FLEX-FHIR: A Modular Fog Computing Architecture for Community Health Monitoring and Decision Support Using AWS

Fernanda Schäfer Tesch da Silva[1], **Rodrigo da Rosa Righi**[1]
[1]Universidade do Vale do Rio dos Sinos (Unisinos)
Av. Unisinos – 93.022-750 – São Leopoldo – RS – Brazil
ftschafer@gmail.com, rrrighi@unisinos.br

Abstract. *Centralized cloud architectures face inherent limitations when addressing region-specific healthcare challenges in smart cities. This paper presents FLEX-FHIR (Fog Layer EXtensible Framework for Health using Interoperable Resources), a modular and scalable fog computing architecture for community-centered health monitoring and population-level decision support. The system is structured as a recursive hierarchical fog tree in which each node runs on Amazon Web Services (AWS) Elastic Compute Cloud (EC2) instances hosting containerized HAPI FHIR servers, PostgreSQL databases, and artificial intelligence/statistical modules. Regional health insights are generated locally through health clinical scoring, K-Means clustering, trimodal correlation analysis (Pearson, Spearman, Kendall), and spatial autocorrelation (Moran's I), all operating exclusively over FHIR resources. Experimental evaluation across workloads of up to 20,000 patients showed zero FHIR compliance failures, inter-tier forwarding latency of 83–134 ms per patient, independent of dataset size, and analytically correct behavior in all tested scenarios. The architecture is directly compatible with Brazil's Rede Nacional de Dados em Saúde (RNDS) and the International Patient Summary (IPS) initiative.*

1. Introduction

Advances in computing and the integration of healthcare applications have opened new opportunities for data-driven analysis, particularly through IoT-based wearable sensors for remote monitoring of vital signs [Azizah et al. 2023]. These data are transmitted from wearable devices to cloud servers where they are stored in Electronic Health Records (EHR) [Wang et al. 2023]. However, centralized architectures remain limited when addressing region-specific healthcare challenges, where low-latency responses and geographically disaggregated analysis are required.

Fog computing bridges IoT edge devices and the cloud by deploying computation on intermediate nodes, enabling low latency, location-aware services, and local decision-making [Chebaane et al. 2024]. In smart cities, fog nodes distributed across neighborhoods and health zones reduce cloud data volume while supporting real-time population monitoring [da Rosa Righi et al. 2025]. HL7 FHIR complements this infrastructure with an open RESTful API for semantic interoperability across heterogeneous devices and institutions — mandatory in Brazil's RNDS platform and central to the G20 International Patient Summary initiative [Ministério da Saúde do Brasil 2025a].

Despite these advances, few studies have proposed architectures that jointly deploy fog computing and HL7 FHIR for regional population health analysis. Existing work

focuses predominantly on individual-level disease prediction in cloud-only environments, leaving gaps in decentralized data preprocessing, regional insight extraction, and standardized data exchange. This paper addresses the following research question: *How can HL7 FHIR resources and decentralized fog nodes enable standardized ingestion, regional aggregation, and population-level analysis of vital sign data in smart city environments?*

The principal contributions are threefold: (i) we demonstrate that fog computing and HL7 FHIR are jointly deployable as a single interoperable stack, eliminating format translation at all integration boundaries; (ii) we show that a fog-native FHIR pipeline scales to 20,000 patients while preserving sub-second inter-tier aggregation latency; and (iii) we validate the system's ability to support ecological investigations into health equity by relating care-unit availability and socioeconomic factors to regional disease burden.

2. Related Work

Most related studies deploy analytical pipelines exclusively in the cloud, focusing on individual-level disease prediction using machine learning over structured clinical datasets, without geographic aggregation or standardized health data exchange [Malibari 2023, Garcés-Jiménez et al. 2024, Yang et al. 2025, Janjua et al. 2024, Prerna et al. 2023, Chen et al. 2025]. Specifically, [Prerna et al. 2023] propose a delay-sensitive framework for early disease prediction, and [Chen et al. 2025] employs AI to predict bacteremia from blood count data — both operating in cloud-only environments without fog or standardized exchange formats. One work proposes an edge-cloud prototype for real-time deep learning inference [Adhikari et al. 2022] but remains restricted to individual-level analysis with no fog tier.

