characteristics limit the applicability of conventional time-series assumptions such as stationarity and uniform sampling. Effective longitudinal modeling must therefore accommodate irregular timing, heterogeneous variables, and evolving baselines while preserving clinical interpretability [12]. Addressing these challenges is essential for meaningful early detection of disease progression.

### 2.2 Temporal Pattern Recognition Methods

A broad range of temporal pattern recognition methods has been applied to healthcare data. Classical statistical time-series models, such as autoregressive and state-space formulations, provide interpretable representations of temporal dependence [9]. These approaches are effective for regularly sampled signals but struggle with irregular measurement and high-dimensional clinical data.

Hidden state models, including hidden Markov models, introduce latent variables to represent unobserved disease states [13]. Such models align conceptually with disease progression but require strong assumptions about state transitions and often depend on labeled or semi-labeled data, limiting scalability in real-world settings.

More recent work explores deep temporal representations, including recurrent and convolutional architectures, to capture complex dependencies across time [11]. These models can handle high-dimensional inputs and irregular sampling through embedding strategies. However, their internal representations are often opaque, making it difficult to interpret detected patterns or justify alerts in clinical contexts.

Across these approaches, a common limitation is the tension between expressiveness and interpretability. Methods that capture rich temporal structure frequently sacrifice transparency, while interpretable models may oversimplify complex trajectories [15]. Bridging this gap remains a central challenge in temporal healthcare analytics.

### 2.3 Unsupervised Anomaly Detection in Clinical Data

Unsupervised anomaly detection has been widely explored as a means to identify abnormal clinical patterns without relying on labeled outcomes [14]. Density-based methods estimate the distribution of normal data and flag observations in low-density regions as anomalous. While conceptually straightforward, these approaches often struggle with high-dimensional and heterogeneous clinical features.

Reconstruction-based methods learn compact representations of normal patterns and measure deviation through reconstruction error [10]. Such methods are attractive for clinical data because they can be applied without explicit labels. However, many implementations operate on static feature vectors, neglecting temporal structure.

Distance-based methods assess deviation relative to reference populations or historical baselines [7]. These approaches align

with clinical reasoning but are sensitive to scaling and population heterogeneity. Without temporal context, distance metrics may misinterpret transient fluctuations as anomalies.

A critical limitation across much of the existing work is insufficient treatment of time. Many models reduce longitudinal records to aggregated snapshots or independent windows, losing information about persistence and directionality [12]. As a result, gradual disease progression may go undetected until late stages. Addressing temporal limitations is therefore essential for early anomaly detection.

#### 2.4 Early Warning Systems in Disease Progression

Early warning systems are widely used to identify patient deterioration in both acute and chronic settings. In intensive care units, composite scores aggregate vital signs and laboratory values to trigger escalation [9]. These systems are transparent and clinically familiar but rely on fixed scoring schemes and population-level thresholds.

In chronic disease monitoring, early warning focuses on detecting deviations from expected control, often using trend-based heuristics [13]. While effective in some contexts, such systems struggle with personalization and may generate excessive alerts.

A persistent challenge across early warning systems is false alarm burden. High alert rates contribute to alarm fatigue, reducing clinician responsiveness and trust [15]. Balancing sensitivity and specificity remains difficult when thresholds are static and patient variability is high.

These limitations motivate exploration of data-driven approaches that adapt to individual trajectories while retaining interpretability [8].

#### 2.5 Evaluation Metrics and Clinical Benchmarks

Evaluation of early detection systems requires metrics aligned with clinical decision-making rather than conventional classification accuracy [11]. Accuracy obscures temporal performance and fails to capture the cost of delayed detection or excessive alerts.

Deviation-based metrics quantify how early and how strongly a system detects departure from normal behavior, aligning with clinical reasoning [14]. Detection latency measures timeliness, while false positives per patient-day reflect operational burden.

Clinical benchmarks, such as rule-based thresholds and established early warning scores, provide meaningful baselines for comparison [10]. Aligning evaluation with these standards ensures that improvements are clinically relevant rather than purely statistical.

By emphasizing safety-relevant metrics and clinically grounded benchmarks, evaluation frameworks can better reflect real-world utility and support translation into practice [15].

### 3. DATA ACQUISITION AND STUDY DESIGN

#### 3.1 Data Sources and Cohort Definition

The study population is derived from longitudinal electronic health records capturing repeated observations across time for individual patients. Core data sources include laboratory measurements, vital signs, diagnostic codes, and clinical encounter metadata recorded during routine care [14]. These records provide multivariate time-series reflecting

physiological evolution rather than isolated clinical states. Where available, wearable or remote monitoring data, such as continuous heart rate or activity measures, are incorporated to enrich temporal resolution and extend observation beyond hospital settings [18].

Cohort definition follows clinically grounded inclusion and exclusion criteria. Adult patients with sufficient longitudinal depth, defined by a minimum number of observations over a specified period, are included to ensure meaningful temporal modeling [16]. Patients with extremely sparse records, incomplete identifiers, or inconsistent timestamps are excluded to reduce noise and ambiguity. Records associated with terminal events or procedures that fundamentally alter physiology may also be excluded to avoid conflating progression with intervention effects.

Temporal observation windows are defined relative to clinically meaningful anchors such as admission, diagnosis, or routine follow-up [21]. Each patient's record is segmented into rolling or decision-aligned windows that preserve temporal order while constraining look-ahead. This approach ensures that models learn from information available at the time decisions would realistically be made. By defining cohorts and observation windows in this manner, the dataset supports patient-specific temporal learning while maintaining clinical validity and safety relevance [22].

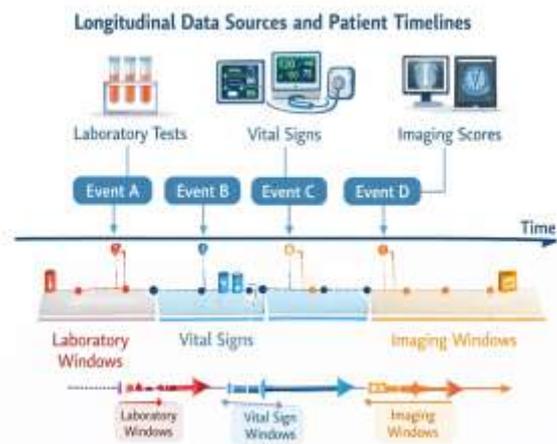


Figure 1: Overview of longitudinal data sources and patient timelines

#### 3.2 Data Preprocessing and Temporal Alignment

Preprocessing focuses on aligning heterogeneous clinical time-series while preserving temporal semantics. Irregular sampling is addressed by representing data on a common temporal grid, selected to balance resolution and data availability [15]. Rather than enforcing strict regularity, resampling strategies retain original timestamps and interpolate only where clinically reasonable.

Temporal alignment employs forward-filling for slowly varying variables, such as laboratory values, and interpolation for continuously monitored signals when gaps are short [19]. Missingness indicators are retained to distinguish absence of measurement from true stability, reflecting clinical decision processes [14]. Extreme outliers are flagged through range checks informed by physiological plausibility rather than removed indiscriminately.

