

VYOMA: A REAL-TIME DIGITAL TWIN FRAMEWORK FOR ASTRONAUT HEALTH AND PERFORMANCE OPTIMIZATION

DR. T. ADILAKSHMI¹, DR. E. SHAILAJA², DR. T. JALAJA³

¹Professor, Department Of Computer Science And Engineering, Vasavi College Of Engineering, Hyderabad, India

²Associate Professor, Department Of Computer Science And Engineering, Vasavi College Of Engineering, Hyderabad, India

³Assistant Professor, Department Of Computer Science And Engineering, Vasavi College Of Engineering, Hyderabad, India

E-mail: ¹t_adilakshmi@staff.vce.ac.in, ²e.shailaja@staff.vce.ac.in, ³jalaja.t@staff.vce.ac.in

ABSTRACT

Long-duration human spaceflight missions introduce complex physiological and cognitive risks that are further exacerbated by communication delays limiting real-time ground intervention. This paper presents VYOMA, a cyber-physical Digital Twin (DT) framework developed for the ISRO Human Space Flight Centre, which enables autonomous astronaut health monitoring through a closed-loop Physical-Digital-Physical (PDP) paradigm. The system fuses multi-modal telemetry, including cardiovascular, neuromuscular, biochemical and environmental signals, into a unified digital replica that is continuously synchronised for real-time state estimation. A hybrid artificial intelligence pipeline combining deep learning, probabilistic inference and reinforcement learning enables predictive analytics under uncertainty, while Verification, Validation and Uncertainty Quantification (VVUQ) mechanisms together with explainable artificial intelligence provide confidence-aware and interpretable decision support. To address deep-space latency, VYOMA performs low-latency edge inference, allowing early detection of critical conditions such as bone demineralisation and radiation exposure up to 48 hours in advance. Validation on statistically calibrated datasets derived from NASA, ESA and PhysioNet repositories demonstrates a predictive accuracy of 0.96, a root mean squared error of 0.28 and an end-to-end latency of 47 ms, which outperforms convolutional, recurrent, gradient boosting and transformer baselines evaluated under identical conditions. The results highlight the potential of scalable Digital Twin systems for autonomous healthcare in future interplanetary missions, and the architecture generalises to terrestrial mission-critical healthcare domains such as telemedicine, military medicine and disaster-response medicine.

Keywords: *Digital Twin, Astronaut Health Monitoring, Cyber-Physical Systems, Predictive Analytics, Edge Artificial Intelligence*

1. INTRODUCTION

Human spaceflight beyond Low Earth Orbit (LEO), particularly in missions targeting lunar habitation and Mars exploration, introduces a convergence of physiological, cognitive, and environmental stressors that challenge the limits of current biomedical monitoring paradigms. Prolonged exposure to microgravity induces progressive bone demineralization and muscle atrophy, while space radiation—comprising galactic cosmic rays (GCR) and solar particle events (SPE)—poses stochastic risks to cellular integrity and long-term health [1], [2], [3]. Concurrently, neuro-ocular conditions such as

Spaceflight-Associated Neuro-ocular Syndrome (SANS), circadian rhythm disruption, and elevated CO₂ levels contribute to cognitive degradation and impaired decision-making [4], [5], [6]. Behavioral stressors including isolation, confinement, and mission duration further exacerbate psychological strain, as evidenced in analog environments such as the MARS-500 study [7], [8].

A critical limitation in deep-space missions arises from communication latency, which can exceed 20 minutes round-trip between Earth and Mars, rendering real-time ground-based intervention infeasible [9], [10]. Existing astronaut health monitoring systems rely heavily on telemetry streaming to Earth-based control centers, where

analysis and decision-making occur post hoc. This paradigm introduces delays that are incompatible with acute anomaly detection and rapid intervention, particularly in life-critical scenarios such as decompression sickness or radiation spikes [11], [12]. Furthermore, current systems lack integrated multi-modal reasoning capabilities, often treating physiological, environmental, and behavioral data streams in isolation, thereby failing to capture emergent cross-domain interactions.

Recent advances in Artificial Intelligence (AI), Cyber-Physical Systems (CPS), and Digital Twin (DT) technologies offer a transformative pathway toward autonomous, onboard health intelligence [13], [14], [15]. A Digital Twin, defined as a dynamic, high-fidelity virtual representation of a physical system continuously updated through real-time data streams, enables predictive modeling, anomaly detection, and decision support within a closed-loop Physical–Digital–Physical (PDP) framework [16], [17]. In aerospace contexts, DTs have been explored for structural health monitoring and mission planning; however, their application to astronaut biomedical systems remains nascent and fragmented [18], [19].

The integration of deep learning architectures such as Long Short-Term Memory (LSTM) networks [20], convolutional neural networks (CNNs) [21], and probabilistic inference models [22] enables the extraction of temporal and nonlinear patterns from high-dimensional biosignals. Complementary techniques such as reinforcement learning [23] and Bayesian uncertainty modeling via Monte Carlo dropout [24] further enhance decision-making under uncertainty, a critical requirement in mission-critical environments. Additionally, Explainable AI (XAI) frameworks including SHAP and Integrated Gradients provide transparency in model predictions, ensuring trustworthiness for clinical decision support [25], [26].

1.1 Research Gap and Motivation

Recent advances in Digital Twin (DT) technology have significantly improved predictive maintenance, industrial automation, and cyber-physical system (CPS) management. Contemporary DT frameworks have increasingly been extended toward healthcare applications, enabling personalized disease monitoring, clinical decision support, and patient-specific modeling. Nevertheless, the majority of these systems remain focused on terrestrial healthcare environments and are not designed to address the unique operational

constraints of long-duration human spaceflight. Recent studies published between 2023 and 2025 have emphasized the growing importance of integrating artificial intelligence, edge computing, explainable AI, and uncertainty-aware decision-making into healthcare digital twins; however, these capabilities are rarely integrated within a unified architecture for astronaut health management.

Existing aerospace Digital Twin implementations primarily concentrate on spacecraft structural integrity, propulsion systems, and mission planning, while astronaut-centric biomedical Digital Twins remain limited to isolated physiological monitoring tasks. Most existing solutions process individual physiological signals independently, lack comprehensive multi-modal sensor fusion, depend heavily on Earth-based computational infrastructure, and provide limited support for autonomous onboard clinical decision-making. Furthermore, uncertainty quantification, explainable artificial intelligence, verification and validation methodologies, and edge-native deployment remain insufficiently addressed despite their critical importance in safety-critical deep-space environments.

These limitations become increasingly significant during lunar and Mars exploration missions, where communication latency can prevent timely medical intervention. Consequently, future astronaut healthcare systems must transition from reactive monitoring frameworks toward intelligent, autonomous Digital Twin ecosystems capable of continuous health prediction, adaptive intervention, and reliable decision support under resource-constrained operational conditions.

Research Gap: Existing studies primarily focus on isolated biomedical monitoring or generic Digital Twin architectures without integrating multi-modal physiological sensing, hybrid artificial intelligence, uncertainty-aware prediction, edge-native intelligence, and closed-loop cyber-physical decision support for astronaut healthcare. Consequently, there remains a lack of a comprehensive Digital Twin framework capable of delivering autonomous, real-time, predictive healthcare under deep-space communication latency constraints. Furthermore, no existing framework specifically addresses the operational requirements of the ISRO Human Space Flight Centre through an integrated, mission-oriented Digital Twin architecture.

Building upon the identified research gap, significant scientific and engineering challenges remain in achieving real-time, scalable, trustworthy, and autonomous Digital Twin systems for astronaut healthcare. Key challenges include multi-modal physiological data fusion across heterogeneous sensing platforms, low-latency edge intelligence under computational constraints, robust Verification, Validation, and Uncertainty Quantification (VVUQ), explainable decision-making, and seamless integration of sensing, modeling, prediction, and adaptive intervention within a unified cyber-physical architecture. Addressing these challenges is essential for enabling reliable onboard healthcare during future deep-space missions.

To address these challenges, this paper presents *VYOMA*, a mission-critical Digital Twin framework developed for the ISRO Human Space Flight Centre (RES-HSFC-2025-001). *VYOMA* operationalizes a five-layer cyber-physical architecture that fuses multi-modal telemetry streams—including cardiovascular, neuromuscular, biochemical, and environmental signals—into a continuously evolving digital representation of the astronaut. The system leverages hybrid AI pipelines integrating deep learning, classical machine learning, and reinforcement learning to enable predictive analytics and adaptive intervention strategies. By deploying edge-based inference pipelines using modern streaming infrastructures such as Apache Kafka and Flink [27], [28], and optimized inference engines such as NVIDIA Triton [29], *VYOMA* achieves near real-time responsiveness with a latency budget below 50 ms.

The proposed framework is validated using statistically calibrated datasets derived from NASA, ESA, and PhysioNet repositories, including large-scale biomedical corpora such as MIMIC-IV, PTB-XL, and the NASA Twin Study [30], [31], [32]. Through comprehensive evaluation, the system demonstrates the capability to predict critical physiological anomalies, including bone density loss and cognitive degradation, up to 48 hours prior to onset, thereby enabling proactive intervention strategies.

The key contributions of this paper are as follows:

- **Novel Digital Twin Architecture:** A five-layer cyber-physical framework integrating sensing, data acquisition, digital modeling, AI-driven processing, and decision

support tailored for astronaut health monitoring.

- **Real-Time Predictive System:** A hybrid AI pipeline combining deep learning, probabilistic inference, and reinforcement learning for low-latency, multi-modal health prediction.
- **VVUQ Integration:** Incorporation of Verification, Validation, and Uncertainty Quantification mechanisms to ensure robustness and reliability in mission-critical environments.
- **Comprehensive Validation:** Demonstration of system performance using both synthetic simulations and real-world biomedical datasets, achieving early anomaly prediction and real-time responsiveness.

1.2 Scope of the Study

This work presents the design, implementation, and evaluation of *VYOMA*, a real-time Digital Twin framework for autonomous astronaut health monitoring and performance optimization. The scope of this study encompasses the development of a five-layer cyber-physical architecture integrating heterogeneous biomedical sensing, multi-modal data fusion, hybrid artificial intelligence, uncertainty-aware predictive analytics, and edge-based decision support for long-duration human spaceflight missions.

The proposed framework is validated using representative biomedical datasets obtained from NASA, ESA, and PhysioNet repositories to evaluate its capability for early physiological anomaly prediction, low-latency inference, and reliable Digital Twin synchronization. Although the framework is motivated by the operational requirements of the ISRO Human Space Flight Centre, its architectural principles are sufficiently general to support future deep-space exploration programs, autonomous space habitats, and next-generation intelligent healthcare systems operating in resource-constrained environments.

The practical implications of this research extend beyond astronaut healthcare. The proposed Digital Twin architecture provides a scalable foundation for autonomous clinical monitoring, remote patient management, telemedicine, military healthcare, disaster-response medicine, and other mission-critical cyber-physical healthcare applications where continuous monitoring, predictive intelligence, and low-latency autonomous decision support are essential.

2. RELATED WORK

The development of Digital Twin (DT) systems for astronaut health monitoring lies at the intersection of aerospace medicine, cyber-physical systems (CPS), and artificial intelligence (AI). This section critically reviews prior work across four domains: (i) digital twin frameworks in aerospace, (ii) biomedical monitoring systems, (iii) AI-driven healthcare analytics, and (iv) system-level limitations.

2.1 Digital Twin Paradigms in Aerospace and CPS

The concept of the Digital Twin was formalized in aerospace engineering as a dynamic, data-driven virtual replica of physical systems [14]. Subsequent work expanded DT into manufacturing and CPS domains, emphasizing real-time synchronization and lifecycle modeling [16], [17], [33]. Worden [15] introduced data-centric engineering perspectives, highlighting uncertainty propagation in DT systems, while Liu et al. [34] proposed CPS-integrated DT frameworks for real-time control.

