From Reactive to Predictive: How Artificial Intelligence May Transform Perioperative Anaesthesia Care

Abstract

Artificial intelligence (AI) is increasingly being explored across healthcare and may significantly influence perioperative anaesthesia practice. Anaesthesia is particularly suited to AI-supported approaches because it involves continuous physiological monitoring, complex decision-making and the interpretation of large volumes of real-time clinical data. Applications are being developed to support preoperative risk assessment, intra-operative monitoring and postoperative management [1].

Emerging machine-learning models have demonstrated the capacity to predict selected perioperative complications, haemodynamic instability and postoperative pain or opioid requirements. These capabilities create opportunities for earlier recognition of risk and more individualised clinical planning. However, much of the supporting evidence remains retrospective, single-centre or focused on predictive accuracy rather than demonstrated improvements in patient-centred outcomes [2].

Safe implementation also presents significant challenges, including algorithmic bias, limited explainability, privacy concerns, medico-legal accountability, automation bias, alert fatigue and inappropriate reliance on decision-support systems [3].

This narrative review examines current and emerging applications of AI across the perioperative pathway and explores the human factors, ethical and governance considerations accompanying its adoption. It proposes the Predict–Interpret–Adapt framework: artificial intelligence predicts risk, clinicians interpret its significance within the context of the individual patient, and multidisciplinary teams adapt care accordingly. The greatest contribution of AI may therefore lie not in autonomous decision-making, but in strengthening clinical anticipation while preserving professional judgement, effective teamwork and human connection. With rigorous validation, thoughtful clinical integration and strong governance, AI may extend the established anticipatory capabilities of anaesthesia towards increasingly predictive, personalised and proactive perioperative care [4].

Keywords: *Artificial intelligence; Machine learning; Anaesthesia; Perioperative medicine; Predictive analytics; Clinical decision support; Human-AI collaboration; Patient safety; Acute pain service; Pre-anaesthesia assessment.*

Introduction

Artificial intelligence (AI) has emerged as one of the most significant technological developments in modern healthcare. Advances in machine learning, predictive analytics and computational processing have enabled AI systems to identify complex relationships within large datasets and generate information that may support clinical decision-making. Although clinical implementation remains uneven and many applications are still being evaluated, AI is increasingly being explored across medical imaging, diagnosis, clinical prediction and personalised care [5][6].

Anaesthesia is particularly suited to AI-supported approaches because it is data-rich, time-sensitive and technologically complex. Anaesthesia professionals continuously integrate haemodynamic trends, respiratory parameters, laboratory results, medication effects and procedural information while making time-critical decisions under conditions of uncertainty. This work already depends heavily upon anticipation: clinicians prepare for foreseeable complications, monitor evolving physiology and adapt care as patient and procedural conditions change [7][8].

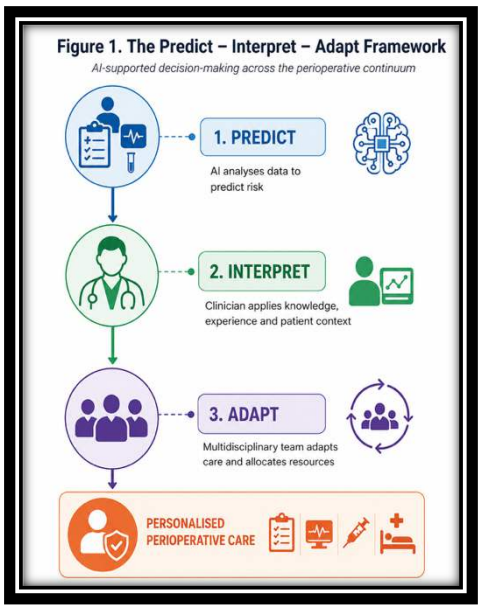
However, many conventional monitoring systems remain threshold-based, drawing attention to physiological disturbance once a predefined limit has been reached. AI may extend anaesthesia's existing anticipatory

[1] Topol, 2019; Rajkomar et al., 2019; Hashimoto et al., 2018; Loftus et al., 2020
[2] Kendale et al., 2018; Wijnberge et al., 2020; Nair et al., 2020; Arina et al., 2024; Mehta et al., 2024
[3] Australian Commission on Safety and Quality in Health Care, 2025; Office of the Australian Information Commissioner, 2024; Therapeutic Goods Administration, 2026; Carayon et al., 2018; Sendelbach and Funk, 2013
[4] Flin et al., 2008; Hollnagel et al., 2015; Hollnagel, 2018; Weick and Sutcliffe, 2015
[5] Topol, 2019; Rajkomar et al., 2019; Hashimoto et al., 2018
[6] Arina et al., 2024; Mehta et al., 2024
[7] Hashimoto et al., 2018; Loftus et al., 2020
[8] Flin et al., 2008; Fletcher et al., 2003; Endsley, 1995

capabilities by analysing continuously generated data, recognising evolving trajectories and identifying patterns associated with deterioration before conventional thresholds are crossed. This creates opportunities for earlier reassessment and intervention rather than waiting for instability to become clinically apparent[9][10].

Prediction alone, however, does not improve patient care. Algorithmic outputs remain probabilistic and may be affected by artefact, incomplete information, unrepresentative training data or changes in clinical populations and practice. Even an accurate prediction may provide little benefit if it is poorly explained, presented at the wrong time or cannot be translated into an effective clinical response. The central question is therefore not only whether AI can predict an event, but whether that information can be safely interpreted and converted into meaningful action[11].

This relationship can be conceptualised through the **Predict–Interpret–Adapt** framework. As illustrated in Figure 1, artificial intelligence may **predict** emerging risks through continuous analysis of clinical and physiological data. Clinicians must then **interpret** those predictions within the context of the individual patient, integrating algorithmic outputs with clinical expertise, situational awareness and information that may not be represented within the dataset. Finally, multidisciplinary teams **adapt** care by modifying management plans, mobilising resources and implementing interventions intended to prevent deterioration or mitigate harm.



The framework deliberately positions AI as one component of clinical decision-making rather than an autonomous authority. Human factors research demonstrates that safe performance in complex environments emerges through interactions between people, technologies, tasks and organisational systems. The value of AI will therefore depend not only on algorithmic performance, but also on appropriate trust, effective communication, clinical oversight and its integration into established workflows and team structures[12].

This narrative review critically examines current and emerging applications of AI across pre-anaesthesia assessment, intra-operative management and postoperative care. It also considers the human factors, ethical, governance and workforce implications of implementation, together with the limitations of the current evidence. It argues that AI's greatest potential lies not in replacing clinical expertise, but in strengthening the ability of anaesthesia professionals and multidisciplinary teams to recognise risk earlier and deliver increasingly predictive, personalised and proactive perioperative care.

Foundations of Artificial Intelligence in Healthcare and Anaesthesia

Artificial intelligence (AI) refers to computational systems designed to perform tasks commonly associated with human intelligence, including pattern recognition, prediction, language processing and decision support[13]. Although AI encompasses numerous technologies, many contemporary healthcare applications use machine learning algorithms to analyse large datasets, identify relationships and generate predictions that may support clinical care[14].

Machine learning models are developed using historical data in which algorithms learn statistical relationships between patient characteristics, clinical information and defined outcomes. Depending on their design, these systems may support tasks such as risk stratification, image interpretation, physiological forecasting and the classification of clinical events. Deep-learning approaches can process particularly complex or high-dimensional data, including medical images, physiological waveforms and continuously generated monitoring information[15].

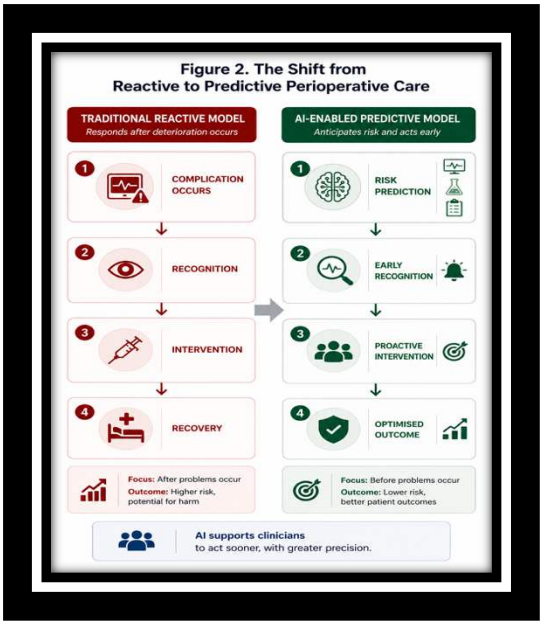
[9] Gaba, 2000; Hollnagel et al., 2015; Weick & Sutcliffe, 2015
[10] Kendale et al., 2018; Wijnberge et al., 2020; Abde et al., 2020
[11] Carayon et al., 2018; Hollnagel et al., 2015; Australian Commission on Safety and Quality in Health Care, 2025
[12] Topol, 2019; Flin et al., 2008; Carayon et al., 2018; Hollnagel, 2018
[13] Topol, 2019
[14] Rajkomar et al., 2019
[15] Topol, 2019; Hashimoto et al., 2018

AI differs from traditional rule-based clinical decision-support systems. Conventional systems follow predefined instructions established by clinicians—for example, generating an alert when a physiological measurement crosses a specified threshold. Machine learning systems instead derive relationships from data and may recognise combinations or trajectories that cannot be represented easily by simple rules[16]. However, deployed models do not necessarily continue learning automatically. Most remain fixed unless they are deliberately monitored, retrained and revalidated using new data. This distinction is important because uncontrolled model updating could alter performance and introduce new risks without appropriate clinical oversight.

