



# Adaptive Clinical Networks for Real-Time, Risk-Aware Care

**Bhasker Katta**

Independent Researcher

## Abstract

Clinical Interaction Networks represent workflows in real-time clinical settings, enabling the generation of up-to-date metrics and risk estimations that monitor the quality and safety of care and provide decision-support capabilities. These networks act as a clinical control tower, providing notifications and recommendations to improve the effectiveness, efficiency, and safety of healthcare delivery. Without the integration of data governance and model feedback loops into their development, the networks risk becoming disconnected from the changing realities in the health environment for which they are designed.


In addressing these considerations, a reliable architecture for ensuring data quality, adapting automatically to changes in data streams and clinical processes, and creating mechanisms for timely clinician feedback and validation has been presented. Self-evolving Clinical Interaction Networks enable data-fusion services that process streaming health data, machine-learning approaches that utilise online or continual learning to produce adaptable network structures, and quality-control procedures capable of integrating clinician judgement into the models, thus enhancing their clinical relevance.

**Keywords :** Clinical interaction networks, care optimization, real-time risk-aware analytics, data stream monitoring, health data integration.

## 1. Introduction

The Connected Health paradigm emphasizes high-quality care at lower costs through real-time monitoring and predictive treatment. Clinical Interaction Networks (CINs) aggregate the interactions of patients, providers, spaces, devices, and other contextual factors to create temporal graphs of real-time care delivered in a transparent and interpretable manner. Optimal care and patient risk under these conditions are often addressed by Quality Indicators, Anomaly Detection or Risk Scoring Methods. However, established CINs may not provide accurate analyses or support decision-making for real-time care, due to the A) complex multi-node nature of CINs, B) demanding requirements on analytics frequency and latencies, and C) validation typically driven by CAUSAL or EXPLAINABLE MACHINE LEARNING.

To fill this gap, a Self-Evolving Clinical Interaction Network (SECIN) is proposed, which provides quality, safety, anomaly detection and risk scoring analytics for real-time care and adapts to the LOCAL CARE OPTIMIZATION AND RISK-AWARE ANALYTICS needs. SECINs support multiple feedback mechanisms to adapt to local care dynamics, detecting changes in SELECTED PROCESS STAGE EFFECTIVENESS OR RISK MODEL DRIFT. A recent SECIN deployment in a Hospital Ward illustrates the approach, including rapid construction, understandability of the streaming analytics, and mapping of error sources. An end-to-end architecture for real-time monitoring and risk scoring of health conditions, medical procedures, and care quality fills a gap in clinical analytics. The methods are assessed within a ward environment and used for predictive monitoring of care quality.



**Fig 1: Self-Evolving Clinical Interaction Networks**

## 1.1. Background and Significance

Clinical interaction networks consist of nodes representing clinical actors, activities, and patients, temporally connected by event logs that detail care processes. Real-time care optimization continuously monitors workflow safety and quality across several dimensions and alerts clinicians to possible quality and safety issues. Seamless integration of these two concepts enables the creation of interactive data streams that provide support for operational decision-making.

Although many real-time analytics are deployed in clinical environments, they are seldom integrated into clinical workflows and decision-making processes. A self-evolving clinical interaction network is proposed that adapts itself based on clinician feedback and continuous analysis of interaction logs through lightweight mechanisms. It supports care quality and safety through alerts, health-related risk scores, and a continuum of care indicator. A pilot deployment in a hospital ward surrounding internal medicine demonstrated the feasibility of the proposed approach, and

lessons learned during this deployment paved the way for future developments and applications.

**Table 1. Core Components of Self-Evolving Clinical Interaction Networks (SECIN)**

| Component                            | Description                                                                                  | Function in Healthcare                           |
|--------------------------------------|----------------------------------------------------------------------------------------------|--------------------------------------------------|
| Clinical Interaction Networks (CINs) | Temporal graphs representing interactions among patients, clinicians, devices, and workflows | Enable real-time monitoring and analytics        |
| Self-Evolving Mechanism              | Adaptive learning process using continual feedback and online learning                       | Handles model drift and workflow evolution       |
| Risk-Aware Analytics                 | Predictive risk scoring and anomaly detection models                                         | Supports early intervention and patient safety   |
| Data Fusion Layer                    | Integration of EHRs, sensors, logs, and external systems                                     | Creates unified healthcare intelligence          |
| Feedback Loops                       | Clinician validation and governance mechanisms                                               | Improves trust, explainability, and adaptability |
| Decision Support Engine              | Generates alerts, recommendations, and workflow insights                                     | Enhances operational and clinical decisions      |

## 1.2. Research design

The study addresses the approach and aims of the present investigation. The overarching goal is to equip clinical environments with self-evolving clinical interaction networks capable of providing actionable feedback aimed at improving the quality, efficiency, and safety of care. The analysis



encompasses real-time data streams, risk models of clinical decisions and actions, and real-time monitoring of care quality and safety. Feedback from clinicians forms a crucial component.

The design is an empirical investigation of real-time monitoring methods with self-evolving networks deployed in a hospital ward environment. Multiple data sources provide direct feeds from the electronic health record system, medical sensors, and clinical logbooks.

#### Equation 1: Risk Score Function for Patient Monitoring

The paper discusses continuous patient risk scoring for adverse outcomes. A generalized clinical risk model can be expressed as:

$$R_i(t) = \sum_{j=1}^n w_j \cdot x_{ij}(t)$$

Where:


- $R_i(t)$ = risk score of patient  $i$  at time  $t$
- $x_{ij}(t)$ = clinical feature  $j$ (vitals, lab values, symptoms)
- $w_j$ = learned importance weight of feature  $j$

### 1.3. Scope and Objective

The study formally applies Self-Evolving Clinical Interaction Networks (SECIN) methodologies to optimize care in two high-stakes clinical settings. A ward-based pilot project, aimed at testing the approach, explores its formal and informal validation mechanisms and makes explicit the ethical, legal, and social implications of such networks. The second SECIN application investigates the emergency department of a major teaching hospital, addressing the CEDC model of queue management, the role of senior clinicians in the triage of CEDC patients, and their putative impact on patient outcomes. The third evaluation of a primary

care SECIN examines chronic disease management, continuity of care, and health service utilization costs.

Performance metrics assess system quality, but the open-loop nature of traditional SECIN also renders them suitable for risk-aware monitoring. Anomaly detection and risk scoring, primary contributions of a recent publication, are now extended to continuously deployed SECIN. Indices draw on machine learning methods previously validated in clinical studies or, when these are lacking, other regulated sectors. Assessment thresholds are calibrated to avoid both false alarms and missed opportunities for intervention. The combination of these innovations into a coherent framework seeks to render SECIN—hitherto academic proposals—fit for serious use by the clinical community.



