



Paper Type: Original Article

AI, Robotics, and Markov Decision Processes: Enhancing Precision and Autonomy in Healthcare Systems

Ikram Dahamou^{1,*} , Cherki Daoui¹ 

¹Laboratory of Information Processing and Decision Support, Sultan Moulay Slimane University, Faculty of Sciences and Techniques Beni-Mellal, Morocco; ikram@agence-grand-voile.fr, cherki@example.com.

Citation:

Received: 16 November 2024	Dahamou, I., & Daoui, Ch. (2025). AI, robotics, and Markov decision processes: Enhancing precision and autonomy in healthcare systems. <i>Annals of healthcare systems engineering</i> , 2(3), 186-207.
Revised: 05 March 2025	
Accepted: 20 May 2025	

Abstract

Robotic systems are increasingly integrated into healthcare to enhance precision, autonomy, and efficiency. This study provides a systematic review of decision-making systems and control architectures for autonomous and social robots in hospitals, with a specific focus on Markov Decision Processes (MDPs) and their variants. A systematic search of ScienceDirect, SpringerLink, and IEEE databases was conducted covering the last three decades. Inclusion criteria focused on studies describing action-selection or decision-making methods for autonomous or semiautonomous healthcare robots. The Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) methodology was applied to identify, screen, and analyze relevant publications. The review identifies major application areas of MDP-based decision-making in healthcare: 1) surgical robotics, predominantly using Completely Observable Markov Decision Processes (COMDPs), 2) rehabilitation, where Partially Observable Markov Decision Processes (POMDPs) combined with deep reinforcement learning are common, 3) telemedicine, using COMDP frameworks with multi-agent coordination, 4) elderly care, leveraging POMDPs with human feedback, and 5) emergency response, applying multi-robot COMDPs enhanced with Bayesian updates. Emerging trends include hybrid COMDP–POMDP approaches and integration with machine learning for real-world deployment. MDP-based decision-making systems demonstrate strong potential to improve autonomy and adaptability in healthcare robotics. While COMDPs are effective in structured environments such as surgery, POMDPs are increasingly preferred for human-centered and uncertain contexts. Key challenges remain, including a lack of standardized benchmarks, limited clinical validation, and computational complexity. Addressing these gaps will be essential for the safe, efficient, and ethical deployment of robotic systems in healthcare.

Keywords: Markov decision processes, Healthcare robotics, Human-robot interaction, Partially observable Markov decision processes, Decision support systems, Deep reinforcement learning.

1 | Introduction

Artificial Intelligence (AI) and robotics have become increasingly prominent in healthcare over the past three decades. Hospitals, as demanding environments requiring efficiency, safety, and precision, represent ideal settings for robotic deployment. Robots are now assisting surgeons in complex procedures [1],

supporting rehabilitation [2], streamlining pharmaceutical tasks [3], and interacting socially with patients [4], thereby reducing clinical workload and improving patient care [5], [6].

Central to these advances is the development of decision-making and control models that enable robots to function autonomously in uncertain hospital environments. These models govern perception, planning, and human interaction. Recent approaches, including biologically inspired architectures [7], machine learning [8], [9], and Markov Decision Processes (MDPs) [10], have been applied to improve adaptability and resilience in healthcare robotics.

This paper provides a systematic review of decision-making systems for autonomous and social robots in hospitals, with a focus on MDPs and their variants. We analyze key architectures, highlight applications across surgery, rehabilitation, telemedicine, and elderly care, and identify persistent challenges such as computational complexity, lack of standardized benchmarks, and limited clinical validation [11], [12]. By consolidating recent advances and research gaps, this review underscores the potential of MDP-based frameworks to optimize resource utilization, support healthcare professionals, and enhance patient outcomes, particularly in contexts where autonomy and contactless interaction have become critical during and after the COVID-19 pandemic [13].

As illustrated in *Fig. 1*, robotics and AI are already impacting multiple healthcare domains, ranging from early detection and diagnosis to treatment, rehabilitation, and decision-making support. These diverse applications demonstrate the breadth of opportunities for intelligent systems in clinical practice.

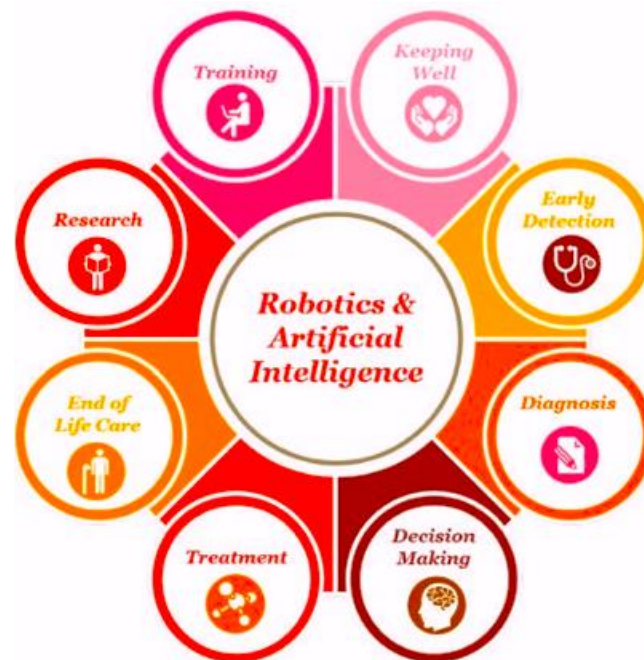


Fig. 1. Domains of application for artificial intelligence and robotics in healthcare, including research, training, diagnosis, treatment, decision-making, and end-of-life care.

2 | Rising Interest in Artificial Intelligence and Robotics in Healthcare

The growing academic interest in AI and robotics in healthcare is reflected in the continuous increase in publications dedicated to this subject. As shown in *Fig. 2*, the number of studies focusing on the theme of "AI and Robots in Value Co-Creation" has followed a strong upward trajectory. Between 2016 and 2024, publications on this topic have increased more than fourfold, highlighting the recognition of the transformative potential of these technologies in clinical practice [8].

2.1 | Factors Driving Robotic Integration in Healthcare

Several factors are driving the adoption of robotics and AI in healthcare environments:

- I. Technological advancements: the increasing sophistication of AI and robotic systems has broadened their application to complex medical tasks, ranging from minimally invasive surgery to advanced imaging and diagnostic support [9].
- II. Improved efficiency and value creation: healthcare institutions face growing pressure to enhance efficiency while reducing costs. Robotic systems can automate repetitive and time-intensive tasks, such as drug dispensing and logistics, thereby allowing healthcare staff to concentrate on patient-centered activities [12].
- III. Personalized care delivery: intelligent robots enable more personalized healthcare by monitoring patientspecific parameters (e.g., vital signs) and adapting interventions or treatment plans to individual needs [2].

As shown in *Fig. 2*, research output in this field has increased sharply since 2016, underscoring the growing recognition of AI and robotics as transformative forces in healthcare.

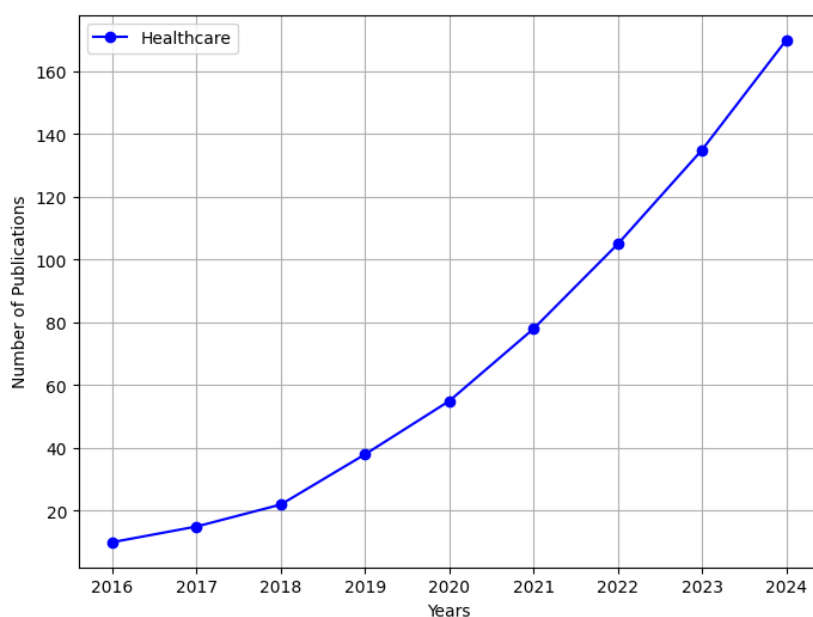


Fig. 2. Growth of publications on AI and robotics in healthcare between 2016 and 2024, showing a more than fourfold increase in less than a decade.

