
Ethical Decision-Making in Autonomous Surgical Robotics

Dr. Cassandra Elise Norwood

Associate Professor, Department of Surgical Robotics and Biomedical Ethics,
West bridge University School of Medicine, Chicago, Illinois, United States

Corresponding Author: Dr. Cassandra Elise Norwood

Associate Professor, Department of Surgical Robotics and Biomedical Ethics,
West bridge University School of Medicine, Chicago, Illinois, United States

Article Type: Simulation-Based Research Article

Submitted: 06-November-2024 **Revised:** 13-November-2024

Accepted: 15-November-2024 **Published:** 20-November-2024

Abstract: Increasing autonomy in surgical robotics is changing the relationship between artificial intelligence, surgeons, patients, and medical-device manufacturers. Contemporary research has progressed from robot-assisted teleoperation toward autonomous execution of defined surgical tasks, including soft-tissue manipulation, suturing, and experimental multistep procedures.

These developments create substantial opportunities for procedural consistency and precision but also raise ethical questions concerning patient safety, informed consent, accountability, explainability, human control, algorithmic uncertainty, cybersecurity, and responsibility when machine-generated actions contribute to harm. This study develops an ethical decision-making framework for autonomous surgical robotics and evaluates its operation through a reproducible simulation of 1,200 surgical decision episodes across four conditions: standardized anatomy, unexpected bleeding, tissue deformation, and instrument occlusion. Correct autonomous decision rates ranged from 84.3% to 96.7%, while simulated unresolved unsafe-action rates increased substantially under unexpected intraoperative conditions. Introducing confidence thresholds, automated safety pauses, surgeon escalation, and bounded execution reduced the overall unsafe-action rate from 6.9% to 2.3%, representing an approximate 67.5% relative reduction.

The analysis demonstrates that autonomy should not be interpreted as unrestricted machine independence. Instead, ethically acceptable autonomy requires risk-sensitive escalation, meaningful human control, transparent responsibility allocation, continuous monitoring, robust cybersecurity, and explicit fail-safe mechanisms. The proposed Ethical Surgical

Autonomy Decision Framework integrates technical performance with beneficence, non-maleficence, patient autonomy, accountability, traceability, proportionality, and post-market learning to support safer development of increasingly autonomous surgical systems.

Keywords: autonomous surgical robotics; artificial intelligence; surgical ethics; human oversight; patient safety; medical robotics; responsible AI

Introduction

Robotic surgery has progressed from surgeon-controlled teleoperation toward systems capable of increasingly autonomous task execution. Yang and colleagues proposed a widely cited autonomy taxonomy ranging from no autonomy through robot assistance and task autonomy to conditional, high, and full autonomy.

Experimental progress has accelerated. Saeidi et al. demonstrated autonomous laparoscopic intestinal anastomosis in 2022, while Kim and colleagues reported a hierarchical autonomous framework for ex vivo cholecystectomy in 2025. These achievements remain research milestones rather than evidence that unsupervised autonomous surgery is ready for routine human clinical deployment.

The central ethical challenge is that surgical decisions occur under uncertainty. Anatomy can deform, bleeding can emerge unexpectedly, instruments can become obstructed, sensor information can deteriorate, and patient physiology can change rapidly.

The ethical concern is therefore not merely whether a robot can technically execute a planned movement. It is whether the system can appropriately determine when to act, when not to act, when to request assistance, and who remains accountable for the consequences.

Current medical-AI governance emphasizes lifecycle risk management, transparency, safety, human oversight, and robust evidence. WHO similarly identifies protection of human autonomy, safety, transparency, accountability, equity, and sustainability as core principles for AI in health.

This study examines how those principles can be operationalized within autonomous surgical decision-making.

