

IMPROVING INVENTORY MANAGEMENT OF A SOUTH AFRICAN HOSPITAL'S OPERATING THEATRE USING A MARKOV DECISION PROCESS

K. Van Dyk^{1*}, M.M. Van Zyl-Cillie² and C. Du Plessis³

1 North-West University, South Africa | 37511459@mynwu.ac.za

2 North-West University, South Africa | Maria.VanZyl@nwu.ac.za

3 North-West University, South Africa | Chantelle.Coetzee@nwu.ac.za

ABSTRACT

South African healthcare is divided into public healthcare, which is available to all South Africans, and private healthcare, which is funded by healthcare insurance or private individuals. Operating theatres are the primary driver of revenue in private hospitals. The operating theatre complex of a private hospital in South Africa experiences challenges with the on-time availability of inventory for certain surgical procedures, leading to operational inefficiencies. This study addresses inefficiencies where inventory shortages occur at the point of use during elective surgeries. Our quantitative data analysis and qualitative investigation identify that orthopaedic elective surgeries occur most frequently, and inadequate inventory policies and demand variability are the primary drivers of inventory shortages. Furthermore, discrepancies occur between items that are prepared, picked and used for orthopaedic surgeries. In this research, a Markov Decision Process (MDP) is used for orthopaedic preference card optimisation to solve the problem. The proposed MDP framework utilises Contextual Multi-Armed Bandits (CMAB) algorithms to optimise surgical preference cards through continuous learning from current and future usage patterns. This approach offers a data-driven solution for dynamic preference card

* Corresponding author

optimisation that adapts to surgeon preferences and procedural variations. Our results show improved efficiencies and availability of inventory at the point of use.

Keywords: inventory management, operating theatre, Markov Decision Process, private healthcare

1 INTRODUCTION AND PROBLEM STATEMENT

In the South African healthcare system, services are divided into two sectors: public and private. Public healthcare facilities are funded by the government, whereas private healthcare facilities operate as businesses, with patients either paying out of pocket or through medical insurance.

Hospital settings are sensitive in nature since human lives are at stake [1]. Operating theatres are a critical component of this setting, serving as the central point for both elective and non-elective surgical procedures. These environments are resource-intensive and inherently high-risk, thus necessitating exceptional levels of clinical precision, time sensitivity, adherence to safety protocols and inventory management. Efficient management of operating theatre consumables and medication is essential in supporting uninterrupted surgical workflows, minimising delays and safeguarding patient outcomes. Items must be accurately tracked and readily available to avoid cancellations or intraoperative complications due to missing or expired supplies. Inefficient inventory management in these critical settings can lead to significant challenges, including stock shortages, excessive waste and overstocking. These challenges not only increase operational costs but also compromise patient care. Additionally, many operating theatre items, especially medication and sterile supplies, require specific storage conditions and time-sensitive handling [2].

This study focuses on the operating theatre complex of a multidisciplinary private hospital located in an urban region. It is one of only two branches of the hospital group in the province and attracts private healthcare patients from across the region. Due to it being one of the only hospitals of its kind in the area, around 70 surgeries are done weekly across six operating rooms at the facility.

Similar to other operating complexes in hospitals, frequent logistical and communication problems lead to logistical challenges between surgeons, nurses, administrative personnel, medical representatives and pharmacists, and discrepancies in supply requests and availability. Manual, paper-based documentation systems are another concern, as they introduce opportunities for human error and inefficiencies, while poor traceability

undermines accountability and compliance with medical standards. This can largely be attributed to poor information flow between surgeons and other clinical operating theatre staff pre-surgery and onwards to the operating theatre dispensary, continuing post-surgery with the restocking of items used. Information on stock requirements for elective surgeries flows between surgeons, nurses, the Unit Administrative Assistant (UAA) and medical representatives from outside the hospital. The process of preparing for elective surgeries, which is scheduled weekly by the UAA, currently relies heavily on the manual preference card system, which creates multiple points where errors and inefficiencies can occur. Surgeons are required to fill out a preference card indicating all consumables and medication needed for a procedure, which is then passed through nurses and the UAA before being used to source, stock, and prepare items for surgery. In practice, however, this process is often inconsistent: preference cards are not always completed or submitted on time, and verbal requests frequently replace formal documentation. This not only delays procurement, especially when medical representatives need to supply items not held in stock, but also results in poor information flow between stakeholders.

During and after surgeries, stock control is further challenged. While nurses document usage on billing sheets, the current system also requires that wrappers of consumables be collected in a bag for post-surgery review to capture missed items. Although this method helps recover some unrecorded usage, it remains error-prone and time-consuming, contributing to discrepancies in billing and restocking. Overall, the reliance on manual preference cards and wrapper collection highlights the inefficiency of the current system and underscores the need for an integrated, data-driven solution to improve accuracy, reduce delays, and ease the administrative burden on operating theatre staff.

The objective of this research is to enhance the availability and management of medication and consumables in the operating theatre complex of a multidisciplinary private hospital in support of elective surgeries. To achieve this, a decision support system that identifies the consumables required for the most frequently performed surgeries has been developed. Given that Markov Decision Processes (MDPs) provide a robust mathematical framework for modelling sequential decision-making

under uncertainty [3], the system is built on the principles of an MDP. In addition, the system streamlines restocking and tracking processes, improving stock availability and reducing discrepancies.

In Section 2, literature that assists in elucidating the requirements for the decision support system is investigated. Section 3 explains the methodology followed to develop the system, and Section 4 briefly demonstrates the system and evaluates the system's performance. A conclusion is provided in Section 5, and recommendations are made for further development of and research into the decision support system.

2 LITERATURE OVERVIEW

This section presents the literature relevant to the control and flow of consumables and medication in operating theatres. The aim is to provide background on the operational challenges faced in surgical environments, particularly regarding inventory management and information flow. This section conducts further analysis while simultaneously gathering insights into the requirements necessary for the design phase.

