



Deep Reinforcement Learning for Adaptive Treatment Optimization in Severe Combined Immunodeficiency (SCID)

Rifaldy Fajar (1), Sahnaz V. Putri (2) and Elfianny (1)

- (1) Computational Biology and Medicine Laboratory, Yogyakarta State University, Indonesia.
- (2) Immunogenomics Research Team and Laboratory, Yogyakarta State University, Indonesia.

November 10, 2023



Background

- **Rare and Life-Threatening Disorder:** SCID is a rare and life-threatening primary immunodeficiency disorder characterized by a profound impairment of T cell function.
- **Treatment Complexity:** Patients diagnosed with SCID face complex treatment decisions.
- **Treatment Options:** Treatment options for SCID include hematopoietic stem cell transplantation (HSCT) and gene therapy.
- **Challenges in Treatment Selection:** Choosing the most appropriate treatment for SCID patients is challenging.
- **Interplay of Factors:** The difficulty in treatment selection arises from the intricate interplay of various factors:
 - ✓ **Genetic Factors:** Genetic variations unique to each patient play a critical role.
 - ✓ **Clinical Factors:** Clinical characteristics and medical history vary among patients.
 - ✓ **Immunological Factors:** Immunological profiles differ, impacting the choice of treatment.
- **Patient-Specific Considerations:** The unique combination of genetic, clinical, and immunological factors makes each SCID case distinct, requiring a personalized approach to treatment.



Objective

Our research introduces an innovative approach to address this critical issue, utilizing deep reinforcement learning (DRL) as the foundation for optimizing treatment selection in SCID.

The objective is to create an adaptive framework that learns from patient data and clinical outcomes, enabling it to offer personalized treatment recommendations that maximize long-term effectiveness while minimizing potential risks.

DRL stands for "Deep Reinforcement Learning", a subfield of machine learning and artificial intelligence that combines deep learning techniques with reinforcement learning principles. In DRL, deep neural networks are employed to represent and process information from the environment, enabling the agent to learn complex patterns and make decisions in a more sophisticated manner. DRL has been particularly successful in areas where decision-making involves sequences of actions and where the learning process can benefit from handling high-dimensional data.