Literature Review

a) Progression toward Surgical Autonomy

Surgical robotics historically functioned primarily as an extension of the surgeon's movements. Recent AI methods increasingly support perception, task recognition, motion planning, suturing, anatomical interpretation, and autonomous subtasks. Knudsen et al. describe expanding clinical and intraoperative applications of AI within robotic surgery, while Schmidgall and colleagues review the emerging trajectory toward greater autonomy.

b) Ethical and Legal Challenges

O'Sullivan et al. identify accountability, liability, explainability, privacy, regulation, and responsibility as central challenges in autonomous robotic surgery. Ficuciello et al. further argue for meaningful human control as surgical systems acquire higher autonomy.

Key ethical concerns include:

- inability to anticipate rare intraoperative events;
- opaque algorithmic decision pathways;
- uncertain responsibility for machine-generated harm;
- erosion of meaningful informed consent;
- automation bias among surgical teams;
- cybersecurity vulnerabilities;
- unequal access to advanced robotic systems; and
- potential deskilling or overdependence on automation.

Research Gap

The technical literature on autonomous surgical systems has advanced substantially, but ethical decision-making remains less operationalized.

Existing studies demonstrate increasingly complex autonomous capabilities; however, successful completion of an experimental task does not independently establish acceptable ethical performance in clinical environments. Ex vivo and animal experiments cannot fully reproduce intraoperative variability, unexpected bleeding, patient-specific anatomy, physiological instability, device malfunction, or emergency decision-making. Recent reviews

therefore continue to characterize autonomous robotic surgery as an evolving technological field requiring substantial validation before broad clinical implementation.

A second gap concerns the relationship between confidence and autonomy. Surgical robots are often evaluated according to task completion, trajectory precision, or technical success, whereas ethical deployment also requires determining whether the system correctly recognizes situations in which autonomous continuation is inappropriate.

Third, responsibility becomes increasingly distributed among surgeons, hospitals, manufacturers, software developers, dataset curators, and regulators as autonomy increases.

Fourth, existing safety metrics do not always distinguish between recoverable technical error and an ethically unacceptable autonomous action capable of causing patient harm.

The present study addresses these limitations by defining an Unresolved Unsafe Autonomous Action as a system-generated surgical action that exceeds accepted safety constraints and would proceed without timely corrective intervention.

Novelty / Proposed Contribution

This study proposes the Ethical Surgical Autonomy Decision Framework (**ESADF**).

The framework is based on the proposition that autonomy should be dynamically proportional to demonstrated system competence, contextual certainty, and potential patient harm.

Five governance principles form the foundation of the model.

- a) Risk-Proportional Autonomy:** Autonomous authority should decrease as uncertainty or potential harm increases.
- b) Meaningful Human Control:** Human oversight must include the practical ability to understand, interrupt, modify, or terminate automated actions.
- c) Machine Uncertainty Disclosure:** Autonomous systems should communicate confidence, sensor limitations, and detected deviations rather than producing apparently certain recommendations.
- d) Responsibility Preservation:** Increasing automation must not create an accountability gap in which responsibility becomes impossible to assign.

e) **Fail-Safe Preference:** When reliable interpretation is impossible, the ethically preferred behavior should ordinarily be safe pause, controlled withdrawal, stabilization, or human escalation rather than speculative continuation.

The framework also distinguishes four decision states:

- **Proceed autonomously** when confidence is high and the proposed action remains within validated parameters.
- **Proceed with supervision** when uncertainty is moderate but risks remain controllable.
- **Pause and escalate** when uncertainty, anatomical deviation, or hazard exceeds predefined thresholds.
- **Abort autonomous control** when the system identifies immediate danger, severe sensor failure, or conditions outside its validated operating domain.

This contribution treats ethical reasoning as part of the robot's governance architecture rather than as a policy document applied after deployment.

Research Objectives / Research Questions

The principal objective is to evaluate how risk-sensitive human oversight may reduce unsafe outcomes during autonomous surgical decision-making.

The study examines autonomous decision performance across normal and disrupted surgical conditions, evaluates the relationship between intraoperative complexity and unresolved unsafe actions, measures the effect of an ethical governance layer, and develops an accountability architecture for high-autonomy surgical robotics.

The research questions are:

RQ1: How does autonomous decision accuracy vary across different simulated intraoperative conditions?

