

Comprehensive Review on Artificial Intelligence-Based Smart Healthcare Monitoring Using IoT and HCI

Swamy TN, Assistant Professor, Department of Electronics & Communication, Dr. Ambedkar Institute of Technology, Bengaluru, Karnataka, India swamyohm@gmail.com

Abstract: The convergence of the Internet of Things (IoT), artificial intelligence (AI), and human-computer interaction (HCI) is reshaping healthcare monitoring from episodic clinical observation toward continuous, context-aware, and increasingly personalized health management. IoT-enabled medical sensors, wearable devices, ambient intelligence systems, mobile platforms, and connected clinical equipment generate heterogeneous physiological, behavioural, environmental, and contextual data at unprecedented temporal resolution. AI techniques transform these data into predictions, classifications, anomaly detection, risk scores, and decision-support information, while HCI provides the mechanisms through which patients, caregivers, and clinicians perceive, interpret, validate, and act upon algorithmic outputs. Recent reviews indicate that AI-enabled remote monitoring has expanded from conventional cloud-based architectures toward edge intelligence, federated learning, explainable AI, and human-centred Healthcare 5.0 architectures [1]–[5]. This review critically examines the integration of AI, IoT, and HCI in smart healthcare monitoring. A layered taxonomy is developed covering sensing and acquisition, communication, edge/cloud intelligence, clinical analytics, interaction, and governance. Machine-learning, deep-learning, multimodal fusion, time-series modeling, explainable AI, reinforcement learning, and federated-learning methodologies are examined with respect to their suitability for continuous monitoring. Particular attention is given to latency, energy consumption, model generalization, data heterogeneity, interpretability, interoperability, privacy, cybersecurity, usability, and clinical validation. The review also compares centralized, edge, cloud, and federated architectures and examines their implications for real-time healthcare applications. Representative applications involving cardiovascular monitoring, diabetes management, respiratory disease, neurological disorders, elderly care, rehabilitation, mental-health monitoring, and emergency detection are discussed. A reference case study is presented for an AI-assisted remote patient-monitoring system integrating wearable sensors, edge analytics, a clinical dashboard, and human-centred alert mechanisms. Finally, the review identifies unresolved challenges involving dataset bias, sensor reliability, model drift, human over-reliance, alarm fatigue, interoperability, regulatory compliance, and equitable access. Future research directions include multimodal foundation models, adaptive edge intelligence, digital twins, privacy-preserving learning, standardized evaluation protocols, explainable multimodal AI, and human-AI collaborative decision-making.

Keywords: Artificial Intelligence; Internet of Things; Internet of Medical Things; Smart Healthcare; Human-Computer Interaction; Remote Patient Monitoring; Edge AI; Federated Learning; Explainable AI; Wearable Sensors; Healthcare 5.0; Clinical Decision Support.

1. Introduction

Healthcare delivery is undergoing a fundamental transition from institution-centric and encounter-based models toward continuous, distributed, and patient-centered monitoring. Traditional healthcare systems primarily depend upon measurements collected during clinical visits, hospital admissions, laboratory investigations, and manually documented observations. Although these mechanisms remain essential, they provide only intermittent representations of physiological state. A patient's blood pressure, heart rate, oxygen saturation, glucose concentration, physical activity, sleep pattern, or respiratory behavior can vary substantially between clinical encounters. Consequently, clinical decisions based exclusively on sparse observations may fail to capture transient abnormalities or gradual deterioration. The emergence of connected sensing technologies has created an opportunity to observe health-related phenomena continuously and outside conventional clinical environments.

The Internet of Things provides the infrastructure required for such continuous observation. Healthcare-oriented IoT systems combine physiological sensors, wearable devices, implantable or bedside equipment, mobile devices, communication networks, gateways, databases, and computational services. The resulting Internet of Medical Things (IoMT) ecosystem can acquire measurements such as electrocardiography (ECG), photoplethysmography (PPG), blood oxygen saturation, body temperature, respiratory rate, blood pressure, glucose, motion, posture, sleep, and environmental variables. IoT and AI-based remote healthcare monitoring has consequently become an important research area, with existing surveys identifying remote monitoring, clinical decision support, chronic disease management, and emergency detection as major application domains [1], [2].

Nevertheless, connectivity alone does not produce intelligent healthcare. Continuous sensing generates high-volume, noisy, heterogeneous, and temporally correlated data that cannot be transformed into clinically meaningful information simply by transmitting measurements to a server. Artificial intelligence provides the analytical layer required to identify patterns, distinguish physiological variation from pathological events, estimate future risk, and support clinical decision-making. Classical machine-learning models such as logistic regression, support vector machines, decision trees, random forests, and gradient-boosting algorithms remain useful for structured clinical data. Deep neural networks, including convolutional neural networks (CNNs), recurrent neural networks (RNNs), long short-term memory networks (LSTMs), temporal convolutional networks, autoencoders,

graph neural networks, and transformers, provide increasingly powerful approaches for physiological time-series, medical images, multimodal data, and longitudinal electronic health records.

The third component—human–computer interaction—is frequently underrepresented when AI and IoT healthcare systems are discussed. An algorithm can achieve high classification accuracy in a laboratory environment while remaining clinically ineffective if its output is presented at an inappropriate time, in an incomprehensible form, or through an interaction mechanism that imposes excessive cognitive burden. Human-centered IoT healthcare research therefore emphasizes the need to integrate patients, clinicians, caregivers, and other stakeholders into the design of monitoring systems rather than treating humans merely as sources of data [3]. Similarly, recent HCI research in AI-enabled healthcare highlights the importance of interaction design, explainability, trust, speech and gesture interfaces, and ethical considerations [4].

The resulting research problem is consequently multidimensional. AI must be sufficiently accurate and robust; IoT infrastructure must be reliable, energy-efficient, secure, and interoperable; and HCI must ensure that information is understandable, actionable, and aligned with clinical workflows. These requirements are interdependent. Increasing model complexity may improve predictive performance but increase computational requirements and reduce interpretability. Moving inference from the cloud to an edge device may reduce latency and protect sensitive data but constrain model size and energy consumption. Increasing the number of alerts may improve sensitivity while simultaneously producing alarm fatigue. An apparently technically superior system may therefore have inferior real-world utility.

Recent research increasingly addresses these trade-offs through edge AI, federated learning, privacy-preserving analytics, explainable AI, and human-centered architectures. Federated-learning research in smart healthcare, for example, emphasizes decentralized training across healthcare institutions and IoT devices while identifying data heterogeneity, model-inversion attacks, poisoning, scalability, and interoperability as unresolved issues [5]. Recent surveys of next-generation smart healthcare similarly identify interoperability, energy efficiency, data quality, security, and ethical governance as persistent barriers to large-scale deployment [6].

This review therefore treats AI-based IoT healthcare monitoring as an integrated cyber-physical-human system rather than as an isolated machine-learning problem. The subsequent sections examine its technological foundations, taxonomy, algorithms, architectures, applications, evaluation methods, limitations, ethical implications, and future research opportunities.

