

Original Article

Real-Time Healthcare Event Processing: Stream Analytics for Clinical Decision Support

Arjun Warriier
Senior Technology Consultant.

Abstract - The growing complexity of modern healthcare, along with the abundance of real-time data in the form of electronic health records (EHRs), wearables, bedside monitors, and the like, poses both a challenge and an opportunity for driving clinical decision support (CDS) to the next level. Traditional batch systems are also often insufficient for controlling the high velocity, volume, and variety of healthcare data, which require timely analysis and response. To address these constraints, this paper presents a new event-driven architecture for real-time healthcare event processing, enabling continuous monitoring and rapid response for informed clinical decision-making. Our proposed approach is based on the use of stream analytics frameworks (e.g., Apache Kafka, Apache Flink) to ingest, transform, and analyze data from various clinical environments in real-time, providing a scalable and low-latency alternative to traditional healthcare information systems. The primary research challenge in this paper is to design and implement real-time clinical decision support mechanisms that respond promptly to streaming health events. We introduce healthcare event streaming patterns that enable systems to identify abnormalities, such as a sudden rise in blood pressure, irregular heartbeats, and early signs of sepsis. These patterns underlie real-time alerting systems that automatically launch clinical interventions, direct information to the appropriate care teams, and record data in the EHR for traceability and compliance. This architecture is completely real-time – clinical decisions are not hampered by processing backlogs or outdated data; all available data is made up-to-date, thereby minimizing errors that can lead to complications.

One of the key technical contributions of this work is a 40% reduction in clinical response times for multiple use cases. These include emergency triage, intensive care unit (ICU) monitoring, and early warnings for chronic diseases. We experimentally verify the superior performance in early detection and response to clinical critical conditions by conducting extensive simulations with synthetic clinical time-series data, as well as real-world healthcare datasets. We further present how stream analytics is integrated with the current instance of the hospital information system (HIS), ensuring backward compatibility and ease of adoption, without disrupting hospital availability. Our system is designed to be HIPAA-compliant, featuring data encryption, audit logging, and access control. In addition, our approach is modular and can be easily deployed in the cloud, at the edge, and in hybrid settings, providing flexibility for a variety of health care environments, from tertiary care hospitals to rural health clinics. We demonstrate, through our implementation results, a significant reduction in time-to-decision, an enhancement in the accuracy of event detection, and a high level of clinician satisfaction with the CDS system.

This paper presents a comprehensive solution for enhancing clinical decision support through Real-time Stream Analytics. Handling healthcare events in real-time not only speeds up clinical response but also enhances the reliability and scalability of decision-making in urgent care events. This investigation has revealed that real-time event processing is not just a theoretical goal, but a strategic necessity within our increasingly digital healthcare environments, one that drives the emergence of responsive, intelligent, and patient-centric care.

Keywords: Real-Time Healthcare Analytics, Clinical Decision Support (CDS), Event-Driven Architecture, Healthcare Event Streaming, Stream Analytics, Apache Kafka, Real-Time Alerting, EHR Integration, Healthcare Interoperability, Clinical Response Optimization.

1. Introduction

The adoption of digital technologies in healthcare practice has led to an increase in the volume, velocity, and variety of information generated within clinical environments—known as the Data Revolution in Healthcare. With the proliferation of electronic health records (EHRs), the integration of Internet of Things (IoT) medical devices, and the arrival of wearable health monitors, the healthcare industry is experiencing a data revolution. However, most HIT systems are still based on batch processing models, which lead to gaps between data collection, processing, and clinical action. In (emergent) high-acuity departments like the (ERs) emergency rooms, ICUs, and operating rooms (surgical departments), the response time delay in seconds could result in fatal

decisions. Given this confirmation, architectures capable of receiving and processing RT data are essential for healthcare professionals to make efficient and rapid medical decisions.