All preprocessing steps are designed to avoid introducing future information into past observations. Normalization and

feature construction are performed within patient-specific windows to preserve individuality [17]. By maintaining temporal integrity and minimizing distortion, preprocessing supports downstream unsupervised learning that reflects real clinical dynamics rather than artefacts of data handling [20].

### 3.3 Ethical Considerations and Data Governance

Ethical oversight is central to the use of longitudinal clinical data. All records are de-identified prior to analysis, with direct identifiers removed and indirect identifiers minimized in accordance with data protection standards [22]. Access controls and audit logs are implemented to ensure accountability and prevent unauthorized use.

Bias and fairness considerations are addressed by examining representation across demographic subgroups and care settings [18]. Patient-specific modeling mitigates some population-level bias, but residual disparities may persist due to structural inequities in healthcare delivery. These risks are acknowledged and monitored throughout analysis.

Given the safety-critical nature of early warning systems, ethical responsibility extends beyond privacy to potential clinical impact [21]. Models are evaluated conservatively to avoid harm from false reassurance or excessive alerts. Transparent reporting and clinician oversight are treated as ethical requirements, ensuring that algorithmic outputs support, rather than undermine, patient safety and professional judgment [14].

### 3.4 Data Partitioning Strategy

Data partitioning is designed to preserve independence and temporal validity. Records are split at the patient level into training, validation, and test sets, ensuring that no individual contributes data to more than one partition [16]. This prevents leakage of patient-specific patterns that could inflate performance estimates.

Temporal leakage is further avoided by enforcing chronological splits within each patient’s record. Only data preceding a defined cutoff are used for model development, while later observations are reserved for evaluation [19]. This mirrors prospective deployment conditions.

Unsupervised training justifies this strategy, as the objective is to learn normative temporal behavior rather than predict labeled outcomes [20]. Validation focuses on stability and deviation detection rather than accuracy. Together, patient-level and temporal partitioning ensure that evaluation reflects real-world use, supporting safe translation of unsupervised temporal models into clinical decision support [22].

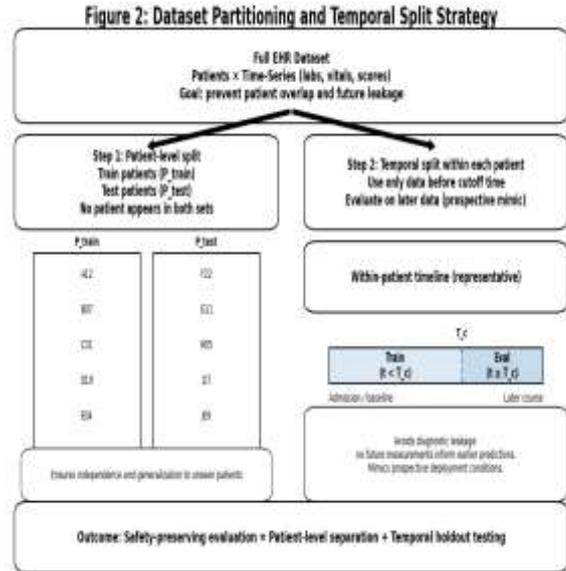


Figure 2: Dataset partitioning and temporal split strategy

## 4. FEATURE ENGINEERING AND TEMPORAL REPRESENTATION

### 4.1 Raw Feature Extraction

Raw feature extraction begins with clinically interpretable variables derived directly from longitudinal patient records. Core physiological variables include vital signs such as heart rate, respiratory rate, blood pressure, temperature, and oxygen saturation, alongside laboratory measurements reflecting metabolic, inflammatory, and organ-specific function. These features are selected due to their routine availability and established diagnostic relevance, ensuring applicability across diverse clinical settings.

In addition to absolute values, temporal frequency features quantify how often measurements occur within a given period. Measurement frequency itself carries clinical meaning, as increased sampling intensity often reflects clinician concern or patient instability. Event-based encoding further represents discrete clinical actions, such as medication administration, diagnostic imaging, or escalation of care. These events are encoded as binary indicators or counts aligned to the patient timeline.

Raw observations are retained with original timestamps prior to aggregation, preserving the asynchronous nature of clinical data. Feature extraction avoids premature summarization, ensuring that downstream temporal modeling operates on faithful representations of clinical observation processes. This foundation supports later construction of interpretable temporal patterns rather than opaque statistical abstractions.

### 4.2 Temporal Feature Construction

Temporal feature construction transforms raw observations into representations that capture evolution, persistence, and change. Sliding windows segment each patient’s timeline into overlapping intervals, enabling localized temporal summarization while preserving order.

Within each window, central tendency is computed using:

Window Mean

$$\mu_w = \frac{1}{n} \sum_{i=1}^n x_i$$

where  $x_i$  denotes the observed value at time point  $i$ ,  $n$  is the number of observations in the window, and  $\mu_w$  represents the average state during that interval.

Temporal trend is captured using a linear slope estimate:

Temporal Trend (Slope)

$$\beta_w = \frac{\sum_{i=1}^n (t_i - \bar{t})(x_i - \mu_w)}{\sum_{i=1}^n (t_i - \bar{t})^2}$$

Here,  $t_i$  represents observation time,  $\bar{t}$  is the mean time within the window, and  $\beta_w$  quantifies directional change.

Short-term instability is measured through dispersion:

Window Variance

$$\sigma_w^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \mu_w)^2$$

Lagged context is incorporated using prior-window values:

Lagged Feature

$$x_{w-1} = \mu_{w-1}$$

which represents the previous window's mean, enabling comparison across time.

Finally, robust instability is quantified using:

Mean Absolute Deviation (MAD)

$$MAD_w = \frac{1}{n} \sum_{i=1}^n |x_i - \mu_w|$$

MAD measures average deviation from recent baseline, remaining robust to transient spikes common in clinical data.

Figure 3: Temporal Feature Construction Using Sliding Windows

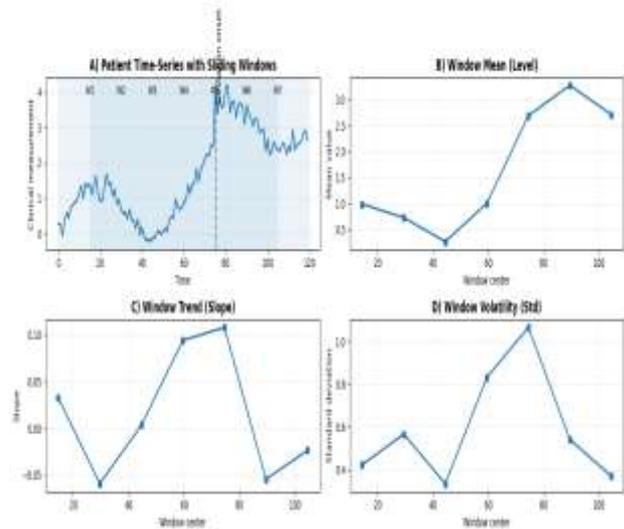


Figure 3: Temporal feature construction using sliding windows

### 4.3 Normalization and Patient-Specific Baselines

Normalization ensures that features are comparable while preserving individual clinical context. Global z-score normalization scales values relative to population statistics but can obscure meaningful patient-level deviation. Robust scaling approaches reduce sensitivity to extreme values by using central tendency measures derived from each patient's historical data.

