

IoT and Artificial Intelligence for Secure Remote Healthcare Monitoring

This research presents a layered approach to remote healthcare by integrating the Internet of Things (IoT), Blockchain, and Reinforcement Learning (RL). The aim is to provide personalised, secure health monitoring, especially for patients in remote or underserved areas.




Figure 1. Overview of the network model.

Figure 1 illustrates the architecture of the proposed secure remote healthcare monitoring system, which integrates three fundamental technological layers. The IoT Layer gathers real-time patient vital signs, including heart rate, temperature, and blood pressure, through connected wearable sensors. This information is subsequently transferred to the blockchain layer, where it is securely stored in immutable blocks, thereby ensuring integrity, transparency, and resistance to tampering. Finally, the Reinforcement Learning Layer analyses the data using a Q-learning algorithm combined with a Long Short-Term Memory (LSTM) network. The model finds the best treatment options—choosing whether to keep an eye on the patient or give medication—by regularly updating the Q-table and predicting results using the LSTM's ability to understand time. This layered integration facilitates smart, secure, and personalised healthcare decision-making.

Dataset, Health Metrics, and Underlying Assumptions

This study uses a simulated healthcare dataset to enable ethical and rapid development of the reinforcement learning model in a controlled environment.

- **Simulated Dataset:** 5,000 synthetic patient records with 13 weeks of time-series health data.
- **Attributes:** Includes vital signs (temperature, BP, heart rate, respiration), BMI, risk classification, and treatment actions.
- **Actions:** Two treatment decisions are modelled: *Monitor* or *Administer Medication*.

Assumptions and Distributions:

To simulate realistic and diverse patient data, the following assumptions and distributions were used:

- **Age:** Uniform distribution — $A \sim \text{Uniform}(16, 65)$
- **Vital Signs:** Modelled as Gaussian — $V \sim \mathcal{N}(\mu, \sigma^2)$
- **Blood Pressure:** Adjusted for age and time — $BP = BP_{base} + \alpha \cdot \text{Age} + \beta \cdot \text{Time}$
- **BMI:** Calculated from height and weight — $\text{BMI} = \frac{\text{Weight}_{kg}}{(\text{Height}_m)^2}$
- **Risk Classification:** Rule-based using thresholds on age, BP, and BMI.

These assumptions enable the model to train on physiologically plausible trends while maintaining ethical integrity through synthetic data.

Model Evaluation and Learning Analysis

To assess the effectiveness of the reinforcement learning (RL) model in making treatment decisions, we employed key evaluation metrics that reflect both learning performance and policy optimisation over time.

1. **Cumulative Reward:** Tracks the total reward accumulated per episode. A rising trend indicates the agent is learning to take better actions—such as deciding when to monitor versus intervene—based on patient state. Higher cumulative rewards signal more effective treatment decisions and policy refinement.
2. **Rolling Average Reward:** This smoothed metric averages the total rewards over a window of episodes to highlight long-term performance improvements while reducing short-term noise.




Figure 2. Rolling Average of the Model

As shown in Figure 2, the blue line represents the total reward per episode, while the orange dashed line illustrates the 10-episode rolling average. Key observations:

- **Initial Variability:** Early training shows fluctuations due to exploration.
- **Learning Trend:** The rolling average steadily increases, indicating improved decision-making.
- **Convergence:** The curve flattens in later episodes, suggesting policy stabilisation.

Agent Treatment Decisions

This plot shows the policy learnt by the reinforcement learning (RL) agent when choosing between *Monitor* and *Administer Medication* actions across all simulated patient episodes.

Distribution of Actions Taken




Figure 3. Distribution of Treatment Decisions by the Agent

Interpretation:

- **Monitor (65%):** The agent most frequently elects to observe rather than intervene. This indicates a conservative, safety-aware policy—useful when overtreatment carries cost or risk. It also reflects the reward structure that penalises unnecessary medication.
- **Administer Medication (35%):** Interventions are reserved for episodes where vitals or risk indicators cross learnt thresholds. This targeted behaviour suggests the model differentiates stable from deteriorating states.

Project Achievements and Future Contributions

This research presents a novel, multi-layered framework that integrates IoT, blockchain, and reinforcement learning to enable secure, intelligent, and personalised remote healthcare monitoring.

Key Achievements:

- Developed a fully simulated RL model using Q-learning and LSTM, capable of making adaptive healthcare decisions based on real-time patient vitals.
- Demonstrated stable policy learning through cumulative and rolling average rewards, with the model converging toward optimal treatment strategies.
- Maintained a safe intervention policy—favouring monitoring (65%) over unnecessary medication (35%)—aligned with ethical healthcare principles.
- Ensured regulatory alignment by designing for HIPAA and GDPR compliance in data handling and patient privacy.

Future Contributions:

- Extend the model to handle diverse chronic conditions and patient demographics using real-world data.
- Integrate the system with real-time wearable device platforms for live deployment in remote settings.
- Advance interoperability between heterogeneous IoT and blockchain systems for seamless healthcare data flow.
- Collaborate with healthcare providers to pilot and validate the system in clinical environments.

This project lays the groundwork for intelligent, privacy-preserving digital health solutions that can scale across regions, ultimately transforming how care is delivered in low-resource and remote communities.

References

- [1] T.T. Kuo, H.E. Kim, and L. Ohno-Machado. 2017. Blockchain distributed ledger technologies for biomedical and health care applications. *Journal of the American Medical Informatics Association* 24, 6 (2017), 1211–1220. Oxford University Press. DOI: <https://doi.org/10.1093/jamia/ocx068>
- [2] Omar Cheikhrouhou, Khaleel Mershad, Faisal Jamil, Rashaad Mahmud, Anis Koubâa, and Sanaz Rahimi Moosavi. 2023. A lightweight blockchain and fog-enabled secure remote patient monitoring system. *IEEE Access* 11 (2023), 17671–17686. [3] Andre Esteva, Katherine Chou, Serena Yeung, Nikhil Naik, Ali Madani, Ali Mottaghi, Yun Liu, Eric Topol, Richard Socher, and Andrew Y. Ng. 2021. Deep learning-enabled medical computer vision. *npj Digital Medicine* 4, 1 (2021), 1–9. Nature Publishing Group. DOI: <https://doi.org/10.1038/s41746-020-00376-2>
- [4] Abdur Rehman, Sagheer Abbas, Kashif Munir, Asad Masood Khattak, Gwanggil Jeon, Kyungroul Lee, Mueen Uddin, and Ali Kashif Bashir. 2022. A secure healthcare 5.0 system based on blockchain technology entangled with federated learning technique. *Sensors* 22, 18 (2022), 6816. MDPI. DOI: <https://doi.org/10.3390/s22186816>