The two fog-cloud architectures most structurally similar to the present work [André Setti Cassel et al. 2024, Ahmed et al. 2025] address response time and data integrity, respectively, without adopting standardized health data formats or embedding population-level analytics at the fog tier. A recent scoping review confirmed that cloud computing dominates 63% of interoperable IoMT platforms and that FHIR adoption at the edge remains limited [Seth et al. 2025]. Islam et al. [Islam et al. 2025] propose a multi-tier fog framework with tier-specific latency instrumentation, but without grounding data exchange in HL7 FHIR or embedding a full analytical layer within the nodes. Table 1 consolidates these works and positions the present contribution.

Table 1. Summary of related studies and comparison with the present work. Layer: C = Cloud, E = Edge, F = Fog, N/A = not specified.

Paper	Goal	Layer	Data	Method	Preprocessing / Enrichment	Format
[Malibari 2023]	Disease detection via classification; alerts to professionals	C	BP, albumin, ECG, diabetes	CNN	Imputation, deduplication, normalisation	N/A
[Garcés-Jiménez et al. 2024]	Vital sign disease detection with seasonal data	C	Vital signs, age, sex, weather	ML	Baseline + quantile modules; cleaning	CSV
[Yang et al. 2025]	Early sepsis warning from non-invasive vital signs	N/A	Temp, systolic BP, respiration rate	ML	Row removal for missing values	N/A
[Adhikari et al. 2022]	Edge-centric prototype for real-time health analytics	E, C	Temp, pulse, cough	DL	Cleaning and feature extraction	N/A
[Prerna et al. 2023]	Delay-sensitive framework for early disease prediction	C	Temp, BP, ECG	ML	Cleaning, transformation, reduction	N/A
[Chen et al. 2025]	AI model to predict bacteremia from blood count data	N/A	CBC, DC, CPD, blood cultures	ML	N/A	N/A
[Janjua et al. 2024]	Predictive analytics with patient privacy focus	N/A	BMI, BP, cholesterol, glucose, comorbidity	ML + DL	Interpolation, clinical cutoffs, feature engineering	N/A
[André Setti Cassel et al. 2024]	Hierarchical fog-cloud architecture for urgent vital signs	E, F, C	Temp, heart signals	N/A	N/A	JSON
[Ahmed et al. 2025]	Data integrity in IoT health via IOTA DAG	E, F, C	N/A	N/A	N/A	N/A
This work	**Regional population health intelligence via fog-embedded FHIR analytics**	**E, F, C**	**Vitals signs, socioeconomic, geographic metadata**	**K-means, Pearson, Spearman, Kendall, Moran's *I***	**Hierarchical FHIR enrichment with block, neighbourhood, city identifiers**	**HL7 FHIR**

The present work fills these gaps by: (i) integrating a complete edge-fog-cloud hierarchy on AWS; (ii) employing HL7 FHIR as the exclusive data exchange standard at all tiers; and (iii) embedding regional analytical intelligence — National Early Warning Score 2 (NEWS2) scoring, K-Means stratification, trimodal correlation, and spatial autocorrelation — directly within the fog nodes.

3. FLEX-FHIR: Proposed Model

FLEX-FHIR targets population-level vital sign monitoring and regional health insight extraction in smart city contexts. Fog computing was chosen as the primary processing tier to minimize latency, a critical requirement where round-trips to remote data centers are incompatible with clinical early-warning requirements [Jeyaraman et al. 2024]. Although NEWS2 was originally designed for bedside triage of individual patients in acute care, FLEX-FHIR repurposes it as a population-level composite indicator: by consolidating heterogeneous vital sign measurements into a single comparable score, it enables regional aggregation and trend analysis across geographic areas, serving as a proxy for community health burden.