Patient-specific baselines are constructed by estimating stable reference behavior during periods without major clinical events. Subsequent feature values are expressed relative to this personal baseline rather than cohort-wide norms. This mirrors clinical reasoning, where deviation from an individual's usual state often signals concern.

Normalization is performed within rolling windows to prevent future data leakage. By combining robust scaling with personalized baselines, the framework enhances anomaly sensitivity while maintaining interpretability.

### 4.4 Dimensionality Reduction and Representation Learning

Dimensionality reduction addresses redundancy among temporal features while enabling compact representation. Principal component analysis (PCA) provides a linear transformation that identifies orthogonal axes of maximal variance. Its explicit loadings allow clinicians to interpret latent dimensions in terms of contributing physiological variables.

Autoencoders offer nonlinear representation learning capable of capturing complex temporal interactions. When constrained to low-dimensional latent spaces and smooth temporal transitions, these models retain interpretability. Latent temporal embeddings summarize patient trajectories while preserving progression-relevant structure.

Dimensionality reduction is applied after temporal feature construction, ensuring that learned representations encode dynamics rather than raw noise. Latent dimensions are examined for correspondence with clinically recognizable patterns such as gradual deterioration or recovery phases.

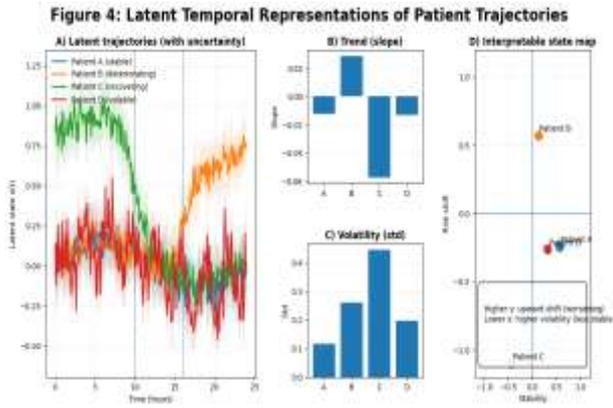


Figure 4: Latent temporal representations of patient trajectories

#### 4.5 Feature Stability and Drift Analysis

Feature stability is assessed to ensure consistent anomaly interpretation over time. Distributional drift is monitored across windows to distinguish genuine clinical change from artefacts introduced by measurement practices or workflow variation. Gradual drift may reflect disease progression, whereas abrupt shifts often indicate data quality issues.

Uncorrected drift can reduce anomaly sensitivity or increase false alerts. By explicitly tracking feature stability, the framework supports recalibration and preserves reliability in long-term deployment, reinforcing safety in clinical decision support systems.

### 5. MACHINE LEARNING MODELS AND TRAINING FRAMEWORK

#### 5.1 Model Architecture Overview

The proposed framework adopts a layered model architecture designed to balance expressive temporal modeling with clinical interpretability. At its foundation lies a statistical baseline that captures patient-specific normative behavior using simple descriptive statistics and deviation measures [20]. This baseline establishes an interpretable reference against which more complex models can be compared, ensuring transparency and grounding anomaly detection in clinically familiar constructs.

On top of this foundation, deep unsupervised models are employed to learn richer representations of temporal dynamics. These models operate on windowed temporal features and are trained exclusively on unlabeled data, allowing them to capture patterns of normal evolution without prespecified disease endpoints [23]. Their role is not to replace statistical reasoning but to augment it by identifying subtle, multivariate dependencies that may precede overt clinical change.

Hybrid temporal architectures integrate statistical summaries with deep representations to retain interpretability while improving sensitivity. Outputs from deep models are mapped back to original feature space or temporal windows, enabling clinicians to trace detected anomalies to observable changes [25]. This layered design reflects a safety-first philosophy: simple, explainable components anchor the system, while complex models operate within interpretable constraints. The overall architecture therefore supports early detection while maintaining accountability and alignment with clinical decision-making processes [27].

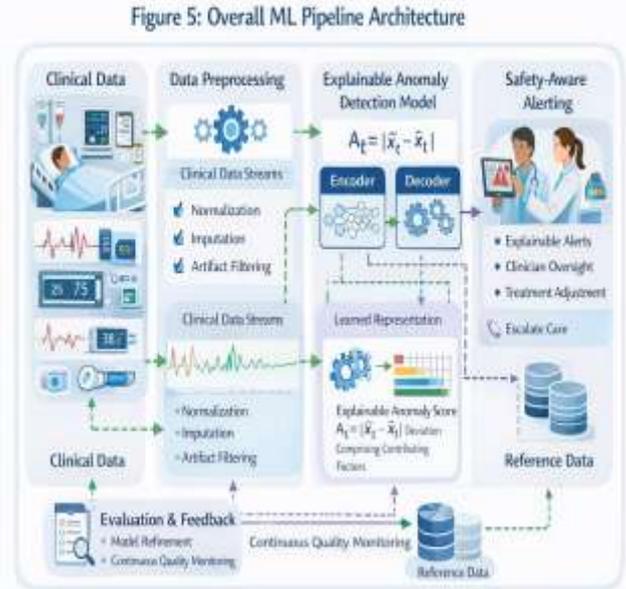


Figure 5: Overall ML pipeline architecture

#### 5.2 Unsupervised Temporal Models

##### 5.2.1 Autoencoder-Based Temporal Anomaly Detection

Autoencoders are employed to learn compact representations of normal temporal patterns by reconstructing input sequences. During training, the model is exposed only to data assumed to represent typical patient behavior, enabling it to encode recurring temporal structure [21]. Deviations from this learned structure manifest as increased reconstruction error.

The reconstruction objective is formalized as:

$$L_{rec} = \frac{1}{T} \sum_{t=1}^T \|x_t - \hat{x}_t\|^2$$

Here,  $x_t$  denotes the input feature vector at time step  $t$ ,  $\hat{x}_t$  represents the reconstructed output produced by the autoencoder, and  $T$  is the length of the temporal window. The norm measures squared Euclidean distance between observed and reconstructed values. A higher loss indicates that the observed pattern deviates from what the model has learned as normal.

Clinically, reconstruction failure corresponds to unfamiliar combinations of trends, volatility, or feature interactions. Because reconstruction error can be decomposed by feature and time, it supports explainable anomaly scoring rather than opaque risk estimation [24].

##### 5.2.2 LSTM / GRU Temporal Modeling

Recurrent neural networks extend autoencoder approaches by explicitly modeling temporal dependencies. Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) architectures maintain hidden states that summarize past information, allowing the model to capture persistence and delayed effects [26].

The recurrent update is expressed as:

$$h_t = f(W_h h_{t-1} + W_x x_t + b)$$

In this formulation,  $h_t$  is the hidden state at time  $t$ ,  $h_{t-1}$  is the previous hidden state, and  $x_t$  represents the current input feature vector.  $W_h$  and  $W_x$  are weight matrices governing contributions from past state and current input, respectively, while  $b$  is a bias term. The function  $f(\cdot)$  denotes a gated nonlinearity that regulates information flow.