2.1.1 | Robots in action: examples

Here are some specific examples of robots already making a difference in healthcare:

Surgical robots: these highly precise robots assist surgeons in complex procedures, minimizing invasiveness and improving patient outcomes.

Rehabilitation robots: robots can support rehabilitation programs for patients recovering from injuries or surgeries. They can provide guided exercises, personalized training, and real-time feedback [1].

Social Assistive robots: these robots offer companionship and emotional support to patients, particularly those experiencing loneliness or isolation. They can also help with tasks like medication reminders and basic communication [4].

Telepresence robots: remotely controlled robots enable healthcare professionals to remotely examine and interact with patients in different locations, expanding access to care in underserved areas [14].

Pharmaceutical robots: robots can automate the dispensing and packaging of medications in pharmacies, reducing errors and improving efficiency [3].

3 | Artificial Intelligence as a Driver of Transformation in Healthcare

AI has emerged as a key enabler of transformation in healthcare over the past decade. Market projections indicate that the global AI healthcare sector is expected to grow substantially, from approximately \$600 million in 2014 to nearly \$150 billion by 2026¹. This convergence of AI and medicine offers the potential to revolutionize disease diagnosis, treatment planning, and patient management. By embedding AI technologies into healthcare systems, clinicians can strengthen decision-making processes, improve patient outcomes, and move toward more predictive, personalized, and efficient models of care.

The COVID-19 pandemic further underscored the need for resilient, technology-enabled healthcare infrastructures. According to the World Health Organization (WHO), approximately 15 million excess deaths were associated with the pandemic during 2020–2021 [15]. Measures such as lockdowns, while necessary for infection control, disrupted access to routine healthcare services and exacerbated mortality from Noncommunicable Diseases (NCDs), which already accounted for nearly 74% of global deaths [16]. These dual pressures highlighted the vulnerabilities of existing health systems and reinforced the role of AI-driven solutions in enabling continuity of care, optimizing resource allocation, and supporting crisis response in public health emergencies.

4 | Artificial Intelligence in Surgery: An Emerging Field

Surgery, one of the most advanced and critical domains of modern medicine, continues to pose significant risks despite centuries of progress. While technological innovations have improved surgical precision and outcomes, complications may still arise due to human error, patient-specific conditions, or equipment malfunction. These complications not only compromise recovery but also contribute to global surgical morbidity and mortality.

According to global health estimates, approximately 310 million surgical procedures are performed annually worldwide, with nearly 50 million associated with major complications and more than 1.5 million resulting in death [17]. Among the most frequent adverse events are bacterial infections and cardiovascular complications, such as perioperative heart attacks, which significantly impact patient survival and healthcare costs.

AI has the potential to mitigate such risks by supporting intraoperative decision-making, enhancing predictive analytics for complications, and enabling real-time monitoring of surgical workflows. As illustrated in *Fig. 3*, AI applications targeting infection control, cardiovascular risk prediction, and other complications represent a promising frontier in reducing surgical risks.

¹ <https://nextgeninvent.com/tag/healthcare>

AI's Potential Impact on Reducing Surgical Complications

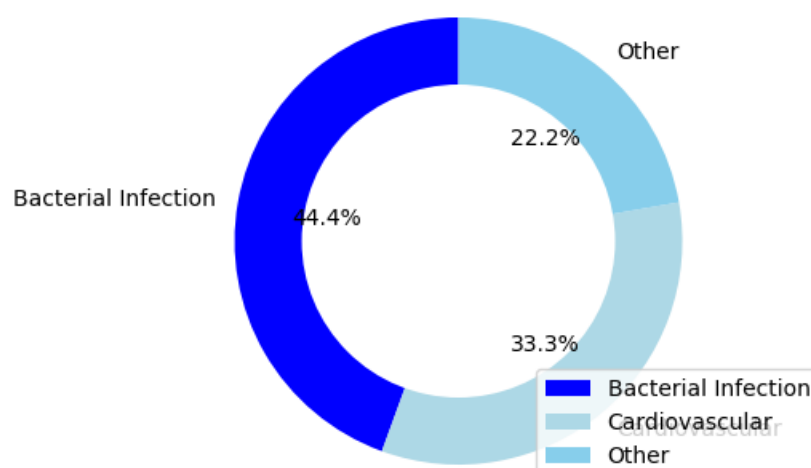


Fig. 3. AI's potential contribution to reducing surgical complications, with major focus areas including bacterial infections, cardiovascular risks, and other intraoperative or postoperative factors.

4.1 | Innovation Emerges: Artificial Intelligence as a Partner for Enhanced Safety

Given the inherent risks of surgery and the possibility of human error, exploring novel approaches to enhance safety has become imperative. AI is increasingly recognized as a critical enabler of surgical innovation, with applications extending far beyond robotic assistance. Instead, AI supports the entire continuum of surgical interventions, spanning the pre-operative, intra-operative, and post-operative phases.

Table 1 summarizes selected applications of AI in each surgical stage. Pre-operatively, AI algorithms analyze medical scans, optimize imaging sequences, and assist in accurate diagnosis and risk stratification. During surgery, intra-operative tools based on AI can track surgical instruments, support malignant tissue identification, and predict complications such as excessive bleeding. In the post-operative stage, AI is leveraged to monitor patient recovery, predict recurrence, and support tailored therapeutic strategies. These integrated applications highlight the potential of AI to reduce complications and improve patient outcomes across the surgical pathway.

Table 1. Representative applications of AI across different surgical phases.

Surgical Phase	AI Applications
Pre-operative	Analyzing medical scans with high accuracy; optimizing imaging sequences; computer-assisted diagnosis; radiomic tumor grade prediction; risk stratification.
Intra-operative	Tracking surgical instruments; malignant tissue identification; predicting intraoperative complications (e.g., bleeding); supporting workflow analysis.
Post-operative	Monitoring recovery; automated histopathological diagnosis; bioinformatics-guided treatment selection; predicting recurrence; tailoring therapies.

4.2 | Recent Advancements

4.2.1 | Artificial Intelligence and Brain Tumor Surgery

AI has shown particular promise in the management of brain tumors, where precision and safety are critical. As illustrated in *Fig. 4*, AI applications extend across pre-operative, intra-operative, and post-operative phases of neurosurgical management. These include radiomic analysis for tumor grading, intra-operative decision support for tissue identification, and post-operative predictive modeling for recurrence and therapy selection [18]. Such advancements highlight the capacity of AI to augment neurosurgeons in delivering safer, more effective, and personalized treatments.

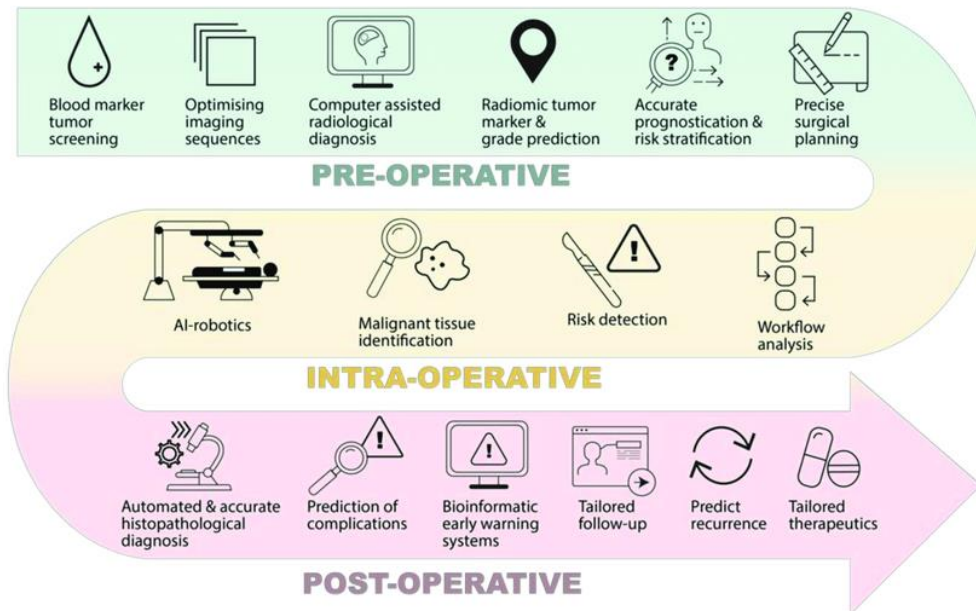


Fig. 4. Potential clinical impacts of AI in neurosurgical management of brain tumors [18].