2. Background and Motivation

2.1 Evolution from Telemedicine to Intelligent Remote Monitoring

Remote healthcare monitoring initially developed around telemedicine, in which clinical information was transmitted from patients or remote facilities to healthcare professionals. Early systems generally transmitted relatively small quantities of data, such as blood pressure readings, glucose measurements, or medical images. The emergence of low-power sensors, smartphones, wireless communication, cloud computing, and wearable electronics substantially changed this paradigm. Instead of collecting isolated measurements, modern systems can construct continuous streams describing physiological and behavioral state.

The development of IoT has been particularly significant because it enables sensing, identification, communication, computation, and actuation to occur within a common ecosystem. A wearable ECG sensor can transmit signals through Bluetooth Low Energy to a smartphone; the smartphone can perform preliminary filtering; an edge gateway can execute anomaly detection; and selected events can subsequently be forwarded to cloud infrastructure for longitudinal analysis. Such architectures reduce the conceptual distinction between medical devices and information systems.

AI introduces another transformation. Rather than presenting raw data to a clinician, an intelligent monitoring system can derive higher-level representations. For example, a continuous ECG signal can be transformed into heart-rate variability features, arrhythmia probabilities, rhythm classifications, and longitudinal risk indicators. Similarly, accelerometer and gyroscope signals can be converted into activity states, gait characteristics, fall probabilities, or rehabilitation metrics. The analytical objective therefore shifts from data collection to computational interpretation.

A significant motivation for this transformation is the growing prevalence of chronic diseases and ageing populations. Chronic cardiovascular, respiratory, metabolic, neurological, and mobility-related conditions frequently require monitoring over months or years rather than during isolated hospital visits. Continuous monitoring may identify changes before symptoms become sufficiently severe to trigger clinical consultation. However, this potential should not be equated with guaranteed clinical benefit. Evidence generated from retrospective datasets or controlled pilot studies does not automatically establish improvements in mortality, hospitalization, quality of life, or healthcare costs. The clinical value of an AI-IoT system depends on whether its predictions produce appropriate and timely actions.

2.2 The Data Deluge Problem

IoT healthcare produces a fundamentally different data environment from conventional clinical informatics. Physiological signals are continuous, multimodal, and highly time dependent. A single patient may simultaneously generate ECG, PPG, oxygen saturation, temperature, accelerometry, sleep, location, medication adherence, and environmental data. These streams differ in sampling rate, noise characteristics, missingness patterns, temporal resolution, and semantic meaning.

For a signal $x(t)$, an observed measurement can be represented as $y(t) = x(t) + n(t) + a(t)$

where $x(t)$ represents the physiological signal of interest, $n(t)$ represents measurement noise, and $a(t)$ represents artifacts generated by motion, sensor displacement, electromagnetic interference, or other disturbances. In real-world monitoring, $a(t)$ can become substantial. Consequently, increasing sensor density does not necessarily increase information quality.

AI is required not only to classify signals but also to determine whether the signal should be trusted. Signal-quality assessment can be formulated as $q_t = f(x_{t-k:t}, s_t, c_t)$ where q_t denotes estimated signal quality, $x_{t-k:t}$ is a recent signal window, s_t represents sensor status, and c_t represents contextual information. A downstream diagnostic model can then condition its prediction on q_t , reducing the probability that a sensor artifact is interpreted as a clinical event.

2.3 Motivation for Human-Centered Intelligence

A clinically useful system must also account for human behavior. Patients may remove wearable devices, ignore notifications, misunderstand risk indicators, or discontinue use because of discomfort. Clinicians may disregard alerts if they occur too frequently or lack sufficient contextual information. Caregivers may require simplified interfaces rather than professional dashboards. Thus, the human element is not an external consideration but part of the system's operational loop.

The human-centered perspective has become increasingly prominent in Healthcare 5.0. A 2024 review of human-centered IoT healthcare monitoring emphasizes that conventional IoT research has often prioritized technical infrastructure while giving insufficient attention to human-centered health-data analytics [3]. This distinction is important because healthcare systems involve heterogeneous users with different levels of technical expertise, health literacy, cognitive load, physical ability, and decision authority.

The motivation for integrating HCI into AI-IoT healthcare can therefore be expressed as a closed-loop objective: $\text{Sensing} \rightarrow \text{Inference} \rightarrow \text{Explanation} \rightarrow \text{Interaction} \rightarrow \text{Action} \rightarrow \text{Feedback}$.

A system that optimizes only the inference stage remains incomplete. The output must be understood and incorporated into a human decision process.

3. Taxonomy / Classification

A comprehensive classification of AI-enabled IoT healthcare monitoring should consider multiple dimensions simultaneously. A device-centric taxonomy based only on wearable, implantable, and stationary devices is insufficient because modern healthcare systems distribute intelligence across sensors, edge nodes, smartphones, cloud platforms, and human interfaces. A more useful taxonomy is therefore constructed around six dimensions: sensing modality, computational location, AI methodology, interaction modality, healthcare objective, and deployment context.

This taxonomy demonstrates that a single application can simultaneously occupy several categories. A smartwatch-based arrhythmia system, for example, is a wearable physiological sensing platform, uses edge/mobile computing, may employ CNN or recurrent models, provides visual and haptic interaction, and performs diagnostic screening in an ambulatory context.

Table 1. Multidimensional taxonomy of AI-enabled IoT healthcare monitoring

Dimension	Major categories	Representative technologies	Primary objective
Sensing	Physiological	ECG, PPG, SpO ₂ , BP, temperature, glucose	Vital-sign monitoring
Sensing	Motion/behavioral	Accelerometer, gyroscope, pressure sensors	Activity, gait, falls
Sensing	Ambient	Cameras, radar, microphones, environmental sensors	Contactless monitoring
Sensing	Biochemical	Glucose, sweat, biomarkers	Metabolic monitoring
Computing	Device	TinyML, embedded inference	Ultra-low latency
Computing	Edge	Gateway, smartphone, edge server	Real-time analytics
Computing	Cloud	Centralized GPU/CPU infrastructure	Large-scale analytics
Learning	Supervised	CNN, RF, SVM, XGBoost	Classification/regression
Learning	Sequential	LSTM, GRU, TCN, Transformer	Time-series prediction
Learning	Unsupervised	Autoencoder, clustering	Anomaly detection
Learning	Federated	FedAvg, personalized FL	Distributed learning
Interaction	Visual	Dashboards, mobile applications	Clinical/patient visualization
Interaction	Conversational	Voice assistants, chat interfaces	Natural interaction
Interaction	Tangible	Wearables, haptic interfaces	Immediate feedback
Interaction	Immersive	AR/VR/MR	Rehabilitation/training
Objective	Diagnostic	Disease/event classification	Detection
Objective	Predictive	Risk forecasting	Early intervention
Objective	Preventive	Behavioral feedback	Risk reduction
Objective	Therapeutic	Closed-loop control	Treatment adaptation
Context	Home	RPM, elderly monitoring	Out-of-hospital care
Context	Hospital	ICU, ward, emergency	Clinical surveillance
Context	Community	Public-health monitoring	Population-level care




Figure 1. Multi-Dimensional taxonomy of IoT-HCI Health care System

The taxonomy also reveals a fundamental methodological shift. Earlier healthcare IoT systems often followed a pipeline of sensing → transmission → cloud storage → classification. Current research increasingly adopts sensing → local preprocessing → edge inference → context-aware decision → human interaction, with cloud resources reserved for longitudinal learning, model management, and population-level analysis. This architectural evolution is driven by latency, privacy, bandwidth, and energy requirements.