Industrial AI-driven DT architectures have incorporated predictive maintenance and autonomy [18], but their extension to biomedical systems remains limited. Existing CPS theories [13] and edge computing paradigms [35] provide the foundational infrastructure for real-time DT deployment, yet lack domain-specific adaptations for astronaut physiology.

2.2 Biomedical Monitoring in Spaceflight

NASA's Human Research Program and Life Sciences Data Archive provide extensive datasets on astronaut physiology, including bone loss, muscle atrophy, and radiation exposure [9], [36], [3]. The NASA Twin Study [32] demonstrated longitudinal multi-omics changes during extended missions, while ESA's MARS-500 experiment highlighted behavioral and psychological degradation under isolation [7].

Physiological monitoring systems have traditionally relied on discrete sensors such as ECG, EEG, and EMG, supported by datasets like MIT-BIH [37], PTB-XL [31], MIMIC-IV [30], TUH EEG [38], Sleep-EDF [39], and Ninapro [40]. These systems employ signal processing techniques such as the Pan-Tompkins algorithm [41] and Beer-Lambert law-based estimation [42], often combined with Kalman filtering for temporal smoothing [43].

Despite these advances, current biomedical systems operate in siloed pipelines, lacking integrated multi-modal reasoning and predictive capabilities necessary for deep-space missions.

2.3 Artificial Intelligence in Healthcare and Time-Series Modeling

AI has significantly advanced healthcare analytics through deep learning and probabilistic modeling. Foundational architectures such as CNNs [21], [44], LSTMs [20], and transformer-based models [45], [46] have demonstrated superior performance in pattern recognition tasks. Temporal models including Temporal Fusion Transformers [47] and Prophet [48] enable multi-horizon forecasting, while Bayesian methods [22], [24], [49] provide uncertainty quantification.

Gradient boosting methods such as XGBoost [50] and LightGBM [51] are widely used for tabular biomedical data, while reinforcement learning frameworks [23], [52] enable adaptive intervention strategies. Optimization techniques such as Hyperopt [53] further enhance model tuning.

Deep learning frameworks including TensorFlow [54] and PyTorch [55] have facilitated scalable deployment, while inference systems such as NVIDIA Triton [29] enable real-time edge processing. Data streaming platforms like Kafka and Flink [27], [28], along with IoT protocols such as MQTT [56], support high-throughput telemetry ingestion. Storage and MLOps systems including TimescaleDB [57], MLflow [58], and DVC [59] enable lifecycle management of AI models.

Explainable AI techniques such as SHAP [25], Captum [60], and Integrated Gradients [26] address interpretability challenges in clinical decision-making.

2.4 Biomedical Deep Learning Architectures

Specialized architectures such as EEGNet [61] for neural signal processing, U-Net [62] for biomedical imaging, and ResNet [21] for feature extraction have been widely adopted in healthcare applications. These models leverage hierarchical feature representations to capture complex physiological patterns.

Classical statistical frameworks [63], [64] and causal inference models [65] further complement deep learning approaches by enabling interpretable and robust predictions.

2.5 Space Medicine and Environmental Modeling

Space-specific biomedical risks, including radiation exposure and microgravity effects, have been extensively studied using datasets from NASA, ESA, and NOAA [66], [67], [3], [12]. TimePix3 detectors [68] and radiation models provide real-time environmental monitoring, while

studies on CO₂ exposure [6], SANS [4], and circadian disruption [5] highlight the multifactorial nature of astronaut health degradation.

Musculoskeletal degradation [1], [2], decompression sickness [11], and cognitive decline [10] further emphasize the need for predictive and integrative monitoring systems.

Table 1: Critical Comparison Of Representative Digital Twin And AI-Based Astronaut Healthcare Studies.

Study	Primary Focus	Limitations	Gap Addressed by VYOMA
NASA Twin Study [32]	Multi-omics astronaut analysis	No Digital Twin implementation, no real-time prediction	Introduces real-time Digital Twin with predictive analytics
Healthcare Digital Twin [19]	Patient monitoring	Limited AI integration and no edge deployment	Hybrid AI with autonomous edge intelligence
Industrial Digital Twins [16], [17]	Manufacturing and predictive maintenance	Not designed for biomedical applications	Astronaut-specific physiological Digital Twin
AI-based Healthcare Systems [69]	Disease prediction	No cyber-physical Digital Twin integration	Closed-loop Digital Twin with continuous synchronization
NASA Human Research Program [9]	Space biomedical monitoring	Ground-dependent monitoring and delayed decision making	Autonomous onboard decision support under communication latency
Proposed VYOMA	Real-time astronaut healthcare	Hybrid AI, VVUQ, explainable AI, edge inference	Unified Digital Twin framework integrating sensing, prediction, uncertainty modelling and autonomous decision support

2.6 Digital Twin in Healthcare

Recent work has explored DT applications in healthcare, focusing on personalized medicine and predictive diagnostics [19]. AI-driven biomedical systems have demonstrated potential in disease prediction and treatment optimization [69], yet their deployment in extreme environments such as space remains limited.

Table 1 provides a critical comparison of the most representative studies discussed above, contrasting their primary focus and limitations with the specific gap addressed by the proposed framework.

2.7 Critical Analysis of Existing Literature

Although previous studies have significantly advanced Digital Twin technologies, healthcare analytics, and astronaut biomedical research, most existing approaches remain domain-specific and fragmented. Industrial Digital Twins primarily focus on predictive maintenance and asset management, whereas healthcare Digital Twins are largely designed for hospital-based patient monitoring. Similarly, astronaut health studies

predominantly emphasize physiological observation rather than autonomous predictive decision support.

A critical limitation observed across the literature is the absence of an integrated framework that combines heterogeneous biomedical sensing, hybrid artificial intelligence, uncertainty-aware prediction, explainable decision-making, and edge-native cyber-physical deployment within a single Digital Twin ecosystem. Existing studies generally address these components independently rather than as an end-to-end intelligent healthcare architecture suitable for long-duration deep-space missions.

Furthermore, contemporary Digital Twin research has increasingly recognized the importance of trustworthy AI, uncertainty quantification, and explainability. However, relatively few studies have simultaneously addressed these requirements while considering communication latency, computational constraints, and autonomous onboard operation required for lunar and Mars exploration missions. These observations clearly identify an important research gap that motivates the development of the proposed VYOMA framework.

2.8 Limitations of Existing Approaches

Despite significant progress, existing approaches exhibit several limitations:

- **Fragmented Data Processing:** Most systems process physiological signals independently, lacking unified multi-modal integration.
- **Lack of Real-Time Autonomy:** Current architectures depend on ground-based analysis, unsuitable for deep-space latency constraints.
- **Limited Uncertainty Modeling:** Few systems incorporate rigorous VVUQ frameworks, leading to unreliable predictions in mission-critical scenarios.
- **Scalability Challenges:** Existing DT implementations are not optimized for continuous, high-frequency telemetry streams.

The literature review demonstrates that existing Digital Twin solutions have not yet achieved a comprehensive integration of multi-modal biomedical sensing, hybrid AI-driven prediction, uncertainty-aware reasoning, explainable decision support, and edge-native autonomous deployment for astronaut healthcare. Most available systems either emphasize industrial Digital Twins, generic healthcare applications, or isolated physiological monitoring, leaving a significant gap in mission-oriented Digital Twin architectures for deep-space exploration.

Consequently, there remains a clear need for a unified cyber-physical framework capable of supporting continuous physiological state estimation, predictive healthcare analytics, autonomous onboard decision support, and reliable operation under communication latency constraints. The VYOMA framework has been specifically designed to address these research challenges.

3. SYSTEM OVERVIEW

The increasing complexity of long-duration human spaceflight missions, including India's Gaganyaan program and prospective interplanetary expeditions, necessitates a paradigm shift from reactive health monitoring to proactive, autonomous biomedical intelligence systems. Traditional telemetry-driven architectures, which rely on Earth-based analysis, are fundamentally constrained by deep-space communication latency and limited onboard decision-making capability [9], [10].

To address these challenges, we introduce *VYOMA*, a real-time Cyber-Physical Digital Twin (DT) platform designed to operate as an onboard

biomedical intelligence layer. The system embodies a closed-loop Physical-Digital-Physical (PDP) framework, where continuous sensor measurements are fused into a dynamic digital replica of the astronaut, enabling predictive analytics and adaptive intervention.

3.1 Mission Context and Design Requirements

Spaceflight environments impose multi-dimensional stressors, including microgravity-induced musculoskeletal degradation [1], [2], radiation exposure from GCR and SPE events [3], [12], and neurocognitive impairments due to isolation and environmental factors [6], [8]. These conditions require continuous monitoring across heterogeneous modalities, including cardiovascular, neuromuscular, biochemical, and environmental domains.

Furthermore, mission constraints such as bandwidth limitations, intermittent connectivity, and computational resource constraints necessitate edge-centric architectures leveraging CPS principles [13], [35]. The system must therefore support:

- Real-time multi-modal data fusion
- Low-latency inference (<50 ms)
- Predictive modeling under uncertainty
- Autonomous decision support

3.2 Digital Twin Conceptualization

The VYOMA Digital Twin extends classical DT formulations [14], [16], [17] into the biomedical domain by modeling the astronaut as a dynamic, multi-state system:

$$x_{t+1} = Ax_t + Bu_t + \epsilon_t \quad (1)$$

where x_t represents the physiological state vector (e.g., HR, SpO₂, bone density), u_t denotes external interventions (exercise, nutrition), and ϵ_t captures stochastic uncertainties.

Unlike static monitoring systems, the DT continuously evolves through real-time synchronization, leveraging probabilistic inference [22] and Bayesian uncertainty modeling [24]. This enables predictive forecasting of physiological trajectories using temporal models such as LSTM [20] and transformer-based architectures [45], [47].

3.3 Real-Time Cyber-Physical Pipeline

The VYOMA system operationalizes a fully integrated end-to-end cyber-physical pipeline that tightly couples sensing, data streaming, real-time analytics, and decision support within a closed-loop

Physical–Digital–Physical (PDP) architecture. Unlike conventional telemetry pipelines, which are primarily passive and Earth-dependent, the proposed system is designed for autonomous, low-latency, onboard intelligence under deep-space constraints [13], [35].

1. **Physical Layer (Sensing and Signal Acquisition):** The system interfaces with heterogeneous, high-frequency biomedical and environmental sensors, including ECG, EEG, EMG, PPG, radiation dosimeters, and cabin environment monitors. These sensors generate continuous multi-modal time-series data streams with sampling rates ranging from 1 Hz to 1 kHz. Signal acquisition incorporates embedded edge preprocessing such as denoising, normalization, and feature extraction (e.g., QRS detection via Pan-Tompkins [41], optical absorption modeling via Beer-Lambert law [42]). The physiological relevance of these signals is grounded in validated biomedical datasets such as MIT-BIH [37], PTB-XL [31], TUH EEG [38], and NASA radiation datasets [3].
2. **Data Ingestion and Communication Layer:** Sensor data streams are transmitted through lightweight IoT communication protocols such as MQTT [56], ensuring reliable low-bandwidth operation in constrained environments. These streams are ingested into distributed event-driven architectures using Apache Kafka [27], which provides fault-tolerant, high-throughput message queuing with partitioned log storage. This design enables horizontal scalability and decoupling between producers (sensors) and consumers (processing engines), a key requirement for CPS deployments [34].
3. **Stream Processing and Feature Engineering:** Incoming telemetry is processed in real-time using distributed stream processing engines such as Apache Flink [28], enabling windowed computations, temporal aggregation, and multi-modal fusion. Advanced signal processing techniques—including Kalman filtering for state smoothing [43], spectral decomposition (FFT-based), and statistical feature extraction—are applied to derive higher-order features (e.g., HRV metrics, EEG band power, respiration variability). These operations transform raw sensor data into structured representations suitable for downstream machine learning pipelines,

aligning with data-centric engineering principles [15].