Screening and diagnosis: machine learning algorithms have been developed to analyze routine blood tests for predicting the presence of brain tumors. In a large-scale study involving 15,176 patients with neurological conditions, a predictive model was trained and subsequently validated on a cohort of 68 patients with brain tumors and 215 control patients. The model achieved a sensitivity of 96% and a specificity of 74%, performance metrics comparable to those of conventional imaging modalities [19]. These results underscore the potential of incorporating machine learning into routine blood testing for early and non-invasive tumor detection.

AI-driven detection tools: Glioblastoma, one of the most aggressive primary brain tumors, remains challenging to diagnose and monitor due to the invasiveness of biopsies and the need for repeated imaging. A recent study applied differential scanning fluorimetry of plasma samples combined with machine learning algorithms to distinguish glioma patients ($n=84$) from healthy individuals ($n=63$). The model achieved an accuracy of 92%, demonstrating the promise of AI-enhanced plasma profiling as a non-invasive diagnostic and monitoring tool [20]. Such approaches may enable risk assessment, facilitate earlier intervention, and support treatment planning across multiple cancer types.

5 | Benefits of Artificial Intelligence and Robotics in Healthcare

The integration of AI and robotics in healthcare is revolutionizing medicine by offering a wide range of advantages. Here's a closer look at some of the most significant benefits:

5.1 | Enhanced Diagnostics and Treatment Planning

AI-powered image analysis: AI algorithms can analyze medical images with exceptional accuracy, assisting doctors in earlier detection of diseases such as cancer. This allows for more timely intervention and potentially improves treatment outcomes.

Personalized medicine: by analyzing vast amounts of patient data (medical history, genetics, etc.), AI can recommend personalized treatment plans. This approach tailors therapies to the individual's specific needs and characteristics, potentially leading to more effective treatment.

5.2 | Improved Surgical Precision and Efficiency

Robotic-assisted surgery has been widely adopted to enhance surgical precision and reduce invasiveness. Robotic arms provide surgeons with improved dexterity, tremor filtration, and high-definition visualization of the operative field. These features translate into more precise surgical maneuvers, smaller incisions, and reduced intraoperative blood loss. Clinical studies have shown that patients undergoing robotic-assisted procedures often experience shorter hospital stays and faster recovery compared to conventional surgery [1], [12].

5.3 | Increased Efficiency and Reduced Costs

AI integration in healthcare contributes substantially to efficiency gains by automating administrative and repetitive tasks. For example, algorithms can manage appointment scheduling, update electronic health records, and generate clinical reports, thereby reducing the workload of healthcare staff and enabling them to devote more time to patient-centered care [8].

Another key advantage lies in the prevention of medication errors. AI-based clinical decision support systems are capable of verifying prescriptions, identifying potential drug interactions, and ensuring correct dosage administration. Such systems not only improve patient safety but also lower the economic burden of preventable adverse drug events, which represent a significant share of healthcare expenditures [9].

5.4 | Improved Patient Care and Outcomes

Beyond cost and efficiency, AI also contributes to improved clinical outcomes by enabling earlier, more accurate, and more personalized care.

- I. AI-enabled diagnostic systems support the early detection of diseases, thereby increasing the chances of successful treatment. Early intervention is particularly critical in conditions such as cancer and cardiovascular diseases, where prognosis strongly depends on timely diagnosis [7].
- II. Remote patient monitoring has been enhanced through AI-powered wearables and sensors, which continuously track vital signs and provide real-time alerts in case of deterioration. This approach is particularly valuable for patients with chronic illnesses, allowing for timely interventions and more effective disease management [2].
- III. AI facilitates the transition toward personalized and preventive medicine. By analyzing patient-specific risk factors and large-scale clinical data, AI systems can help identify high-risk individuals and recommend tailored interventions. This predictive capacity fosters a shift from reactive to proactive healthcare delivery, thereby improving long-term outcomes [9].

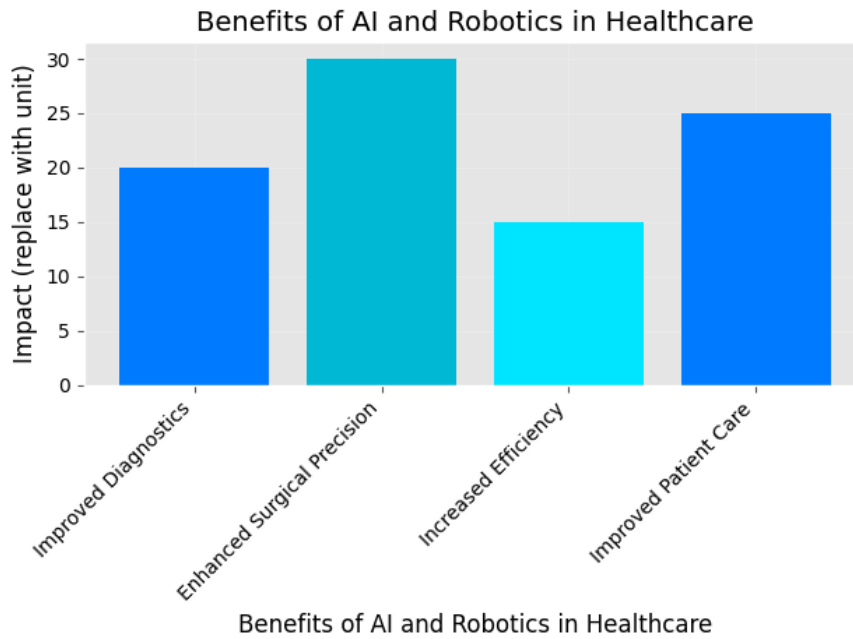


Fig. 5. Growth of Publications on AI and Robots in Healthcare (2016-2024).

6 | Materials and Methods

This systematic review was conducted in accordance with the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines. The objective was to identify, screen, and analyze relevant contributions on control architectures and Decision-Making Systems (DMS) for autonomous and social robots in healthcare.

6.1 | Search Strategy

To compile the dataset, we searched three major electronic databases widely used in engineering and healthcare research: ScienceDirect, Springer Link, and IEEE Xplore. These databases were selected due to their extensive coverage of peer-reviewed publications and global relevance.

Our search strategy relied on predefined keywords and Boolean queries related to AI, robotics, decision-making, and Human-Robot Interaction (HRI). *Table 2* summarizes the queries applied across databases and the number of initial results retrieved.

Table 2. Overview of search queries and initial results.

Database	Keywords Used	Results
ScienceDirect	“AI”, “ML”, “Robotic Systems”	329
Springer Link	“AI in Robotics”, “Automation”	17
IEEE Xplore	HRI	9
Total number of initial results		355

6.2 | Eligibility Criteria

The inclusion criteria were defined as follows: 1) description of the action-selection or decision-making method for autonomous behavior, 2) involvement of humans in the decision-making process, 3) specification of whether the method was implemented on real robots or only tested in simulation, and 4) description of the healthcare application domain.

Exclusion criteria consisted of: 1) papers not written in English, 2) studies focusing exclusively on fully teleoperated robots without autonomous features.

6.3 | Screening and Selection Process

The initial search produced 355 records. After removing duplicates and non-English articles, 290 papers remained for screening. Of these, 42 were excluded based on the eligibility criteria. A total of 257 studies met the inclusion criteria and were retained for full analysis.

The overall process of identification, screening, eligibility assessment, and inclusion is illustrated in the PRISMA flow diagram (Fig. 6).

