Patient Health Response in Emergent and Secure Habitats for Connected Healthcare

D2.1 State Of The Art

Work package	WP2: Use-case Requirements & Architectural Design
Type	Document
Contractual date of deliverable	30/12/2025
Actual date of deliverable	04/03/2026
Dissemination level	Public
Lead partner	Dextromedica S.L.
Status	Submitted

Document history			
Version	**Date**	**Author**	**Description**
0.1	28/02/2026	Saint Patrick Technology	First Version
0.2	02/03/2026	Saint Patrick Technology	Final version

Document review		
Date	**Reviewer**	**Description**
01/03/2026	Makila AI	Document review

Partners involved		
Company	**Name and surname**	**Email address**
TrianecT	Maryse van 't Klooster	m.a.vantklooster@trianect.com
Amsterdam UMC	Jet Veltink	j.m.veltink@amsterdamumc.nl
HIGOE	Manzoor Ahmed	manzoor.ahmed@higoe.com
DEXTROMEDICA	Fabio Román	froman@dextromedica.com
SAINT PATRICK TECHNOLOGY	Jesús Arce	jesus.arce@tecnologiaypersonas.net
IONIDE	David Jandara	david.jandara@ionide.es
MEDrecord	Joanna Morozowska	joanna@medrecord.io
ISEP	Eva Maia, Ivone Amorim	Egm@isep.ipp.pt, ifa@isep.ipp.pt

PHRESH

Table of Contents

PHRESH

Abbreviations

Abbreviation	**Definition**
AI	Artificial Intelligence
AIA	Artificial Intelligence Act (European Union)
APWA	Arterial Pulse Wave Analysis
ASR	Automatic Speech Recognition
BP	Blood Pressure
CNN	Convolutional Neural Network
DP	Differential Privacy
ECG	Electrocardiography
EEG	Electroencephalography
EHR	Electronic Health Record
ER	Emergency Room
FHIR	Fast Healthcare Interoperability Resources
FL	Federated Learning
FPP	Full Project Proposal
F-RNN-LSTM	Fusion Recurrent Neural Network – Long Short-Term Memory
HE	Homomorphic Encryption
HL7	Health Level Seven
HR	Heart Rate
IoT	Internet of Things
ITEA	Information Technology for European Advancement
LLM	Large Language Model
LRP	Layer-wise Relevance Propagation
ML	Machine Learning
MPC / SMPC	(Secure) Multi-Party Computation
PET	Privacy-Enhancing Technologies
PHRESH	Patient Health Response in Emergent and Secure Habitats
PLCO	Prostate, Lung, Colorectal and Ovarian Cancer Model
PPO	Proximal Policy Optimization
PQC	Post-Quantum Cryptography
RNN	Recurrent Neural Network
RL	Reinforcement Learning
RR	Respiratory Rate
SASE	Secure Access Service Edge
SpO_2	Peripheral Capillary Oxygen Saturation
SotA	State of the Art
TREAT	Thoracic Research Evaluation and Treatment
WP	Work Package
XAI	Explainable Artificial Intelligence
ZKP	Zero-Knowledge Proof

PHRESH

1 Executive summary

1.1 Purpose, scope and key outcomes

The purpose of this State of the Art (SotA) analysis is to provide a structured and comprehensive overview of the current market and technological landscape relevant to the PHRESH project. The SotA establishes a clear baseline of existing solutions, ongoing developments, and remaining gaps in digital and connected healthcare, with a particular focus on emergency and pre-hospital scenarios. This analysis demonstrates the consortium's in-depth understanding of both industrial and academic advancements and provides the technical justification for the innovations proposed within PHRESH.

The SotA addressed in this document is structured around two complementary perspectives: the **Market State of the Art**, which analyses trends, growth, and value-chain dynamics in key digital health domains, and the **Technological State of the Art**, which examines current capabilities and limitations in terms of accuracy, robustness, privacy and security, and connectivity. Together, these perspectives define the current context of the digital healthcare sector and highlight the unmet needs that PHRESH aims to address.

From a market perspective, the SotA identifies strong and sustained growth in three strategic areas directly relevant to PHRESH. Privacy-Enhancing Technologies (PET) are experiencing rapid adoption, driven by increasing regulatory pressure, widespread IoT deployment, and rising cybersecurity threats. Artificial Intelligence in healthcare represents one of the fastest-growing digital health segments, supported by demographic changes, the growing burden of chronic diseases, and advances in data-driven medical technologies. In parallel, the health sensor market continues to expand due to the demand for continuous monitoring, remote patient management, and real-time health data acquisition. Despite this growth, current market offerings remain fragmented and insufficiently adapted to high-stakes, time-critical emergency care environments.

From a technological perspective, the SotA analyses the current state of precision and robustness in physiological sensing, artificial intelligence, and real-time data processing. Existing wearable and non-invasive monitoring solutions still face significant limitations in terms of clinical accuracy, regulatory maturity, usability, and their ability to simultaneously capture multiple vital signs under real-world emergency conditions. Continuous blood pressure monitoring, reliable dry-electrode EEG for stroke triage, and robust voice-based medical reporting remain partially solved problems, particularly when deployed in noisy, mobile, and resource-constrained pre-hospital environments such as ambulances.

In the field of artificial intelligence and machine learning, the SotA reviews state-of-the-art diagnostic models, clinical risk assessment tools, reinforcement learning approaches for

ITEA4

personalized medical assistance, and explainable AI techniques. While high diagnostic performance has been demonstrated in controlled settings, major challenges persist with respect to real-time deployment, multimodal data fusion, interpretability, and clinical trust. These challenges limit the effective integration of AI-driven decision support into emergency medical workflows.

The scope of this SotA further includes privacy, cybersecurity, and connectivity aspects, which are critical enablers for trustworthy digital healthcare. Current cryptographic systems face emerging risks from quantum computing, while advanced privacy-enhancing technologies such as federated learning, secure multiparty computation, homomorphic encryption, and zero-knowledge proofs are not yet widely adopted in healthcare due to performance and resource constraints. At the same time, existing connectivity solutions between ambulances and hospitals remain fragmented, often limited to partial or unstructured data exchange, resulting in the well-known “patient black box” situation upon arrival at the emergency room.

The key outcomes of this SotA analysis are threefold. First, it identifies clear market and technological gaps, particularly the absence of integrated, real-time, and clinically robust solutions capable of supporting emergency medical decision-making across organizational and technological boundaries. Second, it highlights the limitations of current approaches when exposed to realistic pre-hospital conditions, including mobility, noise, time pressure, and constrained connectivity. Third, it establishes a solid and transparent baseline against which the PHRESH project can be positioned, clearly demonstrating its differentiation through the integration of advanced sensing, AI-driven decision support, privacy-by-design mechanisms, and next-generation connectivity into a unified, hands-free, and operationally efficient solution tailored for critical emergency care scenarios.

2 Introduction

2.1 Background and context of the deliverable

Digital and connected healthcare systems are increasingly transforming the way medical services are delivered, particularly in time-critical scenarios such as emergency and pre-hospital care. Advances in sensing technologies, wireless communication, and data-driven medical analytics have enabled the continuous collection and processing of physiological data, supporting faster and more informed clinical decision-making. However, despite these advances, significant limitations remain in integration, reliability, and real-time availability of health data across the continuum of care, especially in highly dynamic environments such as ambulances and emergency medical services.

In emergency and pre-hospital settings, medical decisions must often be taken under extreme pressure, based on incomplete or fragmented information, and in the presence of noisy signals, patient movement, and constrained connectivity. In such contexts, **accuracy and robustness** of physiological measures are essential to avoid misinterpretation of patient status, while **real-time data acquisition and transmission** are critical to enable early diagnosis, preparation of hospital teams, and continuity of care upon patient arrival. At the same time, the sensitive nature of medical data requires strict guarantees of **privacy and security**, particularly when data are transmitted across organizational boundaries or processed using advanced digital technologies. **Reliable connectivity and interoperability** between pre-hospital and in-hospital systems are therefore fundamental prerequisites for effective emergency care.

The State of the Art (SotA) analysis presented in this deliverable is a core element of the PHRESH Full Project Proposal (FPP), as defined in Section 2.3.1 of the FPP Annex. Its objective is to systematically review the current market and technological landscape relevant to the project, covering both industrial solutions and academic research. By identifying existing capabilities, limitations, and open challenges, the SotA provides the necessary context to justify the technical objectives, architectural choices, and innovation pathways proposed within PHRESH.

In accordance with ITEA rules and guidelines, this SotA constitutes **public information**. The content is therefore limited to non-confidential descriptions of existing technologies, standards, and research results, without disclosure of proprietary data or sensitive implementation details. As such, this deliverable is intended to be accessible to a broad audience, including evaluators, stakeholders, and the wider research and innovation community, while still providing a rigorous and technically sound foundation for the PHRESH project.

2.2 Relation to the project objectives and work package

The State of the Art (SotA) analysis provides the technical foundation upon which the objectives and structure of the PHRESH project are defined. By systematically reviewing current market solutions and technological developments in connected healthcare, the SotA identifies the limitations of existing approaches and clarifies the unmet needs that PHRESH addresses. These findings directly ground the project's technical objectives, ensuring that they are both relevant to real-world emergency care scenarios and clearly positioned beyond the current state of the art.

In particular, the SotA highlights that existing solutions tend to address individual aspects of emergency healthcare in isolation—such as sensing, analytics, connectivity, or security—without providing an integrated, real-time, and clinically robust system. This observation informs PHRESH's overarching objective of developing a unified solution that combines advanced physiological sensing, AI-driven decision support, secure data handling, and next-generation connectivity within a coherent architecture tailored to pre-hospital and emergency environments.

The results of the SotA also guide key architectural and technological choices within PHRESH. The identified limitations in accuracy and robustness of current sensing technologies justify the project's focus on improved multimodal data acquisition and fusion, particularly under mobile and noisy conditions. Similarly, the analysis of existing AI and machine learning approaches reveals the need for decision-support systems that operate in real time, integrate heterogeneous data sources, and provide explainable outputs that can be trusted by medical professionals. In the areas of privacy, security, and connectivity, the SotA underlines the necessity of privacy-by-design mechanisms and seamless, interoperable communication between ambulances and hospital systems.

These insights are reflected in the structure of the PHRESH work plan and its Work Packages (WPs). Activities related to **sensor technologies and physiological monitoring** are addressed in the data acquisition–oriented WPs, which focus on reliable, continuous, and clinically relevant capture of vital signs in pre-hospital settings. The **artificial intelligence and machine learning** findings of the SotA are translated into analytics and decision-support WPs, where advanced models for diagnosis, risk assessment, and personalized assistance are developed and validated, with explicit attention to explainability and real-time performance. Challenges identified in the SotA regarding **privacy and security** are addressed in dedicated cryptography and privacy-enhancing technology WPs, which aim to ensure compliant, secure, and future-proof data processing across the PHRESH platform. Finally, the connectivity-related gaps highlighted by the SotA are tackled in the integration and demonstration of WPs, which focus on interoperable data exchange, low-latency communication, and end-to-end validation in realistic emergency scenarios.

Overall, this SotA establishes a clear transition from “what exists today” to “what PHRESH must develop.” By explicitly mapping current technological limitations to concrete project objectives and work package activities, the SotA ensures coherence between the identified state of the art, the proposed innovations, and the planned implementation strategy. This alignment strengthens the technical credibility of PHRESH and provides a transparent rationale for its contribution beyond the current state of the art in connected emergency healthcare.

3 Current Technological Landscape

3.1 Industry and academic approaches in connected healthcare

The connected healthcare landscape is currently shaped by a clear divide between commercial, market-ready solutions and academic research initiatives. Commercial systems prioritize usability, scalability, and regulatory compliance, whereas academic research focuses on methodological innovation, advanced analytics, and novel sensing approaches. This section provides an overview of both domains and identifies the technological and operational gaps that PHRESH aims to address through an integrated, end-to-end approach.

3.1.1 Overview of Commercial Solutions

Wearable Technologies and RPM Platforms

The commercial sector has successfully transitioned remote patient monitoring (RPM) from a niche medical application to a mass-market phenomenon. Current industry approaches can be categorized into two main streams:

- **Consumer-Grade Wellness Devices:** Tech giants (e.g., Apple, Garmin, Fitbit) dominate the market with wrist-worn wearables capable of tracking standard vital signs such as Heart Rate (HR), basic Sleep Staging, and intermittent SpO_2. While these devices excel in usability and battery life, they often lack the raw data accessibility required for medical-grade research and clinical decision support. The data is frequently siloed within proprietary ecosystems, limiting interoperability with Hospital Information Systems (HIS).
- **Medical-Grade Remote Monitoring:** Specialized vendors offer FDA/CE-marked devices for specific chronic conditions (e.g., continuous glucose monitors, Holter monitors). These solutions provide high accuracy but typically operate as "single-point" solutions with rigid, closed-loop architectures. They rarely integrate multimodal data streams—such as correlating physiological vitals with acoustic lung sounds or environmental context—restricting their utility for complex comorbidities like Cardiovascular Disease (CVD) and Pulmonary conditions, which are central to the PHRESH Turkish , United Kingdom and Canadian use cases.

Connectivity and Cloud Platforms

Commercially available RPM platforms predominantly rely on standard Bluetooth bridging to patient smartphones. While cost-effective, this architecture introduces reliability issues in pre-hospital and emergency settings where continuous, real-time data streaming is critical. Current "Connected Ambulance" solutions often provide only snapshot transmission (e.g., sending a static

12-lead ECG PDF) rather than a live, continuous telemetry stream required for time-sensitive interventions like stroke triage.