A substantial body of research focuses on improving hospital operations efficiency, recognising its direct impact on patient care quality and cost containment. Recent studies highlight frameworks such as Lean Healthcare and Six Sigma as effective methodologies for streamlining hospital workflows and reducing waste [4]. Specific attention has been given to surgical suite management, where operating theatre scheduling, case cart preparation and turnover time optimisation are critical factors for overall performance.

The healthcare supply chain presents unique challenges due to demand uncertainty and critical service level requirements. Given the complexity of medical supply chains, integrated supply chain strategies that balance cost efficiency with service quality must be implemented. However, the constraints of different healthcare environments must be considered. Supply chains can be significantly supported and improved by implementing patient-oriented, agile and Lean perspectives and practices, while ensuring proper integration of patient needs [5].

Managing inventory in hospitals is characterised by several inherent

complexities. Hospitals must contend with a broad variety of item types and usage patterns, constrained storage space within departments, uncertain demand influenced by both planned and emergency procedures, and strict service level expectations often exceeding ninety-nine percent for critical medical items [6]. Traditional inventory management approaches in hospitals often rely on experience-based ordering practices. While simple to implement, these practices have been shown to result in inconsistencies in inventory parameters, unnecessary holding costs and a frequent reliance on emergency orders. Research highlights these challenges and underscores the need for systematic, data-driven inventory control methods tailored to the specific demands of healthcare environments [7], [8].

The literature also underscores the critical relationship between stockout costs and patient safety. Studies demonstrate that stockouts in healthcare can lead to delayed procedures and increased patient risk, making inventory availability a direct patient safety concern. Vendor-managed inventory (VMI) systems have been proposed as a solution, improving stock availability while reducing administrative burden [9].

Furthermore, while operating rooms are considered a significant revenue source, they are also the main source of waste and cost within hospital departments [10]. This dual nature makes effective inventory policy development critical for healthcare organisations. The procurement and inventory management challenges within the healthcare sector have subsequently been the subject of ongoing debate [11].

Systematic inventory classification methods, such as ABC analysis for item value and VED (Vital, Essential, Desirable) classification for clinical criticality, provide a framework for developing inventory policies that consider both financial impact and clinical importance. The ABC-VED matrix approach ensures that high-value items receive appropriate attention while maintaining service levels for critical, low-cost items [12]. This dual classification system ensures that high-value items receive appropriate attention, whilst critical but low-cost items maintain adequate service levels.

Inventory policies typically employ either periodic or continuous review systems, each with distinct advantages depending on item characteristics and operational constraints. Continuous review systems provide

immediate response to inventory depletion but require sophisticated monitoring infrastructure, making them optimal for high-turnover items where immediate replenishment is critical [13].

Operating theatre inventory policies must address unique challenges, including surgeon preferences, procedure variability and time-critical decision-making. Traditional inventory policies often prove inadequate for managing the complex interactions between scheduled procedures, emergency cases and individual surgeon requirements [7].

The development of procedure-specific inventory policies, supported by accurate preference cards and consumption tracking, represents a new way of healthcare inventory approaches. These policies must account for the sequential nature of surgical decisions, where intraoperative requirements may differ significantly from preoperative planning [7].

Effective surgical inventory policies require integration with scheduling systems, preference card management and real-time consumption tracking. Such policies must be sufficiently flexible to accommodate procedure variations whilst maintaining control over costs and service levels, necessitating sophisticated forecasting models that account for surgeon behaviour, procedure complexity and seasonal demand patterns [14].

Technological integration is increasingly crucial for driving operational efficiencies and particularly for more effective inventory management practices [1]. Systems capable of dynamic reorder points, automated exception reporting, and predictive demand analytics provide the responsiveness required in surgical environments. Case cart management, for example, plays a key role in perioperative efficiency and patient safety. Practical implementations illustrate measurable benefits: at Abbott Northwestern Hospital, case cart accuracy improved from 50% to 93%, labour-related additions decreased by 20%, and annual staff time savings exceeded USD 40,000, alongside improved interdepartmental cooperation [15]. Similarly, Lean Six Sigma interventions in Ireland streamlined case-cart preparation, reducing nurse preparation time by over 55% and cutting packaging waste through standardised preference cards [16].

Post-operative stock management is equally critical to maintaining surgical service continuity. Timely replenishment of instruments and consumables is essential in high-demand operating theatres [17]. Studies

demonstrate that combining data-driven inventory models, Lean methodologies, and digital systems can optimise post-operative processes [16]. For instance, stochastic base-stock models adapted to sterilised surgical instruments using discrete-time Markov chains have enabled more accurate demand forecasting and reduced overstocking, while ensuring the availability of critical items [18].

Structured change management has also proven effective in adopting more efficient approaches to inventory management, as can be seen by the implementation of Kotter's Change Model at a tertiary academic hospital. This resulted in a 37.7% reduction in overall inventory levels, an 18% decrease in tray size, and over 1,300 annual reprocessing hours saved. Surgical cancellation rates due to instrument non-availability dropped from 3.9% to 0.2% [19].

Another notable approach employed by O'Mahony et al. [20] involves the use of Lean Six Sigma and 5S methodologies. In a private hospital setting, applying 5S principles led to a 17.7% reduction in operating theatre stock levels, a 45% decrease in pre-operative nurse preparation time and a dramatic reduction in expired stock, which dropped by more than 90% in value. These outcomes demonstrate that structured organisation and workflow redesign can have a direct impact on reducing waste and improving stock reliability following surgery.

Finally, the integration of supply chain principles into post-operative care environments can enhance overall responsiveness. Research in paediatric perioperative supply chains emphasises the need for value analysis, cross-functional collaboration and evidence-based product selection to avoid stockouts and reduce waste. These strategies ensure that the right products are available when needed and that post-operative care is not compromised by inefficiencies in supply replenishment [21].