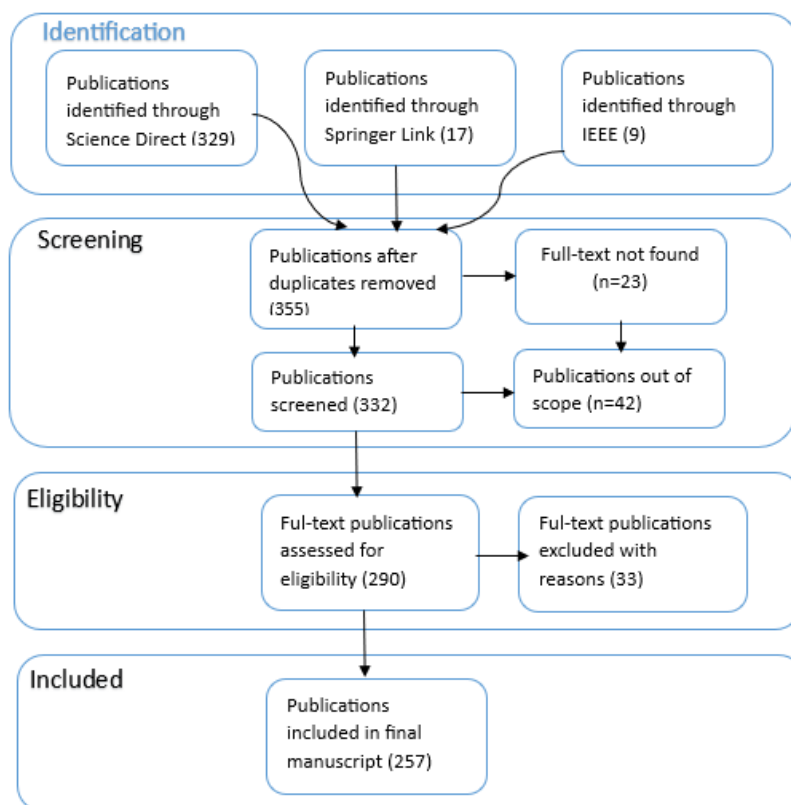


Fig. 6. Preferred reporting items for systematic reviews and meta-analyses flow diagram.

7 | Review

In the domain of decision-making and control systems for autonomous and social robots, our systematic review explores the remarkable advancements witnessed in healthcare over the past three decades. The literature reveals a wide spectrum of applications, which we classify into six healthcare subdomains: patient care, surgical robotics, rehabilitation, telemedicine, elderly care, and emergency response.

7.1 | Patient Care

Decision-making systems have played a pivotal role in enhancing patient care within hospitals and clinical settings. Applications include automated medication management, patient monitoring, and personalized care delivery. Such systems leverage AI-driven models to optimize drug administration and reduce errors, thereby elevating the quality and safety of healthcare services [8], [9].

7.2 | Surgical Robotics

Autonomous decision-making and control architectures have transformed the field of surgical robotics. Robots equipped with advanced perception and planning systems can assist surgeons in complex procedures, minimize invasiveness, and improve patient outcomes [1], [12]. These innovations demonstrate the potential of decision-making frameworks, including MDPs, to guide surgical actions under uncertainty.

7.3 | Rehabilitation

Robotic systems in rehabilitation exploit decision-making algorithms to personalize therapy and accelerate recovery. Intelligent exoskeletons and assistive devices adapt to patients' specific needs, enabling tailored rehabilitation programs that improve mobility and functionality after injury or surgery [2].

7.4 | Telemedicine

In the era of telemedicine, robots powered by decision-making systems facilitate remote consultations, diagnostics, and treatment delivery. Such systems enhance access to healthcare in underserved or remote regions, enabling physicians to extend expertise beyond traditional clinical boundaries [14].

7.5 | Elderly Care

Autonomous robots are increasingly employed to assist older adults with daily activities, medication reminders, and social interaction. Decision-making systems enable these robots to adapt to user preferences and promote independent living, ultimately improving well-being and reducing caregiver burden [4].

7.6 | Emergency Response

Robots equipped with advanced control and decision-making architectures play critical roles in disaster management, hazardous environment exploration, and search-and-rescue operations. These applications illustrate how robotic autonomy can mitigate risks, safeguard lives, and support emergency responders in time-sensitive scenarios [7].

By categorizing and synthesizing contributions across these six domains, our review provides a holistic perspective on decision-making and control systems in healthcare robotics. These systems not only augment patient care but also redefine the delivery of healthcare services, marking a new era of technological advancement.

7.7 | The History of Robotics

A significant milestone in the realm of surgery unfolded approximately a quarter-century ago when the pioneering use of a robot marked its debut in the surgical theater. This groundbreaking event featured the PUMA 200, a creation of Westinghouse Electric based in Pittsburgh, PA. Its inaugural role was in the precise placement of needles during a CT-guided brain biopsy—a watershed moment that ignited the journey of robotic surgery.

Since that pivotal juncture, the trajectory of robotic surgery has been nothing short of remarkable, evolving by leaps and bounds. This transformation can be attributed to the manifold advantages that robotics bestows, virtues often absent in traditional surgical techniques. Notably, robots bring to the table attributes such as unwavering stability, pinpoint accuracy, seamless integration with cutting-edge imaging technologies, expanded ranges of motion, and the prospect of telesurgery, among a plethora of other benefits tailored to specific surgical domains.

In the wake of the advent of robotic surgery, a wave of advancements has surged through the field. To harness the full potential of surgical robotics, it is imperative to embark on a journey through the annals of history, allowing us to chart a course towards an innovative future. This retrospective expedition entails an

exploration of the milestones and accomplishments within various surgical domains, encompassing otolaryngological, neurosurgical, gynecological, cardiothoracic, gastric, urologic, orthopedic, endoscopic, and miniature surgical specialties. By gazing into the past, we gain the insights needed to navigate the horizon of future developments that await our scrutiny.

In the pages that follow, we will embark on an illuminating journey through the annals of robotic history within the hospital environment. Our exploration will traverse various domains, each with its unique tale of robotic innovation and application. We will delve into the chronicles of assistant robots, surgical robots, and a myriad of other robotic companions that have found their purpose in healthcare settings. As we proceed, we will unveil the evolution of these remarkable machines and the profound impact they have had on the revolutionization of healthcare.

7.8 | History of Robotics Timeline:

- I. Unimate (1961): the Unimate, created by George Devol and Joseph Engelberger, is considered the first industrial robot. It was used for tasks like die casting and spot welding.
- II. PUMA (1978): the Programmable Universal Machine for Assembly (PUMA) was one of the first robots designed for more versatile assembly tasks.
- III. Robotic Arms in Space (1980s): robots like the Canadarm were used on space shuttles for tasks like satellite deployment and repair.
- IV. Roomba (2002): the Roomba, a robotic vacuum cleaner developed by iRobot, became a commercial success, popularizing the use of robots in homes.




Fig. 7. Timeline of robotics milestones: from ancient automata to modern innovation.

- V. Boston Dynamics (2010s): Boston Dynamics introduced a series of advanced robots like Spot, Atlas, and Handle, showcasing agility and mobility.

- VI. Emerging Applications (2020s): robotics continues to expand into various fields, including agriculture, logistics, and disaster response.

7.9 | Healthcare

Over the past two decades, autonomous social robots have been increasingly integrated into healthcare, with notable applications targeting children, older adults, and therapeutic support for caregivers. A particularly active research domain has been the use of robots to assist children with autism. Early studies investigated the role of social robots in gaming and educational contexts [21], while subsequent work proposed robot architectures with predefined behaviors based on observed perceptions to engage autistic children in therapeutic sessions [22].

Beyond child-centered applications, efforts have also focused on supporting caregivers during therapy. Hiolle et al. [23] introduced a baby-like robot capable of adapting its emotional responses to evaluate caregiver sensitivity. In parallel, biologically inspired approaches have been proposed, such as a hormonal system for adaptive behavior that integrates external stimuli, valence assessment, and physiological functions [24].

Other researchers have explored adaptive frameworks for chronic disease management. For instance, Lewis and Canamero [25] designed the Robin (NAO) robot to help children learn how to manage diabetes, while subsequent work modeled stress responses in robotic systems to emulate physiological mechanisms and deficits [25]. In elderly care, robots such as Mario have demonstrated the ability to assist individuals with dementia by providing cognitive support and companionship [26]. Similarly, Mini has been applied in cognitive stimulation therapies, employing a fully autonomous decision-making system based on reinforcement learning [27].

