ANESTHESIOLOGY

Openventio PUBLISHERS

Review Article

Evaluation of Digital Decision-Support Tools in Operating Room Anesthesia Management

Marius Bianchi[1*], Bence Sorensen[2] and Renate Fernandez[3]

[1]Department of Anesthesiology, Tyrrhenian University of Health Sciences, Italy
[2]Department of Perioperative Medicine, Nordfjord University of Health Sciences, Norway
[3]Department of Surgical Sciences, Nordfjord University of Health Sciences, Norway

ABSTRACT

Digital decision-support tools have moved from simple rule-based reminders embedded in anesthesia information management systems to predictive algorithms that forecast intraoperative hypotension and closed-loop systems that autonomously titrate anesthetic depth or vasopressor infusion.

This review synthesizes the comparative evidence for these tools across three broad categories: rule-based point-of-care alerting, machine-learning predictive analytics exemplified by the Hypotension Prediction Index, and closed-loop automated drug delivery.

Narrowly scoped alerts tied to a single well-defined action show the most consistent improvement in perioperative process measures, while predictive hypotension indices reliably reduce the physiologic surrogate they target but show weaker and more heterogeneous evidence for reducing hard postoperative complications.

Closed-loop anesthetic-depth control has the deepest randomized evidence base of any category reviewed and generally outperforms manual titration on time-in-target metrics without an apparent safety signal against it, though multi-variable autonomous systems and the question of sustained clinician vigilance during automated control remain comparatively unstudied.

We conclude that closing the gap between well-demonstrated process and surrogate benefits and more rigorously established hard clinical outcomes, rather than adding further predictive parameters, is the field's most pressing next step.

Keywords: *clinical decision support; anesthesia information management systems; hypotension prediction index; closed-loop anesthesia; artificial intelligence in anesthesia*

Citation: Bianchi M, Sorensen B and Fernandez R (2022) Evaluation of Digital Decision-Support Tools in Operating Room Anesthesia Management. Anesthesiology. 2(1):01-06.

DOI: 10.17140/2022.vol2.i01.pp01-06

Corresponding Author: *Marius Bianchi*

E-mail Address: *m.bianchi@tuhs.it*

Article History: *Received 17 August 2022 | Accepted 06 September 2022 | Published 06 October 2022*

Copyright: *© 2022 Marius Bianchi, et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited (**https://creativecommons.org/licenses/by/4.0/**).*

INTRODUCTION

Anesthesiologists have always practiced a kind of continuous decision support in their own heads: watch the trend, anticipate the next five minutes, act before

the vital sign becomes a problem. What has changed over the past decade is that a growing share of that anticipatory work is now shared with, or handed to, software. Anesthesia information management systems (AIMS) accumulate dense physiologic and pharmacologic data streams throughout a case, and that density has made the operating room one of the more fertile environments in medicine for embedded clinical decision support (CDS) — rule-based alerts for drug interactions and antibiotic redosing, statistical or machine-learning models that forecast an impending hypotensive episode, and closed-loop systems that titrate anesthetic depth or vasopressor infusion without waiting for a human to turn a dial. This review surveys the state of digital decision-support tools specifically in the intraoperative anesthesia setting: what has been built, what the comparative evidence says about whether it changes outcomes, and where the unresolved problems — alert fatigue, algorithm transparency, and the diffusion of responsibility when software and clinician disagree — still sit.

It is worth being precise about scope. "Decision support" in anesthesia now spans a wide range of maturity, from decades-old rule-based reminders embedded in an AIMS, through predictive indices trained on retrospective waveform data, to genuinely autonomous closed-loop drug delivery. These are not the same technology and do not deserve the same evidentiary standard, but they are frequently discussed together because they share an underlying premise: that a machine watching the monitor continuously and without fatigue can catch what a busy clinician sometimes misses, or can act faster than a human reaction time allows.

