

(REVIEW ARTICLE)



HYBRID STREAMING AND BATCH INTELLIGENCE FRAMEWORK FOR REAL-TIME HEALTHCARE ANALYTICS ON DISTRIBUTED CLOUD DATA PLATFORMS

Sreenivasa Reddy Vemareddy *

Belhaven University, 1500 Peachtree St, Jackson, MS 39202, USA.

* Corresponding Author

World Journal of Advanced Engineering Technology and Sciences, 2026, 20(03), 157–165

Publication history: Received on 08 June 2026; revised on 16 September 2026; accepted on 18 September 2026

Article DOI: <https://doi.org/10.30574/wjaets.2026.20.3.0361>

Copyright © 2026 Author(s) retain the copyright of this article. This article is published under the terms of the Creative Commons Attribution License 4.0

ABSTRACT

Hybrid streaming and batch intelligence architecture is becoming the backbone of the clinical data platform as it strives to combine low-latency event processing with retrospective learning from very large volumes of longitudinal clinical data. The review examines peer-reviewed journal publications from 2015 to the present on cloud, fog, Internet of Things (IoT), and critical-care analytics architectures and their application to distribute healthcare intelligence. The reviewed evidence indicates that streaming components can improve responsiveness in physiological monitoring, seizure detection, prediction of acute kidney injury, and monitoring of deterioration in intensive care. For cohort construction, model training, calibration, population evaluation and auditability, batch components are still needed. The literature, however, has been split into infrastructure studies, clinical prediction studies, and database-based benchmarking. The major gaps found are limited treatment of concept drift, lack of reporting of end-to-end latency, limited institutional validation and incomplete models of governance for the continuous updating of clinical intelligence. Architectures are therefore required that integrate paradigms of streaming inference, batch retraining, privacy-preserving federation, and interpretable decision support, all in distributed cloud data platforms, with operational monitoring.

Keywords: Batch Intelligence; Cloud Healthcare Analytics; Clinical Time Series; Distributed Data Platforms; Real-Time Streaming

1 INTRODUCTION

Hybrid streaming and batch intelligence are becoming essential for real-time healthcare analytics as clinical data is now delivered via a variety of sources – bedside monitors, wearable sensors, laboratory interfaces, imaging archives, electronic health records and remote patient-monitoring services. The issue concerns not only data volume but also relating event-time processing to learning from historical data. Streaming pipelines can be used to facilitate near-real-time alerting and state estimation, and batch pipelines can be used to facilitate feature construction, retrospective cohort analysis, model training, model evaluation and model governance. Analytic pipelines capable of transforming large, complex, and time-sensitive data into meaningful information to guide action are more important than data storage alone [1]. The same holds true for cloud, Internet-of-Things (IoT), and platform-based health research [2].

The academic importance of hybrid streaming/batch intelligence solutions lies at the intersection of medical informatics, distributed systems, machine learning, data governance, and clinical decision support. Architectures must be able to support continuous telemetry without losing the statistical power of the longitudinal electronic health records (EHRs) that are created by them. While the accuracy of the algorithms is undoubtedly important, studies on healthcare big data analytics have found that the infrastructure, data management, personnel capabilities and decision process also play a role in the capability to use analytics [3]. This is because if data latency, data missingness or alert fatigue are not addressed, then a real-time model an impact in the clinical that is hosted on a distributed, cloud-based platform can have space. The structure is thus not only a prediction engine, but also a socio-technical platform that must be assessed as such.

The healthcare IoT literature highlights this challenge. As streams become more complex and more sensors are added to the environment, streams can become too much for a single process to manage, especially when using wearable sensors or home monitoring systems that are not part of the hospital network [4]. These architectures allow data to be preprocessed in fog and edge gateways, persisted longitudinally in the cloud, and analyzed in batch for population-scale learning. Most of the papers, however, continue to refer to these layers separately. Infrastructure papers tend to be more concerned with scalability, while clinical machine-learning papers tend to be more concerned with discrimination metrics, but without complete information on stream ingestion and failures, update frequency, and platform governance. The result is that there is a gap between clinical goals for real-time healthcare analytics and the operational and distributed cloud data platforms.