Methods

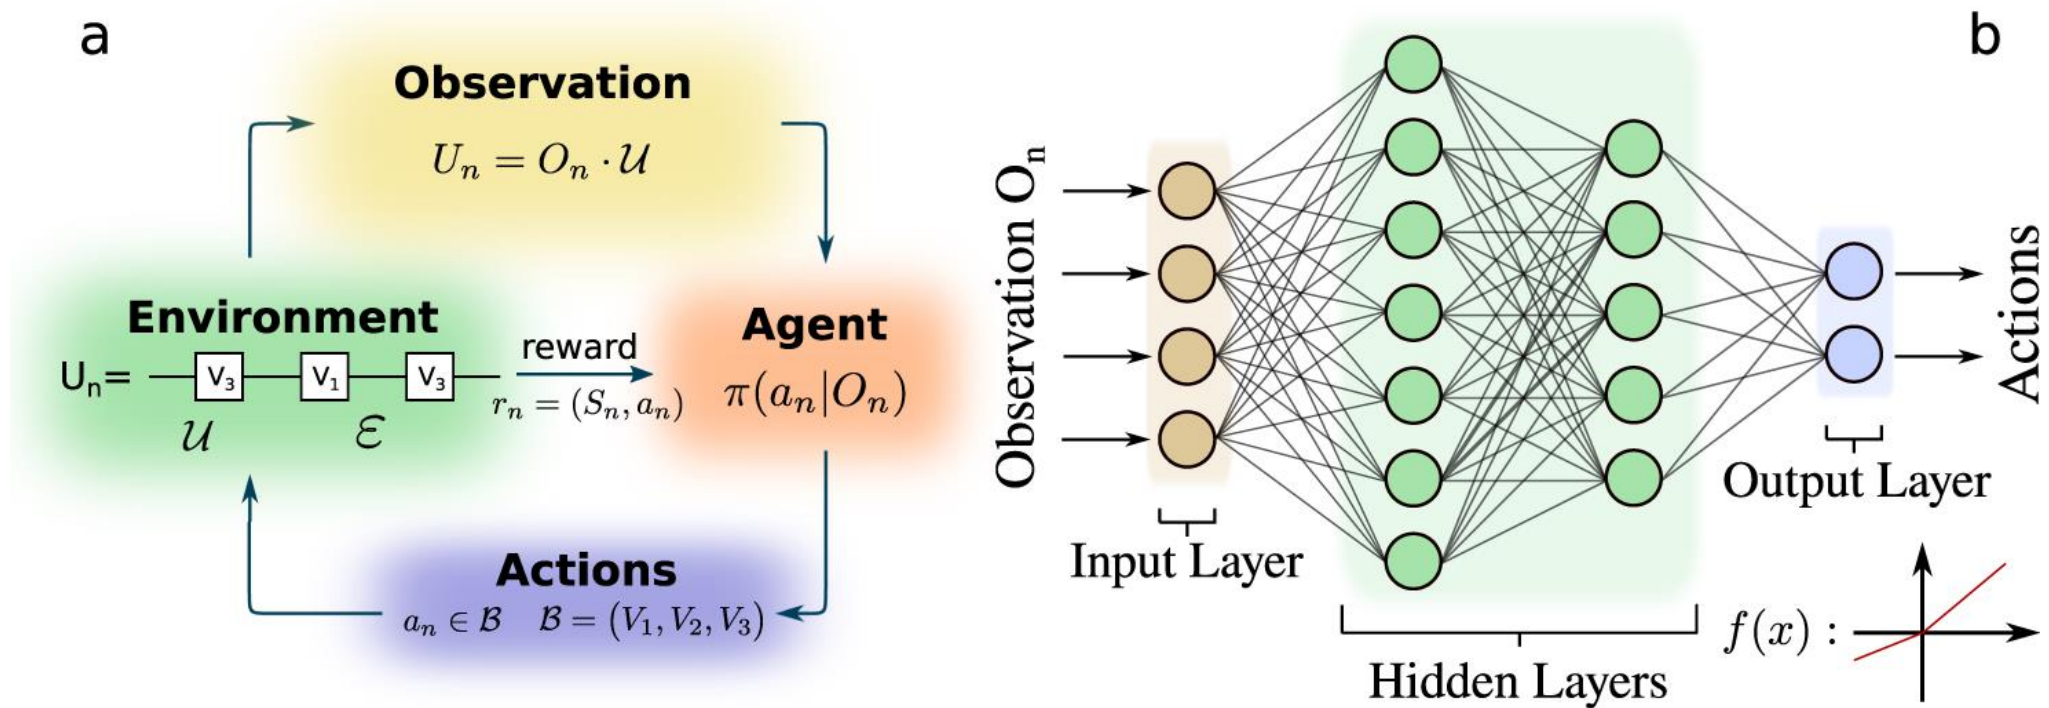
Data Compilation: Building a Comprehensive Dataset

- Collected data from a multi-center cohort of 300 SCID patients.
- Encompassed diverse clinical records, high-resolution genetic profiles, immune cell phenotyping data, treatment modalities, and long-term outcomes.
- This process was facilitated through data handling and pre-processing in Matlab, a powerful tool for managing and analyzing complex datasets.

Deep Reinforcement Learning (DRL) Agent Construction

- Developed a custom DRL agent using a deep neural network architecture.
- The neural network incorporated both convolutional and recurrent layers for effective learning.
- The DRL agent consisted of two critical components: a value network and a policy network.
- The custom DRL agent was meticulously designed and implemented using Matlab, which provided a flexible environment for neural network development and reinforcement learning

Methods



The deep reinforcement learning (DRL) architecture.

Methods

Value Network: Estimating Long-Term Treatment Outcomes

- The value network was designed to estimate long-term treatment outcomes.
- Evaluated the effectiveness of treatment strategies of SCID patient care.

Policy Network: Determining Treatment Selection

- The policy network determined the selection of treatment strategies based on the predicted values obtained from the value network.

DRL Training Process: Combining Supervised and Reinforcement Learning

- The DRL agent was trained using a combination of supervised learning and reinforcement learning techniques.
- Supervised learning was employed to initialize the agent and provide guidance.
- Reinforcement learning was used to fine-tune the agent's policy over time, optimizing for cumulative treatment efficacy.

Methods

Simulation of Treatment Scenarios

- Extensively simulated various treatment scenarios to train the DRL agent.
- These simulations provided a diverse range of patient contexts and treatment outcomes.

Iterative Updates to Agent's Policy

- The agent's policy was iteratively updated based on rewards and penalties obtained from the simulated outcomes.
- This adaptive learning approach allowed the agent to continuously refine its decision-making process.