## 2. Theoretical foundations of clinical interaction networks

Clinical workflows can be represented as networks capturing the main components of care delivery—patients, healthcare providers, clinical services, related resources, and time ordering—collaboratively interacting with each other and the health information systems supporting them. Patient treatment records, sensor logs, and audit trails from the supporting systems continuously generate the data needed to instantiate care delivery networks as temporal graphs, allowing real-time monitoring and analysis of effectiveness,





efficiency, safety, and patient experience. User-defined thresholds on these metrics enable timely detection of problematic situations and the exploration of possible solutions, such as risk scoring or danger warning. These self-evolving clinical interaction networks are thus suitable for real-time care process optimization, with the potential to improve the patient experience when combined with risk-aware analytics.

Self-evolving clinical interaction networks have recently been reported and evaluated in a hospital ward context, where the feature-rich nature of the data enabled a wealth of relevant quality and safety metrics. Here, the focus is on their deployment in other clinical environments. interaction networks for risk-aware analytics that predicts floods and informs new developments in flow-forecasting models. Such analytics require the integration of external data sources with an attention on produce forecasting across space and time. Clinical treatment outcomes can also be positively impacted by the timely and seamless detection of potential issues for patients and healthcare staff, together with decision support and forthcoming intervention guidance. Day-to-day clinical operations, when modelled as a service network, allow scanning alternative routes to explore the bene- ficial effects on response time and patient outcome.

**Table 2. Data Sources Used in SECIN Architecture**

| Data Source                     | Type of Data                               | Purpose                                        |
|---------------------------------|--------------------------------------------|------------------------------------------------|
| Electronic Health Records (EHR) | Clinical notes, lab results, prescriptions | Patient monitoring and workflow reconstruction |
| Medical Sensors                 | Vital signs, physiological streams         | Continuous patient monitoring                  |
| Clinical Logbooks               | Staff actions and workflow events          | Process mining and care analysis               |

| Data Source            | Type of Data                           | Purpose                       |
|------------------------|----------------------------------------|-------------------------------|
| Administrative Systems | Scheduling, admissions, discharge data | Resource optimization         |
| External Data Streams  | Disaster alerts, environmental factors | Context-aware risk prediction |
| Historical Databases   | Archived patient records               | Model training and validation |


## 2.1. Network representation of clinical workflows

A clinical workflow can be grasped as the series of steps and procedures followed by caregivers in managing patients' health conditions. A network-based representation allows explicitly modelling all the actors involved, including patients and the actions they take, as well as their states, needs, and requirements. For each patient, the actors are represented as workflow nodes, and the communications required for the patient's safety and well-being are represented as directed edges, that is, arcs connecting pairs of nodes. The edges are labelled according to the nature of the duties being performed (that is, examinations, prescriptions, therapeutic procedures) and are either activated by the respective reasoners or triggered by other nodes on the network. The resulting multipartite network is either static or dynamic, depending on the length of the examination.

Three main categories of elements are represented: (1) caregiver-patients pairs, (2) clinical sub-processes, and (3) clinical complications within the examination of a patient. Caregiver-patient pairs represent patients within the examination period and all caregiver roles performing actions for the patient. Each node represents a role of nursing staff (or nursing teams) acting for the patient during the examination. Edges may also represent the state of a patient and a nurse's need of drug or laboratory examination results when attending the patient. Secondary and parallel complications of a patient

during a clinical process are also integrated in the network to avoid unnecessary major disruption of the workflow.

The representational framework readily integrates heterogeneous clinical data, capturing both continuous time-series streams (such as blood pressure) and discrete events (such as lab-test results) in a temporal graph format. An EHR acts as a hub for data fusion. Continuously streaming sensor data support process anomaly detection, risk scoring, and other real-time analytics, while historical data are used for longer temporal horizons. Risk models are formulated as landmarks that augment the patient-care continuum with risk-aware decision support for clinicians, enabling timely adjustments of the standard care path in response to rapidly evolving clinical situations.



**Fig 2: Simulation and Improvement of Patients' Workflow**

## 2.2. Real-time data streams in healthcare environments

Real-time data streams arrive from a variety of sources. The most prominent, and in many settings the most data-rich, source is the electronic health record (EHR). The EHR stores temporally organized clinical notes produced during interactions with patients. It also functions as an information gateway, providing clinicians access to test results, images, and reports stored elsewhere. In addition, patients are

sometimes tracked remotely with sensors that provide information streams on their status. Complete clinical workflows can be reconstructed using process-mining techniques on event logs that record the progress of patients through the clinical environment.

Data velocity presents a major quality- and analytical-integration challenge: a steady stream of clinical data is generated while a workflow is ongoing, making quality-control and preprocessing decisions difficult. Node features from risk models that reflect patient characteristics (age, lab test results, etc.) and temporal trends (e.g., declining kidney functions) can be computed in a straightforward manner. However, other aspects are much more intricate, sometimes requiring model output design and validation. The importance of external factors that do not appear in standard EHR data is exemplified by the severely reduced quality of clinical-decision support for patients presenting in the aftermath of a disaster. In these cases, networks model real-time risk-perception and decision-support responses as uncertainty-neutral. For learning purposes, labels for care-quality indicators are derived from the association of demographic characteristics of patients and temporal periods with the performance indicators of interest.

### Equation 2: Anomaly Detection Threshold Equation

The paper emphasizes threshold-based anomaly detection for identifying unsafe clinical patterns. A statistical anomaly rule is:

$$A(t) = \begin{cases} 1, & \text{if } |x(t) - \mu| > k\sigma \\ 0, & \text{otherwise} \end{cases}$$

Where:

- $A(t)$ = anomaly indicator
- $x(t)$ = observed clinical metric
- $\mu$ = mean expected value
- $\sigma$ = standard deviation




- $k$ = sensitivity constant

### 2.3. Risk awareness and decision support

Risk and decision support in health systems relies on the capability of workflows to quantify risks for the patients under treatment and evaluate the associated costs–benefits of applying different treatment options. Arguably, risk-aware decision support is highly desirable in any health setting, yet difficult to realize in practice because it typically requires specialized models tailored for the different health conditions encountered during the care processes delivered by the institution. Risk assessment and decision support should thus be an integral part of clinical interaction networks, with risk models incorporated dynamically into the networks and updated following changes in health conditions, treatment protocols, or clinical staff.