4. **AI/ML Inference Layer:** The processed feature streams are fed into a hybrid AI stack combining deep learning, probabilistic modeling, and ensemble methods. Temporal dependencies are modeled using LSTM networks [20] and transformer-based architectures [45], [47], while spatial and morphological features are extracted using CNN-based models such as ResNet and EEGNet [21], [61]. Tabular and fused features are handled using gradient boosting techniques such as XGBoost and LightGBM [50], [51]. Bayesian inference and uncertainty estimation are incorporated via Monte Carlo dropout [24] and probabilistic frameworks [22], enabling robust predictions under uncertainty. Model training and deployment are supported by modern frameworks such as TensorFlow and PyTorch [54], [55], with hyperparameter optimization via automated search methods [53]. Real-time inference is executed using high-performance serving systems such as NVIDIA Triton [29], ensuring sub-50 ms latency.
5. **Digital Twin State Synchronization:** The Digital Twin core maintains a continuously evolving representation of the astronaut's physiological and cognitive state. This is achieved through recursive state estimation models combining Kalman filtering [43], Bayesian updates [22], and learned transition dynamics. The system integrates multi-source data streams into a unified latent state vector, enabling predictive simulation of future trajectories (e.g., bone density loss, fatigue accumulation). The incorporation of uncertainty quantification techniques ensures reliable forecasting in mission-critical scenarios, addressing limitations identified in prior DT frameworks [15], [17].
6. **Application and Decision Support Layer:** The final stage translates model outputs into actionable insights through an interactive mission-control interface. Predictions, risk scores, and anomaly detections are visualized in real time, supported by explainable AI techniques such as SHAP and Integrated Gradients [25], [26]. This layer enables both autonomous onboard decision-making and human-in-the-loop supervision. Data persistence and lifecycle management are handled using time-series databases and

MLOps frameworks such as TimescaleDB, MLflow, and DVC [57], [58], [59], ensuring reproducibility and continuous model improvement.

The integrated pipeline achieves a tightly coupled, low-latency feedback loop that enables proactive health monitoring and adaptive intervention. By combining distributed streaming architectures with advanced AI models and uncertainty-aware digital twin synchronization, VYOMA establishes a scalable and resilient foundation for autonomous healthcare systems in deep-space missions.

The pipeline achieves near real-time performance through distributed architectures and optimized inference engines, ensuring responsiveness under mission constraints.

3.4 System Architecture Diagram

The overall cyber-physical organisation of the system, including the closed-loop feedback path that realises the Physical-Digital-Physical paradigm, is presented in Figure 1.




Figure 1: Layered Cyber-Physical Architecture Of The VYOMA Digital Twin System. The Vertical Pipeline Represents Real-Time Data Flow, While The Dashed Loop Indicates The Physical-Digital-Physical (PDP) Feedback Mechanism Enabling Adaptive Interventions.

3.5 System-Level Innovations

Compared to existing DT and biomedical monitoring systems, VYOMA introduces several key innovations:

- **Closed-loop PDP architecture** enabling continuous synchronization between physical and digital states.
- **Hybrid AI stack** combining deep learning [70], [71], probabilistic inference [22], and reinforcement learning [23].
- **Edge-native deployment** leveraging streaming systems and inference engines for real-time operation.
- **Explainable decision support** using SHAP and attribution methods for clinical transparency.

These features collectively enable a scalable, autonomous, and predictive healthcare system for next-generation human spaceflight missions.

4. FIVE-LAYER DIGITAL TWIN ARCHITECTURE

The VYOMA platform adopts a hierarchical five-layer Digital Twin (DT) architecture that integrates cyber-physical systems (CPS), real-time data streaming, and AI-driven predictive intelligence into a unified framework for astronaut health monitoring. This layered abstraction aligns with modern DT paradigms [14], [16], [17] and extends them into mission-critical biomedical applications under deep-space constraints [19], [69].

4.1 Physical Layer (Sensors)

The physical layer constitutes the foundational interface between the astronaut and the digital system, comprising a heterogeneous network of biomedical and environmental sensors. These include ECG, EEG, EMG, photoplethysmography (PPG), radiation dosimeters (e.g., TimePix3 [68]), and environmental monitors capturing CO₂, temperature, and pressure.

Signal acquisition operates across multiple temporal resolutions (1 Hz–1 kHz), requiring embedded preprocessing such as filtering, normalization, and feature extraction. Classical signal processing methods—such as Pan-Tompkins QRS detection [41] and Beer-Lambert optical modeling [42]—are combined with modern biomedical datasets (MIT-BIH [37], PTB-XL [31], TUH EEG [38], MIMIC-IV [30]) to ensure clinically validated representations.

Space-specific physiological phenomena, including bone demineralization [1], muscle atrophy [2], radiation exposure [3], [12], circadian disruption [5], and cognitive degradation due to CO₂ [6], are directly observable through these sensing modalities. This layer embodies the “physical twin” within the CPS loop [13].

4.2 Data Acquisition Layer

The data acquisition layer ensures reliable, low-latency transmission of high-frequency telemetry under constrained communication environments. Lightweight IoT protocols such as MQTT [56] enable efficient edge-to-core communication, while Apache Kafka [27] provides distributed, fault-tolerant streaming with partitioned logs for scalability.

This architecture supports asynchronous data ingestion, buffering, and replay capabilities essential for mission resilience. The system aligns with CPS communication frameworks [34] and edge computing paradigms [35], ensuring that critical health signals remain accessible even under intermittent connectivity conditions.

Additionally, telemetry is structured as time-stamped event streams, enabling temporal alignment across heterogeneous modalities. This design facilitates downstream multi-modal fusion and ensures deterministic data flow for real-time analytics.

4.3 Digital Model Layer

The digital model layer represents the core of the Digital Twin, maintaining a continuously updated virtual representation of the astronaut’s physiological and cognitive state. Unlike static models, this layer integrates dynamic system identification, probabilistic inference, and simulation-based modeling.

State estimation is performed using Bayesian filtering techniques [22], Kalman filtering [43], and Monte Carlo methods [49], enabling robust handling of noisy and incomplete sensor data. The modeling framework is informed by NASA datasets such as LSDA [36], OSDR [72], and the NASA Twin Study [32], which provide longitudinal insights into spaceflight-induced physiological changes.

The Digital Twin encapsulates multi-scale representations, from cellular-level omics data to organ-level functional models. It also incorporates causal reasoning frameworks [65], enabling the system to distinguish correlation from causation in physiological changes.

This layer extends traditional DT concepts from manufacturing [16], [17] into biomedical domains, addressing challenges such as personalization, uncertainty quantification, and long-horizon prediction [15].

4.4 Processing Layer (AI/ML Intelligence)

The processing layer operationalizes advanced AI/ML pipelines for real-time inference, prediction, and optimization. It integrates multiple model classes:

- **Deep Learning:** CNN architectures (ResNet [21], EEGNet [61]) extract spatial and morphological features, while LSTM networks [20] and sequence models [73] capture temporal dependencies.

- **Transformer Models:** Attention-based architectures [45], [46] and Temporal Fusion Transformers [47] enable long-range dependency modeling and multi-horizon forecasting.
- **Classical ML:** Gradient boosting models such as XGBoost and LightGBM [50], [51] handle tabular and fused features with high efficiency.
- **Reinforcement Learning:** Adaptive optimization of exercise, sleep, and workload scheduling is achieved using RL frameworks [23], inspired by sequential decision-making paradigms [52].

The AI stack is grounded in foundational machine learning theory [64], [74] and modern deep learning advancements [70], [44], [71]. Uncertainty estimation is incorporated using Bayesian deep learning techniques such as MC dropout [24], enabling risk-aware predictions.

Model training, deployment, and lifecycle management are supported by frameworks such as TensorFlow and PyTorch [54], [55], with optimization strategies [53] and MLOps pipelines (MLflow, DVC) [58], [59]. Real-time inference is executed using high-performance serving systems such as NVIDIA Triton [29].

4.5 Application Layer (Decision Support and Visualization)

The application layer translates model outputs into actionable insights for astronauts and mission control operators. It provides real-time visualization dashboards, predictive alerts, and explainable AI outputs.

Explainability is achieved through SHAP [25], Integrated Gradients [26], and Captum [60], enabling transparent interpretation of model decisions. This is critical for mission-critical environments where trust and interpretability are paramount.

The interface integrates cognitive performance metrics, physiological trends, and environmental risk factors into a unified decision-support system. It also incorporates findings from behavioral and isolation studies such as MARS-500 [7] and NASA HRP [9], ensuring alignment with human factors research.

Furthermore, the application layer supports autonomous interventions, reducing dependence on Earth-based control—a necessity for deep-space missions where communication delays can exceed 20 minutes [10].

4.6 Architecture Diagram




Figure 2: Compact Five-Layer Digital Twin Architecture Of VYOMA Showing Hierarchical Data Flow From Sensing To Decision Support.

The architecture illustrated in Figure 2 represents a hierarchical cyber-physical pipeline that transforms raw physiological signals into actionable intelligence through a tightly integrated Digital Twin framework. Each layer performs a distinct functional role while maintaining strong inter-layer coupling to ensure real-time responsiveness and robustness.

4.6.1 Physical layer (sensors and environment)

This layer constitutes the primary interface with the astronaut's biological and environmental state. It includes wearable and embedded sensing systems such as ECG, EEG, EMG, PPG, and radiation dosimeters, along with cabin environmental sensors (CO₂, temperature, pressure). These sensors operate at heterogeneous sampling rates (ranging from sub-Hz to kHz), generating high-frequency, multi-modal time-series data. Embedded signal conditioning techniques—such as filtering, normalization, and feature extraction—are applied at the edge to reduce noise and transmission overhead. This layer effectively

captures the “ground truth” physiological state and serves as the foundation for all downstream processing.

4.6.2 Data acquisition layer (MQTT and Kafka)

The data acquisition layer is responsible for reliable, low-latency transmission and ingestion of telemetry streams. Lightweight messaging protocols such as MQTT enable efficient communication between edge devices and onboard computing units, particularly under constrained bandwidth conditions. Apache Kafka provides a distributed streaming backbone with partitioned topics, ensuring scalability, fault tolerance, and high-throughput ingestion. This layer also handles buffering, time synchronization, and serialization of heterogeneous data streams, enabling consistent temporal alignment across modalities.

4.6.3 Digital model layer (state estimation)

At the core of the system lies the Digital Twin, which maintains a continuously evolving virtual representation of the astronaut. This layer integrates incoming telemetry with physiological models using state estimation techniques such as Kalman filtering, Bayesian inference, and probabilistic simulation. It fuses multi-modal data to infer latent variables (e.g., fatigue, hydration, cognitive load) that are not directly measurable. The Digital Twin also incorporates prior knowledge from biomedical datasets and mission-specific constraints, enabling personalized modeling and long-term trajectory tracking.

4.6.4 Processing layer (AI/ML and forecasting)

The processing layer performs advanced analytics and predictive modeling on the fused data. Deep learning architectures (e.g., CNNs for signal morphology, LSTMs for temporal dynamics) extract complex patterns from high-dimensional inputs. Gradient boosting models handle structured features, while transformer-based models enable long-horizon forecasting. Reinforcement learning components optimize intervention strategies such as exercise scheduling and sleep regulation. This layer also incorporates uncertainty estimation mechanisms, ensuring that predictions are accompanied by confidence measures essential for mission-critical decision-making.

4.6.5 Application layer (decision support interface)

The application layer provides an interactive interface for astronauts and mission control personnel. It visualizes real-time health metrics, predictive risk scores, and system alerts through a mission-control-style dashboard. Explainable AI

techniques (e.g., SHAP, attribution methods) are integrated to provide transparency into model decisions. This layer supports both monitoring and intervention, enabling users to respond proactively to emerging risks. Importantly, it is designed to operate autonomously, reducing reliance on Earth-based control in scenarios with significant communication latency.

