

Quantum Intelligent Agent of Medical Decision: The Theoretical Machine

EUSTACHE MUTEBA A.¹, NIKOS E. MASTORAKIS²

¹Technical University of Sofia & INTERBIT
Sofia, BULGARIA

²English Language Faculty of Engineering
Technical University of Sofia & INTERBIT
Sofia, BULGARIA

Abstract:— This paper presents the theoretical machine of a Quantum Intelligent Agents of Medical Decision. Introducing the concept of Intelligent Agents in Quantum Deep Q-Learning involves integrating principles of reinforcement learning, quantum computing, and intelligent agent frameworks. The theoretical machine is designed to simulate quantum computer and to explore quantum algorithms without the need for physical hardware (quantum computer). Our approach does not refer directly to any simulator. Oriented to the Sickle Cell Disease, the approach leverages quantum-inspired representations of classical Deep Reinforcement Learning to handle management of related medical data for an optimized treatment recommendations.

Keywords:— Quantum Deep Q-Learning, intelligent agents, theoretical machine, medical decision

Received: June 30, 2025. Revised: August 6, 2025. Accepted: September 2, 2025. Published: September 24, 2025.

1. Introduction

Face to medical decision of complex diseases healthcare professionals often deal with a variety of uncertain, incomplete, and imprecise information derived from patient symptoms, medical imaging, laboratory tests, and historical data. One approach to manage uncertainty and complexity in medical decision, ultimately leading to more accurate and reliable diagnostic processes, involves to incorporate advanced Artificial Intelligence techniques. Among them, there are the Deep Reinforcement Learning (DRL). It enables AI systems to learn optimal decision-making strategies by interacting with medical data and refining predictions over time.

However, implementing DRL in healthcare comes with several challenges such as the data quality and availability [1], and the computational complexity [2]. In one hand, the healthcare data is often incomplete, noisy, or biased, making it difficult for DRL models to learn effectively. In the other one, the DRL requires high computational power.

Nowadays, addressing the limitations of DRL in complex medical decisions are tackling through several innovative approaches: data augmentation, collaborative data sharing, safe exploration techniques, simulated environments, model optimization, hybrid models, transfer learning, domain adaptation, complex reward functions, multi-objective optimization, cloud computing, efficient algorithms, hierarchical reinforcement learning, etc.

Anyways, ideal is to create more efficient and capable learning agents that can make faster, more accurate, and context-aware medical decisions. So, standing down on Deep Reinforcement Learning, it is possible to integrate the principles inspired by Quantum Deep Q-Learning (QDQL) which offers the computing advantages of parametric quantum circuits (Quantum Neural Networks (QNN)) [3, 4].

Parametric quantum circuits, also known as Quantum Neural Networks (QNNs), represent a significant intersection between quantum computing and machine learning. There are

parameterized, meaning they can be adjusted (trained) based on the data input. Thus, they could potentially learn from fewer data points or represent more complex relationships.

Current devices cannot still reach the high expectations that quantum computing promises. However, there exists simulators (e.g., Qiskit, PennyLane, or Cirq) to implement quantum neural networks (QNNs) or quantum-enhanced policies and simulating quantum computing principles [5].

This paper presents the theoretical machine of a Quantum Intelligent Agents of Medical Decision (QIAMED). Introducing the concept of Intelligent Agents in Quantum Deep Q-Learning (Q-DQL) involves integrating principles of reinforcement learning, quantum computing, and intelligent agent frameworks. The theoretical machine is designed to simulate quantum computer and to explore quantum algorithms without the need for physical hardware (quantum computer). Our approach does not refer directly to any simulator.

Oriented to the Sickle Cell Disease (SCD), the approach leverages quantum-inspired of classical Deep Reinforcement Learning (DRL) to handle management of related medical data for an optimized treatment recommendations.

According to [6], Sickle Cell Disease (SCD) is characterized by variable clinical outcomes, with some patients suffering life-threatening complications during childhood, and others living relatively symptom-free into old age. Because of this variability, there is an important potential role for precision medicine, in which particular different treatments are selected for different groups of patients.

2. Methods

The interest in exploring and developing use cases for Quantum Deep Q-Learning in the healthcare field to improve decision-making processes is not new. So there are several studies.

