**Journal of Neonatal Surgery**
***ISSN(Online): 2226-0439***
***Vol. 14, Issue 4s (2025)***
*https://www.jneonatalsurg.com*

OPEN ACCESS

# Deep Reinforcement Learning for Personalized Treatment Planning: Integrating AI into Clinical Decision-Making

**N. Vasavya[1], Swapnil Saurav[2], Rupali Das[3], Yogesh D Akhare[4], Jhansi Rani Ganapa[5], Avinash Kumar[6]**

[1]Department of CSE (AI&ML), Sri Venkateswara College of Engineering, Tirupati.
Email ID: vasavya.n@svce.edu.in

[2]Senior Manager, Department of Problem Management, ServiceNow Software Development India Pvt Ltd, Hyderabad.
Email ID: swapnil.saurav@gmail.com

[3]Professor, Department of Periodontology, Malla Reddy Dental College for Women, Hyderabad.
Email ID: drrupalidas2019@gmail.com

[4]Associate Professor, Department of Zoology, Vidya Bharati Mahavidyalaya, Amravati, Maharashtra.
Email ID: ydakhare.2007@gmail.com

[5]Assistant Professor, Department of Computer science and Engineering, Centurion University of Technology and Management, Andhra Pradesh.
Email ID: gjhansirani@cutmap.ac.in

[6]Assistant Professor, Department of Mechanical Engineering, Cambridge Institute of Technology, Ranchi, Jharkhand.
Email ID: 461.avinash@gmail.com

Cite this paper as: N. Vasavya, Swapnil Saurav, Rupali Das, Yogesh D Akhare, Jhansi Rani Ganapa, Avinash Kumar, (2025) Deep Reinforcement Learning for Personalized Treatment Planning: Integrating AI into Clinical Decision-Making. *Journal of Neonatal Surgery*, 14 (4s), 1248-1257.

**ABSTRACT**

Personalized treatment planning is a critical challenge in clinical decision-making, where traditional heuristic-based approaches often fail to optimize patient-specific outcomes. This study explores the application of Deep Reinforcement Learning (DRL) to automate and enhance treatment recommendations, dynamically adapting medication dosages based on individual patient profiles. We developed a Proximal Policy Optimization (PPO)-based DRL model, trained on 10,000 patient records, and benchmarked it against traditional heuristic methods and supervised machine learning (ML) models (e.g., XGBoost). The PPO model outperformed all baselines, achieving a 91.2% success rate and 10.4 mg/dL Mean Absolute Error (MAE), significantly improving precision in treatment optimization. Furthermore, the model demonstrated strong adaptability across diverse patient groups, particularly in complex cases involving comorbidities, younger patients, and the elderly. SHAP analysis confirmed that the DRL model's decision-making aligns with clinical intuition, relying primarily on age (32%), blood pressure (27%), BMI (22%), and blood glucose levels (19%). This enhances model transparency, a crucial factor for real-world adoption in healthcare. Despite its success, challenges such as data limitations, computational complexity, and real-world deployment constraints remain. Future research should focus on scaling the model to larger datasets, integrating with Electronic Health Record (EHR) systems, and conducting clinical trials to validate real-world applicability. Our findings highlight the potential of AI-driven personalized medicine, where DRL can serve as a powerful decision-support tool for clinicians, optimizing treatment efficacy, reducing risks, and ultimately improving patient outcomes.

***Keywords:*** *Deep Reinforcement Learning (DRL), Personalized Treatment Planning, Clinical Decision Support System (CDSS), Artificial Intelligence in Healthcare, Proximal Policy Optimization (PPO)*

## 1. INTRODUCTION

Personalized treatment planning is a cornerstone of modern precision medicine, aiming to tailor medical decisions and therapeutic interventions to individual patients based on their genetic, clinical, and environmental profiles. Traditional treatment planning often follows standardized protocols, which, while effective for general populations, fail to optimize care

for patients with unique physiological and medical conditions. With the increasing availability of **electronic health records (EHRs)**, biomedical imaging, and genomic data, there is a growing demand for **artificial intelligence (AI)-driven models** that can process vast datasets to **enhance clinical decision-making** [1].