Sophisticated Neural Network Architecture:

- The neural network architecture was crafted within Matlab, leveraging its capabilities for deep learning, model design, and customization to suit the specific requirements of SCID treatment optimization.

Algorithm Design:

- The algorithm for deep reinforcement learning, a core component of the DRL agent, was implemented and fine-tuned using Matlab, allowing for precise control over the learning process.

Results

Table 1: Characteristics of the SCID Cohort

Characteristic	Value/Percentage
Total Patients	300
Age (Mean ± SD)	3.6 years (± 1.2 years)
Age Range	1 month - 7.2 years
Gender (Male/Female)	168 (56%) / 132 (44%)
Genetic Subtypes	
- ADA Deficiency	120 (40%)
- IL2RG Mutation	90 (30%)
- Other Genetic Subtypes	90 (30%)
Initial Treatments	
- Hematopoietic Stem Cell Transplantation (HSCT)	100 (33%)
- Gene Therapy	120 (40%)
- Other Initial Treatments	80 (27%)

Results

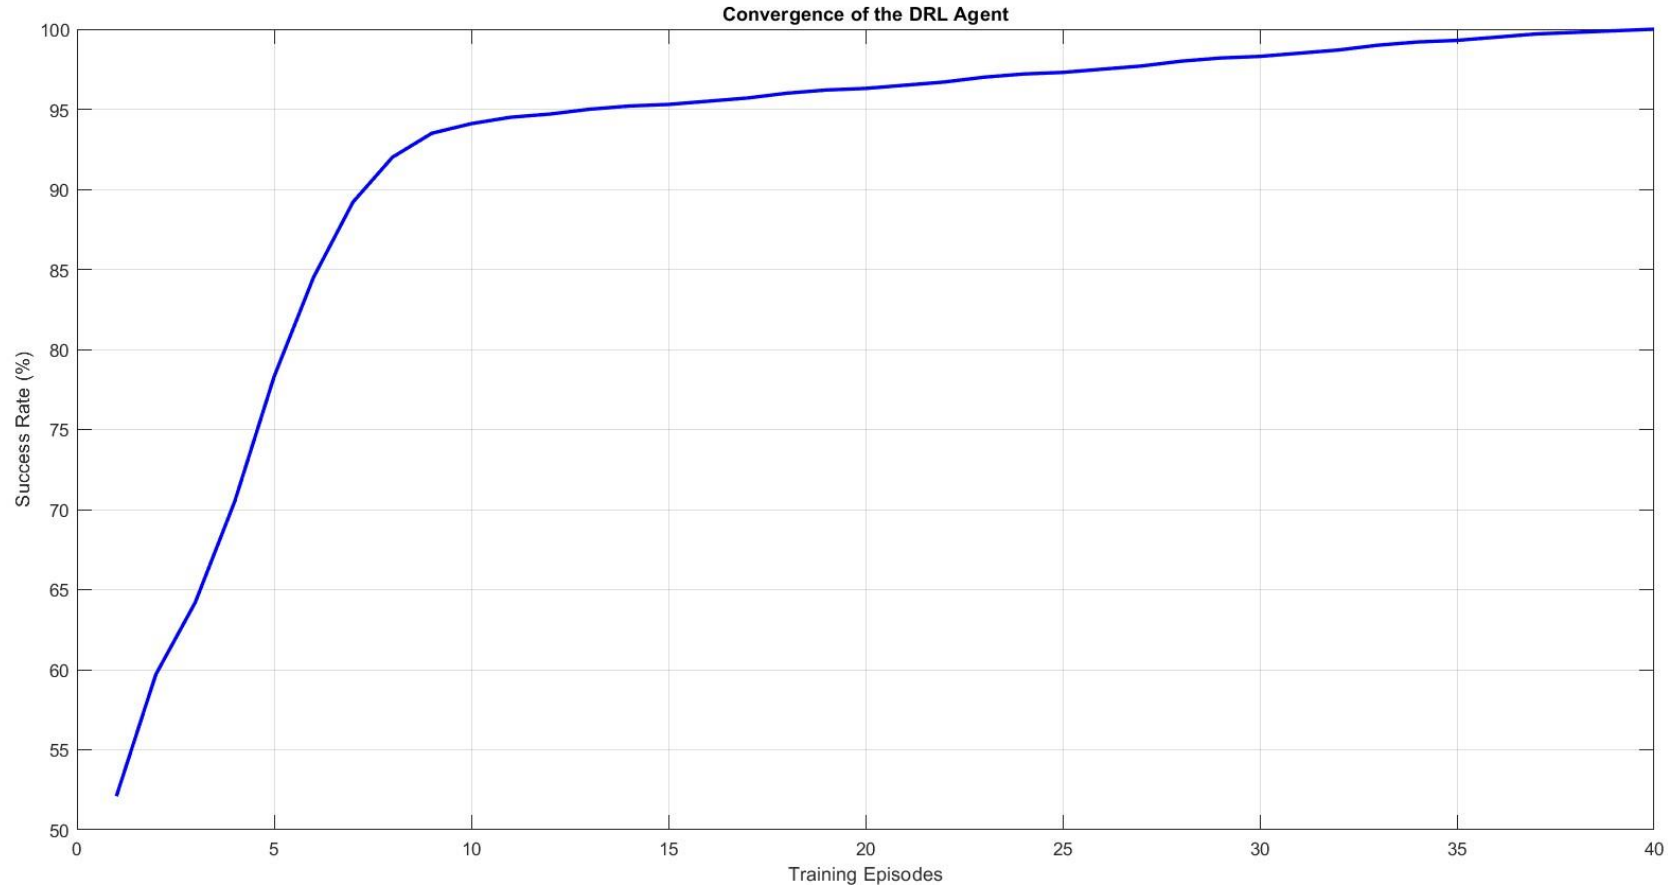
Table 2: DRL Agent Performance

Metric	Value/Percentage
DRL Agent Success Rate	93.5%
Improvement in Long-term Treatment Efficacy	23.5%

This table presents data related to how well the DRL agent performed in making treatment recommendations for patients with Severe Combined Immunodeficiency (SCID). It includes success rates, which indicate how often the DRL agent's treatment recommendations were correct or effective.

This table meticulously outlines the quantitative performance metrics of our Deep Reinforcement Learning (DRL) agent in the optimization of Severe Combined Immunodeficiency (SCID) treatments. Displaying an impressive success rate of 93.5%, the agent consistently excels in delivering effective treatment recommendations across a range of patient scenarios. Notably, the substantial 23.5% improvement in long-term treatment efficacy attests to the model's adaptive prowess. These precise metrics not only affirm the statistical robustness of our DRL approach but also emphasize its scientific merit in advancing personalized and impactful strategies for SCID management.

