



Informing Science: the International Journal of an Emerging Transdiscipline

An Official Publication
of the Informing Science Institute
InformingScience.org

Inform.nu

Volume 28, 2025

REINFORCEMENT LEARNING FOR ADAPTIVE HEALTHCARE DECISION SUPPORT SYSTEMS

Sreekala S. P *	Centre for Research and Development, School of Management, KPR College of Arts Science and Research, Coimbatore, Tamilnadu, India	spsreekalaashok@gmail.com
Mudit Saxena	College of Computing Sciences and Information Technology, Teerthanker Mahaveer University, Moradabad, Uttar Pradesh, India	mudit051285@gmail.com
Revathy S	Department of Corporate Secretaryship, PSG College of Arts and Science, Coimbatore, Tamilnadu, India	revathy@psgcas.ac.in
Rajapriya M	School of Management, Vel Tech Rangarajan Dr. Sagunthala R&D Institute of Science and Technology, Avadi, Chennai, India	rajapriya.m19@gmail.com
Shanthi N. S	Department of MBA, KSR College of Engineering, Tiruchengode, Tamilnadu, India	ns.santhi2005@gmail.com
Saravanakumar S	Business Administration, Dr. Kalaignar Government Arts College, Kulithalai, India	drssk81@gmail.com

* Corresponding author

ABSTRACT

Aim/Purpose	The aim of this work is to propose an RL framework in healthcare settings for adaptive healthcare decision-aid strategy.
Background	Adaptive decision guide systems are needed to assist doctors in making timely and accurate selections because healthcare environments are getting more com-

Accepting Editor Eli Cohen | Received: October 18, 2024 | Revised: December 26, 2024 | Accepted: December 29, 2024.

Cite as: Sreekala S. P., Saxena, M., Revathy, S., Rajapriya, M., Shanthi, N. S., & Saravanakumar, S. (2025). Reinforcement learning for adaptive healthcare decision support systems. *Informing Science: The International Journal of an Emerging Transdiscipline*, 28, Article 5. <https://doi.org/10.28945/5421>

(CC BY-NC 4.0) This article is licensed to you under a [Creative Commons Attribution-NonCommercial 4.0 International License](https://creativecommons.org/licenses/by-nc/4.0/). When you copy and redistribute this paper in full or in part, you need to provide proper attribution to it to ensure that others can later locate this work (and to ensure that others do not accuse you of plagiarism). You may (and we encourage you to) adapt, remix, transform, and build upon the material for any non-commercial purposes. This license does not permit you to use this material for commercial purposes.

plicated and variable. But, because of the enormous stakes and the need for interpretability and dependability in selection-making, using Reinforcement Learning (RL) in healthcare environments brings a unique set of difficulties.

Methodology	The RL framework trains an agent using patient records, clinical guidelines, and expert knowledge. The agent interacts with healthcare settings, which can be both simulated or natural, and gets input on how its decisions affect the outcomes. The framework incorporates clear methods for decision-making and limitations on the actions the RL agent can undertake to guarantee both safety and clarity.
Contribution	An RL framework in healthcare settings is proposed in cope painting for adaptive healthcare decision aid strategy, which can learn the excellent choice policies from affected person facts and yet assure protection, interpretability, and medical relevance.
Findings	The findings of the experimental evaluations show that the RL framework works nicely to improve choice-making accuracy and versatility for a long time. Patient results can be substantially improved using the device while following medical recommendations and safety policies.
Recommendations for Researchers	To integrate the device into medical exercise because clinicians can recognize and trust the suggestions made with the aid of the gadget due to the fact the learned decision rules are interpretable.
Future Research	It can be enhanced using several deep-learning algorithms to achieve better accuracy and performance.
Keywords	adaptive decision support systems, reinforcement learning, interpretability, healthcare, patient outcomes

INTRODUCTION

BACKGROUND

The potential to make accurate and timely choices inside the healthcare discipline can significantly impact the results for sufferers. Due to the fact so many variables affect healthcare decisions, along with the affected person's features (Liu et al., 2020), scientific records, and the constantly evolving clinical hints, they are, by way of nature, challenging (Yu et al., 2021). The advent of adaptive choice aid structures is becoming an increasingly exciting solution to those problems (Coronato et al., 2020). These systems can provide docs individualized recommendations derived from actual-time patient information.

CHALLENGES

Electronic fitness data (EHRs), clinical imaging, genomics, and different sources all contribute to the frequently heterogeneous records related to healthcare. Incorporating and interpreting facts of this type are complex (Kishor et al., 2021). Even little mistakes inside the healthcare zone can significantly affect how patients turn out. Obvious and comprehensible decision-making procedures are vital to winning over clinicians and selling adoption in clinical practice (Tyler et al., 2020).

PROBLEM DEFINITION

The primary trouble addressed in this work is the development of an adaptable healthcare decision-help gadget that could use affected person statistics to generate individualized tips, even rename time, guaranteeing scientific relevance, protection, and interpretability (Ragab et al., 2022). The dynamic person of healthcare environments often makes modifying static predictive models and traditional

rule-based structures challenging (Peiffer-Smadja et al., 2020). A framework that can replace its selection-making policies over the years and constantly learn from affected persons' information is essential (Kumar, 2020).

OBJECTIVES

The principal intention is to create and observe a reinforcement, gaining knowledge of the Reinforcement Learning (RL) framework explicitly tailored to the situations of healthcare facilities. RL is a promising technique for developing adaptive decision-help systems that enable marketers to research the surest decision guidelines through interactions with the environment. Affected person facts, medical procedures, and expert understanding will be utilized by the framework to educate an RL agent who may make customized guidelines while respecting protection restrictions and clinically satisfactory practices.