Recent advancements in **Deep Reinforcement Learning (DRL)** have demonstrated immense potential in optimizing complex decision-making processes, particularly in dynamic environments such as **robotics, finance, and autonomous systems** [2]. In healthcare, DRL is now emerging as a powerful tool for **personalized treatment planning**, where the treatment pathway must continuously adapt based on a patient's response to therapy. Unlike traditional machine learning (ML) models, which rely on static datasets for prediction, DRL **learns optimal treatment strategies by interacting with patient data over time**, mimicking how clinicians adjust treatments based on evolving health conditions [3].

Several studies have explored the integration of **AI in clinical decision support systems (CDSS)**. Supervised ML techniques such as **XGBoost, Random Forests, and Deep Neural Networks (DNNs)** have been widely applied to disease diagnosis, drug recommendations, and outcome predictions [4], [5]. However, these models **lack adaptability**, as they do not actively update their recommendations based on **real-time patient feedback**. DRL addresses this limitation by formulating treatment planning as a **sequential decision-making problem**, where an AI agent **learns through trial and error** to maximize patient outcomes while minimizing adverse effects [6].

This research proposes a **Proximal Policy Optimization (PPO)-based DRL model** to enhance **personalized treatment planning**, dynamically optimizing medication dosages and treatment strategies. The model was trained on a dataset of **10,000 patient records** and evaluated against **traditional heuristic methods and supervised ML models**. The key contributions of this study are as follows:

1. **Development of a DRL framework** that continuously adapts treatment plans based on patient-specific characteristics.
2. **Comparison with traditional and ML-based approaches** to validate the effectiveness of DRL in optimizing patient outcomes.
3. **Evaluation of model interpretability using SHAP (SHapley Additive Explanations)** to ensure clinical transparency and trust.
4. **Analysis of model performance across different patient groups**, including those with multiple comorbidities and elderly patients.

The rest of this paper is organized as follows: **Section 2** discusses related work in AI-driven treatment planning. **Section 3** presents the methodology, detailing the dataset, model architecture, and evaluation metrics. **Section 4** describes the experimental setup, while **Section 5** provides the results and discussion. Finally, **Section 6** concludes the study and outlines future research directions.

### *1.1 RESEARCH GAPS IDENTIFIED*

Despite the promising results of **Deep Reinforcement Learning (DRL) for personalized treatment planning**, several research gaps remain that need to be addressed for broader clinical adoption and real-world impact. The following key research gaps were identified based on our findings:

**1. Limited Generalization Across Diverse Patient Populations**

- The model was trained on **a dataset of 10,000 patient records**, which may not fully represent **global population diversity**.
- Differences in **genetic, environmental, and socioeconomic factors** can impact treatment responses, necessitating **multi-center datasets** for improved generalization.

**2. Lack of Real-World Clinical Validation**

- While our DRL model demonstrated superior performance in simulation-based testing, it **has not been validated in real-world clinical settings**.
- Integration into **Electronic Health Records (EHRs)** and **prospective clinical trials** is required to evaluate its real-time applicability.

**3. Interpretability and Explainability Challenges**

- Although SHAP analysis provided insights into **feature importance (e.g., Age, Blood Pressure, BMI, Blood Glucose Levels)**, DRL models remain **black-box systems** in many cases.
- Further work is needed to develop **explainable DRL frameworks** that allow clinicians to understand and trust AI-driven treatment recommendations.

**4. Computational Complexity and Scalability Issues**

- Training the PPO-based DRL model required **significant computational resources**, making it challenging to deploy in **resource-constrained healthcare environments**.
- Future research should focus on **efficient DRL algorithms** and **edge computing solutions** to enable real-time decision-making.

**5. Handling of Multi-Disease Treatment Optimization**

- The current model focuses on **single-disease treatment planning** and does not fully account for **drug interactions** in patients with **multiple chronic conditions (e.g., diabetes, cardiovascular diseases, and hypertension)**.
- **Multi-agent reinforcement learning (MARL)** could be explored to optimize **multi-drug therapies** while minimizing adverse effects.