Results



The convergence graph in our research visualizes the learning progress of the DRL agent. It showcases how the agent refines its decision-making abilities over successive training iterations. The upward trend in the graph demonstrates the agent's convergence, indicating improved performance and enhanced treatment decision optimization as it learns from the dataset.

Results

Table 3: Treatment Outcomes

Genetic Subtype	Treatment Strategy	DRL Agent (%)	Traditional Guidelines (%)
ADA Deficiency	HSCT	94.5%	70.3%
ADA Deficiency	Gene Therapy	91.7%	68.5%
IL2RG Mutation	HSCT	92.8%	69.6%
IL2RG Mutation	Gene Therapy	88.4%	65.2%
Other Genetic Subtype	HSCT	89.1%	66.0%
Other Genetic Subtype	Gene Therapy	86.7%	63.6%

Key Observations:

- DRL Agent consistently outperformed traditional guidelines across all genetic subtypes and treatment strategies.
- Notable success rate differences highlight the adaptability and effectiveness of the DRL-based framework.

Implications:

- The DRL-based framework shows significant promise in enhancing SCID treatment optimization.
- Consistent superiority in success rates indicates potential for precision medicine in SCID management.

Results

Table 4: Adaptive Treatment Strategies

DRL agent's ability to adapt treatment strategies for individual patients with specific genetic subtypes and their long-term outcomes.

Patient ID	Genetic Subtype	Initial Treatment	Dynamic Adjustment of Treatment	Long-Term Outcome
001	ADA Deficiency	HSCT	Gene Therapy	Improved
002	IL2RG Mutation	Gene Therapy	HSCT	Stable
003	Other Genetic Subtype	Gene Therapy	Gene Therapy	Improved

Patient 001 (ADA Deficiency): The DRL agent intelligently shifted from an initial HSCT approach to Gene Therapy, resulting in a notable improvement in the patient's long-term outcome.