Table 2. AI methodology classification

AI method	Input characteristics	Strengths	Principal limitations	Typical healthcare use
Logistic regression	Structured features	Interpretable, efficient	Limited nonlinear modeling	Risk prediction
SVM	Small/medium feature sets	Strong margins	Scaling and interpretability	Signal classification
Random Forest	Tabular/mixed features	Robust, relatively interpretable	Large ensembles	Clinical prediction
XGBoost	Tabular data	Strong predictive performance	Feature engineering often required	Risk stratification
CNN	Signals/images	Local pattern extraction	Data/computation demand	ECG, imaging
LSTM/GRU	Time series	Temporal dependencies	Sequential training cost	Vital-sign prediction
TCN	Time series	Parallelizable temporal modeling	Receptive-field design	Physiological signals
Transformer	Long sequences/multimodal data	Long-range dependencies	Computational cost	Multimodal monitoring
Autoencoder	Unlabeled signals	Anomaly detection	Threshold sensitivity	Abnormal-event detection
GNN	Relational data	Explicit topology modeling	Graph construction complexity	Patient/device networks
Federated learning	Distributed data	Data-local training	Heterogeneity/security	Multi-hospital learning
Reinforcement learning	Sequential decisions	Adaptive policies	Safety and exploration	Personalized intervention

4. Core Concepts and Methodologies

4.1 IoT-Based Physiological Sensing

The sensing layer constitutes the physical foundation of smart healthcare monitoring. Common sensors include ECG electrodes, optical PPG sensors, pulse oximeters, thermistors, pressure sensors, accelerometers, gyroscopes, microphones, radar systems, and biochemical sensors. Sensor selection determines sampling rate, signal quality, power consumption, placement constraints, and potential clinical interpretation.

Wearable sensing is particularly attractive because it enables continuous measurements without restricting patients to clinical facilities. However, wearable data differ from hospital-grade measurements. Consumer devices can experience motion artifacts, sensor displacement, variable skin contact, environmental interference, and inconsistent sampling. Therefore, models trained exclusively on laboratory-quality signals may perform poorly under ambulatory conditions.

Signal preprocessing generally includes filtering, normalization, segmentation, artifact removal, resampling, and feature extraction. For a discrete physiological signal $\{x[n]\}$, a filtered representation can be expressed as $y[n] = \sum_{k=0}^{M-1} h[k]x[n-k]$ where $h[k]$ represents the digital filter coefficients.

For machine-learning applications, feature extraction may include statistical, frequency-domain, morphological, and nonlinear characteristics. However, deep-learning architectures increasingly learn representations directly from raw or minimally processed signals. This reduces manual feature engineering but increases the importance of data volume, annotation quality, model calibration, and robustness.

4.2 Multimodal Data Fusion

Single-sensor monitoring is frequently insufficient for reliable clinical interpretation. Heart rate alone may increase because of exercise, anxiety, fever, pain, or arrhythmia. Combining heart rate with activity, temperature, oxygen saturation, and contextual data can reduce ambiguity.

Multimodal fusion can be represented as $z = F(z_{\text{ECG}}, z_{\text{PPG}}, z_{\text{motion}}, z_{\text{SpO}_2}, z_{\text{context}})$ where each z_i is a learned modality representation and F is a fusion function.

Three principal strategies are commonly used:

1. **Early fusion**, in which raw or low-level features are concatenated before model inference.
2. **Intermediate fusion**, in which separate modality encoders generate latent representations that are combined.
3. **Late fusion**, in which independent models produce predictions that are subsequently aggregated.

Intermediate fusion is particularly suitable for healthcare because it permits modality-specific preprocessing while retaining cross-modal interactions. Transformer-based architectures further enable attention-driven relationships among modalities, although their computational requirements can complicate deployment on resource-constrained IoT devices.

4.3 Time-Series Intelligence

Healthcare monitoring is intrinsically temporal. A single abnormal measurement may have little clinical meaning without its preceding trajectory. Consequently, models must distinguish isolated outliers from persistent trends. Given a multivariate time series

$X_t = [x_{t1}, x_{t2}, \dots, x_{td}]$, a forecasting model estimates $\hat{X}_{t+h} = f(X_{t-L:t}, C_t, \theta)$ where L is the historical window, h is the prediction horizon, C_t represents contextual information, and θ denotes model parameters.

LSTMs and GRUs remain useful where temporal dependencies are moderate and datasets are relatively small. TCNs provide parallelizable convolutional alternatives. Transformers are increasingly attractive for long-range dependencies and multimodal sequences, although the computational burden of attention mechanisms must be considered for edge deployment.

4.4 Anomaly Detection

Continuous monitoring frequently operates without exhaustive event labels. A patient may produce millions of normal observations but only a small number of clinically significant events. This class imbalance makes supervised classification difficult.

Anomaly detection can therefore model normal physiological behavior: $s_t = D(X_t, \hat{X}_t)$ where D represents a distance or reconstruction function. If $s_t > \tau$, the observation is classified as anomalous, where τ is a threshold determined through validation.

Autoencoders are particularly suitable for this purpose because they learn compact representations of normal data. Nevertheless, anomaly thresholds must account for patient-specific variability. A population-level threshold can generate false alarms for patients whose baseline differs substantially from the training population.

4.5 Explainable Artificial Intelligence

The clinical acceptance of AI depends not only on predictive accuracy but also on whether users can understand and appropriately evaluate its output. Explainable AI methods include feature attribution, surrogate models, saliency maps, counterfactual explanations, example-based explanations, and attention-based visualization.

SHAP and LIME are frequently applied to tabular and machine-learning models, while Grad-CAM and related attribution methods are used for neural networks involving spatial or signal representations. Reviews of healthcare XAI have emphasized that explanation quality should not be reduced to visual attractiveness; explanation fidelity, plausibility, user satisfaction, task performance, and correctness are distinct properties [7], [8].

A useful explanation should answer questions such as:

- Which measurements contributed to the prediction?
- How strongly did each measurement contribute?
- What temporal region was important?
- How uncertain is the prediction?
- What change could alter the classification?

An AI system predicting elevated cardiovascular risk, for example, may provide a probability together with the principal contributing features, recent trend, data quality, and confidence interval. Such an interface supports human verification rather than presenting an unexplained numerical score.