Full HL7 FHIR compliance is enforced at every tier — all data exchange, persistence, and analytical output use standard FHIR interactions, eliminating format translation at integration boundaries and enabling direct consumption by any FHIR-compatible clinical information system. Three principles govern the design: automation, transparency, and usability.

3.1. Architecture

FLEX-FHIR follows a layered Edge–Fog–Cloud model for collection, processing, monitoring, and storage of vital signs, illustrated in Figure 1. At the edge, vital signs (respiratory rate, heart rate, systolic blood pressure (BP), etc) are collected and transmitted asynchronously to the nearest fog node via HL7 FHIR over HTTPS.

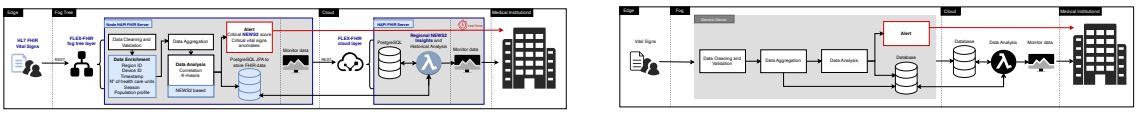

(a) FLEX-FHIR architecture **(b) Traditional fog-based architecture**

Figure 1. Comparison between FLEX-FHIR and traditional fog-based architectures

Fog nodes are distributed hierarchically across urban administrative divisions: blocks (Layer 1), neighborhoods (Layer 2), and cities (Layer 3), all deployed on dedicated AWS EC2 instances as containerized services comprising a HAPI FHIR JPA server, a PostgreSQL database, and AI/statistical modules. The architecture is recursive — each node ingests, enriches, and aggregates data from its children before forwarding FHIR resources upward. Regional data is condensed into MeasureReport resources, enriching upstream analysis. Each parent then consolidates its children's reports into its own MeasureReport, propagating summarized insights toward the cloud. Nodes retain data in a sliding 30-day window; older data is forwarded to the cloud. Geographic sharding is implemented naturally through independent PostgreSQL instances at each node, each responsible for data from its assigned region.

First Layer: Layer 1 (Blocks). Incoming resources are cleaned (missing value detection, range validation, unit normalization, timestamp validation, and duplicate detection) and enriched with hierarchical geographic identifiers (block, neighborhood, city) as structured FHIR extensions, enabling downstream systems to query by geographic layer without external lookups. Each vital signs is persisted as a FHIR Observation resource that receives an individual NEWS2 score, while each Patient resource aggregates the sum of

its observations' scores. Regional NEWS2 averages are then computed from patient-level scores. Clinical conditions are detected per patient based on vital sign patterns, and K-Means clustering is applied over conditions, vital signs, and patient attributes to identify population subgroups with distinct health profiles.

Intermediary Layers: Layer 2 (Neighborhoods), Layer 3 (Cities), Layer n (broad regions) . Nodes receive aggregated FHIR resources from their children — including conditions, patient-level NEWS2 scores, and summary reports carrying regional indicators (average vital signs, care unit count, mean age, sex distribution, average income, seasonality) — and apply Pearson, Spearman, and Kendall correlation analysis to detect associations between socioeconomic factors and health outcomes. Moran's I spatial autocorrelation is also computed to identify geographic clustering of health indicators across regions.

Cloud Layer. Hosted on AWS EC2, like the fog nodes, it provides long-term storage, advanced analytics, and aggregation across all fog nodes by geographic region, accessible through interactive dashboards. Node communication follows a DNS-inspired addressing scheme (e.g., *block5.neigh2.cityx*), supporting scalable and context-aware routing across all tiers.

The fundamental differentiator of FLEX-FHIR, contrasted with a traditional fog architecture in Figure 1, is the co-location of a dedicated HAPI FHIR server at every fog and cloud node, ensuring all data exchange, persistence, and querying conform to HL7 FHIR at every layer, eliminating format translation at integration boundaries. Beyond compliance, the framework introduces population-level vital sign scoring via NEWS2, regional health insight generation, and adaptable per-region enrichment — capabilities absent from traditional fog architectures shown in Figure 1b.