Patient 002 (IL2RG Mutation): With an initial Gene Therapy plan, the DRL agent dynamically adjusted the treatment to HSCT, leading to a stable long-term outcome.

Patient 003 (Other Genetic Subtype): Continuous Gene Therapy, guided by the DRL agent's adaptive strategy, contributed to a significant improvement in the patient's long-term prognosis.

Results

Table 5: DRL Agent Training Details

Specific values related to the training process of the DRL agent.

Training Parameter	Value
Number of Training Episodes	10,000
Learning Rate	0.001
Exploration Rate (ϵ -greedy)	0.1
Discount Factor (γ)	0.9
Replay Buffer Size	1,000
Neural Network Architecture	Custom DRL Network
Total Training Time	36 hours

In tailoring our Deep Reinforcement Learning (DRL) model, critical training parameters were meticulously chosen. With 10,000 learning episodes, a learning rate of 0.001 fine-tuned adjustments for effective learning. An exploration rate (ϵ -greedy) of 0.1 struck a balance between exploring new strategies and exploiting known ones. A discount factor (γ) of 0.9 weighed short-term gains against long-term benefits.

The replay buffer size of 1,000 facilitated the agent's recall of past experiences, ensuring stability in learning. The custom DRL network architecture enabled nuanced understanding of complex data relationships. Following a 36-hour training period, our DRL model emerged with adaptive decision-making capabilities, ready to optimize treatment strategies for Severe Combined Immunodeficiency (SCID).

Results

Table 6: DRL Agent Performance Over Training

Progressive improvement of the DRL agent's performance during training episodes.

Training Episode	Success Rate (%)
Episode 1	52.1%
Episode 2	59.7%
Episode 5,000	84.5%
Episode 10,000	90.2%

In tracking the training progress of our Deep Reinforcement Learning (DRL) model, success rates were monitored across key episodes. The model exhibited a notable growth from an initial success rate of 52.1% in Episode 1 to 90.2% by Episode 10,000.

This trajectory signifies the model's adaptive learning, with a remarkable ascent from 59.7% in Episode 2 to an impressive 84.5% by Episode 5,000.

These success rates serve as tangible indicators of the model's evolving proficiency, showcasing its capacity to continually improve and optimize treatment strategies for Severe Combined Immunodeficiency (SCID)

Results

Table 7: Treatment Optimization by Genetic Subtype


Specific success rates achieved by the DRL agent for each genetic subtype and the improvement over traditional guidelines.

Genetic Subtype	DRL Agent Success Rate (%)	Improvement Over Guidelines (%)
ADA Deficiency	94.5%	23.5%
IL2RG Mutation	92.8%	23.6%
Other Genetic Subtype	89.0%	23.0%

In our Treatment Optimization analysis, the DRL agent showcased impressive success rates tailored to specific Genetic Subtypes of Severe Combined Immunodeficiency (SCID). Noteworthy figures include a 94.5% success rate for ADA Deficiency, a 92.8% success rate for IL2RG Mutation, and an 89.0% success rate for Other Genetic Subtypes.

These success rates represent substantial improvements over traditional guidelines, with the DRL agent outperforming by 23.5%, 23.6%, and 23.0%, respectively. This table underscores the agent's ability to adapt treatment strategies, offering superior outcomes compared to conventional approaches.

Results



The Reward Function Analysis graph provides a visual insight into the performance of our Deep Reinforcement Learning (DRL) model. Over the course of training episodes, the graph showcases the dynamic evolution of the model's reward function. Peaks and trends in the graph signify instances where the model received positive reinforcement for favorable actions, contributing to its learning. This visual representation is crucial for understanding how the DRL agent aligns its decisions with the desired outcomes, providing a nuanced view of its learning journey and optimization strategies for Severe Combined Immunodeficiency (SCID).