3.1.2 Overview of Academic Research

Advanced Algorithmic Approaches

Academic literature demonstrates a strong focus on Deep Learning (DL) and Machine Learning (ML) techniques that far outpace commercial implementations. Research is heavily invested in:

- **Deep Learning for Diagnostics:** The use of Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) for automated anomaly detection in complex signals (e.g., arrhythmia detection from PPG, lung sound classification) is well-established in in silico studies.
- **Privacy-Preserving AI:** Academic prototypes frequently explore advanced cryptographic methods such as Homomorphic Encryption (HE) and Federated Learning (FL) to address data privacy. However, these computationally intensive methods are rarely deployed on resource-constrained edge devices in commercial settings due to latency and power limitations.

Novel Sensing Modalities

Research laboratories are actively developing "invisible" or passive sensing technologies, such as radar-based vital sign monitoring and high-fidelity acoustic arrays for lung sound separation. Unlike commercial wearables, these prototypes prioritize signal fidelity and novel biomarkers over form factor or battery efficiency.

3.1.3 Differences between Market-Ready Solutions and Research Prototypes

The primary divergence between industry and academia lies in the trade-off between robustness and innovation:

Feature	*Market-Ready Commercial Solutions*	*Academic Research Prototypes*	*PHRESH Gap Addressal*
Data Access	Processed, aggregated data (Black Box).	Raw, high-frequency signal access.	Provides secure, raw data streams for transparent AI analysis.
Connectivity	Intermittent syncing; tolerant of latency.	Often assumes stable lab/server connectivity.	Multi-link gateway architecture for continuous 5G/satellite transmission in mobile scenarios.
Security	Standard transport encryption (TLS).	Advanced theoretical cryptography (HE, MPC).	Implements Post-Quantum Cryptography (PQC) and practical Privacy-Enhancing Technologies (PET).

Validation	Clinical trials focused on safety/efficacy.	Validated on static/clean datasets.	Validation in "wild" environments (ambulances, rural homes) with noisy real-world data.

3.2 General challenges in digital health systems

The transition of digital health technologies from controlled pilot environments to widespread clinical deployment is hindered by systemic barriers. While individual components—sensors, algorithms, and transmission protocols—have reached high maturity, their orchestration into cohesive, reliable systems remains a significant challenge. This section outlines the cross-cutting technical and operational constraints that currently limit the efficacy of Connected Health ecosystems, particularly in the pre-hospital and chronic care settings targeted by PHRESH.

3.2.1 Cross-Cutting Challenges

Data Fragmentation and Siloed Ecosystems

A primary structural weakness in current digital health architectures is data fragmentation. Patient health data is frequently sequestered within proprietary "walled gardens" established by device manufacturers. Wearable sensors, Electronic Health Records (EHR), and third-party telehealth platforms often lack semantic interoperability, resulting in disjointed data lakes.

- **Current Limitation:** Clinicians are often forced to interact with multiple disconnected dashboards to view a patient's full history (e.g., one for cardiac vitals, another for historical labs), increasing cognitive load and the risk of oversight.
- **PHRESH Approach:** PHRESH addresses this by implementing a vendor-agnostic data acquisition layer that aggregates heterogeneous streams (vital signs, acoustic markers, patient-reported outcomes) into a unified, standard-compliant (e.g., FHIR-based) structure before platform ingestion.

Interoperability and Standardization Gaps

Beyond simple data access, semantic interoperability remains a critical gap. The integration of high-frequency time-series data (such as continuous ECG or acoustic lung streams) into legacy Hospital Information Systems (HIS) is technically complex. Most HIS infrastructures are designed for transactional, episodic encounters, not the continuous ingestion of IoT telemetry.

- **Current Limitation:** The lack of standardized protocols for streaming "medical-grade" waveforms prevents real-time data from entering the clinical workflow automatically, often necessitating manual entry or static PDF uploads that strip the data of its analytical value.

Reliability in Real-World Conditions (The "Lab-to-Field" Gap)

Digital health systems that perform exceptionally well in validation studies often degrade significantly when deployed in uncontrolled, real-world environments. This degradation is primarily driven by:

- **Signal Noise and Artefacts:** In ambulatory settings, patient motion introduces significant noise into optical (PPG) and acoustic signals. Standard filtering algorithms often struggle to distinguish between motion artefacts and genuine physiological anomalies (e.g., arrhythmias), leading to high false-alarm rates.
- Lack of connectivity: In certain rural locations, there is a lack of internet connectivity for connecting remote health data to systems. This is a huge disadvantage to rural communities and requires solution that allow continuous and reliable connectivity to support health data collection and analysis.
- **Connectivity Instability:** Continuous monitoring relies on unbroken data transmission. However, in pre-hospital mobility scenarios (e.g., patient transport, rural monitoring), cellular coverage is variable.
- **Current Limitation:** Current systems lack sufficient "edge intelligence" to validate data quality locally. They frequently transmit noisy, corrupted data to the cloud, wasting bandwidth and triggering false clinical alerts.
- **PHRESH Approach:** BEWELL implements robust signal pre-processing and quality assessment algorithms directly at the gateway/edge level, ensuring that only validated, high-fidelity data is transmitted to the PHRESH AI models.

Usability in Emergency and High-Stress Scenarios

In emergency medical scenarios, such as pre-hospital care for cardiac or respiratory distress, the "cognitive cost" of technology is a decisive factor. EMTs and clinicians cannot afford to troubleshoot connection pairings or navigate complex software interfaces during a crisis.

- **Current Limitation:** Many digital health tools require active user intervention (e.g., manual pairing, frequent calibration), which renders them unusable in high-pressure emergency workflows where automation and "zero-touch" operation are prerequisites.

3.2.2 Regulatory and Operational Constraints

Regulatory Compliance (MDR and GDPR)

The regulatory landscape for Software as a Medical Device (SaMD) imposes strict requirements on traceability, risk management, and clinical validation (EU MDR 2017/745). Furthermore, the General Data Protection Regulation (GDPR) mandates strict data minimization and sovereignty.

- **Challenge:** Implementing modern AI architectures, such as cloud-based Deep Learning, within these frameworks is difficult. The "black box" nature of many AI models conflicts with the regulatory need for explainability and predictable failure modes.
- **Current Limitation:** There is a conflict between the rapid iteration cycles typical of software development and the rigorous, slow-moving validation cycles required for medical certification.

Operational Scalability and Maintenance

Deploying connected health systems at scale introduces operational hurdles related to device management. Ensuring that thousands of wearable endpoints are charged, updated, and functioning correctly without overwhelming technical support teams is a major barrier to adoption.

- **Current Limitation:** Most RPM pilots fail to scale because they underestimate the operational burden of device logistics, battery management, and remote firmware maintenance.

3.3 Key requirements: accuracy, robustness, privacy, connectivity, and trust

In the context of the PHRESH project, the transition from "wellness tracking" to "clinical-grade remote monitoring" requires the satisfaction of five non-functional requirements. These are not merely desirable features but critical acceptance criteria for the deployment of the PHRESH architecture in pre-hospital and ambulatory settings.

3.3.1 Definition of Requirements in the PHRESH Context

Accuracy (Multi-Modal Precision)

In PHRESH, accuracy extends beyond the statistical error margins of a single sensor. It is defined as the fidelity of the fused physiological profile. It is insufficient for a system to report Heart Rate (HR) with ±1 bpm precision if it fails to correlate that metric with respiratory context.

- **PHRESH Context:** Accuracy requires the synchronized capture of scalar vitals (HR, SpO_2, Temp) and high-bandwidth acoustic markers (lung sounds). The system must discern pathological signals (e.g., wheezing, rales) from background noise with a sensitivity comparable to a clinical auscultation performed in a quiet room.

Robustness (Operational Resilience)

Robustness defines the system’s ability to maintain function despite environmental hostility. Pre-hospital environments—such as moving ambulances or rural homes—are characterized by acoustic noise, electromagnetic interference, and unpredictable patient motion.

- **PHRESH Context:** A robust system must employ "edge-intelligence" to detect signal degradation (e.g., a loose sensor) and automatically compensate or alert the user, rather than silently ingesting corrupted data which could lead to false algorithmic predictions.

Privacy (Sovereignty and Confidentiality)

Privacy in PHRESH is strictly aligned with GDPR (UK GDPR) and EU MDR requirements, focusing on the protection of sensitive health data (SHD) throughout the entire lifecycle—from the sensor tip to the cloud storage.

- **PHRESH Context:** This entails end-to-end encryption (E2EE) using industry standards (TLS 1.3) and secure authentication mechanisms (OAuth 2.0). Furthermore, privacy includes "data minimization," ensuring that only clinically relevant features are transmitted, minimizing the attack surface during transit.

Connectivity (Continuous Telemetry)

Connectivity is defined as the guarantee of data delivery within clinically relevant latency bounds. In emergency use cases, "store-and-forward" (syncing data hours later) is unacceptable.

- **PHRESH Context:** The architecture requires a multi-link capability, utilizing a dedicated gateway to bond multiple backhauls (5G, 4G, Ethernet, potentially Satellite). The requirement is to maintain a continuous data stream even if individual network providers fail, ensuring the hospital receives live patient telemetry before the ambulance arrives.

Trust (Interpretability and Reliability)

Trust is the bridge between technical capability and clinical adoption. It is defined by the system's ability to explain why a specific alert was triggered. A "black box" AI model that predicts a cardiac event with 99% accuracy may still be rejected by clinicians if the rationale is opaque.

- **PHRESH Context:** Trust is achieved through Explainable AI (XAI) methods (SHAP, LIME, etc..) that visualize the specific biomarkers driving a risk score. Additionally, trust encompasses system reliability—the assurance that the system is "always on" and functioning correctly.

3.3.2 The Necessity of Simultaneous Addressal

These five requirements form an interdependent chain of survival. Addressing them in isolation results in critical system failure:

- **Accuracy without Connectivity** renders the data useless for immediate intervention (e.g., a perfect stroke detection algorithm on a device that cannot alert the hospital is clinically futile).
- **Connectivity without Privacy** creates a dangerous vulnerability, where a stable live stream could be intercepted or manipulated, compromising patient safety and institutional liability.
- **Robustness without Trust** leads to "alert fatigue." If a robust system trigger alerts that clinicians cannot understand or verify (due to lack of XAI), they will eventually ignore the system entirely.

Therefore, PHRESH architecture is designed to satisfy these constraints simultaneously. The secure gateway (Privacy/Connectivity) ensures the transport of high-fidelity signals (Accuracy/Robustness) to an interpretable decision support system (Trust).

3.3.3 Consequences of Failure in Emergency Care

The failure of any single requirement in an emergency scenario (e.g., pre-hospital transport of a CVD patient) has cascading clinical consequences:

Failed Requirement	Clinical Consequence	PHRESH Mitigation Strategy
Failure of Accuracy	**Misdiagnosis/Missed Event:** A false negative on a stroke or cardiac arrest algorithm leads to the patient being routed to a non-specialist ward instead of the ICU/Cath Lab.	Multi-modal sensor fusion (Vitals + Acoustic) to cross-validate anomalies.
Failure of Robustness	**Signal Loss:** Motion artefacts from the ambulance render the ECG/PPG signal unreadable, forcing paramedics to revert to manual checks and losing the continuous trend history.	Advanced signal pre-processing and noise-cancellation algorithms at the edge.
Failure of Connectivity	**Treatment Delay:** The receiving hospital has no visibility of the patient’s status until physical arrival ("The Black Box Patient"), delaying preparation of the surgical team.	Multi-backhaul gateway redundancy (5G/4G failover).
Failure of Privacy	**Legal/Safety Breach:** Sensitive patient data is intercepted, leading to potential identity theft or manipulation of medical records, violating GDPR and eroding public trust.	End-to-end encryption and strict identity management (IAM).

Failure of Trust	**Clinical Rejection:** Clinicians disregard valid AI alerts because they suspect a "glitch," potentially missing a life-saving intervention window.	Integration of XAI dashboards providing natural language justifications for alerts.

4 Accuracy and Robustness: Advanced Sensor Technologies

4.1 Physiological monitoring and vital signs (HR, SpO_2, BP, RR)

Accurate monitoring of core vital signs—Heart Rate (HR), Peripheral Oxygen Saturation (SpO_2), Blood Pressure (BP), and Respiratory Rate (RR)—outside the intensive care unit is the fundamental prerequisite for the PHRESH project. Current State of the Art (SotA) technologies in this domain are shifting from intermittent, clinician-performed measurements towards continuous, wearable sensor–based monitoring, enabling high-resolution trend analysis but introducing important trade-offs between data granularity, artefact susceptibility, and measurement fidelity in uncontrolled environments.

4.1.1 State of the Art in Vital Sign Monitoring

Heart Rate (HR) and Heart Rate Variability (HRV):

The industry standard for ambulatory HR monitoring remains Photoplethysmography (PPG), utilized in both consumer wearables and medical-grade devices. While multi-lead Electrocardiography (ECG) provides the reference for rhythm analysis (e.g., QRS complex detection), morphology, mature optical PPG implementations provide accurate pulse rate estimates in sedentary and low-motion conditions. Sophisticated peak-detection algorithms now enable derivation of inter-beat interval (IBI) and surrogate HRV indices from PPG, offering indirect markers of autonomic balance, although these remain more vulnerable to motion and perfusion changes than ECG-based HRV.

Oxygen Saturation (SpO_2):

Reflectance pulse oximetry embedded in wrist- and chest-worn devices has become ubiquitous extending monitoring beyond traditional transmissive (finger-clip) hospital sensors. In controlled settings, wearable SpO_2 typically demonstrates a ±2-–3% compared to bias versus arterial blood gas, which is acceptable for trend monitoring but critical near clinical decision thresholds in respiratory failure. Performance degrades with dark skin pigmentation, low perfusion, all highly relevant in pre-hospital care.