Similarly to the quality of care assessment, risk modeling is not an explicit component of care process optimization. Risk-aware decision support should therefore focus on patient-level optimization, taking the cost–benefit analysis a step further by assessing the hazard of not being seen in time. Risk-calibrated machine-learning models thus combine risk and timing into a single score that summarizes how well an incoming patient is doing relative to the crowd and can be used to trigger clinician intervention even when time conditions are satisfied. Behind the scenes, these models allow the construction of such risk-aware decision support tools by identifying the patient factors that drive treatment decisions and quantifying the impact of the interaction and timing rules on risk.

Clinical Workflow Distribution



## 3. Architecture of self-evolving networks for care optimization

The architecture of self-evolving networks supporting real-time care optimization relies on three elements: data governance and interoperability, learning mechanisms that enable adaptability, and feedback loops that facilitate clinical validation. A cohort of self-evolving networks supports the quality, safety, and efficiency of an intensive care ward, and a dedicated control loop monitors emerging risks.

Supporting the data governance layer is a data ecosystem made up of the data available in the organization and complementary data sources, as well as data streams from patients and continuously learning knowledge representations built upon these sources and streams. Interoperability is achieved through established standards, common formats and vocabularies, privacy controls and access policies, and audit mechanisms that ensure action traceability and data provenance.

Data streaming velocity, timing windows, noise, burstiness, and importance for healthcare decisions are critical factors influencing the end-to-end pipeline architecture. Online





learning addresses model drift by detecting concept drifts and replacing affected models via continual learning. The models serving workflow-processing decision points are regularly validated by clinicians to ascertain their relevance and usefulness. Integrating these validations into clinical governance ensures controlled and transparent updates.

**Table 3. SECIN Functional Capabilities**


| Capability            | Description                                      | Clinical Benefit                 |
|-----------------------|--------------------------------------------------|----------------------------------|
| Real-Time Monitoring  | Continuous observation of workflows and patients | Faster response to deterioration |
| Anomaly Detection     | Identification of abnormal care patterns         | Prevention of adverse events     |
| Risk Scoring          | Predictive estimation of patient risk            | Better triage and prioritization |
| Workflow Optimization | Analysis of bottlenecks and delays               | Improved efficiency              |
| Continual Learning    | Dynamic model updates using new data             | Sustained accuracy               |
| Explainable Analytics | Transparent reasoning behind predictions         | Increased clinician trust        |

### 3.1. Data governance and interoperability

The design and operation of self-evolving clinical interaction networks depend on a data-governance framework that ensures standards-compliant data sharing while respecting the privacy concerns of patients and healthcare professionals. Such a framework covers four aspects: the specification of standards for data formats and protocols, the controls necessary for reliable privacy protection, the definition of access policies specifying the intended uses of the data, and the implementation of mechanisms capable of tracking the provenance of each insight generated by the network. By specifying the conditions under which these insights may be

released, the loop that feeds them back into the networks can be construed as a collaborative learning mechanism.

The large diversity of data types used implies that the network must be built on a service-oriented architecture that comprises various modules, each dedicated to producing one or several data types. Specificities of these modules include not only the required data formats but also the vocabularies or ontologies they conform to, the processing of the data and knowledge they generate, and their readiness to accept incoming data in real-time. Module specifications are associated with privacy enablers or supervisory systems that guarantee that the intended privacy guarantees can be delivered. In turn, these enablers deploy privacy controls (e.g., de-identification, anonymization, access control, disclosive policies) and audit functions that check new data coming to be integrated into the module against the guarantees that can be achieved.



**Fig 3: Self-Evolving Networks (SEN)**

### 3.2. Learning mechanisms and adaptability

Self-evolving clinical interaction networks are designed to operate in dynamic environments characterized by open-ended data and requirements. They therefore need learning mechanisms capable of continually integrating new information, whether from standard real-time data streams or from recorded events that can be retrospectively analyzed. Such learning can also be considered in relation to factors that create discontinuities in clinical workflows or care quality.

#### Equation 3: Care Efficiency Metric



The paper defines healthcare efficiency in terms of resource usage and waiting time. A queue-efficiency equation is:

$$E = \frac{\lambda}{\mu}$$


Where:

- $E$  = system utilization efficiency
- $\lambda$  = patient arrival rate
- $\mu$  = service rate

### 3.3. Feedback loops and clinical validation

The emergence of self-evolving clinical interaction networks enables the development of pragmatic frameworks for real-time care quality and risk-aware safety monitoring, analytics, and optimization. Evidence from an operational network in a hospital ward setting demonstrates the feasibility of real-time data governance, UNIX-like commands for defining, configuring, managing, and monitoring networks, and automatic generation of the feedback loops needed for continual learning. Efforts toward deploying such networks in diverse healthcare environments and use contexts can facilitate setting up the feedback loops required for validation-of-effectiveness, safety, and generalization in the clinical domain and enhance the networks' alignment with clinician experience and the evolving clinical reality.

Real-world clinical environments are populated by expert human clinicians and health workers who continually adapt and respond to subtle and complex disturbances and challenges in a probabilistic manner. In turn, real-world clinical networks can be viewed as a collection of persistent workflows that profoundly and visibly impact the lives of a broad range of people, including patients, clinicians, family and friends, and the entire health system. Self-evolving clinical interaction networks enable the creation of pragmatic frameworks for real-time monitoring and optimization of care quality and safety and considerable spatiotemporal-scale diversity across the healthcare domain.



## 4. Methods for real-time analytics and monitoring

Performance metrics develop care quality and safety, particularly effectiveness, efficiency, and safety. They should directly reflect patient experience, ideally encouraging feedback through channels acceptable in the respective cultural contexts. Continuous monitoring triggers anomaly detection for preventive interventions when feasible. Risk scores relate to adverse outcomes, such as deaths or unanticipated transfers, enabling risk-aware decisions regarding care quality. Model calibrations control score distributions, and thresholds are identified through available outcomes. High-risk patients require particular attention, and low-risk patients require more resources, especially for chronic diseases. Transparency enables prudent resource allocation, avoids clinical drift, or signals potential issues that risk exposure analysis might not capture. End-to-end latency targets are derived from clinical needs, with persistent feedback determining the acceptable degradation of timeliness in the face of latency constraints. Patient experience data influence risk models, anomaly-detection algorithms, and clinical decision support.

Risk-awareness has three major aims. First, it detects patients at higher risk for adverse outcomes, triggering targeted attention or preventive actions. Second, it assesses patients'





clinical or treatment progress relative to expectations given their risk. Third, it informs risk-aware decisions involving trade-offs between quality indicators, such as patient safety, timeliness, continuity of care, and the achievement of clinical targets, and other patient outcomes that indirectly depend on resource allocation, such as overall clinical or functional status and patient experience.