Recent advancements have broadened the scope of healthcare applications. The Pepper robot has been used autonomously to collect patient information [28], while the iCub robot has demonstrated personalized adaptation of its behavior during HRI to maximize user pleasantness [29]. Social robots have also been employed to alleviate children's pain during medical procedures [30], promote healthy diets, and deliver personalized therapies for individuals with dementia. These examples illustrate the increasing specialization of robotic applications in healthcare, emphasizing the role of decision-making systems in tailoring social interaction and therapeutic interventions.

7.10 | Humanoid Robots for Children

QTrobot, developed by LuxAI S.A., represents one of the most advanced humanoid social robots designed for education and therapy. With its expressive facial display and emotionally adaptive interaction, QTrobot has been widely used to support children with Autism Spectrum Disorder (ASD), promoting social skills, communication, and engagement [31]. Its design exemplifies the potential of decision-making systems in tailoring interaction to the specific needs of vulnerable groups.




Fig. 8. General characteristics of QTrobot.

7.11 | Pepper: Redefining Human-Robot Interaction

Pepper, created by Softbank Robotics, has become a benchmark in HRI research. Deployed in diverse domains ranging from education and healthcare to customer service, Pepper demonstrates how social robots can adaptively engage users in emotionally meaningful ways. In healthcare, Pepper has been employed to autonomously collect patient information, thereby enhancing patient experiences and supporting medical staff [28]. Its decision-making systems rely on adaptive dialogue and context-sensitive behavior, which position Pepper as a valuable platform for personalized care.

7.12 | Mario

Mario is a socially assistive robot developed for elderly care, particularly for individuals with dementia. Designed as part of the MARIO project, the robot integrates decision-making capabilities to provide companionship, cognitive stimulation, and support for independent living [26]. Mario demonstrates how decision-making systems can be adapted to long-term HRI, addressing the growing demand for assistive technologies in geriatric care.

7.13 | Social Robots for Alleviating Children's Pain

In pediatric healthcare, social robots have been developed to reduce distress and pain during medical procedures. Moerman et al. [30] introduced a robot capable of autonomously selecting distracting behaviors, such as storytelling or playful interaction, to provide comfort for children undergoing treatment. This example highlights how decision-making systems can be used not only for functional tasks but also for emotional support in sensitive healthcare scenarios.

7.14 | Comparative Analysis of Decision-Making Architectures

Our review reveals distinct trends in the application of MDPs and their variants across healthcare domains. In surgery, robotic systems such as the Da Vinci platform rely primarily on Completely Observable Markov Decision Processes (COMDPs) combined with heuristic planning [1]. Rehabilitation exoskeletons often adopt Partially Observable Markov Decision Processes (POMDPs) integrated with deep reinforcement learning, although most studies remain limited to simulation [11]. Telemedicine robots frequently use

COMDP frameworks with multi-agent coordination [8], while social robots in elderly care (e.g., Pepper) increasingly adopt POMDPs that incorporate human feedback [30]. Emergency response robots leverage multi-robot COMDP frameworks enhanced with Bayesian updates, as described in foundational reinforcement learning studies Sutton and Barto [10].

7.14.1 | Critical observations and trends

From this comparative analysis, several trends emerge: 1) COMDP frameworks dominate in structured clinical environments such as surgery and telemedicine, 2) POMDPs are increasingly adopted in contexts involving uncertainty in human behavior, such as elderly care and rehabilitation, 3) simulation-based evaluations remain prevalent, but real-world deployments are gradually increasing, and 4) integration with deep learning methods is expanding, mitigating some computational limitations of traditional POMDPs.

7.14.2 | Implications for future research

Future work should address key gaps, including the absence of standardized benchmarks for evaluating decision-making architectures across robotic platforms, the limited number of real-world clinical validations, and the need for hybrid approaches combining COMDP and POMDP frameworks. Adaptive architectures that dynamically switch between decision-making paradigms depending on task and environmental conditions may offer promising directions.

8 | Applications of Markov Decision Processes in Healthcare

MDPs have emerged as a powerful mathematical framework for modeling healthcare decision-making under uncertainty. Their versatility has led to applications across multiple domains, ranging from clinical decision support to healthcare policy design.

8.1 | Treatment Planning and Personalized Medicine

MDPs have been widely applied to optimize treatment planning, particularly for chronic and life-threatening diseases. By modeling alternative treatment options, patient responses, and uncertain outcomes, MDPs enable the development of personalized care strategies that adapt dynamically to changes in patient health states [32], [33].

8.2 | Medication Adherence and Chronic Disease Management

Adherence to prescribed therapies is a critical determinant of treatment success. MDPs can represent patient adherence behavior by incorporating probabilities of missed doses, side effects, and preferences, thereby guiding interventions to improve compliance [34]. Similarly, in chronic disease management, MDPs help design optimal long-term care strategies by considering lifestyle modifications, disease progression, and treatment effectiveness [35].

8.3 | Resource Allocation and Patient Flow Optimization

Healthcare systems face constant challenges in allocating limited resources. MDP-based models have been used to optimize staff scheduling, bed allocation, and medical equipment distribution, aiming to maximize patient outcomes while minimizing costs [36]. Applications also extend to patient flow management in hospitals and emergency departments, where MDPs inform strategies to reduce waiting times and improve service delivery [37].

8.4 | Clinical Trial Design

Another promising application is in clinical trial design. MDPs allow adaptive trial protocols by modeling patient cohort selection, treatment arm allocation, and stopping rules, ultimately improving the efficiency and ethical conduct of clinical studies [38].

8.5 | Telemedicine and Decision Support Systems

With the expansion of telemedicine, MDPs have been applied to remote monitoring and dynamic treatment adjustments based on real-time patient data. They also serve as the foundation for intelligent clinical decision support systems, enabling healthcare providers to make evidence-based and optimized decisions regarding diagnosis and care pathways [39].

8.6 | Public Health and Policy Applications

At the population level, MDPs are employed to inform vaccination strategies, screening programs, and public health interventions by modeling disease spread and intervention outcomes [37]. These models guide policymakers in designing cost-effective health strategies that balance resource constraints with societal benefits.

Summary

Overall, MDPs provide a flexible and rigorous framework for addressing healthcare challenges at both individual and system levels. Their capacity to model uncertainty and optimize sequential decisions makes them particularly suited for dynamic healthcare environments. Despite significant progress, challenges remain in translating simulation-based findings into real-world clinical deployment, highlighting the need for standardized benchmarks, large-scale datasets, and interdisciplinary collaboration.

Markov Decision Process

An MDP provides a mathematical framework for sequential decision-making under uncertainty. It is defined by a tuple (S, A, P, R, γ) , where S is the set of states, A the set of actions, P the state transition probabilities, R the reward function, and γ the discount factor.




Fig. 9. Illustrative example of a Markov decision process with three states and six possible actions.

State Space

The set of possible states is defined as $S = \{S_1, S_2, S_3\}$, representing the environment in which the agent operates.

Action Space

The set of available actions is $A = \{a_1, a_2, a_3, a_4, a_5, a_6\}$, where each action leads to transitions between states.

Transition Probabilities

The dynamics of the system are modeled by transition probabilities $P(s'|s, a)$, which specify the likelihood of reaching state s' when action a is taken in state s :

$$P(s' | s, a) = \begin{bmatrix} P(S_1 | S_1, a_1) & P(S_2 | S_1, a_2) & P(S_3 | S_1, a_3) \\ P(S_1 | S_2, a_3) & P(S_2 | S_2, a_4) & P(S_3 | S_2, a_5) \\ P(S_1 | S_3, a_6) & P(S_2 | S_3, a_6) & P(S_3 | S_3, a_6) \end{bmatrix}.$$

Reward Function

The reward function $R(s, a, s')$ defines the immediate gain of transitioning from state s to s' under action a :

$$R(s, a, s') = \begin{bmatrix} R(S_1, a_1, S_1) & R(S_1, a_2, S_2) & R(S_1, a_3, S_3) \\ R(S_2, a_3, S_1) & R(S_2, a_4, S_2) & R(S_2, a_5, S_3) \\ R(S_3, a_6, S_1) & R(S_3, a_6, S_2) & R(S_3, a_6, S_3) \end{bmatrix}.$$

Together, these components enable the computation of policies $\pi(s)$ that maximize expected cumulative rewards.