Fig 1: Growth of Healthcare Data Volume

The purpose of this paper is to address the shortcomings of traditional batch-oriented CDSs by presenting an event-based CDS architecture that utilizes real-time streaming analytics to process clinical data. This architecture, based on event-driven computing, is designed to analyze, correlate, and take action on health-related events as they occur. By shifting from batch aggregation to stream processing, the system we propose would help create a more responsive and intelligent environment that is better suited to the dynamic and unpredictable nature of clinical care.

The concept of real-time analytics has gained significant popularity across various industries, including finance, logistics, and telecommunications. However, adoption in healthcare is limited due to the complexity of medical data, regulations, and clinical decision-making processes. This paper examines the potential application of state-of-the-art stream/analytics platforms, including Apache Kafka, Apache Flink, and Apache Storm, in the clinical domain. These platforms are designed to ingest and parse large-volume, high-velocity data streams, such as physiologic tracings, medication orders, lab results, and clinician notes.

One aspect of our design is the focus on health-specific event streaming patterns. These patterns represent the typical interrelationships between clinical events in time and space. They are defined, for example, among sequences of symptoms leading to septic shock or deterioration in vital signs in post-surgical care. By integrating these patterns into the stream processing pipeline, we can automatically identify clinically relevant situations. Furthermore, the system can include real-time alerts to clinicians via integrated dashboards, mobile alarms, or automatic EHR entries, which expedite response times and enhance care.

Another pertinent result of the study is the proven reduction of response time observed in the clinic. In silico testing and retrospective analysis demonstrated a method that reduces time to intervention by 40% relative to background methods. This advance is not simply a technological milestone, but a critical step forward for the technology, enabling it to decrease patient morbidity and save lives.

This paper is organized as follows. The literature review examines the current state of the art in clinical decision support and stream analytics in healthcare. The Methodology details the design of the event-driven architecture and explains the processing of real-time data using streaming frameworks. Performance and Real-World Deployments are presented in the Results section: the performance metrics and the results of simulated and real-world deployments. The Discussion situates the findings and addresses implementation issues, while the Conclusion presents the contributions and opportunities for future research.

2. Literature Review

The evolution of clinical decision support systems (CDSS) has traditionally been aligned with the development of electronic health records (EHRs), which primarily focused on digitizing patient documentation and enabling retrospective analysis [1]. However, as the volume and complexity of clinical data have grown exponentially, the need for real-time, actionable insights has become critical. The limitations of batch processing in healthcare are well-documented, with studies highlighting delays in intervention, inefficient care coordination, and missed early warnings of deteriorating conditions [2]. This has prompted a shift toward event-driven and real-time analytics architectures, especially in high-stakes clinical environments.

Initial attempts to incorporate real-time features in healthcare systems relied on rule-based CDSS tightly coupled with EHR platforms. For example, Wright et al. [3] emphasized the role of alerting logic embedded in EHRs to support clinical workflows. However, these systems were largely static, unable to scale with the increasing demands of multi-source data, and prone to alert fatigue due to high false-positive rates. More recent studies have explored the integration of clinical ontologies with time-series analysis to improve alert precision [4]; however, these approaches still lack the responsiveness required for real-time applications.

A foundational advancement in this space was the use of stream processing engines, such as Apache Storm and Spark Streaming, in health informatics. Hsieh et al. [5] demonstrated how real-time health monitoring systems based on Apache Storm could process continuous physiological signals from ICU devices, triggering alerts within milliseconds of data pattern recognition. These engines, initially designed for financial and e-commerce applications, have since been adapted to handle the complexity of healthcare data, such as HL7 messages and FHIR resources. However, integrating these platforms into regulated healthcare environments remains a technical and organizational challenge.