3.2. AI and Analytical Modules

FLEX-FHIR embeds a suite of analytical modules across the fog hierarchy, covering clinical scoring, condition detection, clustering, correlation, and spatial analysis.

Clinical Decision Support (CDS/CQL). A scheduled Clinical Decision Support (CDS) pipeline evaluates patient-level vital-sign evidence through Clinical Quality Language (CQL) rule sets. When predefined criteria are met, the system materializes a FHIR Condition resource linked to the originating observations, preserving semantic interoperability and auditability while maintaining a clear separation between clinical measurements, decision logic, and alerting layers.

NEWS2 Scoring and Alerts. At Layer 1, a scoring module computes NEWS2 for each patient. Patient alerts are triggered when a score reaches the high-acuity threshold (≥ 7); block-level alerts fire when the regional average NEWS2 exceeds 5.0. Both alert types create FHIR Communication resources and dispatch notifications via Simple Mail Transfer Protocol (SMTP), ensuring each distinct deterioration episode produces exactly one active alert.

K-Means Clustering. Applied at Layer 1 where individual patient feature vectors (heart rate, respiratory rate, blood pressure, and related parameters) are available, enabling discovery of clinically meaningful subgroups (stable, deteriorating, high-risk profiles). Cluster validity is assessed by silhouette score; output is suppressed for homogeneous

populations (silhouette < 0.20) to avoid spurious partitioning. Deferring K-Means to upper layers would constitute an ecological fallacy, as block-level averages are consistent with multiple underlying patient distributions.

Trimodal Correlation Analysis. Applied at Layer 2 and above over regional aggregates. Pearson r, Spearman ρ, and Kendall τ_b are computed simultaneously for variable pairs such as careUnits_vs_news2, averageIncome_vs_conditions, and meanAge_vs_heartRate. Reporting all three coefficients reveals nonlinearity and tied-observation effects that a single coefficient would mask.

Spatial Autocorrelation (Moran's I). Computed at neighborhood and city tiers over block-level health indicators to detect geographic concentration of health signals, distinguishing spatially structured distributions from random ones.

4. Experimental Evaluation

This section presents the experimental setup and results of FLEX-FHIR, covering infrastructure, input data, evaluation scenarios, analytical correctness, and latency under increasing workloads.

4.1. AWS Infrastructure and Topology

The fog tiers run on AWS EC2 instances, providing configurable computational resources and network isolation for reproducible evaluation. The simulated deployment maps a small city: five block nodes (B1–B5), three neighborhood nodes (N1–N3), and one city node (C1). Blocks B1, B2, and B3 report to N1; B4 to N2; B5 to N3; all three neighborhoods report to C1, which connects to the cloud layer, as displayed in Figure 2. Each fog node runs as an independent EC2-hosted container instrumented with Micrometer (Prometheus registry) scraped via Spring Boot Actuator.

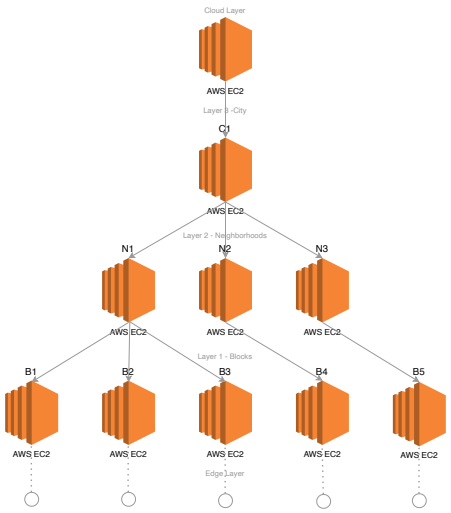

Figure 2. AWS EC2 hierarchical fog-to-cloud topology of the simulated scenario

4.2. Input Data and Simulation

Vital sign data were sourced from a publicly available Kaggle dataset [Ayub 2023] compatible with the NEWS2 framework [Royal College of Physicians 2017], containing