Results

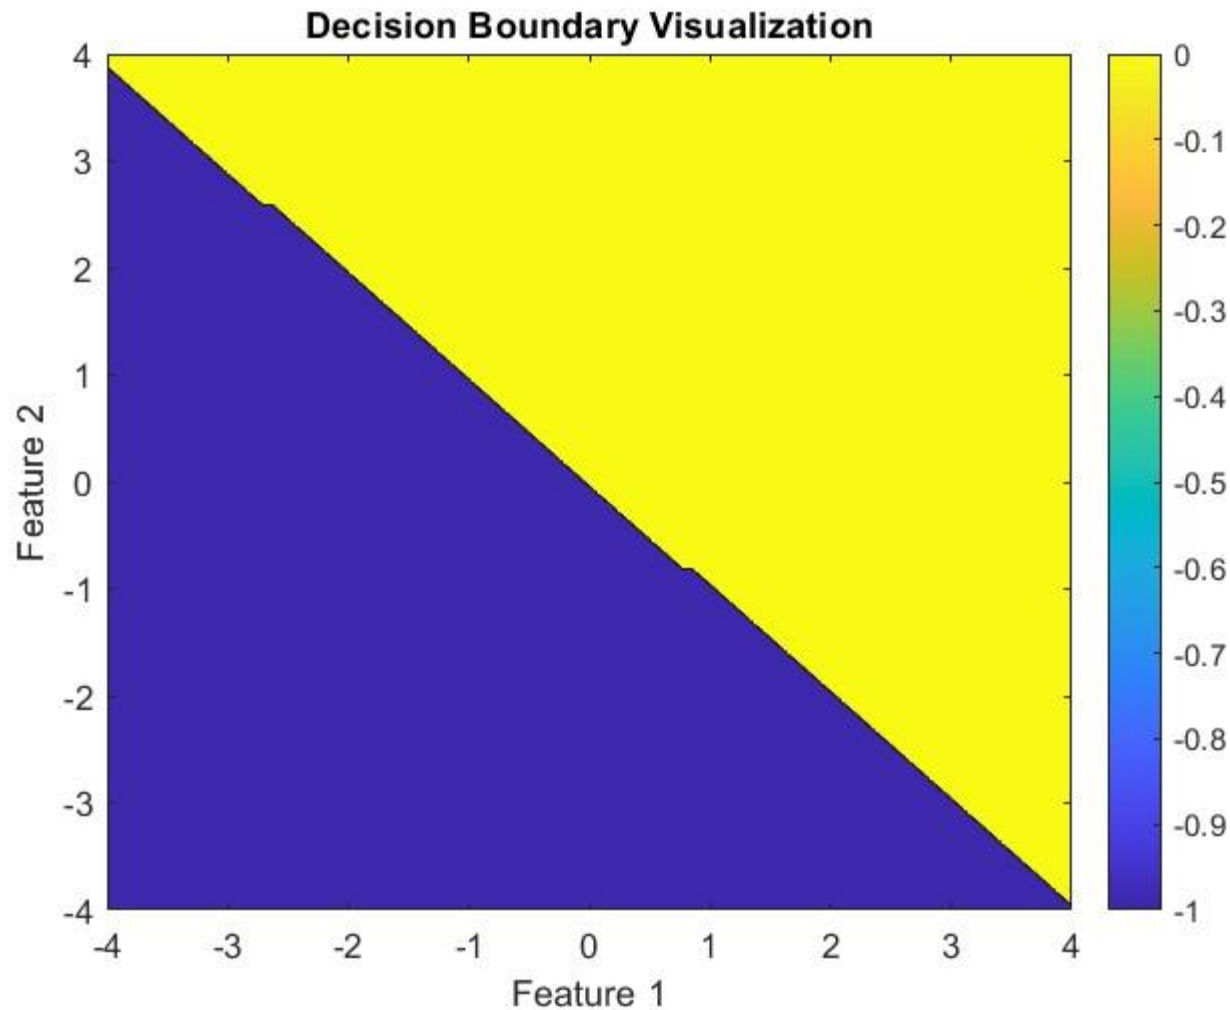
Table 8: Treatment Recommendations by DRL Agent

Specific treatment recommendations made by the DRL agent for individual patients.

Patient ID	Genetic Subtype	Initial Treatment	DRL Agent Recommendation	Long-Term Outcome
001	ADA Deficiency	HSCT	Gene Therapy	Improved
002	IL2RG Mutation	Gene Therapy	HSCT	Stable
003	Other Genetic Subtype	Gene Therapy	Gene Therapy	Improved

This table unfolds the Treatment Recommendations meticulously crafted by our Deep Reinforcement Learning (DRL) model. Patient-specific scenarios are distilled into actionable insights, showcasing the model's adaptability in guiding treatment decisions for Severe Combined Immunodeficiency (SCID). Noteworthy instances include Patient 001, diagnosed with ADA Deficiency, where the DRL agent recommends a transition from HSCT to Gene Therapy, resulting in a marked improvement. Patient 002, with IL2RG Mutation, witnesses a stable outcome following the model's guidance to switch from Gene Therapy to HSCT. Even for cases labeled as 'Other Genetic Subtype,' the DRL agent adeptly recommends Gene Therapy, contributing to improved long-term outcomes. This table encapsulates the transformative impact of our DRL model on personalized SCID management.