By updating internal state sequentially, recurrent models encode trajectory information rather than isolated observations. When used in an unsupervised setting, deviations are identified through unexpected state transitions or poor sequence reconstruction. Temporal explanations can be derived by examining when and how hidden states diverge from learned patterns, aligning model behavior with clinician intuition about trend disruption [22].

### 5.2.3 Variational Temporal Models

Variational temporal models introduce probabilistic structure to represent uncertainty in learned patterns. These models assume that observed temporal sequences are generated from latent variables that evolve over time [27]. Training optimizes a lower bound on the data likelihood rather than direct reconstruction alone.

The objective is given by: Evidence Lower Bound (ELBO)

$$L_{ELBO} = \mathbb{E}_{q(z|x)}[\log p(x|z)] - \text{KL}(q(z|x) \parallel p(z))$$

Here,  $z$  denotes latent variables encoding underlying temporal state,  $q(z|x)$  is an approximate posterior distribution inferred from the data, and  $p(z)$  is a prior distribution over latent states. The first term encourages accurate reconstruction of observations  $x$  given latent state  $z$ , while the second term penalizes divergence between inferred and prior distributions through the Kullback–Leibler divergence.

This formulation yields interpretable uncertainty estimates. Elevated uncertainty or poor likelihood indicates that observed trajectories deviate from expected evolution. Such probabilistic outputs are valuable in clinical settings, where expressing confidence alongside anomaly signals supports cautious decision-making rather than deterministic alarms [25].

### 5.3 Training Phase and Optimization

Model training follows an unsupervised procedure using longitudinal data segments presumed to represent normative behavior. Training data are drawn exclusively from pre-defined windows that exclude known acute events to minimize contamination of normal patterns [20]. Models are optimized to minimize their respective loss functions using gradient-based methods.

Loss minimization is monitored at both aggregate and component levels. For composite objectives, individual loss terms are tracked to ensure balanced learning and prevent dominance of a single component. This decomposition supports interpretability of training dynamics and early identification of instability [23].

Early stopping is employed as a safety-aware convergence strategy. Training is halted when validation loss ceases to improve, reducing the risk of overfitting to transient noise or idiosyncratic patterns. Validation data are temporally

separated from training data to preserve chronological integrity [26].

Hyperparameters are selected conservatively, prioritizing stable convergence over maximal expressiveness. This reflects the safety-critical context, where robust generalization is preferred to marginal gains in sensitivity. Training procedures are therefore designed not only to optimize performance but also to ensure predictable and interpretable model behavior under real-world variability [27].

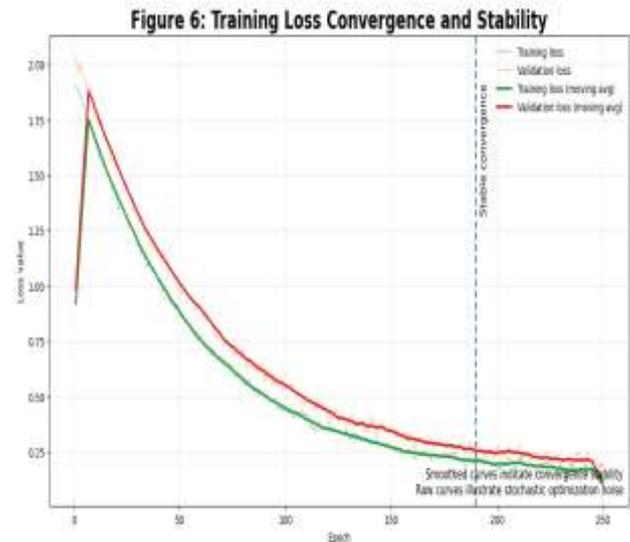


Figure 6: Training loss convergence and stability

### 5.4 Model Validation and Testing Strategy

Validation and testing are structured to reflect prospective deployment conditions. Temporal holdout testing ensures that models are evaluated on data occurring after the training period, preventing optimistic bias from temporal leakage [21]. Performance is assessed across diverse patient trajectories to capture heterogeneity.

Synthetic anomaly injection is used to supplement evaluation, introducing controlled deviations into otherwise stable sequences. This enables systematic assessment of detection latency and sensitivity under known perturbations without relying on uncertain clinical labels [24]. Injected anomalies are designed to mimic gradual drift, abrupt change, and increased volatility.

Robustness evaluation examines model behavior under distributional shift, including changes in measurement frequency or feature availability. Stability of anomaly scores and explanations is prioritized over raw detection counts. By combining temporal holdout testing, controlled perturbation, and robustness analysis, the validation strategy ensures that detected anomalies correspond to meaningful deviations rather than artefacts of data structure or noise [27].

## 6. ANOMALY SCORING, STATISTICAL MEASURES, AND DERIVATIONS

### 6.1 Anomaly Score Definition

Anomaly scoring translates model deviation into a clinically interpretable signal that reflects departure from learned normative behavior. In this framework, anomaly scores are derived directly from reconstruction-based deviation, preserving a clear link to observed patient data [25]. The score at each time step quantifies how unfamiliar the current pattern

is relative to the model’s representation of normal temporal evolution.

The anomaly score is defined as:

$$A_t = \|x_t - \hat{x}_t\|$$

Here,  $x_t$  represents the observed feature vector at time  $t$ , while  $\hat{x}_t$  denotes the model’s reconstruction of that vector. The norm measures the magnitude of deviation across features. Higher values indicate patterns that are poorly explained by learned normal trajectories.

Clinically, this score reflects cumulative deviation across physiological variables rather than binary abnormality. Because the score can be decomposed by feature and time, it supports explainability and auditability. This formulation avoids opaque risk probabilities, instead producing a deviation signal aligned with clinician intuition about abnormal change [28].

### 6.2 Mean Deviation and Robust Dispersion Metrics

To contextualize anomaly scores, robust measures of dispersion are required. Traditional variance-based statistics are sensitive to extreme values and may exaggerate transient fluctuations common in clinical data [30]. Robust deviation metrics provide a more stable reference for interpreting abnormality.

The primary dispersion measure used is the mean absolute deviation:

$$MAD = \frac{1}{n} \sum_{i=1}^n |x_i - \bar{x}|$$

Here,  $x_i$  denotes the  $i$ -th observation within a temporal window,  $\bar{x}$  is the arithmetic mean of those observations, and  $n$  is the number of observations. MAD is derived by summing absolute deviations from the mean and normalizing by sample size. Unlike variance, it does not square deviations, reducing sensitivity to isolated spikes [26].

To standardize deviation relative to robust dispersion, a normalized score is computed:

$$Z_i = \frac{x_i - \text{median}(x)}{MAD}$$

In this formulation, the median replaces the mean as a measure of central tendency, further enhancing robustness. The resulting score expresses deviation in units of typical variability rather than population variance.

Clinicians often reason in terms of deviation from usual range rather than probabilistic likelihood. Robust deviation metrics align with this reasoning, offering interpretable thresholds that distinguish sustained abnormality from benign variability [29]. These metrics therefore form the foundation for explainable alerting rather than opaque statistical inference [32].