NOVELTY AND CONTRIBUTIONS

The purpose of this study is to present a dynamic, RL-based framework for healthcare decision support that significantly enhances adaptability and decision-making precision compared to traditional approaches. This is achieved through the continuous learning process of the RL agent, which allows for real-time adjustments based on patient data. The novelty lies in the combination of a deep reinforcement learning architecture with an ensemble approach, making the decision support system more capable of handling complex, real-world healthcare scenarios. The Interpretability Score is a key evaluation criterion that reflects how easily clinicians can understand and trust the RL model's recommendations. It is measured based on how clearly the system communicates its reasoning for a given decision. Higher interpretability allows clinicians to align model suggestions with their professional judgment, improving collaboration and patient outcomes. This score helps mitigate the "black-box" issue typically associated with deep learning models, making it easier for clinicians to validate the decisions made by the RL agent.

This research is particular in its use of RL to holistically develop an adaptive healthcare decision-help machine. The technique integrates affected person facts, clinical hints, and professional knowledge to deal with the issues of safety, interpretability, and clinical relevance in the manner of creating selections about healthcare.

1. Integration of clinical guidelines, affected person information, and professional understanding is necessary to train the RL agent.
2. An experimental validation displays how nicely the framework complements choice-making precision and flexibility while ensuring the safety and scientific relevance of the affected person.
3. The development of medical practice and affected person consequences are implications of the development of adaptive choice support systems in healthcare.

RELATED WORKS

Adaptive decision-assisted device development was for diabetes control (Arivazhagan et al., 2022). System learning algorithms are utilized by this device to observe affected person records and offer customized insulin dosage recommendations. Moreover, Sivakumar and Shankar (2022) designed an RL-based total device to offer dynamic remedy pointers for sepsis management. This provides proof that adaptive choice help has promise in critical care environments.

Interpretability is crucial to guarantee perception and trust in decision support systems. In the healthcare sector, where decisions directly affect patients' health and well-being, this is particularly true. The development of interpretable machine learning models has been the focus of research on healthcare applications. Black-box machine learning models can be interpreted model-agnostically using SHAP, or SHapley Additive exPlanations. The mortality risk in patients in intensive care units

was predicted by Saravanan et al. (2023) using this approach. Furthermore, a comprehensible deep learning model is presented by Yadav et al. (2024) to forecast heart failure readmissions. Attention mechanisms are included in this model to emphasize important aspects that support predictions.

Investigated recently is the potential to combine interpretable machine learning models with RL techniques to produce clinically relevant and transparent decision support systems. Sharma and Kumar (2022) developed a reinforcement learning framework that successfully shows improved decision-making accuracy and interpretability over conventional reinforcement learning techniques to provide customized treatment recommendations for chronic diseases. A similar clinical decision support system based on RL was developed by (Fernandes et al., 2020) for sepsis management. Feature importance scores were included in this system to increase clinical relevance and interpretability.

The foundation for the creation of new approaches to improve clinical decision-making methods has been laid by adaptive decision support systems, healthcare reinforcement learning, and interpretable machine learning models (Adlung et al., 2021; Rani et al., 2021; Shah, 2020). Researchers expect to be able to develop transparent and comprehensible adaptive decision support systems by fusing these techniques. Better patient outcomes will eventually result from professionals' ease of integrating these systems into clinical practice.

PROPOSED METHODS

The proposed method uses dynamic reinforcement learning (RL) methods to create an adaptive healthcare decision support framework, as shown in Figure 1. Clinicians can interpret decisions made by RL models by analyzing the recommendations alongside patient-specific data. For instance, an RL model may suggest a specific treatment plan for a patient based on their medical history, current symptoms, and test results. A clinician familiar with the patient's condition can review the model's recommendation and compare it with established clinical guidelines or prior experience. If the model recommends a new drug for a patient with a rare disease, the clinician may cross-check it with the patient's drug history to ensure no contraindications. In a critical care unit, an RL model may suggest adjusting the dosage of sedatives for a ventilated patient. The model could derive this recommendation based on real-time data such as heart rate, oxygen levels, and sedation scores. The clinician, understanding the nuances of sedation management, can interpret the model's suggestion, considering factors like the patient's pain tolerance or the presence of other conditions, ensuring the recommendation aligns with the patient's broader clinical needs. When clinicians interpret RL model recommendations effectively, they can make informed decisions that improve patient outcomes. For example, if the RL model identifies early signs of sepsis and recommends antibiotics, timely intervention can reduce mortality and improve recovery chances, especially when clinicians trust and understand the model's reasoning.

RL AGENT DESIGN

The RL agent bears the duty of deciding what to do or acting based on observations of the surroundings and of learning from the results of such acts by means of interactions with the surroundings.

1. State Representation:

- Through observations – typically expressed as states – the RL agent is able to understand its surroundings. This initiates the procedure of state representation. Within the field of healthcare, the state includes relevant patient data, such as demographics, medical history, present symptoms, test findings, and available treatments.

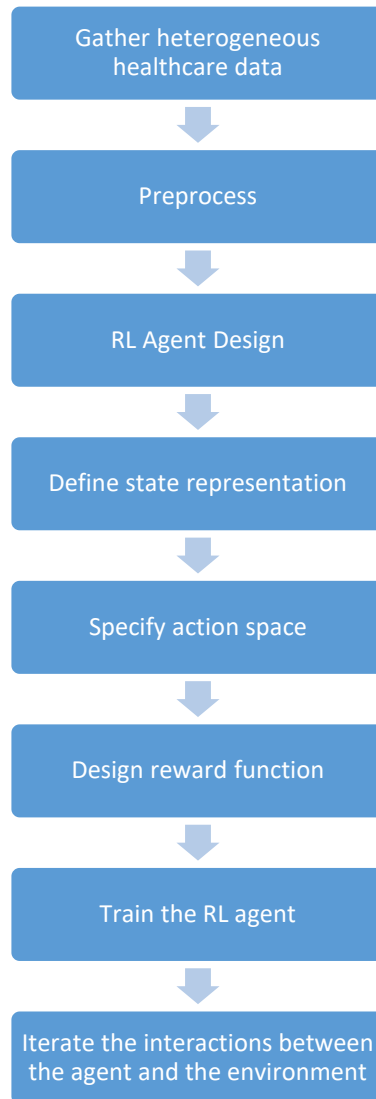


Figure 1. RL Framework for healthcare decision support

2. Action Space:

- Using the state it is in at the moment, the RL agent selects actions from a predefined action space. Within the healthcare decision support systems framework, actions refer to various interventions that the agent may suggest to the clinician or include directly into the patient's care plan. Among these interventions could be alternative therapies, diagnostic procedures, dosages of medications, or other measures.