Respiratory Rate (RR):

RR is a powerful but historically under-measured predictor of clinical deterioration. SotA approaches derive RR indirectly from the modulation of the PPG waveform (respiratory-induced intensity and baseline variations) or ECG-derived respiration (EDR) using R-peak amplitude and

QRS axis changes. Direct methods—thoracic/abdominal inductance or strain bands and acoustic sensors over the trachea or chest—provide higher accuracy but face challenges in comfort, adherence, and long-term wear.

Blood Pressure (BP):

Ambulatory BP monitoring is still dominated by oscillometric cuff-based devices providing intermittent readings Cuffless techniques exploiting pulse transit time (PTT) the delay between an ECG R-wave and the peripheral PPG upstroke-or pulse arrival time (PAT) allow quasi-continuous tracking of relative BP changes but require frequent calibration to a cuff reference and show reduced absolute accuracy during rapid haemodynamic shifts and hypertensive crises. These constraints limit their current suitability as standalone diagnostic tools in high-risk patients.

4.1.2 Intermittent vs. Continuous Monitoring

The clinical value proposition of the PHRESH architecture lies in the shift from intermittent spot-checks to continuous telemetry.

- Intermittent Monitoring (Standard of Care): Conventional remote care is largely based on intermittent spot-checks, with patientsself-measuring HR, BP, and SpO_2 one to three times daily. This approach misses transient but clinically important events such as nocturnal desaturation events in COPD patients episodic tachyarrhythmias or exertional hypotension
- Continuous Monitoring (PHRESH Approach): Provides dense time-series data enabling detection of trajectories rather than isolated values. Composite patterns—such as a gradual SpO_2 decline combined with an increasing RR over several hours may be more predictive of impending decompensation than single-point threshold breaches. Continuous streams enable the calculation of "time-in-range" metrics nocturnal and physical activity, related phenotypes, and automated early warning scores adapted from in-hospital systems

4.1.3 Limitations in Motion, Noise, and Uncontrolled Environments

Despite sensor miniaturisation and improved algorithms, signal quality deteriorates substantially in real-world, mobile conditions.

- **Motion Artefacts (MA):** Physical activity is the primary adversary of optical sensing. Micro-movements of the sensor against the skin disrupt the light path in PPG sensors, creating signal noise that mimics arrhythmias. In RR monitoring, body movement often masks the subtle modulation signals used to derive breath rate.

- **Environmental Noise:** For acoustic-based monitoring (e.g., lung sounds or acoustic RR), ambient noise (traffic, television, conversation) acts as a high-magnitude interference that can drown out physiological signals.
- **Perfusion and Skin Tone:** Optical sensors perform inconsistently across different skin pigmentation levels (melanin absorbs specific wavelengths) and in states of low perfusion (e.g., vasoconstriction due to cold or shock), leading to data gaps during critical health events.

4.1.4 Gap Analysis and PHRESH Innovation

Area	Current Limitation (State of the Art)	PHRESH Technical Innovation
Data Continuity	Systems are predominantly intermittent or suffer from frequent data gaps due to motion/connectivity loss.	**Multi-Link Resilience:** PHRESH employs a buffered, multi-link wearable architecture that ensures continuous data capture even during connectivity dropouts or high patient mobility.
Parameter Fusion	Vital signs are often analyzed in isolation (e.g., SpO_2 drop without context).	**Multi-Modal Context:** The PHRESH Turkish use case integrates **acoustic lung monitoring** alongside HR/SpO_2/RR/Temp, providing the contextual data needed to distinguish genuine respiratory distress from sensor noise.
Signal Robustness	High false-alarm rates due to inability to filter motion artefacts in real-time.	**Edge Quality Assessment:** Implementation of pre-processing algorithms at the gateway level to grade signal quality and filter artefacts before data is fed into the CDSS.

4.2 Wearable devices and current limitations

- Gaps for emergency and pre-hospital use.

The current market for medical wearables is bifurcated into two distinct categories: consumer-grade activity trackers and specialized medical-grade monitors. While the former has democratized access to basic health metrics, the latter remains the standard for clinical decision-making. However, as healthcare shifts towards continuous, pre-hospital monitoring, significant functional gaps have emerged in existing hardware ecosystems.

4.2.1 Capabilities of Current Medical Wearables

State-of-the-art medical wearables are currently capable of reliable "snapshot" monitoring. Devices such as adhesive biosensors (e.g., cardiac patches) and finger-worn pulse oximeters can accurately capture single-modality data streams—primarily Heart Rate (HR) and Electrocardiogram (ECG) waveforms—in sedentary environments.

- **Current Status:** Modern devices have successfully miniaturized photoplethysmography (PPG) and accelerometry sensors, allowing for the passive tracking of rate-based vitals (HR, Respiration Rate) and gross motor activity.
- **Integration:** Leading platforms (e.g., Apple HealthKit, Google Fit) aggregate this data, but often in summarized formats (e.g., "Daily Average HR") rather than the raw, millisecond-resolution timestamps required for arrhythmia detection or detailed sleep apnea analysis.

4.2.2 Key Limitations

1. Limited Number of Vital Signs (Modal Scarcity)

Most wearable devices are optimized for a narrow set of cardiovascular metrics. There is a marked scarcity of devices that can simultaneously capture cardiovascular vitals alongside pulmonary acoustics.

- **Current Limitation:** Clinicians assessing a lung disease patient remotely often receive SpO_2 data without the contextual lung sound auscultation needed to differentiate between a stable lower saturation and an active exacerbation (e.g., wheezing).
- **PHRESH Innovation:** PHRESH addresses this by integrating a multi-modal sensing package that correlates standard vitals (HR, SpO_2, Temp) with digital auscultation (lung sounds), providing a holistic physiological picture unavailable in standard commercial wearables.

2. Usability and Comfort in Continuous Operation

For a wearable to be effective in chronic care, it must be worn continuously ("24/7"). However, power consumption often force a trade-off between sampling frequency and battery life. High-frequency sampling (required for medical-grade HRV) drains batteries rapidly, leading to frequent recharging cycles that create data gaps. Furthermore, rigid form factors often cause skin irritation or discomfort during sleep, leading to patient non-compliance.

- **Current Limitation:** High-fidelity monitoring devices are often too bulky or power-hungry for continuous multi-day use, resulting in fragmented longitudinal data.
- **PHRESH Innovation:** The PHRESH architecture optimizes power management through intelligent edge-based sampling—dynamically adjusting data transmission rates based on patient risk acuity—and utilizes ergonomic, low-profile form factors designed for extended wear in the Turkish pilot use case.

3. Certification Constraints and Data Access

The regulatory barrier (EU MDR / FDA) creates a dichotomy: consumer devices are user-friendly but lack certification for diagnosis, while certified medical devices are often "closed systems." Manufacturers of certified devices typically lock access to raw data streams to protect proprietary algorithms, providing only processed values to third-party platforms.

- **Current Limitation:** Research and CDSS development are stifled because certified devices do not expose the raw signal data necessary to train novel AI algorithms or validate signal quality.
- **PHRESH Innovation:** PHRESH utilizes a "transparent pipeline" approach, selecting hardware and developing APIs that ensure raw data accessibility while maintaining compliance with security standards, enabling the development of verifiable, explainable AI models.

4.2.3 Gaps for Emergency and Pre-Hospital Use

1. Connectivity and the "Last Mile" Problem

In pre-hospital scenarios (e.g., ambulance transport or rural remote monitoring), reliance on a single connectivity protocol (usually Bluetooth to a smartphone) is a single point of failure. If the patient’s phone battery dies or the Bluetooth pairing drops, the monitoring chain is broken.

- **Current Limitation:** Current systems lack "multi-link" redundancy and often fail to buffer data during network outages, leading to permanent data loss during critical transport phases.
- **PHRESH Innovation:** PHRESH implements a Multi-Link Wearable Architecture capable of buffering data on-device during disconnects and utilizing a dedicated gateway with cellular (5G/4G) and Ethernet backhaul options to guarantee zero data loss during the transition from home to hospital.

2. Lack of Contextual Integration

Emergency responders often arrive at a patient's location with no historical context because the wearable data is siloed in a cloud server rather than being pushed to the emergency dispatch system.

Current Limitation: The inability to push real-time wearable telemetry directly to emergency operational dashboards delays informed triage.

PHRESH Innovation: The PHRESH platform is designed to interoperate with emergency response workflows, allowing the secure, authorized "push" of live vital trends to pre-hospital care teams before they reach the patient.

4.3 Continuous blood pressure monitoring: PPG, APWA, and invasive methods

Blood pressure monitoring can be done by any of the following methods, each with their own limitations:

Method	Advantage	Limitation
Oscillometric upper arm cuff	Non-invasive, Reference	Intermitted measurements
Invasive arterial line	Reference, continuous	Invasive
PPG analysis methods	Non-invasive, cheap	Clinical accuracy not guaranteed
Volume clamp method	Non-invasive, continuous, absolute	Not mobile

PHRESH innovation: The PHRESH platform will integrate a mobile version of the volume clamp method for continuous blood pressure monitoring in the pre-hospital setting. Real-time blood pressure data is available to the hospital care teams while the patient is still in the pre-hospital setting.

4.4 Dry-electrode EEG and stroke triage applications

4.4.1. Current EEG-based stroke triage solutions & lack of real-time decision-making

Stroke is responsible for a high disability, mortality, and economic burden worldwide. Being brought to the right hospital with the right treatment, particularly intravenous in ischemic large vessel occlusion (LVO) stroke. Currently, there are no reliable methods to detect LVO in the prehospital setting—except for expensive mobile CT scanners. Stroke recognition by ambulance personnel relies primarily on clinical symptoms, yet only about 7% of these suspected cases are diagnosed as LVO. These patients are usually taken to the nearest primary stroke center, where, after CT diagnosis, these patients are transported again to a comprehensive stroke center where mechanical thrombectomy is available. Electroencephalography (EEG) is a non-invasive clinical tool that can provide an objective and quantitative measure for stroke identification. In 2023 is has been demonstrated for the first time that EEG can provide early identification of LVO in prehospital emergency setting with high diagnostic accuracy [The Dutch ELECTRA-STROKE study (Van Stigt et al. 2023)]. This study confirmed the feasibility of dry-electrode EEG recorded by ambulance personnel. To convert these research results into a medical device, Trianect was founded. Over the past years Trianect has developed StrokePointer, a novel portable EEG device, which enables the acquisition of electrical brain activity (electroencephalography; EEG) with automated artifact detection and LVO stroke decision support, providing emergency medical service (EMS) personnel with the information to enable direct transport of stroke patients to the right hospital for the right treatment.

Although Trianect's StrokePointer EEG solution is already used in research setting and is in the final preparations for market launch, connectivity with the hospital to enable real-time joined decision support with neurologists and/or emergency physicians based on EEG measurements is lacking. Communication of the EEG to the hospital could improve the communication and allow for pre-notification of the hospital that a patient highly suspected of having an LVO stroke is being transported.

Furthermore, although commercial dry-electrode EEG systems are widely available, they are yet to be established in ambulance workflows. Most competing portable EEG devices (e.g. Forest Devices, Ceribell, EDGAR/PLD) are still in the phase of early clinical trials, hospital pilot use or pivoted to other uses.

The main bottleneck in demonstrating feasibility to achieving practical utility lies in the usability and real-time analysis of EEG data. For example, the ELECTRA-STROKE study reported that 32 % of EEG signal recordings were of insufficient quality to meet the threshold for analysis (Van Stigt et al., 2023), so sufficient signal quality is the first barrier to adoption in real-time clinical decision-making. Second, approximately 2% up to> 46% of suspected stroke patients in the pre-hospital setting are ultimately diagnosed with a non-stroke condition or so called “stroke mimics”. Such ambiguity makes autonomous rerouting decisions high risk, as bypassing the nearest hospital in favor of a Comprehensive Stroke Center may prolong transport times and delay essential medical care s who do not have a true stroke (Farid & Naqvi, 2024; Kühne Escolà et al., 2023; Sørensen et al., 2025). Finally, ambulance professionals encounter only few stroke patients per month, making it difficult to train and retrain them on interpreting EEG signals. Real-time AI and communication with the hospital can overcome this bottleneck of pre-hospital EEG.

4.4.2 Single-modality limitations & multimodal integration

Opportunities for pre-hospital EEG systems in the ambulance environment lie in the combination of EEG with other modalities and other clinical information to streamline the workflow of the EEG system and other modalities:

Incorporating additional on-site physiological data (e.g., blood pressure) could even offer additional valuable insights for treatment decisions and might aid in distinguishing other stroke subtypes. While this requires further fundamental research beyond the scope of the current project, integrating data from mutliple modalities at the patient level within a single platform is within the scope of PHRESH and could facilitate new applications.

Critical in the combination of multiple modalities is the matching of various variables and modalities in real-time. Communication with the ambulance electronic patient records as well as the hospital patient records and/or combinations of the two are essential. In the Netherlands we have already identified a standard for this (“*Richtlijn gegevensuitwisseling acute zorg*”), which is build on top of globally accepted HL7 standards. However, this standard is not suitable for real-time integration of modalities with sub-second. PHRESH could offer the real-time integration of modalities in a centralized dashboard/platform that exposes the modalities to the end users at the emergency department.

Currently there is limited information exchange between ambulance and receiving hospital, this requires improved reliability of connectivity (5G, 6G), real-time future-proof data encryption ensuring privacy-enhanced data exchange.