The literature also identifies two primary categories of inventory control systems: periodic review systems and continuous review systems. Periodic review systems in inventory management operate by reviewing the stock level of the inventory at set times which occur sequentially with a fixed delay. These systems are generally more suited to items with lower turnover rates and longer lead times. In contrast, continuous review systems trigger replenishment as soon as inventory levels drop below a specific threshold

and are typically applied to fast-moving, frequently used items. The application of Just-in-Time (JIT) inventory in hospitals has also gained traction, balancing potential cost savings against risks associated with stockouts of critical supplies [22].

A combined approach using two-bin Kanban systems for high-usage items and periodic review for low-frequency purchases has demonstrated improved availability, reduced overstocking, and fewer emergency orders [13]. Successful implementation requires careful classification, regular updates of ordering parameters, accurate data collection, and consideration of storage and packaging constraints.

Accurate data collection and tracking of actual demand and service levels enhance the reliability and effectiveness of inventory models and decision-making processes. Additionally, inventory management models must account for storage space constraints and package sizes to ensure the feasibility of practical implementation within the physical limitations of hospital facilities [23].

MDPs represent a powerful mathematical framework for modelling sequential decision-making problems under uncertainty. MDPs provide an appropriate technique for modelling and solving such stochastic and dynamic decisions [24]. The framework comprises states, actions, transition probabilities and rewards, enabling optimal policy determination through dynamic programming techniques.

Markov models prove useful when a decision problem involves risk that is continuous over time, when the timing of events is important, and when important events may happen more than once [25]. This characteristic makes MDPs particularly suited to healthcare environments where decisions are inherently sequential and uncertain.

The essential stochasticity of patients' response to treatment represents a major consideration for clinical decision-making in radiation therapy. Markov models serve as powerful tools to capture this stochasticity and render effective treatment decisions [26]. This demonstrates the framework's capacity to handle inherent uncertainty, making it relevant for inventory management applications where demand uncertainty exists.

Recent implementations have shown promise in personalised medicine. While MDP models enable the derivation of optimal treatment policies,

they may incur long computational times and generate decision rules that are challenging for physicians to interpret [27], highlighting both opportunities and implementation challenges.

Healthcare inventory systems present unique characteristics well-suited for MDP modelling, including demand uncertainty, critical service requirements and the need for sequential decision-making under time constraints.

Despite their theoretical appeal, MDP implementations face practical challenges. Clinical decisions often have long-term implications, and difficulties can be encountered when employing conventional decision-analytic methods to model these scenarios. These complexities require sophisticated modelling approaches and may necessitate approximate solution methods [24].

The integration of machine learning techniques with MDP frameworks offers potential for adaptive systems that continuously improve performance. Multi-agent MDP approaches present opportunities for modelling complex healthcare supply chains involving multiple decision-makers and interdependent inventory systems [28].

In a case study done by Silva-Aravena et al. [29] a digital twin framework was integrated with a Reinforcement Learning (RL) model, formulated as an MDP, to support dynamic decision-making in surgical scheduling. The state space captured real-time system variables such as patient conditions, waiting lists and operating room availability, while actions represented scheduling and rescheduling decisions. The RL agent was trained through simulation to learn policies that maximise resource efficiency and equity in patient prioritisation. The study employed simulation-based validation using historical hospital data, showing that the MDP-RL hybrid model enhanced throughput by up to 15% and reduced surgical postponements compared to traditional priority-based scheduling. The digital twin component enabled continuous adaptation as new data became available, underscoring the potential of MDP-based RL in optimising surgical workflows.

In implementing an MDP, the choice of RL algorithms is critical to ensure the system can adapt to dynamic, uncertain demand while aligning with operational constraints. RL is thus implemented in this study to

learn optimal policies through interaction with the environment. MDPs rely on solving sequential decision-making problems under uncertainty, which requires algorithms capable of efficiently learning optimal policies from historical and real-time data [30]. Given the constraints of healthcare environments, typically involving 50-300 surgical cases per algorithm training cycle, the selected algorithms must be data-efficient, interpretable and capable of learning from sparse feedback. The algorithm selection spans from traditional RL methods to specialised approaches designed for small-data scenarios with immediate reward feedback. By leveraging these algorithms, the MDP can continuously improve decision accuracy, reduce inventory waste and respond to variability in surgical demand [31].

The literature reveals a clear development from traditional experience-based inventory management toward sophisticated, data-driven approaches in operating theatre environments. While significant progress has been demonstrated through Lean methodologies, case cart optimisation and classification systems, the inherent complexity of surgical demand patterns and critical service requirements necessitates advanced modelling frameworks. The convergence of MDP theory with practical healthcare constraints presents a promising avenue for developing adaptive inventory policies that can simultaneously address cost efficiency, patient safety and operational flexibility. This foundation establishes the theoretical and practical basis for implementing intelligent inventory management systems capable of learning and adapting to the dynamic nature of surgical environments.

3 METHODOLOGY

The development of the decision support system was completed in three phases, which started with an analysis of the current practices for preparing consumables and medication for elective surgeries in the hospital's operating theatre complex. Requirements for a solution that will effectively support decisions in allocating consumables and medication for elective surgeries were then gathered and used as input into the system's design. Thereafter, the system was designed based on a similar methodology to that of the digital twin case study above.

3.1 Phase 1: Analysis

To understand the current system and ensure that the decision support system achieves the main objective of enhancing the availability and management of medication and consumables during surgery, an analysis is performed of the stock-keeping practices in place in the hospital's operating theatre complex. The first step is to develop an affinity diagram as a qualitative data analysis tool to identify categories of problems in the current medication and consumable system. Once the data is organised, relevant concepts can be constructed, and logical connections can be made [32]. These categories and a description of each can be seen in Table 1.