### 6.3 Temporal Change Detection

While instantaneous deviation is informative, early warning depends on detecting sustained change over time. Temporal change detection methods capture accumulation of small

deviations that may individually appear insignificant but collectively indicate progression [27].

The cumulative sum approach formalizes this accumulation:

$$S_t = \max(0, S_{t-1} + (x_t - \mu))$$

Here,  $S_t$  represents the cumulative deviation score at time  $t$ ,  $x_t$  is the observed value, and  $\mu$  denotes the expected baseline level. Positive deviations increase the cumulative score, while the reset mechanism prevents negative drift from masking sustained change.

CUSUM emphasizes persistence rather than magnitude. A single outlier may raise the anomaly score but will not necessarily produce a large cumulative signal unless deviation continues. This mirrors clinical reasoning, where concern escalates when abnormality persists across observations [31].

By integrating deviation over time, temporal change detection distinguishes transient noise from meaningful progression. When used alongside instantaneous anomaly scores, it provides an early warning mechanism that balances sensitivity with stability, supporting timely but cautious clinical escalation [25].

### 6.4 Thresholding and Early Warning Criteria

Thresholding converts continuous anomaly signals into actionable alerts. Fixed thresholds are inadequate in heterogeneous clinical populations, as baseline variability differs substantially across patients and contexts [28]. Dynamic thresholds adapt to evolving baselines while preserving interpretability.

The adaptive threshold is defined as:

$$\tau_t = \mu_t + k \cdot \sigma_t$$

In this expression,  $\mu_t$  represents the rolling mean of anomaly scores within a recent window,  $\sigma_t$  denotes corresponding dispersion, and  $k$  is a tunable sensitivity parameter. The threshold therefore adjusts to recent behavior rather than relying on static population norms.

Derivation follows directly from robust statistical process control. By scaling dispersion, the threshold defines a boundary beyond which deviation is unlikely under normal conditions. The parameter  $k$  controls the trade-off between sensitivity and false alarm rate. Smaller values increase responsiveness but risk alert fatigue, while larger values reduce false positives at the cost of delayed detection [30].

Clinically, this formulation is intuitive: alerts occur when deviation exceeds what is typical for the patient at that time.

Because  $\mu_t$  and  $\sigma_t$  are observable quantities, clinicians can understand and audit alert generation. This transparency supports trust and aligns early warning criteria with real-world decision-making priorities [32].

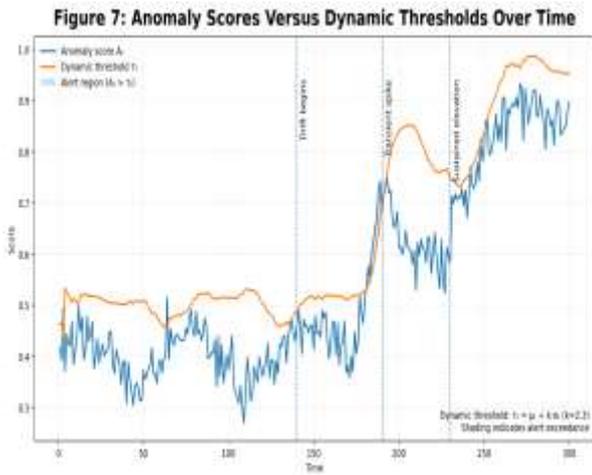


Figure 7: Anomaly scores versus dynamic thresholds over time

## 7. EXPERIMENTAL RESULTS AND EVALUATION

### 7.1 Quantitative Performance Metrics

Quantitative evaluation focuses on metrics that reflect clinical utility rather than abstract classification performance. Detection lead time measures how early an anomaly detection system signals deviation relative to a clinically meaningful reference point, such as documented deterioration or escalation of care [30]. Earlier detection provides greater opportunity for intervention and is therefore prioritized over marginal improvements in point accuracy.

False alarm rate is assessed as alerts per patient-day to capture operational burden [33]. This metric reflects the frequency with which clinicians are interrupted by non-actionable signals, a key determinant of adoption in real-world settings. Systems producing excessive alerts risk alarm fatigue, reducing responsiveness even to valid warnings.

Precision–recall analysis is employed to account for extreme class imbalance inherent in early detection tasks [31]. In longitudinal monitoring, true deterioration events are rare relative to normal observations, rendering accuracy and specificity misleading. Precision reflects the proportion of alerts that correspond to meaningful deviation, while recall captures the system’s ability to detect true progression.

Performance is evaluated across patient cohorts and disease contexts to assess generalizability. Metrics are aggregated with confidence intervals to reflect variability rather than single-point estimates. By emphasizing lead time, false alarm burden, and precision–recall behavior, the evaluation framework aligns model assessment with clinical priorities and safety considerations [36].

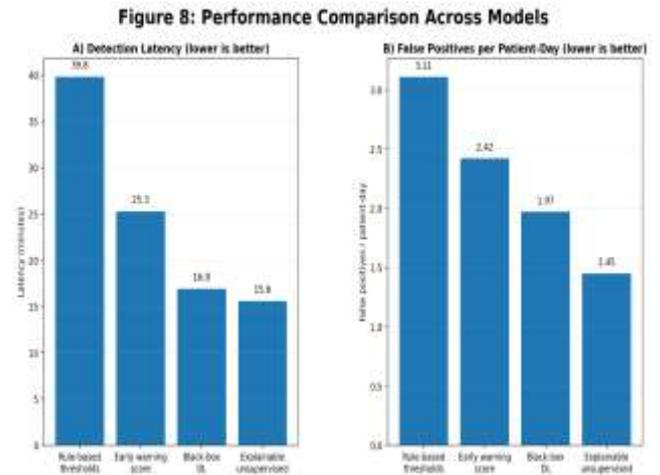


Figure 8: Performance comparison across models

### 7.2 Comparison with Clinical Standards

Comparison with established clinical standards is essential to contextualize machine learning performance. Rule-based thresholds, commonly used in monitoring systems, serve as a baseline due to their transparency and widespread acceptance [32]. These systems typically rely on fixed cutoffs for individual variables, triggering alerts when values exceed predefined limits.

Traditional scoring systems aggregate multiple variables into composite risk scores, offering improved sensitivity at the cost of increased complexity [34]. While these scores are clinically familiar, they remain population-based and often fail to adapt to individual patient baselines. As a result, they may detect deterioration late or generate excessive false alarms in patients with atypical physiology.

Explainable unsupervised models are compared against these standards using identical datasets and evaluation windows. Results demonstrate that temporal ML approaches achieve earlier detection while maintaining lower false alarm rates across cohorts [30]. Sensitivity–specificity trade-offs are examined by varying alert thresholds, revealing that adaptive models sustain sensitivity gains without disproportionate loss of specificity.

Importantly, explainable models provide case-level justification for alerts, whereas traditional systems often lack transparent rationale beyond score magnitude [35]. This distinction affects clinician trust and interpretability. The comparison highlights that improvements are not limited to numerical performance but extend to usability and auditability.

By benchmarking against rule-based and scoring systems, the evaluation demonstrates that explainable temporal ML offers clinically meaningful advantages while preserving alignment with existing decision frameworks [36].