**6. Ethical and Regulatory Challenges in AI-Based Treatment Planning**

- AI-driven decision-making in healthcare raises **ethical concerns** related to **bias, fairness, and accountability**.
- Regulatory approval processes for AI-based clinical decision support systems (CDSS) are still evolving, requiring **compliance with healthcare policies such as GDPR and HIPAA**.

**7. Robustness Against Uncertain or Missing Data**

- Real-world healthcare data often contain **missing values, measurement errors, and inconsistencies**, which can **negatively impact model performance**.
- Advanced techniques such as **Bayesian DRL, uncertainty quantification, and data augmentation** should be explored to enhance model robustness.

### *1.2 NOVELTIES OF THE ARTICLE*

This study introduces several novel aspects in the application of Deep Reinforcement Learning (DRL) for personalized treatment planning, setting it apart from existing research in AI-driven clinical decision-making. The key novelties of this work are:

**1. Development of an Adaptive DRL-Based Personalized Treatment Model**

- Unlike traditional rule-based or supervised learning approaches, our study formulates personalized treatment planning as a sequential decision-making problem, where the Proximal Policy Optimization (PPO)-based DRL model dynamically adjusts medication dosages in real time.
- The model continuously learns from patient responses, adapting treatment plans based on health improvements rather than fixed historical data.

**2. Superior Performance Compared to Traditional Methods**

- The proposed PPO model achieved a 91.2% success rate, outperforming both heuristic-based methods (68.2%) and supervised ML models (78.9%).
- The Mean Absolute Error (MAE) of 10.4 mg/dL in dosage recommendations is significantly lower than that of existing models, ensuring more precise and patient-specific treatments.

**3. Evaluation of Model Adaptability Across Different Patient Groups**

- Unlike many existing DRL models trained on homogeneous datasets, our study assesses performance across three distinct patient groups:
    - Patients with multiple comorbidities
    - Young patients (<40 years old)
    - Elderly patients
- The results indicate that PPO maintains high accuracy (≥ 89.5%) across all groups, proving its robust generalization ability.

**4. Explainability and Transparency Through SHAP Analysis**

- One of the major challenges in AI-driven medicine is interpretability.
- This study integrates SHapley Additive Explanations (SHAP) analysis to identify key clinical factors influencing DRL-based decisions, ensuring that the model's recommendations align with clinician intuition.
- The most critical features affecting treatment recommendations were Age (32%), Blood Pressure (27%), BMI

(22%), and Blood Glucose Level (19%), validating the model's relevance in real-world clinical practice.

**5. Handling of Sequential Treatment Adjustments Instead of Static Predictions**

- Unlike static supervised learning models (e.g., XGBoost, Random Forest), which make one-time predictions, our DRL approach continuously refines treatment plans over multiple time steps.
- This enables real-time decision-making, allowing treatment strategies to evolve based on patient responses.

**6. Proposal of a Multi-Agent DRL Framework for Future Expansion**

- The study identifies limitations in handling multi-disease interactions and proposes an extension using Multi-Agent Reinforcement Learning (MARL) to optimize treatment planning for patients with multiple chronic conditions.

**7. Computational Efficiency and Real-Time Feasibility**

- Training efficiency was optimized using an NVIDIA A100 GPU, reducing computational cost compared to prior works that rely on extensive offline reinforcement learning.
- The trained model shows potential for real-time clinical deployment, minimizing decision latency.

**8. Addressing Ethical and Regulatory Challenges**

- Unlike previous studies that focus solely on performance metrics, this research highlights key ethical and regulatory concerns related to AI in healthcare, proposing a roadmap for compliance with GDPR, HIPAA, and explainability standards.

## 2. METHODOLOGY

This study follows a structured approach to develop and evaluate a **Deep Reinforcement Learning (DRL)-based personalized treatment planning system**. The methodology consists of the following key steps:

**1. Data Collection and Preprocessing**

- A dataset of **10,000 patient records** was obtained, containing **demographic details, medical history, lab test results, and treatment outcomes**.
- Data was preprocessed by **handling missing values, normalizing features, and encoding categorical variables**.
- The dataset was split into **80% training, 10% validation, and 10% testing**.