4.6.6 Inter-layer integration and feedback

Although depicted as a unidirectional pipeline, the architecture inherently supports feedback loops. Outputs from the application and processing layers can influence sensor configurations, model parameters, and intervention strategies, forming a closed-loop Physical–Digital–Physical (PDP) system. This feedback capability is essential for adaptive, personalized health management in dynamic spaceflight environments.

5. METHODOLOGY

The VYOMA Digital Twin framework is formulated as a probabilistic, multi-modal, cyber-physical inference system that integrates dynamical systems theory, Bayesian inference, and deep learning. The methodology combines classical control-theoretic state estimation with modern AI architectures to achieve real-time synchronization between the physical astronaut and its virtual counterpart. The design is grounded in principles from digital twin systems [14], [15], [16], [17], [33], [18], [34], statistical modeling [63], [64], [22], and artificial intelligence [74], [70], [71].

5.1 State Synchronization Model

$$x_{t+1} = Ax_t + Bu_t + \epsilon_t \quad (2)$$

The astronaut’s physiological state is modeled as a discrete-time stochastic dynamical system. Here, $x_t \in \mathbb{R}^n$ represents the latent biomedical state vector at time t , encompassing variables such as cardiovascular dynamics, neural activity, metabolic load, and musculoskeletal integrity. The transition matrix A encodes intrinsic physiological evolution, including microgravity-induced adaptations such as bone demineralization [1] and muscle atrophy [2]. The control matrix B models exogenous inputs u_t , including exercise regimens, nutritional intake, and environmental interventions.

The noise term $\epsilon_t \sim \mathcal{N}(0, \Sigma)$ captures stochastic variability arising from biological uncertainty and sensor noise. This formulation aligns with classical state-space models used in

control theory and Kalman filtering [43], as well as probabilistic robotics frameworks [75].

From a derivation standpoint, this equation arises from discretizing continuous physiological dynamics:

$$\frac{dx}{dt} = F(x, u) + \eta(t) \quad (3)$$

Using first-order Euler discretization:

$$x_{t+1} \approx x_t + \Delta t \cdot F(x_t, u_t) \quad (4)$$

which yields a linearized approximation:

$$A = I + \Delta t \cdot \frac{\partial F}{\partial x}, \quad B = \Delta t \cdot \frac{\partial F}{\partial u} \quad (5)$$

This formulation is consistent with cyber-physical system modeling paradigms [13], [35] and digital twin synchronization mechanisms [34].

5.2 Digital Twin Update Function

$$\hat{x}_t = f_\theta(x_{t-1}, z_t) \quad (6)$$

The Digital Twin state $\hat{x}_t$ is updated through a non-linear function f_θ parameterized by learnable weights θ . The observation vector z_t represents multi-modal sensor measurements (ECG, EEG, EMG, radiation, etc.) derived from large-scale biomedical datasets such as PhysioNet [37], [31], [30], TUH EEG [38], and Ninapro [40].

Function f_θ is implemented as a hybrid architecture combining:

- Convolutional Neural Networks (CNNs) for spatial feature extraction [44], [21]
- Recurrent Neural Networks (RNNs) for temporal modeling [73]
- Transformer-based attention mechanisms for long-range dependencies [45], [46]

The update function approximates the Bayesian filtering equation:

$$P(x_t | z_{1:t}) \propto P(z_t | x_t) \times \int P(x_t | x_{t-1}) P(x_{t-1} | z_{1:t-1}) dx_{t-1} \quad (7)$$

where the neural network serves as a learned approximation to this intractable integral, consistent with deep probabilistic modeling approaches [22].

In the context of space medicine, this allows the system to infer latent risks such as SANS [4], cognitive degradation [6], and circadian disruption [5], integrating evidence from NASA HRP datasets [9] and twin studies [32].

5.3 Uncertainty Quantification (VVUQ)

$$P(y|x) = \int P(y|x, \theta) P(\theta) d\theta \quad (8)$$

Uncertainty quantification is essential for mission-critical decision-making. The predictive distribution is obtained by marginalizing over model parameters θ , following Bayesian inference principles [22], [49]. However, exact computation is intractable for deep neural networks.

VYOMA employs Monte Carlo Dropout as a variational approximation [24], where dropout is applied at inference time:

$$\hat{y} = \frac{1}{K} \sum_{k=1}^K f_{\theta_k}(x) \quad (9)$$

Each forward pass samples a different sub-network, approximating the posterior distribution over θ . This provides both predictive mean and epistemic uncertainty.

This approach is complemented by explainable AI techniques such as SHAP [25], Integrated Gradients [26], and Captum [60], enabling interpretability of model outputs—a critical requirement in healthcare digital twins [19].

5.4 Prediction Model (LSTM)

$$h_t = \sigma(W_h x_t + U_h h_{t-1}) \quad (10)$$

The temporal evolution of physiological signals is modeled using Long Short-Term Memory (LSTM) networks [20]. The above equation represents the hidden state update, where h_t captures temporal dependencies across multiple time scales.

The full LSTM formulation includes gating mechanisms:

$$\begin{aligned} f_t &= \sigma(W_f x_t + U_f h_{t-1}) \\ i_t &= \sigma(W_i x_t + U_i h_{t-1}) \\ o_t &= \sigma(W_o x_t + U_o h_{t-1}) \\ c_t &= f_t \odot c_{t-1} \\ &\quad + i_t \odot \tanh(W_c x_t + U_c h_{t-1}) \\ h_t &= o_t \odot \tanh(c_t) \end{aligned} \quad (11)$$

These gates regulate information flow, enabling the network to capture long-term dependencies—a key requirement for modeling slow physiological drift in spaceflight [10].

Compared to classical models:

- Kalman filters [43] assume linear Gaussian dynamics
- XGBoost [50] and LightGBM [51] handle tabular features but lack temporal memory

- Transformers [45], [46] provide global attention but require high computational resources

LSTMs offer an optimal trade-off between temporal modeling capability and computational efficiency, making them suitable for real-time onboard inference.

5.5 Integrated Learning and Optimization

The system further incorporates reinforcement learning (RL) for adaptive intervention policies [23]. Using policy gradient methods, inspired by breakthroughs such as AlphaGo [52], the system learns optimal strategies for exercise scheduling, sleep regulation, and environmental control.

Hyperparameter optimization is performed using Bayesian optimization frameworks such as Hyperopt [53], while model training leverages modern deep learning platforms including TensorFlow [54] and PyTorch [55]. Deployment is optimized using NVIDIA Triton [29] for low-latency inference.

The streaming architecture (Kafka [27], Flink [28], MQTT [56]) ensures real-time data flow, while storage systems such as TimescaleDB [57] support high-frequency time-series analytics. MLOps pipelines (MLflow [58], DVC [59]) enable reproducibility and continuous retraining.

5.6 Physics-Biological Coupling

The methodology also integrates domain-specific physical models:

- Beer-Lambert law for optical sensing (SpO_2) [42]
- Pan-Tompkins algorithm for ECG signal processing [41]
- Radiation transport models using NSRL datasets [3] and ESA models [12]

These physics-based components ensure that the Digital Twin remains grounded in first-principles modeling, complementing data-driven approaches.

5.7 Holistic System Perspective

The resulting system represents a convergence of:

- Data-centric engineering [15]
- Cyber-physical systems [13]
- AI-driven healthcare [69]

By integrating heterogeneous data sources (NASA LSDA [36], GeneLab [72], NOAA space weather [66], OMNIWeb [67]) with advanced AI models, VYOMA achieves a unified, real-time, uncertainty-aware Digital Twin capable of supporting long-duration human spaceflight missions.

6. DATASETS AND DATA PIPELINE

The VYOMA Digital Twin system is fundamentally data-centric, aligning with modern data-centric engineering paradigms [15]. The robustness, generalization, and reliability of the system are directly dependent on the diversity, scale, and fidelity of the datasets used. Unlike traditional healthcare systems that rely on narrow clinical datasets, VYOMA integrates heterogeneous, multi-modal, and multi-scale datasets spanning physiology, environmental sensing, and behavioral science.

6.1 Dataset Selection Criteria

The datasets incorporated into the VYOMA framework were selected based on five primary criteria: (i) clinical relevance to astronaut physiological monitoring, (ii) diversity of sensing modalities, (iii) data quality and public availability, (iv) temporal characteristics suitable for predictive time-series modeling, and (v) compatibility with Digital Twin synchronization requirements.

Space-specific datasets, including NASA LSDA, GeneLab, the NASA Human Research Program, and the NASA Twin Study, were selected because they provide validated observations of physiological adaptations to microgravity, radiation exposure, and long-duration spaceflight. These datasets enable the modeling of astronaut-specific biomedical phenomena that cannot be adequately represented using conventional terrestrial healthcare datasets.

Clinical datasets obtained from PhysioNet, including MIT-BIH, PTB-XL, MIMIC-IV, TUH EEG, Sleep-EDF, and Ninapro, were selected because they provide high-quality annotated physiological signals covering cardiovascular, neurological, muscular, and sleep-related conditions. Their diversity enables the development and validation of robust multi-modal learning models while improving the generalizability of the proposed framework.

Environmental datasets from NASA, ESA, and NOAA were incorporated to capture space weather conditions, radiation exposure, and other

mission-specific environmental factors that directly influence astronaut health. Integrating these datasets enables the Digital Twin to model interactions between physiological responses and external environmental stressors.

Collectively, the selected datasets satisfy the requirements for constructing a comprehensive cyber-physical Digital Twin by combining validated biomedical observations, heterogeneous sensing modalities, long-duration temporal information, and mission-specific environmental data. This selection strategy ensures that the proposed framework is both scientifically grounded and representative of operational deep-space healthcare scenarios.

6.2 Multi-Source Dataset Ecosystem

The dataset architecture is designed as a hierarchical fusion of biomedical, environmental, and mission-specific data streams, organised into space biomedical repositories, clinical physiological signal archives, and environmental or radiation datasets.

6.2.1 Space biomedical datasets

NASA repositories such as LSDA and GeneLab provide large-scale physiological and omics data [36], [72]. These datasets include measurements of bone density degradation, immune response shifts, and metabolic adaptations observed in astronauts. The NASA Human Research Program (HRP) provides structured risk models for spaceflight hazards [9], including SANS [4], CO₂-induced cognitive decline [6], and long-duration performance degradation [10].

The NASA Twin Study [32] provides longitudinal multi-omics data, enabling validation of temporal drift models. Similarly, ESA's MARS-500 study [7] contributes behavioral and psychological data under prolonged isolation, addressing cognitive and stress modeling [8].

6.2.2 Clinical and physiological signal datasets

High-frequency physiological signals are sourced from PhysioNet repositories, including:

- MIT-BIH Arrhythmia Database for ECG signal morphology [37]
- PTB-XL for large-scale supervised ECG classification [31]
- MIMIC-IV Waveform Database for ICU-grade multi-sensor signals [30]
- TUH EEG Corpus for neural activity modeling [38]

- Sleep-EDF for circadian rhythm and sleep staging [39]
- Ninapro for EMG-based musculoskeletal modeling [40]

These datasets collectively enable the training of deep neural architectures such as CNNs [44], [21], EEGNet [61], and sequence models [73].

6.2.3 Environmental and radiation datasets

Radiation exposure modeling is supported by:

- NASA Space Radiation Lab datasets [3]
- ESA radiation models [12]
- NOAA space weather datasets [66]
- NASA OMNIWeb heliospheric data [67]

These datasets capture solar particle events (SPE) and galactic cosmic radiation (GCR), critical for astronaut safety.

6.3 Data Pipeline Architecture

The end-to-end data pipeline of the VYOMA framework is summarised in Figure 3. Raw multi-source telemetry is ingested through lightweight messaging protocols, processed as continuous streams, persisted in a time-series store, and finally consumed by the hybrid artificial intelligence pipeline that updates the Digital Twin state.




Figure 3: End-To-End Data Pipeline Architecture Of VYOMA Showing Ingestion, Processing, Storage, And AI-Driven Digital Twin Updates.

6.4 Data Preprocessing and Normalization

Raw sensor data is inherently noisy, heterogeneous, and non-stationary. Therefore, rigorous preprocessing is required.

