Journal of Neonatal Surgery
ISSN(Online): 2226-0439
Vol. 14, Issue 13s (2025)
https://www.jneonatalsurg.com

AI-Driven Decision Support Systems for Neonatal Care: A Reinforcement Learning Approach

Sigirisetty Anusha[1], Chopparapu Gowthami[2], Kiran Onapakala[3], V. Suryanarayana[4], Vaddithandra Vijaya[5], Rongala Ravi[6]

[1]Assistant Professor, Department of AI, Shri Vishnu Engineering College for Women, Vishnupur, Bhimavaram, 534202.

[2]Assistant Professor, Department of CSE, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Guntur, Andhra Pradesh, India,

[3]Pacific Source health plans,1500 SW 1st Ave #200, Portland, OR 97201.

[4]CVR College of Engineering, Vastunagar, Hyderabad, 501510, India,

[5]Assistant Professor, Department of Computer Science and Engineering, CMR Institute of Technology, Hyderabad,

[6]ASSISTANT PROFESSOR, Department of CSE-AIML, Malla Reddy Engineering College, Maisammaguda, Secuderabad-500100,

Email ID: raviecet5815@gmail.com

Cite this paper as: Sigirisetty Anusha, Chopparapu Gowthami, Kiran Onapakala, V. Suryanarayana, Vaddithandra Vijaya, Rongala Ravi, (2025) AI-Driven Decision Support Systems for Neonatal Care: A Reinforcement Learning Approach. *Journal of Neonatal Surgery*, 14 (13s), 928-936.

ABSTRACT

Neonatal surgical care involves high-stakes decision-making under conditions of uncertainty, urgency, and limited physiological feedback. Clinicians often rely on experience and generalised guidelines to decide whether to proceed with surgical intervention—an approach that may not always capture the nuances of individual patient presentations. Artificial Intelligence (AI), particularly Reinforcement Learning (RL), offers a promising avenue to improve consistency and adaptiveness in such critical care scenarios. Unlike traditional machine learning models, RL learns optimal decision strategies by balancing rewards and risks through iterative feedback, making it suitable for sequential and high-impact clinical environments. This study presents a conceptual AI-driven Decision Support System (DSS) that leverages RL principles to assist clinicians in binary surgical decisions for neonates. The system employs manually constructed clinical state–action mappings, expert-informed reward logic, and an explainable Q-table rather than relying on patient data or simulations. It features a three-layered architecture, visual decision flowchart, and event tree to support transparent reasoning. Through hypothetical clinical scenarios and clinician-oriented workflow modelling, the system demonstrates potential for low-resource settings and academic prototyping. While preliminary, the model offers a scalable, modifiable foundation for future integration with real-world clinical platforms, aiming to enhance decision quality in neonatal surgical care.

Keywords: *Neonatal care, Reinforcement learning, Clinical decision support, Surgical decision-making, Artificial intelligence*

1. INTRODUCTION

Neonatal surgical care demands timely, precise, and life-saving decisions in high-pressure environments where patient conditions can deteriorate rapidly(Jeong & Kamaleswaran, 2022). The complexity of such care is compounded by the variability in neonates' responses to interventions and the difficulty in objectively quantifying critical thresholds for surgical intervention(Guez-Barber & Pilon, 2024). In these sensitive scenarios, clinicians often rely on their experience, guideline-based protocols, and clinical intuition to decide whether immediate surgical intervention is warranted or if conservative management might suffice.

Traditional decision-making models, although essential, are inherently limited in their adaptability and responsiveness(Catania, 2021). Clinical guidelines often adopt a one-size-fits-all approach, which may not account for subtle yet critical differences in individual patient presentations. Furthermore, human judgment, while invaluable, is susceptible to bias, fatigue, and variability across practitioners. As neonatal outcomes are closely tied to the timing and appropriateness of surgical interventions, there is a compelling need for intelligent systems that can support clinicians by providing consistent, data-informed, and context-aware recommendation(Jyoti et al., 2023).

Artificial Intelligence (AI), particularly Reinforcement Learning (RL), offers a promising avenue to enhance decision-making in such scenarios. Unlike supervised learning, which depends on labelled datasets, RL involves an agent that learns optimal decision policies by interacting with an environment and receiving rewards or penalties based on outcomes. This allows for the development of adaptive systems that refine their strategies through experience, making RL highly suitable for sequential and dynamic clinical settings(Guez-Barber & Pilon, 2024).