9 | Markov Decision Processes and Their Variants

A MDP is formally defined by the tuple (S, A, T, R) , where S is the set of states, A is the set of actions, $T(s, a, s')$ is the transition probability of reaching state s' after taking action a in state s , and $R(s, a)$ is the reward function specifying the immediate return associated with executing a in s [36]. This formalism provides a rigorous framework for modeling sequential decision-making problems under uncertainty, where the objective is to compute a policy $\pi(s)$ that maximizes expected cumulative rewards.

9.1 | Completely Observable Markov Decision Processes

In the classical setting, the environment state is assumed to be fully observable, and the agent can make decisions with complete knowledge of $s \in S$. COMDPs have been widely applied in structured environments such as robotic surgery and hospital logistics, where the system state can be reliably monitored [37].

9.2 | Partially Observable Markov Decision Processes

In many real-world healthcare scenarios, the state of the system cannot be directly observed. POMDPs extend the MDP framework by incorporating a set of observations O and an observation function $Z(o | s, a)$ that describes the likelihood of perceiving observation o given hidden state s and action a . Decision-making thus relies on maintaining a belief state, which is a probability distribution over possible states [40].

Solving POMDPs is significantly more challenging due to the high-dimensional belief space. Classical approaches include value iteration in belief space, point-based methods, and Monte Carlo Tree Search [32], [41]. These methods approximate optimal policies by iteratively refining decision boundaries in the belief space.

9.3 | Partially Observable Markov Decision Processes Augmentation

Recent work has explored augmenting POMDP frameworks with additional computational tools to improve performance. For instance, deep neural networks have been integrated into belief-state estimation and policy learning, enabling scalable solutions for high-dimensional state spaces [42]. Transfer learning has also been employed to adapt POMDP-based decision-making models across related tasks, enhancing generalization in healthcare and robotic applications [43].

9.4 | Applications

MDPs and their variants have found numerous applications in robotics and healthcare, including surgical planning, rehabilitation, elderly care, and resource optimization. In particular, POMDP-based models have

been employed to handle uncertainty in HRI and patient monitoring, where hidden states (e.g., patient discomfort, fatigue) must be inferred from limited observations [37].

Overall, MDPs and their extensions provide a versatile and powerful set of tools for decision-making under uncertainty. Their integration with machine learning, particularly deep reinforcement learning, continues to expand their applicability in healthcare robotics and beyond.

10 | Completely Observable Markov Decision Process

A COMDP is a special case of an MDP where the agent has full knowledge of the system state at every time step [36]. This means that all relevant variables describing the environment are directly accessible to the agent, allowing for precise action selection and policy optimization.

Formally, a COMDP is defined by the tuple (S, A, P, R, γ) , where:

- I. S is the set of possible states of the environment.
- II. A is the set of actions available to the agent.
- III. $P(s'|s, a)$ is the transition probability function specifying the likelihood of reaching state s' when taking action a in state s .
- IV. $R(s, a)$ is the reward function mapping state–action pairs to scalar rewards.
- V. $\gamma \in [0, 1]$ is the discount factor weighting the importance of future rewards relative to immediate ones.

The objective of the agent is to learn a policy $\pi(s)$ that maximizes the expected cumulative discounted reward:

$$V^\pi(s) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t) \mid s_0 = s, \pi \right].$$

COMDPs are widely used in reinforcement learning, where the assumption of full observability allows agents to efficiently estimate value functions and policies from interaction data [10]. In robotics, COMDPs are suitable for structured environments such as surgical robots or automated logistics, where sensors provide accurate and complete state information [37].

11 | Partially Observable Markov Decision Process

In many real-world applications, the system state cannot be observed directly. POMDPs extend the MDP framework by incorporating incomplete and uncertain information [40].

A POMDP is defined by the tuple $(S, A, T, R, Z, O, \gamma, b_0)$, where $S, A, T, R,$ and γ are defined as in COMDPs, and additionally:

- I. Z is the set of possible observations available to the agent.
- II. $O(o|s, a)$ is the observation function specifying the probability of observing $o \in Z$ given state s and action a .
- III. b_0 is the initial belief state, i.e., a probability distribution over S representing the agent's uncertainty about the initial state.

Because the agent cannot directly observe the true state, it maintains a belief state $b(s)$ updated after each action–observation pair using Bayes' rule. The policy $\pi(b)$ maps belief states to actions to maximize expected cumulative rewards [32].

Solving POMDPs is computationally challenging due to the continuous belief space. Approximation methods include particle filters, point-based value iteration, and Monte Carlo tree search [41]. POMDPs have been applied successfully in robotics, autonomous driving, healthcare monitoring, and assistive systems, where uncertainty in perception and human behavior must be taken into account [44], [42].

12 | Completely Observable Markov Decision Processes in Robotics

COMDPs have been widely applied in robotics when the full state of the environment is accessible through reliable sensing and system modeling. In these scenarios, the state space typically includes the robot's position, velocity, and sensor readings, while the action space corresponds to motion primitives or high-level commands [36], [10].

The reward function is designed to encode task-specific goals such as reaching a target, avoiding collisions, or minimizing energy consumption. Dynamic programming, Monte Carlo methods, and reinforcement learning have been successfully applied to compute optimal policies in COMDPs [37]. These approaches have improved robot performance in applications such as autonomous navigation, warehouse automation, and surgical robotics, where the environment can be modeled with sufficient accuracy [1], [3].

13 | Partially Observable Markov Decision Processes in Robotics

In contrast, POMDPs are increasingly employed in robotics when the true system state cannot be directly observed. Instead, the robot must rely on noisy and incomplete sensory inputs and maintain a belief state over possible states of the environment [40].

This framework has been particularly useful for navigation in dynamic or crowded environments, where obstacles or humans may unpredictably occlude the robot's sensors [44]. By updating its belief state through Bayes' rule after each action–observation cycle, the robot can make decisions that maximize expected cumulative rewards under uncertainty.

Solving POMDPs is computationally demanding, and algorithms such as particle filters, point-based value iteration, and Monte Carlo Tree Search have been developed to enable tractable solutions [41], [42]. Applications include autonomous driving, assistive robots for healthcare, and multi-robot coordination in uncertain environments [37].

14 | Human-Robot Interaction

HRI investigates how humans and robots communicate, collaborate, and build trust in shared environments. It is inherently multidisciplinary, drawing from robotics, computer science, psychology, and sociology [45]. Key challenges in HRI include natural language understanding, gesture recognition, adaptive decision-making, and ensuring safe, ethical, and acceptable robot behavior in human-centered contexts [46].

14.1 | Applications in Healthcare

In healthcare, HRI has been applied to a variety of contexts, including surgical assistance, rehabilitation, social companionship, and pediatric support [4]. Robots can help children with autism through socially assistive interactions, provide companionship to elderly individuals with dementia, and assist medical staff with repetitive tasks such as logistics and cleaning.

14.2 | Human-Robot Interaction during the COVID-19 Pandemic

The COVID-19 pandemic accelerated the adoption of HRI in hospitals. Robots were deployed for contactless delivery of food and medication, cleaning and disinfection of rooms, and telepresence consultations, reducing infection risk for healthcare workers [7]. Other robots were used for remote vital sign monitoring, enabling continuous patient care without direct physical contact. These applications demonstrate how AI-powered HRI can improve safety and resilience in healthcare systems during crises.

14.3 | Link with Human-Computer Interaction

Human-Computer Interaction (HCI) shares many principles with HRI, focusing on the design of intuitive, user-centered interfaces. In healthcare robotics, HCI contributes to developing ergonomic and accessible control systems that allow medical professionals to collaborate effectively with robots. By integrating HCI principles, HRI research ensures that robotic systems are not only functional but also acceptable and usable in real-world healthcare settings [47].

15 | Ethical and Privacy Considerations in Healthcare Robotics

While healthcare robotics and AI promise substantial benefits, their adoption raises critical ethical and privacy concerns that must be carefully addressed.