Recent developments also extend beyond predictive modelling. Large language models and other forms of generative AI can synthesise clinical text, extract structured information from narrative records, draft documentation and generate preliminary responses to clinical questions. Early perioperative studies demonstrate both promise and limitation: large language models have extracted structured information from intra-operative echocardiography reports, but ChatGPT-generated preoperative anaesthetic plans have not consistently met minimum clinical standards[17][18]. Ambient documentation systems may reduce the time and cognitive effort associated with clinical note-writing, although most evaluations have occurred outside perioperative practice and do not yet establish their effectiveness or safety in anaesthesia[19]. Multimodal systems may eventually combine narrative records, images, physiological waveforms and device data within a single model, but their clinical value remains dependent upon rigorous validation, secure data handling and effective human oversight.

Anaesthesia is particularly suited to AI-supported approaches because it generates large volumes of continuously evolving information from physiological monitors, anaesthetic workstations, infusion devices, laboratory systems and electronic medical records. Anaesthesia professionals must integrate these data with patient history, procedural requirements and anticipated medication effects while making time-critical decisions under conditions of uncertainty[20][21]. This work requires continuous vigilance, anticipation of future events and the ability to prioritise information as patient and procedural conditions change.

Machine learning systems may support this process by recognising subtle physiological patterns, identifying trajectories associated with deterioration and estimating the likelihood of complications before they become clinically apparent[22][23]. As illustrated in Figure 2, the proposed shift is not from human care to automated care, but from isolated threshold-triggered responses towards earlier, trajectory-based clinical support. An AI system may identify an emerging risk, but clinical benefit occurs only when that information prompts timely reassessment and an appropriate response.



Algorithmic outputs nevertheless remain probabilistic estimates rather than definitive diagnoses or treatment recommendations. Their reliability may be affected by artefact, missing information, unrepresentative training data and differences between the population in which a model was developed and the patients to whom it is subsequently applied[24]. A model may also overlook contextual information that is clinically apparent but absent from the available dataset. Clinical expertise therefore remains necessary to determine whether an output is credible, relevant and actionable for an individual patient[25].

Human factors science further recognises that performance within complex clinical environments emerges from interactions between people, technologies, tasks and organisational systems rather than from individual capability alone[26]. AI should therefore be understood as part of a wider sociotechnical system capable of influencing attention, workload, communication, trust and team

[16] Rajkomar et al., 2019; Topol, 2019
[17] MacKay et al., 2025
[18] Abdel Malek et al., 2025
[19] Stults et al., 2025
[20] Hashimoto et al., 2018; Loftus et al., 2020
[21] Flin et al., 2008; Fletcher et al., 2003; Gaba, 2000
[22] Kendale et al., 2018; Wijnberge et al., 2020
[23] Hollnagel et al., 2015; Weick & Sutcliffe, 2015
[24] Topol, 2019; Rajkomar et al., 2019; Arina et al., 2024
[25] Endsley, 1995; Flin et al., 2008
[26] Carayon et al., 2018

performance[27]. A technically accurate prediction may still fail to improve care if it creates excessive alerts, interrupts workflow, is poorly explained or reaches clinicians too late to support an effective intervention.

Within the **Predict–Interpret–Adapt** framework introduced above, AI contributes primarily by identifying patterns and generating predictions. Clinicians and multidisciplinary teams remain responsible for deciding what those predictions mean and how care should change in response. The foundation of safe perioperative AI is therefore not computational performance alone, but the alignment of reliable prediction with clinical context, human judgement and team adaptability[28].

~~Pre-Anaesthesia Assessment and Risk Prediction~~

Pre-anaesthesia assessment aims to identify patient-specific risks, optimise medical conditions and develop a perioperative plan that minimises avoidable harm. Established assessment combines clinical history, physical examination, investigations, validated risk-stratification tools and clinician judgement[29]. These approaches remain fundamental, but their accuracy may be affected by incomplete information, variation in interpretation and the difficulty of estimating how multiple patient and procedural factors interact[30].

Machine learning may complement this process by analysing large datasets containing demographic characteristics, comorbidities, medications, laboratory results, physiological measurements and previous healthcare utilisation [31].Unlike risk tools based on a relatively small number of predefined variables, some machine learning models can identify complex or non-linear relationships between multiple factors. Models have been developed to estimate the likelihood of postoperative complications, unplanned intensive care admission, prolonged hospitalisation and perioperative mortality[32]. Other applications include predicting difficult airway management, postoperative nausea and vomiting, severe postoperative pain and opioid-related adverse events[33]. These predictions could support earlier intervention, more individualised planning and more appropriate allocation of perioperative resources[34].

However, improved predictive performance within a dataset does not necessarily establish clinical usefulness. A systematic review of 103 perioperative machine learning studies found that 75% relied upon single-centre validation and only 13% had been externally validated across multiple centres. Many studies were retrospective and at high risk of bias[35]. Models developed from a single institution may reflect its particular patient population, documentation practices and models of care, limiting their reliability when transferred elsewhere. Furthermore, studies frequently report measures such as discrimination or accuracy without demonstrating that access to the prediction changes clinical management or improves patient-centred outcomes.

This distinction is illustrated by the 2024 Perioperative ORACLE randomised clinical trial. Although clinicians who viewed machine learning predictions showed greater agreement with the models, the assistance did not significantly improve their ability to predict postoperative mortality or acute kidney injury[36]. This finding does not suggest that predictive models lack value, but it demonstrates that presenting clinicians with an algorithmic risk estimate does not automatically improve clinical judgement. Prospective multicentre validation, comparison with established assessment methods and evaluation of effects on decisions and patient outcomes are therefore required before widespread implementation.

As illustrated in Figure 3, AI could contribute to four connected components of pre-anaesthesia assessment: risk prediction, patient optimisation, planning and resource allocation, and shared decision support. These functions should not be viewed independently. A prediction becomes clinically useful only when it provides information that can meaningfully influence preparation, communication or management.

[27] Carayon et al., 2018; Hollnagel et al., 2015
[28] Topol, 2019; Hollnagel, 2018; Hollnagel et al., 2015; Weick & Sutcliffe, 2015
[29] Kheterpal et al., 2009; Glance et al., 2012
[30] Topol, 2019; Rajkomar et al., 2019; Carayon et al., 2018
[31] Rajkomar et al., 2019; Hashimoto et al., 2018; Loftus et al., 2020
[32] Xue et al., 2021; Arina et al., 2024
[33] Hashimoto et al., 2018; Loftus et al., 2020; Nair et al., 2020
[34] Topol, 2019; Loftus et al., 2020; Arina et al., 2024
[35] Arina et al., 2024
[36] Fritz et al., 2024



For example, an older patient presenting for emergency laparotomy may have interacting risk factors including frailty, cardiovascular disease, anaemia, renal impairment and sepsis. Existing assessment methods will identify many of these risks, while a machine learning system may assist by estimating their combined association with postoperative complications[37]. If appropriately validated, this information could prompt earlier multidisciplinary discussion, preparation for haemodynamic instability, blood-product availability, additional monitoring and consideration of postoperative critical care. The algorithm would not determine the plan; it would provide another source of information for the team to evaluate.

Similarly, identifying an increased probability of severe postoperative pain or opioid-related harm may support earlier planning for multimodal analgesia, regional anaesthesia and acute pain service involvement[38]. However, risk identification is valuable only where an effective and proportionate intervention is available. Poorly calibrated predictions could lead to unnecessary investigations, inappropriate escalation of care or diversion of limited resources from patients who may benefit more.

Algorithmic predictions must also be interpreted alongside information that may be poorly represented within electronic datasets, including frailty, patient preferences, health literacy, social circumstances, cultural considerations, procedural urgency and local service capacity[3940]. Particular attention is required to ensure that model performance remains reliable across demographic and clinical subgroups. A prediction should support informed discussion with the patient rather than replace shared decision-making or be presented as a definitive account of individual risk.

Within the **Predict–Interpret–Adapt** framework, pre-anaesthesia AI supplies a prediction that must still be interpreted within the patient's clinical and personal context before the team determines whether and how care should be adapted[4142]. The central measure of success is therefore not whether an algorithm produces a more sophisticated risk score, but whether it helps patients and perioperative teams make safer, better-informed decisions[43].

Viewed through a resilience perspective, the greatest value of artificial intelligence in preoperative assessment may lie in strengthening the capacity of teams to anticipate vulnerability and proactively prepare for emerging risks rather than simply reacting once deterioration occurs[44]. By identifying risk earlier and supporting proactive planning, AI has the potential to shift pre-anaesthesia assessment from predominantly reactive evaluation towards predictive and preventive perioperative medicine.

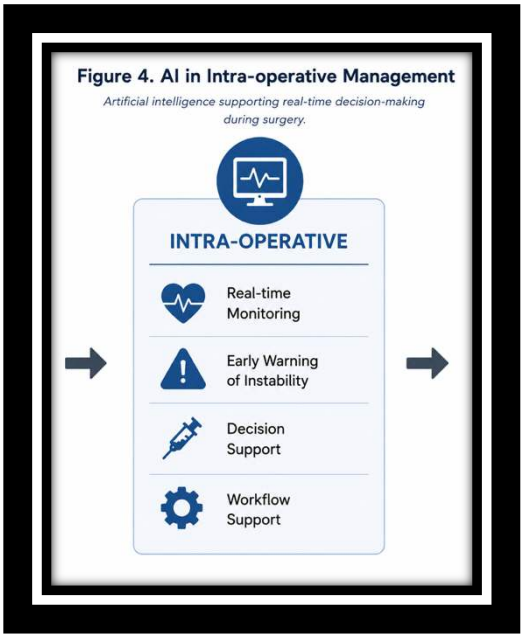
Artificial Intelligence in Intra-operative Management

The intra-operative environment is characterised by continuous physiological monitoring, rapidly evolving clinical conditions and the need for timely decision-making under uncertainty. Anaesthesia professionals must interpret multiple streams of physiological data while administering potent medications, operating sophisticated equipment, anticipating patient responses and coordinating interventions with the wider theatre team [45]. These characteristics create substantial opportunities for appropriately designed AI systems to support clinical performance [46].