The implementation of event-driven architecture in healthcare was further elaborated by Kuo et al. [6], who proposed a publish-subscribe model for continuous EHR updates. Their study showcased how decoupling event producers (e.g., patient monitors) from consumers (e.g., clinical dashboards or notification systems) could enhance system scalability and resilience. By leveraging message queues such as Apache Kafka, real-time data ingestion from multiple clinical endpoints became feasible. Kafka's append-only log structure and distributed architecture also enabled data replay, a critical feature for retrospective diagnosis and audit trails. Similarly, Ahmad et al. [7] explored real-time sepsis detection by embedding event recognition rules into streaming pipelines. Their prototype detected septic conditions nearly an hour earlier than existing hospital alerting systems, underscoring the potential for improving patient outcomes. In addition to technical implementations, several frameworks have been proposed for integrating real-time analytics with clinical workflows. The Mirth Connect engine was frequently cited as an intermediary system for routing, transforming, and filtering HL7 and DICOM messages in near-real time [8]. Integration patterns such as Enterprise Integration Patterns (EIP) [9] were employed to manage message transformation, sequencing, throttling, and exception handling—key concerns in processing asynchronous healthcare events. Despite these advances, the heterogeneity and latency sensitivity of healthcare data still pose substantial obstacles.

Developments also influence real-time clinical decision support in semantic interoperability. The work of Zhang et al. [10] emphasized FHIR (Fast Healthcare Interoperability Resources) as a vital enabler of dynamic data exchange in real-time systems. FHIR's modular resource structure and RESTful APIs make it particularly well-suited for stream-oriented applications, especially when coupled with real-time messaging brokers. This compatibility facilitates integration with cloud-based processing pipelines, enabling elastic scaling and remote collaboration, as evidenced in multiple pilot deployments across hospital systems.

The concept of continuous decision support, where patient context is evaluated dynamically instead of at fixed intervals, has also gained attention. A study by Ghassemi et al. [11] examined the use of continuous learning systems embedded within real-time monitoring infrastructure. Their architecture dynamically updated predictive models based on streaming data, thereby improving accuracy over time and adjusting decision thresholds according to patient-specific trends. This approach is critical for adaptive alerts and feedback mechanisms in future-ready CDSS. The body of literature reviewed highlights the trajectory of healthcare analytics from static, rule-based systems toward dynamic, real-time, and event-driven architectures. While platforms such as Apache Kafka and Apache Flink have proven valuable in enabling scalable, low-latency processing, integration with clinical workflows, standard compliance (e.g., HIPAA, HL7, FHIR), and alerting precision remain focal areas for ongoing research. This paper contributes to the literature by synthesizing these advancements into a unified, deployable architecture tailored for modern healthcare environments, demonstrating measurable improvements in clinical response times and system adaptability.

3. Methodology

The approach used in this work focuses on developing, deploying, and evaluating a real-time healthcare event processing model that incorporates stream analytics through the lens of clinical decision support functions, with actions taken in response to these events. The aim was to build an end-to-end event-driven architecture that enables the ingestion, processing, and reaction to healthcare data arriving in different streams within milliseconds. This was achieved through the alignment of technical components with clinical processes, cooperation with data protection regulations, and the assessment of the system's performance in the context of time-critical decision-making.

The proposed architecture is based on an event-driven framework that supports asynchronous, real-time communication among entities. At its heart is a distributed message broker – Apache Kafka – that underpins the processes for data ingestion and transport. Although we selected Kafka for this purpose, as a high-throughput, low-latency framework that can process millions of events per second [11], we believe it is suitable for healthcare as a control environment where data flows from numerous sources.

These origins could be bedside monitors, IoT wearables, laboratory systems, medication dispensers, and hospital systems. Every device or system sends clinical events, such as abnormal vital signs, medication orders, and lab results, to Kafka topics in real-time.

After the events are ingested into Kafka, stream processing engines are used to consume them, primarily Apache Flink and Apache Storm, which vary in terms of complexity and latency sensitivity for the task. Apache Flink was also employed for complex event pattern recognition, such as sepsis, by comparing heart rate variability, increased white blood cell count, and body temperature behavior. For use cases that required lightweight processing (low latency), Apache Storm was used, with examples including the identification of a decrease in oxygen saturation and a sudden change in heart rate. The stream processors identify and analyze events in continuous time based on predefined event streaming patterns, which encode temporal rules and medical thresholds specified by clinicians and domain experts. These patterns are codified into streaming logic that captures not only the values of the input events but also the context in which they arrive, along dimensions of time, order, and the concomitancy of other clinical events.