**Table 4. Metrics for Care Quality and Safety**

| <b>Metric Category</b> | <b>Definition</b>                                       | <b>Example Indicators</b>         |
|------------------------|---------------------------------------------------------|-----------------------------------|
| Effectiveness          | Degree to which expected clinical outcomes are achieved | Recovery rate, readmission rate   |
| Efficiency             | Optimal use of healthcare resources                     | Waiting time, bed occupancy       |
| Safety                 | Reduction of adverse clinical events                    | Medication errors, ICU transfer   |
| Timeliness             | Speed of clinical response and analytics                | Alert latency, intervention delay |
| Continuity of Care     | Consistency of long-term patient management             | Follow-up adherence               |
| Patient Experience     | Satisfaction and perceived care quality                 | Feedback scores                   |

#### 4.1. Metrics for care quality and safety

Quality and safety of clinical care can be assessed along multiple dimensions such as effectiveness, efficiency, safety, and patient satisfaction and experience. Existing methods for quantifying these dimensions can be adapted and formalized into real-time analytics, supporting the monitoring of the quality of care delivery, and the design of decision-support tools.

Effectiveness of clinical care can be quantified as the degree of achievement of expected outcomes. Many clinical decision-support tools, including acute-deterioration warning systems, are designed to enhance clinical effectiveness by alerting clinicians of an increased probability of an adverse event. A warning with a low positive predictive value is nevertheless useful if timely, as it allows preventive measures to be put in place. By analogy, a low negative predictive value is acceptable if the alerts have a critical time window. The adequacy of any such warning system depends on the particular medical condition and clinical context. Consequently, clinical workflows incorporating such systems need to be deployed and validated by clinicians.

#### Equation 4: Latency Constraint Equation for Real-Time Analytics

The architecture requires strict latency constraints for streaming analytics:

$$L_{total} = L_{ingest} + L_{process} + L_{decision}$$

Where:

- $L_{total}$  = total end-to-end latency
- $L_{ingest}$  = data acquisition latency
- $L_{process}$  = analytics computation latency
- $L_{decision}$  = clinician alert/response latency


#### 4.2. Anomaly detection and risk scoring

The analytics framework supports detection of anomalous care delivery patterns and scores patients' risk profiles with respect to specified clinical endpoints. Multiple algorithms have been employed to identify potential anomalies, including threshold-based methods, control charts, time series abnormality detection techniques, and clustering and classification approaches. Patient-centered metrics are expressed as a composite of multiple indicators, involving data fusion, normalization, and aggregation across clustered



patients. The performance of these scores for risk stratification is evaluated through a risk-score–outcome association test.

Timeliness is critical for models supporting care monitoring and anomalous pattern detection. An end-to-end latency target dictating the frequency of the main analytics operations has been set at 10 minutes for hospital wards. Such short latencies usually necessitate streaming architectures and the use of dedicated online algorithms tailored to each task. When these additive constraints are examined for various tasks, it is further discovered that the absence of a streaming approach in a single task can substantially degrade the end-to-end latency.



**Fig 4: Anomaly-based threat detection in smart health**

#### 4.3. Timeliness and latency considerations


In many healthcare settings, real-time analytics can provide significant clinical value. However, the information provided must be timely. While data streaming technologies have matured to support end-to-end latency on the order of seconds, many analytics methods have been designed for offline use only—these may still be practical if the latency can

be tolerated or if the results can be cached for later use. Therefore, latency targets can be classified as follows.

End-to-end latency targets specify performance goals for the system as a whole—from raw data generation to consumption of the analytics results. Where a patient receives care through multiple stages, such as in an emergency department context, the subprocesses must be executed in order but need not be finished nearly at the same time. Latency targets for each of the subprocesses may therefore be less stringent than would be the case for a parallel process. The targets for individual subprocesses, however, must still be considered, since delays in any one of them can lead to degradation of the overall system.

Real-time analytics check temporal conditions within sliding time windows, either counting occurrences of events or estimating statistics. Other algorithms for time-to-event inference, risk scoring, and anomaly detection may employ a warming-up phase; once the system has entered a steady state, the condition checks can be performed with the desired frequency.

In other contexts such as hospital wards, clinical groupings may change more slowly. Moreover, patients receive care continuously rather than in successive stages. Timeliness is nevertheless important—while care workers may decide not to check their phones, they expect alerts to have been received and acted upon when they next consult the system. For such situations, streaming architectures provide a promising solution.



## 5. Case studies and empirical evaluations

A hospital ward in an educational hospital served as the setting for preparing self-evolving care-optimization networks, followed by deployment experiments. Data from Electronic Health Records (EHR) and other administrative sources were used for network learning, covering 23393 hospitalizations from July 2011 to June 2017. A self-evolving network enabled real-time evaluation support for care safety, effectiveness, and efficiency, with negative feedback refreshing the underlying network. During a one-month deployment, caregiver interaction data provided feedback on the derived care-quality-monitoring streams, and the system was iterated to improve usefulness.

A previously developed clinical-interaction network for the Accreditation Canada accreditation activities of an Emergency Department was capable of evolving through feedback from clinical supervision. The use of queueing theory analysis validated its support on waiting times. A self-evolving triage-model network highlighted the effect that changing the distribution of different triage categories had on the ED throughput and 24-h readmissions. Moreover, negative feedback by caregiver interaction allowed the adaptation of the model for representing triage monitoring and reducing the burden of monitoring.

A chronic-disease-management clinical-interaction network for a Family-Health-Team (FHT) nurse practitioner was created and self-evolved using data from 2014 to 2019. The indicator of continuity of care for adult-care-users increased with the number of previous years and lower cost for the FHT. For pediatric patients, continuity of care was negatively associated with the risk of hospitalization and emergency department visit, whereas the maintenance of prescribed medication for asthma patients, were found to be positively correlated with the risk of hospitalization. The obtained results suggest that assigning a family physician to every community should be the ultimate goal to optimize these services.

**Table 5. Risk-Aware Analytics Methods**

| Method                    | Technique Used                     | Application                  |
|---------------------------|------------------------------------|------------------------------|
| Threshold-Based Detection | Rule-based monitoring              | Critical vital sign alerts   |
| Control Charts            | Statistical monitoring             | Workflow stability tracking  |
| Time-Series Analysis      | Sequential pattern recognition     | Trend prediction             |
| Clustering Algorithms     | Grouping similar patient behaviors | Population stratification    |
| Classification Models     | Predictive machine learning        | Risk scoring                 |
| Explainable AI            | Feature attribution methods        | Transparent decision support |

### 5.1. Hospital ward deployment

The first real-world application was deployed in a hospital ward, leveraging a heterogeneous data stream from the electronic health record (EHR), combined with clinical and nursing logs, sensors, and the contextual signalling of daily work routine elements. Algorithm calibration and tuning





based on clinical expertise remained mandatory for all involved processes, although their presence during operations was a matter of submission to the processes' alert signals or their low-volume monitoring.