This study proposes a conceptual AI-driven Decision Support System (DSS) based on reinforcement learning principles, aimed at assisting clinicians in determining whether surgery is required for neonates under specific clinical conditions(Levin et al., 2024). The approach does not involve patient data or simulation-based training but instead uses manually defined clinical states, reward logic, and state-action mappings to create an explainable and lightweight model. Such a system is particularly valuable for early-stage design, academic validation, and use in low-resource or data-scarce environments(Muntean et al., 2025).

2. LITERATURE REVIEW

2.1 Decision Support Systems in Neonatal Surgery:

Decision Support Systems (DSS) have been developed across various fields of medicine to aid in diagnosis, prognosis, and treatment planning. These systems range from basic rule-based models to advanced machine learning applications(Muntean et al., 2025). However, in neonatal surgery, the integration of DSS remains limited due to several factors, including the scarcity of large, high-quality datasets, the ethical constraints of experimentation, and the variability of neonatal conditions(Jaile et al., 2024).

Most existing DSS tools in neonatal care are focused on monitoring, infection risk prediction, or general ICU support rather than binary decision-making specific to surgical interventions. Moreover, rule-based systems, while easy to implement, struggle with adaptability and cannot adjust to rapidly evolving clinical states. Thus, there is a growing interest in developing intelligent systems that incorporate real-time feedback and offer context-sensitive recommendations(Jyoti et al., 2023).

2.2 Reinforcement Learning in Healthcare:

Reinforcement Learning has gained attention in the healthcare domain for its unique capability to handle sequential decision-making problems(Thakre et al., 2025). It has been successfully applied to areas such as glucose control in diabetic patients, optimal dosing of sedatives in ICU settings, and dynamic treatment strategies for sepsis. The RL framework operates through the interaction of an agent with an environment, learning to select actions that maximise cumulative reward(Roayaei & Soltani, 2025).

In a clinical context, the “state” could represent a patient’s physiological parameters, the “action” could be a treatment decision, and the “reward” could be the outcome (e.g., survival, stability, recovery). This paradigm is particularly suited for critical care and surgical decisions where timing and progression play vital roles(Nadhir et al., 2025). Despite its potential, the application of RL in neonatal surgical care is virtually unexplored. This study addresses that gap by presenting a conceptual model that applies RL logic to the binary decision of surgery versus observation(Lakhan et al., 2024).

3. MATERIALS AND METHODS:

This study was designed as a **15-day conceptual modelling exercise**, intentionally avoiding the use of real patient data or simulation episodes. The focus was to build a lightweight, explainable, and adaptable AI-driven Decision Support System (DSS) using reinforcement learning (RL) principles, specifically tailored to assist clinical decision-making in neonatal surgical care. The system aimed to support a **binary decision framework**: whether to proceed with surgery or opt for continued observation(Nadhir et al., 2025).

The core framework draws inspiration from the fundamentals of reinforcement learning, in which an agent interacts with an environment by taking actions based on its current state and receiving a reward based on the outcome(Onapakala et al., 2024). However, instead of learning from experience, this conceptual model uses manually constructed state-action mappings informed by clinical knowledge and reward logic. These mappings form the basis of a **Q-table**, with each combination of patient condition (state) and potential intervention (action) being assigned a reward value to reflect expected clinical benefit or harm.

To define the relevant clinical states for neonatal surgical decisions, five key indicators were selected based on commonly observed parameters in neonatal intensive care units. These include physiological, biochemical, and clinical markers that typically influence surgical considerations. Each state was assigned a qualitative risk level (moderate, high, very high, or critical) and a “decision relevance” value to reflect how strongly it should influence the decision to operate.