The resultant data is well mediated into a near-real-time clinical alerting engine with multi-modal alerting capabilities. Once an alert is triggered, one function is to determine the appropriate clinical response (i.e., who should be alerted) and to route the alert to the doctor or nurse ordering CC75 at the time, as per escalation protocols. For instance, an emergency alert from ICUs would be delivered to a specialized team. In contrast, an urgency alert from a post-surgery ward would be dispatched to the dedicated care nurse. Notifications are sent through multiple paths: web UI, encrypted native mobile platforms, SMS, or as direct integration to the EHR system through HL7 messages or FHIR APIs. The engine is also conducive to a closed-loop feedback process, whereby clinicians can easily identify instances where an alert was deemed valid or a false positive, and use that information as input to improve future alert accuracy.

The architecture was hosted in a HIPAA-compliant cloud, and all communications were end-to-end encrypted using TLS for data in transit and AES-256 for data at rest. RBAC was applied ubiquitously at every interface by OAuth 2.0 and SAML-based single sign-on to support security and auditing. The system was created to adhere to healthcare messaging standards, ensuring interoperability with protocols such as HL7 v2, FHIR R4, and DICOM whenever possible. Integration with hospital systems was achieved using Mirth Connect, serving as an interface engine for message transformation and routing.

Performance assessment was conducted through simulation using both synthetic and real datasets. We set up a test environment that deployed Kafka, Flink, and related services as containers in Kubernetes. Healthcare cases were represented as a process of early sepsis detection in patients, post-operative complication detection, and emergency triage alerting, with data modeled by representing streaming vital signs and lab results. Metrics such as end-to-end latency, system throughput, alert accuracy, and response time were collected and analyzed. Benchmarking results showed that the IBI-35 could process more than 200,000 events per minute with an average latency of less than 250 milliseconds per event. They demonstrated the ability to trigger more than 90% accurate, clinically actionable alerts. In addition, the feedback loop enabled the system to decrease the number of false positives by 30% for 14 14-day simulation period. The approach taken in this work is a holistic one for developing scalable, secure, and clinically relevant, real-time stream processing systems for healthcare. Its combination of medically motivated technical decisions and operational constraints provides a promising model for rethinking the deployment of data-driven decision support in today's hospitals.

4. Results

The implementation of the proposed real-time event processing architecture for clinical decision support was evaluated across a series of simulated and real-world healthcare data scenarios to determine its effectiveness in improving clinical response times, system throughput, and alert precision. The results indicate a significant enhancement in system responsiveness and decision support reliability when compared to traditional batch-based processing models.

To assess performance, the architecture was deployed in a containerized Kubernetes environment using synthetic EHR data streams that simulated real-time patient monitoring in critical care units. Each synthetic stream included multivariate data such as heart rate, blood pressure, oxygen saturation (SpO₂), temperature, white blood cell count, and clinical notes. Events were generated at one-second intervals to reflect realistic ICU and emergency department conditions. A total of 10 million events were processed across test scenarios that replicated three use cases: early detection of sepsis, postoperative complication monitoring, and emergency room triage.

The first key outcome was a measurable reduction in clinical response time. Compared to a benchmark batch-processing system that scanned clinical data every 15 minutes, the event-driven system demonstrated a 40% decrease in response latency. For sepsis detection, alerts were triggered an average of 8.5 minutes earlier using stream analytics, a critical improvement given that

early intervention in sepsis can significantly reduce mortality rates. In the post-operative care simulation, the system detected anomalies such as hypoxia and irregular heart rates approximately 6 minutes earlier than static CDS dashboards. Emergency triage use cases showed even greater improvement due to rapid real-time ingestion and alert delivery within 1 minute of the first critical vital deviation.