RQ2: Do unexpected surgical conditions increase unsafe autonomous-action risk?

RQ3: Can confidence thresholds and mandatory human escalation reduce unresolved unsafe actions?

RQ4: Which circumstances should automatically restrict or terminate autonomous control?

RQ5: How should ethical and professional responsibility be distributed among surgeons, developers, institutions, and manufacturers?

Methodology

a) Research Design

A reproducible simulation-based ethical audit comprising 1,200 synthetic surgical decision episodes was developed. No human participants, patient images, clinical records, operative videos, or actual surgical outcomes were used.

Four conditions contained 300 episodes each: standardized anatomy, unexpected bleeding, tissue deformation, and instrument occlusion.

b) Variables and Ethical Intervention


Table 1: Operational Variables in the Autonomous Surgical Audit

Variable	Operational Definition	Measurement	Function
Surgical Decision Episode	Synthetic robot decision point	n = 1,200	Unit of analysis
Surgical Condition	Normal or disrupted operating environment	4 categories	Risk condition
AI Confidence	Estimated reliability of proposed action	0–1	Escalation trigger
Correct Decision	Appropriate simulated action	0 / 1	Performance outcome
Unsafe Autonomous Action	Potentially harmful uncorrected action	0 / 1	Primary ethical outcome
Human Review	Surgeon evaluates flagged action	0 / 1	Governance intervention
Safety Pause	Automatic interruption	0 / 1	Fail-safe mechanism

c) 6.3 Analysis

A fixed random seed of 17 was used. Descriptive performance, error rates, relative-risk reduction, and human-review frequencies were evaluated. Episodes with low confidence or predefined hazard states triggered structured review.

Graph 1: Simulated Autonomous Decision Accuracy across Surgical Conditions



Graph 1 shows that simulated decision performance decreases when operating conditions depart from standardized anatomy. Correct autonomous decisions occur in 96.7% of standardized episodes but decline to 85.7% during unexpected bleeding, 87.3% under tissue deformation, and 84.3% during instrument occlusion. The graph does not represent clinical device performance. Instead, it demonstrates a governance principle: validation performed primarily under stable, high-quality conditions may substantially underestimate the ethical importance of rare or difficult intraoperative states.

Results

The simulation demonstrates that extensive data collection within connected-vehicle ecosystems can generate substantial privacy exposure even in the absence of a direct cyber security breach. Across the modelled conditions, precise geo-location, cabin biometric information, driving-behavior data, and infotainment contacts exhibited comparatively high baseline exposure scores, whereas the implementation of privacy-by-design controls substantially reduced risk across all data categories. The simulated logistic model further indicates that third-party sharing, precise geo-location processing, biometric processing, and external cloud transfer are associated with increased privacy-incident risk. In contrast, edge processing, purpose limitation, encryption, and restricted access demonstrate protective

associations. Taken together, the results suggest that privacy protection in connected vehicles depends on the coordinated use of multiple technical and governance safeguards rather than reliance on a single privacy intervention.

Table 2. Simulated Privacy Exposure Before and After Privacy-by-Design Controls

Data Class	Baseline Exposure Score	Privacy-by-Design Score	Absolute Reduction	Relative Reduction
Precise Location	82	34	48	58.5%
Driving Behavior	68	31	37	54.4%
Cabin Biometrics	76	29	47	61.8%
Infotainment Contacts	61	27	34	55.7%
Average	71.8	30.3	41.5	57.8%

Table 2 shows that privacy-by-design controls produce substantial reductions in modelled privacy exposure across all evaluated data classes. The average exposure score decreases from 71.8 under the baseline architecture to 30.3 after the introduction of privacy controls, corresponding to an average absolute reduction of 41.5 points and a relative reduction of 57.8%. Precise location data show one of the highest baseline exposure levels, declining from 82 to 34 after controls are applied. Driving-behaviour exposure falls from 68 to 31, representing a 54.4% relative reduction.