Table 1: Affinity diagram categories

Group	Description
Staff	Refers to the operating theatre personnel involved in surgical procedures, including scrub nurses, UAA, care workers and surgeons. This category examines their roles, responsibilities, workflow practices and how human factors such as training, workload and compliance affect consumable control and usage.
Inventory Policies	Encompasses the rules, procedures and standards that govern how surgical consumables and medication are ordered, stored, replenished and tracked.
Resource Planning	Focuses on how resources are allocated and scheduled for surgical procedures. It includes the planning of consumable needs based on case types and surgeon preferences and how well these plans align with actual usage.
Demand Variability	Addresses the unpredictability in surgical case types, volumes and intraoperative needs.
Stakeholders	Includes all individuals or groups affected by or involved in the surgical supply chain and their needs, influence and perspectives.
Communication	Examines the flow of information between departments and roles, such as from the consultation rooms to the operating theatres.

Interrelationships are drawn between categories in Table 1 to determine the main category from which most problems stem. These relationships are visually represented in Figure 1 below.



Figure 1: Interrelationship diagram

It is found that inventory policies and demand variability are the categories of problems that have the highest influence on the current system. These two factors can thus be identified as the key root causes. Staff and Communication are the most dependent on other factors and will thus most likely reflect the outcome or impact any changes might have on the system.

Poorly defined inventory policies often fail to account for the unique, high-stakes nature of surgical environments, leading to inaccurate stock levels, over-preparation of rarely used items or critical shortages of essential supplies. At the same time, demand variability, driven by differences in surgical procedures, surgeon preferences, emergency interventions and seasonal case fluctuations, makes it difficult to predict supply needs with accuracy. Without systems that can dynamically respond to these

variations, operating theatres are forced into reactive modes of operation, often resulting in last-minute retrievals, delayed surgeries, inefficient resource use and increased stress on staff. Together, these factors undermine operational efficiency, increase costs and pose risks to patient outcomes.

To further enhance our understanding of the current system, quantitative analysis is done based on the Pareto Principle, which serves as a valuable tool to concentrate efforts on the few key sources that contribute to most of the challenges [33].

The Pareto Chart in Figure 2 is an analysis of the number of cases of different elective surgical disciplines. This chart helps to identify which category of elective surgical disciplines are most frequently performed and which specific elective surgeries are to be incorporated in our decision support tool for maximum impact. Data from surgical procedures over a two-month period is analysed, and the number of cases for each surgical category is converted to indicate the percentage of total elective surgeries performed over the two months. A total of 312 surgeries were performed over the two-month period and almost 80% of all elective surgeries performed occurred in three surgical categories, namely: Orthopaedic, Urology and General Surgery.




Figure 2: Pareto chart of surgical disciplines

Further analysis was done to investigate which of the orthopaedic elective surgeries contributed to 80% of the performed surgeries. A total of 16

surgeries were identified, and the stock flow and usage of these surgeries were examined.

Through this systematic qualitative and quantitative analysis, two key root causes were identified as driving the majority of medication and consumable stockout at the point of use. These are inadequate inventory policies and demand variability. Therefore, the inability to predict and prepare for these variations leads to emergency supply retrievals and operational disruptions.

3.2 Phase 2: System requirements

The analysis from Section 3.1 and the literature discussed in Section 2 are used as a basis to determine the requirements for the proposed solution. These requirements and an explanation of each can be seen in Table 2 below. These requirements provide the guidelines for the decision support system’s design.

Table 2: Requirements for proposed solution

Requirement	Explanation
Identification of items used more than prepared for in elective surgeries	This refers to tracking and flagging when surgical procedures consume more supplies or instruments than originally anticipated in pre-operative planning. This helps identify patterns of under-preparation, improve future case planning accuracy and prevent related operational inefficiencies.
Possibility of future integration with the current AS400 system and the existing operating theatre workflows	This involves ensuring any new inventory or management system can seamlessly connect with the existing AS400 enterprise system and current operating room processes. This prevents workflow disruption and allows data to flow between systems without requiring staff to learn entirely new processes or duplicate data entry.

Real-time adaptability and responsiveness to changing demand	The system should automatically adjust inventory levels and supply allocation based on current usage patterns. This ensures adequate supplies are available when needed without excessive overstocking during quieter periods.
Ease of use	The interface and processes should be intuitive for clinical and administrative staff, requiring minimal training and reducing the likelihood of user errors. Simple navigation, clear displays and streamlined workflows are essential for busy healthcare environments.
Compliance with current inventory policies	The system must align with the existing hospital policies regarding stock rotation (FIFO), expiration date management, controlled substance tracking, vendor agreements and regulatory requirements.
Reduction of the impact of demand variability	This involves implementing strategies like safety stock optimisation or predictive analytics to minimise how fluctuations in surgical schedules or case complexity affect inventory availability and costs. The goal is to maintain service levels while controlling inventory investment.

3.3 Phase 3: System design

The decision support system design is based on an MDP that will be the foundation of preference cards for surgery types. The MDP consists of States (S), representing the different surgeons and procedures; Actions (A), such as adding, removing or modifying quantities and items on the preference cards; Transition Probabilities (P), which capture the likelihood that a change will lead to a new state based on observed usage patterns; and Rewards (R), which quantify outcomes such as cost savings, waste reduction and avoidance of delays.

For operating theatres, an MDP can be implemented by collecting historical data on actual versus requested supply usage, defining states that incorporate current card contents and contextual factors (for example, emergency item usage), modelling transitions based on usage probabilities and assigning rewards that balance efficiency gains with the risk of supply shortages. Optimisation techniques such as value iteration, policy iteration

or RL can then be used to determine the optimal set of actions.

Over time, the MDP adapts as new data become available, allowing the preference cards to reflect evolving surgical techniques, changing supply availability and surgeon-specific preferences. This results in a dynamic, data-driven optimisation process that directly reduces waste, minimises costs and supports seamless surgical workflows while also addressing demand variability and inventory policies.

A structured plan for designing and implementing the MDP approach for preference card optimisation is followed as shown in Figure 3.




Figure 3: MDP design plan

The MDP is designed in five steps. The first step, completed in Phase 1, is to collect data on the usage patterns of inventory for elective orthopaedic surgeries. Steps two and three inform the different parts of the MDP framework, with step two defining the States, Actions, Rewards and Transitions and step three choosing the RL algorithms to be used as well as simulating the framework. Steps four and five are focused on the User Interface (UI) design and the refinement of the framework when tested on surgeon-procedure combinations.