The review critically analyzes the peer-reviewed journal articles published in the last ten years (2015-2023) that are directly related to the aforementioned hybrid streaming and batch intelligence for real-time analytics on distributed cloud data platforms in the healthcare domain. The phrase in the title is not yet a "journal field convention" so the evidence base is spread across areas focused on cloud-assisted health monitoring, fog-enabled healthcare Internet-of-Things, smart-healthcare big-data ecosystems, scalable electronic health record learning, continuous clinical prediction and critical-care time-series benchmarking. This review aims to clarify the conceptual differences between streaming and batch, review reported results, and outline the challenges and specific recommendations for future platform research.

2 LITERATURE REVIEW

The literature on hybrid streaming and batch intelligence can be viewed as three strands of real-time healthcare analytics. The first one is about cloud distributed, fog, and Internet-of-Things (IoT) platforms for health monitoring. Forkan et al. [5] proposed BDCaM, a context-aware framework for personalized monitoring and cloud-based knowledge discovery from sensor and contextual data in a context-aware healthcare system. Hossain and Muhammad introduced an Industrial Internet of Things (IIoT) framework that showed the benefit of distributed processing for the health monitoring of multimedia and physiological data streams, with cloud support [6]. Rahmani et al. extended this idea to smart e-health gateways, which can cut down latency, allow local pre-processing and secure bandwidth prior to analytics in the cloud [7]. Manogaran et al. presented a safe smart-healthcare monitoring and alerting system based on IoT data flow and big data storage and intelligence security [8]. Further work by Alhussein et al. demonstrated cognitive IoT-cloud integration for epileptic seizure detection and monitoring, illustrating domain-specific inference in a distributed monitoring system [9].

The second thread is on predictive healthcare analytics using large, heterogeneous datasets. Sahoo et al. explored the future health condition prediction using healthcare big-data for the assessment of health condition over future years and found that the big analytic model can combine various clinical parameters and perform long-term risk assessment [10]. Rajkomar et al. showed that deep learning models can use standardized EHR representations to predict multiple clinical outcomes [11]. Tomašev et al. [12] reported the continuous prediction of AKI, an important case in which the inference has to be repeated over time. Similarly, Hyland et al. studied early prediction of circulatory failure in an intensive care unit and showed that the use of temporally aware modeling improves the performance of tasks related to the prediction of circulatory failure [13]. Based on these studies, it is clear that a batch-trained model can produce clinically meaningful predictions, but using the model to power a real-time intelligence application requires reliable inputs at stream time, consistent timestamp handling, and it requires calibration surveillance and alert integration.

The third strand has to do with benchmarking and database foundations. Harutyunyan et al. provided multitask benchmarks, modeling time series clinical data for mortality task, decompensation task, length-of-stay task and phenotype task, such that the performance of these tasks can be compared based on the same critical-care data structure [14]. Che et al. demonstrated that patterns of missingness can themselves provide predictive information in multivariate clinical time series [15]. The MIMIC-III database was introduced by Johnson et al. as a large critical-care database that can be used for retrospective batch experimentation on high-resolution hospital data [16]. Pollard et al. created a new eICU Collaborative Research Database as a preliminary step to multi-center evaluation of critical-care analytics [17]. In

a review of deep learning in EHRs, Xiao et al. pointed out the limitations of interpretability, data quality, temporal representation and clinical deployment [18]. The articles are essential to the adoption of these hybrid intelligence frameworks as they need to use batch data to train and validate and then use real-time clinical data to determine whether these static benchmarks will be applicable in real platforms. Table 1 gives an overview of typical studies directly informing the architecture, the intelligence layer or the evaluation base of hybrid streaming and batch healthcare analytics.

Table 1: Summary of key findings

Ref	Focus	Key Findings
[5]	Context-aware assisted healthcare on cloud infrastructure	Personalized monitoring benefits from combining sensor context, patient profiles, and cloud-scale knowledge discovery.
[6]	Cloud-assisted IIoT health monitoring	Cloud support enables intensive data handling for physiological and multimedia health streams, although latency governance remains underdeveloped.
[7]	Fog-enabled e-health gateways	Gateway-level preprocessing can reduce latency and bandwidth dependence before cloud analytics.
[8]	Secured IoT and big-data smart-healthcare architecture	Monitoring and alerting require data security, distributed storage, and analytic integration rather than isolated sensors.
[9]	Cognitive IoT-cloud seizure monitoring	Disease-specific inference can operate across IoT-cloud layers, but external clinical validation remains limited.
[10]	Healthcare big-data prediction	Longitudinal prediction benefits from large data integration, although operational streaming requirements are weakly specified.
[11]	Scalable EHR deep learning	Standardized EHR representations can support accurate prediction across tasks, but deployment infrastructure is separate from model training.
[12]	Continuous acute kidney injury prediction	Repeated temporal inference can anticipate acute kidney injury, making it relevant to streaming clinical intelligence.