Fig 2: Comparison of Clinical Response Times for Batch vs Real-Time Systems

The system's throughput was another primary performance metric. Benchmarks showed that the pipeline processed over 3,500 events per second per node, scaling linearly across a five-node cluster without bottlenecks. This scalability was maintained even under high-load conditions and stress testing. Kafka handled ingestion from up to 1,200 concurrent event producers with stable latency. Apache Flink performed continuous evaluation of complex event patterns with sub-second decisioning latency (mean latency of 243 milliseconds per event). These figures demonstrate the architecture's viability in live hospital environments, where thousands of devices and applications generate continuous data.

Alert accuracy was measured in terms of precision, recall, and false positive rate. In sepsis detection, the system achieved a precision of 91.2% and a recall of 88.7%, significantly outperforming the baseline threshold-based alerting system. In the postoperative monitoring simulation, alert precision improved by 28%, mainly due to the contextual filtering provided by pattern-based event correlation. The feedback loop mechanism integrated into the clinician dashboard enabled continuous learning from user feedback, resulting in a 30% reduction in false alerts within 14 days of system use. This addresses the critical issue of alert fatigue, a common problem in conventional CDSS that contributes to clinician burnout.

Further, system reliability was evaluated through uptime monitoring and error tracking. Over a 72-hour continuous test cycle, system uptime was recorded at 99.96%, with only three minor processing interruptions caused by Kafka topic misconfiguration, which were automatically resolved by the platform's failover mechanisms. No data loss was observed during rebalancing or broker failures, validating the robustness of the system's message durability and distributed coordination.

Interoperability and EHR integration were assessed through HL7 message parsing and FHIR resource mapping. The system demonstrated seamless compatibility with Epic and Cerner systems via Mirth Connect, confirming practical deployment potential without requiring vendor-specific adaptations. This interoperability ensures that the event processing engine can be embedded into existing hospital workflows, further lowering adoption barriers.

Overall, the results validate that the proposed real-time stream analytics framework enhances decision-making capabilities in clinical environments by improving responsiveness, reducing time to intervention, minimizing false positives, and delivering consistent, scalable performance under real-time constraints. These improvements make a compelling case for the shift from batch processing to event-driven systems in modern healthcare infrastructure.

5. Discussion

The results obtained from implementing and evaluating the proposed real-time stream analytics architecture for healthcare reveal significant insights into both the opportunities and challenges associated with deploying event-driven systems in clinical settings. The demonstrated 40% improvement in clinical response time, along with substantial gains in system scalability and alert accuracy, underscores the transformative potential of adopting real-time event processing in critical care environments. However,

these outcomes must be interpreted within the broader context of healthcare operations, infrastructure readiness, and clinical workflow dynamics.

One of the most notable implications of this research is the operational efficiency gained through immediate alerting and data-driven responsiveness. The system's ability to detect clinically significant patterns in streaming data and trigger timely interventions addresses one of the most persistent limitations of traditional CDSS—latency. In environments such as intensive care units and emergency departments, delays in detecting early signs of deterioration can result in preventable complications or mortality. By minimizing these delays, the proposed architecture directly contributes to quality improvement and patient safety objectives that healthcare providers strive to achieve.

Another critical insight is the role of event streaming patterns in enhancing the accuracy of decision support. Rather than relying on static thresholds that often fail to account for contextual factors, the use of dynamic temporal and multi-signal correlations allows the system to identify complex conditions with higher fidelity. This capability not only improves the relevance of alerts but also reduces false positives—a primary contributor to clinician alert fatigue. The feedback loop integrated into the system further enables continuous tuning of alerting logic based on user validation, supporting the transition from a static to a learning decision support system. Such adaptive capabilities align with broader trends in precision medicine and personalized care. Scalability and reliability, as observed in the results, are critical for deployment in large healthcare systems where thousands of concurrent data streams must be processed without performance degradation. The ability to maintain low latency and high throughput under stress confirms the viability of distributed stream processing frameworks such as Apache Kafka and Apache Flink in mission-critical healthcare applications. Furthermore, the architecture's modularity and containerized deployment model enable elasticity in cloud or hybrid environments, making it adaptable to a variety of hospital infrastructures, ranging from urban tertiary centers to rural health networks.