Despite being mainly a validating pilot test of the connection between data flow structure and quality monitoring indication, results showed areas where early warning of deterioration (within the rediagnosis time frame of a few hours) could produce better-than-usual outcomes in terms of readmissibility, pressure sore incidents, psychosis-related outcomes, prolonged stay duration, and dead multilocalization. A queuing network connected to the ward and representing its only admission point showed timeliness, triage-dependence on patient choice, and monitoring of both throughput and typical-service-duration outcomes, providing indications on when the number of patients waiting was producing a negative effect on care quality.

#### Equation 5: Continual Learning / Model Drift Update Equation

Self-evolving networks continuously adapt using online learning:

$$\theta_{t+1} = \theta_t - \eta \nabla J(\theta_t)$$

Where:

- $\theta_t$  = model parameters at time  $t$
- $\eta$  = learning rate
- $\nabla J(\theta_t)$  = gradient of loss function

#### 5.2. Emergency department workflow optimization

The case study illustrates optimization of emergency department workflows through modeling and real-time monitoring of a clinical process. High admission rates and care delays in the clinical center from the study would benefit from system performance improvement, yet timeliness and risk-awareness indicators have not been previously

considered together. A queueing model evaluated the impact of triage configuration on waiting times. Comparative evaluation also focused on patient safety and quality of care: a high ratio of abandoned treatment in the emergency department highlighted the need for mitigating—rather than just monitoring—this phenomenon. Insights led the unit manager to reconfigure the service, temporarily shifting an available physician from specialized care to ensure immediate first contact. The adaptation, supported by observational data, was reflected in EHRs, and the corresponding feedback indicated improved patient experience. Continuous monitoring subsequently confirmed the decrease in waiting times. These analyses validate computational support for decision making in emergency departments and their incorporation into the clinical workflow.

Operational decisions in emergency departments rely on the timely interpretation of complex patterns in multiple information sources, such as waiting queues, staffing ratios, abandoned treatment rates, acuity levels, workload, and clinical outcomes. The literature proposes static analysis approaches—e.g., capacity allocation and service scheduling models—but does not examine the real-time monitoring of such indicators nor the integration of prediction and optimization tools.




Fig 5: Emergency Department Simulator



### 5.3. Chronic disease management in primary care

In the third case study, the Self-Evolving Clinical Interaction Networks are deployed within a Primary Healthcare setting, where it provides indicators of patient treatment continuity and disease closeness, detecting high-risk patients. The healthcare flow is characterized by permanent electronic records centralized in the primary disease health service agencies in a national public health system. These diseases require permanent monitoring, and the analysis considers mainly the integration of the same clinical teams performing the consultations, examinations, interventions and maintaining the electronic records. The seeking of continuity in their use is more considered as an indicator of opportunity in the use of the system. It is in serious consideration the analysis of the closeness of the patients, in the searching of the control of the disease and indication to return to the clinician. However, in the final analysis, it considers also the decrease of the cost of the health assistance, evidencing the increase of the consultations of the patients in the year after using the indication system.

In chronic diseases such as diabetes mellitus, systematic annual checkups with appropriate examinations serve as important indicators of disease management. Regular monitoring and assisting patients with pressure ulcers also reduce costs. The analysis indicates the ability to identify patients at risk of reduced continuity of healthcare. Based on two consecutive years of data, the calibration of suggestions for the decrease in expected costs shows agreement between the model and reality: the direction of the financial impact is verified. All of these demonstrate the multivariate outcome of the self-evolving networks: the aid in risk using the healthcare service is improved, and at the same time, it gives unified and relevant evidence of the continuity of use of the service.



## 6. Ethical, legal, and social implications

Privacy concerns, data stewardship and security of surfaces remain paramount. The utility of data for optimising service delivery needs to be balanced with privacy considerations. During an ethics workshop, recommendations were made for privacy, consent and data stewardship. A layered model of consent for sharing patient data with a Care Optimisation System — based on broad consent, service-related transaction flow consent and ad-hoc consent for secondary studies — was proposed. Provided privacy and identity-protection measures are strictly observed, the benefits of continued analyses and broad utility of historical patient data far outweigh the risks and costs of such analyses.

Since the models inferred from the historical data undergo continuous validation and small-scale testing, there is sufficient transparency provided to the users of the models to ensure that model support does not lead towards a black-box effect. Traceability of all decisions is also considered critical for preventing black-box effects. Transparency is further enhanced by empowering clinicians to add relevant information at the time of model use. The integration of these elements with access logs is of paramount importance for building trust in such models, and is along expected lines. Appropriate explainability measures beyond these actions remain important aspects in such decision-aiding solutions.



Evaluation of data questioned the risk of discrimination with respect to race and sex, revealing no negative correlations. Nevertheless, sensitivity analysis remains of utmost importance since even small biases can induce discrimination. Incorporating bias detection fairness filters in the model-building pipeline also assists in preventing introduction of such bias. Additional monitoring during the testing phase, and the flexibility of involving Minority and Women-owned Business Enterprises in determining model-exposure levels, further reduce bias and discrimination risks. The cognitive challenges involved in the interpretation of algorithms were also duly considered, and solutions along these areas explored.

**Table 6. Architecture Layers of SECIN**

| Architecture Layer    | Main Responsibilities                 | Technologies/Methods             |
|-----------------------|---------------------------------------|----------------------------------|
| Data Governance Layer | Privacy, interoperability, provenance | Access control, anonymization    |
| Integration Layer     | Data ingestion and fusion             | APIs, HL7/FHIR standards         |
| Analytics Layer       | Risk scoring and anomaly detection    | ML models, streaming analytics   |
| Learning Layer        | Continual learning and adaptation     | Drift detection, online learning |
| Validation Layer      | Clinical feedback and governance      | Human-in-the-loop systems        |
| Visualization Layer   | Dashboards and alerts                 | Real-time monitoring interfaces  |

### 6.1. Privacy, consent, and data stewardship

Privacy and data stewardship are paramount when designing decision support systems in healthcare environments, particularly those using personal data and offering automated

decisions. Addressing ethical, legal, and social implications must be an integral part of the design, and not an ad-hoc consideration after technology demonstration. Key aspects of privacy, de-identification, consent and usage policies, access control, and auditability are outlined below.