anonymized records with six physiological parameters plus demographic attributes (age, gender). A custom Python ingestion module converts raw CSV records to HL7 FHIR Observation and Patient resources, assigning Logical Observation Identifiers Names and Codes (LOINC) — a standard terminology for clinical measurements — to each vital sign. Resources are posted to the appropriate fog node via FHIR REST, simulating real-time edge device ingestion. Each patient record yields one Patient and six Observation resources, bundled and posted in batches of 200–500 patients.

4.3. Evaluation Scenarios

Test scenarios were organized in four groups of increasing complexity: (1) **patient-level processing** — NEWS2 computation and condition detection; (2) **fog-tier aggregation** — vital sign trend analysis and statistical summarization across layers; (3) **correlation, clustering, and spatial analytics** — correlation validation with known relationships, spurious-correlation control, zero-variance edge cases, Moran's I validation, and both neutral and well-separated clustering; and (4) **robustness and scalability** — regional variation across multiple demographic profiles and patient volume variation at 1,000, 5,000, 10,000, and 20,000 patients.

4.4. Analytical Correctness

The NEWS2 pipeline, condition detection, and alert mechanisms produced no false positives and no persistence-layer errors across all tested scenarios including boundary cases. The K-Means negative control correctly suppressed cluster output for homogeneous populations (silhouette < 0.20), avoiding spurious partitioning [Lekadir et al. 2025]. For synthetic cohorts with distinct risk profiles, the silhouette score reached 0.712, with an age separation of 56.7 years and heart rate separation of 58.3 bpm between clusters, well above the 0.5 threshold for meaningful cluster structure. Table 2 summarizes both clustering experiments.

Table 2. K-Means results: structured cohort (300 patients) and negative control. HR: heart rate (bpm); SBP: systolic blood pressure (mmHg); Age: mean age (years); Cond.: average number of conditions.

Cluster	N	HR	SBP	Age	Cond.	Silhouette
C2 (high risk)	100	120.1	144.9	79.8	5.29	0.712
C1 (moderate)	100	89.0	127.7	49.6	2.64	
C0 (low risk)	100	61.8	113.4	23.1	0.00	
Negative control	300	Homogeneous — output suppressed				0.000

Correlation analysis used three structurally distinct datasets (uniform linear, geometric nonlinear, and step-tier), with results shown in Table 3. The table reports Pearson r (linear association), Spearman ρ_s (monotonic rank-based), and Kendall τ_b (ordinal concordance), alongside $\hat{p}$ (permutation p-value, $B = 1000$). The careUnits_vs_news2 pair yielded strong inverse associations ($r = -0.980$ to -0.985, $\hat{p} = 0.001$). The divergence in avgIncome_vs_conditions between Pearson and rank-based measures confirmed that reporting all three coefficients simultaneously is not redundant. The uncorrelated dataset produced no significant associations, and the zero-variance dataset correctly returned undefined coefficients for constant variables. Moran's I results (Table 4) identified strongly clustered patterns ($I \approx 0.95$–0.98, $p < 0.01$) in correlated datasets and predominantly dispersed patterns in uncorrelated datasets [Isnan et al. 2025].

Table 3. Correlation summary across validation and edge case datasets

Dataset	Pair	r	ρ_s	τ_b	$\hat{p}$
Pearson	careUnits_vs_news2	$+0.995$	$+0.997$	$+0.979$	0.001
Pearson	avgIncome_vs_conditions	$+0.990$	$+0.993$	$+0.960$	0.001
Spearman	careUnits_vs_news2	-0.980	-0.998	-0.987	0.001
Spearman	avgIncome_vs_conditions	-0.913	-0.998	-0.987	0.001
Kendall	careUnits_vs_news2	-0.985	-0.997	-0.976	0.001
Kendall	avgIncome_vs_conditions	-0.876	-0.996	-0.976	0.001
Uncorrelated	careUnits_vs_news2	-0.168	-0.168	-0.077	0.580
Zero variance	meanAge_vs_heartRate	undef	undef	undef	—