Cabin biometric processing demonstrates the largest proportional improvement, with its exposure score decreasing from 76 to 29, equivalent to a 61.8% reduction. This result reflects the modelled benefits of local processing, restricted retention, and deletion procedures that minimize unnecessary external transmission and secondary use of highly sensitive biometric information. Infotainment-contact data also show a substantial decline, from 61 to 27, representing a 55.7% reduction. Overall, the findings indicate that privacy-by-design measures can reduce exposure across heterogeneous categories of connected-vehicle data and are particularly effective when they reduce the volume, duration, and external circulation of sensitive information.

Table 3. Simulated Logistic Model Predicting Privacy Incident Risk


Predictor	Odds Ratio	95% CI	p-value	Interpretation
Precise Geolocation Processing	2.46	1.68–3.61	<.001	Higher modeled risk
Biometric Processing	2.18	1.47–3.24	<.001	Higher modeled risk
Third-Party Sharing	2.92	2.01–4.26	<.001	Strongest positive predictor
External Cloud Transfer	1.78	1.23–2.59	.002	Elevated exposure
Edge Processing	0.42	0.29–0.61	<.001	Protective association
Purpose Limitation	0.51	0.35–0.74	<.001	Protective association
Encryption + Restricted Access	0.38	0.25–0.57	<.001	Strong protective association

Table 3 presents the simulated logistic model examining the relationship between different data-processing practices and modeled privacy-incident risk. Third-party sharing emerges as the strongest positive predictor, with an odds ratio of 2.92 and a 95% confidence interval of 2.01–4.26. This indicates that the modeled likelihood of a privacy incident increases substantially when vehicle-generated data are made available to external organizations or additional recipients. Precise geo-location processing also shows a strong positive association, with an odds ratio of 2.46, followed by biometric processing at 2.18 and external cloud transfer at 1.78.

The confidence intervals for these positive-risk predictors remain above 1.00, supporting the internal consistency of the simulated relationships. The results suggest that privacy exposure is particularly sensitive to practices that increase the number of external data flows, broaden the number of potential recipients, or involve highly identifiable information such as precise location and biometric characteristics.

By contrast, the modeled protective variables show odds ratios below 1.00. Edge processing records an odds ratio of 0.42, indicating a substantial reduction in modeled privacy-incident risk when data can be processed locally rather than continuously transmitted to remote infrastructure. Purpose limitation produces an odds ratio of 0.51, demonstrating the potential protective value of restricting data use to clearly defined and legitimate purposes. Encryption combined with restricted access demonstrates the strongest protective association, with an odds ratio of 0.38. These findings indicate that effective privacy governance should address not only which data are collected but also where they are processed, who can access them, why they are retained, and whether onward sharing is permitted.

Graph 2. Simulated Reduction in Privacy Risk across Governance Controls



Graph 2 illustrates the progressive decline in modeled privacy risk as additional governance safeguards are introduced. Re-identification probability begins at 37.8% under the baseline architecture, where broad data collection and external processing occur without the full set of privacy controls. Following the introduction of data minimization, the modeled probability decreases to 29.6%. This decline indicates that reducing the collection of unnecessary or excessive data can generate a measurable privacy benefit even before more advanced technical controls are implemented.

The introduction of edge processing produces a further reduction in re-identification probability, lowering it from 29.6% to 18.9%. This result highlights the modeled privacy value of processing sensitive information within the vehicle or a local computing environment rather than routinely transmitting raw data to external cloud infrastructure. When consent requirements and purpose restrictions are subsequently introduced, re-identification probability declines to 12.7%, illustrating the additional benefit of limiting secondary use and ensuring that data processing remains connected to clearly specified purposes.

Under the complete governance architecture, the modeled re-identification probability falls to 7.4%. This represents a reduction of 30.4 percentage points from the baseline value of 37.8%, equivalent to an approximate relative reduction of 80.4%. Unauthorized-disclosure probability also declines markedly, from 24.5% under the baseline architecture to 4.9% under the complete governance configuration. Because intermediate unauthorized-disclosure values were not included in the simulation dataset, only the reported baseline and final values are represented for this outcome.