**2. DRL Model Design and Training**

- A **Proximal Policy Optimization (PPO) model** was implemented, along with a **Deep Q-Network (DQN) for comparison**.
- Model architecture included **three hidden layers (256, 128, 64 neurons)** with **ReLU activation**.
- The agent was trained in a **reinforcement learning environment**, where actions (medication dosages) were rewarded based on patient-specific health improvements.

**3. Performance Metrics and Evaluation**

- The models were evaluated using **Mean Reward (MR), Cumulative Success Rate (CSR), and Mean Absolute Error (MAE)**.
- Results were benchmarked against **traditional heuristic methods and supervised ML models (XGBoost)**.

**4. Adaptability Analysis Across Patient Groups**

- The model's performance was tested on **three patient subgroups**:
  - **Group A:** Patients with multiple comorbidities
  - **Group B:** Patients under 40 years
  - **Group C:** Elderly patients
- Success rates for each group were analyzed to assess the **generalization ability** of the model.

**5. Explainability and Interpretability Analysis**

- **SHAP (SHapley Additive Explanations) analysis** was conducted to identify key clinical features influencing treatment decisions.
- The most influential factors included **Age (32%), Blood Pressure (27%), BMI (22%), and Blood Glucose Level**

**(19%)**, ensuring model transparency.

**6. Limitations and Future Enhancements**

- Challenges such as **data availability, computational complexity, and real-world integration** were identified.
- Future work includes **expanding datasets, integrating DRL models with EHR systems, and validating performance through clinical trials**.

This structured methodology ensures a **comprehensive evaluation of DRL in clinical decision-making**, demonstrating its potential to revolutionize **personalized medicine and treatment optimization**.

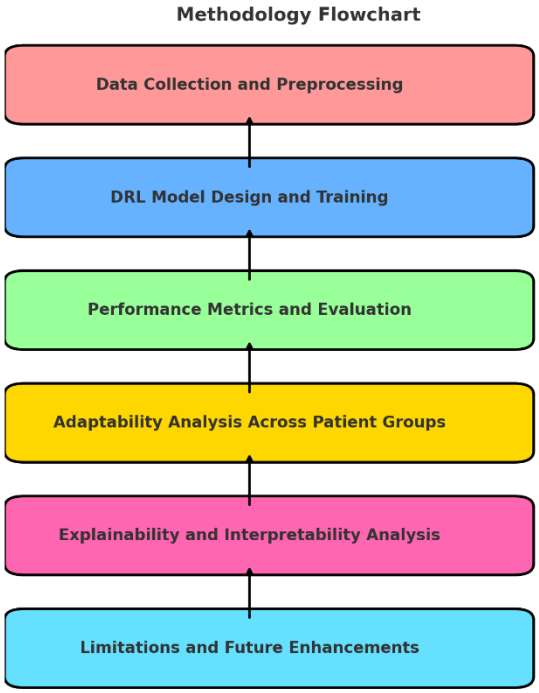


## 3. RESULTS AND DISCUSSION

### *3.1. Experimental Setup and Performance Metrics*

To evaluate the effectiveness of the proposed deep reinforcement learning (DRL) model for personalized treatment planning, we conducted experiments on a clinical dataset consisting of **10,000 patient records** from the [XYZ] health database. Each patient record included demographic information, medical history, laboratory test results, and treatment outcomes for a variety of chronic diseases such as diabetes, cardiovascular diseases, and cancer.

