International Journal of Engineering Research and Science & Technology

www.ijerst.org

ISSN : 2319-5991

Vol. 22 No. 1(s) (2026)



ijerst.editor@gmail.com

editor@ijerst.com

www.ijerst.com
© 2026 IJERST. All Rights Reserved
Peer Reviewed, Referred & Indexed Journal

Research Paper

Enhancing Clinical Decision Support: A Hybrid Framework Integrating Markov Decision Processes and Explainable Neurosymbolic Reasoning

Name Of Author

Wabale Pallavi Ajinath

Assistant Professor,Department Of Computer Science , Jijamata college Of Science and Arts Bhende

Landepallavi123@gmail.com

1.Abstract

The integration of Artificial Intelligence (AI) into healthcare promises to revolutionize diagnostic accuracy and treatment optimization. However, the deployment of these systems faces significant hurdles related to trust, explainability, and the ability to reason through complex, sequential clinical scenarios. Current "black-box" deep learning models often fail to provide the transparency required for medical and legal accountability, while traditional rule-based systems lack the flexibility to learn from vast biomedical datasets. This paper proposes a hybrid framework that combines Markov Decision Processes (MDPs) for temporal treatment planning with Neurosymbolic AI to integrate logical reasoning with data-driven learning. By prioritizing explainability and the rigorous management of uncertainty, this approach aims to support clinicians in making optimal, evidence-based decisions. We discuss the theoretical underpinnings of this method, analyse its relation to existing literature, and outline a robust evaluation strategy to measure both performance and clinical helpfulness.

2.Introduction

The modern healthcare ecosystem is characterized by rapidly expanding complexity, exploding information streams, and a growing myriad of treatment options that often exceed the cognitive capacity of individual practitioners (Bennett & Hauser, 2013). Artificial Intelligence (AI) has emerged as a critical tool to address these challenges, offering the potential to analyse high-dimensional patient data and suggest interventions. Despite the success of AI in specific tasks like image segmentation, the broader adoption of AI for high-stakes clinical decision-making remains limited. This hesitation stems largely from the "black-box" nature of deep neural networks, where the internal logic driving a prediction is opaque to the user.

Existing approaches to AI in healthcare are insufficient for widespread clinical integration for several key reasons. First, standard supervised learning models typically treat diagnosis as a static classification problem, failing to account for the sequential and temporal nature of patient care, where a decision made today influences the patient's state tomorrow (Bennett & Hauser, 2013). Second, pure data-driven models often lack the ability to provide meaningful justifications for their outputs. As noted in recent surveys, legal and ethical requirements mandate effective Explainable AI (XAI), yet many current methods prioritize model performance over human-understandable reasoning (Labara et al., 22024) (harati et al., 2023). Third, standard models struggle with "off-nominal" or anomalous cases; they perform well on

distribution data but often fail when encountering novel patient presentations that require creative adaptation of existing medical knowledge (Gizzi et al., 2022).

3.Problem Statement

Modern Clinical Decision Support Systems (CDSS) are increasingly driven by data-intensive machine learning models to assist clinicians in diagnosis, prognosis, and treatment planning. While these systems demonstrate strong predictive performance, they suffer from three major limitations:

1. **Limited Sequential Decision Modeling**
   Many existing CDSS treat medical decisions as isolated predictions rather than sequential, dynamic processes. However, clinical care (e.g., chronic disease management, ICU treatment, oncology protocols) involves step-by-step decisions over time under uncertainty. Traditional models fail to explicitly optimize long-term patient outcomes.
2. **Lack of Explainability and Transparency**
   Deep learning models operate largely as "black boxes," providing limited interpretability. In high-stakes healthcare environments, clinicians require transparent reasoning to trust and validate system recommendations. Regulatory requirements (e.g., explainable AI mandates) further emphasize this need.
3. **Poor Integration of Domain Knowledge**
   Purely data-driven models often ignore established medical guidelines, pathophysiological knowledge, and causal reasoning. This can result in recommendations that are statistically accurate but clinically questionable or unsafe.