Despite these advantages, the research also exposes important challenges that must be addressed for successful implementation in real-world scenarios. Data standardization remains a substantial barrier, particularly in environments where EHR data formats are inconsistent or proprietary. Although HL7 and FHIR standards provide a foundation for interoperability, variations in vendor implementation can introduce integration complexity. Achieving seamless bidirectional communication between the stream processing system and clinical applications requires a robust interface engine and skilled configuration management.

Security and compliance are also paramount, especially given the sensitivity of real-time clinical data. The architecture must maintain rigorous HIPAA compliance, which includes data encryption, access control, audit trails, and secure transmission. Any system interruption or data breach in such a framework could have direct implications for patient safety. While the current implementation incorporated security best practices, ongoing monitoring and third-party audits would be necessary to ensure long-term compliance.




Fig 3: Improvement in Alert Precision Over 14 Days

Another challenge lies in the organizational readiness for adopting event-driven systems. Real-time CDS requires not just technological adaptation but also cultural and procedural shifts in how clinicians interact with data. Hospitals must provide adequate training, revise workflows to integrate alerting into decision-making pathways, and ensure that clinicians trust the system's outputs. The initial learning curve and change management efforts can be significant, especially in environments where digital health maturity is still in its early stages of development.

Ethical considerations also surface when real-time alerts influence clinical decisions. Overreliance on automated alerts, particularly in high-stakes decisions such as medication administration or triage prioritization, could result in unintended

consequences if the system misclassifies events. Therefore, the system must be designed to complement, rather than replace, human judgment, with clear guidelines for alert verification and override mechanisms.

In terms of broader impact, this research contributes to the growing discourse on real-time health informatics and its role in value-based care delivery. By demonstrating tangible benefits such as reduced time-to-intervention and improved alert relevance, it supports the business case for investing in modern analytics infrastructure. Moreover, as healthcare shifts toward more data-driven and patient-centered models, real-time stream analytics can serve as a foundation for advanced capabilities such as continuous risk scoring, real-time cohort monitoring, and predictive triage systems.

While the proposed event-driven architecture offers substantial benefits in terms of responsiveness, scalability, and clinical impact, its success hinges on thoughtful implementation, interoperability planning, and alignment with clinical workflows. Future work should explore the integration of AI-driven predictive analytics with streaming data and evaluate long-term patient outcomes across diverse care settings to establish further the value proposition of real-time clinical decision support systems.

6. Conclusion

The research presented in this paper validates the growing necessity and feasibility of integrating real-time stream analytics into clinical decision support systems (CDSS) to enhance the responsiveness and efficiency of healthcare delivery. By shifting from traditional batch-oriented data processing models to an event-driven architecture, this study addresses a critical gap in time-sensitive medical environments where every minute can significantly affect patient outcomes. The proposed framework, designed using open-source stream processing platforms such as Apache Kafka and Apache Flink, successfully demonstrates how healthcare data can be processed and acted upon in real-time, rather than as static records.

A key contribution of this work is the incorporation of healthcare-specific event streaming patterns, enabling intelligent and context-aware decision-making. These patterns capture the nuanced relationships between various clinical parameters and time-series events, enabling the early detection of critical conditions such as sepsis, hypoxia, or post-surgical deterioration. Through continuous data ingestion and pattern recognition, the system supports proactive intervention and prioritizes alerts for clinical teams based on real-time data streams. This approach fundamentally improves upon legacy CDSS that rely on retrospective analysis or periodic polling, which often leads to delayed responses and reduced clinical utility.

The results obtained through simulation and testing confirm the system's effectiveness. The framework achieved a 40% reduction in average clinical response time across multiple use cases, including emergency triage, post-operative monitoring, and ICU alerting. Additionally, improvements in alert precision and a substantial drop in false positives reflect the system's ability to deliver meaningful notifications, mitigating the risk of alert fatigue. These improvements have implications not just for patient safety and satisfaction, but also for clinician trust and adoption of digital decision support tools. The architecture's scalability and interoperability were also demonstrated through successful integration with multiple EHR systems and standardized data exchange protocols such as HL7 and FHIR. Its deployment in a containerized, HIPAA-compliant environment further establishes its readiness for real-world implementation in diverse healthcare infrastructures, whether on-premises, in the cloud, or across hybrid networks. This flexibility makes it suitable for institutions of varying sizes and digital maturity levels, from urban academic medical centers to regional clinics and remote care settings.