For EEG we focus on including prehospital EEG reports that include both human- and machine-readable information for downstream processing – further development is within scope of this project. Vital signs, and the treatment administered during transport can be transferred from the ambulance to the hospital staff, providing with immediate and early access to the latest information, without the risks of information loss associated with repeated verbal handovers.

4.5 Data transmission challenges in pre-hospital settings

Pre-hospital care—specifically within the dynamic environment of an ambulance or emergency transport—represents one of the most hostile environments for reliable data transmission. While static telemedicine (e.g., home-to-hospital video calls) has matured, the requirement to transmit continuous, high-fidelity physiological telemetry from a moving vehicle to a central hospital platform introduces severe technical constraints. Current "Connected Ambulance" solutions often fail to maintain the Quality of Service (QoS) necessary for diagnostic-grade monitoring. The connected ambulance could also function as the central platform for other modalities to present real-time information.

4.5.1 Constraints in Ambulances: Bandwidth, Latency, and Reliability

1. **Bandwidth Fluctuations and Throughput**

Theoretical maximums for 4G/LTE and 5G networks often mask the reality of shared cellular infrastructure. In a moving ambulance, the available uplink bandwidth fluctuates largely due to distance from cell towers and network congestion.

- **Constraint:** Transmitting multimodal data—specifically high-bandwidth streams such as continuous acoustic lung sounds (WAV/FLAC formats) or high-frequency ECG waveforms—requires stable throughput. Standard cellular connections frequently throttle this data, forcing systems to downsample or drop packets, which degrades the diagnostic quality of the signal (e.g., aliasing in acoustic files).
- **PHRESH Innovation:** The PHRESH architecture utilizes an intelligent IoT Gateway (developed in WP3) capable of adaptive compression. When bandwidth drops, the system prioritizes critical features (e.g., alerts and low-frequency vitals) while buffering high-fidelity acoustic raw data for burst transmission once signal quality improves.

2. **Latency and Jitter in Mobile Networks**

Latency in pre-hospital settings is exacerbated by the "handover" process as the modem switches between cellular towers. These handovers often induce jitter (variance in packet arrival time) and temporary connection drops.

- **Constraint:** For real-time Clinical Decision Support Systems (CDSS) running in the cloud, high latency renders the "real-time" aspect null. A delay of 30–60 seconds in receiving patient vitals means the hospital team is looking at the patient’s past state, not their current state, potentially leading to delayed instructions during a cardiac or respiratory crisis.
- **PHRESH Innovation:** PHRESH implements Edge-Computing protocols where critical anomaly detection (e.g., sudden drop in SpO_2) is processed locally on the gateway. This ensures that the alert is generated immediately, independent of the cloud uplink latency.

3. **Reliability and Coverage Dead Zones**

Reliability is the binary state of "connected" vs. "disconnected." In rural or urban-canyon environments (common in the Turkish pilot regions), ambulances inevitably pass through coverage dead zones.

- **Constraint:** Standard IoT wearables rely on a single communication path (typically Bluetooth to a phone, then phone to cloud). If the phone loses signal, the monitoring chain breaks completely. Current systems often lack "store-and-forward" logic, resulting in permanent data gaps in the patient's record during transport.
- **PHRESH Innovation:** PHRESH employs a Multi-Link Resilience Strategy. The gateway bonds multiple connections (e.g., combining two cellular carriers or switching to satellite/ethernet where available) to maximize uptime. Crucially, the system includes local storage to buffer all data during outages, automatically re-syncing with the PHRESH SaaS platform upon reconnection to ensure zero data loss.

4.5.2 Missing Real-Time Hospital Feedback Loops

The current standard of care is predominantly unidirectional: the ambulance sends limited data (voice or static ECG) to the hospital, but the hospital lacks a digital channel to "talk back" or control the sensors.

- **Operational Gap:** Receiving physicians cannot trigger remote measurements or adjust sampling rates based on what they see. For example, if a doctor suspects a pulmonary embolism based on vital signs, they cannot remotely command the system to "record lung sounds now" or "increase ECG sampling frequency." Communication is limited to voice calls, which distract paramedics from patient care.
- **Impact:** The receiving trauma team remains reactive, waiting for the patient to arrive before initiating advanced diagnostics, rather than preparing the Cath Lab or CT scanner based on confirmed real-time telemetry.

4.5.3 Gap Analysis and PHRESH Innovation Summary

Challenge Area	Current Limitation (State of the Art)	PHRESH Technical Innovation
Transmission Stability	Single-modem connections suffer from packet loss and	**Multi-Link Gateway:** Aggregation of multiple backhaul connections

	drops during cell tower handovers in moving vehicles.	(5G/4G/Ethernet) with automatic failover to maintain a stable data pipe.
Data Fidelity	High-bandwidth signals (acoustics/video) are often compressed to the point of clinical uselessness.	**Adaptive Serialization:** Intelligent prioritization of data packets; critical alerts are sent immediately, while raw waveforms are buffered or compressed losslessly.
Bi-Directionality	Systems are "read-only" for the hospital; doctors cannot query sensors or send digital commands to the ambulance.	**Bi-Directional Command Layer:** The PHRESH API allows the CDSS or hospital staff to send configuration commands back to the gateway (e.g., "Start Acoustic Recording"), closing the feedback loop.

4.6 Voice recognition and automatic reporting in medical environments

Voice recognition and automated clinical reporting technologies represent an increasingly relevant area within the state of the art of digital healthcare, particularly in high-pressure and documentation-intensive environments such as emergency medicine and mental health services. These technologies aim to reduce administrative burden, improve data accuracy, and enhance the timeliness of clinical documentation by transforming spoken interactions into structured medical records in real time.

In emergency and pre-hospital scenarios, voice-based systems offer the potential to enable hands-free operation, allowing paramedics and healthcare professionals to focus on patient care while clinical information is captured automatically. However, achieving reliable speech transcription in chaotic and acoustically challenging environments—such as ambulances—remains a significant technical challenge. Background noise, sirens, overlapping speech, patient distress, and constant movement complicate audio acquisition and processing. The state of the art therefore includes the development of advanced acoustic capture techniques, such as microphone arrays and spatial filtering, to isolate relevant speech signals and suppress interfering sounds. These approaches are complemented by Natural Language Processing (NLP) models adapted to complex medical terminology to generate accurate and context-aware transcriptions suitable for clinical documentation.

Beyond transcription, voice-based technologies are increasingly being explored for qualitative analysis of speech content and vocal characteristics. In mental health and hospital settings, automated systems are being investigated not only to structure clinical notes, but also to detect patterns in verbal expressions that may indicate emotional distress, anxiety, or changes in mental state. The integration of speech analysis with language models enables the generation of structured summaries, draft clinical reports, and decision-support cues, thereby supporting clinicians while preserving human oversight.

Recent advances in Large **Language Models (LLMs)** have further expanded the capabilities of automated medical documentation, enabling summarization of lengthy clinical conversations and synthesis of discharge notes or outpatient reports. Nevertheless, challenges remain regarding model adaptation to domain-specific terminology, performance in cases of speech impairment (e.g., dysarthria), latency requirements for real-time use, bias related to accent or gender, and seamless interoperability with Electronic Health Records (EHR) systems. Additionally, privacy constraints necessitate secure processing architectures, including on-device or edge-based computation where feasible.

Overall, while voice recognition and automated reporting technologies demonstrate substantial promise within digital healthcare, their reliable deployment in critical medical environments requires further advances in robustness, contextual understanding, interoperability, and privacy-aware design. These challenges define the technological baseline against which PHRESH positions its innovative efforts in this domain.

Limitations in noisy, chaotic environments

Current state-of-the-art speech recognition systems demonstrate high accuracy under controlled acoustic conditions; however, their performance degrades significantly in **noisy, dynamic, and unstructured clinical environments**. Emergency settings such as ambulances present particularly challenging conditions, including sirens, medical equipment alarms, engine noise, overlapping conversations, and continuous movement of personnel and patients. These factors introduce non-stationary background noise and signal distortions that conventional Automatic Speech Recognition (ASR) systems are not fully designed to handle.

In addition, chaotic environments often involve fragmented speech, stress-induced variations in tone and articulation, and rapid exchanges of critical information. Such characteristics reduce transcription reliability and increase the risk of omission or misinterpretation of clinically relevant data. Commercial speech recognition tools, typically trained on general-domain datasets, may further struggle with complex medical terminology, domain-specific abbreviations, and multilingual or accented speech patterns common in healthcare contexts.

Another limitation concerns real-time performance constraints. High computational requirements for advanced noise suppression and speech enhancement techniques may introduce latency, which is unacceptable in time-critical medical workflows. Furthermore, robust speaker identification and tracking in multi-speaker scenarios remain technically demanding, particularly when combined with mobility and spatial variability inside ambulances.

These limitations highlight a significant gap between laboratory-level ASR performance and operational deployment in emergency healthcare environments. Addressing this gap requires integrated solutions combining advanced acoustic capture, domain-adapted language models, low-latency processing architectures, and robust validation under realistic field conditions.

Challenges with medical terminology and accents.

Despite recent advances in Automatic Speech Recognition (ASR) and Natural Language Processing (NLP), significant challenges remain in accurately processing **medical terminology and diverse speech patterns** within clinical environments. Medical communication is characterized by highly specialized vocabulary, domain-specific abbreviations, drug names, anatomical terms, and procedural references that are often absent from general-purpose language models. As a result, off-the-shelf ASR systems may produce transcription errors that compromise clinical accuracy and reduce the reliability of automated documentation.

The problem is further compounded by the use of shorthand expressions, incomplete sentences, and context-dependent phrasing typically in emergency care. Clinical dialogue frequently includes rapid exchanges, overlapping speech, and implicit references to prior medical history, all of which require contextual understanding beyond simple word recognition. Without domain-specific adaptation and structured medical lexicons, language models may misinterpret critical information or fail to capture clinically relevant nuances.

Accent and linguistic diversity introduce additional complexity. Healthcare environments often involve multilingual staff and patients with varying accents, dialects, and speech patterns. Commercial speech recognition systems have demonstrated measurable bias in performance across demographic groups, leading to disparities in transcription accuracy. These biases can negatively affect both documentation quality and equitable access to AI-supported tools.

Moreover, speech impairments—such as dysarthria or stress-induced articulation changes—present further difficulties for ASR systems trained primarily on standard speech corpora. In mental health and emergency contexts, emotional distress can significantly alter vocal characteristics, affecting recognition performance.

Addressing these challenges requires domain-adapted language models, continuous vocabulary expansion, accent-robust acoustic modelling, and targeted fine-tuning strategies. Ensuring high accuracy across diverse linguistic and demographic profiles is essential for safe and equitable deployment of voice-based clinical support systems.

To address the limitations of speech recognition in noisy, chaotic, and linguistically complex medical environments, a multi-layered technological approach is required. This includes the use of advanced acoustic capture systems such as microphone arrays with beamforming and spatial filtering techniques to isolate relevant speech and suppress background noise (e.g., sirens, alarms, engine sounds), combined with adaptive noise cancellation and speaker-tracking mechanisms. At the software level, domain-adapted Automatic Speech Recognition (ASR) models are fine-tuned using medical corpora and structured clinical vocabularies to improve recognition of specialized terminology, abbreviations, and drug names, while Large Language Models (LLMs) are employed for contextual correction, structured summarization, and automated report generation. Additional strategies include bias-aware and accent-robust acoustic modelling trained on diverse datasets, support for impaired or stress-affected speech, and the integration of edge or on-device processing architectures to meet real-time performance and privacy requirements. Finally, human-in-the-loop validation mechanisms and confidence scoring are incorporated to ensure clinical reliability, transparency, and compliance with regulatory frameworks, thereby enabling safe and effective deployment in emergency and hospital settings.

4.7.2 Challenges with using AI for speech recognition and note taking

There are various issues with using Ai based note takers within a healthcare environment and these include accuracy and reliability issues related to hallucinations where AI models may create information, such as documenting physical exams that never happened or misrepresenting statements between patients and doctors. In addition, the AI algorithm may struggle with complex, medical based conversations often leading to missed or misinterpreted diagnoses and treatment plans. Also, the use of AI could affect privacy and security as using AI involves transmitting sensitive patient data, which could risk breaches if not properly secured, especially if data is sent to external servers.

4.7 Integration with Electronic Health Records (EHR): constraints and opportunities

The integration of wearable-derived data into Electronic Health Records (EHR) and Hospital Information Systems (HIS) represents the "last mile" challenge in connected health. While the acquisition of physiological data has been democratized by IoT devices, the meaningful ingestion of this data into the clinical system of records remains a significant bottleneck. Without this integration, remote monitoring data remains an "orphan" asset—viewable on a separate dashboard but disconnected from the patient’s medical history, medication list, and care plan.

4.7.1 Interoperability Issues

1. The "Velocity Mismatch" Problem

The fundamental constraint in EHR integration is the structural mismatch between the nature of IoT data and the architecture of legacy HIS.

- **IoT Nature:** Wearable sensors generate high-velocity, high-volume time-series data (e.g., continuous SpO_2 sampled at 1Hz, resulting in 86,400 data points per day).
- **EHR Architecture:** Traditional EHRs are transactional databases designed for episodic, discrete encounters (e.g., one blood pressure reading per clinic visit).
- **Constraint:** Attempting to pipe raw wearable streams directly into an EHR inevitably crashes the system or creates unmanageable "data noise." Most EHRs lack the schema flexibility to store continuous waveforms (e.g., acoustic lung recordings) alongside standard vital signs.