The data steward of a deployed self-evolving clinical interaction network governs privacy and data usage. If the produced analytic modules have no legal or operational authority, decision-making may remain with clinicians, but stewards still need to inform patients of the purposes and scope of the data usage and decide on consent models. When personal data is used to train or refine a module with decision-making authority, protective measures should ensure that patient exposure during model application is minimized, by detecting sensitive patient groups and applying risk-averse strategies. De-identification is crucial, as the high re-identification risk of relational data may endanger not only the patients involved, but also their relatives and acquaintances.

Therefore, patient data enters the system through a clearly defined pipeline and with the approval of the data steward. The data provenance trace allows auditing to verify that data has been used only for legitimate and authorized purposes. Depending on the authorizations granted by the steward, different data ensemble policies can be designed. A standard policy requires that patient data can only be used for their own care and for staff training. More permissive policies are also possible, with appropriate matching of patient exposure and risk.



**Fig 6: Privacy, consent, and data stewardship**






## 6.2. Transparency and explainability

Automated decision-support systems are often considered “black boxes,” and some citizens may reasonably question the ethics of employing these systems for important societal functions such as driving, hiring, and criminal sentencing. Explanations of how a model makes predictions may increase its perceived fairness, and the risk of unanticipated or harmful consequences can be reduced if users understand how recommendations were derived. The selection of advisable actions may also be better informed when human decision-makers understand the underlying computational process, and this can increase their confidence in the decisions. Consequently, explainability emerges as a key ELSI theme.

Technical strategies include the automatic generation of model narratives that summarize system operation; recording model inputs, predictions, and advice; retention of decision traces that explain source element relevance; and visualizations that highlight the effect of different data elements on predictions. Exposure to the decision points and predictive feature relevance of a model may build understanding and trust among recipients. Biased training data can lead to biased models, and information about the populations represented during training can help alleviate these concerns. Transparency and explainability can therefore be enhanced through an inclusive design process that ensures representation of different groups during development and deployment. Approaches for detecting these issues in risk models are being investigated, and opportunities can be seized to bring any undesirable properties of models to the attention of domain experts and request corrective action. These principles can be implemented in a range of domains, implementing equity and bias mitigation measures as needed.




## 6.3. Equity and bias mitigation

Concern for equity in AI presents a range of challenges at different stages of a model’s life cycle. Initial data collection must ensure that individuals belonging to different social groups have a chance of being exposed to the phenomenon under study; for example, people belonging to a non-majority ethnic/cultural group must nevertheless suffer the disease/condition in significant numbers if their medical needs are to be addressed by a risk model. Once a model is developed, tools are becoming available for measuring bias across different social groups, thereby allowing developers to safeguard against perpetuating systemic disadvantages. In a further step, bias can be actively mitigated through appropriate preprocessing of the training data, in-processing adjustments during model building, and postprocessing realignment of outputs. All these approaches are complementary and can be applied when necessary, as the specific means of addressing bias can differ across problems and social contexts.

A specific aspect of equity in risk-aware monitoring systems concerns sensitivity to external risk factors not recorded in the data. For example, being a member of a specific ethnic or cultural group can increase the risk score for a condition, independently of how accurately the overall risk model is trained. In addition, model predictions should be expressed in terms of diagnosis or treatment considerations that, if not addressed, could lead to worse treatment or outcomes for



minority patients. Evaluation of Real-World Health Technologies In recent years, and especially since the COVID-19 pandemic, the demand for natural language processing (NLP) models, such as ChatGPT, capable of assisting patients and medical specialists across thousands of requests in multiple languages has grown tremendously. With its launch, OpenAI simultaneously published two papers showing that the AI—trained on vast corpora of web data, books, and human interactions— demonstrated impressive zero-shot performance on several NLP benchmarks, including the Stanford Question Answering Dataset. Subsequent releases of the model have also sought to enhance capabilities for visualizing, analyzing, and converting multimodal query responses.





## Self-Evolving Feedback Layer

### 7. Challenges, limitations, and future directions

Scalability and integration with existing systems present difficult challenges for data-centric approaches. The successful deployment of practical analytics usually requires end-to-end data governance, yet the absence of standard formats and vocabularies for exchange, access, and synthesis of patient journey data remains a key obstacle. Migration paths that incrementally integrate evolved monitoring and analytics on top of existing legacy systems would lower barrier to entry. Highly modular system designs facilitate phased implementations. Open-source efforts such as the Knowledge Graph Toolkit and their role in fostering seamless integration should be further pursued and built upon.

The robustness of data-centric adaptive systems to natural variations remains to be seen. Data streams in clinical settings possess a unique temporal characteristic: they are simultaneously non-stationary, with many changing distributional or relationship-oriented properties, and highly variable in data quality (e.g., missingness, noise, class imbalance, class rarity). Mitigating these natural properties—and in particular the drop in data quality during a large event such as an epidemic—remains central to many system designs. In particular, a proper longitudinal monitoring mechanism should examine both treatment and patient outcomes over time to assess whether the developed analytics or monitoring add genuine clinical value.

#### 7.1. Scalability and integration with existing systems

Commercial solutions for specific tasks, such as anomaly detection and risk scoring, are already widely available. Future developments will therefore focus on building large-scale, modular networks for real-time monitoring and decision support that integrate these individual components. A major challenge is the availability of timely and high-quality data. Despite current efforts to leverage EHR data for

automated clinical process monitoring and configuration using private hospital data, the labor-intensive nature of such techniques has restricted their use to a limited number of clinical processes that are monitored without major configuration. Deployment for public services in the context of a regional health system will likely face even greater hurdles in this regard. However, establishing an efficient and effective configuration of the approach for a given process is a necessary step before any operational deployment.

Integration with existing systems in order to use them as data suppliers is thus a priority. Beyond the obvious gain in timeliness and potentially quality of the data, this would ease the burden of adoption for prospective users. Support for real-time analytics in areas such as monitoring or decision support could then be provided to such interconnected systems, potentially with added feedback loops that take the data configuration of the system into account. The developed architectural design, built upon formal independent components, aims to enable such an integration while providing a migration path toward a full networking solution.





## Real-Time Care Optimization Layer

### 7.2. Robustness to data quality variability

Data quality in real-time healthcare systems varies due to factors such as incompleteness, noise, and sample imbalance. Missing data arise when patients do not undergo certain measurements, sensor failures occur, or actions are not logged/recorded. Stream data may be labeled as incomplete or erroneous or contain incorrect values, resulting in noisy training samples. Time-dependent models may receive imbalanced samples at different times, for example, limited numbers of patients in a certain risk range. Moreover, the proportions of individuals in different risk groups may vary significantly over time. Learning systems must therefore be robust to these factors.