RULE-BASED AND ALERT-DRIVEN DECISION SUPPORT IN THE AIMS ERA

The earliest operating-room decision support was unglamorous and rule-based — checklist logic embedded in perioperative information systems that flagged a missed antibiotic redose, an unusually low or high drug dose, or a documentation gap. Sanderson and colleagues' early work on a decision-support system to assist anesthesiologists intraoperatively [1] anticipated much of what followed: the goal was never to replace clinical judgment but to catch the lapses that occur when attention is divided across a monitor, a surgical field, and a chart. A 2011 study using real-time CDS embedded in an AIMS to improve performance on perioperative quality and process measures [2] demonstrated one of the clearer wins in this category — measurable improvement in adherence to process metrics such as timely antibiotic administration and glucose management when the system fired point-of-care reminders rather than relying on retrospective audit and feedback.

A systematic review of near-real-time and point-of-care CDS in anesthesia information management systems [3] found that the tools with the most consistent effect were narrowly scoped ones — a single reminder tied to a single, well-defined action, delivered at the moment the action was still possible to take. Broader, more ambitious alerting systems fared less well, and the review's authors flagged what has become a recurring theme in this literature: an alert's effectiveness depends less on its clinical accuracy than on its timing, its specificity, and whether it asks the clinician to do something they were not already about to do anyway. A more recent implementation study describing a CDS tool built to curtail high-global-warming-potential anesthetic gas use across a large health network [4] extended this rule-based tradition into an environmental-stewardship application, showing that simple, well-targeted point-of-care nudges — in this case a real-time low-flow and agent-selection reminder — can shift practice patterns measurably even outside the traditional safety-and-quality framing, though the same review noted that sustained behavior change required more than a one-time alert.

The broader meta-analytic picture is more equivocal. A 2024 systematic review and meta-analysis of outcomes from clinical decision support systems in real-world perioperative care [5] pooled a heterogeneous set of studies spanning alerting, predictive, and closed-loop tools and found genuine but modest aggregate benefit on process measures, with much weaker and less consistent signal on hard clinical endpoints like postoperative complications or mortality. That gap — reliable improvement in whether the "right thing" gets documented or triggered, against a much shakier case that patients actually do better — recurs across nearly every category of tool discussed in this review, and it is the central methodological challenge facing the field.

PREDICTIVE ANALYTICS FOR INTRAOPERATIVE HYPOTENSION

The single most heavily studied predictive tool in contemporary anesthesia practice is the Hypotension Prediction Index (HPI), a machine-learning algorithm derived from arterial waveform analysis that outputs a probability score for an impending hypotensive episode, typically with several minutes of lead time before mean arterial pressure actually falls below a defined threshold. The appeal is obvious: intraoperative hypotension is common, frequently asymptomatic in an anesthetized patient, and — per a substantial body of prior observational literature — associated with acute kidney injury, myocardial injury, and stroke in a duration- and depth-dependent fashion. A tool that gives the clinician a warning before the event, rather than a notification after it has already happened, changes the nature of the intervention from reactive to preventive.

A narrative review of the HPI's predictive and preventive rationale [6] lays out the physiologic basis for the algorithm and the early single-center evidence

that HPI-guided protocols reduce both the frequency and cumulative time-weighted burden of intraoperative hypotension compared with standard hemodynamic monitoring. That early promise has since been tested in a run of systematic reviews and meta-analyses. A 2025 meta-analysis of HPI-guided management specifically in abdominal surgery [7] found a consistent reduction in intraoperative hypotension exposure and some signal toward improved postoperative outcomes, though the authors were candid that the trials feeding the pooled estimate were heterogeneous in hypotension definition, threshold, and the co-interventions bundled alongside the HPI display. A broader systematic review and meta-analysis of the HPI across surgical populations [8] reached a similar conclusion — real reduction in the physiologic surrogate, weaker and less consistent evidence for downstream clinical benefit — and an ongoing randomized trial protocol in brain tumor surgery [9] illustrates that the field itself recognizes the evidence base still needs population-specific confirmation before the tool can be considered validated for high-risk neurosurgical patients specifically.

Two further meta-analyses examined the machine-learning methodology underlying HPI and comparable algorithms more critically. One pooling analysis of machine-learning algorithms for predicting intraoperative hypotension [10] found respectable discriminative performance by conventional statistical metrics but noted substantial between-study variability in how "prediction" was validated — many studies tested the algorithm on data resembling its own training distribution rather than genuinely prospective, external populations. A companion systematic review of the HPI and machine-learning methods more broadly [11] echoed this caution, and a single-center clinical trial evaluating HPI software specifically during prolonged prone-position spine surgery [12] is a useful illustration of the direction the field needs to move — testing the tool in a specific, physiologically distinct surgical position rather than treating "intraoperative hypotension" as a single undifferentiated entity across all case types.