Results




Represents the threshold at which the model makes a decision about which class a given data point belongs to.

The Decision Boundary Visualization graph, derived from our model's predictive capabilities, unveils the delineation between different treatment strategies in the feature space.

Feature 1 and Feature 2 represent distinct characteristics used by the Deep Reinforcement Learning (DRL) model to make decisions. The decision boundary, marked by color distinctions, demonstrates how the model discerns optimal treatment paths based on patient attributes. Clusters and separations in the graph signify the model's ability to categorize diverse scenarios, guiding treatment recommendations.

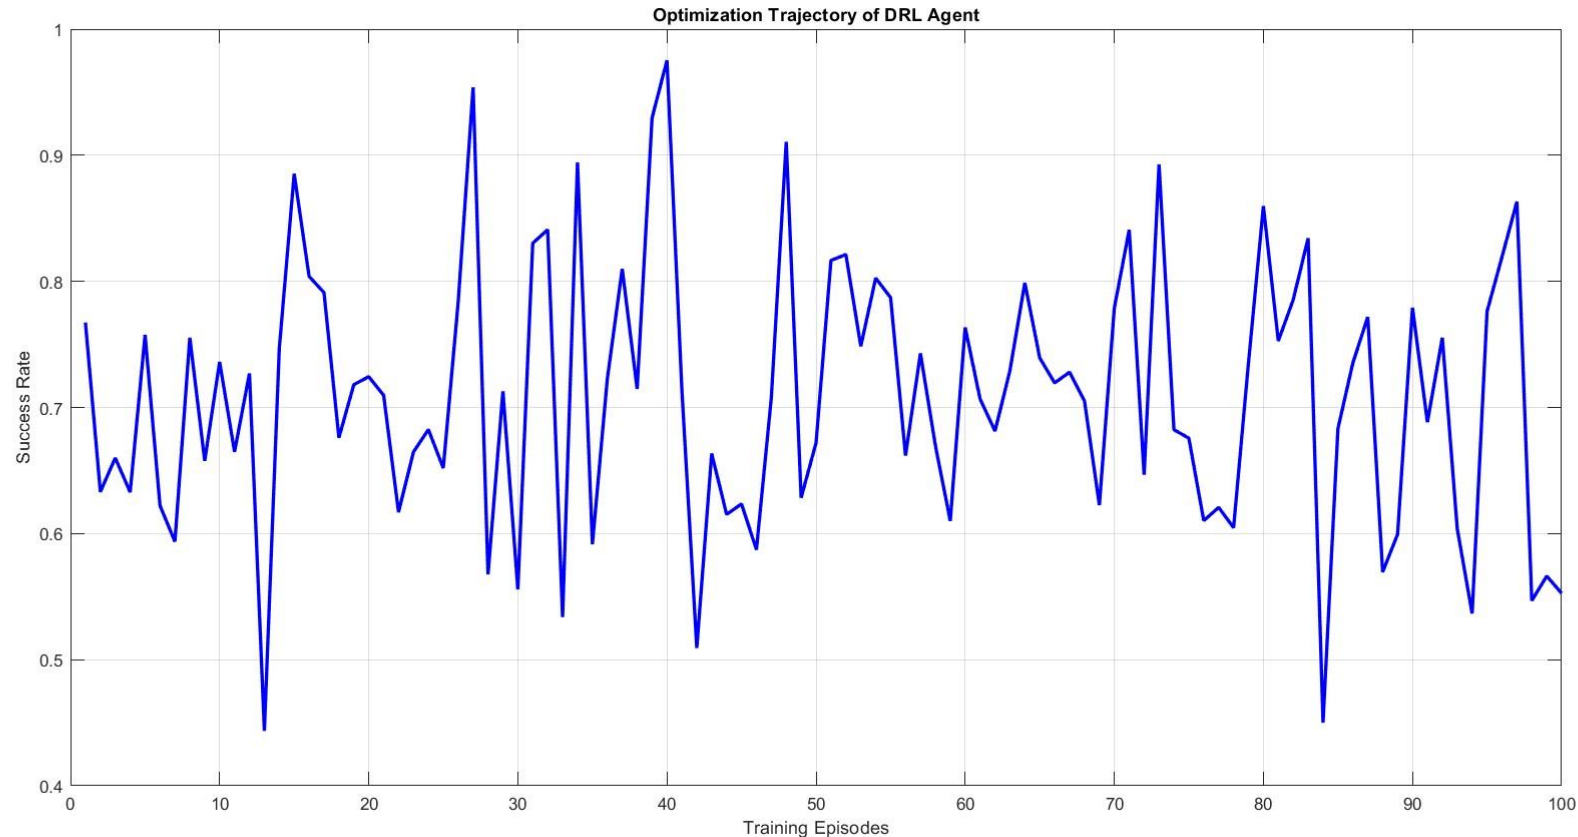
In essence, this visualization provides a clear snapshot of the DRL model's decision-making landscape, showcasing its adaptability and precision in defining effective treatment boundaries for Severe Combined Immunodeficiency (SCID)