The table highlights that studies focused on platforms tend to focus on data movement, storage, location of gateways and security, while clinical prediction studies focus on model performance. The mismatch between these focus points is the research gap – little has been done to evaluate end-to-end hybrid systems for which the following steps are measured: Streaming ingestion, Batch retraining, Distributed storage, Clinical alerting, Outcome evaluation.

3 METHODOLOGY

There are four major methods that have been presented in the literature reviewed. Cloud-based methods move storage, analytics and decision support to distributed systems that are concentrated. This allows scaling, longitudinal data aggregation, and batch training but can lead to latency, network dependency, and potential privacy risks associated with continuous streams of raw physiological data. Second, fog and edge-gateway technologies move the processing, compression or anomaly screening, and/or anomaly finding/feature extraction closer to the sensor and care setting. There are resource constraints and problems of updating models reliably and getting them in sync with local action; but these designs make the system more reliable by reducing transmission load and enable quicker local action to be taken. Thirdly, Internet of Things streams, security controls, distributed databases and alert engines are examples of components in the big data models for the ecosystem. They are similar to the "hybrid streaming" and "batch intelligence" concepts, although, conceptually, they are the same, but the assessment is not always clinical and is often still architectural [8], [9]. Fourth, models are trained from the retrospective EHR and critical care studies, and the data used in the platform is typically abstracted from the data and only reported based on discrimination or early warning performance.

The hybrid streaming and batch intelligence framework can be conceptualized as a multilayered system architecture. Streaming layers are responsible for ingesting sensor events, bedside observations, orders, medication events and lab changes and device messages. The longitudinal features are stored in a feature store within the batch layer, retrain feature models, evaluate drift, to audit predictions, and provide population analytics. Distributed cloud platforms then provide all the storage, computation, identity management, access control and monitoring. Complementarity is the

methodological power of this configuration: timely predictions from stream processing, historical context and model governance from batch processing. The biggest challenge is that of integration. Even a highly accurate batch model can cause harm if a live data schema changes, device clocks drift, or clinical workflows disable alerts. Rapid model updating can introduce instability if it is not supported by robust retrospective calibration and validation. Table 2 shows comparison of the approaches found in the literature reviewed and how they can be incorporated into a hybrid framework.

Table 2: Method comparison

Ref	Method	Strengths	Limitations
[5]	Context-aware cloud knowledge discovery	Links patient context with monitoring and personalized inference.	Limited reporting of clinical deployment latency and update governance.
[7]	Fog gateway architecture	Reduces bandwidth burden and supports local preprocessing.	Gateway heterogeneity complicates maintenance and validation.
[8]	Secured IoT-big-data ecosystem	Treats monitoring, alerts, storage, and security as connected components.	Clinical outcome evaluation is less mature than architectural design.
[11]	Scalable EHR deep learning	Demonstrates learning from large longitudinal records across tasks.	Live stream ingestion and operational alert safety are not central.
[12]	Continuous temporal prediction	Aligns prediction with repeated clinical risk estimation.	External deployment, drift response, and infrastructure metrics need deeper reporting.
[14]	Critical-care multitask benchmarking	Enables comparative evaluation across standardized time-series tasks.	Retrospective benchmarks do not fully represent distributed live operation.




Figure 1: Hybrid streaming and batch intelligence framework for distributed cloud healthcare analytics

The concept map in Figure 1 illustrates the various layers in the distributed cloud healthcare analytics, and how they are interconnected. This is important because the studies reviewed tend to study only part of this architecture, and the process of ingesting events to providing clinical feedback must be coordinated to meet real-time healthcare intelligence requirements. The diagram shows a multi-layered architecture which facilitates the streaming of events to support low-latency inference, and batch data feeding to support training, recalibration, audit and population analytics. Data/source basis: author-coded conceptual interpretation of platform and clinical patterns of analytics reported in [5]-[15].