CLOSED-LOOP AND AUTONOMOUS DRUG DELIVERY SYSTEMS

Predictive alerting still leaves the therapeutic action — adjusting a vasopressor rate, giving a fluid bolus, deepening or lightening anesthesia — in human hands. Closed-loop systems remove that step for a defined subset of decisions, using continuous feedback from a monitored variable (processed EEG for anesthetic depth, arterial pressure for vasopressor titration) to drive automated infusion pump adjustments without an intervening clinician action. The technology has now accumulated over a decade of randomized evidence, most extensively for closed-loop control of anesthetic depth using bispectral index or similar EEG-derived targets.

A multicenter randomized controlled trial of a closed-loop anesthesia delivery system [13] found that automated propofol titration to a target depth-of-anesthesia index kept patients within the target range a significantly greater proportion of case time than manual titration by an attending anesthesiologist, without any signal of harm. Subsequent work extended this comparison into specific high-risk populations: a trial in patients with moderate to severe left ventricular systolic dysfunction comparing closed-loop delivery against target-controlled infusion [14] found the closed-loop approach maintained more stable anesthetic depth in a population where hemodynamic swings from over- or under-dosing carry particular risk, and a randomized controlled study of automated propofol delivery during thoracic surgery [15] reported similar findings in a case type where one-lung ventilation and positioning changes complicate manual titration.

A 2026 systematic review and meta-analysis in the British Journal of Anaesthesia pooling closed-loop hypnotic drug delivery trials [16] synthesized this now-substantial trial base and found a consistent benefit in time-in-target-range for anesthetic depth, alongside secondary signals toward reduced total drug consumption and faster early recovery in some trials, though — consistent with the pattern seen for HPI — the meta-analysis found the evidence for reduced postoperative complications or length of stay considerably less robust than the evidence for the physiologic surrogate outcome. A broader review of advances in closed-loop anesthesia and automated drug delivery [17] situates the depth-of-anesthesia literature alongside emerging closed-loop work on neuromuscular blockade reversal and vasopressor titration, arguing that the field is moving from single-variable control toward multi-parameter systems that could eventually manage several concurrent physiologic targets, though multi-variable closed-loop control has, as of this writing, considerably less randomized trial support than single-variable depth-of-anesthesia control does.

INTEGRATED AND AI-ASSISTED SYSTEMS IN PRACTICE

The most recent generation of tools attempts to move beyond single-variable prediction or single-variable control toward integrated decision support spanning multiple concurrent aspects of an anesthetic — a dashboard, rather than a single alert. A comparative clinical evaluation of an AI-assisted decision support system for general anesthesia management, tested across six centers [18], reported improved adherence to protocolized targets across several concurrent domains (hemodynamics, depth of anesthesia, ventilation parameters) when clinicians had the integrated system available, compared with usual practice at the same centers before implementation — though as a comparative rather than randomized

design, the study cannot fully separate the effect of the software from secular improvement in practice patterns over the study period, a limitation the authors acknowledged.

A narrative review of the current landscape and future trajectory of artificial intelligence in anesthesia [19] takes a wider view, situating operating-room decision support within the broader arc of perioperative AI that also includes preoperative risk stratification and postoperative deterioration prediction, and argues that the intraoperative applications discussed in this review are the most mature because the data streams are richest and the feedback loop between prediction and action is shortest. A parallel review focused specifically on cardiac anesthesia [20] and a pragmatic assessment of AI integration in cardiac surgery [21] both make a related point about subspecialty adoption: cardiac anesthesiologists already work with dense continuous data streams from invasive monitoring and often cardiopulmonary bypass, which makes the marginal burden of adding an algorithmic layer lower than in lower-acuity cases, and correspondingly these authors report faster uptake and a somewhat richer evidence base in cardiac practice than in general surgical anesthesia. A broader review of smart technologies and digital innovation for perioperative patient safety [22] and a focused discussion of AI algorithms for next-generation intraoperative hemodynamic monitoring [23] both flag the same open technical question: as these systems begin drawing on more input channels simultaneously, the interpretability of any single alert or recommendation becomes harder for the bedside clinician to audit, which is precisely the property that determines whether an alert gets trusted or ignored.