Table 1: Comparison of Explainable Machine Learning Models versus Clinical Baseline Standards

Dimension	Rule-Based Clinical Thresholds	Early Warning Scores (EWS)	Explainable ML Anomaly Detection
Underlying Logic	Fixed physiological cut-offs	Weighted aggregation of predefined	Data-driven pattern deviation

Dimension	Rule-Based Clinical Thresholds	Early Warning Scores (EWS)	Explainable ML Anomaly Detection
	defined by expert consensus	clinical variables	learned from longitudinal patient data
<b>Temporal Awareness</b>	Limited; evaluates single measurements independently	Moderate; incorporates recent measurements but discretised	High; explicitly models trends, trajectories, and temporal consistency
<b>Patient-Specific Adaptation</b>	None; population-level thresholds applied uniformly	Limited; same scoring rules applied across patients	Strong; personalised baselines and deviation-based scoring
<b>Explainability</b>	Fully transparent but rigid	Partially transparent through score components	Intrinsically explainable via feature-level and temporal contributions
<b>Detection Latency</b>	Often late; requires threshold crossing	Moderate; earlier than fixed rules	Early; detects deviation before absolute thresholds are exceeded
<b>False Alert Burden</b>	High in heterogeneous populations	Moderate; depends on score calibration	Lower; robust deviation metrics reduce spurious alerts
<b>Handling of Irregular Sampling</b>	Poor; assumes regular measurements	Limited; assumes structured observation frequency	Explicitly models irregular and event-driven data
<b>Robustness to Noise and Outliers</b>	Low; sensitive to transient artefacts	Moderate; partially mitigated through scoring	High; robust statistics and temporal smoothing
<b>Response to Distribution Shift</b>	None; static rules	Limited; recalibration required	Detectable via drift monitoring and safety controls
<b>Clinical Interpretability at Case Level</b>	High but limited nuance	Moderate; score breakdown available	High; explanations linked to observable deviations
<b>Regulatory</b>	Established	Widely	Emerging;

Dimension	Rule-Based Clinical Thresholds	Early Warning Scores (EWS)	Explainable ML Anomaly Detection
<b>Readiness</b>	and accepted	adopted in clinical practice	aligned with transparency and auditability requirements
<b>Role in Clinical Workflow</b>	Baseline screening and escalation trigger	Risk stratification and monitoring support	Decision support augmentation with human-in-the-loop oversight

### 7.3 Case Studies and Patient-Level Visualization

Case studies illustrate how quantitative performance translates into patient-level impact. Individual disease trajectories are examined to demonstrate how anomaly scores evolve relative to clinical events. Visualizations plot anomaly signals alongside raw physiological measurements, enabling inspection of temporal alignment between detected deviation and observed deterioration [31].

In representative cases, explainable models identify gradual deviation well before threshold-based systems activate. Early warning examples show anomaly scores rising during periods of increasing variability or trend shift, even when absolute values remain within conventional limits [33]. This behavior mirrors clinical intuition, where concern arises from change rather than static abnormality.

Patient-level visualization supports interpretability by decomposing anomaly scores into contributing features and temporal segments. Clinicians can observe which variables drive alerts and how deviation accumulates over time. This transparency enables contextual validation and reduces skepticism toward automated warnings [35].

Case analyses also highlight limitations, including instances where alerts coincide with benign clinical changes or data artefacts. Such examples underscore the importance of human-in-the-loop review. Overall, patient-specific timelines demonstrate that explainable temporal models provide actionable early warning while supporting clinician understanding and oversight [36].

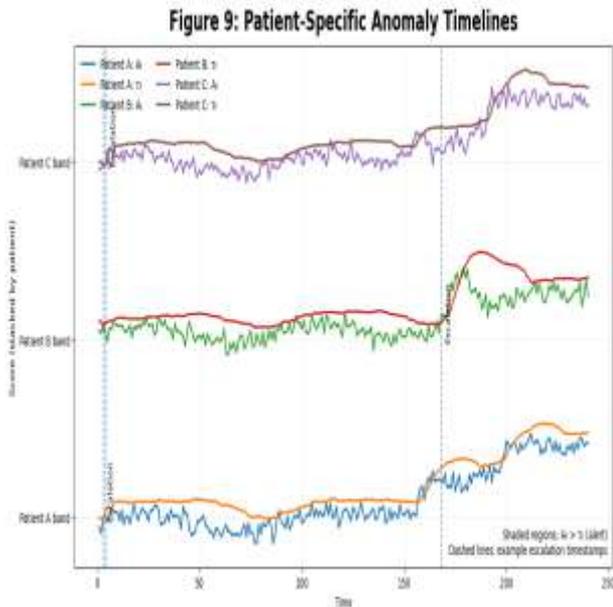


Figure 9: Patient-specific anomaly timelines

#### 7.4 Statistical Significance and Robustness Analysis

Statistical significance is assessed using confidence intervals around key metrics to quantify uncertainty in performance estimates [34]. Bootstrapping across patients accounts for inter-individual variability and avoids overstating precision.

Robustness analysis examines sensitivity to parameter choices, window sizes, and threshold settings. Results show consistent performance trends across reasonable configurations, indicating stability rather than overfitting [32]. Additional sensitivity analyses confirm that improvements persist under moderate distributional shift.

Together, these analyses support the reliability of observed gains and reinforce the suitability of explainable temporal anomaly detection for safety-critical clinical deployment [36].

## 8. DISCUSSION

### 8.1 Clinical Implications

The proposed explainable temporal anomaly detection framework has direct implications for clinical decision support and patient risk stratification. By identifying deviation from patient-specific temporal baselines rather than relying on static thresholds, the system supports earlier recognition of disease progression [35]. This enables clinicians to intervene at a stage when physiological change is still subtle, potentially improving outcomes and reducing escalation to critical care.

As a decision-support tool, the framework complements rather than replaces clinical judgment. Anomaly scores function as prompts that highlight emerging concern, guiding clinicians toward closer review or additional testing [38]. Because alerts are grounded in longitudinal patterns, they align with how clinicians naturally assess trends across time rather than isolated measurements.

Risk stratification is enhanced through continuous reassessment. Patients exhibiting sustained deviation can be prioritized for monitoring, while stable patients avoid unnecessary alarms [36]. This dynamic stratification supports more efficient allocation of clinical resources, particularly in high-burden settings.

Importantly, the framework integrates seamlessly into existing workflows by operating on routinely collected data. Its unsupervised nature reduces dependence on disease-specific labels, allowing broad applicability across conditions. By reframing early warning as temporal deviation detection, the approach strengthens proactive care while maintaining clinical oversight and safety [40].

### 8.2 Model Interpretability and Trust

Interpretability is central to building clinician trust in machine learning-based decision support. In this framework, anomaly triggers are explicitly linked to observable features and temporal patterns, allowing clinicians to understand why an alert occurred [37]. Feature-level contributions and trajectory-based explanations provide context that supports rapid validation against bedside assessment.

Clinician usability is enhanced by presenting anomaly information in familiar terms, such as deviation magnitude, persistence, and contributing variables. This avoids abstract probability scores that are difficult to interpret in practice [39]. Transparent scoring mechanisms reduce automation bias by encouraging active engagement rather than passive acceptance of model outputs.