Table 1: Key Clinical States for Surgical Decision-Making

S.no	Clinical State	Indicator	Risk Level	Decision Relevance
1.	Abdominal distension	Visible swelling or radiologic signs	High	+
2.	Lactate > 5 mmol/L	Blood test	Very High	++
3.	Oxygen saturation < 90%	Pulse oximeter	Moderate	+
4.	Elevated CRP or Procalcitonin	Inflammatory marker	High	+
5.	Hypotension unresponsive to fluids	Systolic < 60 mmHg	Critical	++

The identification of critical clinical states is essential in neonatal surgical decision-making to ensure timely interventions and improve patient outcomes. The categorization of risk levels and decision relevance in Table 1 provides a structured framework for evaluating key clinical indicators in a systematic manner. Abdominal distension, lactate levels, and hypotension are among the crucial factors that influence surgical decisions, as outlined in Table 1. A high lactate level, particularly when exceeding 5 mmol/L, is a strong indicator of metabolic distress and is marked as "very high risk" in Table 1. The presence of inflammatory markers such as CRP and Procalcitonin, as detailed in Table 1, suggests a potential infectious or inflammatory response that may necessitate further clinical evaluation. Hypotension that does not respond to fluid resuscitation is classified as "critical" in Table 1, highlighting its significance in urgent surgical decision-making. The structured classification of clinical states in Table 1 allows for a logical mapping of patient conditions to potential surgical actions, reducing ambiguity in clinical assessment. By integrating the risk levels from Table 1 into the decision-making framework, the model enhances the objectivity and consistency of surgical recommendations. The parameters listed in Table 1 align with established neonatal intensive care unit (NICU) guidelines, reinforcing their clinical relevance in determining the necessity for surgical intervention. The decision relevance scores in Table 1 help prioritize conditions that have the highest impact on neonatal outcomes, ensuring that high-risk cases receive immediate attention.

These clinical states were then integrated into the overall system architecture. The architecture was conceptualised as a three-layered framework: an input layer, a reinforcement learning logic engine, and an output module. The input layer captures clinical data from the user (vitals, laboratory findings, and physical symptoms). The processing layer applies reinforcement logic to the input data, referencing the Q-table to determine the best action. Finally, the output layer provides a decision recommendation—either to proceed with surgery or to continue observation—along with an optional confidence score or justification text.

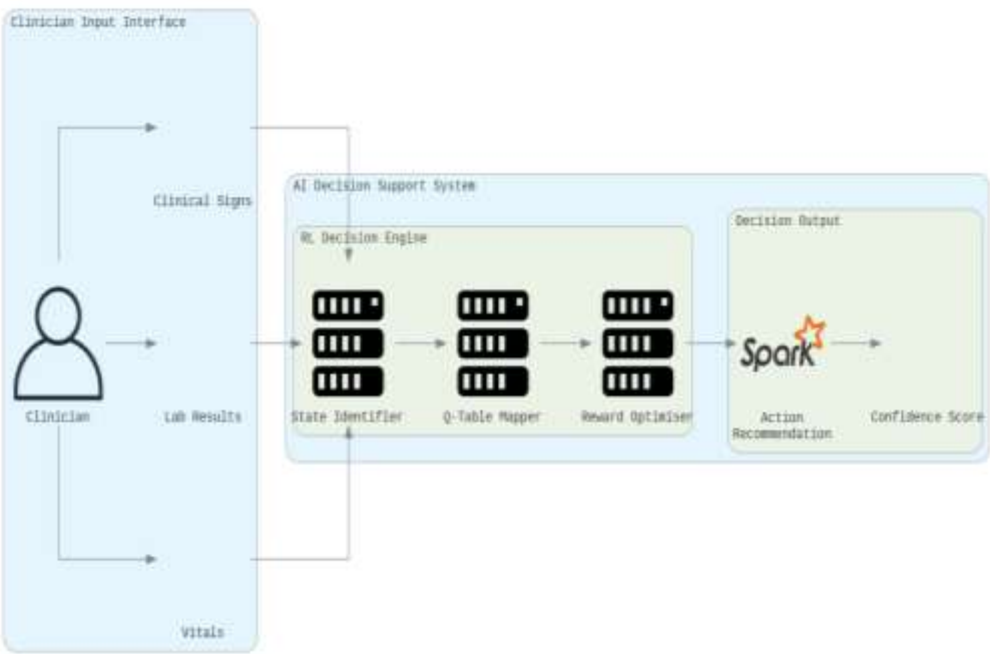

Figure 1: Architecture of the AI-DSS

The AI-driven Decision Support System (DSS) is designed around a structured, three-layered architecture. This architecture, visually represented in Figure 1, facilitates the logical flow of data from input to decision output. The system's architecture, as shown in Figure 1, comprises an input layer, a reinforcement learning (RL) decision engine, and an output module. The input layer's primary function is to capture relevant clinical data, such as vitals, laboratory results, and observed symptoms, from the clinician user.