Data Privacy and Confidentiality: robotic systems often rely on continuous patient monitoring and the processing of sensitive health data. This raises risks of data breaches and unauthorized use of personal health information. Robust data protection protocols, encryption standards, and compliance with regulations such as GDPR and HIPAA are essential to safeguard patient privacy [48], [49].

Bias and Fairness: machine learning algorithms used in healthcare robotics may inherit biases from training data, leading to unequal treatment across different demographic groups. This not only undermines fairness but may exacerbate existing healthcare disparities. Mitigating bias requires diverse datasets, fairness-aware algorithms, and continuous auditing [50].

Accountability and Responsibility: determining liability in cases where robotic systems cause harm remains a major challenge. Questions arise as to whether responsibility lies with manufacturers, healthcare providers, or the algorithms themselves. Establishing transparent accountability frameworks and legal standards is crucial [51].

16 | Challenges and Limitations of Artificial Intelligence and Robotics in Healthcare

16.1 | Data and Bias

AI and robotic systems depend heavily on large-scale, high-quality datasets. Limitations in data availability or representativeness can compromise performance. Moreover, privacy regulations may restrict data sharing, slowing innovation [9].

16.2 | Technical Challenges

A persistent challenge is the lack of explainability in complex deep learning models, often considered “black boxes”. Without interpretability, it is difficult for clinicians to trust AI-driven recommendations. Interoperability with existing hospital infrastructures also remains a barrier to large-scale deployment [7].

16.3 | Ethical and Social Concerns

Automation raises fears of job displacement, particularly in repetitive tasks. At the same time, unequal access to advanced robotic technologies may exacerbate healthcare inequities. Ensuring affordability, accessibility, and retraining programs for healthcare workers is essential for sustainable adoption [13].

16.4 | Safety and Liability

Robotic systems deployed in surgical or diagnostic contexts must be rigorously validated to avoid malfunction risks. Furthermore, liability for errors involving autonomous systems remains a contested issue. Clear regulatory guidelines and safety certifications are required before widespread adoption [46].

17 | Conclusion

The integration of AI and robotics in healthcare represents a paradigm shift toward more personalized, efficient, and precise medical care. Applications range from minimally invasive robotic surgery and rehabilitation to AI-enabled diagnostics and patient monitoring. These innovations have the potential to enhance patient outcomes, reduce clinician workload, and optimize resource utilization.

However, realizing this potential requires addressing significant challenges. Ethical considerations such as fairness, transparency, and accountability must guide system design. Technical barriers, including explainability and interoperability, need innovative solutions, while equitable access and workforce reskilling must remain policy priorities.

In conclusion, AI and robotics in healthcare should not be viewed as a replacement for human expertise but rather as a complement that augments clinicians' capabilities. With careful governance, responsible innovation, and inclusive deployment, these technologies can transform healthcare delivery while safeguarding patient rights, equity, and trust.

References

- [1] Lanfranco, A. R., Castellanos, A. E., Desai, J. P., & Meyers, W. C. (2004). Robotic surgery: A current perspective. *Annals of surgery*, 239(1), 14–21. <http://dx.doi.org/10.1097/01.sla.0000103020.19595.7d>
- [2] Loh, H. W., Ooi, C. P., Seoni, S., Barua, P. D., Molinari, F., & Acharya, U. R. (2022). Application of explainable artificial intelligence for healthcare: A systematic review of the last decade (2011–2022). *Computer methods and programs in biomedicine*, 226, 107161. <https://doi.org/10.1016/j.cmpb.2022.107161>
- [3] Paul, D., Sanap, G., Shenoy, S., Kalyane, D., Kalia, K., & Tekade, R. K. (2021). Artificial intelligence in drug discovery and development. *Drug discovery today*, 26(1), 80–93. <https://doi.org/10.1016/j.drudis.2020.10.010>
- [4] Broadbent, E., Stafford, R., & MacDonald, B. (2009). Acceptance of healthcare robots for the older population: Review and future directions. *International journal of social robotics*, 1(4), 319–330. <https://doi.org/10.1007/s12369-009-0030-6>
- [5] Alsabah, M., Naser, M. A., Albahri, A. S., Albahri, O. S., Alamoodi, A. H., Abdhussain, S. H., & Alzubaidi, L. (2025). A comprehensive review on key technologies toward smart healthcare systems based IoT: technical aspects, challenges and future directions. *Artificial intelligence review*, 58(11), 343. <https://doi.org/10.1007/s10462-025-11342-3>
- [6] Olaronke, I., Ojerinde, O., & Ikono, R. (2017). State of the art: A study of human-robot interaction in healthcare. *International journal of information engineering and electronic business*, 3(3), 43–55. <https://doi.org/10.5815/ijieeb.2017.03.06>
- [7] Yu, K.-H., Beam, A. L., & Kohane, I. S. (2018). Artificial intelligence in healthcare. *Nature biomedical engineering*, 2(10), 719–731. <https://doi.org/10.1038/s41551-018-0305-z>
- [8] Jiang, F., Jiang, Y., Zhi, H., Dong, Y., Li, H., Ma, S., ... , & Wang, Y. (2017). Artificial intelligence in healthcare: Past, present and future. *Stroke and vascular neurology*, 2(4), 230-243. <https://doi.org/10.1136/svn-2017-000101>
- [9] Topol, E. (2019). *Deep medicine: How artificial intelligence can make healthcare human again*. Hachette UK. <https://www.amazon.com/Deep-Medicine-Artificial-Intelligence-Healthcare/dp/1541644638>
- [10] Sutton, R. S., Barto, A. G., & others. (1998). *Reinforcement learning: An introduction* (Vol. 1). MIT press Cambridge. <https://web.stanford.edu/class/psych209/Readings/SuttonBartoIPRLBook2ndEd.pdf>
- [11] Birkhoff, D. C., van Dalen, A. S. H. M., & Schijven, M. P. (2021). A review on the current applications of artificial intelligence in the operating room. *Surgical innovation*, 28(5), 611–619. <https://doi.org/10.1177/1553350621996961>
- [12] Yang, G. Z., Cambias, J., Cleary, K., Daimler, E., Drake, J., Dupont, P. E., ... , Taylor, R. H. (2017). Medical robotics—Regulatory, ethical, and legal considerations for increasing levels of autonomy. *Science Robotics*, 2(4). <https://doi.org/10.1126/scirobotics.aam8638>
- [13] Köbis, N., Starke, C., & Rahwan, I. (2022). The promise and perils of using artificial intelligence to fight corruption. *Nature machine intelligence*, 4(5), 418–424. <https://doi.org/10.1038/s42256-022-00489-1>