[37] Kheterpal et al., 2009; Glance et al., 2012; Xue et al., 2021
[38] Kehlet et al., 2006; Gan et al., 2014; Nair et al., 2020
[39] Flin et al., 2008; Fletcher et al., 2003; Endsley, 1995
[40] Topol, 2019; Rajkomar et al., 2019; Australian Commission on Safety and Quality in Health Care, 2025
[41] Topol, 2019; Rajkomar et al., 2019
[42] Endsley, 1995; Flin et al., 2008
[43] Topol, 2019; Rajkomar et al., 2019; Australian Commission on Safety and Quality in Health Care, 2025
[44] Hollnagel et al., 2015; Hollnagel, 2018; Weick and Sutcliffe, 2015
[45] Hashimoto et al., 2018; Flin et al., 2008; Fletcher et al., 2003; Endsley, 1995
[46] Hashimoto et al., 2018; Loftus et al., 2020

Many conventional monitoring alarms are threshold-based, alerting clinicians when a physiological variable has already moved beyond a predefined limit. These systems remain essential and have contributed to the safety of modern anaesthesia, but an isolated threshold breach may provide limited warning of an evolving problem[47]. Machine learning systems may complement conventional monitoring by analysing high-frequency data, identifying interactions between variables and recognising trajectories that precede clinically apparent deterioration[48].

As illustrated in Figure 4, potential intra-operative applications include real-time monitoring, early warning of instability, clinical decision support and workflow support. These functions are interconnected. Continuous monitoring provides the data from which an emerging risk may be identified, but the warning becomes clinically useful only if it is understandable, timely and linked to an appropriate response. AI should therefore support the clinician's attention and reasoning rather than simply add another alarm to an already information-dense environment.



The most extensively studied example is prediction of intra-operative hypotension. Machine learning algorithms analysing characteristics of the arterial pressure waveform can identify patterns associated with an increased likelihood of hypotension before the mean arterial pressure crosses a conventional threshold[49]. The preliminary HYPE randomised clinical trial found that an early-warning system combined with a structured haemodynamic assessment and treatment protocol reduced the depth and duration of intra-operative hypotension. However, this was a small, single-centre study designed primarily to evaluate hypotension burden rather than improvements in patient-centred outcomes. The findings therefore demonstrate that prediction can influence a physiological process, but do not independently establish reductions in complications, mortality or length of stay.

This distinction is reflected in the broader evidence. A 2024 systematic review of 13 randomised trials found that Hypotension Prediction Index-guided interventions reduced exposure to intra-operative hypotension, while nociception-guided systems reduced pain scores in the post-anaesthesia care unit. However, significant improvements were not consistently demonstrated across other outcomes, including opioid consumption and hospital or post-anaesthesia care unit length of stay[50]. Predicting and reducing a physiological disturbance is potentially valuable, but the clinical importance of that reduction must still be demonstrated.

For example, a patient undergoing major abdominal surgery may develop subtle changes in arterial pressure variability, stroke volume and heart-rate trends before significant haemodynamic instability becomes apparent. An AI-supported monitor may identify the trajectory and prompt earlier reassessment[51]. The appropriate response, however, cannot be determined by the warning alone. The clinician must assess whether the change reflects vasodilation, altered anaesthetic depth, hypovolaemia, impaired cardiac function, surgical manipulation, artefact or another cause before selecting an intervention. Earlier warning may create time for better reasoning, but it does not remove the need for diagnostic judgement.

Other emerging applications include analysis of electroencephalographic and nociception data, ventilatory decision support and closed-loop delivery of anaesthetic agents, fluids or vasoactive medications[52]. Closed-loop systems can automatically adjust therapy towards a specified physiological target under defined conditions. Nevertheless, maintaining a variable within a target range is not equivalent to understanding the patient's condition. Such systems may perform poorly when confronted with artefact, unusual physiology, rapid surgical

[47] Gaba, 2000; Sendelbach and Funk, 2013; Loftus et al., 2020
[48] Topol, 2019; Hashimoto et al., 2018; Kendale et al., 2018; Wijnberge et al., 2020
[49] Wijnberge et al., 2020
[50] Mehta et al., 2024
[51] Wijnberge et al., 2020; Loftus et al., 2020
[52] Hashimoto et al., 2018; Topol, 2019; Loftus et al., 2020

events or situations not adequately represented during development. Autonomous adjustments therefore require clearly defined limits, fail-safe mechanisms and immediate clinician oversight.

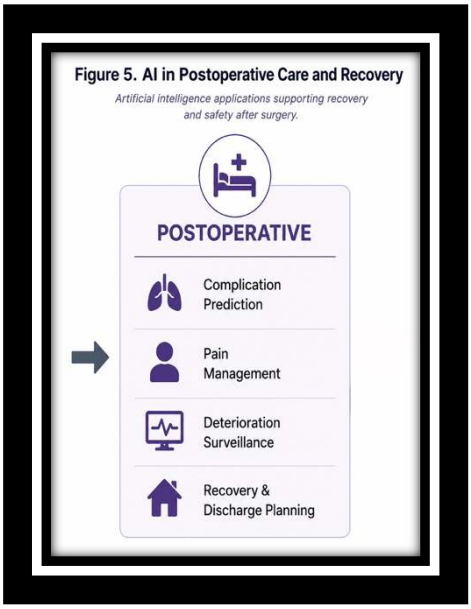
The quality of intra-operative predictions also depends upon the quality of the underlying signal. Arterial-line damping, transducer problems, electrocautery interference, patient movement, arrhythmias and missing data may all produce misleading inputs[53].Poorly designed alerts could increase cognitive workload, contribute to alarm fatigue or direct attention away from more important clinical information. Conversely, excessive confidence in automated recommendations may create automation bias, in which clinicians become less likely to independently verify an output or consider alternative explanations[54].

AI may strengthen situational awareness when it presents reliable information in a way that helps clinicians perceive, understand and anticipate an emerging problem[55]. It may undermine situational awareness when the reasoning behind an alert is unclear, when notifications are excessive or when clinicians become detached from the underlying physiological data. Successful implementation must therefore evaluate not only algorithmic accuracy, but also usability, alert burden, workflow integration and effects on team communication.

Within the **Predict–Interpret–Adapt** framework, intra-operative AI may provide earlier indication of physiological instability, but clinicians must determine whether the prediction is credible and what it means for the individual patient before the team adapts care[56] [57]. The principal value of intra-operative AI is therefore not that it reacts faster than a clinician, but that it may create additional time for clinicians and teams to recognise, investigate and respond to deterioration before it progresses to harm.

~~Artificial Intelligence in Postoperative Care and Acute Pain Management~~

Postoperative recovery is highly variable. Patients undergoing similar procedures may experience substantially different levels of pain, analgesic requirements, complications and recovery times[58]. This variability can make it difficult to determine which patients require intensified monitoring or earlier intervention, particularly after transfer from the post-anaesthesia care unit, when clinical observation may become less continuous.



Artificial intelligence may support a more anticipatory and individualised approach by integrating demographic information, comorbidities, medication history, procedural characteristics and evolving physiological data. These systems can estimate the likelihood of outcomes such as severe pain, elevated opioid requirements, respiratory compromise, delirium and unplanned critical-care admission[59]. As illustrated in Figure 5, these applications extend across complication prediction, pain management, deterioration surveillance, and recovery and discharge planning. Their purpose is not to determine care independently, but to make emerging risks visible early enough for clinicians to respond.

Postoperative pain is a particularly relevant application because inadequate analgesia may delay mobilisation, prolong hospitalisation and contribute to persistent postsurgical pain[60].Procedure-specific protocols provide a useful starting point, but they cannot fully account for differences in previous opioid exposure, psychological distress, comorbidities, pain sensitivity or individual responses to analgesic interventions[61]. Machine-learning models may help identify patients whose likely analgesic requirements differ from the expected pathway.

[53] Wijnberge et al., 2020; Rajkomar et al., 2019
[54] Sendelbach and Funk, 2013; Carayon et al., 2018; Topol, 2019; Rajkomar et al., 2019
[55] Endsley, 1995
[56] Rajkomar et al., 2019; Topol, 2019
[57] Flin et al., 2008; Fletcher et al., 2003; Hollnagel et al., 2015
[58] Kehlet et al., 2006; Gan et al., 2014
[59] Rajkomar et al., 2019; Topol, 2019; Hashimoto et al., 2018; Nair et al., 2020
[60] Kehlet et al., 2006; Gan et al., 2014
[61] Gan et al., 2014 ; Nair et al., 2020

However, predictive performance should be interpreted cautiously. Nair et al. developed models using data from 13,700 ambulatory surgical procedures and reported an overall accuracy of approximately 72% for the best-performing model. The model was considerably better at identifying patients with lower opioid requirements than those requiring higher doses, with sensitivities of 89% and 43%, respectively[62]. This distinction is clinically important because the patients most likely to benefit from early identification were also those detected least reliably. The study demonstrated the feasibility of prediction, but did not establish that model use improved pain, function, opioid-related harm or other patient-centred outcomes. More broadly, much of the evidence in this field remains based on retrospective model development or internal validation rather than prospective evaluation across multiple institutions[63].

The clinical value of a prediction therefore depends on whether it leads to an appropriate and proportionate response. For example, two patients undergoing total knee arthroplasty may have very different postoperative pain trajectories. If one patient is identified as being at increased risk of severe pain or elevated opioid requirements, the clinical team could review the planned multimodal regimen, consider an appropriate regional anaesthesia technique, discuss expectations with the patient and arrange earlier acute pain service involvement[64]. The prediction should initiate assessment and shared decision-making rather than automatically determine treatment.