4.Primary Objective

To develop and evaluate a hybrid Clinical Decision Support framework that integrates Markov Decision Processes (MDPs) with Explainable Neurosymbolic Reasoning to optimize sequential clinical decision-making while ensuring interpretability, transparency, and adherence to medical domain knowledge.

5.Literature Review

1. Clinical Decision Support Systems (CDSS)

Clinical Decision Support Systems have evolved from simple rule-based systems to advanced machine learning-driven platforms. Early CDSS relied heavily on expert-defined rules and clinical guidelines, which ensured interpretability but lacked adaptability and scalability. Modern AI-driven CDSS leverage deep learning and predictive analytics to improve diagnostic accuracy and risk prediction.
However, despite improved performance, many contemporary systems suffer from limited interpretability, reduced clinician trust, and challenges in regulatory compliance.

Gap Identified: Existing CDSS either prioritize performance (black-box models) or interpretability (rule-based systems), but rarely both simultaneously.

2. Markov Decision Processes in Healthcare

Markov Decision Processes (MDPs) and reinforcement learning have been widely applied in healthcare for modeling sequential treatment decisions under uncertainty. Applications include:

- Chronic disease management
- ICU treatment optimization
- Personalized treatment planning

MDPs allow modeling of state transitions, action policies, rewards, and long-term outcome optimization. They are particularly effective in scenarios where decisions must account for future consequences.

Limitation:
Most MDP-based healthcare models lack transparency. Clinicians often cannot understand why a specific policy is recommended, reducing trust and adoption.

6.Explainable Artificial Intelligence (XAI) in Healthcare

Explainable AI techniques aim to make AI systems transparent and interpretable. Methods include:

- Rule extraction
- Attention mechanisms
- Feature importance scoring
- Counterfactual explanations

In healthcare, explainability is critical due to ethical, legal, and safety requirements. However, many explainability techniques are post-hoc, meaning explanations are generated after prediction rather than being inherently built into the model.

Gap Identified: There is a need for inherently interpretable architectures rather than post-hoc explanation methods.

7.Neurosymbolic AI and Knowledge Integration

Neurosymbolic AI combines neural networks (data-driven learning) with symbolic reasoning (logic-based rules). This approach allows:

- Integration of domain knowledge
- Logical constraint enforcement
- Improved interpretability
- Better generalization

In medical contexts, symbolic rules can encode clinical guidelines, disease ontologies, and causal relationships. When combined with neural learning, this results in systems that are both adaptive and knowledge-aware.

Limitation:
Few studies have integrated neurosymbolic reasoning with sequential decision frameworks like MDPs for clinical applications.