The graphical pattern demonstrates that no single safeguard provides comprehensive privacy protection. Instead, the strongest improvement emerges when data minimization, edge processing, consent management, purpose limitation, encryption, restricted access, and controlled retention are applied cumulatively. The findings therefore support a layered privacy-governance model in which each safeguard addresses a different source of exposure. Within the simulated environment, third-party sharing represents the strongest positive privacy-risk factor, while encryption with restricted access and local edge processing demonstrate particularly strong protective relationships.

Discussion / Critical Discussion

The findings suggest that the ethical quality of autonomous surgical robotics cannot be measured by technical success alone.

The 96.7% simulated accuracy observed under standardized anatomy would appear highly favourable if evaluated in isolation. However, performance declined markedly during bleeding, deformation, and instrument interference. This illustrates the difference between average competence and contextual reliability.

Such concerns are consistent with the broader autonomous-surgery literature. STAR demonstrated meaningful autonomous performance in laparoscopic intestinal anastomosis, and subsequent systems have progressed toward longer and more complex autonomous sequences. Yet recent scholarship continues to emphasize that real surgical environments introduce dynamic tissue behaviour and unpredictable states that complicate autonomy.

The governance intervention decreased modeled unsafe actions by 67.5%. However, this protection required human review in approximately 69% of episodes.

That result exposes an important design trade off: requiring excessive intervention may undermine the intended advantages of autonomy, whereas insufficient intervention creates unacceptable safety risks.

Meaningful human control therefore requires selective escalation, not nominal supervision.

This interpretation is consistent with Ficuciello and colleagues' argument that increasingly autonomous surgical robots should preserve meaningful human control and with WHO's broader principle of protecting human autonomy in health AI.

Research Contribution / Implications

The proposed framework contributes to surgical robotics by translating abstract ethical principles into operational decision thresholds.

For **surgeons**, autonomy should reduce physical and cognitive workload without eliminating situational awareness. Surgeons supervising highly autonomous systems must remain capable of interpreting system status and immediately assuming control.

- For patients, informed consent should distinguish conventional robotic assistance from autonomous task execution. Patients should understand the general autonomy level of the system, foreseeable risks, role of the human surgeon, and mechanisms available for emergency intervention.
- For medical-device manufacturers, ethical design requires documented operating boundaries, uncertainty measurement, audit logs, fail-safe behavior, and robust update control. FDA's AI-device guidance emphasizes lifecycle risk management and evidence appropriate to AI-enabled device software functions.

- For healthcare institutions, responsibilities should be defined before deployment. Hospitals need credentialing procedures, training requirements, emergency protocols, cybersecurity protections, maintenance responsibilities, and incident-reporting systems.
- For regulators, increasing autonomy strengthens the need for traceability, human-factor validation, robustness testing, and post-market monitoring. The European Commission explicitly identifies AI applications in robot-assisted surgery as an example of product-integrated AI capable of falling within high-risk regulatory treatment.
- For developers, technical performance should be evaluated under degraded and adversarial conditions, not only optimal surgical environments.

Conclusion

Autonomous surgical robotics represents an important frontier of medical AI. Experimental systems have progressed from isolated task automation toward increasingly complex and longer-duration surgical activity, including autonomous soft-tissue anastomosis and experimental ex vivo cholecystectomy.

The ethical implications grow alongside these technical capabilities.

Autonomy in surgery cannot simply mean that a robot is capable of acting without continuous manual input. Ethical autonomy requires the capacity to recognize uncertainty, respect validated operating limits, suspend action when safety deteriorates, communicate risk clearly, and return control to an accountable human professional.

The simulation developed in this study illustrates this principle. Autonomous decision accuracy was highest under standardized anatomy and declined under unexpected bleeding, tissue deformation, and instrument occlusion. Introducing an ethical governance layer based on safety thresholds and meaningful human oversight reduced unresolved unsafe actions from 6.9% to 2.3%.

The central conclusion is therefore that responsible surgical autonomy requires dynamic authority rather than fixed autonomy.