A separate strand of literature, less algorithmically ambitious but methodologically important, has examined decision support and computing infrastructure in anesthesia more broadly. A review of recent advances in system management, decision support, and computing in anesthesia [24] catalogs the diversity of tools now deployed — from AIMS-embedded rule engines through predictive analytics to early autonomous systems — and makes the point that most operating rooms today run several of these tools concurrently, layered on top of one another, without much empirical study of how they interact. An anesthesiologist managing a case with an AIMS medication-interaction alert, an HPI hypotension warning, and a closed-loop depth-of-anesthesia system all running at once is operating in an environment no single validation trial has actually tested.

ALERT FATIGUE, TRUST, AND THE LIMITS OF AUTOMATION

Every category of tool surveyed above shares a common vulnerability: an alert or recommendation only changes behavior if the clinician receiving it still trusts it enough to act. Alert fatigue — the well-documented phenomenon in which clinicians habituate to and begin overriding or ignoring frequent, low-specificity warnings — is not unique to anesthesia, but the operating room's combination of task density, time pressure, and continuous physiologic monitoring makes it a particularly acute risk. The systematic review of point-of-care CDS in AIMS discussed earlier [3] found that alert specificity, not sensitivity, was the strongest predictor of whether an alerting system actually changed clinical behavior over time — a system tuned to catch every possible deviation quickly trains its users to dismiss it.

The HPI and machine-learning hypotension-prediction literature illustrates a second, subtler version of the same problem: several of the meta-analyses discussed above [10,11] found that reported predictive performance metrics varied considerably depending on how a "true positive" prediction was defined and validated, and that internally validated performance figures did not always reproduce in external or prospective testing. A clinician who has seen a predictive alert prove wrong a few times in a given case is unlikely to weight the next alert as heavily, regardless of the tool's published aggregate accuracy — which means the psychological reality of trust calibration, not just the algorithm's technical performance, determines real-world effectiveness.

Closed-loop systems raise a related but distinct question, less about trust calibration and more about the allocation of attention and responsibility. When an automated system is titrating anesthetic depth or vasopressor dose without moment-to-moment human input, the clinician's task shifts from continuous manual control to intermittent supervision — a mode of work that skilled-performance literature outside anesthesia has repeatedly shown humans are poor at sustaining, because vigilance for rare deviations degrades when the routine case is running itself. None of the closed-loop trials reviewed here were designed or powered to detect this kind of supervisory-attention degradation, since they measured time-in-target-range and short-term outcomes rather than clinician behavior during automated control, and this remains, in our view, one of the more consequential blind spots in the current evidence base — not whether the algorithm performs well when everything goes as expected, but what happens to human readiness to intervene during the rare moment it does not.

CONCLUSION

The decade's evidence on digital decision support in operating room anesthesia management supports a layered rather than a uniform conclusion. Narrowly scoped, rule-based alerts tied to a single well-defined action have the most consistent track record of changing process measures, and that is not a modest achievement — missed antibiotic redoses and overlooked drug interactions are exactly the kind of

lapse continuous software attention is well suited to catch. Predictive tools like the Hypotension Prediction Index reliably reduce the physiologic surrogate they were built to detect, but the meta-analytic evidence for translation into reduced postoperative organ injury remains weaker than the enthusiasm around the technology sometimes suggests, and validation in specific high-risk populations and positions is still catching up to broader clinical deployment. Closed-loop drug delivery has the deepest randomized evidence base of any category discussed here, at least for single-variable anesthetic-depth control, and appears to outperform manual titration on the metrics it targets without an apparent safety signal against it — though multi-variable, more autonomous systems have outrun their own evidence base and questions about sustained clinician vigilance during automated control remain largely unstudied. What ties all of these together is a shared methodological gap: process and physiologic surrogate outcomes have been measured far more rigorously than the hard clinical endpoints that ultimately justify bringing more software into the operating room, and closing that gap — through adequately powered, externally validated, prospective trials — is the field's next necessary step, not the addition of yet another predictive parameter to the display.

DECLARATIONS

Acknowledgements

The authors thank their respective departments for institutional support during the preparation of this review.

Funding

This review received no external funding.

Conflict of Interest

The authors declare no conflicts of interest.