Niraula [7], develop, in oncology, “a novel quantum deep reinforcement learning (qDRL) framework for clinical decision support that can estimate an individual patient’s dose

response mid-treatment and recommend an optimal dose adjustment". Their model was trained in the IBMQ quantum processor. Chow in [8] proposes a review paper that "examines the foundational concepts, key applications, and challenges of these technologies in healthcare, explores their potential synergy in solving clinical problems, and outlines future directions for quantum-enhanced ML in medical decision-making". From a general theoretical point of view on quantum computing, we mention the tutorial proposed in [9], that "introduces the mathematical framework of quantum algorithms ranging from basic elements including quantum bits and quantum gates to more advanced concepts such as variational quantum algorithms and quantum errors."

What emerges from all these articles is a set of concepts necessary to understand the contribution of quantum computing theories to classical computation.

It therefore follows that our theoretical machine refers to a computational or abstract model of a QDQL-based agent whose learning process and decision-making are governed by theoretical principles of quantum circuits and reinforcement logic.

This approach leverages the strengths of both quantum and classical computing, aiming to enhance the efficiency and capability of reinforcement learning agents.

2.1 Intelligent Agent

Intelligent agent is a computer system or entity that can perceive their environment, reason about it, and act upon it to achieve specific goals.

Thus, the mathematical formalization of intelligent agents is based on defining their environment, actions, perceptions, and goals. This includes their strategies, optimizing their decision-making processes, and applying learning and adaptation algorithms.

Intelligent agents are often modeled using Markov Decision Processes (MDPs), where:

(S): Set of states.

(A): Set of actions.

($T(s, a, s')$): Transition probability function, denoting the probability of moving to state (s') after taking action (a) in state (s).

($R(s, a)$): Reward function.

(γ): Discount factor (between 0 and 1) that models the importance of future rewards.

2.2 Deep Q-learning

Deep Q-Learning is an extension of the traditional Q-learning algorithm that uses deep neural networks to approximate the Q-value function.

Q-learning is a model-free reinforcement learning algorithm used to learn the value of an action in a given state.

It is particularly useful for solving Markov Decision Processes (MDPs) where the agent learns to make decisions by interacting with its environment.

Q-learning is guaranteed to converge to the optimal Q-values, given sufficient exploration of the state-action space and appropriate learning parameters. The convergence is typically achieved through repeated interaction with the environment and updating the Q-values based on the observed rewards.

2.3 Quantum Computing

1) Qubits

The basic unit of a quantum circuit is a qubit (quantum bit), which can exist in a superposition of states (0, 1, or both), represented mathematically as: $\{|0\rangle, |1\rangle, |0\rangle$ and $|1\rangle\}$.

2) Quantum Gates

Quantum gates operate on qubits and are the building blocks of quantum circuits. They manipulate the state of qubits through unitary transformations.

Common quantum gates include:

- Hadamard Gate (H): Creates superposition.
- Pauli Gates (X, Y, Z): Perform bit-flip and phase-flip operations.
- $CNOT$ Gate: A two-qubit gate that flips the state of a target qubit if the control qubit is in the state $|1\rangle$.
- Phase Gates: Introduce relative phase shifts between states.

3) Quantum Circuit

A quantum circuit is a sequence of quantum gates applied to a set of qubits. The circuit operates from left to right, with the input state on the left and the output state on the right.

Quantum circuits manipulate qubits through quantum gates, which are the quantum analogs of classical logic gates.

4) Measurement

At the end of a quantum circuit, qubits are typically measured to obtain classical bits. The measurement collapses the quantum state into one of the basis states with probabilities determined by the state vector.

2.4 Framework of Quantum Intelligent Agent (QIA)

5) Quantum State Space

The quantum states or quantum-inspired probabilistic representations of an agent can be characterized as a set $S = \{s_1, s_2, \dots, s_j\}$. At any given instant, the environment env is assumed to be in one these states.

Let S be the set of quantum-inspired states, represented as vectors $|s\rangle$ in a Hilbert space \mathcal{H} [10, 11], encoding superpositions of classical states with amplitudes $c_i \in \mathbb{C}$, such that:

$$|s\rangle = \sum_i c_i |\varphi_i\rangle \quad (1)$$

where $|\varphi_i\rangle$ are basis vectors of \mathcal{H} , c_i are complex coefficients satisfying $\sum |c_i|^2 = 1$ (normalization condition).