Artificial intelligence may also support postoperative deterioration surveillance. Algorithms analysing changing observations, laboratory results, medication administration and clinical documentation may identify patterns associated with respiratory compromise, delirium, sepsis, acute kidney injury or unplanned intensive-care admission[65]. This may be particularly valuable on surgical wards, where intermittent observations can leave periods during which deterioration is not immediately recognised.

Nevertheless, a risk score is not equivalent to an improved outcome. A clinically useful system must provide information early enough to alter management, identify a clear recipient for the alert and connect the prediction to an agreed escalation pathway. Poorly calibrated systems may generate false-positive alerts, increase workload and contribute to alert fatigue. Predictions may also become unreliable when monitoring data are incomplete, affected by artefact or drawn from patient populations that differ from those used to develop the model. Postoperative AI should therefore be evaluated not only by discrimination or accuracy, but also by its effects on clinical decisions, timeliness of intervention, patient outcomes and staff workload.

Remote monitoring may extend postoperative surveillance beyond hospital discharge. Wearable devices and digital platforms can collect information about physiological observations, activity, symptoms and recovery patterns, potentially allowing complications to be recognised earlier and follow-up to be adapted to individual need[66]. However, successful implementation requires more than continuous data collection. Systems need reliable connectivity, patient engagement, clinically meaningful alert thresholds and a clearly resourced process for reviewing and responding to abnormal findings. Digital access and confidence also vary between patients, creating a risk that remote monitoring could widen existing inequities if it becomes a substitute for accessible clinical follow-up.

Importantly, postoperative recovery remains patient-centred and relational. Pain, anxiety, functional goals and expectations cannot be understood through physiological data or algorithmic prediction alone[67]. A patient's reported experience may appropriately outweigh a model's estimate. Safe postoperative care also depends on effective communication and shared situational awareness among anaesthetic, surgical, nursing, pain and allied-health teams[68].

Within the Predict–Interpret–Adapt framework, AI can predict emerging risks, but clinicians must interpret those predictions alongside symptoms, procedural factors and patient preferences before the multidisciplinary team adapts the recovery plan[69]. The principal value of AI in postoperative care is therefore not its ability to forecast every adverse event. It is its potential to give patients and clinical teams more time to recognise vulnerability, coordinate an appropriate response and prevent a predicted risk from becoming avoidable harm[70].

[62] Nair et al., 2020
[63] Nair et al., 2020; Arina et al., 2024; Mehta et al., 2024
[64] Nair et al., 2020; Gan et al., 2014
[65] Xue et al., 2021; Topol, 2019; Loftus et al., 2020
[66] Topol, 2019; Loftus et al., 2020
[67] Gan et al., 2014; Topol, 2019
[68] Flin et al., 2008; Fletcher et al., 2003
[69] Endsley, 1995; Flin et al., 2008
[70] Hollnagel et al., 2015; Hollnagel, 2018; Weick & Sutcliffe, 2015

Human Factors and Human–AI Teaming in Anaesthesia

The integration of artificial intelligence into anaesthesia is not simply a technological change. Anaesthesia is a human-centred specialty that depends on clinical judgement, situational awareness, communication and adaptive decision-making under uncertainty[71]. The safety of AI therefore depends not only on what an algorithm can predict, but also on how its outputs affect attention, decisions, workflow and team performance[72].

Human factors science recognises that performance emerges through interactions between people, technologies, tasks and organisational systems[73]. Artificial intelligence should therefore be considered part of a broader sociotechnical system rather than an isolated decision-support tool [74]. Even an algorithm with strong predictive performance may provide little clinical benefit if its output is poorly timed, difficult to interpret or disconnected from the team responsible for acting on it.

One potential benefit is cognitive support. Anaesthesia professionals must continuously interpret changing physiological data while administering medications, preparing equipment, coordinating procedural requirements and anticipating complications. These simultaneous demands place considerable pressure on attention and working memory[75]. AI may assist by analysing complex trends in the background and directing attention towards subtle changes that might otherwise remain unrecognised.

For example, during major abdominal surgery, a system detecting physiological patterns associated with impending hypotension may prompt the team to reassess blood loss, anaesthetic depth, fluid status and vascular tone before clinically significant instability develops[76]. Its value lies in supporting earlier enquiry rather than prescribing a predetermined intervention. The same prediction may require very different responses depending on the patient, procedure and reliability of the available data.

Cognitive support can also produce unintended consequences. Automation bias occurs when users place excessive trust in technological outputs and become less likely to verify information or challenge recommendations[77]. This may produce a commission error, in which an inappropriate recommendation is followed, or an omission error, in which clinicians fail to act because the system has not generated an alert. A patient may still be deteriorating when an algorithm reports low risk, particularly if relevant information is missing, delayed or outside the variables the model was designed to analyse.

The appropriate goal is therefore calibrated trust rather than unquestioning acceptance or wholesale rejection. Clinicians need enough confidence to use a reliable prediction while remaining prepared to override it when it conflicts with the patient's presentation. Interfaces should make uncertainty and changes in risk visible, while training should include examples of model failure, missing data and misleading outputs. This preserves independent clinical reasoning and reduces the risk that prediction replaces assessment.

Predictive notifications may also increase cognitive workload. Anaesthesia environments already contain numerous alarms and competing sources of information. If AI systems prioritise sensitivity without sufficient attention to clinical relevance, they may produce frequent low-value alerts and contribute to alarm fatigue[78]. Repeated false-positive notifications can desensitise clinicians, interrupt critical tasks and make genuinely important warnings more difficult to recognise. Safe design therefore requires careful consideration of alert thresholds, timing, prioritisation and the person best placed to respond.

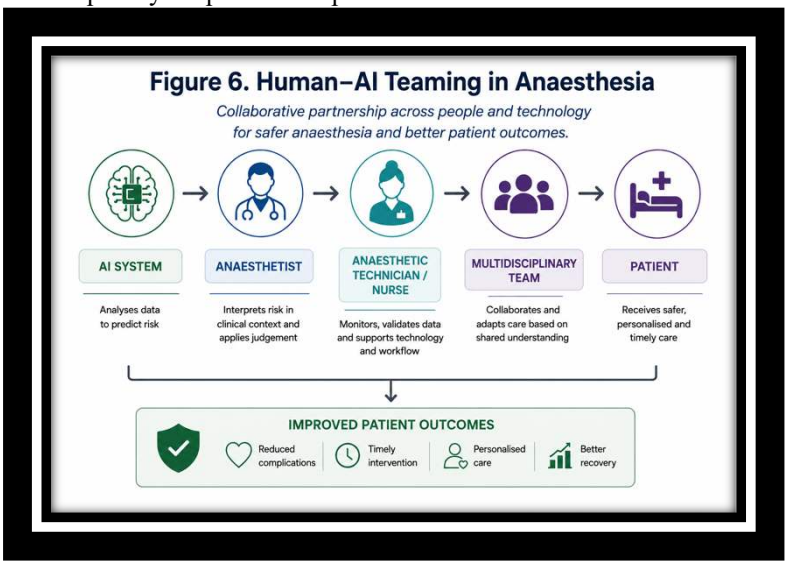
Artificial intelligence may similarly strengthen or weaken situational awareness. Situational awareness involves perceiving relevant information, understanding its significance and anticipating how the situation may develop[79]. Well-designed systems can support this process by presenting trends clearly and drawing attention to emerging risk. Poorly integrated systems may instead narrow attention, encourage fixation on a numerical risk score or distract clinicians from contradictory information such as changes in the surgical field, equipment performance, capnography or the patient's physical presentation.

---

[71] Flin et al., 2008; Fletcher et al., 2003; Gaba, 2000
[72] Topol, 2019; Carayon et al., 2018
[73] Carayon et al., 2018
[74] Carayon et al., 2018; Hollnagel et al., 2015
[75] Flin et al., 2008; Young et al., 2014
[76] Wijnberge et al., 2020
[77] Topol, 2019; Rajkomar et al., 2019
[78] Sendelbach & Funk, 2013; Carayon et al., 2018
[79] Endsley, 1995

These effects extend across the anaesthesia team. Anaesthetists, anaesthetic technicians or practitioners, nurses and other perioperative professionals may each observe different elements of the clinical situation. AI-generated information must therefore be communicated in a way that supports a shared understanding of what has changed, how reliable the prediction appears and what response is required[80]. An alert visible only to one team member may have limited value if the concern is not verbalised, acknowledged and incorporated into the wider management plan. As illustrated in Figure 6, effective human–AI teaming links algorithmic prediction with clinical interpretation, data validation, multidisciplinary adaptation and patient-centred outcomes.

Psychological safety is particularly important within human–AI systems. Team members must feel able to question an algorithmic output, identify an apparent inconsistency and escalate concerns without fear of dismissal[81]. An anaesthetic technician who identifies a discrepancy between an AI prediction and the monitoring equipment, for example, should be able to raise that concern immediately regardless of professional hierarchy. Safe performance may depend on team members being willing to challenge both human assumptions and technological recommendations.



Implementation should consequently be evaluated by more than model accuracy. Prospective testing should examine whether the system improves recognition and decision-making in the intended clinical environment, while usability assessment should consider workload, communication, false-alert burden and unintended changes in team behaviour. Simulation may also help teams practise responding when AI advice is correct, uncertain or demonstrably wrong. This allows weaknesses to be identified before clinicians and patients become dependent on the technology.

Within the Predict–Interpret–Adapt framework, AI contributes prediction, but interpretation and adaptation remain distributed across the clinical team [82]. From a resilience perspective, safety is strengthened when technology helps professionals anticipate problems while preserving their capacity to question, coordinate and recover when predictions are incomplete or incorrect. The future of human–AI teaming in anaesthesia should therefore be judged not by whether AI appears intelligent, but by whether the combined system makes sound clinical action easier and patient harm less likely [83].