6.4.1 Signal processing

Signal conditioning includes filtering, denoising, and feature extraction:

- ECG signals processed using Pan-Tompkins algorithm [41]
- Optical signals modeled using Beer-Lambert law [42]
- State smoothing using Kalman filters [43]

The Kalman filter update equations are:

$$\begin{aligned}\hat{x}_t^- &= A\hat{x}_{t-1} \\ K_t &= P_t^- H^T (H P_t^- H^T + R)^{-1} \\ \hat{x}_t &= \hat{x}_t^- + K_t (z_t - H\hat{x}_t^-)\end{aligned}\quad (12)$$

where K_t is the Kalman gain, optimally balancing prediction and observation uncertainty.

6.4.2 Normalization

To ensure numerical stability and model convergence, all inputs are normalized:

$$\tilde{x} = \frac{x - \mu}{\sigma} \quad (13)$$

This transformation standardizes each feature to zero mean and unit variance, a requirement for gradient-based learning [70].

Derivation:

Given a dataset $x_1, x_2, \dots, x_n$,

$$\begin{aligned}\mu &= \frac{1}{n} \sum_{i=1}^n x_i \\ \sigma &= \sqrt{\frac{1}{n} \sum_{i=1}^n (x_i - \mu)^2}\end{aligned}\quad (14)$$

Thus, normalized data:

$$\tilde{x}_i = \frac{x_i - \mu}{\sigma} \quad (15)$$

ensures isotropic feature scaling, improving optimization efficiency.

6.4.3 Temporal windowing

Time-series data is structured into sliding windows:

$$X_t = [x_{t-w+1}, \dots, x_t] \quad (16)$$

This enables sequence modeling using LSTM [20] and Temporal Fusion Transformers [47].

6.5 Feature Engineering and Fusion

Feature extraction integrates both domain knowledge and data-driven representations:

- HRV metrics derived from RR intervals (statistical features)
- Spectral features from EEG using Fourier transforms
- Learned embeddings via CNNs and Transformers [45], [46]

Fusion strategies include:

- Early fusion: concatenation of raw features
- Late fusion: ensemble predictions using XGBoost [50] and LightGBM [51]

6.6 Comparison with Existing Pipelines

Conventional biomedical data pipelines primarily rely on offline data acquisition, centralized processing, and retrospective analysis, making them unsuitable for mission-critical deep-space healthcare applications where autonomous decision-making is required. Most existing healthcare pipelines process individual physiological modalities independently and provide limited support for continuous synchronization within a Digital Twin environment.

In contrast, the proposed VYOMA pipeline integrates heterogeneous biomedical, environmental, and behavioral datasets through a

unified real-time streaming architecture based on Apache Kafka and Apache Flink [27], [28]. Furthermore, the incorporation of probabilistic inference, causal reasoning, reinforcement learning, and uncertainty-aware prediction enables continuous Digital Twin synchronization rather than conventional episodic data analysis.

Unlike industrial Digital Twin pipelines, which predominantly model mechanical assets and manufacturing systems [16], [17], the proposed architecture is specifically designed to handle highly dynamic, non-linear physiological processes under uncertain operational conditions encountered during long-duration human spaceflight. These characteristics distinguish VYOMA from existing biomedical and industrial Digital Twin implementations.

6.7 MLOps and Data Lifecycle

The pipeline is governed by modern MLOps frameworks:

- MLflow for experiment tracking [58]
- DVC for dataset versioning [59]
- Hyperopt for hyperparameter tuning [53]

Deep learning training is executed using TensorFlow [54] and PyTorch [55], while inference is optimized via NVIDIA Triton [29].

6.8 Holistic Perspective

The data pipeline represents a convergence of:

- Cyber-physical systems [13]
- Edge computing paradigms [35]
- AI-driven healthcare systems [69]

By integrating multi-source datasets with real-time processing and advanced AI models, VYOMA achieves a scalable, adaptive, and mission-critical Digital Twin infrastructure capable of supporting long-duration human spaceflight.

7. RESULTS

7.1 Experimental Setup (Extended)

The VYOMA system was evaluated under a hybrid cyber-physical simulation integrating biomedical datasets (PhysioNet [37], [31], [30]), space physiology archives (NASA LSDA [36], NASA Twin Study [32]), and behavioral isolation datasets (ESA MARS-500 [7]).

Data Characteristics:

- Total samples: $> 3.2 \times 10^7$ time steps
- Modalities: ECG, EEG, EMG, radiation, cognitive metrics
- Temporal resolution: 1 Hz – 1 kHz

Training Regime:

- Sliding window size: 60 timesteps
- Forecast horizons: 1h, 24h, 7d
- Normalization: z-score normalization

7.2 Metric Derivations and Interpretation

7.2.1 Mean squared error (MSE)

$$\text{MSE} = \mathbb{E}[(Y - \hat{Y})^2] \quad (17)$$

Derived from maximum likelihood estimation assuming Gaussian noise:

$$P(y|x, \theta) \sim \mathcal{N}(\hat{y}, \sigma^2) \quad (18)$$

Minimizing negative log-likelihood leads to MSE [64], [22].

7.2.2 Bias-variance decomposition

$$\text{MSE} = \text{Bias}^2 + \text{Variance} + \text{Noise} \quad (19)$$

This explains why hybrid models reduce error:

- CNN → reduces bias
- LSTM → reduces temporal variance
- Ensemble → reduces total error

7.3 Extended Performance Table

Table 2 reports the performance of the proposed framework against convolutional, gradient boosting, recurrent, and transformer baselines evaluated on identical data splits.

7.4 Comparison with Existing Studies

To demonstrate the practical significance of the proposed framework, VYOMA was compared with representative Digital Twin and AI-driven astronaut healthcare systems reported in the literature. Table 3 summarizes the comparison across predictive performance, latency, system capabilities, and deployment characteristics.

Table 2: Extended Performance Comparison Of The Proposed Framework Against Baseline Models.

Model	Acc	Prec	Rec	F1	RMSE	MSE	AUC	Latency (ms)
CNN [21]	0.91	0.89	0.88	0.885	0.42	0.176	0.93	18
XGBoost [50]	0.89	0.87	0.86	0.865	0.45	0.202	0.91	9
LightGBM [51]	0.90	0.88	0.87	0.875	0.44	0.193	0.92	8
LSTM [20]	0.92	0.91	0.90	0.905	0.39	0.152	0.94	14
TFT [47]	0.93	0.92	0.91	0.915	0.36	0.129	0.95	21
VYOMA	0.96	0.95	0.94	0.945	0.28	0.078	0.98	12

Table 3: Comparison Of Vyoma With Representative Digital Twin Healthcare Frameworks Reported In The Literature.

Study	Accuracy	Latency	Multi-modal	Edge AI	VVUQ	Explainable AI
NASA Twin Study [32]	—	Offline	Partial	No	No	No
Healthcare DT [19]	0.90	Offline	Limited	No	Partial	Partial
AI Healthcare Framework [69]	0.92	25 ms	Yes	Limited	No	Yes
Industrial DT [16]	—	Real-time	No	Yes	No	No
VYOMA	0.96	12 ms	Yes	Yes	Yes	Yes

7.5 Additional Contributions over Existing Studies

Unlike previous Digital Twin frameworks, the proposed VYOMA system integrates multi-modal physiological sensing, hybrid artificial intelligence, uncertainty-aware prediction, explainable AI, edge-native inference, and cyber-physical feedback within a single operational architecture. Existing astronaut health monitoring studies primarily focus on isolated physiological analysis or offline biomedical assessment, whereas VYOMA enables continuous real-time state synchronization and predictive decision support suitable for long-duration autonomous space missions.

The proposed framework further incorporates Verification, Validation, and Uncertainty Quantification (VVUQ), enabling confidence-aware clinical recommendations that are largely absent in existing Digital Twin implementations. Collectively, these capabilities provide a unified solution for autonomous astronaut healthcare under deep-space communication latency constraints.

7.6 Confidence Intervals

$$CI = \bar{x} \pm z \frac{\sigma}{\sqrt{n}} \quad (20)$$

For VYOMA:

- Accuracy: 0.96 ± 0.01
- RMSE: 0.28 ± 0.02

Indicating stable generalization.

7.7 ROC Curve and AUC Analysis




Figure 4: Receiver Operating Characteristic Curve Of The Proposed Framework Compared With A Random Classifier.

7.8 Confusion Matrix

The structure of the classification outcome space used throughout the evaluation is shown in Figure 5, where the reported accuracy, precision, recall and F1-score of Table 2 are derived from the four outcome counts.

		Predicted Class	
		Positive	Negative
Actual Class	Positive	True Positive (TP)	False Negative (FN)
	Negative	False Positive (FP)	True Negative (TN)

Figure 5: Confusion Matrix Structure Used For Deriving The Classification Metrics. Rows Correspond To The Actual Class And Columns To The Predicted Class.

7.9 Ablation Study

The contribution of each architectural component was isolated through the ablation study reported in Table 4.

Table 4: Ablation Study Of The Individual Architectural Components.

Configuration	Accuracy	RMSE
Full model	0.96	0.28
Without transformer	0.93	0.34
Without LSTM	0.91	0.39
Without XGBoost	0.92	0.36

Insight: Each component contributes significantly to the overall predictive performance, with the removal of the recurrent branch causing the largest degradation in both accuracy and RMSE.

7.10 Residual Error Distribution

$$\epsilon = y - \hat{y} \quad (21)$$

The residuals follow a near-Gaussian distribution centred on zero, as illustrated in Figure 6.




Figure 6: Residual Error Distribution Showing Minimal Systematic Bias.

7.11 Bone Density Case Study (Deep Analysis)

The prediction model integrates:

- EMG-derived muscle load [40]
- Calcium metabolism drift [1]
- Bayesian filtering [22]

State update:

$$x_{t+1} = Ax_t + \epsilon_t \quad (22)$$

Kalman gain:

$$K = PH^T(HPH^T + R)^{-1} \quad (23)$$

This enables early detection of bone degradation.

7.12 Latency Analysis (Queue Theory)

Using Little's Law:

$$L = \lambda W \quad (24)$$

Where:

- λ = arrival rate
- W = latency

Measured:

- $\lambda = 50\,000$ events per second
- $W = 47$ ms

System remains stable under CPS constraints [13].

7.13 Uncertainty Calibration

Using Monte Carlo Dropout [24]:

$$\hat{y} = \frac{1}{T} \sum_{t=1}^T f_{\theta_t}(x) \quad (25)$$

Variance:

$$\text{Var}(y) = \frac{1}{T} \sum_{t=1}^T (y_t - \bar{y})^2 \quad (26)$$

Ensures reliable confidence estimation.

7.14 Discussion (Extended)

The experimental results confirm:

- Hybrid AI reduces both bias and variance
- Multi-modal fusion significantly improves robustness
- Real-time inference satisfies mission-critical constraints

Compared to existing digital twin systems [14], [16], [19], VYOMA achieves:

- Higher predictive accuracy
- Lower latency
- Better uncertainty calibration

This positions the system at the forefront of AI-driven space medicine [69].

8. DISCUSSION

8.1 Why Digital Twin Paradigm Works (Theoretical Foundation)

The VYOMA system operates under the Digital Twin paradigm originally formalized by Glaessgen and Stargel [14] and further extended in cyber-physical systems literature [13], [34]. The core principle is the continuous synchronization between physical and virtual states:

$$x_t^{\text{DT}} \approx x_t^{\text{real}} \quad (27)$$

This approximation is achieved through recursive Bayesian estimation:

$$P(x_t | z_{1:t}) \propto P(z_t | x_t) P(x_t | x_{t-1}) \quad (28)$$

Where:

- x_t = latent physiological state
- z_t = observed sensor data

This formulation aligns with probabilistic robotics frameworks [75] and Bayesian inference models [22].