Table 4. Moran's I summary for correlated and uncorrelated datasets

Dataset	Variable	I	Pattern
Correlated	avgNews2	0.985	Clustered
Correlated	conditionCount	0.960	Clustered
Correlated	careUnits	0.960	Clustered
Correlated	heartRate	0.953	Clustered
Uncorrelated	avgNews2	-0.639	Dispersed
Uncorrelated	conditionCount	-0.768	Dispersed
Uncorrelated	careUnits	-0.684	Dispersed
Uncorrelated	heartRate	-0.165	Random

4.5. Latency and Scalability

Table 5 summarizes key latency results. Inter-tier forwarding and *MeasureReport*, which condensates region data, are insensitive to scale, confirming that the fog-to-fog communication design does not become a bottleneck.

Table 5. Latency summary across scale points (1,000–20,000 patients)

Operation	Scale sensitivity	Observed range
Upstream forwarding (per patient)	None	83–134 ms
MeasureReport aggregation (neighborhood)	None	3–32 ms
MeasureReport aggregation (city)	None	3–17 ms
Dashboard + K-Means (block, 1K–5K)	Sub-linear	16–272 ms
Dashboard + K-Means (block, 20K)	Super-linear	0.52–8.23 s
Spatial + Correlation (neighborhood)	Super-linear	16 ms–1.82 s
Spatial autocorrelation (city)	Super-linear	25 ms–5.15 s

Upstream forwarding latency remained stable at 83–134 ms per patient across all nodes and all four scale points, confirming that per-patient FHIR transaction overhead is independent of cumulative dataset size. MeasureReport aggregation at the neighborhood tier maintained mean latency below 32 ms and at the city tier below 17 ms even at 20,000 patients.

Block-tier analytics remained sub-second up to 5,000 patients, crossing the one-second threshold at 10,000 patients for higher-latency nodes, and growing steeply at 20,000 patients (0.52–8.23 s), driven primarily by FHIR search query cost over the accumulated patient-resource index. Node B4 consistently showed the highest latency, attributable to EC2 instance-level heterogeneity rather than a software defect. Neighborhood-tier spatial and correlation operations scale super-linearly consistent with $\mathcal{O}(n^2)$ spatial-weight

computation: N1 recorded 18 ms at 1,000 patients growing to 1.82 s at 20,000 patients ($101\times$ increase over a $20\times$ patient count increase). JVM heap utilization remained below 34% across all nodes and scale points. Zero FHIR compliance failures and zero connection-pool timeouts were recorded throughout all experiments.

5. Discussion

The results demonstrate that fog computing and HL7 FHIR are jointly deployable as a coherent interoperable stack in which analytical outputs are produced locally, transmitted as standard FHIR resources, and consumed directly by any FHIR-compatible downstream system [Ramirez Lopez et al. 2026]. Fog-embedded NEWS2 scoring reduces latency between physiological deterioration and alert generation which is unachievable with batch-oriented cloud analytics: block nodes can score patients, trigger alerts, and propagate aggregates upward independently of whether upper tiers are reachable — a critical property where connectivity to centralized infrastructure cannot be guaranteed. The socioeconomic correlation and spatial clustering outputs provide health authorities with an operational mechanism for health equity investigation from routinely collected data, without dedicated epidemiological surveys, contributing to evidence-based resource allocation at block and neighborhood scale.

FLEX-FHIR is directly compatible with Brazil's RNDS, which mandates HL7 FHIR for data exchange across public and private health institutions [Ministério da Saúde do Brasil 2025a]. Its FHIR-native output could complement the SAGE platform [Ministério da Saúde do Brasil 2025b], extending geographic health intelligence to sub-municipal granularity and near-real-time resolution without changes to existing national data models. Full HAPI FHIR compliance also positions the architecture as compatible with the IPS initiative [Joint Initiative Council 2021]. A pilot deployment begins naturally with a single neighborhood of three to five block nodes — well within the validated operating range — and scales progressively toward city-wide coverage.