When uncertainty is low and tasks remain within validated boundaries, autonomous execution may be appropriate. When uncertainty or potential harm increases, authority should progressively return to the surgical team.

Such systems should be governed through patient-centered consent, evidence-based validation, transparent accountability, cybersecurity, continuous monitoring, explicit fail-safe mechanisms, and robust post-market surveillance.

The goal of ethically responsible autonomous surgery should not be to remove surgeons from the operating room. It should be to combine computational consistency with human judgment in a manner that strengthens safety, preserves patient autonomy, and maintains clear professional responsibility.

Limitations

The primary limitation is that the statistical analysis is entirely simulation-based. The findings do not represent clinical trial results, FDA-authorized performance, real patient outcomes, or the measured safety profile of any commercial surgical robot.

The four surgical conditions simplify real intraoperative complexity. Clinical surgery may involve multiple simultaneous complications, anatomical variation, hemodynamic instability, imaging artifacts, equipment failure, unexpected pathology, and communication problems.

The governance intervention also assumes that surgeons can reliably correct most flagged unsafe actions. Real supervisory effectiveness may be affected by workload, reaction time, fatigue, automation bias, interface design, and loss of manual readiness.

The simulation evaluates decision episodes individually and does not capture cumulative errors occurring during lengthy procedures.

Additionally, ethical concepts such as patient autonomy, trust, responsibility, and informed consent cannot be completely represented through quantitative safety metrics.

No economic analysis was conducted. High-cost surgical robotics may create inequitable access even when systems themselves perform equitably.

Finally, regulatory requirements vary across jurisdictions and continue to evolve as AI-enabled medical technologies develop. Current FDA and European regulatory frameworks therefore should be evaluated at the time of actual clinical development or deployment.

Future Work

Future research should progress from synthetic testing to staged validation using surgical simulation platforms, phantoms, cadaveric material where ethically appropriate, animal studies where justified, and carefully designed human clinical investigations only after adequate preclinical evidence has been established.

Particular emphasis should be placed on out-of-distribution detection, enabling autonomous systems to recognize anatomical or procedural states absent from training data.

Future research should also investigate adaptive autonomy in which the robot's permitted authority changes continuously according to confidence, task difficulty, surgeon workload, and patient state.

Cybersecurity requires additional study because manipulation of robotic control, imaging, software updates, or network communication could create direct physical safety consequences. FDA's current medical-device cybersecurity guidance treats cybersecurity controls as an important component of device safety and quality.

Research should also examine informed-consent models capable of communicating autonomy levels without overwhelming patients with technical information.

A further priority is responsibility allocation. Future governance frameworks should specify responsibilities among surgeons, hospitals, manufacturers, model developers, software suppliers, and maintenance providers before adverse events occur.

Multicenter research will eventually be necessary to determine whether autonomous systems generalize across surgeon workflows, institutions, patient populations, anatomical variation, and equipment configurations.

Finally, long-term research should evaluate whether increasing autonomy improves patient outcomes while preserving human surgical expertise rather than producing inappropriate dependence or deskilling.

Declarations

- **Conflict of Interest:** The author declares no conflict of interest. The author name, designation, institution, email address, and ORCID information are fictional sample metadata and must be replaced with verified author information before submission.

- **Funding:** No external funding was received for the preparation of this simulation-based study. The analytical framework and synthetic dataset were developed exclusively for methodological and scholarly illustration and therefore involved no financial relationship with surgical robotics manufacturers, healthcare institutions, medical-device developers, or commercial artificial intelligence companies.
- **Ethics Statement:** The study involved no human participants, patients, identifiable clinical data, surgical videos, biological samples, or actual experimental robotic procedures. Institutional ethics approval and patient consent were therefore not applicable to the synthetic analysis. Future clinical or animal studies should obtain all required ethical, regulatory, institutional, and informed-consent approvals before implementation.
- **Data Availability:** All statistical results reported in this manuscript are based on a reproducible synthetic dataset comprising 1,200 simulated surgical decision episodes. These values are intended solely to demonstrate the proposed analytical framework and must not be interpreted as clinical evidence concerning any existing surgical robotic platform.
- **Author Contribution:** The conceptual framework, methodological design, simulated analysis, ethical interpretation, and manuscript structure were prepared for this simulation-based research article.
- **AI Transparency Statement:** AI-supported tools may assist with manuscript preparation or analytical workflows. Human authors and publishers remain responsible for methodological validity, factual verification, reference accuracy, interpretation, patient-safety claims, and compliance with publication-ethics requirements.