Ethical, Legal and Governance Considerations

Artificial intelligence may improve risk prediction and support more personalised perioperative care, but clinical usefulness cannot be separated from ethical, legal and governance responsibilities. AI systems can influence decisions about whether surgery proceeds, how patients are monitored and how limited resources are allocated. Their implementation must therefore protect safety, equity, privacy, autonomy and professional accountability rather than treating these concerns as secondary to technological performance[84].

Algorithmic bias is a central concern. Machine-learning systems learn from historical data, which may contain incomplete documentation, unequal access to care or patterns of clinical decision-making that reflect existing disadvantage[85]. A model trained within one institution may also perform differently when applied to patients with different demographic characteristics, comorbidities or patterns of healthcare access. Overall accuracy can conceal poor performance within smaller or under-represented groups, meaning fairness cannot be assumed from a single performance measure.

[80] Flin et al., 2008; Fletcher et al., 2003; Hollnagel et al., 2015; Weick and Sutcliffe, 2015
[81] Edmondson, 1999; Edmondson, 2019
[82] Endsley, 1995; Flin et al., 2008
[83] Topol, 2019; Hollnagel et al., 2015; Hollnagel, 2018; Weick and Sutcliffe, 2015
[84] Topol, 2019; Rajkomar et al., 2019; Australian Commission on Safety and Quality in Health Care, 2025
[85] Rajkomar et al., 2019; Topol, 2019; Australian Commission on Safety and Quality in Health Care, 2025

In anaesthesia, biased prediction could influence whether a patient receives additional optimisation, postoperative monitoring or access to critical-care resources. Governance should therefore require evaluation across relevant patient subgroups and clinical settings, with particular attention to calibration, false-negative rates and the consequences of error[86]. Where meaningful performance differences are identified, organisations should modify, restrict or discontinue use rather than expecting clinicians to compensate for an unsuitable model.

Explainability presents a related challenge. Complex models may generate predictions without providing clinicians with a readily understandable account of how those conclusions were reached[87]. Clinicians do not necessarily require access to every computational step, but they need sufficient information to determine the intended use, relevant inputs, known limitations and degree of uncertainty. Explanations should support appropriate challenge rather than create a false impression that a prediction is reliable simply because it has been presented alongside a list of influential variables.

Accountability must also remain clear. AI systems can generate recommendations but do not accept professional responsibility for their consequences. Clinicians remain responsible for determining whether an output is appropriate for the individual patient, while healthcare organisations remain responsible for procurement, implementation, staff training and ongoing oversight[88]. Legal responsibility may vary according to the circumstances and continues to evolve, but uncertainty does not justify an accountability gap in which responsibility is informally transferred to the technology.

For example, an algorithm may identify a patient as being at high risk of postoperative respiratory compromise and recommend intensive monitoring. The team must still determine whether the prediction is credible, whether the proposed response is proportionate and how it aligns with the patient's preferences and the available resources. An inaccurate recommendation followed without adequate assessment does not become clinically defensible merely because it originated from an approved system.

Patient autonomy and informed consent also require a proportionate approach. Not every background use of automation will require a separate formal consent process. However, patients should receive understandable information when AI materially influences a significant clinical decision, uses their information in an unexpected way or introduces risks relevant to their choices[89]. Communication should describe AI as decision support, acknowledge uncertainty and provide an opportunity for questions. This helps preserve shared decision-making without overstating either the system's capabilities or the clinician's ability to predict every possible outcome.

Privacy and data governance are particularly important because perioperative AI may process sensitive health information collected from electronic records, monitoring systems, imaging, audio or wearable devices. In Australia, privacy obligations apply to personal information entered into an AI system and to identifiable information generated by it[90]. Organisations must understand where data are stored, who can access them, whether information is used for secondary purposes and whether data may be retained for further model training. Identifiable patient information should not be entered into publicly available AI tools without an approved system and a lawful basis for its use.

Clinical AI may also fall within medical-device regulation. In Australia, software intended to diagnose, monitor, predict or guide treatment may meet the definition of a medical device. Unless excluded or exempt, it must undergo the relevant assessment and be included in the Australian Register of Therapeutic Goods before being legally supplied[91]. Regulatory approval is important, but it does not replace local governance or demonstrate that a system will improve outcomes within every clinical environment.

Effective governance must therefore extend across the entire AI lifecycle. Before implementation, organisations should define the intended clinical purpose, identify accountable owners, examine the evidence base, confirm regulatory status, conduct privacy and security assessments, and validate performance within the intended

[86] Rajkomar et al., 2019; Hashimoto et al., 2018; Australian Commission on Safety and Quality in Health Care, 2025
[87] Topol, 2019; Rajkomar et al., 2019
[88] Australian Commission on Safety and Quality in Health Care, 2025; Topol, 2019
[89] Australian Commission on Safety and Quality in Health Care, 2025 ; Topol, 2019
[90] Office of the Australian Information Commissioner, 2024
[91] Therapeutic Goods Administration, 2026

population and workflow. Clinicians and other users require education about appropriate use, uncertainty, known failure modes and escalation processes.

After implementation, performance should be monitored for model drift, subgroup disparities, false alerts, automation bias and unintended effects on workload or communication[92]. Systems should maintain auditable version histories and enable clinicians to report questionable outputs or adverse events. Significant software updates, changes in clinical practice or shifts in patient populations should trigger reassessment. Governance must also include criteria for restricting or withdrawing a system when its safety, usefulness or equity can no longer be demonstrated.

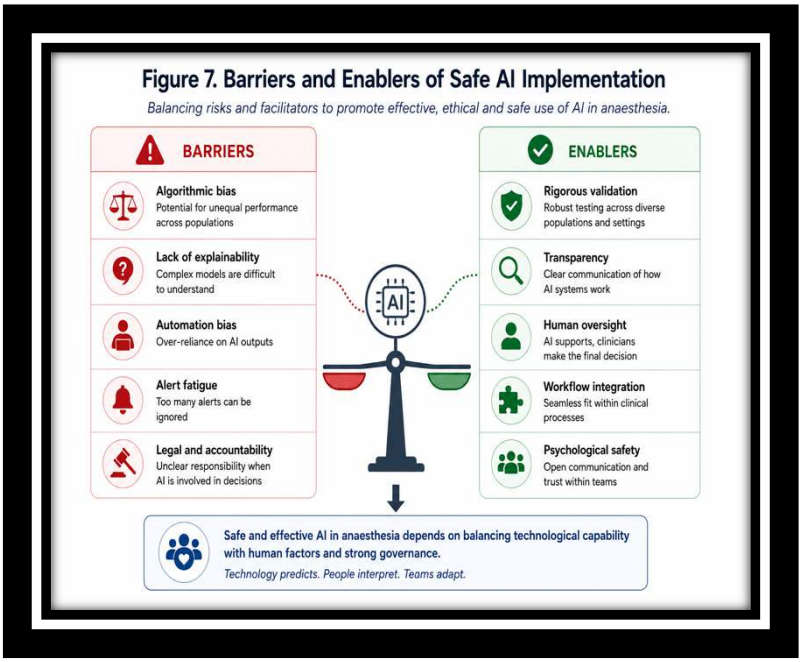


Figure 7 summarises the principal barriers and enablers of safe AI implementation, demonstrating that technological capability must be balanced by rigorous validation, transparency, human oversight, workflow integration and psychological safety. Human oversight must be meaningful rather than symbolic. A clinician cannot safely oversee an algorithm if its operation is opaque, its recommendation arrives too late to be challenged or organisational policy pressures staff to follow it automatically[93]. Oversight requires sufficient information, time, authority and psychological safety to question the technology and choose an alternative course of action.

The Predict–Interpret–Adapt framework therefore applies not only to individual clinical decisions but also to governance. AI predicts risk; clinicians interpret its relevance and limitations; and organisations adapt implementation through monitoring, learning and corrective action. Responsible perioperative AI is not defined by the absence of error, but by transparent systems that identify emerging problems, preserve accountability and respond before technological limitations become patient harm[94].

Practical Implications for Anaesthesia Professionals

The integration of artificial intelligence into perioperative care is likely to change how information is gathered, interpreted and communicated across anaesthesia teams. These changes will require capabilities extending beyond traditional clinical and technical expertise to include digital literacy, critical appraisal, data interpretation and safe interaction with AI-supported systems[95]. The objective is not for every practitioner to become a data scientist, but for users to understand what a system is designed to do, what information it relies upon and when its output may be unreliable.

The effect on workforce demand remains uncertain. Although AI may automate elements of documentation, monitoring analysis and risk calculation, automation does not necessarily reduce workload. It may redistribute work towards verifying information, responding to alerts, managing exceptions and correcting inaccurate outputs[96]. Poorly integrated systems may create additional documentation or troubleshooting demands, whereas well-designed systems may release time for direct patient care and complex decision-making. Claims that AI will either replace anaesthesia professionals or universally improve efficiency are therefore premature.

[92] Australian Commission on Safety and Quality in Health Care, 2025; Therapeutic Goods Administration, 2026; Hollnagel, 2018
[93] Hollnagel, 2018; Edmondson, 2019
[94] Hollnagel, 2018; Edmondson, 2019; Weick & Sutcliffe, 2015
[95] Australian Commission on Safety and Quality in Health Care, 2025; Topol, 2019; Hashimoto et al., 2018
[96] Carayon et al., 2018; Rajkomar et al., 2019; Topol, 2019

In practice, safe use can be considered across three stages. Before using an AI tool, professionals should understand its intended purpose, evidence base, relevant inputs, known limitations and escalation process. During use, they should verify the quality of the underlying data, interpret the output within the clinical context and communicate material uncertainty to the team. After use, significant discrepancies, unexpected outputs and suspected adverse events should be documented and reported through appropriate governance processes[97].