The RL decision engine then processes this input data, utilizing the Q-table to match the clinical state with appropriate actions. Following this processing, the output module generates an action recommendation, suggesting either surgery or observation. This output may also include a confidence score to indicate the certainty of the recommendation. Effectively, Figure 1 illustrates the system's operational sequence, clarifying how data is transformed into a decision recommendation.

This diagram represents the logical flow of data within the DSS, from input to decision. It includes:

- Input: Vitals, labs, symptoms
- RL Decision Engine: Matches state–action pairs
- Output: Action recommendation + confidence score

The heart of the system is a **manually created Q-table**, in which representative clinical state combinations are mapped to one of the two actions (Surgery or Observe), along with an assigned reward value. This value reflects how beneficial the action is considered, given the state, and is based on existing clinical guidelines and neonatal care principles.

Table 2: Manually Created Q-Table for Surgical Decision Logic

S.no	State ID	Clinical State Combination	Action	Reward	Justification
1..	S1	High lactate + low SpO_2 + distension	Surgery	10	Life-threatening infection suspected
2.	S2	Mild SpO_2 drop + no distension	Observe	6	Monitoring preferred before escalation
3.	S3	CRP elevated + borderline vitals	Surgery	7	Early intervention could improve outcome
4.	S4	Normal vitals + mild distension	Observe	5	Supportive care may suffice
5.	S5	Hypotension + SpO_2 < 85% + distension	Surgery	9	Likely surgical emergency

The manually created Q-table is a fundamental part of the AI-driven Decision Support System. This table, presented as Table 2, plays a crucial role in mapping specific combinations of clinical states in neonates to the system's recommended actions, which are either to proceed with surgery or to continue with observation. A core function of Table 2 is the assignment of reward values to each potential decision, providing a structured approach to evaluating the potential benefits and risks associated with surgical intervention. For instance, when the system identifies a clinical state involving high lactate levels, low oxygen saturation, and abdominal distension, a high reward is assigned to the "Surgery" action, reflecting the urgency of the situation. In contrast, cases with mild oxygen desaturation and no abdominal distension receive lower reward values for surgery, suggesting that observation may be more appropriate. To enhance transparency, the "Justification" column in Table 2 explains the reasoning behind these reward assignments. This reward-based approach in Table 2 inherently balances the risks of delaying necessary intervention against the potential complications of unnecessary surgery. Furthermore, the design of Table 2 allows for modifications and refinements based on clinical feedback, making the model adaptable. Ultimately, the structured mappings within Table 2 provide a foundation for a reliable and interpretable decision support system, with its alignment to neonatal care principles enhancing its potential for real-world application.

To visualise how the system processes data and reaches a recommendation, a decision flowchart was developed. This flowchart captures the step-by-step logic embedded in the RL system. Starting from clinical input, the system classifies the current state, consults the Q-table, and returns the highest-reward action.