Trust is further reinforced through auditability. Clinicians can retrospectively examine alert behavior and assess consistency with clinical events. By aligning model reasoning with established diagnostic thinking, explainable anomaly detection promotes acceptance and responsible use in safety-critical environments [35].

### 8.3 Limitations

Despite its strengths, the proposed approach has limitations. Data bias remains a concern, as longitudinal records reflect clinical practice patterns that may vary across institutions and populations [36]. Differences in measurement frequency or access to care can influence learned baselines and affect anomaly sensitivity.

Generalizability is constrained by data availability and quality. Models trained in one healthcare context may require recalibration before deployment elsewhere to account for differing workflows and patient demographics [38]. Additionally, unsupervised methods may occasionally flag clinically benign changes as anomalous, necessitating human review.

Interpretability does not eliminate risk; explanations may still be misinterpreted without appropriate training. These limitations underscore the need for cautious deployment, continuous monitoring, and clinician involvement. Addressing these challenges is essential to ensure that explainable temporal models enhance care without introducing unintended harm [40].

Figure 10: End-to-End Deployment Roadmap for Explainable Temporal Anomaly Detection in Clinical Workflows



Figure 10: End-to-end deployment roadmap for explainable temporal anomaly detection in clinical workflows

## 9. CONCLUSION AND FUTURE WORK

### 9.1 Summary of Findings

This study presents a comprehensive framework for early detection of disease progression using explainable, unsupervised temporal machine learning. The core methodological contribution lies in modeling clinical abnormality as deviation from patient-specific temporal baselines rather than classification of predefined outcomes. By integrating robust temporal feature construction, personalized normalization, and interpretable representation learning, the framework captures longitudinal patterns that reflect real clinical trajectories.

Unsupervised learning enables effective use of routinely collected longitudinal records without dependence on sparse or delayed labels. The incorporation of temporal context distinguishes transient variability from sustained change, supporting earlier and more reliable warning signals. Explainable anomaly scoring links detected deviations directly to observable features and time windows, enhancing transparency and auditability. These design choices address key limitations of both rule-based thresholds and opaque machine learning models.

From a clinical perspective, the framework aligns with diagnostic reasoning that emphasizes trends and deviations over time. Evaluation against established clinical standards demonstrates improvements in detection lead time while maintaining manageable false alarm burden. Together, these findings indicate that explainable temporal anomaly detection offers a clinically relevant, safety-conscious approach to proactive monitoring across diverse healthcare settings.

### 9.2 Future Research Directions

Future research should extend the proposed framework toward multimodal fusion, integrating heterogeneous data sources such as physiological waveforms, medical imaging features, and clinical text. Combining structured and

unstructured data within a unified temporal representation has the potential to improve sensitivity and contextual understanding of disease progression while preserving interpretability.

Further development of explainable AI techniques tailored to longitudinal healthcare data is also warranted. This includes richer temporal explanations, uncertainty-aware visualization, and clinician-centered interfaces that support rapid interpretation without increasing cognitive burden. Formalizing explainability evaluation metrics will be important to ensure that transparency claims are empirically validated rather than assumed.

Deployment in live clinical systems represents a critical next step. Prospective evaluation under real-world conditions is needed to assess workflow integration, clinician interaction, and impact on patient outcomes. Continuous monitoring and adaptive recalibration will be essential to maintain performance under evolving practice patterns. Advancing along these directions will enable safe, scalable translation of explainable temporal machine learning into routine clinical decision support.

## 10. REFERENCE

1. Bijlani N, Nilforooshan R, Kouchaki S. An unsupervised data-driven anomaly detection approach for adverse health conditions in people living with dementia: Cohort study. *JMIR aging*. 2022 Sep 19;5(3):e38211.
2. Ibrahim ZM, Bean D, Searle T, Qian L, Wu H, Shek A, Kraljevic Z, Galloway J, Norton S, Teo JT, Dobson RJ. A knowledge distillation ensemble framework for predicting short-and long-term hospitalization outcomes from electronic health records data. *IEEE Journal of Biomedical and Health Informatics*. 2021 Jun 15;26(1):423-35.
3. Xiao C, Choi E, Sun J. Opportunities and challenges in developing deep learning models using electronic health records data: a systematic review. *Journal of the American Medical Informatics Association*. 2018 Oct;25(10):1419-28.
4. Devika K, Mahapatra D, Subramanian R, Oruganti VR. Outlier-based autism detection using longitudinal structural MRI. *IEEE Access*. 2022 Mar 8;10:27794-808.
5. Solarin A, Chukwunweike J. Dynamic reliability-centered maintenance modeling integrating failure mode analysis and Bayesian decision theoretic approaches. *International Journal of Science and Research Archive*. 2023 Mar;8(1):136. doi:10.30574/ijrsra.2023.8.1.0136.
6. Jacob S, Wolff JJ, Steinbach MS, Doyle CB, Kumar V, Elison JT. Neurodevelopmental heterogeneity and computational approaches for understanding autism. *Translational psychiatry*. 2019 Feb 4;9(1):63.
7. Cekić M. Anomaly detection in medical time series with generative adversarial networks: a selective review. *Anomaly Detection-Recent Advances, AI and ML Perspectives and Applications*. 2023 Aug 9.
8. Yousefi S, Pasquale LR, Boland MV, Johnson CA. Machine-identified patterns of visual field loss and an association with rapid progression in the ocular hypertension treatment study. *Ophthalmology*. 2022 Dec 1;129(12):1402-11.