The model is framed as an MDP with states, actions, rewards and transitions. States represent a surgeon–procedure pair, while actions are item–quantity recommendations. Confidence for each recommendation is calculated by averaging probabilities from supervised learning (logistic

regression) and Thompson Sampling. Rewards capture both usage and quantity accuracy: a binary reward if the item was used, and a continuous reward measuring how close the suggested quantity was to actual usage. A weighted combination of these rewards guides learning. Transitions occur after each surgery, forming a contextual bandit setup.

In this study, the states consist of surgeon–procedure pairs, represented as $s = (\textit{surgeon}, \textit{procedure})$, which inform the model about the surgical context. The actions correspond to item and quantity recommendations, expressed as $a = (\textit{select item } i \textit{ with quantity } q)$. To generate these recommendations, probabilities from supervised learning and Thompson Sampling are averaged. The supervised learning probability is defined as

$$P_{SL} = P(R = 1 \mid \textit{surgeon}, \textit{procedure}, \textit{item}) \#(1)$$

while the Thompson Sampling probability is drawn from a Beta distribution,

$$P_{TS} \sim \textit{Beta}(1 + \textit{successes}, 1 + \textit{failures}) \#(2)$$

$$\begin{aligned} \textit{where successes} &= \sum \textit{reward} & \textit{where failures} &= n - \sum \textit{reward} \\ & & &= \textit{number of times items were observed} \end{aligned}$$

The combined confidence score used for ranking recommendations is then given by

$$\textit{Confidence} = \frac{P_{SL} + P_{TS}}{2} \#(3)$$

The reward function captures both usage and quantity accuracy. The usage reward assigns a value of one if the recommended item is used and zero otherwise,

$$R_{\textit{usage}} = \{1 \quad \textit{if item is used} \quad 0 \quad \textit{otherwise} \} \#(4)$$

while the quantity reward measures how close the recommended quantity is to the actual quantity used.

$$R_{quantity} = 1 - \frac{|q_{rec} - q_{used}|}{q_{used} + 1} \#(5)$$

The overall reward is a weighted sum of these two components,

$$R = w_1 \times R_{usage} + w_2 \times R_{quantity} \#(6)$$

where the weights are currently set to 0.5 each but can be adjusted to prioritise either usage accuracy or quantity accuracy. After each surgery, the model transitions to a new case. Since future states are only weakly affected by current actions, the process can be simplified to a contextual bandit formulation that optimises expected rewards without explicitly modelling transition probabilities. The transition probability is thus represented as

$$T(s, a, s') = P(s, a) \#(7)$$

This formulation allows the system to continuously learn and refine recommendations by balancing historical usage patterns with adaptive decision-making.

A Contextual Multi-Armed Bandit (CMAB) framework underpins the system while the Thompson Sampling adapts to surgeons' preferences by sampling from Beta distributions of item usage as logistic regression predicts probabilities for new surgeon-procedure combinations. Averaging these produces a confidence score. Quantities are estimated from historical statistics and adjusted through the reward function. This hybrid of supervised learning and CMAB improves accuracy by balancing pattern recognition with adaptive exploration.

For the pilot study, data were collected and analysed in Excel, exported as CSV, and stored in SQLite for preprocessing and model training in Python. Streamlit provides the interface for recommendations. In the long term, direct integration with the hospital's AS400 system will automate data retrieval and ensure scalability.

Items are recommended if their confidence exceeds a threshold

optimised using the F1 score, which balances precision and recall. This ensures that both correctly recommended and necessary items are prioritised. The UI in the form of a Streamlit app allows users to select a surgeon and procedure, then view ranked recommendations with suggested quantities. Items can be displayed above the threshold through the “Recommend Items” option, or all historical items through the “Show All Items” option. Tables are interactive and downloadable, supporting ease of use in the operating theatre.

4 SYSTEM DEMONSTRATION AND VALIDATION

This section will show a step-by-step run-through of the UI as well as explain the digital flow of the UI.

When the app is opened, you are welcomed to the landing page depicted below in Figure 4. No login page is necessary since users will already have had to log into their account on the AS400 system when opening the UI, if scaled and implemented in the future. On this landing page, instructions are given after which the user will select a surgeon and procedure type from the dropdown menu. Once the surgeon and procedure type are selected, the user can click on the “Recommend Items” button to show the list of items and respective quantities recommended for the surgery. The maximum number of items that are allowed in the dispensary according to inventory policies is also listed for each item.




Figure 4: UI landing page

As can be seen in Figure 5 below, two tables will then be displayed, one

indicating the surgeon-specific items and the other the procedure-specific items needed for the surgery. The items are ranked based on their confidence. The recommended items list consists of all items that have a confidence above the calculated confidence threshold.




Figure 5: UI recommended item lists

The confidence threshold is indicated at the bottom of the page, as shown in Figure 6 below, and it is evident that all items displayed when clicking on “Recommend Items” either meet or exceed this confidence threshold. All tables also have the option, when hovered over, to be downloaded as a CSV file, or to be enlarged to full screen. When enlarged to full screen, items can be rearranged in ascending or descending order based on different criteria, such as alphabetical order. Furthermore, it is also possible to search a table for particular information.

Rank	Item	Confidence	Avg Qty	Inventory Policy	
1	I174	0.71	1	1	1
2	I151	0.68	1	1	1
3	I155	0.68	1	1	1
4	I156	0.66	1	1	3
5	I165	0.64	1	1	1
6	I172	0.63	1	1	1
7	I152	0.62	0	0	0
8	I175	0.6	1	1	1
9	I170	0.6	1	1	2
10	I177	0.6	1	1	2
11	I169	0.57	1	1	1
12	I162	0.56	1	1	1
13	I178	0.56	2	2	2
14	I173	0.56	1	1	1
15	I163	0.55	0	0	0
16	I167	0.55	0	0	0
17	I161	0.54	0	0	0
18	I171	0.54	0	0	0
19	I160	0.47	1	1	2

Calculated best F1 threshold for your model: 0.33

Figure 6: UI table specific view

When clicking on the “Show All Items” button, the confidence threshold is disregarded and all items previously prepared, picked or used for the surgeon and/or procedure type are listed in a table in the same style and format as the “Recommend Items” tables.