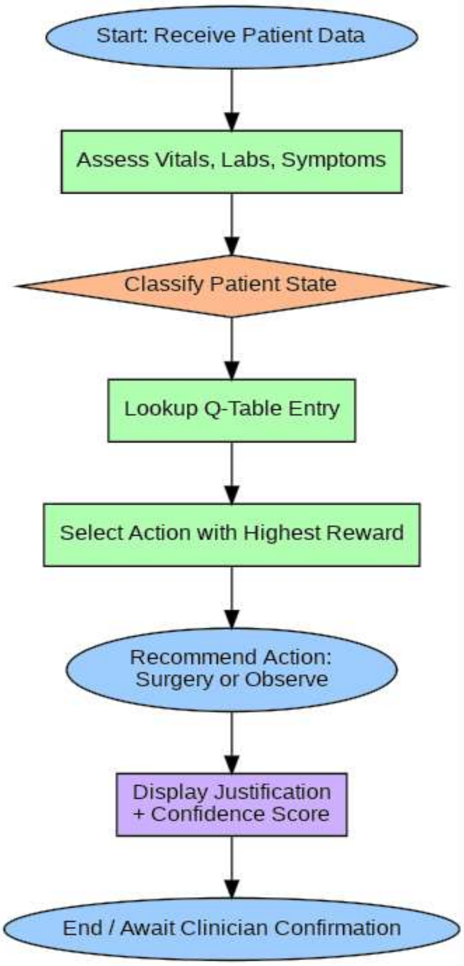

Figure 2: RL Decision Flowchart

To clearly illustrate the system's decision-making process, a detailed decision flowchart was developed. This flowchart, presented as Figure 2, effectively captures the step-by-step logic embedded within the reinforcement learning system. The flowchart in Figure 2 begins with the system's assessment of the input clinical state, reflecting the patient's condition. Following this assessment, the system proceeds to identify the specific clinical conditions present. The system then consults the Q-table, matching the identified conditions with corresponding actions, either surgery or observation. Subsequently, the system suggests the action that is associated with the highest reward value, indicating the most favourable decision. As shown in Figure 2, the process culminates in the output of the recommended action to the clinician. This visual representation in Figure 2 aids in understanding the system's operational flow and reasoning.

Flowchart Logic*:*

1. Assess input state
2. Identify condition
3. Match to Q-table
4. Suggest action
5. Output recommendation

To support deeper reasoning, an **event tree model** was created. This event tree outlines the range of possible outcomes following each of the two actions—surgery or observation. It visually presents potential risks such as post-operative complications or deterioration due to delayed surgery, which are useful for understanding the logic behind the reward assignments.

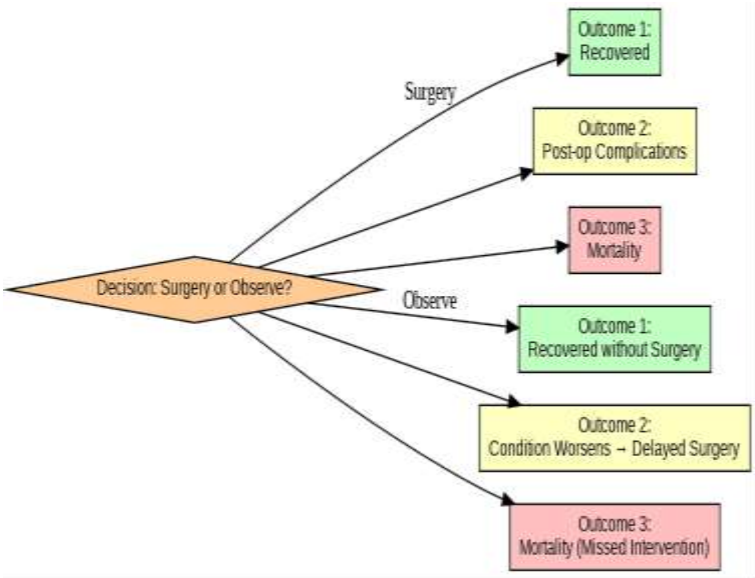

Figure 3: Event Tree – Surgical vs. Non-Surgical Outcomes

To support more in-depth reasoning about potential outcomes, an event tree model was developed. This event tree, illustrated in Figure 3, outlines the range of possible outcomes following each of the two actions: surgery or observation. As shown in Figure 3, it visually represents potential risks associated with each decision, such as post-operative complications following surgery. The event tree in Figure 3 also depicts potential negative outcomes like deterioration due to delayed surgery in the observation arm. These visual representations are useful for understanding the logic behind the reward assignments within the decision-making framework. The branches in Figure 3 detail the possible trajectories, including recovery, complications, or death after surgery. Similarly, Figure 3 shows that observation can lead to recovery, the need for delayed surgery, or death. Ultimately, this event tree in Figure 3 provides a clear visualization of the potential consequences of each decision.

Branches:

- Surgery → Recovery, Complication, Death
- Observe → Recovery, Delayed Surgery, Death

To further test and validate the model, a set of hypothetical scenarios was developed. Each scenario simulated a real-world clinical presentation and was mapped to a reward-based outcome depending on whether surgery or observation was chosen. These scenarios were useful in stress-testing the consistency and rationale of the decision logic.