Author Contributions

MB conceived the review and drafted the initial manuscript. BS performed the literature search and synthesized the sections on predictive and closed-loop systems. RF critically revised the manuscript for intellectual content and finalized the reference list. All authors approved the final version submitted.

Ethics Statement

This is a literature review that did not involve new human or animal data collection; no ethics committee approval was required.

Data Availability

No new data were generated or analyzed for this review. All sources discussed are publicly available published literature, cited in the references list.

REFERENCES

1. Development of a decision support system to assist anesthesiologists in operating room (1997) Journal of Medical Systems, DOI:10.1023/A:1005535224082.
2. Using real-time clinical decision support to improve performance on perioperative quality and process measures (2011) Anesthesia & Analgesia, PMID:21295752.
3. A systematic review of near real-time and point-of-care clinical decision support in anesthesia information management systems (2016) Journal of Clinical Monitoring and Computing, DOI:10.1007/s10877-016-9921-x.
4. An accessible clinical decision support system to curtail anesthetic greenhouse gases in a large health network: implementation study (2022) JMIR Perioperative Medicine, DOI:10.2196/40831.
5. Outcomes of clinical decision support systems in real-world perioperative care: a systematic review and meta-analysis (2024) systematic review and meta-analysis.
6. Prediction and prevention of intraoperative hypotension with the hypotension prediction index: a narrative review (2022) narrative review, PMID:36132057.
7. Impact of hypotension prediction index-guided management on intraoperative hypotension and postoperative outcomes in abdominal surgery: a meta-analysis of randomized controlled trials (2025) meta-analysis, PMID:41175934.
8. Hypotension prediction index in surgical patients: a systematic review and meta-analysis of randomized controlled trials (2025) systematic review and meta-analysis.
9. Comparison of intraoperative hypotension outcomes with and without hypotension prediction index tool in patients undergoing brain tumour surgery: a study protocol for a randomised controlled trial (2025) study protocol, PMID:41510867.
10. Clinical performance of machine-learning algorithms to predict intraoperative hypotension: a meta-analysis (2025) BMC Surgery, DOI:10.1186/s12893-025-03412-8.
11. Predictive ability of hypotension prediction index and machine learning methods in intraoperative hypotension: a systematic review and meta-analysis (2024) systematic review and meta-analysis, PMID:38911129.
12. Effectiveness of hypotension prediction index software in reducing intraoperative hypotension in prolonged prone-position spine surgery: a single-center clinical trial (2025) single-center clinical trial.
13. A multicenter evaluation of a closed-loop

anesthesia delivery system: a randomized controlled trial (2015) Anesthesia & Analgesia, PMID:25902324.

14. Anaesthetic depth control using closed loop anaesthesia delivery system vs. target controlled infusion in patients with moderate to severe left ventricular systolic dysfunction (2017) PMID:28865915.

15. Evaluation of automated delivery of propofol using a closed-loop anesthesia delivery system in patients undergoing thoracic surgery: a randomized controlled study (2020) Journal of Cardiothoracic and Vascular Anesthesia.

16. Closed-loop systems for automated hypnotic drug delivery during general anaesthesia: a systematic review and meta-analysis (2026) British Journal of Anaesthesia, PMID:42014224.

17. Closed-loop anesthesia: advances in automated drug delivery and perioperative monitoring (2026) review.

18. Clinical evaluation of an AI-assisted decision support system for general anesthesia management based on data from 6 centers: comparative study (2025) comparative study.

19. Transforming perioperative care: the current landscape and future trajectory of artificial intelligence in anesthesia - a narrative review (2026) narrative review.

20. Applications of artificial intelligence and machine learning in cardiac anesthesia across the continuum of perioperative care (2025) Academic Medicine & Surgery.

21. Beyond the hype: the pragmatic integration of intraoperative artificial intelligence in cardiac surgery (2026) review.

22. Smart technologies and digital innovations for improving perioperative patient safety: a review (2025) Patient Safety in Surgery, DOI:10.1186/s13037-025-00454-y.

23. AI algorithms - intraoperative hemodynamic monitoring of tomorrow? (2025) PMID:42009031.

24. Recent advances in system management, decision support systems, artificial intelligence and computing in anaesthesia (2023) systematic review, PMID:36967935.