Key Insight: Digital Twins succeed because they transform:

- Static monitoring → Dynamic state estimation
- Reactive healthcare → Predictive + proactive systems

8.2 Why Hybrid AI Outperforms Single Models

The superiority of VYOMA arises from bias-variance minimization:

$$\text{Error} = \text{Bias}^2 + \text{Variance} + \text{Noise} \quad (29)$$

Component Contributions:

- CNN (ResNet [21]) → spatial feature extraction
- LSTM [20] → temporal dependencies
- Transformer [45] → global attention
- XGBoost [50] → non-linear tabular refinement

Each model addresses a different failure mode, as summarised in Table 5.

Table 5: Model Capability Comparison Across Temporal, Multi-Modal And Interpretability Dimensions.

Model	Temporal	Multi-modal	Interpretability
CNN [44]	No	Limited	Low
LSTM [20]	Yes	Limited	Medium
Transformer [45]	Yes	High	Low
XGBoost [50]	No	High	High
VYOMA	Yes	Very High	High

This confirms ensemble learning theory [64].

8.3 Multi-Modal Fusion Advantage

Human physiology is inherently multi-modal:

$$x_t = f(\text{ECG, EEG, EMG, Radiation, Cognitive}) \quad (30)$$

Single-modality models fail because:

- ECG alone cannot detect cognitive decline
- EEG alone cannot predict bone loss

VYOMA solves this using feature fusion:

$$F = \bigoplus_{i=1}^n f_i(x_i) \quad (31)$$

Where $\bigoplus$ denotes learned fusion.

This aligns with multimodal AI principles [70], [74].

8.4 Uncertainty and Reliability (VVUQ)

Using Monte Carlo Dropout [24]:

$$\text{Var}(y) = \mathbb{E}[y^2] - (\mathbb{E}[y])^2 \quad (32)$$

Low variance in predictions indicates:

- High confidence
- Reliable decision-making

This is critical for mission-critical systems.

8.5 Real-Time Performance (CPS Constraints)

The system satisfies cyber-physical constraints [13]:

$$L_{\text{total}} = L_{\text{DSP}} + L_{\text{network}} + L_{\text{inference}} \quad (33)$$

Measured:

$$L_{\text{total}} = 47 \text{ ms}$$

This meets real-time safety thresholds.

8.6 Comparison with Existing Digital Twin Systems

A qualitative comparison against representative industrial and healthcare Digital Twin systems is given in Table 6.

Table 6: Comparison With Existing Digital Twin Systems.

System	Real-time	AI-driven	Multi-modal
Industrial DT [16]	Yes	Limited	No
Healthcare DT [19]	Partial	Moderate	Limited
VYOMA	Yes	Advanced	Full

8.7 Limitations

Despite strong performance, limitations exist:

- Data scarcity in extreme space conditions
- Domain shift (Earth → Deep space)
- Computational overhead of hybrid models

These challenges align with observations in AI healthcare systems [69].

8.8 Implications for Space Medicine

The system enables:

- Early detection of SANS [4]
- Bone loss prediction [1]
- Cognitive decline monitoring [6]

This represents a shift toward autonomous medical systems.

8.9 Strengths of the Proposed Framework

The proposed VYOMA framework demonstrates several strengths that directly support the research objectives defined in this work.

- **Unified Multi-modal Digital Twin:** The framework integrates cardiovascular, neurological, musculoskeletal, biochemical, and environmental data into a continuously synchronized Digital Twin, enabling holistic astronaut health assessment.
- **Hybrid Artificial Intelligence:** The combination of CNNs, LSTMs, Transformers, gradient boosting models, and Bayesian inference allows accurate prediction across heterogeneous physiological modalities.

- **Real-Time Edge Intelligence:** The integration of Apache Kafka, Apache Flink, and NVIDIA Triton enables low-latency inference suitable for deep-space environments with limited communication bandwidth.
- **Trustworthy Decision Support:** The incorporation of Verification, Validation, and Uncertainty Quantification (VVUQ), together with Explainable AI techniques, improves the reliability and interpretability of clinical recommendations.
- **Scalable Cyber-Physical Architecture:** The modular five-layer architecture facilitates future integration with additional sensing modalities, onboard computing platforms, and autonomous healthcare systems.

Collectively, these strengths demonstrate that the proposed framework successfully satisfies the primary research objectives of enabling autonomous, reliable, and predictive astronaut healthcare.

8.10 Limitations and Weaknesses

Despite encouraging results, several limitations should be acknowledged.

- The framework was evaluated primarily using publicly available biomedical and spaceflight datasets rather than operational astronaut mission data.
- Physiological responses observed under terrestrial simulation conditions may differ from those encountered during prolonged lunar or Martian missions.
- The hybrid AI architecture introduces additional computational complexity, which may require further optimization for deployment on resource-constrained onboard hardware.
- Although multiple public datasets were integrated, complete end-to-end validation using continuously synchronized astronaut Digital Twins has not yet been performed.
- Long-term adaptation of the Digital Twin under evolving physiological conditions requires continual learning mechanisms that remain outside the scope of the present study.

These limitations provide opportunities for further refinement and practical deployment in future space missions.

8.11 Threats to Validity

Several potential threats to validity should be considered when interpreting the results of this study.

Internal Validity: The experimental evaluation relies on publicly available datasets collected from multiple sources. Differences in acquisition protocols, sensor characteristics, and annotation standards may introduce variability into the integrated dataset.

External Validity: Although the selected datasets represent diverse physiological conditions, they may not fully capture the unique biomedical responses experienced during long-duration deep-space missions. Consequently, additional validation using future astronaut mission data is required.

Construct Validity: Performance was evaluated using established machine learning metrics including accuracy, precision, recall, F1-score, RMSE, AUC, and latency. While these metrics adequately measure predictive performance, they may not completely represent operational clinical usefulness in real mission scenarios.

Conclusion Validity: The reported results were obtained under controlled experimental conditions using statistically calibrated datasets. Although confidence intervals and uncertainty estimation improve the reliability of the findings, additional real-world validation is necessary before operational deployment.

9. CONCLUSION

The VYOMA system establishes a comprehensive, real-time Digital Twin framework for astronaut health monitoring and performance optimization, unifying multi-modal physiological sensing, hybrid artificial intelligence, and cyber-physical system design into a single operational paradigm. Unlike conventional monitoring systems that operate in a reactive and fragmented manner, VYOMA demonstrates that continuous state synchronization, probabilistic inference, and predictive modeling can transform astronaut healthcare into a proactive, anticipatory, and autonomous discipline. By integrating heterogeneous biosignals—ranging from cardiovascular and neuromuscular activity to environmental and cognitive indicators—the system constructs a dynamically evolving representation of the astronaut's latent

physiological state, enabling both high-resolution monitoring and forward-looking risk estimation.

The empirical results confirm that the hybrid architecture—combining convolutional neural networks for spatial feature extraction [44], [21], long short-term memory networks for temporal modeling [20], attention-based transformers for global dependency learning [45], and gradient boosting frameworks for structured data refinement [50], [51]—achieves superior predictive performance across all evaluation metrics. The observed accuracy of 96%, coupled with a low RMSE of 0.28 and a real-time inference latency of 47 milliseconds, demonstrates not only statistical superiority but also operational feasibility within mission-critical constraints. These findings are consistent with ensemble learning theory and bias-variance trade-off optimization principles [64], wherein the integration of complementary learning paradigms yields a more robust and generalizable predictive system than any single-model approach.

From a systems perspective, VYOMA advances the Digital Twin paradigm beyond its traditional industrial applications [16], [17] into the domain of biomedical cyber-physical systems, aligning with emerging frameworks in digital health [19] and data-centric engineering [15]. The architecture satisfies stringent real-time requirements defined in cyber-physical system theory [13], ensuring that sensing, inference, and decision support operate within latency bounds compatible with human physiological response times. Furthermore, the incorporation of uncertainty quantification through Bayesian approximations and Monte Carlo methods [24], [49] enhances the reliability and interpretability of predictions, which is essential for high-stakes environments such as human spaceflight.

The implications of this work extend significantly to deep space missions, where communication delays between Earth and spacecraft can range from 4 to 24 minutes, effectively eliminating the possibility of real-time ground intervention. In such scenarios, autonomous decision-making becomes not merely advantageous but indispensable. VYOMA addresses this requirement by embedding intelligence directly within the onboard system, enabling continuous monitoring, early anomaly detection, and context-aware intervention without reliance on Earth-based support. This capability is particularly critical for mitigating known spaceflight risks, including microgravity-induced bone

demineralization [1], muscle atrophy [2], neuro-ocular syndromes [4], radiation exposure [3], [12], and cognitive degradation under elevated CO₂ conditions [6]. By predicting these conditions before clinical thresholds are reached, the system enables preventive countermeasures, thereby enhancing mission safety and crew longevity.

Beyond space applications, the conceptual framework introduced in this work contributes to the broader vision of autonomous healthcare systems on Earth. The convergence of artificial intelligence, digital twins, and real-time data streams represents a paradigm shift in medical practice, moving from episodic diagnosis to continuous health intelligence. In this paradigm, healthcare is no longer confined to clinical settings but becomes an always-on, adaptive system capable of personalized intervention and long-term optimization. Such a transformation aligns with recent advances in AI-driven medicine [69][76], where predictive analytics and intelligent systems are increasingly integrated into clinical workflows.

Despite its strengths, the system faces several challenges that warrant further investigation. Data scarcity in extreme space environments limits the availability of ground-truth labels for rare physiological events, while domain shifts between terrestrial and extraterrestrial conditions introduce uncertainties in model generalization. Additionally, the computational complexity of hybrid architectures necessitates efficient deployment strategies, particularly for resource-constrained onboard systems. Addressing these challenges will require advancements in federated learning, allowing models to be updated across distributed spacecraft without centralized data aggregation, as well as the integration of causal inference frameworks [65] to move beyond correlation-based predictions toward mechanistic understanding.

Future research directions include the development of fully autonomous learning systems capable of adapting to novel environments in real time, the incorporation of genomic and pharmacological data for precision medicine applications, and the exploration of quantum-inspired optimization techniques for large-scale model training. These directions aim to further enhance the robustness, adaptability, and intelligence of digital twin systems, ultimately enabling a new class of self-evolving biomedical platforms.

In conclusion, VYOMA demonstrates that the integration of hybrid artificial intelligence with

digital twin technology constitutes a viable and powerful approach for addressing the complex challenges of astronaut health monitoring and performance optimization. The system not only achieves state-of-the-art predictive performance but also establishes a scalable and extensible framework for autonomous healthcare in extreme environments. By bridging the domains of artificial intelligence, cyber-physical systems, and space medicine, this work lays the foundation for the next generation of intelligent biomedical systems, where continuous, data-driven insight replaces episodic observation, and autonomous decision-making becomes the cornerstone of human survival beyond Earth.

10. FUTURE WORK

While VYOMA demonstrates a robust and scalable Digital Twin framework for astronaut health monitoring, several advanced research directions remain to be explored to elevate the system toward fully autonomous, next-generation space healthcare.

10.1 Federated Learning in Space Environments

One of the most critical extensions of the current system is the incorporation of federated learning for distributed model training across spacecraft and ground stations. In deep space missions, bandwidth constraints and communication delays make centralized data aggregation impractical. Federated learning enables decentralized model updates, where individual spacecraft or crew units locally train models on onboard data and share only model parameters instead of raw data.

Mathematically, this can be expressed as:

$$\theta^{\text{global}} = \sum_{i=1}^N \frac{n_i}{n} \theta_i \quad (34)$$

where:

- θ_i = local model parameters from spacecraft i
- n_i = number of local samples
- n = total samples across all nodes

This approach preserves data privacy, reduces communication overhead, and enables continuous learning in isolated environments. It is particularly relevant for long-duration missions such as Mars expeditions, where each spacecraft or habitat may encounter unique physiological patterns and environmental conditions.

10.2 Genomics and Precision Medicine Integration

Future iterations of VYOMA will incorporate genomic and multi-omics data to enable precision medicine capabilities. By integrating genetic, transcriptomic, and proteomic information from repositories such as NASA GeneLab [72], the Digital Twin can be personalized at the molecular level.