Model-agnostic resilience to sample imbalance can be achieved by two methods: over-sampling from under-represented categories and under-sampling from over-represented ones. Sensitivity to missing data can be reduced using methods such as mean/mode imputation and interpolation or employing robust algorithms such as tree-based methods for prediction tasks. A simple disaster recovery technique can be used even when the underlying anomaly detection mechanism is not robust to missing data: risk scores should ideally be re-evaluated when late or incomplete data appear. However, it may be difficult to recover valuable information on the importance of certain factors. Deep learning-based approaches may further improve robustness to imbalanced or noisy data samples.

### 7.3. Longitudinal impact assessment

Despite growing interest in examining these risks, little work has rigorously assessed the application of the general NLP model to specific patient-centered health-domain queries. Primarily descriptive analyses have raised concerns such as unreliable or incorrect answers, presentation of harmful clinical advice, generation of biased or prejudiced responses, and nonsensical visual outputs. The findings underline the need for further longitudinal evaluation of the model's



performance over time, addressing a range of medical queries, hypotheses, and applied objectives guided by recognized ethical principles for health technologies. These principles assert that for a technology to be considered responsible and thus trustworthy, it must be efficacious and supportive of the patient's best interests; patient-centered health outcomes must be continually monitored; and support for minimizing inequities must be both explicit and active.

## 8. Conclusion

Numerous innovations and discoveries in information and communication technologies and artificial intelligence have improved healthcare delivery. Nevertheless, increasing demand, rising costs, and chronic workforce shortages present unprecedented challenges. One promising response is to leverage the vast amount of real-time data generated in clinical environments to optimize care processes and facilitate evidence-based decision making. Research concentrates primarily on improving performance in specific areas such as quality, efficiency, and safety. Self-evolving clinical interaction networks seeking to monitor all aspects of care in real time are an attempt to systematically address such demands across diverse attributes and care categories.

Following an objective, scholarly synthesis, the work identifies several critical components that must be addressed for networks to be effective in real-world settings: careful governance of data generation, storage, and sharing; continuous learning of network representation; incorporation of feedback mechanisms to facilitate clinical validation and acceptance; definitions of data streams, quality indicators, and performance metrics to meet real-time requirements; and validation through deployment in diverse clinical situations. Addressing these aspects represents a step toward the development of self-evolving clinical interaction networks that optimize care processes, improve patient outcomes, and enhance the performance of clinical support systems such as risk assessment models and decision-support tools.

## References

1. Acosta, J. N., Falcone, G. J., Rajpurkar, P., & Topol, E. J. (2022). Multimodal biomedical AI. *Nature Medicine*, 28, 1773–1784.
2. Afshar, M., Adelaine, S., Resnik, F., Mundt, M. P., Long, J., Leaf, M., Ampian, T., Wills, G. J., Schnapp, B., Chao, M., & Brown, R. (2023). Deployment of real-time natural language processing and deep learning clinical decision support in the electronic health record: Pipeline implementation for an opioid misuse screener in hospitalized adults. *JMIR Medical Informatics*, 11, e44977.
3. Boulitsakis Logothetis, S., Green, D., Holland, M., & Al Moubayed, N. (2023). Predicting acute clinical deterioration with interpretable machine learning to support emergency care decision making. *Scientific Reports*, 13, 13563.
4. Challa, K. (2023). Transforming Travel Benefits through Generative AI: A Machine Learning Perspective on Enhancing Personalized Consumer Experiences. *Educational Administration: Theory and Practice*. Green Publication. Educational Administration: Theory and Practice. Green Publication. <https://doi.org/10.53555/kuey.v29i4.9241>.



5. Choi, A., Choi, S. Y., Chung, K., Chung, H. S., Song, T., Choi, B., & Kim, J. H. (2023). Development of a machine learning-based clinical decision support system to predict clinical deterioration in patients visiting the emergency department. *Scientific Reports*, 13, 8561.
6. Coelho, K. K., Nogueira, M., Vieira, A. B., Silva, E. F., & Nacif, J. A. M. (2023). A survey on federated learning for security and privacy in healthcare applications. *Computer Communications*, 207, 113–127.
7. Coley, R. Y., Walker, R. L., Cruz, M., Simon, G. E., & Shortreed, S. M. (2021). Clinical risk prediction models and informative cluster size: Assessing the performance of a suicide risk prediction algorithm. *Biometrics*, 77, 1375–1388.
8. Paleti, S., Burugulla, J. K. R., Pandiri, L., Pamisetty, V., & Challa, K. (2022). Optimizing digital payment ecosystems: AI-enabled risk management, regulatory compliance, and innovation in financial services. *Regulatory Compliance. And Innovation In Financial Services* (June 15, 2022).
9. Corny, J., Rajkumar, A., Martin, O., Dode, X., Lajonchère, J.-P., Billuart, O., Bézie, Y., & Buronfosse, A. (2020). A machine learning-based clinical decision support system to identify prescriptions with a high risk of medication error. *Journal of the American Medical Informatics Association*, 27, 1688–1694.
10. Çubukçu, H. C., Topcu, D. İ., & Yenice, S. (2023). Machine learning-based clinical decision support using laboratory data. *Clinical Chemistry and Laboratory Medicine*, 62, 793–823.
11. Du, Y., McNestry, C., Wei, L., Antoniadi, A. M., McAuliffe, F. M., & Mooney, C. (2023). Machine learning-based clinical decision support systems for pregnancy care: A systematic review. *International Journal of Medical Informatics*, 173, 105040.
12. Gottimukkala, V. R. R. (2020). Energy-Efficient Design Patterns for Large-Scale Banking Applications Deployed on AWS Cloud. *International Journal of Advanced Research in Computer and Communication Engineering (IJARCCE)*, DOI, 10.
13. Haroz, E. E., Grubin, F., Goklish, N., Pioche, S., Cwik, M., Barlow, A., Waugh, E., Usher, J., Lenert, M. C., & Walsh, C. G. (2021). Designing a clinical decision support tool that leverages machine learning for suicide risk prediction: Development study in partnership with Native American care providers. *JMIR Mental Health*, 7, e24377.