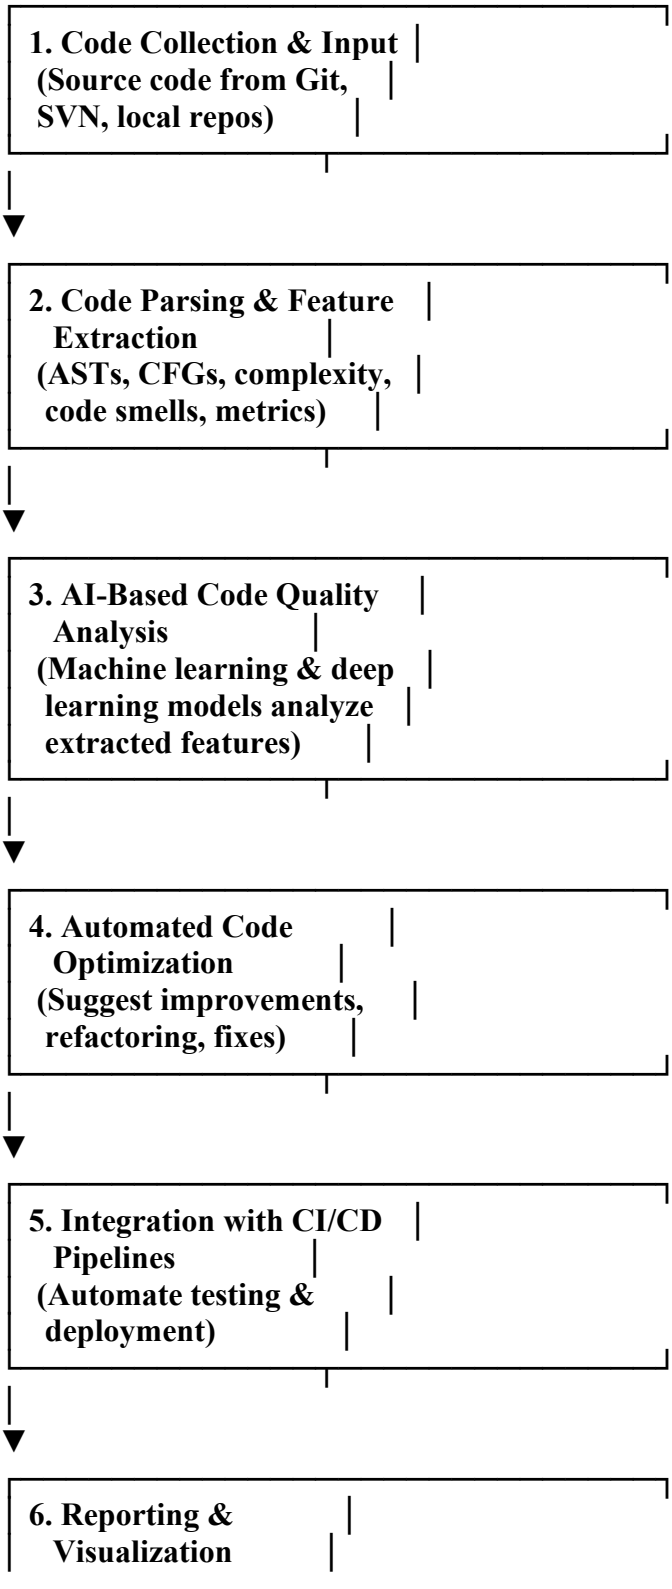
8.Research Gap

From the reviewed literature, the following gaps emerge:

1. Limited integration of sequential decision optimization and explainable reasoning.

2. Lack of unified frameworks combining MDPs with neurosymbolic models.
3. Over-reliance on black-box reinforcement learning in clinical settings.
4. Insufficient embedding of clinical knowledge constraints into adaptive decision models.

9.Proposed System Architecture:



(Dashboards, defect summaries, severity)

1.Code Collection and Input

- Collect source code from multiple repositories such as Git, SVN, or local environments.
- Supports multiple languages (Java, Python, C++, JavaScript).
- Standardizes formatting, removes comments, and tokenizes code to prepare clean data **for analysis.**

2. Code Parsing and Feature Extraction

- Parse code to understand its structure using Abstract Syntax Trees (ASTs) and Control Flow Graphs (CFGs).
- Extract features like cyclomatic complexity, nesting depth, code smells (e.g., duplicate code), and performance bottlenecks.
- These features serve as input for AI models.

3. AI-Based Code Quality Analysis

- Apply machine learning and deep learning models trained on historical datasets to assess code quality.
- Detect bugs, security vulnerabilities, and code smells automatically.
- Identify poor design patterns or anti-patterns affecting maintainability

4. Automated Code Optimization

- Suggest code improvements such as refactoring options or fixing detected issues.
- May include automated patching or recommendations to developers.

5. Integration with CI/CD Pipelines

- Seamlessly integrate analysis into continuous integration and deployment workflows.
- Automate quality checks, testing, and deployment in the development lifecycle.

6. Reporting and Visualization

- Provide dashboards summarizing code quality metrics, defect predictions, and severity indicators.
- Enable teams to monitor and prioritize improvements efficiently

10.Framework Architecture

The system is composed of three primary modules:

1.**The State Representation Module (Probabilistic):**

Patient data (vitals, history, labs) is ingested to form a "belief state." Recognizing that medical data is often incomplete (partially observable), we adopt the perspective of constructive probability, where judgments are subjective fits to canonical examples (Shafer, 2013). This module uses belief functions to represent uncertainty explicitly, distinguishing between "risk" (known probability) and "ambiguity" (missing information).