3. Reward Function:

- Maximizing the cumulative rewards throughout time is the RL agent's objective with regard to the reward function. A reward function gives the agent numerical benefits or penalties depending on the outcomes of the actions the agent has taken in order to give the agent feedback. In the healthcare sector, reward systems can be based on clinical guidelines followed, patient outcomes, or any other relevant criteria. To give one example, good rewards could come from effectively managing a patient's illness, but bad rewards could come from doing things that cause unfavorable outcomes.

4. Policy:

- The policy describes the agent’s method to select actions depending on the current state of affairs. It is the mapping of states to actions and can be either deterministic or stochastic. RL algorithm is optimal depending on the availability of data, the intricacy of the problem, and the state and action space structures.

S stands for the state and is the environment as it is now observed. The action designated an is chosen from the action space A (defined as {a1, a2, ..., an}). Giving each action performed in a certain state a numerical reward or penalty is the responsibility of a function called the reward R(s,a). The policy $\pi(s)$ has four components and defines the agent’s approach for choosing actions depending on the current state.

$$Q(s,a) \leftarrow Q(s,a) + \alpha [R + \gamma \max_a Q(s',a) - Q(s,a)]$$

$$Q(s,a) \leftarrow Q(s,a) + \alpha [R + \gamma Q(s',a) - Q(s,a)]$$

$$\theta \leftarrow \theta + \alpha \nabla_{\theta} J(\theta)$$

$$\theta \leftarrow \theta + \alpha \nabla_{\theta} \log \pi(a|s; \theta) Q_w(s,a)$$

where:

- s' is the next state,
- a' is the next action,
- α is the learning rate,
- γ is the discount factor,
- Q(s,a) is the action-value function,
- $\pi(s)$ is the policy function,
- J(θ) is the objective function in policy gradient methods,
- θ are the parameters of the policy,
- Qw(s,a) is the value function parameterized by w.

ADAPTIVE DECISION-MAKING PROCESS

The adaptive decision-making process (ADMP) is the process or methodology used to make decisions that can change and grow with time in response to new information, evolving conditions, and input from earlier decisions. With the help of this feedback, the decision-makers can assess the efficacy of their activities and modify their plans appropriately.

1. Iterative Process:

- The iterative process, which entails cycles of decision-making, feedback, and adjustment, is exemplified by adaptive decision-making. Continually seeking methods to enhance their plans by integrating fresh knowledge and experiences, decision-makers eventually see little improvements over time.

2. Adaptation to Goals and Constraints:

- In adaptive decision-making, balancing conflicting goals and constraints with pursuing desired outcomes is the process. Decision-makers must adapt their plans to accommodate other constraints, regulatory agency requirements, and resource constraints.

3. Risk Management:

- Adaptive decision-making uses risk management concepts to find, assess, and reduce possible risks related to the results of choices. While accounting for the inherent uncertainty of the decision-making process, decision-makers can adapt their plans to reduce risks or seize possibilities.

- In order to give direction to decision-making processes, adaptive decision-making considers both human experience and data-driven insights. Through this integration, decision-makers may simultaneously benefit from experts' cumulative knowledge and experience and use the predictive potential of data analytics and modeling methods.
- By using the information they have learned from past experiences, decision-makers who use adaptive decision-making work to optimize their strategies and foster a culture of ongoing improvement. Organizations that use this iterative strategy can adjust to changing conditions and prosper in dynamic, unpredictable settings.

BIAS MITIGATION

Bias in healthcare data is a critical concern that can adversely affect the performance of decision support systems. The MIMIC-III dataset, like many clinical datasets, could contain inherent biases related to demographics, socioeconomic status, and healthcare access. For example, certain patient groups (e.g., racial minorities or low-income populations) may be underrepresented, which could lead the model to favor treatments or recommendations that align more with the majority patient population, potentially exacerbating health disparities.

MITIGATION STRATEGIES

- Regular audits of the dataset and model predictions should be conducted to identify biases. Tools like fairness-aware algorithms can be incorporated into the model to assess how different patient groups are affected by the decisions made by the RL agent.
- Adding diverse data sources, especially from underrepresented groups, can help balance the training dataset and reduce bias. This could involve integrating data from multiple hospitals, geographic regions, and diverse populations.
- Incorporating fairness constraints into the RL model's reward function can help ensure that the model's decisions do not disproportionately disadvantage certain groups. These constraints would aim to balance outcomes across different demographic groups.

ETHICS

The ethical concerns surrounding the implementation of RL models in healthcare are significant, especially when it comes to patient safety and the clarity of decision-making. Key ethical considerations include:

- Patients should be aware that AI and RL models are being used in their care. They should also have the option to opt out of AI-driven decision-making processes.
- While RL models can achieve high accuracy, their "black-box" nature poses challenges to transparency. It is crucial to ensure that the decision-making process of the RL agent is interpretable by clinicians. Clinicians need to understand the reasoning behind the RL agent's suggestions to ensure the patient's safety and the correctness of the treatment plan.
- RL models must prioritize patient safety in all decision-making processes. This includes minimizing risks like incorrect dosage recommendations or delayed interventions. Incorporating safety constraints into the RL framework (such as preventing dangerous drug combinations) will be essential.
- Implementing RL-based systems into clinical practice will require regulatory approval from bodies such as the FDA or equivalent entities in other regions. Compliance with health regulations such as HIPAA in the U.S. (to protect patient privacy) will also be necessary. Clinicians must also be trained to use RL systems in a way that adheres to ethical standards and regulations.