2. Semantic Interoperability

Even when connectivity is established, semantic disconnects persist. A wearable might tag a value as *heart_rate_variability_rmssd*, while the receiving HIS expects a LOINC code (e.g., *80404-7*). Without an intermediate semantic translation layer, the data lands in the "Notes" or "Attachments" section of the patient record as an unsearchable PDF or blob, rendering it useless for automated risk scoring.

4.7.2 Fragmentation of Standards

The Standard Wars: HL7 V2 vs. FHIR vs. Proprietary APIs The healthcare IT landscape is fragmented across multiple generations of standards.

- **HL7 V2.x:** The dominant standard in many legacy institutions (including parts of the Turkish healthcare infrastructure). It is message-based and rigid, ill-suited for modern RESTful API interactions required by cloud-native IoT platforms.
- **HL7 FHIR (Fast Healthcare Interoperability Resources):** While FHIR represents the State of the Art for modern interoperability, its adoption is uneven. Many "FHIR-ready" EHRs only support a limited subset of resources (e.g., reading patient demographics) but block the writing of external observation data due to security policies.
- **Proprietary Device Clouds:** Commercial wearable vendors often force developers to pull data from their proprietary clouds (e.g., Fitbit Web API, Garmin Health), each with unique authentication methods and data formats, creating a massive integration burden for hospital IT teams.

4.7.3 Potential Clinical Value if Integration is Achieved

Overcoming these constraints unlocks the "Closed Loop" of care, transforming passive monitoring into active management:

- **Automated Safety Netting:** If a patient's home SpO_2 drops below a threshold, the integrated system can automatically trigger a "Task" in the nurse's EHR worklist, ensuring the alert is part of the formal clinical workflow rather than a missed email notification.
- **Contextualized AI Analysis:** Integrating wearable data (Real-World Evidence) with the EHR's static data (medications, comorbidities) allows CDSS models to filter out false positives. For example, a high heart rate alert might be suppressed if the EHR confirms the patient is on a medication known to cause tachycardia, or if the patient profile indicates a baseline arrhythmia.
- **Longitudinal Disease Trajectory:** Instead of isolated snapshots, the EHR becomes a repository of the patient's continuous health journey, allowing physicians to visualize trends (e.g., gradual respiratory decline) weeks before an acute exacerbation occurs.

4.7.4 Gap Analysis and PHRESH Innovation

Area	**Current Limitation (State of the Art)**	**PHRESH Technical Innovation**
Ingestion Strategy	Direct raw data dumps overwhelm EHR databases; mostly manual PDF uploads are used.	**Middleware Abstraction (WP3):** Bewell implements an intelligent API layer that aggregates and summarizes continuous streams into clinically significant "episodes" (e.g., "Daily Average" or "Alert Event") before pushing the HIS.
Standardization	Fragmentation between HL7 V2 and diverse proprietary API formats.	**FHIR-First Architecture:** The PHRESH platform acts as a normalizing bridge, converting diverse sensor payloads into standard **FHIR Observation** resources to ensure compatibility with modern HIS interfaces.
Workflow Integration	Data lives in a separate portal ("swivel chair interoperability"), requiring clinicians to log into multiple systems.	**Unified Workflow:** By targeting direct HIS integration (Turkish Use Case), PHRESH aims to inject alerts directly into the clinician's primary workspace, eliminating the need for parallel monitoring dashboards.

5 Accuracy and Robustness: Artificial Intelligence and Machine Learning (AI/ML)

Artificial Intelligence (AI) and Machine Learning (ML) technologies have become central enablers of digital healthcare, supporting medical diagnosis, risk assessment, decision support, and personalized care. In recent years, advances in deep learning, reinforcement learning, and multimodal data processing have demonstrated significant potential to improve clinical outcomes. Nevertheless, when deployed in real-world healthcare environments—particularly in emergency, pre-hospital, and highly dynamic contexts—current AI/ML solutions still face substantial challenges in terms of accuracy, robustness, interpretability, and real-time applicability.

Artificial Intelligence (AI) in the medical sector has experienced significant progress in recent years, driven by advances in data availability, computational power, and algorithmic techniques. Current state-of-the-art approaches demonstrate strong potential across several application domains, particularly in medical diagnosis, treatment personalization, and clinical decision support.

In the area of **diagnostic techniques**, deep learning models such as Convolutional Neural Networks (CNNs) are widely applied to medical imaging tasks, including the detection of pulmonary diseases such as lung cancer, where reported accuracy can reach up to 94% under controlled conditions. In addition, hybrid architectures, including combinations of recurrent neural networks (e.g. F-RNN-LSTM), have been proposed to analyze chest X-rays and time-series medical data with reduced computational complexity, enabling more efficient deployment in resource-constrained environments.

Reinforcement Learning (RL) has emerged as a promising paradigm for the development of personalized treatment strategies. By modelling medical decision-making as a sequential optimization problem, RL-based approaches enable the adaptation of treatment policies over time, with the objective of improving long-term patient outcomes. Within this domain, **Proximal Policy Optimization (PPO)** has gained attention due to its stability and efficiency, allowing policies to be refined through iterative interaction with clinical data rather than relying exclusively on static, rule-based systems.

At the same time, the increasing complexity of AI models has reinforced the importance of **Explainable Artificial Intelligence (XAI)** in healthcare. Techniques such as Layer-wise Relevance Propagation (LRP) are used to provide insights into the internal decision-making processes of neural networks, supporting transparency, and fostering trust among healthcare professionals. Despite these advances, the integration of explainability mechanisms into real-time clinical workflows remains limited in current solutions.

5.1 AI applications in medical diagnosis

Artificial Intelligence (AI) has become an increasingly important enabler of medical diagnostic processes, supporting clinicians in the detection, classification, and assessment of a wide range of pathologies. Current state-of-the-art diagnostic systems are predominantly based on machine learning and deep learning techniques, which are capable of identifying complex patterns in medical data that may not be readily observable through conventional analysis. These approaches are applied across multiple diagnostic domains, including medical imaging, physiological signal analysis, laboratory data interpretation, and multimodal clinical assessment.

In medical imaging, deep learning models—particularly Convolutional Neural Networks (CNNs)—are widely used for tasks such as the detection of lung cancer, cardiovascular abnormalities, neurological disorders, and other pathologies from radiological data (e.g. X-rays, CT scans, MRI). In controlled evaluation settings, such systems have demonstrated high diagnostic accuracy, often comparable to or exceeding that of human experts for specific, well-defined tasks. Beyond imaging, AI techniques are also employed for the analysis of time-series physiological signals, such as electrocardiography (ECG), electroencephalography (EEG), and respiratory signals, as well as for the interpretation of structured and unstructured data from Electronic Health Records (EHRs).

Despite these advances, the **maturity level of existing AI-based diagnostic solutions varies significantly**. While some applications have reached a relatively high level of technological readiness and have been validated in clinical studies, most solutions remain task-specific and operate under constrained conditions. Many systems are developed and evaluated using curated datasets that do not fully reflect the variability, noise, and incompleteness of data encountered in real-world clinical environments. As a result, generalization across institutions, patient populations, and acquisition of devices remains a key challenge. Furthermore, regulatory approval and clinical certification processes represent additional barriers to widespread adoption.

The adoption of AI in **real-time clinical settings**, particularly in emergencies and pre-hospital contexts, is further constrained by several practical and technical factors. These include limitations in computational resources, latency requirements for time-critical decision-making, dependency on stable data streams, and integration of challenges with existing clinical workflows and hospital information systems. In addition, the lack of transparent and explainable decision mechanisms can reduce clinician trust and hinder acceptance, especially in high-risk diagnostic scenarios where accountability and interpretability are essential.

Overall, while AI-based diagnostic technologies demonstrate strong potential and measurable benefits in specific use cases, their current deployment is often limited to decision-support roles rather than fully integrated, real-time clinical systems. These constraints highlight the need for more robust, explainable, and operationally validated AI solutions capable of functioning reliably under realistic clinical conditions—an identified gap that directly motivates the research and innovation objectives addressed within the PHRESH project.

5.2 Deep learning models (CNNs, RNN, F-RNN-LSTM)

In the context of digital healthcare, deep learning models have become a central component of state-of-the-art diagnostic systems, particularly for the analysis of medical images and physiological signals. The PHRESH SotA highlights the evolution of these models and their increasing relevance for transforming medical diagnostics through data-driven approaches.

5.2.1 Models used in medical imaging and signal analysis

Among the most widely adopted architectures, **Convolutional Neural Networks (CNNs)** are the dominant approach for medical image analysis. CNN-based models are extensively used for tasks such as disease detection and classification in radiological data, including chest X-rays and computed tomography images. In the specific case of pulmonary disease detection, CNNs have demonstrated reported accuracies of up to **94% for lung cancer detection** when trained on annotated datasets of carcinoma images. Due to their ability to automatically extract hierarchical spatial features, CNNs are considered one of the most suitable techniques for chest radiograph analysis in clinical research settings.

For the analysis of sequential and temporal patterns in medical data, **Recurrent Neural Networks (RNNs)** and their variants, particularly **Long Short-Term Memory (LSTM)** networks, are widely applied. These models can capture temporal dependencies in physiological signals and longitudinal patient data, making them relevant for signal-based diagnostics and disease progression analysis.

More recently, **hybrid and fusion-based architectures**, such as **F-RNN-LSTM (Fusion Recurrent Neural Network – LSTM)** models, have been proposed to combine the strengths of convolutional and recurrent approaches. In the state of the art, F-RNN-LSTM models are used to detect and classify pulmonary diseases from chest X-rays by integrating normalized feature representations and efficient feature vector fusion. These architectures have been shown to improve key performance metrics such as accuracy, F1-score, and specificity when compared to single-model approaches. In addition, the combination of **transfer learning techniques with CNN-based models** is widely recognized as an effective strategy to enhance predictive performance in clinical imaging tasks, particularly when labeled medical data are limited.

5.2.2 Performance versus computational cost

While deep learning models have demonstrated high diagnostic performance, their **computational cost** remains a critical factor for real-world deployment. Large CNN architectures and deep recurrent models often require substantial computational resources, which can limit their applicability in time-critical or resource-constrained environments.

In this context, hybrid models such as **F-RNN-LSTM** are particularly relevant, as they are designed to improve classification performance while maintaining **reduced computational complexity**. Compared to standard deep CNN-based approaches, such as Chest Disease Detection CNN (CDD-CNN), F-RNN-LSTM models have been reported to achieve **higher diagnostic accuracy with significantly lower computational effort**. This balance between performance and efficiency is especially important for real-time processing and deployment in constrained infrastructures, including edge-based systems and connected ambulance environments, which are directly relevant to the PHRESH use cases.

5.2.3 Data dependency and generalization issues

Despite their strong performance, deep learning models in medical diagnostics remain highly **data-dependent**, and their generalization capabilities are often limited by the availability, quality, and diversity of training datasets. Many existing AI models are developed using datasets that do not fully represent the heterogeneity of real patient populations, leading to biases and reduced robustness when applied across different demographic groups, institutions, or clinical settings.

In addition, the performance of advanced architectures such as F-RNN-LSTM is closely linked to the availability of **large, diverse, and well-annotated datasets**. For example, extending model capabilities to assess disease severity or rare clinical conditions typically requires the incorporation of additional datasets with new classes and distributions. In sensitive domains such as mental health or emotion-related analysis, data collection and sharing are further constrained by strict privacy regulations and the highly sensitive nature of the information, which complicates model training and validation.

These limitations contribute to a persistent gap between theoretical model performance and practical clinical adoption. Models that perform well in experimental settings may fail to adapt to individual patient cases or to integrate seamlessly into real clinical workflows without extensive validation and contextual adaptation.

To mitigate these challenges, the PHRESH project explores **hybrid data strategies** that combine real patient data with **context-aware synthetic data**, enabling increased diversity and representativeness of training datasets while preserving patient privacy. This approach aims to improve model generalization, reduce bias, and enhance robustness, thereby addressing key limitations identified in the current state of the art.

5.3 Clinical risk models (PLCO, Mayo, McWilliams, TREAT)

Clinical risk models are widely used in routine medical practice to support decision-making in the assessment and management of pulmonary nodules, particularly in the context of lung cancer. These models are typically based on statistical methods and predefined clinical variables, and they provide structured risk estimates that guide clinicians at different stages of the diagnostic and treatment pathway. Within the state of the art, several conventional models are commonly referenced, including PLCO, Mayo, McWilliams, and TREAT.

Each of these models serves a specific purpose within the clinical workflow. The **PLCO (Prostate, Lung, Colorectal, and Ovarian Cancer)** model is primarily designed to support **screening decisions**, identifying individuals who are likely to benefit from lung cancer screening programmes based on demographic and clinical risk factors. Its focus is on population-level risk stratification and early detection.

The **Mayo** and **McWilliams** models are mainly applied during the **diagnostic evaluation of pulmonary nodules**. They aim to estimate the probability that a detected nodule is malignant, supporting clinicians in determining appropriate follow-up strategies, additional imaging, or diagnostic procedures. These models are typically used once a nodule has already been identified and require structured clinical inputs.

The **TREAT (Thoracic Research Evaluation and Treatment)** model has a more **interventional focus**, supporting surgical decision-making by identifying patients for whom invasive diagnostic procedures, such as biopsy, are most appropriate. It is generally applied at later stages of the diagnostic pathway, where treatment or invasive confirmation is being considered.

5.3.1 Strengths and limitations

A key strength of these clinical risk models is their ability to provide **standardised, transparent, and clinically accepted decision support** across critical points of the patient's pathway, from screening to diagnosis and intervention. Their rule-based nature and reliance on established clinical variables facilitate interpretability and regulatory acceptance, which has contributed to their adoption in clinical practice.