2.The Neurosymbolic Reasoning Engine:

This core module integrates logical constraints with a learning mechanism. Following the dimensions outlined by Marra et al., this system utilizes a "proof-based" inference approach where neural networks guide the symbolic search for a diagnosis (Marra et al., 2021).

- ***Symbolic Component:*** Encodes clinical guidelines (e.g., "If sepsis is suspected, antibiotics must be administered within one hour").
- ***Neural Component:*** Learns patterns from unstructured data (e.g., radiological images) to populate the symbols used in the logic.

3.The Sequential Decision Optimizer (MDP):

Once the state is defined and constraints are applied, the MDP calculates the optimal policy. The goal is to maximize a reward function defined by patient health utility (quality of life, survival). This component simulates alternative sequential decision paths to predict long-term outcomes (Bennett & Hauser, 2013).

11.Method/Approach

We propose a framework titled **Clinical-NeSy-MDP** (Clinical Neurosymbolic Markov Decision Process). This framework is designed to function as a simulation environment for exploring treatment policies and as a decision support tool (Bennett & Hauser, 2013).

12.Future Work

Future research should focus on two avenues. First, incorporating mechanisms for **Creative Problem Solving (CPS)** is essential. Agents must be designed to recognize when a problem is outside their scope and attempt to adapt knowledge from different medical domains to the current context (Gizzi et al., 2022). Second, we aim to expand the probabilistic judgment module to better integrate disparate "languages" of probability—reconciling Bayesian updates with Dempster-Shafer belief functions to better model conflicting medical evidence (Shafer, 2013).

13.Conclusion

This paper presented a hybrid framework for AI in healthcare that addresses the critical shortcomings of current purely data-driven approaches. By coupling the temporal planning capabilities of Markov Decision Processes with the transparency and logic of Neurosymbolic AI, we offer a path toward systems that can "think" more like physicians—balancing evidence, protocol, and uncertainty. While challenges in data representation and ethical deployment remain, the integration of reasoning and learning is the necessary next step to move AI from a passive analytical tool to an active, trustworthy partner in clinical decision-making.

References

Bennett, Casey C., & Hauser, Kris (2013). *Artificial Intelligence Framework for Simulating Clinical Decision-Making: A Markov Decision Process Approach*. Artificial Intelligence in Medicine. 57(1): 9-19. (2013). https://doi.org/10.1016/j.artmed.2012.12.003

Labarta, Tobias, Kulicheva, Elizaveta, Froelian, Ronja, Geißler, Christian, Melman, Xenia, & Klitzing, Julian von (2024). *Study on the Helpfulness of Explainable Artificial Intelligence*. Longo, L., Lapuschkin, S., Seifert, C. (eds) Explainable Artificial Intelligence. xAI 2024. Communications in Computer and Information Science, vol 2156. https://doi.org/10.1007/978-3-031-63803-9_16

Bharati, Subrato, Mondal, M. Rubaiyat Hossain, & Podder, Prajoy (2023). *A Review on Explainable Artificial Intelligence for Healthcare: Why, How, and When?*. IEEE Transactions on Artificial Intelligence, 2023. https://doi.org/10.1109/TAI.2023.3266418

Gizzi, Evana, Nair, Lakshmi, Chernova, Sonia, & Sinapov, Jivko (2022). *Creative Problem Solving in Artificially Intelligent Agents: A Survey and Framework*. Journal of Artificial Intelligence Research 2022. https://doi.org/10.1613/jair.1.13864

Marra, Giuseppe, Dumančić, Sebastijan, Manhaeve, Robin, & Raedt, Luc De (2021). *From Statistical Relational to Neurosymbolic Artificial Intelligence: a Survey*. https://arxiv.org/pdf/2108.11451v4

Shafer, Glenn (2013). *Probability Judgement in Artificial Intelligence*. https://arxiv.org/pdf/1304.3429v1