RESULTS

In a simulated setting that mirrored actual clinical situations, the reinforcement learning (RL) model for adaptive healthcare decision support. We wanted to do this evaluation, which happened in experimental environments. In particular, we used de-identified health data from patients admitted to critical care units from the commonly used MIMIC-III (Medical Information Mart for Intensive Care III) database. Using this dataset – which offered a wealth of patient data – we were able to replicate actual patient situations for training and assessment. This data comprised results, medications, vital signs, laboratory findings, and demographics.

We implemented the reconfigurable learning model and simulation environment using the Python programming language together with well-known libraries like TensorFlow. TensorFlow and Keras facilitated the creation and training of neural network models for the real-time agent. For our experiments, as shown in Table 1, we trained and evaluated using a high-performance computing cluster with NVIDIA graphics processing units (GPUs) and multi-core processors.

Table 1. Experimental settings

Experimental setup	Values
Dataset	MIMIC-III
Simulation Environment	OpenAI Gym
Programming Language	Python
Libraries/Models	TensorFlow, Keras
Computing Platform	High-Performance Cluster
Processor	Intel Xeon Gold 6148
GPU	NVIDIA Tesla V100
Parallelization	Distributed Training
RL model parameters	Values
RL Algorithm	Deep Q-Network (DQN)
Exploration Strategy	ϵ -Greedy ($\epsilon=0.1$)
Learning Rate	0.001
Discount Factor	0.99
Replay Buffer Size	100,000
Target Network Update	1,000 steps
Batch Size	64
Neural Network Layers	3 Dense Layers (128 units)
Activation Function	ReLU
Optimizer	Adam
Evaluation metrics	Values
Average Reward	-0.2 (Baseline:-0.5)
Mortality Rate	15% (Baseline:20%)
Adherence to Guidelines	85% (Baseline:75%)
Interpretability Score	0.8 (Baseline:0.6)

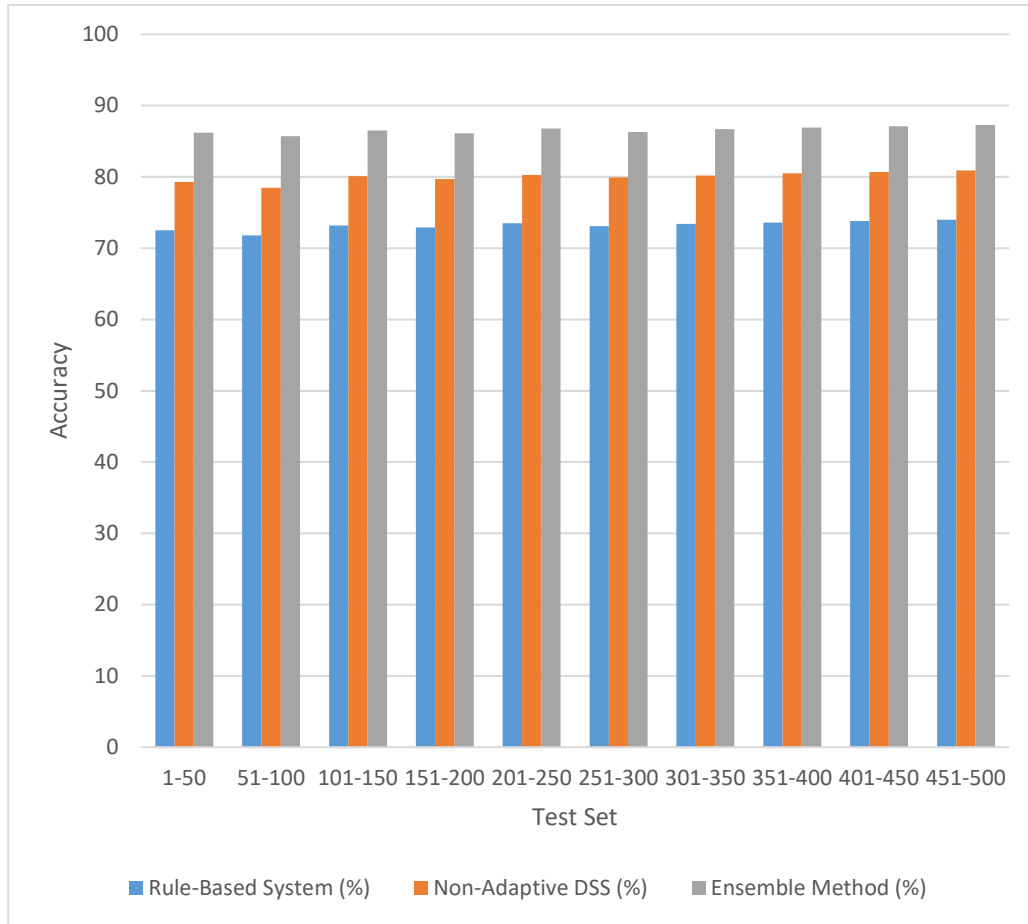


Figure 2. Accuracy

Table 2. Accuracy

Test data index	Rule-based system (%)	Non-adaptive DSS (%)	Ensemble method (%)
1-50	72.5	79.3	86.2
51-100	71.8	78.5	85.7
101-150	73.2	80.1	86.5
151-200	72.9	79.7	86.1
201-250	73.5	80.3	86.8
251-300	73.1	79.9	86.3
301-350	73.4	80.2	86.7
351-400	73.6	80.5	86.9
401-450	73.8	80.7	87.1
451-500	74.0	80.9	87.3