4.6 Human-Computer Interaction

HCI determines how AI-generated information enters human cognition and decision-making. Patient-facing interfaces generally prioritize simplicity, reassurance, actionable recommendations, and accessibility. Clinician interfaces require richer contextualization, longitudinal trends, uncertainty indicators, and integration with electronic health records.

A human-centered design process can be represented as \[Requirements \rightarrow Prototype \rightarrow User\ Evaluation \rightarrow Iteration \rightarrow Validation. \]

This process is especially important in healthcare because the consequences of poor interaction design can be clinical rather than merely usability-related. Recent HCI research emphasizes multimodal interaction, including speech, gesture, touch, visual analytics, haptic feedback, and immersive environments [4]. Such interaction mechanisms may be useful for users with visual, motor, or cognitive limitations, but they introduce additional concerns regarding privacy, accessibility, reliability, and user fatigue.

5. System Architecture / Algorithms

An integrated AI-IoT-HCI healthcare monitoring architecture can be organized into seven layers: sensing, communication, edge processing, AI analytics, interoperability/data management, interaction, and governance.




Figure 2. Proposed AI-IoT-HCI healthcare architecture

5.1 Sensing Layer

The sensing layer captures physiological and contextual information. Local preprocessing is increasingly performed before transmission to reduce bandwidth and energy consumption. Compression, denoising, feature extraction, and signal-quality assessment can therefore occur directly on microcontrollers or wearable processors.

5.2 Communication Layer

The communication layer must balance range, bandwidth, latency, reliability, and energy consumption. Bluetooth Low Energy is appropriate for body-area networks, Wi-Fi for high-throughput local communication, cellular and 5G for mobile connectivity, and low-power wide-area technologies for distributed sensors. The choice of protocol should reflect the application rather than assuming that the highest available bandwidth is always desirable.

5.3 Edge Intelligence

Edge computing places computation close to the patient. This reduces communication latency and limits the amount of raw medical data transmitted to remote infrastructure. Recent AI-IoT healthcare research identifies edge intelligence as a response to cloud latency, security fragmentation, and limited real-world evaluation of centralized systems [9].

A lightweight inference model can be represented as $\hat{y} = f_{\theta_q}(Q(x))$, where Q represents quantization and θ_q denotes compressed model parameters.

Model compression may include quantization, pruning, knowledge distillation, low-rank approximation, or architecture redesign. The optimization problem can be formulated as $\min_{\theta} L(\theta) + \lambda_1 C(\theta) + \lambda_2 E(\theta) + \lambda_3 M(\theta)$, where L represents prediction loss, C computational cost, E energy consumption, and M memory requirements.

5.4 Cloud and Federated Intelligence

Cloud infrastructure remains valuable for longitudinal analytics, population-level modeling, model training, and large-scale storage. However, transferring all raw patient data to a centralized server can increase privacy and bandwidth concerns.

Federated learning provides an alternative: $\theta_{t+1} = \sum_{k=1}^K \frac{n_k}{N} \theta_{t+1}^{(k)}$, where K represents participating clients, n_k represents local sample size, and $N = \sum n_k$.

The classical FedAvg mechanism enables local models to be trained without directly transferring raw datasets.

However, federated learning does not automatically guarantee privacy. Model updates may still expose information, and malicious participants may conduct poisoning or inference attacks.

Recent reviews specifically identify data heterogeneity, adversarial attacks, model inversion, scalability, and interoperability as significant issues in healthcare federated learning [5].

5.5 Human-AI Decision Loop

The final architectural component is the decision loop. Instead of $AI \rightarrow \text{Alert}$,

The proposed approach is $AI \rightarrow \text{Confidence} \rightarrow \text{Explanation} \rightarrow \text{Human Review} \rightarrow \text{Action}$.

For high-risk clinical events, the system should expose uncertainty and data quality rather than treating every prediction as equally reliable.

Table 3. Architecture comparison

Architecture	Latency	Privacy	Computational capacity	Energy requirement	Scalability	Suitable use
Cloud-centric	High–moderate	Lower	Very high	Moderate device-side	High	Longitudinal analytics
Edge-centric	Low	Higher	Moderate	Moderate	Moderate	Real-time monitoring
Device-centric	Very low	High	Low	High constraint	High	Simple event detection
Hybrid edge-cloud	Low	High	High	Moderate	High	Clinical monitoring
Federated	Moderate	High	Distributed	Moderate–high	High	Multi-institution learning
Hierarchical FL	Low–moderate	High	Distributed	Moderate	Very high	Large IoMT ecosystems

6. Comparative Analysis

Comparison of AI-enabled healthcare monitoring systems should extend beyond classification accuracy. A model achieving 98% accuracy on a balanced laboratory dataset may be clinically less useful than a model achieving 94% accuracy on heterogeneous ambulatory data if the latter has lower false-negative rates, better calibration, lower latency, and stronger external validation.

Table 4. Comparative analysis of AI methodologies

Method	Data volume	Temporal modeling	Interpretability	Edge suitability	Major limitation
Logistic Regression	Low–medium	Limited	High	Excellent	Linear assumptions
Random Forest	Medium	Indirect	Moderate–high	Good	Model size
XGBoost	Medium	Feature-based	Moderate	Good	Feature engineering
CNN	Medium–high	Local temporal/spatial	Low–moderate	Good after compression	Requires training data
LSTM	Medium–high	Strong	Low	Moderate	Computational cost
TCN	Medium–high	Strong	Moderate	Good	Architecture sensitivity
Transformer	High	Excellent	Low	Limited without compression	High computation
Autoencoder	Medium	Moderate	Low	Good	Threshold selection
GNN	Medium–high	Relational	Low	Moderate	Graph construction
Federated model	Distributed	Depends on base model	Depends on base model	Moderate	Non-IID data

The comparison demonstrates that no single AI or HCI method is universally optimal. Healthcare monitoring requires application-specific combinations. A wearable fall-detection device prioritizes low latency and energy efficiency, whereas an ICU decision-support system prioritizes accuracy, interpretability, interoperability, and clinical validation.

A broader review of next-generation smart healthcare likewise emphasizes that integration problems arise across device interoperability, energy management, data quality, security, and ethics rather than from AI model performance alone [6]. This finding suggests that future benchmarking should adopt system-level evaluation.

Table 5. HCI modalities for smart healthcare

Interaction	Advantages	Limitations	Healthcare examples
Mobile GUI	Familiar and accessible	Screen dependence	RPM
Web dashboard	Rich visualization	Less mobile	Clinical monitoring
Voice	Hands-free	Privacy/noise concerns	Elderly care
Haptic	Immediate feedback	Limited information bandwidth	Alerts
Gesture	Contactless	Recognition errors	Rehabilitation
AR	Contextual visualization	Hardware requirements	Surgery/rehabilitation
VR	Immersive	Motion sickness/accessibility	Therapy/training
Conversational AI	Natural interaction	Hallucination and safety risks	Patient assistance

7. Applications

7.1 Cardiovascular Monitoring

Cardiovascular monitoring represents one of the most mature applications of IoT and AI. Wearable ECG and PPG devices can continuously measure cardiac activity and detect rhythm abnormalities. AI models can identify arrhythmias, estimate heart rate, analyze heart-rate variability, and detect deviations from personalized baselines. CNNs are effective for morphological analysis of ECG segments, while recurrent and transformer architectures can capture temporal dependencies. However, ambulatory ECG analysis presents significant challenges because motion artifacts and electrode variability can mimic pathological patterns.