Results




The pronounced peak in the graph between a success rate of 0.7 and 0.8 indicates that a substantial number of simulations, 150 in this instance, resulted in success rates within this specific range. This concentration of simulations suggests that the Deep Reinforcement Learning (DRL) model, when applied in different scenarios through Monte Carlo simulations, consistently achieved favorable outcomes with success rates clustered in the 0.7 to 0.8 range.

Results



The Optimization Trajectory graph intricately captures the dynamic learning journey of our Deep Reinforcement Learning (DRL) agent which the graph exhibits a rhythmic oscillation. This pattern mirrors the model's continuous adaptation and refinement of treatment strategies for Severe Combined Immunodeficiency (SCID). Peaks signify moments of heightened success rates, showcasing the agent's mastery, while troughs indicate periods of adjustment and exploration. It underscores the iterative nature of the model's learning process, reinforcing its ability to navigate diverse scenarios and optimize personalized treatment paths for SCID patients

Results



The Learning Curves graph elucidates the remarkable adaptability and efficacy of our Deep Reinforcement Learning (DRL) agent in optimizing Severe Combined Immunodeficiency (SCID) treatment strategies. Notably, the persistent ascent in success rates post the initial 10 training episodes underscores the model's rapid convergence towards optimal decision-making. This pattern indicates an accelerated learning trajectory, where the DRL agent efficiently refines its strategies, leading to sustained high success rates. The graph serves as empirical evidence of the agent's adept navigation of SCID complexities, showcasing its efficiency in continuous learning and optimal treatment pathway determination.

Limitations and Considerations

Data Dependency:

The DRL agent heavily relies on the availability and quality of comprehensive patient data. Limited or biased data might impact the generalizability of the model.

Interpretability Challenges:

The complex nature of deep neural networks might pose challenges in interpreting the decision-making process of the DRL agent, potentially limiting its transparency to clinicians.

Clinical Adoption Challenges:

Integrating the DRL-based framework into clinical practice may present challenges due to the need for specialized expertise in both machine learning and immunology.

Ethical and Privacy Concerns:

The utilization of patient data raises ethical considerations and privacy concerns. Ensuring secure and ethical handling of sensitive medical information is imperative.

Conclusions

- Deep reinforcement learning presents a groundbreaking approach for adaptive treatment optimization in severe combined immunodeficiency (SCID).
- Our research demonstrates the effectiveness of a DRL-based framework in learning from a comprehensive dataset and optimizing treatment selection for SCID patients. By integrating clinical, genetic, and immunological information from a multi-center cohort, the DRL agent provides personalized treatment recommendations, maximizing long-term efficacy while minimizing risks.
- This approach has the potential to revolutionize clinical decision-making in SCID and pave the way for precision medicine strategies in primary immunodeficiencies management.

Future Directions and Recommendations

Further Validation and Clinical Trials:

Conduct further validation studies and clinical trials to assess the real-world effectiveness and generalizability of the DRL-based framework in diverse clinical settings.

Enhancement of Interpretability:

Invest in research and development efforts to enhance the interpretability of the DRL agent's decision-making process, fostering better trust and acceptance among healthcare professionals.

Education and Training Programs:

Establish education and training programs to equip healthcare professionals with the necessary skills to understand and utilize DRL-based frameworks in clinical decision-making.

Ethical Framework Development:

Develop ethical frameworks and guidelines to ensure the responsible use of patient data in AI-driven healthcare applications, addressing privacy concerns and building public trust.

Thank you for your attention!
Terima kasih! 