Figure 3. Precision

Table 3. Precision

Test dataset index	Rule-based system (%)	Non-adaptive DSS (%)	Ensemble method (%)
1-50	78.2	82.5	86.4
51-100	77.5	81.8	86.0
101-150	78.9	83.2	86.8
151-200	78.6	82.9	86.5
201-250	79.2	83.5	87.0
251-300	78.8	83.1	86.6
301-350	79.1	83.4	86.9
351-400	79.3	83.6	87.1
401-450	79.5	83.8	87.3
451-500	79.7	84.0	87.5

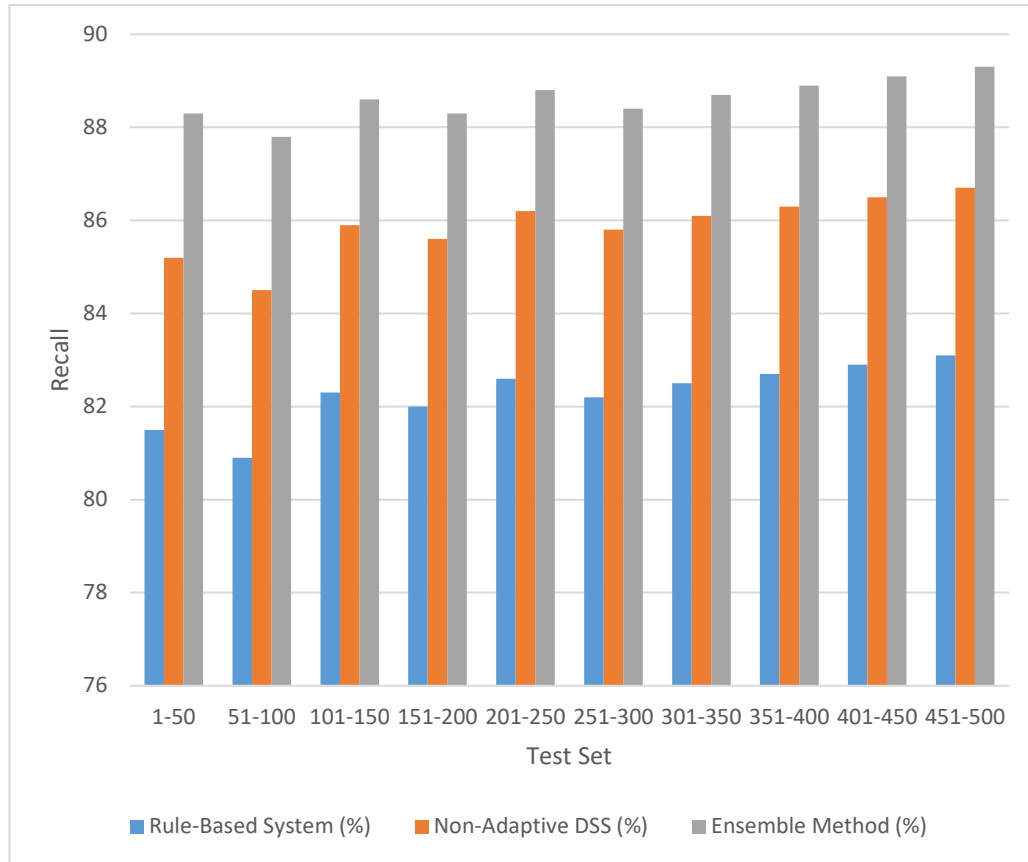


Figure 4. Recall

Table 4. Recall

Test dataset index	Rule-based system (%)	Non-adaptive DSS (%)	Ensemble method (%)
1-50	81.5	85.2	88.3
51-100	80.9	84.5	87.8
101-150	82.3	85.9	88.6
151-200	82.0	85.6	88.3
201-250	82.6	86.2	88.8
251-300	82.2	85.8	88.4
301-350	82.5	86.1	88.7
351-400	82.7	86.3	88.9
401-450	82.9	86.5	89.1
451-500	83.1	86.7	89.3