A clinically relevant system should therefore combine rhythm classification with signal-quality assessment and contextual information. The output should distinguish between "high-confidence abnormal rhythm," "possible abnormality requiring review," and "low-quality signal."

7.2 Diabetes and Metabolic Monitoring

IoT-enabled glucose monitoring can provide continuous or semi-continuous metabolic information. AI models can forecast glucose trajectories and detect potential hypo- or hyperglycemic events. Integration with dietary, activity, sleep, and medication data creates opportunities for personalized prediction.

However, predictive models must address delayed physiological responses and missing observations. A model that predicts glucose levels accurately over short horizons may not necessarily produce clinically meaningful long-term intervention recommendations.

7.3 Respiratory Disease

Respiratory monitoring can use oxygen saturation, respiratory rate, chest motion, airflow, cough sounds, and environmental conditions. Multimodal models are particularly valuable because respiratory deterioration may manifest across several channels simultaneously.

AI can detect deviations from personal baselines and prioritize patients for clinical review. Edge processing is advantageous because respiratory emergencies may require rapid detection and local alerts.

7.4 Elderly Care and Fall Detection

Accelerometers, gyroscopes, cameras, radar, pressure sensors, and ambient sensors can support fall detection and activity recognition. Deep-learning models can classify movement patterns and distinguish falls from normal transitions such as sitting or lying down.

HCI is particularly important because elderly users may have visual, auditory, or cognitive limitations. Interfaces should therefore use simple visual signals, audio prompts, haptic feedback, or caregiver notification mechanisms.

7.5 Neurological Monitoring

IoT systems can monitor tremor, gait, movement symmetry, sleep, and other behavioral indicators relevant to neurological disorders. Longitudinal monitoring may reveal changes that are difficult to capture during short clinical appointments.

The principal challenge is distinguishing disease-related changes from normal day-to-day variability. Personalized baselines and longitudinal models are consequently more appropriate than static population classifiers.

7.6 Rehabilitation

Wearable inertial sensors can quantify range of motion, gait velocity, repetition count, posture, and exercise adherence. AI can transform these measurements into individualized rehabilitation metrics.

HCI extends this capability by providing real-time feedback. Visual, auditory, or haptic feedback can inform users whether an exercise is being performed correctly. Such systems create a closed loop between measurement, interpretation, and behavioral correction.

7.7 Mental Health and Behavioral Monitoring

Smartphones and wearables can capture activity, sleep, mobility, communication patterns, and physiological indicators. AI can identify behavioral changes associated with stress, sleep disruption, or other conditions.

This application presents particularly strong privacy concerns because behavioral data can reveal sensitive information even when explicit medical variables are absent. Consequently, privacy-preserving computation and strong consent mechanisms are essential.

8. Case Study: AI-Assisted Remote Patient Monitoring System

A representative architecture can be constructed for remote monitoring of a patient with cardiovascular and respiratory risk factors. The system includes an ECG/PPG wearable, pulse oximeter, temperature sensor, accelerometer, smartphone gateway, edge inference engine, cloud platform, clinician dashboard, and patient-facing mobile application.




Figure 3. Case-study workflow

The system continuously collects $X_t = [HR_t, SpO_{2,t}, T_t, ECG_t, A_t]$ where (HR_t) denotes heart rate, $(SpO_{2,t})$ oxygen saturation, (T_t) temperature, (ECG_t) electrocardiographic information, and (A_t) acceleration.

The first stage performs signal-quality assessment. Measurements that fail predefined quality criteria are flagged and excluded from clinical inference. The second stage extracts temporal features. The third stage executes an AI classifier.

A risk score may be represented as $R_t = P(Y_t = 1 | X_{t-L:t}, C_t)$ where $(Y_t = 1)$ denotes a clinically relevant event.

The system uses three thresholds $R_t < \tau_1 \rightarrow \text{Normal}$, $\tau_1 \leq R_t < \tau_2 \rightarrow \text{Monitor}$, $R_t \geq \tau_2 \rightarrow \text{Clinical Alert}$.

The intermediate category is important because binary normal/abnormal classification may produce unnecessary alerts. A "monitor" state permits additional data collection and trend confirmation.

For example, a temporary increase in heart rate during walking should not necessarily generate a clinical alert. If the accelerometer indicates inactivity while heart rate remains elevated, the system may increase the posterior probability of an abnormal event.

The HCI layer translates the prediction into different representations for different stakeholders. The patient may receive a simple message such as "Please remain seated while your readings are reassessed." A clinician may see a more detailed visualization containing heart-rate trajectory, oxygen saturation, signal quality, prediction probability, confidence, and relevant historical measurements.

The system should avoid presenting the AI probability as a definitive diagnosis. Instead, the clinician interface should clearly distinguish measurement, prediction, explanation, and recommended action.

Table 6. Case-study component design

Component	Technology	Processing
ECG/PPG wearable	BLE-enabled sensor	Local filtering
Pulse oximeter	Optical sensor	SpO ₂ estimation
IMU	Accelerometer/gyroscope	Activity recognition
Smartphone	Android/iOS	Gateway + preprocessing
Edge AI	TensorFlow Lite/PyTorch Mobile-class deployment	Real-time inference
Cloud	GPU/CPU infrastructure	Training and longitudinal analysis
Interoperability	HL7 FHIR	Clinical data exchange
Dashboard	Web application	Clinician visualization
Patient interface	Mobile application	Feedback
Privacy	Encryption + federated learning	Data protection

FHIR is particularly relevant to interoperability because it provides standardized resources for exchanging healthcare information electronically [10]. Its adoption can reduce dependence on proprietary data structures, although semantic interoperability remains more difficult than simple syntactic data exchange.

9. Tools, Technologies and Implementation

Implementation of an AI-IoT healthcare platform requires coordinated hardware, communication, software, data, AI, and interface technologies.

MQTT is particularly suitable for constrained IoT environments because it supports lightweight publish-subscribe communication. Healthcare deployments, however, require authentication, authorization, encryption, certificate management, device identity, and secure update mechanisms in addition to protocol selection.

FHIR provides a standardized framework for healthcare information exchange, including structured resources suitable for interoperability between AI systems and clinical information systems [10]. Nevertheless, FHIR does

not itself solve all interoperability challenges. Differences in terminology, coding systems, workflow semantics, device metadata, and institutional practices remain significant.