6) Quantum Action Space

The set of action space $A = \{a_1, a_2, \dots, a_i\}$ consists of predefined quantum operations, such as applying specific quantum gates (e.g., Hadamard, Pauli-X, or controlled operations) [12].

7) Quantum Q-Function Approximation

Quantum Q-Function Approximation is used to estimate the Q-values that guide an agent's decision-making process.

In classical reinforcement learning, the Q-function represents the expected reward for taking an action in a given state.

Quantum approaches aim to enhance this approximation using parameterized quantum circuits (PQCs) [13].

A parameterized function $\hat{Q}: \mathcal{H} \times \mathcal{A} \rightarrow \mathbb{R}$, approximates the expected cumulative reward of taking action a in quantum state $|s\rangle$:

$$\hat{Q}(|s\rangle, a; \theta) \approx Q^*(|s\rangle, a) \quad (2)$$

where $\theta \in \mathbb{R}$ are learnable parameters.

8) Policy

The agent's policy $\pi: \mathcal{H} \rightarrow \mathcal{P}(\mathcal{A})$ maps quantum states to probability distributions over actions. This means that given a quantum state, the policy determines the likelihood of selecting each possible action.

Typically derived via softmax over \hat{Q} :

$$\pi(a|s) = \frac{e^{(\hat{Q}(|s\rangle, a; \theta)/\tau)}}{\sum_{a'} e^{(\hat{Q}(|s\rangle, a'; \theta)/\tau)}} \quad (3)$$

where $\tau > 0$ is a parameter controlling exploration.

9) Temporal Difference (TD) Learning Objective

The agent updates Q-values using TD learning. Minimizing the Bellman error [14, 15] in Quantum Deep Q-Learning (QDQL) ensures accurate Q-value estimation, improving reinforcement learning performance.

Given a transition $(|s\rangle, a, r, |s'\rangle)$, minimize the Bellman error:

$$L(\theta) = \mathbb{E}_{|s\rangle, a, r, |s'\rangle} [r + \gamma \max_{a'} \hat{Q}(|s'\rangle, a'; \theta^-) - Q(|s\rangle, a; \theta)]^2 \quad (4)$$

where $\gamma \in [0, 1]$ is the discount factor, and θ^- are parameters of a target network or previous iteration.

10) Parameter Update

Update θ by gradient descent:

$$\theta \leftarrow \theta - \alpha \nabla_{\theta} L(\theta),$$

where $\alpha > 0$ is the learning rate.

11) Quantum-Inspired State Encoding

Classical observations are encoded into quantum state vectors $|s\rangle$ via feature maps $\Phi: X \rightarrow \mathcal{H}$, enabling representation of uncertainty and correlations as quantum superpositions.

12) Measurement and Action Selection

Action selection corresponds to a quantum measurement collapsing $|s\rangle$ to an action outcome with probability $\pi(a|s)$.

3. Results

As mentioned in the introduction, this study aims to create a framework for simulating quantum computing principles, specifically the application of Deep Reinforcement Learning (Q-Learning) by a so-called “Quantum Intelligent Agent” system to medical decisions. And so we evoked about optimizing the treatment of Sickle Cell Disease.

3.1 Problem Formulation

As long as the patient's condition is known to be sickle cell disease, medical staff can quickly conduct a clinical examination. For most patients, the management of sickle cell disease revolves around prevention of complications and regular medical monitoring.

Sickle cell disease manifests itself by various signs and each person may experience symptoms differently. Symptoms and complications as stated in [16], [17] can be classified in 5 groups of manifestation of the disease: Sickle cell anemia; Vaso-occlusive complications; Frequent episodes of acute pain; Significant proteinuria; Low oxygen saturations or Hypoxemia.

Furthermore the different categories of treatment of sickle cell diseases are the following: “Medicine to prevent the sickling of red blood cells, Medicine to reduce vaso-occlusive and pain crises, Medicine to reduce or prevent multiple complications, Medicine to treat pain, Medicine to reduce risk of infection, Transfusions and Potential genetic therapy treatments” [18].

According to [19], “there are no standard treatments that cure sickle cell disease. However, there are treatments that help people manage and live with the disease”.

3.2 Architecture of QIMed

The architecture in figure 1 shows a step-by-step modular view of QIMed system.