Validation is an essential step in the development of decision-support systems, ensuring that the final solution not only complies with the specified requirements but also addresses the real-world challenges identified during the problem analysis. Various methods of validation are applied to ensure that all aspects of the MDP framework are implemented correctly.

The first method of validation is scenario testing. The aim of the scenario testing is to ensure that the recommended lists are comparable to the actual usage data available during this study. The UI is prompted to recommend items for known surgeon-procedure combinations, and this output is then compared to the data set available for this surgeon-procedure combination. An example of this is shown in Table 3 below, where a snippet of the recommended items for Surgeon S003 and Procedure I is compared to that of the actual usage data. An orange flag indicates where the recommendation

differs from the actual usage.

Table 3: Scenario testing

Item	Recommendation	Actual Usage	Flag
I045	2	2	
I053	1	1	
I061	3	3	
I095	1	1	
I156	1	1	
I162	1	1	
I163	1	0	
I166	1	0	
I174	1	0	
I177	1	1	
I178	1	1	
I188	0	0	
I189	1	1	

In the above surgeon-procedure combination a total of 27 items are listed. Three of the 27 items’ recommendations differ from the actual usage. Items I163, I166, and I174 are all surgeon-specific items, and the higher recommendation is due to surgeon S003 also using these items in many other cases and procedure types. These items were nevertheless not used for the particular case the scenario testing was done on. It is, however, better to be over-prepared for a few items than it is to be under-prepared, as over-prepared items are easier to account for post-surgery compared to under-prepared items.

To evaluate the predictive performance of the logistic regression model, k-fold cross-validation was applied to the dataset in Python. Specifically, the dataset was divided into five folds and the model was trained and tested five times, each time using a different fold as the test set while the remaining

folds were used for training. This method ensured that all data points were used for both training and evaluation, providing a robust assessment of model generalisation. The cross-validation results were as follows:

Accuracy scores across folds: [0.736, 0.7435, 0.7467, 0.7139, 0.7182]

Mean accuracy: 0.7317

Variance of accuracy: 0.000177

F1-scores across folds: [0.7112, 0.7317, 0.7309, 0.6965, 0.7006]

Mean F1 score: 0.7142

Variance of F1 score: 0.000218

These metrics indicate that the model achieves moderately high predictive performance. With the F1-scores being close to the accuracy values, a balanced reflection of both precision and recall for predicting whether a surgical item will be used is shown.

The relatively low variance in both accuracy and F1 score suggests that the model's performance is stable across different subsets of the data. However, the scores are not perfect, indicating room for improvement, particularly in the handling of variability in item usage patterns across surgeons and procedures.

Increasing the dataset size is expected to enhance model performance in several ways. It should generalise predictions better as more training examples reduce overfitting and allow the model to capture a wider range of surgeon-procedure-item interactions. Higher stability of metrics should also be achieved with larger datasets, as it typically reduces variance in cross-validation scores, making predictions more reliable. If rare items or procedures are underrepresented, additional data can provide more examples, improving F1 scores, especially for less common cases. With more observations, dynamic thresholds for classifying item usage can also be better estimated per surgeon or procedure, further enhancing recommendation precision.

The current cross-validation metrics are thus acceptable when considering the smaller data set this study works on. This then subsequently indicates that the supervised learning model works effectively by utilising logistic regression.

To assess the performance of the whole surgical item recommendation model, a combination of global classification metrics and per surgeon-procedure regression metrics was used. These metrics provide a comprehensive view of both the model's ability to predict whether an item is used and how accurately it predicts the quantities needed.

The performance of the model was evaluated using four global metrics, each providing a different perspective on predictive quality. Accuracy measures the proportion of all predictions that were correct, offering a general view of performance. The F1-score combines precision and recall into a single measure, ensuring a balance between false positives and false negatives. The AUC/ROC captures the model's ability to distinguish between items that are likely to be used and those that are not, reflecting its ranking capability. Finally, Log Loss assesses how well the predicted probabilities match actual outcomes, penalising confident but incorrect predictions more heavily. Together, these metrics give a comprehensive view of both classification performance and probability calibration. To further clarify the AUC/ROC, it is visualised by means of a plot in Figure 7 below. The larger the area under the curve the better the model is able to discriminate between items that are used and not used.




Figure 7: AUC/ROC

For predicting the quantity of items used per surgery, Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) were calculated for each surgeon–procedure combination. MAE reflects the average size of the prediction error, while RMSE emphasises larger mistakes more strongly. Lower values indicate better prediction accuracy and consistency. Table 4 below shows different examples and interpretations of these results.

Table 4: Explanation of MAE and RMSE results

Example Cases	Interpretation
Low error (S001–A: 0.38 / 0.53)	Predictions are very close to actual quantities, with small deviations.
Moderate error (S001–H: 0.60 / 0.92)	Predictions are reasonably accurate but show some variability.
High error (S001–B: 1.16 / 1.19)	Predictions struggle due to high variability in usage or limited data.

Overall, accuracy is procedure-dependent: frequent and consistent procedures achieve low MAE/RMSE, while less common or more variable ones lead to higher errors.

5 CONCLUSION

This study effectively addressed the critical challenge of medication and consumable management inefficiencies in a private hospital’s operating theatre through the development and validation of an MDP-based recommendation system. The findings demonstrate that systematic analysis of operational challenges, combined with data-driven solution development, can generate tangible improvements in healthcare supply chain management.