Table 7. Representative implementation technologies

Layer	Technologies
Microcontrollers	ESP32, STM32, Nordic nRF series
Embedded AI	TensorFlow Lite Micro, CMSIS-NN
Mobile	Android, iOS
Edge	Raspberry Pi, NVIDIA Jetson, industrial gateways
AI frameworks	PyTorch, TensorFlow, scikit-learn
Data processing	Python, NumPy, Pandas, Apache Spark
Databases	PostgreSQL, InfluxDB, MongoDB
Messaging	MQTT, AMQP, WebSockets
Interoperability	HL7 FHIR, DICOM
Cloud	AWS, Azure, Google Cloud
Visualization	Grafana, Plotly, custom dashboards
Security	TLS, OAuth 2.0, PKI
Federated learning	Flower, TensorFlow Federated, PySyft
Explainability	SHAP, LIME, Captum
Deployment	Docker, Kubernetes

9.1 Dataset Selection

Public datasets are essential for reproducible research. MIMIC-IV is one of the major critical-care resources used for healthcare AI research. Version 3.0 contains data associated with 364,627 individuals, 546,028 hospitalizations, and 94,458 ICU stays, with hospital and ICU modules derived from different clinical information sources [11]. Access to credentialed versions is controlled through PhysioNet policies.

The eICU Collaborative Research Database provides a complementary multicentre dataset containing more than 200,000 ICU admissions from U.S. hospitals [12]. Its multicentre structure makes it useful for studying generalization, although missingness can vary between institutions and interfaces.

Table 8. Representative healthcare datasets

Dataset	Primary domain	Characteristics	Main limitation
MIMIC-IV	ICU/EHR	Large longitudinal clinical database	Credentialed access; retrospective
eICU-CRD	ICU	Multicenter	Interface-dependent missingness
PhysioNet waveform datasets	Physiological signals	ECG/PPG and related signals	Dataset-specific populations
UCI HAR	Human activity	Wearable motion	General activity rather than clinical disease
MIT-BIH Arrhythmia	ECG	Annotated arrhythmias	Limited population diversity
PTB-XL	ECG	Large clinical ECG collection	ECG-focused
WESAD	Stress/activity	Wearable multimodal data	Experimental setting

Dataset choice should be aligned with the research question. A model intended for wearable arrhythmia detection should not be validated solely on ICU data because ICU populations and signal characteristics differ considerably from ambulatory populations.

10. Performance Evaluation

Evaluation of smart healthcare monitoring systems requires a multidimensional framework. Conventional accuracy is inadequate because healthcare datasets frequently exhibit class imbalance and asymmetric costs.

For binary classification $\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}}$ $\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}$ $\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}$ $F1 = 2 \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$.

For clinical monitoring, sensitivity may be more important than accuracy for dangerous events, while excessive false positives can generate alarm fatigue. Therefore, evaluation should include sensitivity, specificity, precision, negative predictive value, AUROC, AUPRC, calibration, and decision-curve analysis.

Table 9. Performance dimensions

Dimension	Metric	Relevance
Classification	Accuracy	General discrimination
Detection	Sensitivity	Missing-event risk
Detection	Specificity	False-alarm control
Precision	PPV	Alert usefulness
Ranking	AUROC	Discrimination
Imbalanced data	AUPRC	Positive-event detection
Probability	Brier score	Calibration
Edge deployment	Latency	Real-time response
Energy	Joules/inference	Battery life
Communication	Bytes/event	Network efficiency
HCI	Task completion time	Usability
Clinical use	Alert burden/hour	Workflow impact
Trust	User confidence	Human-AI adoption

Calibration is particularly important. A model producing a risk score of 0.8 should ideally correspond to approximately 80% event probability within an appropriately defined population and time horizon. High discrimination with poor calibration can lead to inappropriate clinical decisions.

10.1 Cross-Patient and Cross-Institution Validation

Randomly splitting observations from the same patient into training and test sets can cause information leakage. In longitudinal healthcare datasets, patient-level splitting is therefore preferable $D = D_{\text{train}} \cup D_{\text{validation}} \cup D_{\text{test}}$ With $P_i \cap P_j = \emptyset$ for patient sets across the partitions.

External validation is even more important. A model trained in one hospital should be tested on another institution where patient demographics, devices, workflows, and measurement practices differ.

The multicenter nature of resources such as eICU makes them useful for examining this issue [12]. Federated-learning studies similarly emphasize non-IID data and institutional heterogeneity as important challenges [5].

10.2 HCI Evaluation

HCI evaluation should include usability, workload, trust, comprehension, error recovery, and task performance. A healthcare dashboard that requires excessive navigation may delay response despite excellent visualization quality. Human-centered evaluation can therefore define: $U = w_1P + w_2T + w_3C - w_4L$ where P represents task performance, T trust/calibration, C comprehension, and L cognitive load.

The weights should be determined according to the application rather than assumed universally.

Explainability evaluation also requires special attention. Systematic reviews have shown that healthcare XAI studies often use inconsistent measures of explanation effectiveness, including satisfaction, trust, task performance, and correctness [7]. Consequently, future benchmarking should report standardized human-centered metrics alongside technical metrics.

11. Challenges and Limitations

The principal limitation of AI-enabled IoT healthcare is not the absence of algorithms but the difficulty of deploying algorithms reliably in heterogeneous real-world environments.

11.1 Data Quality and Sensor Reliability

Wearable sensors are affected by motion artifacts, battery degradation, environmental interference, poor placement, and individual anatomical differences. AI models may unintentionally learn sensor-specific artifacts rather than physiological patterns. This creates a distribution-shift problem $\mathbb{P}_{\text{train}}(X,Y) \neq \mathbb{P}_{\text{deployment}}(X,Y)$. A model trained on hospital-grade signals may therefore degrade when deployed on consumer wearables.

11.2 Dataset Bias

Healthcare datasets frequently underrepresent certain demographic groups, geographic regions, age categories, and socioeconomic populations. A high-performing model can therefore exhibit unequal error rates across subgroups. Fairness should be evaluated explicitly $\Delta_{\text{TPR}} = \text{TPR}_A - \text{TPR}_B$ and $\Delta_{\text{FPR}} = \text{FPR}_A - \text{FPR}_B$. However, fairness metrics can conflict with each other. Therefore, fairness cannot be reduced to a single numerical criterion.

11.3 Explainability Limitations

XAI can improve transparency but does not guarantee correctness. An explanation can be faithful to a model while still reflecting spurious correlations. Reviews have emphasized the continuing gap between explainability research and validated clinical implementation [7], [8].

11.4 Privacy and Security

Healthcare IoT creates a large attack surface consisting of sensors, gateways, mobile applications, APIs, cloud infrastructure, and clinical systems. Threats include device compromise, eavesdropping, ransomware, model poisoning, membership inference, model inversion, and denial of service.

Federated learning reduces the need to centralize raw data but introduces its own security vulnerabilities [5]. Differential privacy, secure aggregation, trusted execution environments, and homomorphic encryption can strengthen privacy but may increase computational overhead or reduce model utility.

11.5 Interoperability

Healthcare ecosystems contain legacy EHRs, proprietary medical devices, laboratory systems, imaging systems, mobile applications, and cloud platforms. FHIR provides an important interoperability foundation [10], but implementation differences and semantic inconsistencies remain.