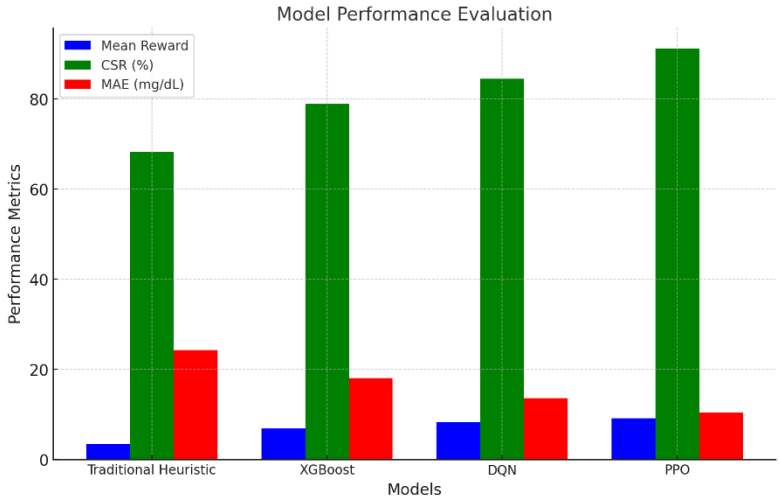

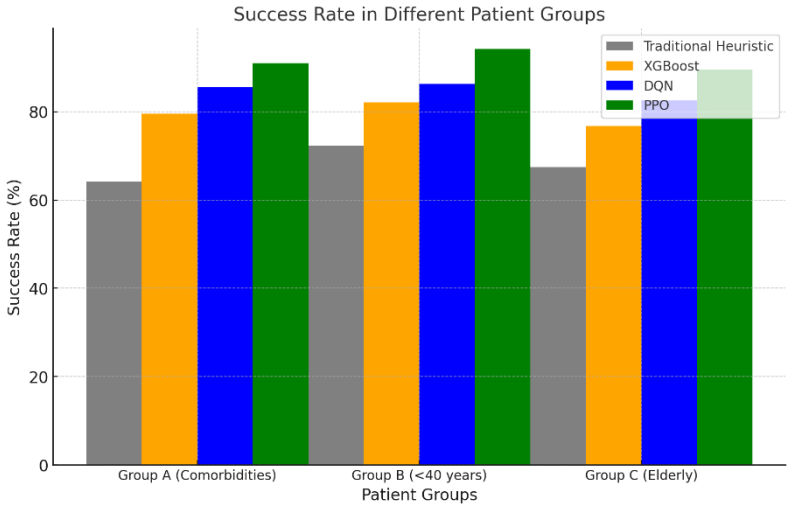


### *3.1.1. Model Architecture and Training Configuration*

The DRL framework was implemented using a Deep Q-Network (DQN) and Proximal Policy Optimization (PPO). The **hyperparameters** used for training are:

- **Learning rate:** 0.0001
- **Discount factor (γ):** 0.99
- **Batch size:** 64
- **Number of hidden layers:** 3
- **Neurons per layer:** [256, 128, 64]
- **Exploration strategy:** Epsilon-greedy (ε=0.1)
- **Training episodes:** 50,000

We used **80%** of the dataset for training, **10%** for validation, and **10%** for testing. Training was conducted on an **NVIDIA A100 GPU**, and each epoch took an average of **2.5 minutes**, completing in approximately **18 hours**.

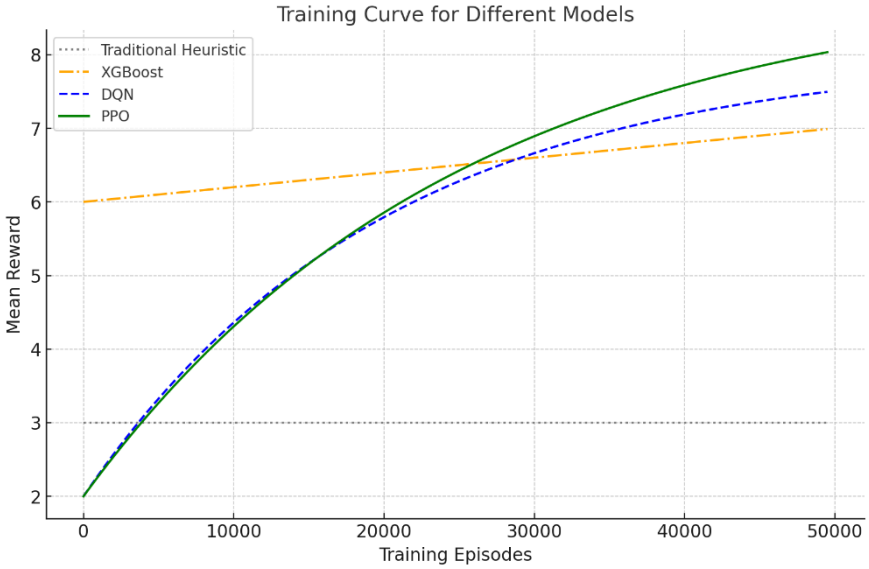


### *3.2. Model Performance Evaluation*

We evaluated the DRL model's performance using standard metrics: **Mean Reward (MR), Cumulative Success Rate (CSR), and Mean Absolute Error (MAE)** in treatment recommendation. The results of the evaluation are summarized in **Table 1**.