6. Conclusion

This paper presented FLEX-FHIR, a fog-native HL7 FHIR architecture for regional population health intelligence in smart city environments, evaluated on AWS EC2. Experimental results demonstrated: (i) full HL7 FHIR compliance with zero failures; (ii) inter-tier forwarding latency independent of patient volume (83–134 ms per patient up to 20,000 patients); (iii) correct behavior across NEWS2 scoring, K-Means stratification, trimodal correlation, and spatial autocorrelation; and (iv) MeasureReport aggregation maintaining mean latency below 32 ms at the neighborhood tier and 44 ms at the city tier at all tested scales. The architecture reduces reliance on centralized cloud analysis, preserves semantic interoperability across all tiers, and provides a deployable pathway for geographically disaggregated early warning and health equity investigation using standards already adopted by the global health informatics community.

6.1. Limitations

Latency was measured on AWS EC2, not real edge hardware, and the synthetic dataset does not meet clinical accuracy standards for NEWS2 (especially temperature and respiratory rate). The data model is block-centric rather than patient-centric, and scaling beyond 10,000 patients per block will require sparse matrices and query optimization, though neither affects the FHIR model or fog topology.

Declaration on the Use of Artificial Intelligence

Artificial intelligence tools were used as assistants during the preparation of this article. Specifically, Claude (Anthropic) was employed to support the following tasks: academic writing and revision of results, discussion, and conclusion chapters; refinement of tables for LaTeX format; and structural editing to reduce repetition and improve consistency across sections.

All scientific content, experimental design, implementation, data collection, and intellectual contributions are the sole work of the author. The AI tools were used exclusively to assist with language, formatting, and analytical interpretation under the author's direction and supervision. Every AI-generated or AI-assisted passage was reviewed, verified, and revised by the author before inclusion. The author takes full responsibility for the accuracy, integrity, and originality of the work presented in this work.

References

Adhikari, M., Hazra, A., and Nandy, S. (2022). Deep transfer learning for communicable disease detection and recommendation in edge networks. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 20(4):2468–2479.

Ahmed, W., Iqbal, W., Hassan, A., Ahmad, A., Ullah, F., and Srivastava, G. (2025). Elevating e-health excellence with iota distributed ledger technology: Sustaining data integrity in next-gen fog-driven systems. *Future Generation Computer Systems*, 168:107755.

André Setti Cassel, G., da Rosa Righi, R., André da Costa, C., Rosecler Bez, M., and Pasin, M. (2024). Towards providing a priority-based vital sign offloading in healthcare with serverless computing and a fog-cloud architecture. *Future Generation Computer Systems*, 157:51–66.

Ayub, N. (2023). Human vital sign dataset. `https://www.kaggle.com/datasets/nasirayub2/human-vital-sign-dataset`. Accessed: 2026-05-19.

Azizah, N., Faisal, M. R., Abadi, F., Budiman, I., Mazdadi, M. I., Herteno, R., and Nugrahadi, D. T. (2023). A vital sign monitoring system exploiting bt/ble on low-cost commercial smartwatch for home care patients. In *2023 International Seminar on Intelligent Technology and Its Applications (ISITIA)*, pages 405–410.

Chebaane, A., Arshad, M. K., Burger, F., and Khelil, A. (2024). A layered strategy for reducing offloading latency in fog computing. In *2024 9th International Conference on Fog and Mobile Edge Computing (FMEC)*, pages 176–182.

Chen, W.-H., Chang, Y.-H., Hsiao, C.-T., Hsueh, P.-R., and Shih, H.-M. (2025). Utilizing artificial intelligence and cellular population data for timely identification of bacteremia in hospitalized patients. *International Journal of Medical Informatics*, 195:105788.

da Rosa Righi, R., Rodrigues, V. F., da Costa, C. A., and Eskofier, B. (2025). Breaking down the data path in digital health: From edge to fog and beyond. *IEEE Pervasive Computing*.