11.6 Energy Constraints

Continuous sensing and communication can substantially reduce wearable battery life. The energy budget can be represented as $E_{\text{total}} = E_{\text{sense}} + E_{\text{compute}} + E_{\text{communication}} + E_{\text{idle}}$. Communication can dominate the energy budget in some systems. Local feature extraction and event-triggered transmission can therefore provide significant savings.

11.7 Alarm Fatigue

False-positive alerts represent a major operational problem. If a system generates too many low-value alerts, clinicians may become desensitized. The objective should therefore be maximizing clinical value rather than maximizing sensitivity alone.

11.8 Human Over-Reliance

The opposite problem is automation bias. Users may accept AI predictions without sufficient verification. Human-AI systems should therefore communicate uncertainty, data quality, and model limitations.

12. Ethical and Societal Considerations

AI-enabled healthcare monitoring raises ethical issues extending beyond conventional data protection. Continuous monitoring changes the relationship between patients and healthcare systems because data can be collected outside explicit clinical encounters.

12.1 Informed Consent

Consent must address what data are collected, how frequently they are collected, where they are stored, who can access them, how AI models use them, and whether data may be reused for research or model training. Consent for a single clinical interaction may not adequately cover continuous longitudinal monitoring.

12.2 Privacy

Physiological data are sensitive, but behavioral and contextual information can be equally revealing. Location, sleep, activity, voice, and interaction patterns may expose lifestyle and health characteristics.

Privacy should therefore be incorporated at the architecture level rather than added after deployment. Techniques such as local processing, federated learning, differential privacy, secure aggregation, encryption, and data minimization should be considered according to the threat model.

12.3 Transparency and Accountability

When AI contributes to a healthcare decision, responsibility should remain identifiable. A monitoring system should maintain auditable records of the input data, model version, prediction, confidence, explanation, human action, and subsequent outcome.

12.4 Equity

Smart healthcare may unintentionally widen healthcare inequalities if effective systems require expensive smartphones, reliable broadband, modern wearables, or high levels of digital literacy. Rural and low-resource populations may experience poorer connectivity and limited access to technical support. Thus, equitable design should consider device affordability, offline operation, multilingual interfaces, accessibility, low-bandwidth communication, and community healthcare infrastructure.

12.5 Trustworthy Human-AI Interaction

Trust should be calibrated rather than maximized. Excessive distrust prevents useful technology adoption, while excessive trust encourages automation bias. The appropriate goal is evidence-based reliance. Recent XAI research emphasizes that explanation effectiveness must be evaluated according to stakeholder objectives rather than assuming that every explanation increases trust or clinical usefulness [7]. Human-centered design is therefore essential for trustworthy healthcare AI.

13. Future Research Directions

The next generation of AI-enabled IoT healthcare systems is likely to move beyond isolated predictive models toward adaptive, multimodal, personalized, and human-collaborative intelligence.

13.1 Multimodal Foundation Models

Large foundation models capable of processing physiological signals, clinical text, imaging, laboratory data, and contextual information could provide unified representations of patient state. However, medical foundation models require domain-specific validation, hallucination control, calibration, privacy safeguards, and efficient deployment.

13.2 Personalized AI

Population-level models may fail to capture individual physiological baselines. Future systems should combine population knowledge with personalized adaptation $\theta_i = \theta_{\text{global}} + \Delta\theta_i$. Here, θ_i represents patient-specific parameters and $\Delta\theta_i$ represents personalization. Meta-learning, continual learning, and personalized federated learning provide promising mechanisms.

13.3 Continual Learning and Model Drift

Healthcare data distributions change over time due to ageing, disease progression, medication changes, device replacement, and environmental factors. Static models therefore require continuous monitoring. A future system should detect drift $D(P_t(X), P_{t-1}(X)) > \delta$ and initiate model reassessment when divergence exceeds a predefined threshold.

13.4 Digital Twins

Patient digital twins could combine physiological observations, medical history, behavioral data, and predictive models into dynamic computational representations. Such systems could simulate possible interventions and support personalized decision-making. However, digital twins should not be interpreted as exact representations of human physiology. Uncertainty propagation and validation remain major research problems.

13.5 Privacy-Preserving Edge Intelligence

Future architectures should combine edge inference with federated learning, secure aggregation, and adaptive privacy mechanisms. The objective is not simply to avoid transmitting data but to minimize the overall exposure of sensitive information.

13.6 Human-AI Collaborative Decision-Making

AI should increasingly be designed as a collaborative assistant rather than an autonomous replacement for clinical judgment. Research should examine how clinicians modify decisions when AI provides uncertainty, explanations, alternative hypotheses, and counterfactual information.

13.7 Healthcare 5.0 and Human-Centered IoT

Human-centered IoT research indicates that future healthcare systems must incorporate the human element explicitly rather than treating users as passive recipients of technology [3]. This suggests a transition from "AI for healthcare" toward "human-AI-IoT healthcare ecosystems."

Table 10. Priority research directions

Research direction	Current limitation	Proposed development
Edge AI	Limited resources	Adaptive model compression
Federated learning	Non-IID data	Personalized/hierarchical FL
XAI	Inconsistent evaluation	Human-centered explanation benchmarks
Wearables	Sensor artifacts	Context-aware sensor fusion
Foundation models	High computation	Medical edge foundation models
Digital twins	Validation uncertainty	Probabilistic patient twins
HCI	Generic interfaces	Adaptive stakeholder-specific interfaces
Security	Expanding attack surface	Zero-trust IoMT architectures
Interoperability	Fragmented systems	FHIR-native AI pipelines

14. Conclusion

Artificial intelligence-based smart healthcare monitoring represents the convergence of sensing, communication, computational intelligence, clinical informatics, and human-centered interaction. IoT provides the physical and communication infrastructure through which physiological and contextual information can be continuously acquired, while AI transforms high-volume heterogeneous data into predictions, classifications, anomaly scores, and longitudinal insights. HCI provides the essential mechanism through which these computational outputs become understandable and actionable for patients, caregivers, and healthcare professionals.

The review demonstrates that the field has moved beyond simple cloud-connected monitoring toward distributed edge intelligence, multimodal learning, federated architectures, explainable AI, and human-centered Healthcare 5.0 paradigms. Nevertheless, technical progress has not eliminated the principal barriers to clinical deployment. Sensor artifacts, dataset bias, distribution shift, energy limitations, interoperability problems, privacy threats, model opacity, alarm fatigue, and inadequate external validation continue to restrict real-world effectiveness.

A central conclusion is that predictive accuracy should not be treated as the sole criterion of success. A clinically useful AI-IoT monitoring system must simultaneously achieve acceptable predictive performance, calibration, latency, energy efficiency, security, interoperability, usability, and clinical workflow compatibility. Its explanations must support human reasoning rather than simply expose model internals, and its alerts must prioritize clinically meaningful events rather than maximizing detection counts.