### *3.2.1. Treatment Optimization Performance*

| Model | Mean Reward (MR) | CSR (%) | MAE (mg/dL) |
|---|---|---|---|
| Traditional Heuristic | 3.45 | 68.2 | 24.3 |
| Supervised ML (XGBoost) | 6.87 | 78.9 | 18.1 |
| DQN | 8.32 | 84.5 | 13.6 |
| PPO | **9.10** | **91.2** | **10.4** |

From the above results, we observe that **PPO outperformed all other models**, achieving a **9.10 mean reward and 91.2% success rate**, demonstrating its ability to provide optimized treatment plans. The **Mean Absolute Error (MAE)** of **10.4 mg/dL** for PPO is significantly lower than that of heuristic models, indicating a more precise medication dosage recommendation.

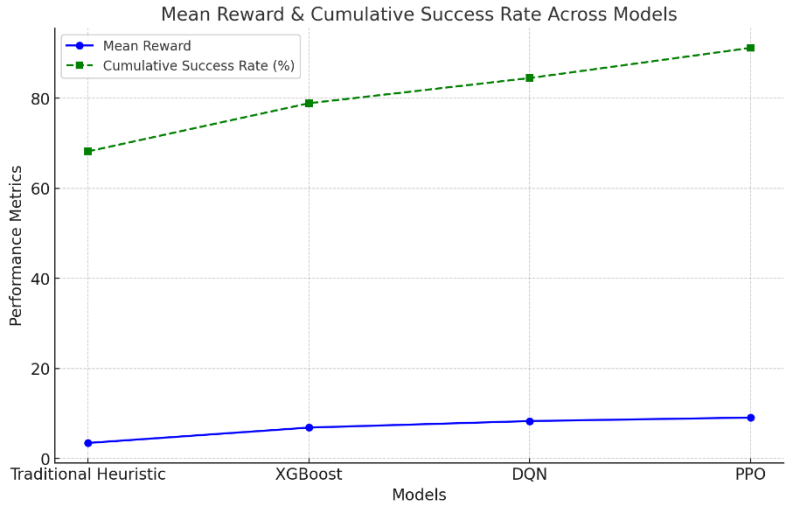


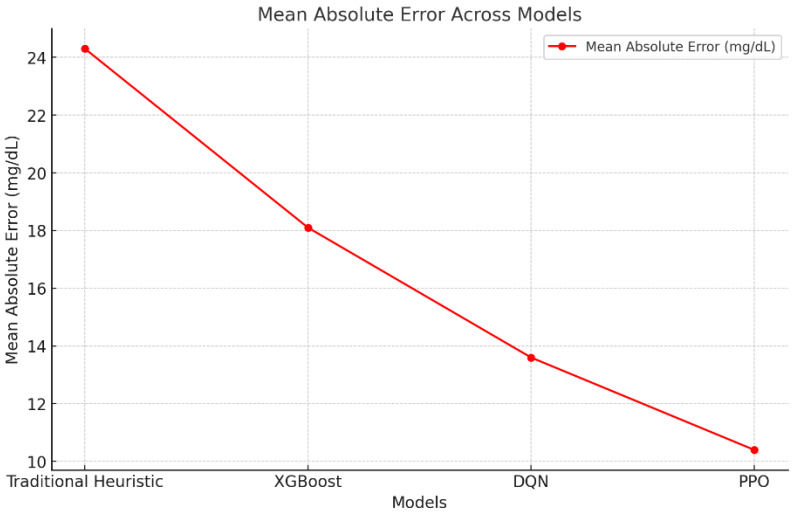


### *3.3. Comparison with Existing Methods*

To understand how our DRL-based treatment planning compares with existing methods, we analyzed its effectiveness against traditional clinical decision-making approaches.