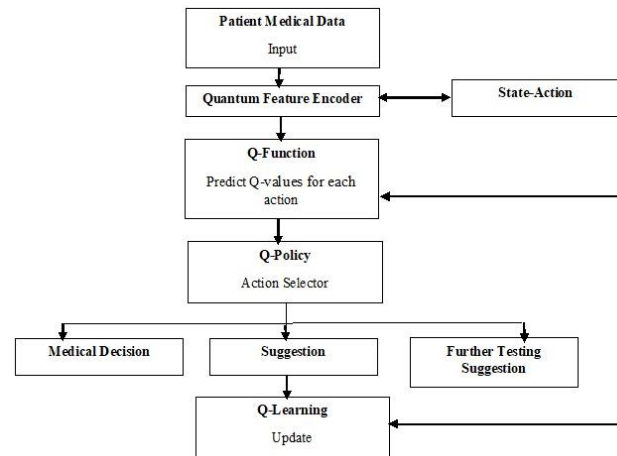


Fig. 1: Architecture of QIMed System

3.3 Implementation of System

The project's implementation follows essential steps that are closely linked to the programming environment namely the “Python compiler”. And of course, it is related to the proposed approach “Quantum Deep Q-Learning”, which is shown in Figure 1.

1) Datasets for QIMed

Main datasets and their features for a sickle cell disease medical decision are related to symptom, disease and treatment. They are structured as arrays.

symptom = ['anemia', 'painful episodes', 'painful swelling of hands and feet', 'frequent infections', 'pain', 'swelling', 'fever', 'fatigue', 'sickle cell pain crisis', 'acute inflammatory arthritis', 'septic arthritis', 'swelling', 'swelling', 'more frequent urination', 'shortness of breath', 'vomiting', 'rapid heart rate', 'coughing', 'wheezing', 'confusion', 'bluish color']

disease = ['Sickle cell anemia', 'Vaso-occlusive complications', 'Frequent episodes of acute pain', 'Significant proteinuria', 'Low oxygen saturations']

treatment = ['Penicillin V', 'Hydroxyurea', 'Overnight oxygen', 'Regular blood transfusions', 'Crizanlizumab', 'Opioid analgesics', 'L-glutamine', 'Anti-inflammatory drug', 'Nonsteroidal', 'ACE inhibitors', 'Dietary changes', 'Angiotensin-converting-enzyme', 'Overnight oxygen']

2) Data Processing

To perform the simulation, remember that we programmed it in Python and without using libraries such as: PennyLane, TensorFlow Quantum and PyQLearning.

DATA STRUCTURE REPRESENTATION:

Bit representation of state-action is used for efficient encoding of states and faster computation.

The first state-action is composed of 20 symptoms and 5 disease types under the environment_symptom_disease = [

```
# Sickle cell anemia (class 0)
([1,1,1,1,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0], 0),
# Vaso-occlusive complications (class 1)
([0,0,0,0,1,1,1,1,0,0,0,0,0,0,0,0,0,0,0,0], 1),
# Frequent episodes of acute pain (class 2)
([0,0,0,0,0,0,0,0,1,1,1,0,0,0,0,0,0,0,0,0], 2),
# Significant proteinuria (class 3)
([0,0,0,0,0,0,0,0,0,0,0,0,1,1,1,0,0,0,0,0], 3),
# Low oxygen saturations' (class 4)
([0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,1,1,1,1], 4)].
```

Thus for example, the Sickle cell anemia (class 0) is represented as ([1,1,1,1,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0], 0). Sickle cell anemia is manifested by the following symptoms: 'anemia', 'painful episodes', 'painful swelling of hands and feet', 'frequent infections', ...]. The bit 1 corresponds to the name of medicine in the list of treatments.

The second state-action is composed of 5 disease types and 13 medicines under the env_disease_treatment = [

```
# Sickle cell anemia
([1,1,1,1,0,0,0,0,0,0,0,0,0], 0),
# Vaso-occlusive complications
([1,0,0,0,1,1,0,0,0,0,0,0,0], 1),
# Frequent episodes of acute pain
([1,0,0,0,0,0,1,1,1,0,0,0,0], 2),
# Significant proteinuria
([1,0,0,0,0,0,0,0,0,1,1,1,0], 3),
# Low oxygen saturations
([1,0,0,0,0,0,0,0,0,0,0,0,1], 4)].
```

INPUT DATA

The simulation

```
test_input1 = [0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,1,1,1,1],
test_input2 = [1,1,1,1,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0],
test_input3 = [0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,1,0,1,1].
```

PARAMETERS

Parameters play a central role in shaping how the quantum agent learns and adapts. The parameters are adjusted during training to approximate the optimal Q-function, just like weights in a neural network.