14. Hong, N., Liu, C., Gao, J., Han, L., Chang, F., Gong, M., & Su, L. (2022). State of the art of machine learning-enabled clinical decision support in intensive care units: Literature review. *JMIR Medical Informatics*, 10, e28781.
15. Inala, R. Designing Scalable Technology Architectures for Customer Data in Group Insurance and Investment Platforms.
16. Jahandideh, S., Ozavci, G., Sahle, B. W., Kouzani, A. Z., Magrabi, F., & Bucknall, T. (2023). Evaluation of machine learning-based models for prediction of clinical deterioration: A systematic literature review. *International Journal of Medical Informatics*, 175, 105084.
17. Kim, M. S., Ryu, J. M., & Choi, B. K. (2022). Development and effectiveness of a clinical decision support system for pressure ulcer prevention care using machine learning: A quasi-experimental study. *CIN: Computers, Informatics, Nursing*, 40, 698–706.
18. Miyachi, Y., Ishii, O., & Torigoe, K. (2023). Design, implementation, and evaluation of the computer-aided clinical decision support system based on learning-to-rank: Collaboration between physicians and machine learning in the differential diagnosis process. *BMC Medical Informatics and Decision Making*, 23, 26.
19. Narmadha, K., & Varalakshmi, P. (2022). Federated learning in healthcare: A privacy preserving approach. *Studies in Health Technology and Informatics*, 294, 194–198.
20. Pfohl, S. R., Foryciarz, A., & Shah, N. H. (2021). An empirical characterization of fair machine learning for clinical risk prediction. *Journal of Biomedical Informatics*, 113, 103621.
21. Rajpurkar, P., Chen, E., Banerjee, O., & Topol, E. J. (2022). AI in health and medicine. *Nature Medicine*, 28, 31–38.
22. Rahman, A., Hossain, M. S., Muhammad, G., Kundu, D., Debnath, T., Rahman, M., Khan, M. S. I., Tiwari, P., & Band, S. S. (2022). Federated learning-based AI approaches in smart healthcare: Concepts, taxonomies, challenges and open issues. *Cluster Computing*, 26, 1–41.
23. Aitha, A. R. (2021). Optimizing Data Warehousing for Large Scale Policy Management Using Advanced ETL Frameworks.
24. Rubins, D., McCoy, A. B., Dutta, S., McEvoy, D. S., Patterson, L., Miller, A., Jackson, J. G., Zuccotti, G., & Wright, A. (2022). Real-time user feedback to support clinical decision support system improvement. *Applied Clinical Informatics*, 13, 1024–1032.



25. Sadilek, A., Liu, L., Nguyen, D., Kamruzzaman, M., Serghiou, S., Rader, B., Ingerman, A., Mellem, S., Kairouz, P., Nsoesie, E. O., MacFarlane, J., Vullikanti, A., Marathe, M., Eastham, P., Brownstein, J. S., Aguera y Arcas, B., Howell, M. D., & Hernandez, J. (2021). Privacy-first health research with federated learning. *NPJ Digital Medicine*, 4, 132.
26. Sharma, S., & Guleria, K. (2023). A comprehensive review on federated learning based models for healthcare applications. *Artificial Intelligence in Medicine*, 146, 102691.
27. Singer, S. J., Kellogg, K. C., Galper, A. B., & Viola, D. (2022). Enhancing the value to users of machine learning-based clinical decision support tools: A framework for iterative, collaborative development and implementation. *Health Care Management Review*, 47, E21–E31.
28. Susanto, A. P., Lyell, D., Widyantoro, B., Berkovsky, S., & Magrabi, F. (2023). Effects of machine learning-based clinical decision support systems on decision-making, care delivery, and patient outcomes: A scoping review. *Journal of the American Medical Informatics Association*, 30, 2050–2063.
29. Vasey, B., Nagendran, M., Campbell, B., Clifton, D. A., Collins, G. S., Denaxas, S., Denniston, A. K., Faes, L., Weber, W., Wheatstone, P., McCulloch, P., & DECIDE-AI Expert Group. (2022). Reporting guideline for the early-stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. *Nature Medicine*, 28, 924–933.
30. Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., & Wang, F. (2021). Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 5, 1–19.
31. Paleti, S., Burugulla, J. K. R., Pandiri, L., Pamisetty, V., & Challa, K. (2022). Optimizing digital payment ecosystems: AI-enabled risk management, regulatory compliance, and innovation in financial services. *Regulatory Compliance. And Innovation In Financial Services* (June 15, 2022).
32. Li, R., Chen, Y., & Ritchie, M. D. (2020). Electronic health records and polygenic risk scores for predicting disease risk. *Nature Reviews Genetics*, 21, 493–502.
33. Acosta, J. N., Falcone, G. J., Rajpurkar, P., & Topol, E. J. (2022). Multimodal biomedical AI: Opportunities and challenges for clinical care. *Nature Medicine*, 28, 1773–1784.
34. Miyachi, Y., Ishii, O., & Torigoe, K. (2023). Machine learning-supported differential diagnosis and clinical



- decision-making in primary care. *BMC Medical Informatics and Decision Making*, 23, 26.
35. Kim, M. S., Ryu, J. M., & Choi, B. K. (2022). Machine learning-based clinical decision support for prevention and improvement of nursing care outcomes. *CIN: Computers, Informatics, Nursing*, 40, 698–706.
36. Çubukçu, H. C., Topcu, D. İ., & Yenice, S. (2023). Machine learning applications for clinical laboratory decision support and patient risk assessment. *Clinical Chemistry and Laboratory Medicine*, 62, 793–823.
37. Jahandideh, S., Ozavci, G., Sahle, B. W., Kouzani, A. Z., Magrabi, F., & Bucknall, T. (2023). Machine learning for early detection of clinical deterioration in hospital settings: A systematic review. *International Journal of Medical Informatics*, 175, 105084.
38. Logothetis, S. B., Green, D., Holland, M., & Al Moubayed, N. (2023). Interpretable machine learning for prediction of acute clinical deterioration and emergency care decision support. *Scientific Reports*, 13, 13563.
39. Ramgopal, S., Sanchez-Pinto, L. N., Horvat, C. M., Carroll, M. S., Luo, Y., & Florin, T. A. (2023). Artificial intelligence-based clinical decision support in pediatrics. *Pediatric Research*, 93, 334–341.
40. Coelho, K. K., Nogueira, M., Vieira, A. B., Silva, E. F., & Nacif, J. A. M. (2023). Federated learning for privacy-preserving healthcare intelligence: A survey. *Computer Communications*, 207, 113–127.
41. Afshar, M., Adelaine, S., Resnik, F., Mundt, M. P., Long, J., Leaf, M., Ampian, T., Wills, G. J., Schnapp, B., Chao, M., & Brown, R. (2023). Real-time clinical decision support using natural language processing and deep learning in electronic health records. *JMIR Medical Informatics*, 11, e44977.