Anaesthetists will retain responsibility for integrating algorithmic predictions with clinical assessment, patient preferences and procedural considerations[98]. AI may support earlier recognition of risk, but decisions concerning anaesthetic technique, haemodynamic management, airway strategy or escalation of care require contextual judgement. As predictive systems become more sophisticated, clinical expertise may increasingly involve determining when a forecast is actionable, when additional information is required and when the prediction should be disregarded.

Anaesthetic Technicians and Anaesthetic Healthcare Practitioners may make particularly important contributions to the safe operation of AI-enabled monitoring and anaesthesia systems. Existing responsibilities for equipment preparation, monitoring integrity, troubleshooting and anticipation of clinical requirements provide a strong foundation for identifying technical conditions that could distort algorithmic output[99]. These may include incorrectly positioned sensors, signal artefact, disconnected monitoring, delayed data transfer or incompatibility between integrated systems.

Their role may also include recognising discrepancies between the patient's presentation, the monitoring equipment and an AI-generated prediction; communicating those discrepancies to the anaesthetist; and supporting recovery when technology fails. These responsibilities must remain within defined scopes of practice, local policies and competency frameworks. An AI-generated recommendation does not independently authorise a practitioner to undertake an intervention outside their education, competence or professional responsibilities.

Anaesthesia and recovery nurses will similarly contribute through continuous patient assessment, recognition of symptom changes and escalation when AI outputs conflict with clinical observations. On postoperative wards, the value of deterioration prediction may depend on nursing staff being able to determine whether an alert is credible, reassess the patient and initiate the appropriate response.

Predictive systems will also become additional sources of information within multidisciplinary communication. Teams will require a shared language for communicating the meaning and urgency of an AI output. Rather than stating that "the AI says the patient will deteriorate", a clearer exchange would identify the predicted outcome, timeframe, degree of uncertainty, relevant physiological changes and required reassessment. This supports shared situational awareness and reduces the risk that probabilistic information is treated as an established fact[100].

Psychological safety remains essential. Any member of the team must be able to question an algorithmic recommendation, identify a data-quality concern or escalate a mismatch between the prediction and the patient's condition[101]. Hierarchical barriers may be particularly hazardous when technological authority is added to existing professional authority. Teams should explicitly establish that AI outputs are open to challenge and that raising a concern is an expected component of safe practice.

Education will therefore need to extend beyond instructions for operating a particular device. Anaesthesia professionals should develop a practical understanding of probability, sensitivity, specificity, false-positive and false-negative results, calibration, model drift and the influence of missing or unrepresentative data. Education should also address automation bias, alert fatigue, privacy, informed consent and the distinction between prediction and causation. These concepts are necessary for appropriately calibrated trust: neither unquestioning reliance nor automatic rejection of AI-generated information[102].

Simulation provides a useful method for developing these competencies. Scenarios could include accurate early warnings, false-positive alerts, silent system failures and predictions generated from corrupted monitoring data.

---

[97] Australian Commission on Safety and Quality in Health Care, 2025
[98] Australian Commission on Safety and Quality in Health Care, 2025; Topol, 2019
[99] Australian Anaesthesia Allied Health Practitioners, 2025; Medical Sciences Council of New Zealand, 2023; Carayon et al., 2018
[100] Endsley, 1995; Flin et al., 2008; Fletcher et al., 2003; Hollnagel et al., 2015
[101] Edmondson, 1999; Edmondson, 2019
[102] Australian Commission on Safety and Quality in Health Care, 2025; Topol, 2019; Rajkomar et al., 2019; Carayon et al., 2018; Sendelbach and Funk, 2013

Training should assess whether teams verify information, communicate uncertainty, maintain situational awareness and recover safely when technological advice is incorrect. It should also include downtime procedures so that essential clinical functions can continue when AI systems or network connections are unavailable.

Frontline professionals should participate in procurement, implementation and evaluation because they understand how technologies interact with real clinical workflows. Anaesthetists, anaesthetic technicians or practitioners, nurses and educators can identify practical hazards that may not be apparent during technical validation, including poorly positioned displays, ambiguous alerts, duplicated work and unclear responsibility for escalation. Their involvement can help determine whether a system supports clinical work or merely adds another layer of information[103].

Within the Predict–Interpret–Adapt framework, AI may strengthen prediction, but safe interpretation and adaptation depend upon a knowledgeable, coordinated and appropriately supported workforce[104]. The practical measure of success is therefore not how much work an algorithm performs independently, but whether the combined human–technology system improves anticipation, preserves professional judgement and enables safer care.

Evidence Gaps and Limitations of Artificial Intelligence in Anaesthesia

Despite rapid growth in artificial intelligence research, the evidence supporting its routine use in anaesthesia remains uneven. Numerous models demonstrate promising predictive performance, but a smaller proportion have progressed through external validation, prospective evaluation and assessment of their effects on clinical outcomes[105]. This distinction is important because successful prediction within a research dataset does not establish that a system will improve care when introduced into a complex clinical environment.

Much of the existing literature remains retrospective and based on data from individual institutions. Models are frequently developed and tested using different portions of the same underlying dataset, meaning the validation population may closely resemble the population used for development[106]. Performance may also be described primarily through discrimination measures such as the area under the receiver operating characteristic curve. A favourable discrimination statistic can obscure poor calibration, low positive predictive value or clinically important false-negative results, particularly when the outcome being predicted is uncommon.

Generalisability therefore remains a major concern. A model developed in one hospital may perform differently when introduced into another with different patient demographics, documentation practices, monitoring equipment, surgical specialties or perioperative pathways[107]. For example, a system predicting postoperative respiratory compromise may perform well in a centre with one analgesic protocol and level of ward monitoring but become less reliable where opioid practices, staffing or escalation thresholds differ. Prospective external validation should therefore occur within the intended clinical environment before clinicians rely on its predictions.

Evidence of improved patient outcomes is more limited than evidence of predictive accuracy. A 2024 systematic review and meta-analysis identified only 13 randomised controlled trials evaluating machine-learning-driven perioperative interventions. These trials examined only three intervention types: the Hypotension Prediction Index, the Nociception Level Index and a scheduling system. The Hypotension Prediction Index reduced exposure to intra-operative hypotension, while the Nociception Level Index produced lower pain scores in the post-anaesthesia care unit. However, no significant effects were demonstrated for intra-operative or postoperative opioid consumption, post-anaesthesia care unit length of stay or hospital length of stay[108].

These findings illustrate the difference between influencing an intermediate measure and improving a patient-centred outcome. Reducing the duration of hypotension may be clinically valuable, but the wider benefit depends on whether intervention prevents organ injury, improves recovery or reduces mortality without introducing treatment-related harm[109]. Likewise, identifying a patient as high risk has limited value if the

[103] Australian Commission on Safety and Quality in Health Care, 2025; Carayon et al., 2018; Hollnagel, 2018
[104] Topol, 2019; Hollnagel et al., 2015; Hollnagel, 2018; Weick and Sutcliffe, 2015
[105] Topol, 2019; Rajkomar et al., 2019; Hashimoto et al., 2018; Loftus et al., 2020; Mehta et al., 2024
[106] Rajkomar et al., 2019; Hashimoto et al., 2018; Loftus et al., 2020
[107] Rajkomar et al., 2019; Nair et al., 2020; Loftus et al., 2020
[108] Mehta et al., 2024
[109] Mehta et al., 2024; Wijnberge et al., 2020

prediction arrives too late, does not connect to an effective intervention or cannot be acted upon because staffing and monitoring resources are unavailable.

Data quality creates additional limitations. Electronic health records contain missing values, inconsistent terminology, delayed documentation and measurements affected by artefact. Missing information is also not necessarily random; the decision to order a test or increase monitoring may itself reflect clinician concern[110]. Models may inadvertently learn these local patterns rather than the underlying clinical condition. Data leakage may further inflate reported performance if information recorded after the intended prediction point is included during model development.

Representativeness and equity must also be considered. Training datasets may inadequately represent smaller demographic groups, rural or resource-limited services, patients with disability, or populations experiencing unequal access to healthcare. A model can demonstrate acceptable overall performance while producing systematically less accurate predictions for particular groups. Evaluation should therefore report subgroup calibration and error rates rather than assuming that aggregate accuracy demonstrates equitable performance.

Explainability does not completely resolve these problems. Presenting influential variables may help clinicians understand what contributed to a prediction, but a plausible explanation does not prove that the model is accurate or that the relationship is causal[111]. Clinicians need information about the system's intended purpose, validation population, uncertainty and known failure modes. Reporting should also make clear when a model should not be used and what action is expected when its prediction conflicts with clinical assessment.

Human-factors effects remain insufficiently evaluated. AI may reduce cognitive workload by directing attention towards emerging risk, but it may also introduce automation bias, alert fatigue, fixation and loss of vigilance[112]. These effects may not be apparent during retrospective model testing. Prospective evaluation should examine how clinicians actually interact with the system, including whether alerts are understood, whether recommendations are challenged and whether new information improves or disrupts team communication.

Organisational conditions are equally influential. Implementation requires integration with existing monitoring systems, electronic records, escalation pathways and clinical responsibilities[113]. A system may create additional work if alerts are duplicated, responsibility for responding is unclear or outputs must be manually transferred between platforms. Differences in digital infrastructure and workforce capacity may also mean that evidence from large academic hospitals cannot be directly applied to smaller regional or resource-constrained services.

Performance may change after implementation. Patient populations, clinical practices, data definitions, software integrations and model versions evolve over time, creating the possibility of model drift[114]. Ongoing governance should therefore monitor calibration, subgroup performance, false-alert burden, overrides, adverse events and changes in clinical outcomes. Significant software updates or workflow changes should trigger reassessment rather than assuming that the original validation remains applicable indefinitely.