- [14] Kristoffersson, A., Coradeschi, S., & Loutfi, A. (2013). A review of mobile robotic telepresence. *Advances in human-computer interaction*, 2013, 1–17. <https://doi.org/10.1155/2013/902316>
- [15] World Health Organization (WHO). (2021). 14.9 million excess deaths associated with the COVID-19 pandemic in 2020 and 2021. <https://www.who.int/news/item/05-05-2022-14.9-million-excess-deaths-were-associated-with-the-covid-19-pandemic-in-2020-and-2021>
- [16] World Health Organization (WHO). (2021). *Noncommunicable diseases: Mortality*. <https://www.who.int/data/gho/data/themes/topics/noncommunicable-diseases-mortality>
- [17] Weiser, T. G., Regenbogen, S. E., Thompson, K. D., Haynes, A. B., Lipsitz, S. R., Berry, W. R., & Gawande, A. A. (2008). An estimation of the global volume of surgery: A modelling strategy based on available data. *The lancet*, 372(9633), 139–144. [https://doi.org/10.1016/S0140-6736\(08\)60878-8](https://doi.org/10.1016/S0140-6736(08)60878-8)
- [18] Williams, S., Layard Horsfall, H., Funnell, J. P., Hanrahan, J. G., Khan, D. Z., Muirhead, W., ... , & Marcus, H. J. (2021). Artificial intelligence in brain tumour surgery – An emerging paradigm. *Cancers*, 13(19), 1–25. <https://doi.org/10.3390/cancers13195010>
- [19] Podnar, S., Kukar, M., Gunčar, G., Notar, M., Gošnjak, N., & Notar, M. (2019). Diagnosing brain tumours by routine blood tests using machine learning. *Scientific reports*, 9(1), 14481. <https://doi.org/10.1038/s41598-019-51147-3>
- [20] Eyraud, R., Ayache, S., Tsvetkov, P. O., Kalidindi, S. S., Baksheeva, V. E., Boissonneau, S., ... , & Tabouret, E. (2023). Plasma nanodsf denaturation profile at baseline is predictive of glioblastoma EGFR status. *Cancers*, 15(3), 1–9. <https://doi.org/10.3390/cancers15030760>
- [21] Autenbahn, K., & Billard, A. (2022). The impact of interactive robots on autism therapy. *Proceedings of the IEEE international symposium on robot and human interactive communication*. IEEE. <https://doi.org/10.1109/ROMAN.2002.1138485>
- [22] Feil-Seifer, D., & Matarić, M. (2009). Toward socially assistive robotics for augmenting interventions for children with autism spectrum disorders. In *Experimental robotics* (Vol. 54, pp. 201–210). Springer Tracts in Advanced Robotics. https://doi.org/10.1007/978-3-642-00196-3_24
- [23] Hiolle, A., Lewis, M., & Cañamero, L. (2014). Arousal regulation and affective adaptation to human responsiveness by a robot that explores and learns a novel environment. *Frontiers in neurobotics*, 8. <https://doi.org/10.3389/fnbot.2014.00017>
- [24] Lones, J., Lewis, M., & Cañamero, L. (2018). A hormone-driven epigenetic mechanism for adaptation in autonomous robots. *IEEE transactions on cognitive and developmental systems*, 10(2), 445–454. <https://doi.org/10.1109/TCDS.2017.2775620>
- [25] Lewis, M., & Canamero, L. (2016). Hedonic quality or reward? A study of basic pleasure in homeostasis and decision making of a motivated autonomous robot. *Adaptive behavior*, 24(5), 267–291. <https://doi.org/10.1177/1059712316666331>
- [26] Casey, D. (2016). MARIO Managing active and healthy aging with use of caring service robots. <https://researchrepository.universityofgalway.ie/server/api/core/bitstreams/f2032b66-6112-4857-9fcb-f56341a4261b/content>
- [27] Kelasidi, E., Moe, S., Pettersen, K. Y., Kohl, A. M., Liljebäck, P., & Gravidahl, J. T. (2019). Path following, obstacle detection and obstacle avoidance for thrusted underwater Snake robots. *Frontiers in robotics and ai*, 6. <https://doi.org/10.3389/frobt.2019.00057>
- [28] Pandey, A. K., & Gelin, R. (2018). A mass-produced sociable humanoid robot: Pepper: The first machine of its kind. *IEEE robotics & automation magazine*, 25(3), 40–48. <https://doi.org/10.1109/MRA.2018.2833157>
- [29] Tanevska, A., Rea, F., Sandini, G., Cañamero, L., & Sciutti, A. (2020). A socially adaptable framework for human-robot interaction. *Frontiers in robotics and AI*, 7. <https://doi.org/10.3389/frobt.2020.00121>
- [30] Moerman, C. J., van der Heide, L., & Heerink, M. (2019). Social robots to support children’s well-being under medical treatment: A systematic state-of-the-art review. *Journal of child health care : for professionals working with children in the hospital and community*, 23(4), 596–612. <https://doi.org/10.1177/1367493518803031>
- [31] Pairet, È., Ardón, P., Liu, X., Lopes, J., Hastie, H., & Lohan, K. S. (2019). A digital twin for human-robot interaction. 2019 14th ACM/IEEE international conference on human-robot interaction (HRI) (pp. 372). IEEE. <https://doi.org/10.1109/HRI.2019.8673015>

- [32] Hauskrecht, M. (2000). Value-function approximations for partially observable markov decision processes. *Journal of artificial intelligence research*, 13, 33–94. <https://doi.org/10.1613/jair.678>
- [33] Shortreed, S. M., Laber, E., Lizotte, D. J., Stroup, T. S., Pineau, J., & Murphy, S. A. (2011). Informing sequential clinical decision-making through reinforcement learning: An empirical study. *Machine learning*, 84(1), 109–136. <https://doi.org/10.1007/s10994-010-5229-0>
- [34] Blankart, K. E., & Lichtenberg, F. R. (2020). Are patients more adherent to newer drugs? *Health care management science*, 23(4), 605–618. <https://doi.org/10.1007/s10729-020-09513-5>
- [35] Shechter, S., Bailey, M., & Schaefer, A. (2008). The optimal time to initiate HIV therapy under ordered health states. *Operations research*, 56(1), 20–33. <https://doi.org/10.1287/opre.1070.0480>
- [36] Puterman, M. L. (1994). *Markov decision processes: Discrete stochastic dynamic programming*. John Wiley & Sons. <https://doi.org/10.1002/9780470316887>
- [37] Alagoz, O., Hsu, H., Schaefer, A. J., & Roberts, M. S. (2010). Markov decision processes: A tool for sequential decision making under uncertainty. *Medical decision making*, 30(4), 474–483. <https://doi.org/10.1177/0272989X09353194>
- [38] Thall, P. F., & Wathen, J. K. (2007). Practical Bayesian adaptive randomisation in clinical trials. *European journal of cancer*, 43(5), 859–866. <https://doi.org/10.1016/j.ejca.2007.01.006>
- [39] Liu, S., See, K. C., Ngiam, K. Y., Celi, L. A., Sun, X., & Feng, M. (2020). Reinforcement learning for clinical decision support in critical care: Comprehensive review. *Journal of medical internet research*, 22(7), e18477. <https://doi.org/10.2196/18477>
- [40] Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1), 99–134. [https://doi.org/10.1016/S0004-3702\(98\)00023-X](https://doi.org/10.1016/S0004-3702(98)00023-X)
- [41] Silver, D., & Veness, J. (2010). Monte-carlo planning in large pomdps. *Advances in neural information processing systems* (pp. 2164–2172). Curran Associates, Inc. <https://www.researchgate.net/publication/221620445>
- [42] Igl, M., Zintgraf, L., Le, T. A., Wood, F., & Whiteson, S. (2018). Deep variational reinforcement learning for pomdps. *Proceedings of the 35th international conference on machine learning* (Vol. 80, pp. 2117–2126). PMLR. <https://proceedings.mlr.press/v80/igl18a.html>
- [43] Daumé III, H. (2007). Frustratingly easy domain adaptation. *Proceedings of the 45th annual meeting of the association of computational linguistics* (pp. 256–263). Association for Computational Linguistics (ACL). <https://aclanthology.org/P07-1033.pdf>
- [44] Pineau, J., Gordon, G., Thrun, S. (2003). Point-based value iteration: An anytime algorithm for POMDPs. *Proceedings of the 18th International joint conference on artificial intelligence (IJCAI)* (Vol. 3, pp. 1025–1032). Morgan Kaufmann Publishers Inc. <http://www.thrun.org/papers/Pineau03a.pdf>
- [45] Goodrich, M. A., & Schultz, A. C. (2008). Human-robot interaction: A survey. *Foundations and trends in human-computer interaction*, 1(3), 203–275. <https://www.emerald.com/fthci/article/1/3/203/1321642>
- [46] Sheridan, T. B. (2016). Human-robot interaction: status and challenges. *Human factors*, 58(4), 525–532. <https://doi.org/10.1177/0018720816644364>
- [47] Dautenhahn, K. (2007). Socially intelligent robots: Dimensions of human-robot interaction. *Philosophical transactions of the royal society b: Biological sciences*, 362(1480), 679–704. <https://royalsocietypublishing.org/rstb/article/362/1480/679/20947>
- [48] Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big data & society*, 3(2). <https://doi.org/10.1177/2053951716679679>
- [49] Floridi, L., Cowls, J., Beltramini, M., Chatila, R., Chazerand, P., Dignum, V., ... , & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, Risks, Principles, and Recommendations. *Minds and machines*, 28(4), 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
- [50] Rajkomar, A., Hardt, M., Howell, M. D., Corrado, G., & Chin, M. H. (2018). Ensuring fairness in machine learning to advance health equity. *Annals of internal medicine*, 169(12), 866–872. <https://doi.org/10.7326/M18-1990>
- [51] Matthias, A. (2004). The responsibility gap: Ascribing responsibility for the actions of learning automata. *Ethics and information technology*, 6(3), 175–183. <https://doi.org/10.1007/s10676-004-3422-1>