The research identified inadequate inventory policies and demand variability as primary drivers of operational inefficiencies within surgical environments. Through comprehensive root cause analysis employing both qualitative and quantitative methods, orthopaedic surgeries were highlighted as representing the greatest opportunity for improvement,

accounting for 33% of all procedures performed during the study period.

The proposed system aims to address communication and inventory challenges by integrating adaptive, learning-based mechanisms that personalise surgical item recommendations for each surgeon and procedure. The system employs a hybrid prediction framework combining Contextual Multi-Armed Bandits, Thompson Sampling, and Supervised Learning. This approach enables continuous learning from historical data while dynamically adapting to new usage patterns, thereby improving both the accuracy of item selection and of the recommended quantities.

An intuitive Streamlit-based UI ensures accessibility for operating theatre staff without requiring technical expertise, addressing a critical barrier to technology adoption in clinical settings. The modular system architecture, designed with potential integration into enterprise systems such as AS400 in mind, reflects practical consideration of real-world implementation challenges.

Validation results indicate that the decision-support system is functional and aligned with the requirements defined during the design phase. Scenario-based testing confirmed that the system generates recommendations comparable to historical usage data, with minor deviations attributable to surgeon-specific preferences. Cross-validation demonstrated the model's generalisation capability, with stable accuracy and F1 scores across folds. Global evaluation metrics, including AUC and log loss, further highlighted the system's ability to distinguish between used and unused items and to generate well-calibrated probability predictions. Per surgeon-procedure MAE and RMSE values confirmed that the system provides reliable quantity recommendations, although performance varied with the availability of historical data and the variability of item usage.

The study faced certain limitations, including a relatively small dataset spanning two months (312 surgeries, including 102 orthopaedic cases), a focus limited to elective orthopaedic procedures at a single hospital, and moderate validation results (73% accuracy, F1 score 0.626). Simulated testing rather than live implementation also constrained confidence in real-world performance. Future research should expand data collection to at least one year, incorporate additional surgical disciplines such as urology and general surgery, integrate the system with enterprise platforms

like AS400, and enhance predictive models by incorporating contextual factors. Despite these limitations, the study provides a robust foundation for transforming surgical supply chain management through intelligent decision-support systems.

The broader implications of this research extend to healthcare efficiency and patient safety, as improved inventory management supports uninterrupted surgical workflows and reduces the risk of procedure delays or complications due to supply shortages. Furthermore, potential cost savings from reduced waste and optimised inventory levels highlight the strategic value of data-driven supply chain management in healthcare.

Looking forward, this work establishes a framework for continued innovation in healthcare supply chain optimisation, with clear pathways for expansion across additional surgical disciplines, integration with emerging technologies, and scaling to larger healthcare systems.

6 REFERENCES

- [1] M. M. van Zyl-Cillié, D. H. van Dun, and H. Meijer, "Toward a roadmap for sustainable lean adoption in hospitals: a Delphi study," *BMC Health Serv Res*, vol. 24, no. 1, Sep. 2024, doi: 10.1186/s12913-024-11529-4.
- [2] S. A. Khokhar, "The challenges of inventory management in medical supply chain," *South Asian Journal of Operations and Logistics*, vol. 2, no. 2, pp. 1–17, Dec. 2023, doi: 10.57044/SAJOL.2023.2.2.2306.
- [3] O. Alagoz, H. Hsu, A. J. Schaefer, and M. S. Roberts, "Markov decision processes: A tool for sequential decision making under uncertainty," *Medical Decision Making*, vol. 30, no. 4, pp. 474–483, Jul. 2010, doi: 10.1177/0272989X09353194.
- [4] D. B. McLaughlin, J. R. Olson, and L. Sharma, *Healthcare Operations Management*, 4th ed. Washington, DC.: Foundation of the American College of Healthcare Executives, 2022.
- [5] V. Nabelsi and S. Gagnon, "Information technology strategy for a patient-oriented, lean, and agile integration of hospital pharmacy and medical equipment supply chains," *Int J Prod Res*, vol. 55, no. 14, pp. 3929–3945, Jul. 2017, doi: 10.1080/00207543.2016.1218082.

- [6] L. Vanbrabant, L. Verdonck, S. Mertens, and A. Caris, "Improving hospital material supply chain performance by integrating decision problems: A literature review and future research directions," *Comput Ind Eng*, vol. 180, Jun. 2023, doi: 10.1016/j.cie.2023.109235.
- [7] E. Ahmadi, D. T. Masel, A. Y. Metcalf, and K. Schuller, "Inventory management of surgical supplies and sterile instruments in hospitals: a literature review," *Health Syst*, vol. 8, no. 2, pp. 134–151, 2019, doi: 10.1080/20476965.2018.1496875.
- [8] M. Bijvank and I. F. A. Vis, "Lost-sales inventory theory: A review," *Eur J Oper Res*, vol. 215, no. 1, pp. 1–13, Nov. 2011, doi: 10.1016/j.ejor.2011.02.004.
- [9] S. R. Kochi, S. Mostofa, and Dr. S. Shaheen, "Inventory management of medical store in a selected tertiary public hospital," *International Journal of Medical Science and Clinical Research Studies (IJMSCRS)*, vol. 02, no. 01, pp. 48-55, Jan. 2022, doi: 10.47191/ijmscrs/v2-i1-09.
- [10] J. Volland, A. Fügner, J. Schoenfelder, and J. O. Brunner, "Material logistics in hospitals: A literature review," *Omega*, vol. 69, pp. 82–101, Jun. 2017, doi: 10.1016/j.omega.2016.08.004.
- [11] K. Moons, G. Waeyenbergh, and L. Pintelon, "Measuring the logistics performance of internal hospital supply chains – A literature study," *Omega*, vol. 82, pp. 205-217, Jan. 01, 2019, doi: 10.1016/j.omega.2018.01.007.
- [12] R. Nigah, M. Devnani, and A. K. Gupta, "ABC and VED analysis of the pharmacy store of a tertiary care teaching, Research and Referral Healthcare Institute of India," *J Young Pharm*, vol. 2, no. 2, pp. 201–5, Apr. 2010, doi: 10.4103/0975-1483.63170.
- [13] M. B. Kloosterman, "Increase the on-hand availability of surgical supplies at the operating room department by improving inventory management," Thesis, University of Twente, Zwolle, 2023. [Online]. Available: www.utwente.nl. [Accessed: Jul. 14, 2025].
- [14] E. Ahmadi, D. T. Masel, and S. Hostetler, "A robust stochastic decision-making model for inventory allocation of surgical supplies to reduce logistics costs in hospitals: A case study," *Oper Res Health Care*, vol. 20, pp. 33–44, Mar. 2019, doi: 10.1016/j.orhc.2018.09.001.