Future research should move beyond demonstrating that AI can predict an outcome. Priority should be given to prospective multicentre studies, standardised perioperative outcomes and comparisons against existing clinical practice rather than against no decision support[115]. Studies should evaluate patient experience, functional recovery, morbidity, mortality, equity, workload and cost-effectiveness. Qualitative and simulation research will also be required to understand how AI changes communication, professional roles and adaptive performance.

Current evidence therefore supports cautious progression rather than either uncritical adoption or wholesale rejection. Within the Predict–Interpret–Adapt framework, research has concentrated most heavily on whether AI can **predict** risk. Considerably more evidence is required to establish how predictions are **interpreted** in practice and whether healthcare systems can **adapt** in ways that reliably improve outcomes[116]. The central evidence gap is no longer simply whether an algorithm can forecast deterioration, but whether its implementation helps clinical teams prevent it.

[110] Rajkomar et al., 2019; Hashimoto et al., 2018; Carayon et al., 2018
[111] Topol, 2019; Rajkomar et al., 2019; Australian Commission on Safety and Quality in Health Care, 2025
[112] Carayon et al., 2018; Sendelbach & Funk, 2013; Topol, 2019
[113] Flin et al., 2008; Fletcher et al., 2003; Endsley, 1995; Hollnagel et al., 2015
[114] Australian Commission on Safety and Quality in Health Care, 2025; Therapeutic Goods Administration, 2026; Hollnagel, 2018
[115] Mehta et al., 2024; Rajkomar et al., 2019; Australian Commission on Safety and Quality in Health Care, 2025
[116] Carayon et al., 2018; Hollnagel et al., 2015; Hollnagel, 2018; Weick and Sutcliffe, 2015

Future Directions and Conclusion: Towards Predictive and Personalised Perioperative Care

Artificial intelligence in anaesthesia has reached an important translational threshold. The challenge is no longer simply demonstrating that algorithms can identify patterns within perioperative data, but determining whether those predictions improve clinical decisions and patient outcomes. Recent reviews show that model development continues to outpace external validation and clinical-impact evaluation. The Perioperative ORACLE randomised trial similarly found that providing clinicians with machine-learning risk estimates increased agreement with the models but did not significantly improve their prediction of postoperative mortality or acute kidney injury[117]. Making a prediction available is therefore not equivalent to making clinical judgement more accurate.

Future research should prioritise prospective, multicentre and adequately powered evaluations across diverse patient populations and healthcare settings. Studies should compare AI-supported care with existing clinical practice and measure outcomes that matter to patients, including morbidity, functional recovery, pain, patient experience, length of stay and mortality[118]. Evaluation should also examine equity, workload, false-alert burden and cost-effectiveness. This will help determine not merely whether a model performs accurately, but whether its implementation provides sufficient benefit to justify the resources and new risks it introduces.

Future perioperative AI will also extend beyond conventional predictive models. Large language models, ambient documentation systems and multimodal platforms may combine clinical notes, physiological signals, imaging, medication data and spoken communication. Early studies demonstrate potential applications in preoperative planning, clinical documentation and extraction of structured information from perioperative reports[119]. However, generative systems introduce additional risks, including hallucinated information, loss of context, privacy breaches and outputs that appear authoritative despite being incorrect. Their evaluation will therefore require clinical verification and careful assessment of how generated information influences decisions.

Human-factors research should remain central to this process. Future studies must examine how AI affects cognitive workload, situational awareness, communication, trust calibration and multidisciplinary teamwork[120]. Evaluation should include circumstances in which predictions are correct, uncertain or wrong, because safety depends on the capacity of clinicians and teams to recognise the difference. Technologies should be designed around clinical work rather than requiring clinical work to be reorganised around technological limitations.

Education and workforce development will be equally important. Anaesthesia professionals will require practical competencies in digital health, data interpretation, algorithmic limitations and human–AI interaction[121]. Anaesthetists will continue to provide clinical interpretation and decision-making, while anaesthetic technicians or practitioners, nurses and other perioperative professionals will contribute through data-quality assurance, equipment integration, monitoring, troubleshooting, patient assessment and coordinated responses to emerging risk. These roles should be developed through multidisciplinary education and simulation rather than isolated technical training.

Governance must also evolve with the technology. Healthcare organisations will require transparent processes for approving AI systems, confirming regulatory status, evaluating evidence, protecting patient information and monitoring performance after implementation[122]. Governance should account for model drift, software updates, cybersecurity, subgroup performance and changes in clinical workflows. It must also provide mechanisms for reporting questionable outputs, investigating patient-safety events and withdrawing systems whose benefit or reliability cannot be maintained.

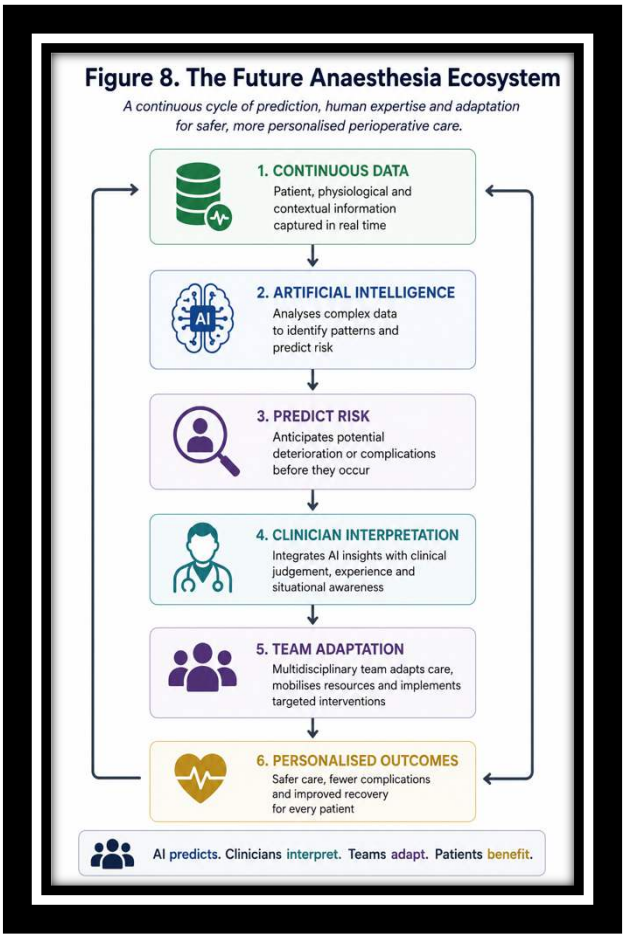
Across the perioperative continuum, the evidence supports a consistent but qualified direction of travel. Before surgery, AI may assist with risk stratification and optimisation. During anaesthesia, predictive systems may identify physiological vulnerability before conventional thresholds are crossed. Following surgery, models may

[117] Arina et al., 2024; Fritz et al., 2024; Mehta et al., 2024
[118] Rajkomar et al., 2019; Arina et al., 2024; Mehta et al., 2024; Australian Commission on Safety and Quality in Health Care, 2025
[119] Abdel Malek et al., 2025; MacKay et al., 2025; Stults et al., 2025
[120] Flin et al., 2008; Fletcher et al., 2003; Endsley, 1995; Carayon et al., 2018
[121] Australian Commission on Safety and Quality in Health Care, 2025; Topol, 2019; Australian Anaesthesia Allied Health Practitioners, 2025
[122] Australian Commission on Safety and Quality in Health Care, 2025; Therapeutic Goods Administration, 2026; Office of the Australian Information Commissioner, 2024; Hollnagel, 2018

support pain planning, deterioration surveillance and individualised recovery pathways[123]. These applications do not eliminate uncertainty; they create an opportunity to recognise it earlier.

The Predict–Interpret–Adapt framework captures the conditions required for that opportunity to become clinically useful. Prediction provides an estimate of what may happen, but clinicians must interpret its relevance within the patient's circumstances, and multidisciplinary teams must adapt care through proportionate and coordinated action[124]. If interpretation is weak or the system cannot support an effective response, predictive accuracy alone offers limited protection.

As illustrated in Figure 8, the future anaesthesia ecosystem may operate as a continuous cycle in which clinical data support prediction, clinicians interpret risk, teams adapt care and patient outcomes inform further improvement. From a resilience perspective, AI is most valuable when it increases the time and information available for teams to anticipate, monitor, respond and learn[125]. It should strengthen the adaptive capacity of perioperative systems without weakening vigilance, professional judgement or the ability to recover when technology fails. The goal is not to create an error-free automated environment, but to build a safer clinical system that can recognise developing risk and respond before harm becomes inevitable.



Artificial intelligence may fundamentally change how information is processed across the perioperative continuum, supporting a shift from reactive responses towards more predictive and personalised care. However, the defining capabilities of anaesthesia—clinical judgement, adaptability, situational awareness, communication and human connection—remain essential. The future of perioperative AI will depend not merely on smarter algorithms, but on combining technology with sound clinical judgement, effective teamwork and strong governance[126].

Predict risk. Interpret complexity. Adapt care.

This is the promise of artificial intelligence in anaesthesia.