The simulation works with the following parameters:

```
alpha = 0.5 # learning rate
gamma = 0.95 # discount factor
epsilon = 0.2 # exploration rate
episodes = 500
```

3.4 Output - Medical Decision

The printed outputs below show finding symptom lists, suggested diagnosis, similarity scores when uncertain and selected treatments.

Input value 1 (Finding symptoms): coughing, wheezing, confusion, bluish color
Suggested diagnosis: Low oxygen saturations
Selected treatments: Penicillin V, Overnight oxygen

Input value 2 (Finding symptoms): anemia, painful episodes, painful swelling of hands and feet, frequent infections
Suggested diagnosis: Sickle cell anemia
Selected treatments: Penicillin V, Hydroxyurea, Overnight oxygen, Regular blood transfusions

Input value 3 (Finding symptoms): coughing, confusion, bluish color
Suggested diagnosis: Sickle cell anemia, Vaso-occlusive complications, Frequent episodes of acute pain, Significant proteinuria and Low oxygen saturations
Similarity 75% (Jaccard coefficient) with: Low oxygen saturations
Selected treatments: Penicillin V, Hydroxyurea, Overnight oxygen, Regular blood transfusions

TABLE I. Q-VALUES IN % FOR EACH STATE-ACTION PAIR

Input value	Class 1	Class 2	Class 3	Class 4	Class 5
test_input1	0%	0%	0%	0%	100%
test_input2	100%	0%	0%	0%	0%
test_input3	20%	20%	20%	20%	20%

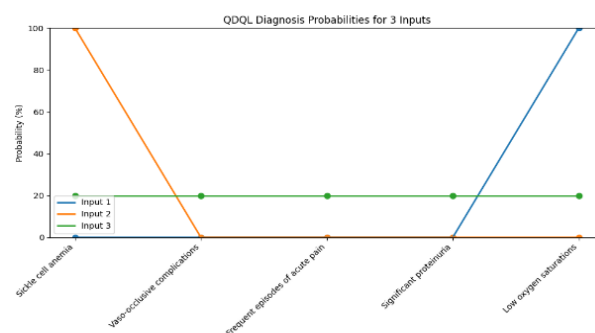


Fig. 2: QDQL diagnosis for the 3 inputs

4. Discussion

The result shows that the Quantum Deep Q-Learning agent has perfectly learned to associate symptom patterns with the correct sickle cell disease classes in the small test set, achieving 100% accuracy. The classification report confirms perfect precision, recall, and F1-score for all disease classes, indicating no misclassifications.

This high performance is expected given the small, well-defined dataset and the agent's training on the same patterns. The model uses softmax probabilities over Q-values to predict diseases, and when uncertain, it leverages Jaccard similarity to suggest likely diagnoses. Treatment suggestions correspond to the learned mappings from diseases to treatments.

Overall, the result validates that the agent can correctly predict disease classes from symptom inputs and suggest appropriate treatments, demonstrating the effectiveness of the Q-learning approach combined with similarity measures in this controlled scenario.

The contribution of our study lies in the methodology. Simulating Quantum Deep Q-Learning (QDQL) without a framework is ambitious but possible.

The approach proposed simulates quantum-inspired probabilistic decision-making in a classical Q-learning framework, improving diagnosis under uncertainty by combining learned Q-values and similarity metrics.

Tackling Deep Reinforcement Learning (DRL) challenges using Quantum Deep Learning (QDL) is an emerging area of research that leverages the principles of quantum computation to enhance traditional machine learning techniques.