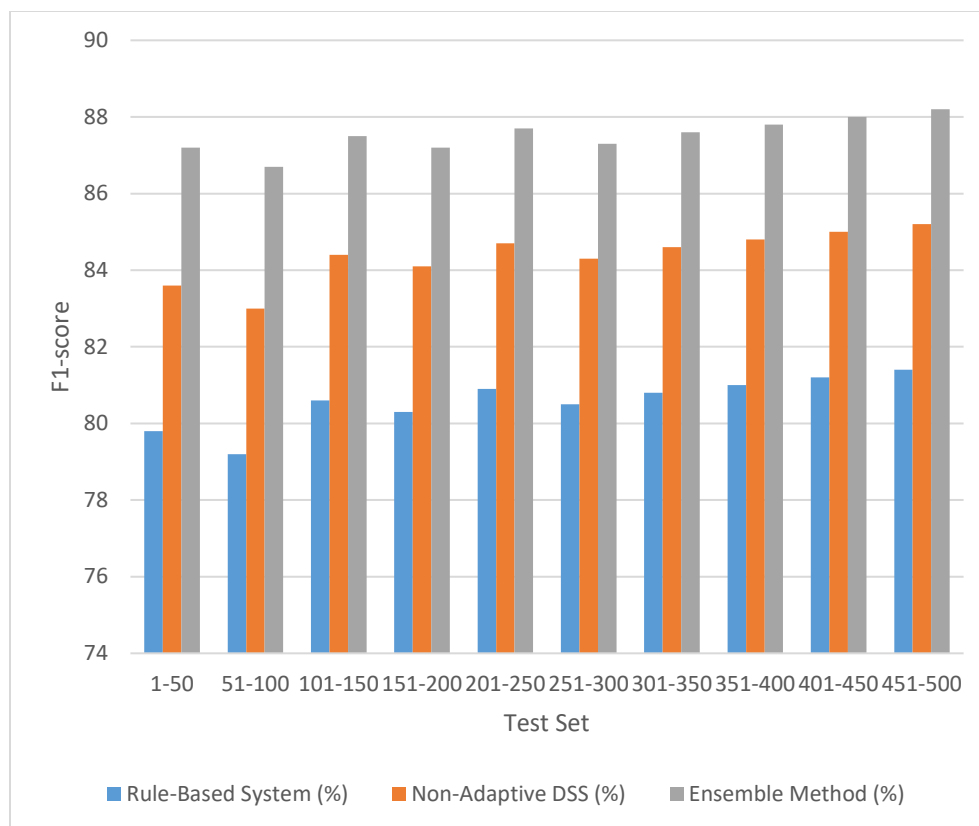


Figure 5. F-Score

Table 5. F-Score

Test dataset index	Rule-based system (%)	Non-adaptive DSS (%)	Ensemble method (%)
1-50	79.8	83.6	87.2
51-100	79.2	83.0	86.7
101-150	80.6	84.4	87.5
151-200	80.3	84.1	87.2
201-250	80.9	84.7	87.7
251-300	80.5	84.3	87.3
301-350	80.8	84.6	87.6
351-400	81.0	84.8	87.8
401-450	81.2	85.0	88.0
451-500	81.4	85.2	88.2

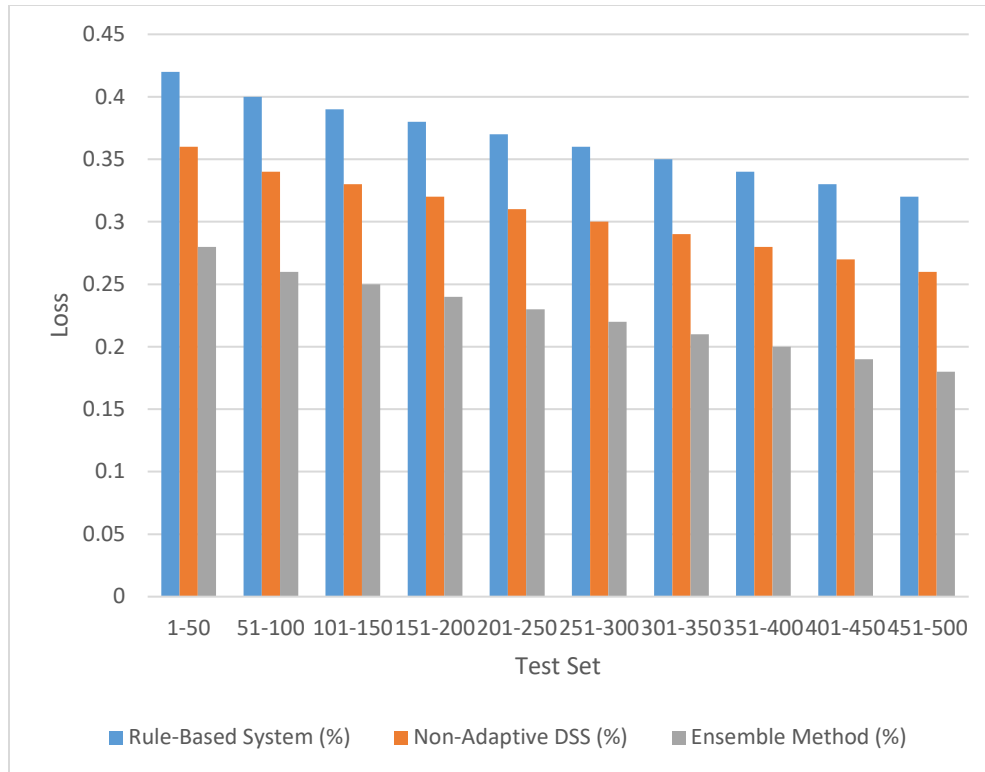


Figure 6. Loss

Table 6. Loss

Test dataset index	Rule-based system	Non-adaptive DSS	Ensemble method
1-50	0.42	0.36	0.28
51-100	0.40	0.34	0.26
101-150	0.39	0.33	0.25
151-200	0.38	0.32	0.24
201-250	0.37	0.31	0.23
251-300	0.36	0.30	0.22
301-350	0.35	0.29	0.21
351-400	0.34	0.28	0.20
401-450	0.33	0.27	0.19
451-500	0.32	0.26	0.18

DISCUSSION OF RESULTS

The results, as shown in Table 2 and Figure 2, show that the Ensemble approach routinely outperforms non-adaptive decision support systems and rule-based systems in terms of accuracy and precision over all test datasets. This implies that the Ensemble approach can provide precise and accurate recommendations, eventually improving decision-making in clinical situations such as those illustrated in Table 3 and Figure 3. Using the Ensemble method as a comparable point of reference, one

can see that it can efficiently find pertinent cases and reduce false negatives compared to other current techniques. In healthcare decision support systems, where patients' outcomes may suffer greatly if important diagnoses or treatments are missed, as Table 4 and Figure 4 illustrate, this is crucial.

The F-score, a measure that balances precision and recall, also consistently improves when the Ensemble approach is used, as Table 5 and Figure 5 show. This shows, therefore, that the Ensemble method achieves a more advantageous balance between detecting relevant cases and decreasing false positives, eventually improving decision-making performance.

The Ensemble method outperforms the other methods in reducing the error or discrepancy between the expected results and the real ground truth, as demonstrated by the comparison of losses in Table 6 and Figure 6.

The Ensemble method seems to have the potential to improve decision-making in a variety of clinical settings, as evidenced by the consistent performance improvements that have been observed across a wide range of evaluation metrics and test datasets. Lower loss values indicate a better alignment between the outcomes that were predicted and those that actually occurred. Its ability to dynamically adjust to changing patient conditions and optimize decision strategies based on observed outcomes makes applications in the real world of healthcare a good fit.

INFERENCES

The outcomes show that the DVM technique offers higher choice help than the rule-based totally and non-adaptive structures currently in use. The ability to optimize choice techniques and dynamically adapt to changing affected person situations leads to a development in accuracy, precision, consideration, and essential selection-making overall performance.