4 DISCUSSION

The most evident conclusion from reviewing the literature is that the paradigm for using analytics is moving from a single central point (retrospective analytics) to a distributed intelligence paradigm (distributed intelligence), with intelligence distributed across the edge, fog, cloud and clinical workflow layers. Cloud-supported research paved the way for the possibility of gathering, storing and transmitting huge amounts of data generated by sensors for analysis in the field of healthcare [5,6]. Fog-gateway work is being developed on top of this concept, as in the case of health streams, it is vital to process them in advance and take the right action locally [7]. In smart healthcare ecosystem studies, security and alerts are integrated, as well as security requirements are embedded into the design of data-platforms, including its confidentiality, access control and operational resilience [8] [9]. But these studies are mainly based on architecture, and usually refer to feasibility, conceptual integration or prototype behavior, but not necessarily clinical endpoints. Thus, there is more evidence that infrastructure is adequate than there is evidence that patients benefit.

More literature is available on clinical prediction with more outcome-oriented evaluations, and less transparency of the platforms. Data from large-scale longitudinal clinical EHR deep-learning studies demonstrate that high-dimensional longitudinal data could be applied to predict a range of clinical tasks with high accuracy, as shown in [11]. Recent studies have demonstrated the feasibility of dynamic risk estimation to predict deterioration prior to its conventional identification, including in the case of continuous acute kidney injury prediction and circulatory failure prediction [12, 13]. Multitask critical-care benchmarks, as well as the use of missingness-aware recurrent models, also illustrate the influence of temporal structure and irregular sampling and observation patterns on performance [14, 15]. The results provide evidence of the intelligent aspect of hybrid streaming and batch architecture. However, there is limited evidence on how predictive analytics can be used when deploying cloud capabilities in real time. A data set that has been deliberately curated for its similarity to current data does not guarantee safe inference when the data is ingested after a delay, the data interface is down, workflow operates differently at each unit, or the ordering pattern is different from that in the current data set. The results presented or extrapolated from the literature that are similar to hybrid streaming and batch healthcare analytics are listed in Table 3. The intent is not to directly rank models, but to highlight areas of evidence and operational gaps.

Table 3: Results comparison

Ref	System / Context	Metric / Basis of Comparison	Outcome
[7]	Fog-enabled e-health gateway	Architectural latency and bandwidth logic	Local preprocessing reduces dependence on cloud-only routing for time-sensitive monitoring.
[9]	IoT-cloud seizure monitoring	Disease-specific event detection context	Cognitive cloud integration supports targeted inference, but broader clinical validation remains limited.
[11]	Multi-task EHR deep learning	Predictive accuracy across EHR tasks	Scalable learning improves prediction from longitudinal records but leaves live-platform behavior partly unresolved.
[12]	Continuous AKI prediction	Early warning across repeated time points	Continuous inference can anticipate acute kidney injury, supporting streaming clinical surveillance.
[13]	ICU circulatory failure prediction	Early deterioration prediction	Dynamic machine learning improves anticipation of circulatory failure in intensive care contexts.
[14]	Critical-care benchmark tasks	Mortality, decompensation, length-of-stay, phenotyping	Standardized benchmarking improves comparability but remains retrospective.
[15]	Missingness-aware RNN time-series modeling	Handling irregular multivariate clinical time series	Missingness patterns improve predictive modeling, but live data-quality shifts require monitoring.
[17]	Multi-center ICU database	Cross-site critical-care data availability	Multi-center data improves external evaluation potential for distributed healthcare intelligence.

There are a number of comparisons that can be made. The solution offered most significantly from an architectural point of view is the presence of the fog and cloud health-monitoring studies, where the latency, storage, security and distributed processing are highlighted [5]–[9]. EHR and ICU time-series studies give the best clinical contribution, where

large clinical datasets are used to investigate risk prediction [11]–[15]. The problem is that these bodies of work do not often overlap within the same evaluation design. A mature hybrid streaming intelligence solution would evaluate ingestion latency, event-time correctness, throughput, model accuracy, model calibration, clinical alert burden, security controls, and update governance. A significant proportion of the reviewed articles report only a subset of the metrics required for comprehensive end-to-end evaluation.