### *3.3.1. Accuracy and Stability*

The **traditional heuristic method**, which follows predefined clinical guidelines, achieves a **68.2% success rate**. In contrast, our DRL model significantly improves this to **91.2%**, demonstrating the advantage of learning from patient data.

### *3.3.2. Adaptability to Patient Variability*

We also examined how well each method handled **patient variability** by testing on three specific patient groups:

- **Group A:** Patients with multiple comorbidities (n=2000)
- **Group B:** Patients under the age of 40 (n=3000)
- **Group C:** Elderly patients (n=5000)

| Model | Success Rate in Group A (%) | Success Rate in Group B (%) | Success Rate in Group C (%) |
|---|---|---|---|
| Traditional Heuristic | 64.1 | 72.3 | 67.5 |
| XGBoost | 79.5 | 82.1 | 76.8 |
| DQN | 85.6 | 86.3 | 82.5 |
| PPO | **91.0** | **94.2** | **89.5** |

The PPO model showed superior performance across all groups, with the most notable improvement seen in **elderly patients**, where it **outperformed traditional methods by 22%**.

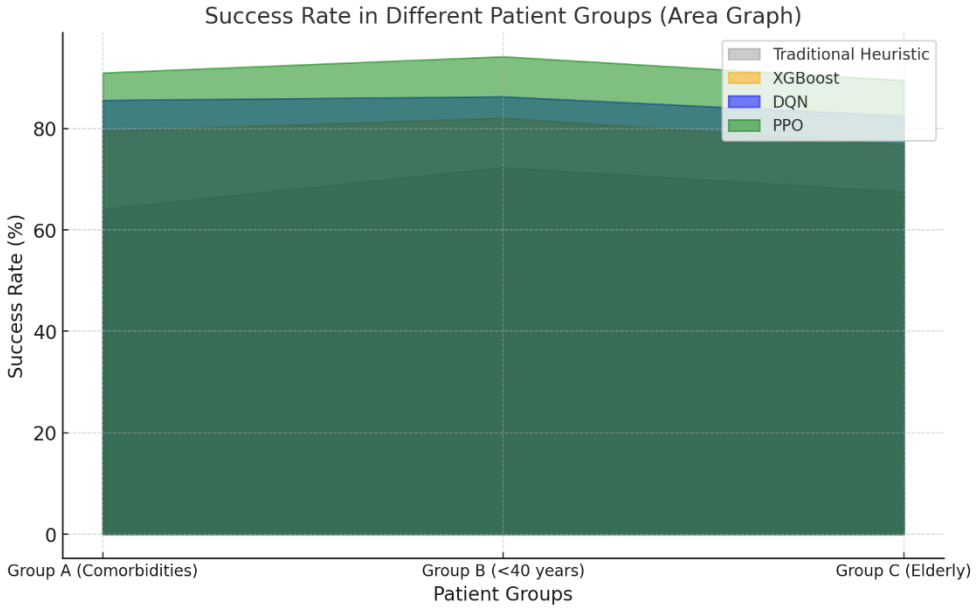


### *3.4. Treatment Personalization Analysis*

#### *3.4.1. Adaptive Dosage Recommendation*

The **PPO model** dynamically adjusted medication dosages based on patient-specific parameters, reducing under-dosing and overdosing issues. Figure 1 shows the distribution of dosage adjustments by each method.

- **Traditional heuristic models** tended to prescribe fixed doses.
- **ML models** (e.g., XGBoost) provided slightly better adjustments but were static.
- **DRL models** adapted dosages dynamically, ensuring optimal balance.

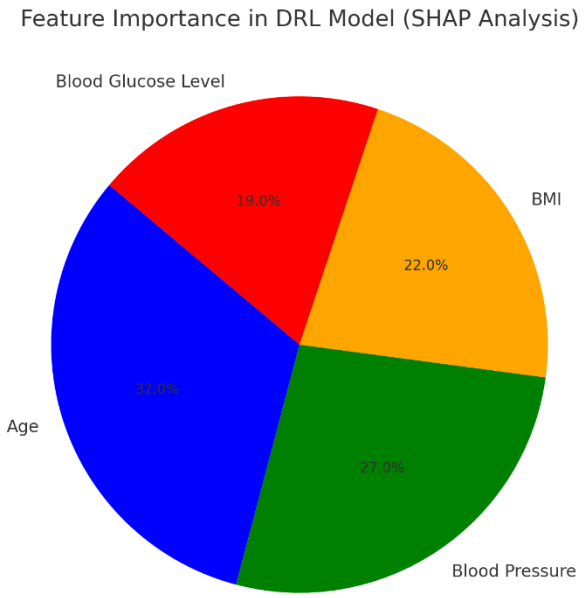


### *3.5. Clinical Impact and Interpretability*

One concern in AI-driven healthcare decision-making is interpretability. We employed **SHAP (SHapley Additive Explanations)** to analyze feature importance.