Despite these strengths, the study also identifies challenges that must be addressed before full-scale implementation. These include the need for consistent data standardization, robust security protocols, seamless integration with existing hospital workflows, and adequate training for clinicians. Cultural acceptance within clinical environments will be essential for the successful adoption of real-time systems, as these necessitate changes in how healthcare professionals consume, interpret, and act upon information. Additionally, long-term evaluation of patient outcomes and system performance in live clinical settings will be critical to establishing sustained value.

Looking ahead, the convergence of real-time stream analytics with artificial intelligence (AI) and machine learning (ML) opens promising avenues for future research and system evolution. By embedding predictive models into the event stream, it becomes possible to anticipate adverse events before any clinical signs are apparent, transforming decision support from reactive to anticipatory. Furthermore, federated data architectures could extend this real-time intelligence across institutions, supporting public health monitoring and collaborative clinical research without compromising patient privacy.

The proposed real-time healthcare event processing architecture represents a significant advancement in digital health infrastructure. It aligns with global objectives for improving clinical outcomes, reducing care delays, and optimizing resource

utilization. Through strategic deployment and ongoing refinement, such systems can help shape the future of responsive, intelligent, and patient-centric healthcare delivery. This paper not only lays a technical foundation for that vision but also contributes a practical, standards-based model that healthcare organizations can adapt and build upon to meet the demands of modern clinical practice.

References

- [1] D. W. Bates et al., “Ten commandments for effective clinical decision support: making the practice of evidence-based medicine a reality,” *J. Am. Med. Inform. Assoc.*, vol. 10, no. 6, pp. 523–530, Nov. 2003.
- [2] A. Wright et al., “Clinical decision support capabilities of commercially available clinical information systems,” *J. Am. Med. Inform—Assoc.*, vol. 18, no. 5, pp. 387–392, Sept. 2011.
- [3] A. Wright and D. F. Sittig, “A four-phase model of the evolution of clinical decision support architectures,” *Int—10*, pp. 641–649, Oct. 2008.
- [4] H. Haug et al., “Combining health information models and ontologies for clinical decision support,” *Stud. Health Technol. Inform.*, vol. 136, pp. 265–270, 2008.
- [5] C.-Y. Hsieh, J.-Y. Chung, and C.-F. Lin, “Development of a real-time health monitoring system using Apache Storm,” in *Proc. IEEE Int. Conf. Healthcare Informatics (ICHI)*, New York, NY, 2017, pp. 494–497.
- [6] A. M. Kuo, “Opportunities and challenges of cloud computing to improve health care services,” *J. Med. Internet Res.*, vol. 13, no. 3, p. e67, 2011.
- [7] M. Ahmad et al., “Real-time sepsis prediction using streaming analytics,” in *Proc. ACM Int. Conf. Bioinformatics, Computational Biology, and Health Informatics*, 2018, pp. 179–186.
- [8] NextGen Healthcare, “Mirth Connect: Open-source healthcare integration engine,” [Online]. Available: <https://www.nextgen.com>. [Accessed: Oct. 2019].
- [9] G. Hohpe and B. Woolf, *Enterprise Integration Patterns: Designing, Building, and Deploying Messaging Solutions*, Addison-Wesley, 2003.
- [10] Y. Zhang et al., “FHIRChain: Applying blockchain to securely and scalably share clinical data,” *Comput. Struct. Biotechnol. J.*, vol. 16, pp. 267–278, 2018.
- [11] M. Ghassemi et al., “A multiscale dynamic modeling approach for real-time prediction of sepsis,” *IEEE J. Biomed. Health Inform.*, vol. 22, no. 2, pp. 453–462, Mar. 2018.