Table 3: Hypothetical Scenario Action–Outcome–Reward Mapping

S.no	Scenario ID	Action	Clinical Outcome	Assigned Reward	Notes
1	A1	Surgery	Infection resolved	10	Early intervention successful
2	A2	Observe	Deterioration, late surgery	4	Delay caused worsening
3	A3	Observe	Stable condition maintained	8	Correct non-invasive choice
4	A4	Surgery	Complication post-op	5	Risk balanced with benefit
5	A5	Observe	Sudden collapse	0	Missed opportunity for early action

The hypothetical scenario action-outcome-reward mapping is essential for evaluating the practical application of the reinforcement learning model. This crucial evaluation tool, presented in Table 3, details several simulated clinical scenarios alongside their assigned action, outcome, and reward values. Each scenario within Table 3 pairs a potential clinical action, either surgery or observation, with a corresponding projected clinical outcome and a numerical reward. These pairings enable a systematic assessment of the reinforcement learning logic's performance under various conditions. For example, when surgery leads to a positive resolution, such as an infection being resolved, the scenario in Table 3 assigns a high reward. Conversely, if observation results in a negative outcome like deterioration requiring delayed surgery, Table 3 assigns a lower reward. These varied reward assignments within Table 3 serve to highlight instances where the reinforcement learning logic might exhibit tendencies toward more conservative or aggressive surgical recommendations. Ultimately, the collection of manually constructed scenarios found in Table 3 contributes significantly to the validation of the model's clinical soundness and its potential to aid in informed clinical decision-making.

These manually constructed scenarios helped assess the practical utility of the reward system and highlighted cases where RL logic was either conservative or aggressive in recommending surgery. Overall, the model is designed to support clinical reasoning with transparent, consistent, and modifiable logic, suitable for academic prototyping and future integration with real-world clinical platforms.

4. CLINICIAN WORKFLOW USING THE SYSTEM

To ensure clinical usability, the AI-driven Decision Support System (DSS) was designed with a straightforward interaction flow that complements existing hospital workflows(Khan, 2025). The system is intended for use at the bedside or within an electronic medical record interface, requiring only basic patient inputs to deliver a surgical recommendation.

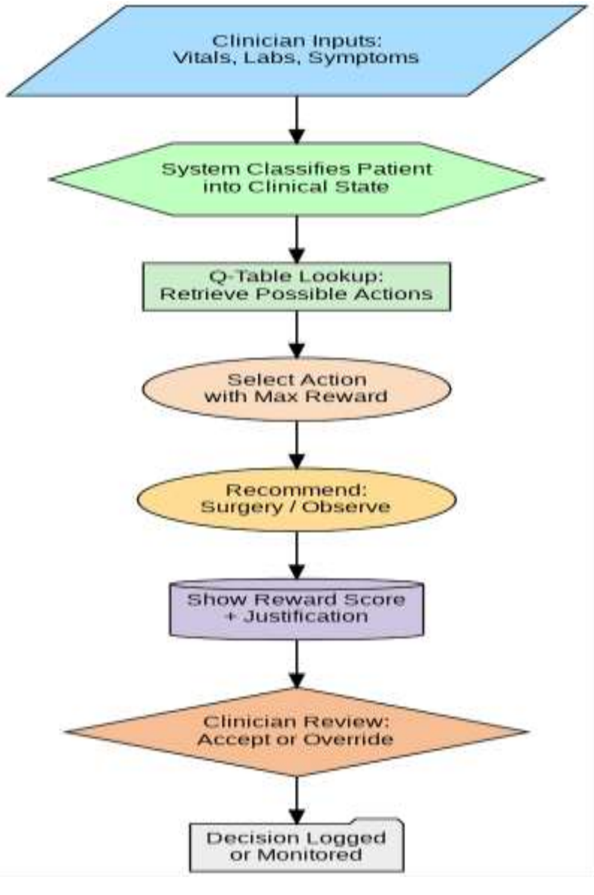

Figure 4: Workflow Diagram

To ensure the AI-driven Decision Support System's (DSS) practical application in clinical settings, a clear workflow diagram was developed. This diagram, presented in Figure 4, illustrates the step-by-step interaction between the system and the clinician user. As shown in Figure 4, the workflow commences with the clinician providing essential patient data, including parameters like blood lactate levels and oxygen saturation. The system then processes these inputs to accurately identify the corresponding clinical state of the neonate. Following this, the system consults the Q-table to determine the most appropriate action, whether it be surgery or observation. The recommended action is then presented to the clinician, accompanied by a justification based on reward values and a confidence indicator. Importantly, as detailed in Figure 4, the clinician retains the autonomy to review, accept, override, or provide feedback on the system's recommendation. This cyclical process, visually represented in Figure 4, embodies a human-in-the-loop approach, facilitating continuous improvement of the model through user interaction.