Secondly, there is a wealth of literature that confirms that in real-time healthcare analytics, batch intelligence is a must. Here, batch operations decide on the training population, select variables to be included, impute or encode missing data, set risk thresholds, validate by sub-population/subgroup and monitor the subsequent activities after deployment. Large critical-care databases are still central to research into real-time analytics, allowing for retrospective investigations and providing common examples for the temporal modelling research (as done in MIMIC-III and eICU) [16, 17]. Batch datasets also create the illusion of stability, however. It's a simplified version with some of the interface disruptions, delayed charting, device disconnects, timestamp errors and policy changes in measurement frequency that happen in the real clinical stream. Hybrid frameworks, however, also call for batch retraining and batch auditing of the framework, and ongoing observability of the platform to detect when the live stream is no longer comparable to the training distribution.

Missingness is not merely a nuisance variable in clinical time-series analysis. The fact that multivariate time series are subject to non-uniform observation times has been leveraged in the past for recurrent modeling [15] and benchmarks have been conducted to show that the ability of an ICU is affected by these non-uniform schedules [14]. The implication of this finding is for hybrid streaming and batch configuration. Stream processors should not only maintain the timing of the observation, but the absence of data, as well as the frequency of the observation, and not inevitably remake all the data to a coarse grid of too fixed a resolution. Values should be stored in feature stores (batch), together with data lineage, delay, source system, and acquisition context. This metadata is critical when inferring clinical improvement or deterioration in real time, to prevent errors due to incorrect reading of the signal that a patient is improving when they are actually reducing measurement intensity, or that they are deteriorating when they are actually reducing measurement intensity.

Fifth, there has been inconsistent information in literature about centralization. Cloud computing provides scale, elasticity of storage and integration with other institutions [2] [5] [6] while fog and edge computing provide local responsiveness and low latency. In a typical clinical machine-learning study, a centralized location is used for collecting retrospective datasets which are then filtered and harmonized [11] – [16] before being transferred to and analyzed on a clinical machine-learning platform. However, today, health systems are increasingly required to operate across distributed locations (e.g., hospitals, community settings and home monitoring). Any viable framework should thus not endorse a binary approach of cloud versus edge. For batch training: use cloud layers, for registry management, longitudinal storage, and cross-site audit: use edge or fog layers, for signal screening, local failover, and immediate event handling. Without building isolated intelligence, there is no easy way to coordinate model versions, security policies, threshold changes and audit logs across these layers.

Fifth, however, the interpretability and workflow integration are still less developed than the predictive performance. According to reviews [18], the successful implementation of EHR deep learning is constrained by the black box characteristics of the model, the differences in data representations, and the limited clinical explainability. But interpretability has a time component when it comes to real-time healthcare analytics. Apart from the reasons for prediction being raised, a clinician would want to be aware of new streaming events that have influenced the estimate of risk, whether the model has been updated, and whether there are any delays or missing information in the data input. If the risk is evolving rapidly (second by second), then batch explainability techniques may not be enough. The framework should not be restricted to simply feature importance but should also include event-level traceability and patient-level temporal narratives.

Figure 2 presents the distribution of the reviewed articles across the main contribution categories of the hybrid framework. It is notable because of an imbalance in the evidence distribution, where there is evidence of infrastructure, evidence of prediction, and evidence of benchmarking, but not as much evidence of end-to-end hybrid evaluation. The figure shows a qualitative summary of the various reviewed journal articles based on an author-coded classification of the articles by dominant contribution: platform architecture, edge or fog processing, clinical prediction, benchmark data, and governance or interpretability. The basis of data/source is a set of coded references [5]–[18] with one predominant category assigned to each article.




Figure 2: Evidence distribution across hybrid healthcare analytics components

Figure 3 combines technical and clinical aspects of three popular design patterns: cloud-centric batch analytics, streaming edge or "fog" analytics and cloud-stream mix. It is important because there is no one right pattern for all criteria; the right architecture is a mix of latency, scalability, model governance, privacy exposure and integration of clinical workflows. The radar chart shows three architecture patterns and the author-derived scores (rated from 1 to 5) based on the common strengths and limitations mentioned in [5]–[18]. The scores represent qualitative, literature-informed assessments and are not the results of a meta-analysis.