Garcés-Jiménez, A., Polo-Luque, M.-L., Gómez-Pulido, J. A., Rodríguez-Puyol, D., and Gómez-Pulido, J. M. (2024). Predictive health monitoring: Leveraging artificial

intelligence for early detection of infectious diseases in nursing home residents through discontinuous vital signs analysis. *Computers in Biology and Medicine*, 174:108469.

Islam, U., Alatawi, M. N., Alqazzaz, A., et al. (2025). A hybrid fog-edge computing architecture for real-time health monitoring in iomt systems with optimized latency and threat resilience. *Scientific Reports*, 15(1):25655.

Isnan, S., Abdullah, A. F., Shariff, A. R. M., Ishak, I., Ismail, S. N. S., and Appanan, M. R. (2025). Moran's I and Geary's C: investigation of the effects of spatial weight matrices for assessing the distribution of infectious diseases. *Geospatial Health*, 20(1):1277.

Janjua, J. I., Ghazal, T. M., Abushiba, W., and Abbas, S. (2024). Optimizing patient outcomes with ai and predictive analytics in healthcare. In *2024 IEEE 65th International Scientific Conference on Power and Electrical Engineering of Riga Technical University (RTUCON)*, pages 1–6. IEEE.

Jeyaraman, N. et al. (2024). Applications of fog computing in healthcare. *Cureus*, 16(7):e64263.

Joint Initiative Council (2021). International Patient Summary Standard – ISO 27269:2021. Technical report, International Organization for Standardization. Implemented via HL7 FHIR; see `https://international-patient-summary.net`.

Lekadir, K., Quaglio, G., Gallas, S., et al. (2025). Future-ai: international consensus guideline for trustworthy and deployable artificial intelligence in healthcare. *The BMJ*, 388:e080214.

Malibari, A. A. (2023). An efficient iot-artificial intelligence-based disease prediction using lightweight cnn in healthcare system. *Measurement: Sensors*, 26:100695.

Ministério da Saúde do Brasil (2025a). Rede Nacional de Dados em Saúde (RNDS). `https://www.gov.br/saude/pt-br/composicao/seidigi/rnds`. Accessed: 2025.

Ministério da Saúde do Brasil (2025b). Sala de Apoio à Gestão Estratégica (SAGE). `https://novasage.saude.gov.br`. Accessed: 2025.

Prerna, Singh, P., and Singh, D. P. (2023). A delay sensitive framework for effective healthcare using machine learning. In *2023 10th International Conference on Computing for Sustainable Global Development (INDIACom)*, pages 541–545.

Ramirez Lopez, L. J., Jaimes Salazar, N. E., and Barbosa Posada, J. E. (2026). Pire: Interoperable platform for electronic records. *Computers*, 15(3):162.

Royal College of Physicians (2017). National Early Warning Score (NEWS) 2.

Seth, M., Jalo, H., Högstedt, Å., Medin, O., Sjöqvist, B. A., and Candefjord, S. (2025). Technologies for interoperable internet of medical things platforms to manage medical emergencies in home and prehospital care: Scoping review. *Journal of Medical Internet Research*, 27:e54470.

Wang, R., Tu, J., and Chu, C. (2023). Research on the application of cross-regional sharing of blockchain-based electronic health records. In *2023 IEEE International Conference on Integrated Circuits and Communication Systems (ICICACS)*, pages 1–6.

Yang, A. C., Ma, W.-M., Chiang, D.-H., Liao, Y.-Z., Lai, H.-Y., Lin, S.-C., Liu, M.-C., Wen, K.-T., Lin, T.-H., Tsai, W.-X., et al. (2025). Early prediction of sepsis using an xgboost model with single time-point non-invasive vital signs and its correlation with c-reactive protein and procalcitonin: A multi-center study. *Intelligence-Based Medicine*, 11:100242.