However, these models also present important limitations. As conventional, largely static approaches, they are often applied **in isolation** and rely on a limited set of predefined variables. They typically do not incorporate real-time or dynamic clinical data streams, such as continuously monitored physiological signals, imaging updates, or contextual information that modern AI-based systems can process. In addition, their adaptability to individual patient variability and rapidly evolving clinical conditions is limited, particularly in complex or time-critical scenarios.

5.3.2 Absence of unified decision-support selection

An important limitation identified in the current state of the art is the **absence of a unified clinical decision-support framework** capable of determining which risk model should be applied in each clinical context. At present, the selection of PLCO, Mayo, McWilliams, or TREAT is largely dependent on clinician preference, local protocols, or institutional practices, rather than on an integrated, context-aware decision-support mechanism.

This fragmentation results in potential inconsistencies in clinical decision-making and limits the ability to optimally combine or sequence risk assessments across the patient pathway. Existing literature indicates that additional work is required to develop systems that can assist clinicians in selecting, adapting, or integrating risk models based on patient-specific characteristics, clinical stage, and available data.

Addressing this gap is directly relevant to the objectives of PHRESH, which aims to move beyond isolated risk estimation tools toward an **integrated, AI-assisted decision-support framework**. By contextualizing and orchestrating existing clinical risk models within a connected and data-driven ecosystem, PHRESH seeks to enhance consistency, robustness, and clinical relevance in decision-making, while preserving transparency and clinician control.

5.4 Reinforcement Learning (RL, PPO) for personalized medical assistance

Reinforcement Learning (RL) is an artificial intelligence paradigm that has gained increasing attention in healthcare research due to its suitability for modelling complex, sequential decision-making processes. In contrast to supervised learning approaches, which are typically limited to pattern recognition or classification tasks, RL focuses on learning optimal decision strategies through interaction with an environment. This makes RL particularly relevant for the development of **personalized medical assistance and treatment regimens**, where clinical decisions must be adapted over time based on patient response and evolving medical information.

Within the state of the art, RL has been explored in applications such as treatment planning, therapy optimization, and adaptive clinical workflows. A commonly adopted RL approach in healthcare research is **Proximal Policy Optimization (PPO)**, a model-free algorithm that enables stable and efficient learning of decision policies through iterative trial-and-error. PPO is designed to balance exploration and exploitation while maintaining learning stability, which is a key requirement in safety-critical domains such as healthcare.

5.4.1 Advantages over static models

A principal advantage of RL-based approaches over conventional, static clinical models is their **adaptability**. Traditional decision-support systems are often governed by fixed rules or predefined thresholds that do not evolve as new patient data becomes available. In contrast, RL and PPO-based systems are capable of continuously refining their decision policies in response to new clinical evidence, enabling more up-to-date and context-aware treatment recommendations.

RL also supports a higher degree of **personalization** compared to rule-based or population-level models. By learning from individual patient trajectories and outcomes, RL-based systems can move beyond "one-size-fits-all" approaches and support individualized care strategies tailored to patient-specific characteristics, disease progression, and response to treatment.

In addition, RL has the potential to contribute to **workflow optimization** in clinical environments. By learning optimal sequences of actions, RL-based systems can assist in identifying efficient treatment pathways or care processes, reducing unnecessary interventions, and supporting more effective allocation of clinical resources. Such capabilities are particularly relevant in complex disease management scenarios, including pulmonary conditions, where treatment decisions evolve over time.

5.4.2 Barriers to clinical deployment

Despite their potential, RL-based approaches face significant barriers to deployment in real clinical settings. One major challenge relates to **computational complexity**. Algorithms such as PPO require substantial computational resources to explore and evaluate decision policies, which can limit their applicability in real-time or resource-constrained environments.

Another critical barrier is **transparency and clinical trust**. RL models often operate as black-box systems, making it difficult for clinicians to understand the rationale behind specific recommendations. This lack of interpretability poses a significant obstacle to clinical acceptance, particularly in high-risk decision-making contexts. As a result, the integration of **Explainable AI (XAI)** mechanisms—such as Layer-wise Relevance Propagation (LRP) and related techniques—is increasingly recognized as essential to make RL-based decisions accessible and interpretable to healthcare professionals.

Finally, the deployment of RL in healthcare requires robust **human oversight and validation mechanisms**. Fully autonomous decision-making is generally not acceptable in clinical practice, and AI-assisted recommendations must remain subject to professional judgment. Effective deployment therefore depends on the implementation of supervision frameworks that ensure AI-generated insights are reviewed and validated by qualified clinicians, maintaining human responsibility and accountability in medical decision-making.

5.5 Explainable Artificial Intelligence (XAI): LRP and interpretability methods

Explainable Artificial Intelligence (XAI) has emerged as a critical requirement for the adoption of AI-based systems in healthcare, particularly in clinical decision-support contexts where accountability, transparency, and trust are essential. As AI models increase in complexity, many state-of-the-art approaches operate as so-called "black-box" systems, limiting clinicians' ability to understand, validate, and confidently rely on automated recommendations. Within the current state of the art, XAI is therefore recognized as a foundational element for enabling effective collaboration between AI technologies and medical expertise.

5.5.1 Importance of explainability for clinician trust

In clinical practice, the acceptance of AI-assisted recommendations depends strongly on the ability of healthcare professionals to interpret and assess the reasoning underlying system outputs. Explainability supports clinician trust by ensuring that AI-generated insights are **understandable, traceable, and clinically meaningful**, rather than opaque or purely statistical. This transparency enables clinicians to integrate AI recommendations into their own decision-making processes without relinquishing professional judgment or responsibility.

The absence of explainability represents a major barrier to real-world deployment of AI in healthcare, particularly in critical or time-sensitive scenarios. Without clear insight into how decisions are produced, it becomes difficult for clinicians to validate results, identify potential errors, or assess the suitability of AI outputs for individual patient cases. As a result, XAI is increasingly considered a prerequisite for the operational use of AI in professional healthcare environments.

5.5.2 Overview of LRP and related interpretability techniques

Among the interpretability methods proposed in the state of the art, **Layer-wise Relevance Propagation (LRP)** is a widely adopted technique for explaining predictions made by deep neural networks. LRP works by decomposing a model's output and propagating relevance scores backward through the network, thereby identifying which input features contributed most significantly to a specific prediction. This enables clinicians and developers to gain insight into the factors influencing diagnostic or decision-support outputs, enhancing transparency, and facilitating clinical validation.

In addition to LRP, other interpretability approaches are increasingly used, particularly in the context of large language models (LLMs). **Chain-of-Thought reasoning techniques** provide structured, step-by-step explanations of how a model arrives at a conclusion, presenting reasoning processes in a form that can be aligned with clinical logic and terminology. Furthermore, the use of **intrinsically interpretable ("white-box") models**, as well as post-hoc explanation techniques, is actively explored to balance model performance with transparency and human oversight.

5.5.3 Regulatory and ethical relevance

Beyond its technical benefits, explainability has significant **regulatory and ethical implications** in healthcare. Under the European **AI Act (AIA)**, AI systems classified as high-risk—including those used in medical diagnosis and decision support—must meet strict requirements related to transparency, traceability, auditability, and reliability. XAI mechanisms are therefore essential to demonstrate compliance with these regulatory obligations and to support systematic auditing and monitoring of AI behaviour.

From an ethical perspective, explainable AI contributes to the identification and mitigation of algorithmic bias, supporting fairness and non-discrimination in medical decision-making. By enabling visibility into model behaviour, XAI allows developers and clinicians to detect unintended biases and to take corrective actions where necessary. In addition, the integration of explainability within broader governance frameworks, such as responsible AI and LLMOps practices, supports accountability and risk mitigation, including the management of issues such as erroneous or misleading model outputs.

In summary, the current state of the art recognises XAI as one of the key enablers for trustworthy, compliant, and ethically responsible AI in healthcare. Its integration is essential for bridging the gap between advanced AI capabilities and their safe, acceptable, and effective use in real clinical environments.

5.6 Multimodal data integration and real-time decision support

Multimodal Artificial Intelligence (AI) represents one of the most advanced directions in current digital healthcare research, aiming to integrate heterogeneous data sources to support more accurate and robust clinical decision-making. In the state of the art, multimodal AI systems seek to combine structured clinical information, such as patient records and clinical scores, with unstructured data originating from medical imaging, physiological signals, voice interactions, and sensor-based monitoring devices. This integration enables a more comprehensive understanding of patient condition and context, particularly in complex and time-critical healthcare scenarios.

Recent developments demonstrate the potential of multimodal approaches in applications such as stroke diagnosis and triage, where diagnostic decision support benefits from the simultaneous analysis of imaging data, physiological parameters, and contextual information. In parallel, multimodal frameworks are being investigated for emotion-aware healthcare applications, combining voice and text analysis to derive insights into patient emotional state and support more personalised and responsive care. These approaches illustrate the growing recognition that isolated data streams are insufficient to fully capture the complexity of real clinical situations.

Despite their potential, multimodal AI systems face significant technical and architectural challenges that limit their adoption in operational healthcare environments. A primary challenge is **data fragmentation and interoperability**. Current healthcare ecosystems lack standardised mechanisms for seamless data exchange across stakeholders, resulting in fragmented patient records and inefficient coordination. In pre-hospital settings, ambulance-based solutions often function as isolated technological systems that are incompatible with hospital Electronic Health Record (EHR) platforms, further exacerbating information discontinuity.

From an architectural perspective, designing platforms capable of supporting multimodal integration in real time is highly complex. Systems must be **device-agnostic** in mobile environments such as ambulances, while simultaneously enabling hospital staff to access and visualise information through standard, lightweight interfaces, such as web-based dashboards. Ensuring low latency, high reliability, and scalability across heterogeneous devices and networks remains a key challenge.

Model-level limitations also persist. Large language models (LLMs) and other advanced AI components often struggle to accurately capture **contextual and emotional subtleties**, particularly in chaotic and unstructured clinical environments. Furthermore, multimodal systems must balance the demand for high accuracy and robustness against stringent **privacy and security requirements** associated with sensitive health data, which complicates architectural design and deployment.

5.6.1 Rationale for the PHRESH approach

The limitations identified in the current state of the art underline the need for an integrated and context-aware approach such as that proposed by PHRESH. Rather than relying on isolated technologies, PHRESH adopts a **holistic system perspective**, integrating AI-driven analytics, voice-based interaction, and sensing technologies within a cohesive framework designed for hands-free operation and operational efficiency in emergency care scenarios.

PHRESH introduces an ecosystem of **synchronised sensors and wearable devices**, enabling collaborative and on-demand acquisition of multimodal data streams that are temporally aligned and clinically meaningful. This integrated data foundation supports real-time decision support and reduces information loss during patient transfer between pre-hospital and in-hospital settings.

A key differentiator of the PHRESH approach is its focus on **validation under critical conditions**. Solutions are subjected to rigorous stress testing in acute care scenarios, where real-time data exchange is essential, and delays can have significant clinical consequences. In addition, PHRESH addresses data scarcity and bias through a **hybrid data model**, combining real patient data with context-aware synthetic data to enhance training diversity, improve algorithmic robustness, and preserve patient privacy.

Overall, PHRESH responds directly to the shortcomings of current multimodal AI solutions by providing an integrated, scalable, and clinically oriented framework for real-time decision support, advancing beyond the fragmented and experimental approaches characterising the current state of the art.

5.7 Secure and collaborative learning: Cryptographic Technologies for Security, Trust and Privacy

One of the key components of the PHRESH framework is a secure system architecture, set of technologies and policies that prevent unauthorized access (to data, AI algorithms, software code etc) and at the same time allows only those who are verified individual users or organisations. The common targets of attackers are sensitive data such as healthcare data, which can be used to make ransom demands or to target individuals. When devising an architecture that comprises various frameworks and technologies, its critical to look at:

- Best Practices
- Frameworks and policies
- Technologies

Also equally important is the goal of the system and data security. There are many threats that healthcare systems face and they include:

- Malware
- Phishing attacks
- Credential theft and abuse
- Insider threats
- Prompt Injection
- AI attacks (or attacks against AI)
- Harvest now decrypt later (HNDL), storage of data so that it can be decrypted later by quantum computers

In addition, technological trends are driving cyber threats:

- Cloud computing
- Multi-cloud
- Distributed work
- IoT device integrations
- Use of AI

Therefore, system development needs to incorporate various types of solutions to cover the complete threat surface and include AI, critical infrastructure, network, end points, application, cloud, data and identity. Within the PHRESH project, the SoTA will look at cryptographic and privacy enhancing technologies.

5.8 Federated Learning

5.8.1 General concept

The concept of Federated Learning (FL) was introduced by Google researchers in 2016. This paradigm allows participants to collaboratively train a Machine Learning (ML) model without sharing raw data. In a typical FL workflow, a coordinating server initializes and distributes a shared model to multiple clients (e.g., hospitals, home gateways, or wearable devices). Each client trains the model locally using its private data and sends back only the resulting model updates. The server then aggregates these updates, commonly using algorithms such as Federated Averaging (FedAvg), to produce an improved global model. This process is repeated over multiple rounds until convergence is reached or a certain stopping criterion is met. By keeping raw data local, FL achieves a performance comparable to traditional centralized training while mitigating many of the privacy risks associated with data centralization. It reduces significantly the risk of large-scale data breaches and supports data minimization principles, making it particularly attractive for healthcare applications where privacy, trust, and regulatory compliance are paramount.

FL is commonly categorized into three main settings based on the overlap in feature and sample spaces between clients.