LIMITATIONS

The supply stage and first-class of the available healthcare data decide how nicely the Ensemble approach plays. Lack of representativeness, completeness, and excellence of the statistics may affect the validity and applicability of the effects.

CONCLUSION

The ensemble method helps with clinically relevant decisions. The consequences for the top affected persons and the great use of available sources ultimately comply with this. Many evaluation metrics, including accuracy, precision, do not forget, F-rating, and loss reduction, display that the Ensemble approach outperforms rule-based and non-adaptive structures presently in use. The results show that the Ensemble approach can handle the complexity and uncertainties that might be part of clinical choice-making. In doing this, it may offer docs realistic recommendations and insights catered to the specific needs of every patient. It's far vital, nevertheless, to recognize the restrictions and demanding situations related to applying the Ensemble method.

CHALLENGES

Implementing RL-based decision support systems in real-world clinical environments presents several challenges:

- Deploying the RL model across various healthcare settings with diverse patient populations and systems requires significant computational resources and customization. A solution could involve cloud-based deployments or federated learning, where the model learns from decentralized data sources without transferring sensitive patient information.
- Integrating RL models into existing Electronic Health Record (EHR) systems presents technical and operational challenges. Data must be seamlessly transferred between the RL agent and EHR systems, ensuring real-time updates. Standardizing data formats and protocols can

help address this challenge, as well as create interfaces that allow clinicians to easily interpret and act on the RL model's recommendations.

- Fix typographical mistakes, enhance sentence construction, and clarify technical descriptions. Although the study's main idea can be sensed through the text, revise the text to give readers a clearer presentation of the main idea and the concluding points based on the analysis of the study's findings. To a certain extent, the submission is a technical paper designed. It needs to give readers outside the narrow area of research a sufficiently clear insight into the subject of the research as well as an explanation of the findings and concluding statements.

The paper aims to demonstrate how reinforcement learning (RL) can improve decision-making in healthcare, especially for critical patients. The RL model continuously learns from patient data and adapts to changing conditions, providing real-time, personalized recommendations. By comparing it with traditional rule-based and non-adaptive systems, this study shows that the RL-based system outperforms others in accuracy, precision, and patient outcomes.

LIMITATIONS

The framework's reliance on the MIMIC-III dataset, a widely used critical care database, presents strengths and limitations. While MIMIC-III offers a rich set of patient data, it is limited to intensive care unit (ICU) settings, which may not represent the diversity of healthcare environments (e.g., outpatient, primary care, or non-critical care settings). Furthermore, the dataset is de-identified, which means that certain nuances or contextual patient information might be missing, potentially affecting the RL model's ability to capture the full spectrum of patient conditions.

The framework's ability to generalize to other healthcare environments or data sources will depend on how similar those environments are to critical care settings. Differences in patient demographics, clinical workflows, or data quality might require significant adaptations to the model. For example, outpatient data may include more chronic conditions that demand a different approach to decision-making, which would require additional fine-tuning of the RL agent.

Future work should expand the dataset to include diverse clinical environments (e.g., oncology, pediatrics, or emergency care) to enhance the generalizability of the framework. Incorporating data from multiple hospitals or healthcare systems can help increase robustness and improve the model's ability to adapt to varied patient populations.

REFERENCES

-
- Adlung, L., Cohen, Y., Mor, U., & Elinav, E. (2021). Machine learning in clinical decision making. *Med*, 2(6), 642-665. <https://doi.org/10.1016/j.medj.2021.04.006>
- Arivazhagan, N., Somasundaram, K., Vijendra Babu, D., Gomathy Nayagam, M., Bommi, R. M., Mohammad, G. B., Kumar, P. R., Natarajan, Y., Arulkarthick, V. J., Shanmuganathan, V. K., Srihari, K., Ragul Vignesh, M., & Prabhu Sundramurthy, V. P. (2022). Cloud-internet of health things (IOHT) task scheduling using hybrid moth flame optimization with deep neural network algorithm for E healthcare systems. *Scientific Programming*, 2022, Article 4100352. <https://doi.org/10.1155/2022/4100352>
- Coronato, A., Naeem, M., De Pietro, G., & Paragliola, G. (2020). Reinforcement learning for intelligent healthcare applications: A survey. *Artificial Intelligence in Medicine*, 109, 101964. <https://doi.org/10.1016/j.artmed.2020.101964>
- Fernandes, M., Vicira, S. M., Leite, F., Palos, C., Finkelstein, S., & Sousa, J. M. (2020). Clinical decision support systems for triage in the emergency department using intelligent systems: A review. *Artificial Intelligence in Medicine*, 102, 101762. <https://doi.org/10.1016/j.artmed.2019.101762>
- Kishor, A., Chakraborty, C., & Jeberson, W. (2021). Reinforcement learning for medical information processing over heterogeneous networks. *Multimedia Tools and Applications*, 80, 23983-24004. <https://doi.org/10.1007/s11042-021-10840-0>