Figure 3: Comparative capability profile of cloud-centric, edge-streaming, and hybrid architectures

When taken together, Tables 1-3 and Figures 1-3 illustrate that hybrid streaming and batch intelligence is more of a convergence than a discrete invention, the merging of distributed engineering of healthcare platforms and clinical machine learning. The most advanced research has demonstrated that real-time prediction is technically achievable, but not entirely operationally ready – distributed clinical intelligence. There is a need to integrate the infrastructure metrics and clinical metrics in the field within the same evaluation design. An incomplete prediction platform lacks freshness, model version, an alert pathway and recalibration status. Similarly, a cloud-IoT architecture without clinical endpoints that have been validated is infrastructural capability and not healthcare intelligence.

5 FUTURE DIRECTIONS

Looking ahead, the integration and evaluation of end-to-end hybrid streaming and batch intelligence frameworks in realistic distributed healthcare scenarios will be emphasized. The event-ingestion latency, event-time error, throughput, missing-data behavior, model inference delay, calibration drift, alert volume, clinician response, and clinical

outcome proxies should be measured using the same design. This would fill in the current disconnect between papers evaluating the architecture and papers evaluating the models [5]–[15] employed in clinical prediction. Single system validation of the streaming behavior should be a standard, as local charting practices, monitoring devices, laboratory workflows and alert policies can impact the behavior of a single system [16] [17].

Also, temporal governance methods should be improved. Batch retraining should not be treated as a sporadic offline task that is not related to the streaming inference. Future frameworks should include a model registry, automatic drift detection, shadow deployment, rollback logic and monitoring of subgroup calibration. The absence-aware framework should be extended to real-time data quality monitoring because of the possibility of behavior changes in predictions due to absence of data, disconnection of devices, and delayed documentation [14, 15]. The explainability research should focus on temporal explanations: events that recently occurred, missing observations or changes in the context that would impact risk estimates [18].

Interdisciplinary collaboration among distributed systems engineering, clinical informatics, nursing workflow, implementation science, cybersecurity, and biostatistics should be represented. However, for cloud platforms that involve home monitoring networks and hospital networks, it is especially critical to implement technical privacy measures, yet only technical privacy measures are not sufficient. Governance should clarify who is responsible for updating the model, adjusting thresholds, handling data-quality exceptions and suppressing inappropriate alerts. Going forward, future studies will move from prototype architecture and past performance to a continually evolving clinical intelligence platform that is deployed within auditable healthcare infrastructure.

6 CONCLUSION

In healthcare, the integration of streaming and batch intelligence is crucial for implementing real-time analytics on distributed cloud data platforms. The reviewed literature indicates that streaming components play a crucial role in timely physiological monitoring, event detection and dynamic risk estimation, and batch components will continue to play an important role in learning, calibration, audit, and population-based evaluation. The central conclusion is that neither stream processing alone nor retrospective machine learning alone is sufficient to provide safe real-time clinical intelligence. Value lies in a combination of coordinated and distributed data ingestion, data modelling in time, governance and workflow integration.

There continue to be significant gaps. A number of platform studies demonstrate the viability of cloud, fog and/or Internet of Things (IoT) architectures, although often without adequate clinical validation. Numerous prediction studies report strong model performance, but they often do not describe the operational infrastructure required for real-time deployment. Future advances depend on evaluating hybrid systems as complete clinical data platforms, including latency, reliability, privacy, calibration, interpretability, and clinical action. The field is moving toward distributed intelligence that can support hospital, community, and home-based services but requires more developed scholarship to deliver distributed robust evidence, governance that is transparent, and ongoing monitoring after deployment.