The future of smart healthcare monitoring is therefore unlikely to be determined by a single algorithm or sensing technology. Instead, progress will emerge from integrated architectures combining multimodal sensing, edge intelligence, privacy-preserving learning, interoperable clinical data, adaptive models, and human-centered interaction. Personalized and federated learning can help address patient heterogeneity and data governance; explainable AI can improve transparency; digital twins can support individualized simulation; and multimodal foundation models may provide more comprehensive representations of patient state. These advances must, however, be accompanied by rigorous prospective validation, standardized benchmarks, cybersecurity engineering, ethical governance, and equitable deployment.

Ultimately, AI-enabled IoT healthcare should be understood as a **human-AI cyber-physical healthcare ecosystem**. Its success depends not merely on whether an algorithm can predict an event, but on whether the complete system can sense reliably, reason responsibly, communicate clearly, preserve privacy, integrate with clinical practice, and enable humans to make better-informed healthcare decisions.

References

- [1] M. Alshamrani, "IoT and artificial intelligence implementations for remote healthcare monitoring systems: A survey," *Journal of King Saud University—Computer and Information Sciences*, vol. 34, no. 8, pp. 4687–4701, 2022, doi: 10.1016/j.jksuci.2021.06.005. ScienceDirect
- [2] A. Motwani, P. K. Shukla, and M. Pawar, "Ubiquitous and smart healthcare monitoring frameworks based on machine learning: A comprehensive review," *Artificial Intelligence in Medicine*, vol. 134, Art. no. 102431, 2022, doi: 10.1016/j.artmed.2022.102431. ScienceDirect
- [3] S. Rashid and A. Nemati, "Human-centered IoT-based health monitoring in the Healthcare 5.0 era: Literature descriptive analysis and future research guidelines," *Discover Internet of Things*, 2024. Springer
- [4] E.-L. Pop and A. Rațiu, "Human-Computer Interaction in Artificial Intelligence with Applications in Healthcare: A Review," *Machine Learning and Intelligent Systems*, 2024, doi: 10.3233/FAIA241213. Sage Journals
- [5] S. R. Abbas, Z. Abbas, A. Zahir, and S. W. Lee, "Federated Learning in Smart Healthcare: A Comprehensive Review on Privacy, Security, and Predictive Analytics with IoT Integration," *Healthcare*, vol. 12, no. 24, Art. no. 2587, 2024. PubMed Central (PMC)
- [6] "IoT for Next-Generation Smart Healthcare: A Comprehensive Survey," *IEEE Internet of Things Journal*, vol. 12, no. 16, pp. 32616–32639, 2025, doi: 10.1109/JIOT.2025.3570188. IEEE Xplore
- [7] J. Jung, H. Lee, H. Jung, and H. Kim, "Essential properties and explanation effectiveness of explainable artificial intelligence in healthcare: A systematic review," *Heliyon*, vol. 9, no. 5, Art. no. e16110, 2023, doi: 10.1016/j.heliyon.2023.e16110. PubMed
- [8] H. W. Loh, C. P. Ooi, S. Seoni, P. D. Barua, F. Molinari, and U. R. Acharya, "Application of explainable artificial intelligence for healthcare: A systematic review of the last decade (2011–2022)," *Computer Methods and Programs in Biomedicine*, vol. 226, Art. no. 107161, 2022, doi: 10.1016/j.cmpb.2022.107161. PubMed
- [9] "From Sensors to Edge Intelligence: A Critical Review of AI-Enabled IoT for Healthcare," in *Proc. 6th International Conference on Electronics and Sustainable Communication Systems (ICESC)*, 2025, doi: 10.1109/ICESC65114.2025.11212327. IEEE Xplore
- [10] HL7 International, "FHIR Overview, Release 5," *Fast Healthcare Interoperability Resources*, 2023. HL7
- [11] A. E. W. Johnson, L. Bulgarelli, T. Pollard, S. Horng, L. A. Celi, and R. G. Mark, "MIMIC-IV," *PhysioNet*, version 3.0, 2024, doi: 10.13026/07hj-2a80. PhysioNet
- [12] T. J. Pollard, A. E. W. Johnson, J. D. Raffa, L. A. Celi, O. Badawi, and R. G. Mark, "The eICU Collaborative Research Database, a freely available multi-center database for critical care research," *Scientific Data*, vol. 5, 2018, Art. no. 180178. PhysioNet
- [13] D. Jin, E. Sergeeva, W.-H. Weng, G. Chauhan, and P. Szolovits, "Explainable deep learning in healthcare: A methodological survey from an attribution view," *WIREs Mechanisms of Disease*, vol. 14, no. 3, Art. no. e1548, 2022, doi: 10.1002/wsbm.1548. PubMed

- [14] B. H. M. van der Velden, H. J. Kuijf, K. G. A. Gilhuijs, and M. A. Viergever, "Explainable artificial intelligence (XAI) in deep learning-based medical image analysis," *Medical Image Analysis*, vol. 79, Art. no. 102470, 2022, doi: 10.1016/j.media.2022.102470. PubMed
- [15] "Explainable, trustworthy, and ethical machine learning for healthcare: A survey," *Computers in Biology and Medicine*, 2022. PubMed
- [16] S. Ali, F. Akhlaq, A. S. Imran, Z. Kastrati, S. M. Daudpota, and M. Moosa, "The enlightening role of explainable artificial intelligence in medical & healthcare domains: A systematic literature review," *Computers in Biology and Medicine*, vol. 166, Art. no. 107555, 2023, doi: 10.1016/j.compbiomed.2023.107555. PubMed
- [17] "Advancing explainable AI in healthcare: Necessity, progress, and future directions," PubMed-indexed review, 2025. PubMed
- [18] "Explainable Artificial Intelligence in Healthcare: Current Landscape, Challenges, and Future Directions," PubMed-indexed review, 2026. PubMed
- [19] A. E. W. Johnson, T. J. Pollard, L. H. Shen, L. H. Lehman, M. Feng, M. Ghassemi, et al., "MIMIC-III, a freely accessible critical care database," *Scientific Data*, vol. 3, Art. no. 160035, 2016.
- [20] A. E. W. Johnson, T. J. Pollard, J. S. Shen, L. H. Lehman, M. Feng, M. Ghassemi, et al., "MIMIC-IV (version 1.0)," *PhysioNet*, 2021, doi: 10.13026/s6n6-xd98. PhysioNet
- [21] T. J. Pollard, A. E. W. Johnson, J. D. Raffa, L. A. Celi, O. Badawi, and R. G. Mark, "eICU Collaborative Research Database," *PhysioNet*, version 2.0, 2019, doi: 10.13026/C2WM1R. PhysioNet
- [22] "IoT-Based Smart Healthcare Monitoring Systems: A Comprehensive Review of Technologies, Challenges, and Innovations," in *Proc. IEEE 2nd International Conference on Computing, Applications and Systems (COMPAS)*, 2025, doi: 10.1109/COMPAS67506.2025.11381748. IEEE Xplore