- [15] K. Shields and T. Voigt, "Improving the case cart process at Abbott Northwestern Hospital," *Qual Lett Healthc Lead*, vol. 6, no. 6, pp. 75–78, July 1994.
- [16] P. Egan, A. Pierce, A. Flynn, S. P. Teeling, M. Ward, and M. McNamara, "Releasing operating room nursing time to care through the reduction of surgical case preparation time: A Lean Six Sigma pilot study," *Int J Environ Res Public Health*, vol. 18, no. 22, Nov. 2021, doi: 10.3390/ijerph182212098.
- [17] A. Bhosekar, S. Ekşioğlu, T. Işık, and R. Allen, "A discrete event simulation model for coordinating inventory management and material handling in hospitals," *Ann Oper Res*, vol. 320, pp. 603–630, Jan. 2023, doi: 10.1007/s10479-020-03865-5.
- [18] A. Diamant, J. Milner, F. Quereshy, and B. Xu, "Inventory management of reusable surgical supplies," *Health Care Manag Sci*, vol. 21, no. 3, pp. 439–459, Sep. 2018, doi: 10.1007/s10729-017-9397-3.
- [19] J. Toor *et al.*, "Inventory optimization in the perioperative care department using Kotter's Change Model," *The Joint Commission Journal on Quality and Patient Safety*, vol. 48, no. 1, pp. 5–11, Jan. 2022, doi: 10.1016/j.jcjq.2021.09.011.
- [20] L. O'Mahony, K. McCarthy, J. O'Donoghue, S. P. Teeling, M. Ward, and M. McNamara, "Using Lean Six Sigma to redesign the supply chain to the operating room department of a private hospital to reduce associated costs and release nursing time to care.," *Int J Environ Res Public Health*, vol. 18, no. 21, Oct. 2021, doi: 10.3390/ijerph182111011.
- [21] J. L. Davis and R. Doyle, "Supply chain optimization for pediatric perioperative departments," *AORN J*, vol. 96, no. 6, pp. S73–S83, Dec. 2012, doi: 10.1016/j.aorn.2012.09.022.
- [22] P. K. Muppala, "JIT inventory management: reducing waste and cost in medical infrastructure," *Global Journal of Engineering and Technology Advances*, vol. 22, no. 1, pp. 149–161, Jan. 2025, doi: 10.30574/gjeta.2025.22.1.0256.
- [23] B. V. Neve and C. P. Schmidt, "Point-of-use hospital inventory management with inaccurate usage capture," *Health Care Manag Sci*, vol. 25, pp. 126–145, Mar. 2022, doi: 10.1007/s10729-021-09573-1.

- [24] A. J. Schaefer, M. D. Bailey, S. M. Shechter, and M. S. Roberts, "Modeling medical treatment using Markov Decision Processes," in *Operations Research and Health Care*, M. L. Brandeau, F. Sainfort, and W. P. Pierskalla, Eds. Boston, MA: Springer, 2005, pp. 593–612. doi: 10.1007/1-4020-8066-2_23.
- [25] F. A. Sonnenberg and J. R. Beck, "Markov models in medical decision making," *Medical Decision Making (MDM)*, vol. 13, no. 4, pp. 322–338, Dec. 1993, doi: 10.1177/0272989X9301300409.
- [26] L. B. McCullum *et al.*, "Markov models for clinical decision-making in radiation oncology: A systematic review," *J Med Imaging Radiat Oncol*, vol. 68, no. 5, pp. 610–623, Aug. 2024, doi: 10.1111/1754-9485.13656.
- [27] H. Y. Shi, C. H. Li, Y. C. Chen, C. C. Chiu, H. H. Lee, and M. F. Hou, "Quality of life and cost-effectiveness of different breast cancer surgery procedures: a Markov decision tree-based approach in the framework of Predictive, Preventive, and Personalized Medicine," *EPMA Journal*, vol. 14, pp. 457–475, Jun. 2023, doi: 10.1007/s13167-023-00326-4.
- [28] C. Yu, J. Liu, S. Nemati, and G. Yin, "Reinforcement learning in healthcare: A survey," *ACM Comput Surv*, vol. 55, no. 1, pp. 1–36, Nov. 2021, doi: 10.1145/3477600.
- [29] F. Silva-Aravena, J. Morales, and M. Jayabalan, "e-Health Strategy for Surgical Prioritization: A Methodology Based on Digital Twins and Reinforcement Learning," *Bioengineering*, vol. 12, no. 6, Jun. 2025, doi: 10.3390/bioengineering12060605.
- [30] R. S. Sutton and A. G. Barto, *Reinforcement Learning : An Introduction*, 2nd ed. Cambridge, MA: MIT Press, 2018.
- [31] R. P. Pantha, "Demand prediction and inventory management of surgical supplies," M.A. Thesis, University of Arkansas, Fayetteville, 2023. [Online]. Available: <https://scholarworks.uark.edu/etd> [Accessed: Oct 27, 2025].
- [32] U. Flick, *The SAGE Handbook of Qualitative Data Analysis*. London: Sage Publications, 2013.
- [33] T. Pyzdek, *The Lean Healthcare Handbook. Management for Professionals*, Cham, Switzerland: Springer, 2021, pp. 157–164. doi: 10.1007/978-3-030-69901-7_14.