This allows the system to model individual variability in:

- Drug metabolism (pharmacogenomics)
- Radiation sensitivity
- Bone density loss susceptibility
- Immune response under microgravity

The Digital Twin state can be extended as:

$$x_t = f(\text{physiology, environment, genomics})(35)$$

Such integration transforms the system from population-level modeling to truly individualized healthcare, enabling tailored interventions and optimized treatment strategies for each astronaut.

10.3 VR/AR-Based Mission Surgeon Interface

Another promising direction is the development of immersive Virtual Reality (VR) and Augmented Reality (AR) interfaces for mission surgeons and onboard crew. By leveraging spatial computing platforms, the Digital Twin can be visualized as an interactive 3D entity, allowing clinicians to:

- Explore organ-level simulations in real time
- Visualize stress, radiation, and physiological heatmaps
- Simulate interventions before execution
- Monitor temporal evolution of health parameters in 3D space

This transforms the traditional dashboard into a fully immersive diagnostic environment, enhancing situational awareness and decision-making accuracy. Integration with real-time telemetry streams ensures that the virtual representation remains continuously synchronized with the physical astronaut.

10.4 Causal AI for Explainable Decision-Making

Current models primarily rely on correlation-based learning. Future work will integrate causal inference frameworks [65] to identify true cause-effect relationships between physiological variables and spaceflight stressors.

Using structural causal models:

$$Y = f(X, U) \quad (36)$$

where:

- X = observed variables
- U = latent confounders

This enables:

- Counterfactual reasoning ("What if exercise increased?")
- Intervention planning
- Improved interpretability for clinical decisions

10.5 Self-Adaptive and Continual Learning Systems

To handle non-stationary environments in space, future systems will incorporate continual learning mechanisms that adapt to evolving physiological conditions without catastrophic forgetting. Drift detection techniques combined with automated retraining pipelines will ensure that models remain accurate throughout the mission duration.

10.6 Edge AI and Energy-Efficient Inference

Given the constraints of onboard computational resources, optimizing models for edge deployment is essential. Techniques such as model pruning, quantization, and knowledge distillation will be explored to reduce computational overhead while maintaining performance.

10.7 Multi-Crew Digital Twin Federation

For future space stations such as the Bharatiya Antariksh Station (BAS), VYOMA can be extended to support multi-crew Digital Twin federation. This enables:

- Group-level health analytics
- Interaction modeling between crew members
- Collective stress and performance monitoring

Such a system can model not only individual health but also team dynamics, which are critical in confined and high-stress environments.

10.8 Autonomous Medical Intervention Systems

The ultimate goal of this research is to evolve toward fully autonomous medical systems capable of not only predicting health risks but also executing interventions. This includes:

- AI-driven exercise recommendations
- Adaptive sleep scheduling

- Automated drug dosing guidance
- Robotic-assisted medical procedures (future scope)

10.9 Final Outlook

The future evolution of VYOMA lies in its transformation from a monitoring and prediction system into a fully autonomous, self-learning, and self-intervening healthcare ecosystem. By integrating federated intelligence, molecular-level personalization, immersive visualization, and causal reasoning, the system has the potential to redefine human healthcare not only in space but also on Earth.

11. DECLARATIONS

11.1 Ethical Considerations

This study did not involve the recruitment of human participants or the use of animals, and no new clinical data were collected by the authors. All physiological, environmental and behavioural data used in this work were obtained from publicly available, de-identified repositories maintained by NASA, ESA, NOAA and PhysioNet, and were used strictly in accordance with the access conditions and data-use policies published by those repositories. No personally identifiable information was accessed, stored or reproduced at any stage of the study. The proposed framework is presented as a research prototype and is not intended, in its present form, for autonomous clinical decision-making without further verification, validation and regulatory approval.

11.2 Funding

This research was carried out under the ISRO Human Space Flight Centre research initiative RES-HSFC-2025-001. The funding body had no role in the study design, in the collection, analysis and interpretation of data, in the writing of this manuscript, or in the decision to submit the article for publication.

11.3 Conflict of Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

11.4 Data Availability

All datasets analysed in this study are publicly available from the NASA Life Sciences Data Archive, the NASA Open Science Data Repository, the NASA Human Research Program,

the ESA data services, the NOAA Space Weather Prediction Center and the PhysioNet repository, as cited in the references. No proprietary or restricted datasets were used.

11.5 Declaration on the Use of Artificial Intelligence in Writing

The authors declare that generative artificial intelligence tools were used only for language editing, grammatical correction and formatting of the manuscript. No artificial intelligence tool was used to generate scientific content, to design or execute the experiments, to analyse or interpret the results, or to produce the conclusions reported in this work. All AI-assisted text was subsequently reviewed, verified and edited by the authors, who take full responsibility for the integrity, originality and accuracy of the entire manuscript.

REFERENCES

- [1] T. Lang, A. LeBlanc, H. Evans, Y. Lu, H. Genant, A. Yu, "Cortical and Trabecular Bone Mineral Loss from the Spine and Hip in Long-Duration Spaceflight", *Journal of Bone and Mineral Research*, Vol. 19, No. 6, 2004, pp. 1006–1012.
- [2] R.H. Fitts, S.W. Trappe, D.L. Costill, P.M. Gallagher, A.C. Creer, R.A. Colloton, J.R. Romatowski, J.L. Bain, D.A. Riley, "Prolonged Space Flight-Induced Alterations in the Structure and Function of Human Skeletal Muscle Fibres", *The Journal of Physiology*, Vol. 588, No. 18, 2010, pp. 3567–3592.
- [3] NASA Space Radiation Laboratory (NSRL), Brookhaven National Laboratory, Upton (USA), <https://www.bnl.gov/nsrl>, 09:30 GMT, 15/07/2026.
- [4] T.H. Mader, C.R. Gibson, A.F. Pass, L.A. Kramer, A.G. Lee, J. Fogarty, W.J. Tarver, J.P. Dervay, D.R. Hamilton, A. Sargsyan, J.L. Phillips, D. Tran, W. Lipsky, J. Choi, C. Stern, R. Kuyumjian, J.D. Polk, "Optic Disc Edema, Globe Flattening, Choroidal Folds, and Hyperopic Shifts Observed in Astronauts after Long-Duration Space Flight", *Ophthalmology*, Vol. 118, No. 10, 2011, pp. 2058–2069.
- [5] L.K. Barger, E.E. Flynn-Evans, A. Kubey, L. Walsh, J.M. Ronda, W. Wang, K.P. Wright Jr., C.A. Czeisler, "Prevalence of Sleep Deficiency and Use of Hypnotic Drugs in Astronauts Before, During, and After Spaceflight: An Observational Study", *The Lancet Neurology*, Vol. 13, No. 9, 2014, pp. 904–912.
- [6] J. Law, M. Van Baalen, M. Foy, S.S. Mason, C. Mendez, M.L. Wear, V.E. Meyers, N. Alexander, "Relationship Between Carbon Dioxide Levels and Reported Headaches on the International Space Station", *Journal of Occupational and Environmental Medicine*, Vol. 56, No. 5, 2014, pp. 477–483.
- [7] M. Basner, D.F. Dinges, D. Mollicone, A. Ecker, C.W. Jones, E.C. Hyder, A. Di Antonio, I. Savelev, K. Kan, N. Goel, B.V. Morukov, J.P. Sutton, "Mars 520-d Mission Simulation Reveals Protracted Crew Hypokinesia and Alterations of Sleep Duration and Timing", *Proceedings of the National Academy of Sciences (PNAS)*, Vol. 110, No. 7, 2013, pp. 2635–2640.
- [8] G.M. Sandal, G.R. Leon, L. Palinkas, "Human Challenges in Polar and Space Environments", *Reviews in Environmental Science and Bio/Technology*, Vol. 5, No. 2-3, 2006, pp. 281–296.
- [9] NASA Human Research Program, "Human Research Roadmap", National Aeronautics and Space Administration, Houston (USA), <https://humanresearchroadmap.nasa.gov>, 09:00 GMT, 15/07/2026.
- [10] NASA Human Research Program, "Evidence Report: Risk of Adverse Cognitive or Behavioral Conditions and Psychiatric Disorders", National Aeronautics and Space Administration, Houston (USA), 2016, <https://humanresearchroadmap.nasa.gov/evidence>, 09:10 GMT, 15/07/2026.
- [11] J. Conkin, "Evidence Report: Risk of Decompression Sickness (DCS)", NASA Human Research Program, Houston (USA), 2016, <https://humanresearchroadmap.nasa.gov/evidence>, 09:15 GMT, 15/07/2026.
- [12] European Space Agency, "SPENVIS – Space Environment Information System, Radiation Environment Models", ESA and Royal Belgian Institute for Space Aeronomy, Brussels (Belgium), <https://www.spenvis.oma.be>, 09:20 GMT, 15/07/2026.
- [13] E.A. Lee, "Cyber Physical Systems: Design Challenges", *Proceedings of the 11th IEEE International Symposium on Object and Component-Oriented Real-Time Distributed Computing (ISORC)*, Orlando (USA), May 5-7, 2008, pp. 363–369.
- [14] E.H. Glaessgen, D.S. Stargel, "The Digital Twin Paradigm for Future NASA and U.S. Air Force Vehicles", *Proceedings of the 53rd*

- AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, Honolulu (USA), April 23-26, 2012, pp. 1–14.
- [15] D.J. Wagg, K. Worden, R.J. Barthorpe, P. Gardner, “Digital Twins: State-of-the-Art and Future Directions for Modeling and Simulation in Engineering Dynamics Applications”, *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B: Mechanical Engineering*, Vol. 6, No. 3, 2020, pp. 1–18.
- [16] F. Tao, J. Cheng, Q. Qi, M. Zhang, H. Zhang, F. Sui, “Digital Twin-Driven Product Design, Manufacturing and Service with Big Data”, *The International Journal of Advanced Manufacturing Technology*, Vol. 94, No. 9-12, 2018, pp. 3563–3576.
- [17] Q. Qi, F. Tao, “Digital Twin and Big Data Towards Smart Manufacturing and Industry 4.0: 360 Degree Comparison”, *IEEE Access*, Vol. 6, 2018, pp. 3585–3593.
- [18] J. Lee, H. Davari, J. Singh, V. Pandhare, “Industrial Artificial Intelligence for Industry 4.0-Based Manufacturing Systems”, *Manufacturing Letters*, Vol. 18, 2018, pp. 20–23.
- [19] K. Bruynseels, F. Santoni de Sio, J. van den Hoven, “Digital Twins in Health Care: Ethical Implications of an Emerging Engineering Paradigm”, *Frontiers in Genetics*, Vol. 9, 2018, pp. 1–11.
- [20] S. Hochreiter, J. Schmidhuber, “Long Short-Term Memory”, *Neural Computation*, Vol. 9, No. 8, 1997, pp. 1735–1780.
- [21] K. He, X. Zhang, S. Ren, J. Sun, “Deep Residual Learning for Image Recognition”, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas (USA), June 27-30, 2016, pp. 770–778.
- [22] D. Barber, “Bayesian Reasoning and Machine Learning”, Cambridge University Press, Cambridge (UK), 2012.
- [23] R.S. Sutton, A.G. Barto, “Reinforcement Learning: An Introduction”, 2nd Edition, MIT Press, Cambridge (USA), 2018.
- [24] Y. Gal, Z. Ghahramani, “Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning”, *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, New York (USA), June 19-24, 2016, pp. 1050–1059.
- [25] S.M. Lundberg, S.I. Lee, “A Unified Approach to Interpreting Model Predictions”, *Advances in Neural Information Processing Systems (NIPS)*, Long Beach (USA), December 4-9, 2017, pp. 4765–4774.
- [26] M. Sundararajan, A. Taly, Q. Yan, “Axiomatic Attribution for Deep Networks”, *Proceedings of the 34th International Conference on Machine Learning (ICML)*, Sydney (Australia), August 6-11, 2017, pp. 3319–3328.
- [27] J. Kreps, N. Narkhede, J. Rao, “Kafka: A Distributed Messaging System for Log Processing”, *Proceedings of the 6th International Workshop on Networking Meets Databases (NetDB)*, Athens (Greece), June 12, 2011, pp. 1–7.
- [28] P. Carbone, A. Katsifodimos, S. Ewen, V. Markl, S. Haridi, K. Tzoumas, “Apache Flink: Stream and Batch Processing in a Single Engine”, *IEEE Data Engineering Bulletin*, Vol. 38, No. 4, 2015, pp. 28–38.
- [29] NVIDIA Corporation, “Triton Inference Server Documentation”, Santa Clara (USA), <https://docs.nvidia.com/deeplearning/triton-inference-server>, 10:00 GMT, 15/07/2026.
- [30] A.E.W. Johnson, L. Bulgarelli, L. Shen, A. Gayles, A. Shammout, S. Horng, T.J. Pollard, S. Hao, B. Moody, B. Gow, L.H. Lehman, L.A. Celi, R.G. Mark, “MIMIC-IV, a Freely Accessible Electronic Health Record Dataset”, *Scientific Data*, Vol. 10, 2023, pp. 1–9.
- [31] P. Wagner, N. Strodthoff, R.D. Bousseljot, D. Kreiseler, F.I. Lunze, W. Samek, T. Schaeffter, “PTB-XL, a Large Publicly Available Electrocardiography Dataset”, *Scientific Data*, Vol. 7, 2020, pp. 1–15.
- [32] F.E. Garrett-Bakelman et al., “The NASA Twins Study: A Multidimensional Analysis of a Year-Long Human Spaceflight”, *Science*, Vol. 364, No. 6436, 2019, pp. 1–20.
- [33] R. Rosen, G. von Wichert, G. Lu, K. Bettenhausen, “About the Importance of Autonomy and Digital Twins for the Future of Manufacturing”, *IFAC-PapersOnLine*, Vol. 48, No. 3, 2015, pp. 567–572.
- [34] Y. Liu, L. Zhang, Y. Yang, L. Zhou, L. Ren, F. Wang, R. Liu, Z. Pang, M.J. Deen, “A Novel Cloud-Based Framework for the Elderly Healthcare Services Using Digital Twin”, *IEEE Access*, Vol. 7, 2019, pp. 49088–49101.
- [35] M. Satyanarayanan, “The Emergence of Edge Computing”, *Computer*, Vol. 50, No. 1, 2017, pp. 30–39.
- [36] NASA Life Sciences Data Archive (LSDA), Johnson Space Center, Houston (USA),

- https://lsda.jsc.nasa.gov, 10:10 GMT, 15/07/2026.
- [37] G.B. Moody, R.G. Mark, "The Impact of the MIT-BIH Arrhythmia Database", *IEEE Engineering in Medicine and Biology Magazine*, Vol. 20, No. 3, 2001, pp. 45–50.
- [38] I. Obeid, J. Picone, "The Temple University Hospital EEG Data Corpus", *Frontiers in Neuroscience*, Vol. 10, 2016, pp. 1–5.
- [39] B. Kemp, A.H. Zwiderman, B. Tuk, H.A.C. Kamphuisen, J.J.L. Obery, "Analysis of a Sleep-Dependent Neuronal Feedback Loop: The Slow-Wave Microcontinuity of the EEG", *IEEE Transactions on Biomedical Engineering*, Vol. 47, No. 9, 2000, pp. 1185–1194.
- [40] M. Atzori, A. Gijsberts, C. Castellini, B. Caputo, A.G.M. Hager, S. Elsig, G. Giatsidis, F. Bassetto, H. Muller, "Electromyography Data for Non-Invasive Naturally-Controlled Robotic Hand Prostheses", *Scientific Data*, Vol. 1, 2014, pp. 1–13.
- [41] J. Pan, W.J. Tompkins, "A Real-Time QRS Detection Algorithm", *IEEE Transactions on Biomedical Engineering*, Vol. BME-32, No. 3, 1985, pp. 230–236.
- [42] D.F. Swinehart, "The Beer-Lambert Law", *Journal of Chemical Education*, Vol. 39, No. 7, 1962, pp. 333–335.
- [43] R.E. Kalman, "A New Approach to Linear Filtering and Prediction Problems", *Journal of Basic Engineering*, Vol. 82, No. 1, 1960, pp. 35–45.
- [44] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, "Gradient-Based Learning Applied to Document Recognition", *Proceedings of the IEEE*, Vol. 86, No. 11, 1998, pp. 2278–2324.
- [45] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L. Kaiser, I. Polosukhin, "Attention Is All You Need", *Advances in Neural Information Processing Systems (NIPS)*, Long Beach (USA), December 4-9, 2017, pp. 5998–6008.
- [46] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, N. Houlsby, "An Image Is Worth 16x16 Words: Transformers for Image Recognition at Scale", *Proceedings of the International Conference on Learning Representations (ICLR)*, Vienna (Austria), May 3-7, 2021, pp. 1–21.
- [47] B. Lim, S.O. Arik, N. Loeff, T. Pfister, "Temporal Fusion Transformers for Interpretable Multi-Horizon Time Series Forecasting", *International Journal of Forecasting*, Vol. 37, No. 4, 2021, pp. 1748–1764.
- [48] S.J. Taylor, B. Letham, "Forecasting at Scale", *The American Statistician*, Vol. 72, No. 1, 2018, pp. 37–45.
- [49] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, E. Teller, "Equation of State Calculations by Fast Computing Machines", *The Journal of Chemical Physics*, Vol. 21, No. 6, 1953, pp. 1087–1092.
- [50] T. Chen, C. Guestrin, "XGBoost: A Scalable Tree Boosting System", *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, San Francisco (USA), August 13-17, 2016, pp. 785–794.
- [51] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, T.Y. Liu, "LightGBM: A Highly Efficient Gradient Boosting Decision Tree", *Advances in Neural Information Processing Systems (NIPS)*, Long Beach (USA), December 4-9, 2017, pp. 3146–3154.
- [52] D. Silver, A. Huang, C.J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, D. Hassabis, "Mastering the Game of Go with Deep Neural Networks and Tree Search", *Nature*, Vol. 529, No. 7587, 2016, pp. 484–489.
- [53] J. Bergstra, D. Yamins, D.D. Cox, "Making a Science of Model Search: Hyperparameter Optimization in Hundreds of Dimensions for Vision Architectures", *Proceedings of the 30th International Conference on Machine Learning (ICML)*, Atlanta (USA), June 16-21, 2013, pp. 115–123.
- [54] M. Abadi et al., "TensorFlow: A System for Large-Scale Machine Learning", *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, Savannah (USA), November 2-4, 2016, pp. 265–283.
- [55] A. Paszke et al., "PyTorch: An Imperative Style, High-Performance Deep Learning Library", *Advances in Neural Information Processing Systems (NeurIPS)*, Vancouver (Canada), December 8-14, 2019, pp. 8026–8037.
- [56] OASIS, "MQTT Version 5.0 – OASIS Standard", Organization for the Advancement of Structured Information Standards, 2019,

- <https://docs.oasis-open.org/mqtt/mqtt/v5.0/mqtt-v5.0.html>, 10:20 GMT, 15/07/2026.
- [57] Timescale Inc., “TimescaleDB Documentation”, New York (USA), <https://docs.timescale.com>, 10:25 GMT, 15/07/2026.
- [58] M. Zaharia, A. Chen, A. Davidson, A. Ghodsi, S.A. Hong, A. Konwinski, S. Murching, T. Nykodym, P. Ogilvie, M. Parkhe, F. Xie, C. Zumar, “Accelerating the Machine Learning Lifecycle with MLflow”, *IEEE Data Engineering Bulletin*, Vol. 41, No. 4, 2018, pp. 39–45.
- [59] Iterative Inc., “DVC: Data Version Control – Git for Data and Models”, <https://dvc.org>, 10:30 GMT, 15/07/2026.
- [60] N. Kokhlikyan, V. Miglani, M. Martin, E. Wang, B. Alsallakh, J. Reynolds, A. Melnikov, N. Kliushkina, C. Araya, S. Yan, O. Reblitz-Richardson, “Captum: A Unified and Generic Model Interpretability Library for PyTorch”, *arXiv preprint arXiv:2009.07896*, 2020, pp. 1–13.
- [61] V.J. Lawhern, A.J. Solon, N.R. Waytowich, S.M. Gordon, C.P. Hung, B.J. Lance, “EEGNet: A Compact Convolutional Neural Network for EEG-Based Brain-Computer Interfaces”, *Journal of Neural Engineering*, Vol. 15, No. 5, 2018, pp. 1–17.
- [62] O. Ronneberger, P. Fischer, T. Brox, “U-Net: Convolutional Networks for Biomedical Image Segmentation”, *Proceedings of Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, Munich (Germany), October 5-9, 2015, pp. 234–241.
- [63] G.E.P. Box, “Science and Statistics”, *Journal of the American Statistical Association*, Vol. 71, No. 356, 1976, pp. 791–799.
- [64] C.M. Bishop, “Pattern Recognition and Machine Learning”, Springer, New York (USA), 2006.
- [65] J. Pearl, “Causality: Models, Reasoning and Inference”, 2nd Edition, Cambridge University Press, Cambridge (UK), 2009.
- [66] National Oceanic and Atmospheric Administration, “Space Weather Prediction Center Datasets”, NOAA, Boulder (USA), <https://www.swpc.noaa.gov>, 10:35 GMT, 15/07/2026.
- [67] NASA OMNIWeb Plus, “Heliospheric and Near-Earth Solar Wind Data Service”, Goddard Space Flight Center, Greenbelt (USA), <https://omniweb.gsfc.nasa.gov>, 10:40 GMT, 15/07/2026.
- [68] T. Poikela, J. Plosila, T. Westerlund, M. Campbell, M. De Gaspari, X. Llopart, V. Gromov, R. Kluit, M. van Beuzekom, F. Zappone, V. Zivkovic, C. Brezina, K. Desch, Y. Fu, A. Kruth, “Timepix3: A 65K Channel Hybrid Pixel Readout Chip with Simultaneous ToA/ToT and Sparse Readout”, *Journal of Instrumentation*, Vol. 9, No. 5, 2014, pp. 1–22.
- [69] P. Rajpurkar, E. Chen, O. Banerjee, E.J. Topol, “AI in Health and Medicine”, *Nature Medicine*, Vol. 28, No. 1, 2022, pp. 31–38.
- [70] I. Goodfellow, Y. Bengio, A. Courville, “Deep Learning”, MIT Press, Cambridge (USA), 2016.
- [71] J. Schmidhuber, “Deep Learning in Neural Networks: An Overview”, *Neural Networks*, Vol. 61, 2015, pp. 85–117.
- [72] NASA Open Science Data Repository (OSDR/GeneLab), Ames Research Center, Moffett Field (USA), <https://osdr.nasa.gov>, 10:45 GMT, 15/07/2026.
- [73] A. Graves, “Supervised Sequence Labelling with Recurrent Neural Networks”, *Studies in Computational Intelligence*, Vol. 385, Springer, Berlin (Germany), 2012.
- [74] S. Russell, P. Norvig, “Artificial Intelligence: A Modern Approach”, 4th Edition, Pearson Education, Harlow (UK), 2021.
- [75] S. Thrun, W. Burgard, D. Fox, “Probabilistic Robotics”, MIT Press, Cambridge (USA), 2005.
- [76] E.A. Katsoulakis, Q. Wang, H. Wu, L. Shahriyari, R. Fletcher, J. Liu, L. Achenie, H. Liu, P. Jackson, Y. Xiao, T. Syeda-Mahmood, R. Tuli, J. Deng, “Digital Twins for Health: A Scoping Review”, *npj Digital Medicine*, Vol. 7, 2024, pp. 1–11.