- Horizontal FL, also known as sample-based FL, applies when client datasets share the same feature space but consist of different samples. For instance, in healthcare, hospitals may store patient records containing similar attributes, such as age, gender, height, and weight, but for entirely different patients.
- Vertical FL, or feature-based FL, is applied when clients have datasets with the same samples but distinctive features. For example, in health care, one client might hold the clinical records of a patient, while other might hold the insurance data.
- Federated Transfer Learning is used in scenarios where both samples and features differ. In such cases, traditional horizontal or vertical federated learning approaches are insufficient due to the lack of alignment between clients' data. In healthcare it might be useful to link clinical and home-collected data.

FL shifts computation toward the edge of the network. In home care scenarios, wearable devices, smartphones, or home IoT hubs can perform local training and inference, reducing reliance on continuous cloud connectivity and enabling real-time responses, an essential property for safety-critical applications such as emergency health alerts.

5.8.2 Applications of FL in Home Care Systems

Federated learning is being applied to a variety of home care and remote health monitoring scenarios. In the United Kingdom Federated learning platforms are gaining traction, particularly in the healthcare and public sectors, to enable collaborative AI training without moving sensitive raw data, thereby ensuring GDPR compliance. Major initiatives and platforms include the NHS Federated Data Platform, academic projects like the Federated Learning Interoperability Platform. Also, for instance, c**hronic disease** requires continuous monitoring of physiological signals and lifestyle factors. Cardiovascular disease management is a good example, where wearable heart monitors, blood pressure cuffs, and smartphone apps have been used to collaboratively train risk prediction models. Analogously, researchers have applied FL to diabetes management, aggregating models from continuous glucose monitoring devices and insulin pump data to predict glycemic events.

Elderly care and assisted living are other areas that benefit from FL to allow activity recognition, fall detection, and behavioural monitoring. A recent study by Vijay *et al.* (2026) focuses on federated learning for elderly healthcare, addressing challenges like fall detection and chronic disease monitoring in seniors. They leveraged datasets such as WESAD (wearable stress and activity data) and UCI HAR (human activity recognition) in a federated manner, using deep learning models (CNNs, RNN/LSTMs) to detect unusual patterns and health events. Beyond physical health, mental health and well-being applications increasingly rely on personal data such as mood logs, activity patterns, and physiological stress indicators. FL supports collaborative model training directly on personal devices, which is critical for user trust in highly sensitive mental health contexts.

5.8.3 Privacy, Security, and Trust in FL

Federated Learning, as originally introduced, is now commonly referred to as Centralized Federated Learning (CFL), as it relies on a central coordinating server to orchestrate the training process. This reliance in a central server constitutes the principal limitation of CFL, as the server represents both a single point of failure and a potential trust bottleneck. If the server is compromised, malicious, or misconfigured, the privacy of all participating clients may be jeopardized. Moreover, the server may become a communication bottleneck when scaling to large numbers of clients, since it is responsible for coordinating and aggregating updates from potentially thousands or even millions of nodes. From a privacy perspective, even if the server does not observe raw data, it still has access to individual model updates, which may leak sensitive information through inference attacks such as membership inference or model inversion. In healthcare applications, where model updates may encode information about diagnoses, physiological patterns, or behavioural traits, even partial leakage can have serious ethical and legal implications.

To mitigate the structural weaknesses of CFL, Decentralized Federated Learning (DFL) has been proposed. In DFL, the central orchestration server is removed, and nodes communicate directly using peer-to-peer protocols. Each node trains a local model, exchanges updates with a subset of neighbours, and performs local aggregation. By eliminating the central coordinator, DFL removes the single point of failure and distributes trust across the network. This decentralization also improves robustness against server compromise and can reduce communication bottlenecks in large-scale systems.

However, decentralization does not automatically imply stronger privacy guarantees. Although DFL avoids a malicious central server, it introduces new privacy risks by expanding the attack surface. Model updates are now shared with multiple neighbouring nodes, each of which may be honest-but-curious or actively malicious. As a result, privacy in DFL depends not only on the learning algorithm, but also on the network topology and the aggregation protocol.

Recent studies show that the structure of the communication graph plays a decisive role in determining privacy leakage. One line of work demonstrates that privacy loss in DFL is closely linked to the size of the so-called *honest component*, defined as the connected subgraph of nodes that participate honestly in the learning process. When this honest component is large, individual contributions are more effectively masked within aggregated updates, leading to improved privacy. These studies further indicate that network topologies with uniform degree distributions tend to provide better privacy protection than highly skewed graphs, as no single node is disproportionately exposed to observation. In contrast, other work argues that decentralization can increase privacy risks by enabling adversaries to exploit structural properties of the network. In sparse topologies, where nodes have few neighbours, model updates propagate slowly, causing nodes to overfit to their local data during early training rounds. This phenomenon, referred to as *local generalization*, increases vulnerability to membership inference attacks. Conversely, in dense topologies, adversaries may gain *system knowledge* by observing updates from many neighbours. This broader visibility allows attackers to isolate and reconstruct a victim's contribution more effectively, especially when colluding nodes are present. These findings highlight a fundamental trade-off: increasing connectivity improves convergence and model performance but can also strengthen adversarial inference capabilities.

Further comparative analyses between centralized and decentralized learning suggest that neither paradigm is universally superior in terms of privacy. Without secure aggregation, dense DFL topologies can approach the privacy leakage of centralized learning, as nodes effectively expose their updates to many peers. However, when secure aggregation or masking mechanisms are employed, increasing neighbourhood size can instead enhance privacy by strengthening the "hiding in the crowd" effect, where individual updates are concealed within aggregated contributions. This indicates that the privacy impact of topology is highly conditional on the presence and design of protection mechanisms.

These results show that there is no single network topology that guarantees optimal privacy in DFL. Instead, privacy emerges from a complex interaction between graph density, degree distribution, the size of the honest component, adversarial behaviour (passive or active), and the aggregation protocol.

In healthcare monitoring systems, namely, in remote patient monitoring architectures that combine wearables, home IoT/IoMT devices, and cloud analytics, CFL is the most used approach in the literature, mainly because it aligns with existing cloud-centric coordination and model lifecycle management (e.g., global model distribution, aggregation, and versioning). On the other hand, DFL is increasingly investigated for edge-dominated and intermittently connected environments, which are common in-home care and mobile contexts.

To meet the privacy and security requirements in remote patient monitoring and safety-critical alert pipelines, FL must be combined with privacy-enhancing technologies. Secure aggregation is used to ensure that the server (in CFL) or other participants (in DFL) cannot directly inspect individual client updates in the clear, thereby reducing the exposure of per-client gradients or parameters. Differential privacy (DP) is also commonly incorporated, either at the client side (local DP) or during aggregation (central DP), to bound the information leakage attributable to any single patient's data, acknowledging an accuracy–privacy trade-off. Finally, Homomorphic Encryption (HE) has been proposed and evaluated as a cryptographic mechanism that enables aggregation and certain computations on encrypted updates, so that intermediate values remain confidential even during transmission and processing. Nevertheless, the HE can introduce non-trivial computational and communication overhead, which motivates hybrid approaches, such as Hybrid Homomorphic Encryption (HHE).

5.9 Homomorphic Encryption and Hybrid Homomorphic Encryption

HE is a cryptographic technique that enables computations to be performed directly on encrypted data, producing ciphertexts that, once decrypted, yield the same result as if the operations had been applied to the plaintext. This property makes HE particularly attractive for privacy-sensitive domains such as healthcare, where sensitive data must be processed without being revealed to intermediate parties.

HE schemes are commonly classified according to the types and number of operations they support. Partially Homomorphic Encryption (PHE) schemes allow an unlimited number of a single arithmetic operation, typically either addition or multiplication. Somewhat Homomorphic Encryption (SWHE) schemes support both addition and multiplication, but only for a limited number of operations due to noise accumulation. Fully Homomorphic Encryption (FHE) schemes remove this limitation and enable arbitrary sequences of additions and multiplications, thereby supporting general-purpose computation on encrypted data.

Several FHE schemes have been proposed in the literature. Among the most widely used are TFHE (Fast Fully Homomorphic Encryption over the Torus), which operates at the bit level; BGV and BFV, which support exact arithmetic over integers and share similar algebraic foundations; and CKKS, which enables approximate arithmetic over real or complex numbers and is particularly suited for machine learning workloads. Despite their expressiveness, FHE schemes remain computationally expensive. Their security relies on hard lattice-based problems, which leads to significant overheads in terms of computation, communication, and memory.

To address these limitations, HHE has been proposed as an alternative that combines HE with symmetric encryption (SE), which offers fast encryption, efficient communication, and constant ciphertext expansion. In typical HHE designs, data or model parameters are encrypted using a symmetric cipher, while the symmetric keys themselves are protected using HE. This allows computationally intensive homomorphic operations to be applied selectively, reducing overhead on client devices while preserving strong confidentiality guarantees. However, HHE has its limitations and for commercial applications, it still faces significant challenges that prevent widespread, immediate adoption, which include high implementation complexity and lack of standardisation for incorporating into software systems, accuracy issues, higher costs and efficiecy.

Several works have explored the integration of HE into Federated Learning protocols to enhance privacy. Some approaches rely on simplified HE schemes, such as PHE, to protect aggregation operations, trading expressiveness for efficiency. Other works employ FHE as the primary cryptographic primitive, most commonly in centralized FL architectures, where the server performs most of the homomorphic computation. The main advantage of these FHE-based approaches is flexibility, as arbitrary aggregation functions can be supported. However, even in centralized settings, the computational and communication burden imposed on clients remains significant, which limits applicability in resource-constrained environments such as wearable devices and home IoT systems.

To address these challenges, recent research has focused on HHE-based FL systems that aim to reduce client-side overhead. One of the earliest HHE-based FL frameworks proposed a centralized topology combining a lattice-based HE scheme with an HE-friendly symmetric cipher. While this design demonstrated that model accuracy could be preserved and client communication costs significantly reduced, it also some limitations. In particular, server-side computation increased substantially due to the homomorphic encryption–symmetric decryption (HESD) step, and all clients shared a common HE key pair. This shared-key assumption weakens the threat model, as a single malicious client could exploit key reuse to decrypt other participants' encrypted model updates and subsequently perform inference attacks.

Subsequent work proposed alternative centralized HHE-based designs that introduce a Trusted Key Dealer (TKD) to manage encryption keys. In these systems, model parameters are encrypted under a symmetric cipher, while the corresponding symmetric keys are encrypted under HE and distributed by the TKD. Although this approach partially mitigates the shared-key issue, it introduces a strong trust assumption in the TKD and shifts computational burden to the server. Experimental evaluations of these systems consistently show that HHE can preserve model accuracy and significantly reduce client communication costs, but at the expense of increased server-side computation and more complex key management. In the context of federated learning for healthcare monitoring systems, HHE represents a promising compromise: it reduces client-side overhead and communication costs while maintaining confidentiality of sensitive updates. At the same time, existing HHE-based FL designs highlight open challenges related to key management, trust assumptions, and server-side scalability.

5.10 Cryptography and Privacy Enhancing Technologies

As dealing with patient data is a privacy-critical task, encryption has to be employed to make sure that information is not disclosed to third parties. Digital signatures can be used to make sure that communication partners can be sure that they are indeed talking to the right entity. Therefore, encryption and digital signature algorithms are key to implement a network of sensors, IoT devices, (cloud) servers and clients. For cryptography, two variants of encryption have to be distinguished, both aiming at different scenarios: symmetric and asymmetric (public key) cryptographic algorithms. While symmetric algorithms are fast and can be used to safely encrypt large chunks of information, they have to deal with the so-called key distribution problem, because all involved parties need to know the secret key. This is why asymmetric algorithms are also important because they can be used to solve the key distribution problem by creating a hybrid cipher, which employs both, symmetric and asymmetric algorithms. Additionally, asymmetric algorithms can be used to create digital signature schemes, which can be used to verify the identity of a communication partner. The security of the currently implemented public key encryption algorithms and digital signature schemes such as RSA and ECC relies on the difficulty of mathematical problems, namely integer factorization and discrete logarithm problems, which are considered to be hard problems on classical computer architectures. However, on quantum computers algorithms exists, which could be used to break those types of encryption algorithms, e.g. Shor's Algorithm. Shor's algorithm, once implemented over a large scale and practical quantum computer, therefore poses a threat to the security of our current communication systems. Post-quantum cryptography (PQC) refers to cryptographic algorithms that are believed to be secure against quantum computing attacks. Recently, NIST announced, as part of a 6-year international PQC standardisation effort, draft standards for key encapsulation mechanisms (KEM) and digital signature algorithms (DSA), based on lattices and hash functions, namely Kyber for KEM, and Dilithium, FALCON, and SPHINICS for DSA. Even though these PQC algorithms have been widely studied in academia and industry, there remains the challenge of implementing these algorithms for resource-constrained environments and their integration over existing (classical) systems.

Secure Access Service Edge (SASE) plays a pivotal role in healthcare systems by merging network security and connectivity into a unified, cloud-native service. It ensures regulatory compliance in North America and the EU by integrating advanced security measures that protect sensitive patient data and adhere to stringent privacy standards like HIPAA and GDPR. At the forefront of innovation, SASE platforms leverage advanced machine learning algorithms to deliver real-time threat protection for all connected devices. This approach surpasses traditional sandbox-based methods by making faster, more accurate threat prevention decisions, thereby enhancing the security posture of healthcare systems and ensuring the uninterrupted delivery of critical care services.