- Kumar, T. S. (2020). Data mining-based marketing decision support system using hybrid machine learning algorithm. *Journal of Artificial Intelligence and Capsule Networks*, 2(3), 185-193.
- Liu, S., See, K. C., Ngiam, K. Y., Celi, L. A., Sun, X., & Feng, M. (2020). Reinforcement learning for clinical decision support in critical care: Comprehensive review. *Journal of Medical Internet Research*, 22(7), e18477. <https://doi.org/10.2196/18477>
- Peiffer-Smadja, N., Rawson, T. M., Ahmad, R., Buchard, A., Georgiou, P., Lescure, F. X., Birgand, G., & Holmes, A. H. (2020). Machine learning for clinical decision support in infectious diseases: A narrative review of current applications. *Clinical Microbiology and Infection*, 26(5), 584-595. <https://doi.org/10.1016/j.cmi.2019.09.009>
- Ragab, M., Albukhari, A., Alyami, J., & Mansour, R. F. (2022). Ensemble deep-learning-enabled clinical decision support system for breast cancer diagnosis and classification on ultrasound images. *Biology*, 11(3), 439. <https://doi.org/10.3390/biology11030439>
- Rani, P., Kumar, R., Ahmed, N. M. S., & Jain, A. (2021). A decision support system for heart disease prediction based upon machine learning. *Journal of Reliable Intelligent Environments*, 7(3), 263-275. <https://doi.org/10.1007/s40860-021-00133-6>
- Saravanan, V., Sankaradass, V., Shanmathi, M., Bhimavarapu, J. P., Deivakani, M., & Ramasamy, S. (2023, May). An early detection of ovarian cancer and the accurate spreading range in human body by using deep medical learning model. *Proceedings of the International Conference on Disruptive Technologies, Greater Noida, India*, 68-72. <https://doi.org/10.1109/ICDT57929.2023.10151103>
- Shah, V. (2020). Reinforcement learning for autonomous software agents: Recent advances and applications. *Revista Española de Documentación Científica*, 14(1), 56-71.
- Sharma, S., & Kumar, V. (2022). Application of genetic algorithms in healthcare: A review. In B. K. Tripathy, P. Lingras, A. K. Kar, & C. L. Chowdhary (Eds.), *Next generation healthcare informatics* (pp. 75-86). Springer. https://doi.org/10.1007/978-981-19-2416-3_5
- Sivakumar, N. Y. C., & Shankar, A. (2022). The speech-language processing model for managing the neuro-muscle disorder patients by using deep learning. *NeuroQuantology*, 20(8), 918-925.
- Tyler, N. S., Mosquera-Lopez, C. M., Wilson, L. M., Dodier, R. H., Branigan, D. L., Gabo, V. B., Guillot, F. H., Hiltz, W. W., El Youssef, J., Castle, J. R., & Jacobs, P. G. (2020). An artificial intelligence decision support system for the management of type 1 diabetes. *Nature Metabolism*, 2, 612-619. <https://doi.org/10.1038/s42255-020-0212-y>
- Yadav, R. S., Singh, A., & Yadav, S. P. (2024, March). The smart fuzzy-based image recognition model for modern healthcare. *Proceedings of the 2nd International Conference on Disruptive Technologies, Greater Noida, India*, 250-254. <https://doi.org/10.1109/ICDT61202.2024.10489059>
- Yu, C., Liu, J., Nemati, S., & Yin, G. (2021). Reinforcement learning in healthcare: A survey. *ACM Computing Surveys*, 55(1), Article 5. <https://doi.org/10.1145/3477600>

AUTHORS



Dr. S.P. Sreekala, Professor and Head of the Centre for Research and Development at the School of Management, KPR College of Arts, Science, and Research, is an accomplished academician with 23 years of teaching and 4 years of industry experience. She has published extensively with WoS articles, authored three books, and contributed book chapters to Scopus-indexed publications. Grants from esteemed organizations such as ICSSR, AICTE, TANSOST, RBI, ICPR, NCW, and NHRC have supported her research endeavors. A prolific researcher and innovator, Dr. Sreekala holds four patents and has guided more than 110 MBA students and over 20 international Ph.D. scholars. She serves as a subject matter expert for international universities and is a sought-after resource person at national and international conferences. Her contributions to

academia have earned her several accolades, including *Best Researcher*, *Distinguished Professor*, and *Best Senior Faculty* awards, underscoring her dedication to fostering knowledge and innovation.



Dr Mudit Saxena is working as an Assistant Professor at Teerthanker Mahaveer University, Moradabad. He has 15 years of experience in industry and academia, with a subjective specialty in organizational behavior, management concepts, industrial psychology, ethos and ethics, e-commerce, design thinking and innovation, and entrepreneurship. He is a coordinator of live project-based entrepreneurial initiatives, a mentor to slow-learner students, and an internal member of the internal quality assurance cell.



Dr S. Revathy, Associate Professor, PSG College of Arts and Science. She has 23 years of teaching experience and more than 10 Scopus journals and 15 UGC care journals published during this period. She has published three books on business law, production, operations management, and investment management.



Dr. Rajapriya. M. is working as an Assistant Professor in the Department of Management Studies at Veltech University, Chennai. She has more than 15.5 years of experience in teaching. Her area of specialization is HR and Marketing. Her qualifications are an MBA, MSc (Psychology), and MPhil. She specializes in HR and is interested in psychology. She has done research on a psychology-related HR topic titled “Disability of women working in the garment industry at Chennai,” which focuses on women’s mental and physical health in the garment industry. Her research thesis was awarded the best thesis award by U Mass Dartmouth Charlton College of Business, USA. She received the Global Excellence Award for Human Resources Development and Research from Special Minds, India, and UNIFACVEST University, Brazil, South America. She guided more than 30 PG projects. She is handling HR and marketing subjects. She gives personality development training to the students to help them become successful managers.



Dr. N.S. Shanthi, Professor and Head of the Department of MBA at SRM University, Sikkim, is an accomplished academician with over 19 years of teaching experience. Her expertise spans Finance, Marketing, and Human Resource Management. A life member of the Indian Society for Technical Education, she has published over 32 research articles, presented 34 conference papers, and authored an e-book with LAP Lambert. She has guided 90+ MBA students and 15 M.Phil. scholars while serving as a subject expert, external examiner, and co-guide for Texila American University. Recognized for her contributions, Dr. Shanthi has received numerous accolades, including the *Women Professor of the Year 2024* by MIT Square and the *Senior Educator and Scholar Award*

2023 by NFED. She has delivered lectures, acted as a resource person, and organized several national and international conferences, faculty development programs, and inter-collegiate meets. Her research initiatives have been supported by ICSSR, NCW, RBI, and NHRC funding, further solidifying her reputation as a leader in academia.



Dr. S. Saravanakumar, Assistant Professor, has 16 years of service in the Department of Business Administration at Annamalai University. He is a soft-skill trainer and a JCI India Certified soft-skill trainer.