Workflow Steps:

1. **Clinician Input:** The system prompts the clinician to enter a minimal set of inputs—such as blood lactate levels, oxygen saturation, inflammatory markers (CRP or Procalcitonin), blood pressure, and signs of abdominal distension.
2. **State Identification:** The DSS maps the input values to a predefined clinical state (e.g., "S1: high lactate + distension").
3. **Action Matching:** The mapped state is then referenced in the manually constructed Q-table, where actions (surgery or observe) are compared by their assigned rewards.
4. **Recommendation Output:** The system displays the recommended action with a reward-based justification and confidence indicator (e.g., "Surgery recommended – reward: +10, high confidence").
5. **Clinician Review:** The clinician may accept, override, or comment on the recommendation. This human-in-the-loop design ensures clinical autonomy is retained and allows the model to be refined over time based on user feedback.

This process encourages transparency in decision-making and serves as a training and documentation tool, particularly useful in settings where junior clinicians or resource limitations make decision support essential(Kovalchuk et al., 2022).

5. DISCUSSION

The conceptual model developed in this study demonstrates the feasibility of applying reinforcement learning logic to one of the most critical binary decisions in neonatal care—whether or not to perform surgery. While traditional DSS models rely heavily on large datasets or patient simulations, this approach proves that meaningful, interpretable decision support systems can be developed using logic-driven structures, manual Q-tables, and expert-informed reward assignments.

One of the key strengths of this model is **explain ability**. Unlike opaque neural networks or black-box machine learning tools, this system allows clinicians to clearly see how each decision is made. Every output can be traced back to a specific state, action, and reward—thereby enhancing trust and encouraging use in clinical practice. Furthermore, its **modularity** allows the Q-table to be updated as clinical guidelines evolve or as feedback is collected from end users.

Another strength lies in its **flexibility for low-resource environments**. The model does not require high computational power, internet access, or EHR integration to function in its current form. As a result, it can be adopted as a paper-based tool, a mobile app, or a plugin for hospital information systems without significant infrastructure investment.

Despite these advantages, the system has **several limitations**. First, its logic is not derived from real-world clinical data, which may reduce its accuracy or generalisability in diverse patient populations. Second, the model currently lacks dynamic learning or simulation capabilities—it does not improve with use or adapt to rare cases unless manually modified. Third, while expert logic was used to assign rewards, these assignments are subjective and may differ across institutions or specialists.

Finally, **ethical and regulatory implications** must be considered. Although the system supports decisions, it does not replace clinical judgment. As with any AI system, ensuring that accountability, transparency, and oversight are maintained is critical to responsible deployment.

6. CONCLUSIONS

This study presented the conceptual design of an AI-driven Decision Support System that uses reinforcement learning principles to assist clinicians in making high-stakes decisions regarding surgery in neonates. Without the use of clinical data or simulation episodes, the model demonstrates how logic-based state–action mappings and reward-driven recommendations can mimic clinical reasoning and support transparent, binary decision-making.

The manually constructed Q-table, decision flowchart, and event tree provided a foundational structure for an intelligent yet explainable decision system. Scenario testing and comparative tables validated the clinical soundness of the model, and a human-in-the-loop workflow was proposed to ensure safe and ethical integration into clinical practice.

Though conceptual, the framework offers a **scalable prototype** for future development. It can be extended through expert validation, integration with hospital data systems, or enhancement using synthetic patient data and simulation. Ultimately, this work lays the groundwork for a new generation of AI-supported tools tailored to the sensitive, time-critical, and high-risk decisions that define neonatal surgical care.

REFERENCES

[1] Catania, L. J. (2021). 4 - AI applications in the business and administration of health care. In L. J. Catania (Ed.), Foundations of Artificial Intelligence in Healthcare and Bioscience (pp. 79–123). Academic Press. https://doi.org/https://doi.org/10.1016/B978-0-12-824477-7.00003-1

[2] Guez-Barber, D., & Pilon, B. (2024). Parental impact during and after neonatal intensive care admission. Seminars in Perinatology, 48(5), 151926. https://doi.org/https://doi.org/10.1016/j.semperi.2024.151926

[3] Jaile, J. C., Brady, J. D., Nelson, P., Sourour, W., Almodovar, M. C., Macicek, S., Pettitt, T. W., & Pigula, F. A. (2024). Cardiac Resynchronization Therapy for Pacing-Related Dysfunction Post Cardiac Surgery in Neonates. Annals of Thoracic Surgery Short Reports, 2(4), 825–828. https://doi.org/https://doi.org/10.1016/j.atssr.2024.05.007