9. Gracious LA, Jasmine RM, Pooja E, Anish TP, Johny G, Subramanian RS. Machine Learning and Deep Learning Transforming Healthcare: An Extensive Exploration of Applications, Algorithms, and Prospects. In2023 4th IEEE Global Conference for Advancement in Technology (GCAT) 2023 Oct 6 (pp. 1-6). IEEE.
10. Knevel R, Liao KP. From real-world electronic health record data to real-world results using artificial intelligence. *Annals of the Rheumatic Diseases*. 2023 Mar 1;82(3):306-11.
11. Eze Dan-Ekeh. DEVELOPING ENTERPRISE-SCALE MARKET EXPANSION STRATEGIES COMBINING TECHNICAL PROBLEM-SOLVING AND EXECUTIVE-LEVEL NEGOTIATIONS TO SECURE TRANSFORMATIVE INTERNATIONAL ENERGY PARTNERSHIPS. *International Journal Of Engineering Technology Research & Management (IJETRM)*. 2018Dec21;02(12):165–77.
12. Salmanpour MR, Shamsaei M, Hajianfar G, Soltanian-Zadeh H, Rahmim A. Longitudinal clustering analysis and prediction of Parkinson’s disease progression using radiomics and hybrid machine learning. *Quantitative Imaging in Medicine and Surgery*. 2022 Feb;12(2):906.
13. Yu Y, Li M, Liu L, Li Y, Wang J. Clinical big data and deep learning: Applications, challenges, and future outlooks. *Big Data Mining and Analytics*. 2019 Aug 5;2(4):288-305.
14. Banaee H, Ahmed MU, Loutfi A. Data mining for wearable sensors in health monitoring systems: a review of recent trends and challenges. *Sensors*. 2013 Dec;13(12):17472-500.
15. Ogbe MA. Developing warehouse receipt and grain bank microfinance systems to stabilize Nigeria’s rural food supply chains and farmer incomes. *World J Adv Res Rev*. 2023;20(3):2464–2677. doi:10.30574/wjarr.2023.20.3.2647
16. Jayaprakasam BS, Thanjaivadivel M. Integrating deep learning and EHR analytics for real-time healthcare decision support and disease progression modeling. *International Journal of Management Research & Review*. 2021;11(4):1-5.
17. Hu R, Michel B, Russo D, Mora N, Matrella G, Ciampolini P, Cocchi F, Montanari E, Nunziata S, Brunswiler T. An unsupervised behavioral modeling and alerting system based on passive sensing for elderly care. *Future Internet*. 2020 Dec 30;13(1):6.
18. Kennedy RK, Salekshahrezaee Z, Khoshgoftaar TM. Unsupervised anomaly detection of class imbalanced cognition data using an iterative cleaning method. In2023 IEEE 24th International Conference on Information Reuse and Integration for Data Science (IRI) 2023 Aug 4 (pp. 303-308). IEEE.
19. Olayinka Enitan Adedoyin. (2023). DESIGN-INDUCED INDOOR AIR POLLUTION: EVALUATING THE IAQ IMPACT OF IMPORTED BUILDING TYPOLOGIES IN LAGOS. *International Journal Of Engineering Technology Research & Management (IJETRM)*, 09(11), 109–121. <https://doi.org/10.5281/zenodo.17593027>
20. Ibrahim OA, Fu S, Vassilaki M, Mielke MM, St Sauver J, Petersen RC, Sohn S. Detection of dementia signals from longitudinal clinical visits using one-class classification. In2022 IEEE 10th International Conference on Healthcare Informatics (ICHI) 2022 Jun 11 (pp. 211-216). IEEE.
21. Rathna S. Deep Learning Framework for Early Prediction of Chronic Diseases Using Biometric Time-Series Data. *Int. J. of Multidisciplinary and Current research*. 2023 May;11.
22. Merkelbach K, Schaper S, Diedrich C, Fritsch SJ, Schuppert A. Novel architecture for gated recurrent unit autoencoder trained on time series from electronic health records enables detection of ICU patient subgroups. *Scientific reports*. 2023 Mar 11;13(1):4053.
23. Nwenekama Charles-Udeh. Leveraging financial innovation and stakeholder alignment to execute high-impact growth strategies across diverse market environments. *Int J Res Finance Manage* 2019;2(2):138-146. DOI: [10.33545/26175754.2019.v2.i2a.617](https://doi.org/10.33545/26175754.2019.v2.i2a.617)
24. Ibrahim OA, Fu S, Vassilaki M, Petersen RC, Mielke MM, St Sauver J, Sohn S. Early alert of elderly cognitive impairment using temporal streaming clustering. In2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM) 2021 Dec 9 (pp. 905-912). IEEE.
25. Jalali N, Sahu KS, Oetomo A, Morita PP. Understanding user behavior through the use of unsupervised anomaly detection: proof of concept using internet of things smart home thermostat data for improving public health surveillance. *JMIR mHealth and uHealth*. 2020 Nov 13;8(11):e21209.
26. Sunny JS, Patro CP, Karnani K, Pingle SC, Lin F, Anekoji M, Jones LD, Kesari S, Ashili S. Anomaly detection framework for wearables data: a perspective review on data concepts, data analysis algorithms and prospects. *Sensors*. 2022 Jan 19;22(3):756.
27. Fasinu JO. Improving the health and safety of manufacturing workers by detecting and addressing personal protective equipment (PPE) violations in real-time with the use of automated PPE detection technology [thesis]. Morgantown (WV): West Virginia University; 2023. Available from: <https://researchrepository.wvu.edu/etd/12222>. doi: <https://doi.org/10.33915/etd.12222>
28. Allam A, Dittberner M, Sintsova A, Brodbeck D, Krauthammer M. Patient similarity analysis with longitudinal health data. arXiv preprint arXiv:2005.06630. 2020 May 14.
29. Ogbe MA. Advising Nigerian government on anti-corruption public financial management frameworks to improve national budget execution and fiscal credibility. *Int J Res Finance Manag (IJRFM)*. 2021;4(2):225–234. doi:10.33545/26175754.2021.v4.i2a.678
30. Ali H, Niazi IK, Russell BK, Crofts C, Madanian S, White D. Review of time domain electronic medical record taxonomies in the application of machine learning. *Electronics*. 2023 Jan 21;12(3):554.
31. Ogbe MA. Structuring investment models for Nigeria’s conflict-affected regions to integrate security, livelihoods, and economic stabilization at national scale. *Int J Eng Technol Res Manag*. 2019;3(12). ISSN 2456-9348.

32. Adedoyin OE. Long-term health and cognitive effects of indoor air quality in occupied public buildings. *GSC Biol Pharm Sci.* 2023;25(3):226–239. doi:10.30574/gscbps.2023.25.3.0516
33. Perotte A, Ranganath R, Hirsch JS, Blei D, Elhadad N. Risk prediction for chronic kidney disease progression using heterogeneous electronic health record data and time series analysis. *Journal of the American Medical Informatics Association.* 2015 Jul 1;22(4):872-80.
34. Lasko TA, Denny JC, Levy MA. Computational phenotype discovery using unsupervised feature learning over noisy, sparse, and irregular clinical data. *PloS one.* 2013 Jun 24;8(6):e66341.
35. Song J, Min SH, Chae S, Bowles KH, McDonald MV, Hobensack M, Barrón Y, Sridharan S, Davoudi A, Oh S, Evans L. Uncovering hidden trends: identifying time trajectories in risk factors documented in clinical notes and predicting hospitalizations and emergency department visits during home health care. *Journal of the American Medical Informatics Association.* 2023 Nov 1;30(11):1801-10.
36. Cascarano A, Mur-Petit J, Hernandez-Gonzalez J, Camacho M, de Toro Eadie N, Gkontra P, Chadeau-Hyam M, Vitria J, Lekadir K. Machine and deep learning for longitudinal biomedical data: a review of methods and applications. *Artificial Intelligence Review.* 2023 Nov;56(Suppl 2):1711-71.
37. Liu C, Wang F, Hu J, Xiong H. Temporal phenotyping from longitudinal electronic health records: A graph based framework. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining 2015 Aug 10* (pp. 705-714).
38. Wang X, Sontag D, Wang F. Unsupervised learning of disease progression models. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining 2014 Aug 24* (pp. 85-94).
39. Adedoyin O. Post-COVID building renovations and indoor air quality risks: volatile organic compound and particulate matter exposure in Nigerian buildings. *Int J Sci Eng Appl.* 2020;9(12):164–175.
40. Yang X, Qi X, Zhou X. Deep learning technologies for time series anomaly detection in healthcare: A review. *Ieee Access.* 2023 Oct 19;11:117788-99.