REFERENCES

- [1] Dash, S., Shakyawar, S. K., Sharma, M., & Kaushik, S. (2019). Big data in healthcare: Management, analysis and future prospects. *Journal of Big Data*, 6, 54.
- [2] Aceto, G., Persico, V., & Pescapé, A. (2020). Industry 4.0 and health: Internet of Things, big data, and cloud computing for healthcare 4.0. *Journal of Industrial Information Integration*, 18, 100129.
- [3] Wang, Y., Kung, L., & Byrd, T. A. (2018). Big data analytics: Understanding its capabilities and potential benefits for healthcare organizations. *Technological Forecasting and Social Change*, 126, 3–13.
- [4] Islam, S. M. R., Kwak, D., Kabir, M. H., Hossain, M., & Kwak, K. S. (2015). The Internet of Things for health care: A comprehensive survey. *IEEE Access*, 3, 678–708.
- [5] Forkan, A. R. M., Khalil, I., Ibaida, A., & Tari, Z. (2017). BDCaM: Big data for context-aware monitoring—a personalized knowledge discovery framework for assisted healthcare. *IEEE Transactions on Cloud Computing*, 5(4), 628–641.
- [6] Hossain, M. S., & Muhammad, G. (2016). Cloud-assisted industrial Internet of Things-enabled framework for health monitoring. *Computer Networks*, 101, 192–202.
- [7] Rahmani, A. M., Gia, T. N., Negash, B., Anzanpour, A., Azimi, I., Jiang, M., & Liljeberg, P. (2018). Exploiting smart e-health gateways at the edge of healthcare Internet-of-Things: A fog computing approach. *Future Generation Computer Systems*, 78, 641–658.

- [8] Manogaran, G., Thota, C., Lopez, D., & Sundarasekar, R. (2018). A new architecture of Internet of Things and big data ecosystem for secured smart healthcare monitoring and alerting system. *Future Generation Computer Systems*, 82, 375–387.
- [9] Alhussein, M., Muhammad, G., Hossain, M. S., & Amin, S. U. (2018). Cognitive IoT-cloud integration for smart healthcare: Case study for epileptic seizure detection and monitoring. *Mobile Networks and Applications*, 23(6), 1624–1635.
- [10] Sahoo, P. K., Mohapatra, S. K., & Wu, S. L. (2016). Analyzing healthcare big data with prediction for future health condition. *IEEE Access*, 4, 9786–9799.
- [11] Rajkumar, A., Oren, E., Chen, K., Dai, A. M., Hajaj, N., Hardt, M., Liu, P. J., Liu, X., Marcus, J., Sun, M., Sundberg, P., Yee, H., Zhang, K., Zhang, Y., Flores, G., Duggan, G. E., Irvine, J., Le, Q., Litsch, K., ... Dean, J. (2018). Scalable and accurate deep learning with electronic health records. *npj Digital Medicine*, 1, 18.
- [12] Tomašev, N., Glorot, X., Rae, J. W., Zielinski, M., Askham, H., Saraiva, A., Mottram, A., Meyer, C., Ravuri, S., Protsyuk, I., Connell, A., Hughes, C. O., Karthikesalingam, A., Cornebise, J., Montgomery, H., Rees, G., Laing, C., Baker, C. R., Peterson, K., ... Mohamed, S. (2019). A clinically applicable approach to continuous prediction of future acute kidney injury. *Nature*, 572(7767), 116–119.
- [13] Hyland, S. L., Faltys, M., Hüser, M., Lyu, X., Gumbsch, T., Esteban, C., Bock, C., Horn, M., Moor, M., Rieck, B., Zimmermann, M., Bodenham, D., Borgwardt, K., Rätsch, G., & Merz, T. M. (2020). Early prediction of circulatory failure in the intensive care unit using machine learning. *Nature Medicine*, 26(3), 364–373.
- [14] Harutyunyan, H., Khachatrian, H., Kale, D. C., Ver Steeg, G., & Galstyan, A. (2019). Multitask learning and benchmarking with clinical time series data. *Scientific Data*, 6, 96.
- [15] Che, Z., Purushotham, S., Cho, K., Sontag, D., & Liu, Y. (2018). Recurrent neural networks for multivariate time series with missing values. *Scientific Reports*, 8, 6085.
- [16] Johnson, A. E. W., Pollard, T. J., Shen, L., Lehman, L. W. H., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Celi, L. A., & Mark, R. G. (2016). MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3, 160035.
- [17] Pollard, T. J., Johnson, A. E. W., Raffa, J. D., Celi, L. A., Mark, R. G., & Badawi, O. (2018). The eICU Collaborative Research Database, a freely available multi-center database for critical care research. *Scientific Data*, 5, 180178.
- [18] Xiao, C., Choi, E., & Sun, J. (2018). Opportunities and challenges in developing deep learning models using electronic health records data: A systematic review. *Journal of the American Medical Informatics Association*, 25(10), 1419–1428.