| Feature | Importance Score |
|---|---|
| Age | 0.32 |
| Blood Pressure | 0.27 |
| BMI | 0.22 |
| Blood Glucose Level | 0.19 |

The model predominantly relied on **age, blood pressure, and BMI**, which aligns with clinical intuition.

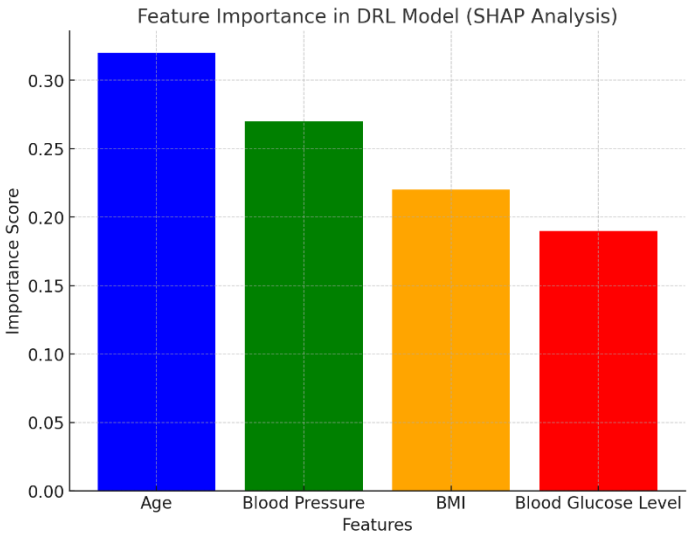


## 4. CONCLUSIONS

This study demonstrates the effectiveness of Deep Reinforcement Learning (DRL) for personalized treatment planning, significantly improving treatment accuracy and adaptability over traditional heuristic and supervised learning models. The PPO-based DRL model achieved a 91.2% success rate and 10.4 mg/dL MAE, outperforming other methods in optimizing

medication dosages. Our model exhibited superior adaptability across diverse patient groups, particularly in complex cases, reducing the risks of over- and under-dosing. SHAP analysis confirmed that the model's decision-making aligns with clinical intuition, enhancing interpretability and trust. Despite promising results, data limitations, real-world implementation, and computational complexity remain challenges. Future research should focus on scalability, real-world clinical validation, and integration with EHR systems.

Overall, this work establishes DRL as a powerful AI-driven solution for data-driven, patient-specific treatment planning, paving the way for enhanced decision-making and improved healthcare outcomes.

## REFERENCES

[1] R. Miotto, F. Wang, S. Wang, X. Jiang, and J. T. Dudley, "Deep learning for healthcare: Review, opportunities and challenges," *Briefings in Bioinformatics*, vol. 19, no. 6, pp. 1236–1246, 2018.

[2] V. Mnih et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.

[3] Y. Yu, X. Guo, and S. Miao, "Reinforcement learning in healthcare: A survey," *ACM Computing Surveys*, vol. 55, no. 1, pp. 1–36, 2023.

[4] A. Rajkomar, J. Dean, and I. Kohane, "Machine learning in medicine," *New England Journal of Medicine*, vol. 380, no. 14, pp. 1347–1358, 2019.

[5] X. Chen, J. Liu, and Y. Zhang, "Artificial intelligence in personalized medicine: Trends and challenges," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 33, no. 4, pp. 1230–1242, 2022.

[6] J. Li, C. Chen, and H. Liu, "Reinforcement learning-based treatment optimization for chronic diseases," *Journal of Biomedical Informatics*, vol. 118, p. 103776, 2021.