[123] Kendale et al., 2018; Wijnberge et al., 2020; Nair et al., 2020; Mehta et al., 2024
[124] Endsley, 1995; Flin et al., 2008; Topol, 2019
[125] Hollnagel et al., 2015; Hollnagel, 2018; Weick and Sutcliffe, 2015
[126] Flin et al., 2008; Edmondson, 2019; Hollnagel, 2018; Australian Commission on Safety and Quality in Health Care, 2025

References

Abdel Malek, M., van Velzen, M., Dahan, A., Martini, C. H., Sitsen, E., Sarton, E. Y., & Boon, M. (2025). Generation of preoperative anaesthetic plans by ChatGPT-4.0: A mixed-method study. *British Journal of Anaesthesia, 134*(5), 1333–1340. https://doi.org/10.1016/j.bja.2024.08.038

Arina, P., Kaczorek, M. R., Hofmaenner, D. A., Pisciotta, W., Refinetti, P., Singer, M., Mazomenos, E. B., & Whittle, J. (2024). Prediction of complications and prognostication in perioperative medicine: A systematic review and PROBAST assessment of machine learning tools. *Anesthesiology, 140*(1), 85–101. https://doi.org/10.1097/ALN.0000000000004764

Australian Anaesthesia Allied Health Practitioners. (2025). *Anaesthesia technician national scope of practice*. https://www.aaahp.org.au/aaahp-cn/uploads/2026/03/AAAHP-National-Scope-of-Practice-for-Anaesthesia-Technicians.pdf

Australian Commission on Safety and Quality in Health Care. (2025). *AI clinical use guide: Guidance for clinicians* (Version 1.0). https://www.safetyandquality.gov.au/sites/default/files/resources/attachments/ai-clinical-use-guide.pdf

Carayon, P., Wooldridge, A., Hose, B.-Z., Salwei, M., & Benneyan, J. (2018). Challenges and opportunities for improving patient safety through human factors and systems engineering. *Health Affairs, 37*(11), 1862–1869. https://doi.org/10.1377/hlthaff.2018.0723

Edmondson, A. C. (1999). Psychological safety and learning behavior in work teams. *Administrative Science Quarterly, 44*(2), 350–383. https://doi.org/10.2307/2666999

Edmondson, A. C. (2019). *The fearless organization: Creating psychological safety in the workplace for learning, innovation, and growth*. Wiley.

Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors, 37*(1), 32–64. https://doi.org/10.1518/001872095779049543

Fletcher, G., Flin, R., McGeorge, P., Glavin, R., Maran, N., & Patey, R. (2003). Anaesthetists' non-technical skills (ANTS): Evaluation of a behavioural marker system. *British Journal of Anaesthesia, 90*(5), 580–588. https://doi.org/10.1093/bja/aeg112

Flin, R., O'Connor, P., & Crichton, M. (2008). *Safety at the sharp end: A guide to non-technical skills*. Ashgate.

Fritz, B. A., King, C. R., Abdelhack, M., Chen, Y., Kronzer, A., Abraham, J., Tripathi, S., Ben Abdallah, A., Kannampallil, T., Budelier, T. P., Helsten, D., Montes de Oca, A., Mehta, D., Sontha, P., Higo, O., Kerby, P., Gregory, S. H., Wildes, T. S., & Avidan, M. S. (2024). Effect of machine learning models on clinician prediction of postoperative complications: The Perioperative ORACLE randomised clinical trial. *British Journal of Anaesthesia, 133*(5), 1042–1050. https://doi.org/10.1016/j.bja.2024.08.004

Gaba, D. M. (2000). Crisis resource management and teamwork training in anaesthesia. *British Journal of Anaesthesia, 85*(1), 26–31. https://doi.org/10.1093/bja/85.1.26

Gan, T. J., Habib, A. S., Miller, T. E., White, W., & Apfelbaum, J. L. (2014). Incidence, patient satisfaction, and perceptions of postsurgical pain: Results from a US national survey. *Current Medical Research and Opinion, 30*(1), 149–160. https://doi.org/10.1185/03007995.2013.860019

Glance, L. G., Lustik, S. J., Hannan, E. L., Osler, T. M., Mukamel, D. B., Qian, F., & Dick, A. W. (2012). The Surgical Mortality Probability Model: Derivation and validation of a simple risk prediction rule for noncardiac surgery. *Annals of Surgery, 255*(4), 696–702. https://doi.org/10.1097/SLA.0b013e31824b45af

Hashimoto, D. A., Rosman, G., Rus, D., & Meireles, O. R. (2018). Artificial intelligence in surgery: Promises and perils. *Annals of Surgery, 268*(1), 70–76. https://doi.org/10.1097/SLA.0000000000002693

Hollnagel, E. (2018). *Safety-II in practice: Developing the resilience potentials*. Routledge. https://doi.org/10.4324/9781315201023

Hollnagel, E., Wears, R. L., & Braithwaite, J. (2015). *From Safety-I to Safety-II: A white paper*. The Resilient Health Care Net. https://www.england.nhs.uk/signuptosafety/wp-content/uploads/sites/16/2015/10/safety-1-safety-2-whte-papr.pdf

Kehlet, H., Jensen, T. S., & Woolf, C. J. (2006). Persistent postsurgical pain: Risk factors and prevention. *The Lancet, 367*(9522), 1618–1625. https://doi.org/10.1016/S0140-6736(06)68700-X

Kendale, S., Kulkarni, P., Rosenberg, A. D., & Wang, J. (2018). Supervised machine-learning predictive analytics for prediction of postinduction hypotension. *Anesthesiology, 129*(4), 675–688. https://doi.org/10.1097/ALN.0000000000002374

Kheterpal, S., Tremper, K. K., Heung, M., Rosenberg, A. L., Englesbe, M., Shanks, A., & Campbell, D. A. (2009). Development and validation of an acute kidney injury risk index for patients undergoing general surgery. *Anesthesiology, 110*(3), 505–515. https://doi.org/10.1097/ALN.0b013e3181979440

Loftus, T. J., Tighe, P. J., Filiberto, A. C., Efron, P. A., Brakenridge, S. C., Mohr, A. M., Rashidi, P., Upchurch, G. R., Jr., & Bihorac, A. (2020). Artificial intelligence and surgical decision-making. *JAMA Surgery, 155*(2), 148–158. https://doi.org/10.1001/jamasurg.2019.4917

MacKay, E. J., Goldfinger, S., Chan, T. J., Grasfield, R. H., Eswar, V. J., Li, K., Cao, Q., & Pouch, A. M. (2025). Automated structured data extraction from intraoperative echocardiography reports using large language models. *British Journal of Anaesthesia, 134*(5), 1308–1317. https://doi.org/10.1016/j.bja.2025.01.028

Medical Sciences Council of New Zealand. (2023). *Review of the scope of practice for anaesthetic technology practitioners* [Consultation document]. https://www.mscouncil.org.nz/assets_mlsb/Consultation-Review-of-the-Scope-of-Practice-for-Anaesthetic-Technology-Practitioners.pdf

Mehta, D., Gonzalez, X. T., Huang, G., & Abraham, J. (2024). Machine learning-augmented interventions in perioperative care: A systematic review and meta-analysis. *British Journal of Anaesthesia, 133*(6), 1159–1172. https://doi.org/10.1016/j.bja.2024.08.007

Nair, A. A., Velagapudi, M. A., Lang, J. A., Behara, L., Venigandla, R., Velagapudi, N., Fong, C. T., Horibe, M., Lang, J. D., & Nair, B. G. (2020). Machine learning approach to predict postoperative opioid requirements in ambulatory surgery patients. *PLOS ONE, 15*(7), e0236833. https://doi.org/10.1371/journal.pone.0236833

Office of the Australian Information Commissioner. (2024, October 21). *Guidance on privacy and the use of commercially available AI products*. https://www.oaic.gov.au/privacy/privacy-guidance-for-organisations-and-government-agencies/guidance-on-privacy-and-the-use-of-commercially-available-ai-products

Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. *The New England Journal of Medicine, 380*(14), 1347–1358. https://doi.org/10.1056/NEJMra1814259

Sendelbach, S., & Funk, M. (2013). Alarm fatigue: A patient safety concern. *AACN Advanced Critical Care, 24*(4), 378–386. https://doi.org/10.1097/NCI.0b013e3182a903f6

Stults, C. D., Deng, S., Martinez, M. C., Wilcox, J., Szwerinski, N., Chen, K. H., Driscoll, S., Washburn, J., & Jones, V. G. (2025). Evaluation of an ambient artificial intelligence documentation platform for clinicians. *JAMA Network Open, 8*(5), e258614. https://doi.org/10.1001/jamanetworkopen.2025.8614

Therapeutic Goods Administration. (2026, February 5). *Artificial intelligence (AI) and medical device software regulation*. https://www.tga.gov.au/products/medical-devices/software-and-artificial-intelligence-ai/manufacturing/artificial-intelligence-ai-and-medical-device-software-regulation

Topol, E. (2019). *Deep medicine: How artificial intelligence can make healthcare human again*. Basic Books.

Weick, K. E., & Sutcliffe, K. M. (2015). *Managing the unexpected: Sustained performance in a complex world* (3rd ed.). Wiley.

Wijnberge, M., Geerts, B. F., Hol, L., Lemmers, N., Mulder, M. P., Berge, P., Schenk, J., Terwindt, L. E., Hollmann, M. W., Vlaar, A. P., & Veelo, D. P. (2020). Effect of a machine learning-derived early warning system for intraoperative hypotension versus standard care on depth and duration of intraoperative hypotension during elective noncardiac surgery: The HYPE randomized clinical trial. *JAMA, 323*(11), 1052–1060. https://doi.org/10.1001/jama.2020.0592

Xue, B., Li, D., Lu, C., King, C. R., Wildes, T., Avidan, M. S., Kannampallil, T., & Abraham, J. (2021). Use of machine learning to develop and evaluate models using preoperative and intraoperative data to identify risks of postoperative complications. *JAMA Network Open, 4*(3), e212240. https://doi.org/10.1001/jamanetworkopen.2021.2240

Young, J. Q., van Merriënboer, J., Durning, S., & ten Cate, O. (2014). Cognitive load theory: Implications for medical education: AMEE Guide No. 86. *Medical Teacher, 36*(5), 371–384. https://doi.org/10.3109/0142159X.2014.889290