[4] Jeong, H., & Kamaleswaran, R. (2022). Pivotal challenges in artificial intelligence and machine learning applications for neonatal care. Seminars in Fetal and Neonatal Medicine, 27(5), 101393. https://doi.org/https://doi.org/10.1016/j.siny.2022.101393

[5] Jyoti, J., Laing, S., Spence, K., Griffiths, N., & Popat, H. (2023). Parents' perspectives on their baby's pain management in a surgical neonatal intensive care unit: The parents' awareness and involvement in pain management (PAIN-PAM) study - Part 2. Journal of Neonatal Nursing, 29(6), 839–845. https://doi.org/https://doi.org/10.1016/j.jnn.2023.06.007

[6] Khan, S. M. (2025). Chapter 14 - AI-enabled decision support systems in clinical practice. In S. M. Khan (Ed.), Fundamentals of AI for Medical Education, Research and Practice (pp. 305–330). Academic Press. https://doi.org/https://doi.org/10.1016/B978-0-443-33584-6.00014-1

[7] Kovalchuk, S. V, Kopanitsa, G. D., Derevitskii, I. V, Matveev, G. A., & Savitskaya, D. A. (2022). Three-stage intelligent support of clinical decision making for higher trust, validity, and explainability. Journal of Biomedical Informatics, 127, 104013. https://doi.org/https://doi.org/10.1016/j.jbi.2022.104013

[8] Lakhan, A., Nedoma, J., Mohammed, M. A., Deveci, M., Fajkus, M., Marhoon, H. A., Memon, S., & Martinek, R. (2024). Fiber-optics IoT healthcare system based on deep reinforcement learning combinatorial constraint scheduling for hybrid telemedicine applications. Computers in Biology and Medicine, 178, 108694. https://doi.org/https://doi.org/10.1016/j.compbiomed.2024.108694

[9] Levin, C., Kagan, T., Rosen, S., & Saban, M. (2024). An evaluation of the capabilities of language models and nurses in providing neonatal clinical decision support. International Journal of Nursing Studies, 155, 104771. https://doi.org/https://doi.org/10.1016/j.ijnurstu.2024.104771

[10] Muntean, A., Marsland, L., Sikdar, O., Harris, C., Ade-Ajayi, N., Patel, S. B., Cook, J., Sellars, M., Greenough, A., Nicolaides, K., & Davenport, M. (2025). Neonatal Surgery for Congenital Lung Malformations: Indications, Outcomes and Association With Malignancy. Journal of Pediatric Surgery, 60(5), 162253. https://doi.org/https://doi.org/10.1016/j.jpedsurg.2025.162253

[11] Nadhir, A. M., Mounir, B., Abdelkader, L., & Hammoudeh, M. (2025). Enhancing Cybersecurity in Healthcare IoT Systems Using Reinforcement Learning. Transportation Research Procedia, 84, 113–120. https://doi.org/https://doi.org/10.1016/j.trpro.2025.03.053

[12] Onapakala, K., Varma, M. N., Nagesh, M. A., Maturi, S., Kumari, P. L., Saibaba, C. M. H., Bommisetty, J., & Donthi, R. (2024). Multi-Task Deep Learning Approaches For Named Entity Recognition, Sentiment Analysis, And Summarization In Natural Language Processing. African Journal of Biomedical Research, 27(4S), 5713–5720. https://doi.org/10.53555/AJBR.v27i4S.4670

[13] Roayaei, M., & Soltani, Z. (2025). Chapter 2 - Advancing healthcare: Reinforcement learning applications for personalized healthcare. In S. Mahajan & A. K. Pandit (Eds.), Innovations in Biomedical Engineering (pp. 33–86). Academic Press. https://doi.org/https://doi.org/10.1016/B978-0-443-30146-9.00002-2

[14] Thakre, B., Yadav, U., & Bondre, S. V. (2025). Chapter 8 - Deep reinforcement learning in healthcare and biomedical application. In S. Mahajan & A. K. Pandit (Eds.), Innovations in Biomedical Engineering (pp. 241–299). Academic Press. https://doi.org/https://doi.org/10.1016/B978-0-443-30146-9.00008-3