Privacy-enhancing technologies (PET) refer to information-based and cryptography-based techniques to protect the security and privacy of data while it is being used in applications such as training and evaluating machine learning models, linking records over different databases, and federating data coming from various sources. It can be distinguished between soft PET and hard PET. Soft PET aims to not disclose information to outsiders but allow trusted third parties to access the data for processing. This can be achieved using the previously described cryptographic algorithms. However, hard PET aim to not disclose information to anyone. Hard PET can be classified under "Federated Learning (FL)", "Secure Multiparty Computation (MPC)", "Differential Privacy (DP)", and "Homomorphic Encryption (HE)", and they offer different features while having their own restrictions. For example, HE is equipped with end-to-end encryption, and it is somewhat hardware independent while it is not efficient to implement for evaluating large depth circuits such as in the training of advanced ML models over encrypted data. MPC can overcome some of the efficiency challenges at the cost of a less ideal security model. FL and DP are widely used in training collaborative ML models while protecting the privacy of users' data, but their applications are rather limited. The need for adopting PET and its challenges has been recently published by several organisations and that PET has not been fully utilised in the health-sector .

Zero knowledge proof (ZKP) is a cryptographic technique that allows one party (the prover) to convince another party (the verifier) that a statement is true without revealing any information beyond the validity of the statement. For example, it can be used to prove that someone knows a password without revealing the password itself.

ZKP has many applications in various domains, such as privacy, security, integrity, and scalability of blockchain and Web3 systems. They can enable confidential transactions, verifiable computation, decentralized identity, and more. Furthermore, they can also enhance the performance and efficiency of networks by reducing the amount of data that needs to be stored and verified in the nodes of the network.

The state of the art of ZKP in 2023 is characterized by rapid innovation and adoption. Several breakthroughs have been made in the design and implementation of ZK proofs, such as:

- Zero-Knowledge-powered smart contract blockchains: A new generation of blockchains that support full on-chain smart contract functionality with ZKP has emerged, such as Aleo , Zexe , and Zcash . These blockchains enable developers to create applications that are both programmable and private, leveraging the power of ZKP to ensure data confidentiality and integrity.
- Zero-Knowledge-rollups: Scaling solutions that use ZKP to aggregate and compress transactions off-chain and then submit them to the main chain as a single proof. ZeroKnowledge-rollups can significantly increase the throughput and lower the cost of transactions on

blockchains, such as Ethereum, Polygon, and Arbitrum. ZK rollups also preserve the security and decentralization of the main chain, unlike other scaling solutions that rely on trusted intermediaries or validators.

- ZK-SNARKs: A specific type of ZKP that is succinct, non-interactive, and an argument of knowledge. ZK-SNARKs have the advantage of being very compact and fast to verify, making them suitable for applications that require high efficiency and low latency. ZK-SNARKs have been widely used in blockchains such as Zcash[3], Celo, and Mina, as well as in protocols such as Starkware, Zama, and Polybase.

- ZK-STARK: Another type of ZK proof that is scalable, transparent, and an argument of knowledge. ZK-STARKs have the advantage of being quantum-resistant and not relying on any trusted setup or cryptographic assumptions, making them more secure and robust. ZK-STARKs have been adopted in blockchains such as Ethereum and Tezos, as well as in protocols such as Gensyn, Rarimo, and Cryptonet.

These are some of the highlights of the state of the art of ZKP in 2023. ZKP is expected to continue to evolve and expand in the coming years, as more research, development, and investment are dedicated to this promising technology. However, a big issue for the adoption of ZKPs is their resource efficiency. Existing approaches require too much memory and energy depending on the application context to be deployed on microcontrollers or other edge devices. Hence, they cannot be efficiently used in many IoT applications.

PHRESH

6 Conclusions

This State of the Art (SotA) analysis has provided a comprehensive and structured overview of the current market and technological landscape relevant to digital and connected healthcare, with a specific focus on emergency and pre-hospital care scenarios addressed by the PHRESH project. By systematically reviewing existing industrial solutions, academic research, and emerging technological trends, the SotA establishes a clear and transparent baseline against which the objectives and innovations of PHRESH can be positioned.

From a market perspective, the analysis confirms strong growth and increasing investment in key domains directly related to PHRESH, including Privacy-Enhancing Technologies (PET), Artificial Intelligence (AI) in healthcare, and health sensor technologies. Despite this growth, current market offerings remain largely fragmented and insufficiently adapted to time-critical, high-risk environments such as emergency medical services and connected ambulances. Existing solutions typically address isolated aspects of care—such as sensing, analytics, or connectivity—without delivering an integrated, real-time, and clinically robust system capable of supporting end-to-end emergency workflows.

From a technological perspective, the SotA highlights persistent limitations across several critical dimensions. In advanced sensor technologies, current wearable and non-invasive monitoring solutions face challenges in achieving the accuracy, robustness, and multi-parameter coverage required under real-world conditions characterized by motion, noise, and constrained connectivity. Continuous, non-invasive blood pressure monitoring, reliable dry-electrode EEG for stroke triage, and robust voice-based medical reporting remain only partially solved, particularly in mobile pre-hospital environments.

In the domain of Artificial Intelligence and Machine Learning, the SotA confirms that state-of-the-art models—such as deep learning architectures for medical imaging, clinical risk models, reinforcement learning approaches for personalized assistance, and explainable AI techniques—demonstrate high potential and promising performance in controlled settings. However, significant barriers remain for their effective deployment in operational clinical environments. These barriers include limited generalisation across populations and devices, high computational demands, lack of real-time performance guarantees, insufficient explainability, and challenges in integrating AI-driven insights into existing clinical workflows and information systems.

The analysis further emphasizes that privacy, security, and connectivity are not auxiliary concerns but foundational enablers for trustworthy digital healthcare. Emerging threats, such as those posed by future quantum computing, combined with the limited practical adoption of advanced PETs, highlight the need for privacy-by-design approaches that are both secure and operationally feasible. At the same time, current connectivity solutions between ambulances and hospitals remain fragmented, often resulting in incomplete or unstructured data exchange and the persistent “patient black box” problem upon arrival at the emergency department.

A key overarching conclusion of this SotA is that accuracy, robustness, privacy, connectivity, and trust must be addressed simultaneously rather than in isolation. Failure in any one of these dimensions can undermine the entire clinical value chain in emergency care, leading to delayed treatment, misinterpretation of patient status, legal and ethical risks, or rejection of digital tools by clinicians. This interdependence reinforces the need for holistic, system-level solutions rather than incremental or siloed technological improvements.

In this context, SotA clearly identifies and justifies the necessity of the PHRESH approach. PHRESH is positioned to move beyond the current state of the art by integrating advanced multimodal sensing, AI-driven and explainable decision support, privacy-by-design mechanisms, and resilient, low-latency connectivity into a unified architecture tailored for emergency and pre-hospital scenarios. By addressing validated gaps identified in this SotA and by targeting rigorous validation under realistic and demanding conditions, PHRESH aims to contribute to a scalable, clinically relevant, and trustworthy framework for connected emergency healthcare.

Overall, this SotA provides a solid technical and contextual foundation for the subsequent phases of the PHRESH project. It ensures coherence between identified market and technological gaps, the defined project objectives, and the planned work package activities, thereby strengthening the technical credibility and innovation potential of PHRESH within the ITEA framework.

PHRESH

7 References

- Litjens, G., et al. (2017). *A survey on deep learning in medical image analysis*. **Medical Image Analysis**, 42, 60–88.
- Esteva, A., et al. (2019). *A guide to deep learning in healthcare*. **Nature Medicine**, 25(1), 24–29.
- Rajpurkar, P., et al. (2018). *Deep learning for chest radiograph diagnosis: A retrospective comparison of the CheXNeXt algorithm to practicing radiologists*. **PLoS Medicine**, 15(11).
- Shen, D., Wu, G., & Suk, H.-I. (2017). *Deep learning in medical image analysis*. **Annual Review of Biomedical Engineering**, 19, 221–248.
- Hochreiter, S., & Schmidhuber, J. (1997). *Long short-term memory*. **Neural Computation**, 9(8), 1735–1780.
- Karim, F., et al. (2019). *Multivariate LSTM-FCNs for time series classification*. **Neural Networks**, 116, 237–245.
- Ullah, A., et al. (2021). *Fusion-based deep learning models for pulmonary disease detection from chest X-rays*. **Computers in Biology and Medicine**, 134.
- Tammemägi, M. C., et al. (2013). *Selection criteria for lung-cancer screening*. **New England Journal of Medicine**, 368(8), 728–736.
- Swensen, S. J., et al. (1997). *Probability of malignancy in solitary pulmonary nodules*. **Archives of Internal Medicine**, 157(8), 849–855.
- McWilliams, A., et al. (2013). *Probability of cancer in pulmonary nodules detected on first screening CT*. **New England Journal of Medicine**, 369(10), 910–919.
- Gould, M. K., et al. (2013). *Evaluation of individuals with pulmonary nodules*. **Chest**, 143(5), e93S–e120S.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press.
- Gottesman, O., et al. (2019). *Guidelines for reinforcement learning in healthcare*. **Nature Medicine**, 25(1), 16–18.
- Schulman, J., et al. (2017). *Proximal Policy Optimization Algorithms*. **arXiv preprint arXiv:1707.06347**.
- Komorowski, M., et al. (2018). *The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis*. **Nature Medicine**, 24, 1716–1720.
- Samek, W., et al. (2017). *Explainable Artificial Intelligence: Understanding, visualizing and interpreting deep learning models*. **ITU Journal**, 1(4).
- Bach, S., et al. (2015). *On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation*. **PLoS ONE**, 10(7).
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). *"Why should I trust you?" Explaining the predictions of any classifier*. **KDD 2016**.
- Doshi-Velez, F., & Kim, B. (2017). *Towards a rigorous science of interpretable machine learning*. **arXiv:1702.08608**.

- Baltrušaitis, T., Ahuja, C., & Morency, L.-P. (2019). *Multimodal machine learning: A survey and taxonomy*. **IEEE TPAMI**, 41(2), 423–443.
- Huang, S.-C., et al. (2020). *Fusion of medical imaging and electronic health records using deep learning*. **Scientific Reports**, 10.
- Zadeh, A., et al. (2017). *Tensor fusion network for multimodal sentiment analysis*. **EMNLP 2017**.
- Topol, E. (2019). *High-performance medicine: the convergence of human and artificial intelligence*. **Nature Medicine**, 25(1), 44–56.
- European Commission (2021). *Proposal for a Regulation laying down harmonised rules on Artificial Intelligence (Artificial Intelligence Act)*.
- European Data Protection Board (2024). *Opinion 28/2024 on data protection aspects of AI models*.
- ISO/IEC TR 24028:2020. *Artificial intelligence — Overview of trustworthiness in AI*.
- Van Stigt, M. N., Groenendijk, E. A., Van Meenen, L. C., Van De Munckhof, A. A., Theunissen, M., Franschman, G., Smeekes, M. D., Van Grondelle, J. A., Geuzebroek, G., Siegers, A., Visser, M. C., Van Schaik, S. M., Halkes, P. H., Majoie, C. B., Roos, Y. B., Koelman, J. H., Koopman, M. S., Marquering, H. A., Potters, W. V., & Coutinho, J. M. (2023b). Prehospital detection of large vessel occlusion stroke with EEG. *Neurology*, *101*(24), e2522–e2532. https://doi.org/10.1212/wnl.0000000000207831
- Farid, H. A., & Naqvi, A. (2024). The Burden of Stroke Mimics Among Hyperacute Stroke Unit Attendees with Special Emphasis on Migraine: A 10-Year Evaluation. *Cureus*, *16*(5), e59700. https://doi.org/10.7759/cureus.59700
- Kühne Escolà, J., Bozkurt, B., Brune, B., Chae, W. H., Milles, L. S., Pommeranz, D., Brune, L., Dammann, P., Sure, U., Deuschl, C., Forsting, M., Kill, C., Kleinschnitz, C., Köhrmann, M., & Frank, B. (2023). Frequency and Characteristics of Non-Neurological and Neurological Stroke Mimics in the Emergency Department. *Journal of clinical medicine*, *12*(22), 7067. https://doi.org/10.3390/jcm12227067
- Sørensen, V. F. I., Eddelien, H. S., Butt, J. H., & Kruuse, C. (2025). Patient-reported symptoms and admission pathways in stroke mimics versus confirmed stroke or transient ischaemic attack: a cross-sectional observational study. *BMJ open*, *15*(2), e088014. https://doi.org/10.1136/bmjopen-2024-088014
- McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2023). *Communication-efficient learning of deep networks from decentralized data*. arXiv. https://arxiv.org/abs/1602.05629
- S. Garst, J. Dekker, and M. Reinders, “A comprehensive experimental comparison between federated and centralized learning,” Database, vol. 2025, baaf016, Mar. 2025, issn: 1758-0463. doi: 10.1093/database/baaf016.
- J. Wu, F. Dong, H. Leung, Z. Zhu, J. Zhou, and S. Drew, Topology-aware federated learning in edge computing: A comprehensive survey, New York, NY, USA, 2024. https://dl.acm.org/doi/full/10.1145/3659205.

- Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). *Federated machine learning: Concept and applications*. ACM Transactions on Intelligent
- Arul Antran Vijay, S., Jebson, J., Subairkhan, A., Vasanthakumar, A., Vinith Kumar, V. (2026). Elderly Healthcare Using Federated Learning Approach. In: Bansal, J.C., Borah, S., Hussain, S., Salhi, S. (eds) Computing and Machine Learning. CML 2025. Lecture Notes in Networks and Systems, vol 1614. Springer, Singapore. https://doi.org/10.1007/978-981-95-2878-3_47