References

- [1] M. M. Morovati, F. Tambon, M. Taraghi, A. Nikanjam, F. Khomh, "Common challenges of deep reinforcement learning applications development: an empirical study, Empirical Software Engineering", vol. 29, Issue 4, June 2024. Doi: <https://doi.org/10.1007/s10664-024-10500-5>
- [2] D. M. B. Simão, R. P. P. Almeida, A. Abrantes, "Challenges in Implementing Clinical DRL's in a Radiology Department - A Review", ECR 2024. Doi: <https://dx.doi.org/10.26044/ecr2024/C-20961>
- [3] A. A. Piovesana, An Introduction to Quantum Reinforcement Learning, Monography, Campinas, 2022.
- [4] Y. Cao, G. G. Guerreschi, A. Aspuru-Guzik, "Quantum Neuron: an elementary building block for machine learning on quantum computers", arXiv:1711.11240v1, November 2017.
- [5] P. R. Griffin, M. Boguslavsky, J. Huang, R. J. Kauffman, B. R. Tan, "Quantum computing: Computational excellence for Society 5.0", Data Science and Innovations for Intelligent Systems: Computational Excellence and Society 5.0. pp. 1-32, 2021. Available: https://ink.library.smu.edu.sg/sis_research/6851
- [6] S. El Hoss, W. El Nemer, D. C. Rees, "Precision Medicine and Sickle Cell Disease", Hemasphere, vol 6(9), 2022.
- [7] D. Niraula, J. Jamaluddin, M. M. Matuszak, et al. "Quantum deep reinforcement learning for clinical decision support in oncology: application to adaptive radiotherapy", Sci Rep 11, 23545, 2021. Available: <https://doi.org/10.1038/s41598-021-02910-y>
- [8] J. C. L. Chow, "Quantum Computing and Machine Learning in Medical Decision-Making: A Comprehensive Review", Algorithms 2025, vol. 18(3), pp. 156. Available: <https://doi.org/10.3390/a18030156>
- [9] J. Berberich, D. Fink, "Quantum computing through the lens of control: A Tutorial Introduction", arXiv:2310.12571v2, December 2024. Available: <https://arxiv.org/pdf/2310.12571>
- [10] R. B. Griffiths, "Hilbert Space Quantum Mechanics", Course Notes at the Physics Department of Carnegie-Mellon University, 2014. Available: <https://quantum.phys.cmu.edu/QCQI/qitd14.pdf>
- [11] H. Hinrichsen, "Introduction to Quantum Information Theory", 3rd Workshop on Quantum Information and Thermodynamics, September 2017. Available: <https://home.kias.re.kr/MKG/upload/The3rdQIT/Haye.pdf>
- [12] V. Wu, S. Jin, D. Wen, D. Han, X. Wang, "Quantum reinforcement learning in continuous action space", Quantum 9, p. 1660, Doi: <https://doi.org/10.22331/q-2025-03-12-1660>
- [13] .Z. Perić, A. Marković, N. Kontrec, J. Nikolić, M. D. Petković, A. Jovanović, "Two Interval Upper-Bound Q-Function Approximations with Applications", Mathematics, vol. 10(19), 3590, 2022. Doi: <https://doi.org/10.3390/math10193590>
- [14] H. Li, Z. Zhang, W. Luo, C. Han, Y. Hu, T. Guo, S. Liao, "Towards Optimal Adversarial Robust Q-learning with Bellman Infinity-error", ICML, 2024. Doi: <https://doi.org/10.48550/arXiv.2402.02165>
- [15] H. Y. Chen, Y. J. Chang, S. W. Liao, C. R. Chang, "Deep-Q Learning with Hybrid Quantum Neural Network on Solving Maze Problems", arXiv:2304.10159v3, December 2023. Available: <https://arxiv.org/pdf/2304.10159>
- [16] Johns Hopkins Medicine, "Sickle Cell Disease". Johns Hopkins Medicine, 2024. Available: <https://www.hopkinsmedicine.org/health/conditions-and-diseases/sickle-cell-disease>
- [17] S. El Hoss, W. El Nemer, D. C. Rees, "Precision Medicine and Sickle Cell Disease", Hemasphere, vol. 6(9), 2022.
- [18] NHLBI. "Sickle Cell Disease – Treatment", The National Heart, Lung, and Blood Institute, Last updated on August 08, 2023. Available: <https://www.nhlbi.nih.gov/health/sickle-cell-disease/treatment>
- [19] Sickle Cell Disease, "American Society of Hematology", 2024. Available: <https://www.hematology.org/education/patients/anemia/sickle-cell-disease>