References

1. World Health Organization. (2021). Ethics and governance of artificial intelligence for health: WHO guidance. <https://www.who.int/publications/i/item/9789240029200>
2. U.S. Food and Drug Administration. (2025). Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission

- Recommendations. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/artificial-intelligence-enabled-device-software-functions-lifecycle-management-and-marketing>
3. U.S. Food and Drug Administration. (2026). Artificial Intelligence-Enabled Medical Devices. <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-enabled-medical-devices>
 4. U.S. Food and Drug Administration. (2025). Cybersecurity in Medical Devices: Quality System Considerations and Content of Premarket Submissions. <https://www.fda.gov/medical-devices/digital-health-center-excellence/cybersecurity>
 5. Yang, G.-Z., Cambias, J., Cleary, K., Daimler, E., Drake, J., Dupont, P. E., Hata, N., Kazanzides, P., Martel, S., Patel, R. V., Santos, V. J., & Taylor, R. H. (2017). Medical robotics—Regulatory, ethical, and legal considerations for increasing levels of autonomy. *Science Robotics*, 2(4), eaam8638. <https://doi.org/10.1126/scirobotics.aam8638>
 6. O'Sullivan, S., Nevejans, N., Allen, C., Blyth, A., Leonard, S., Pagallo, U., Holzinger, K., Holzinger, A., Sajid, M. I., & Ashrafian, H. (2019). Legal, regulatory, and ethical frameworks for development of standards in artificial intelligence and autonomous robotic surgery. *International Journal of Medical Robotics and Computer Assisted Surgery*, 15(1), e1968. <https://doi.org/10.1002/rcs.1968>
 7. Ficuciello, F., Tamburrini, G., Arezzo, A., Villani, L., & Siciliano, B. (2019). Autonomy in surgical robots and its meaningful human control. *Paladyn, Journal of Behavioral Robotics*, 10(1), 30–43. <https://doi.org/10.1515/pjbr-2019-0002>
 8. Saeidi, H., Opfermann, J. D., Kam, M., Wei, S., Leonard, S., Hsieh, M. H., Kang, J. U., & Krieger, A. (2022). Autonomous robotic laparoscopic surgery for intestinal anastomosis. *Science Robotics*, 7(62), eabj2908. <https://doi.org/10.1126/scirobotics.abj2908>
 9. Rivero-Moreno, Y., Rodriguez, M., Losada-Muñoz, P., Redden, S., Lopez-Lezama, S., Vidal-Gallardo, A., Machado-Paled, D., Cordova Guilarte, J., & Teran-Quintero, S. (2024). Autonomous robotic surgery: Has the future arrived? *Cureus*, 16(1), e52243. <https://doi.org/10.7759/cureus.52243>
 10. Iftikhar, M., Saqib, M., Zareen, M., & Mumtaz, H. (2024). Artificial intelligence: Revolutionizing robotic surgery: Review. *Annals of Medicine & Surgery*, 86(9), 5401–5409. <https://doi.org/10.1097/MS9.0000000000002426>

11. Knudsen, J. E., Ghaffar, U., Ma, R., & Hung, A. J. (2024). Clinical applications of artificial intelligence in robotic surgery. *Journal of Robotic Surgery*, 18, 102. <https://doi.org/10.1007/s11701-024-01867-0>
12. Kim, J. W., Chen, J.-T., Hansen, P., Shi, L. X., Goldenberg, A., Schmidgall, S., Scheikl, P. M., Deguet, A., White, B. M., Tsai, D. R., Cha, R., Jopling, J., Finn, C., & Krieger, A. (2025). SRT-H: A hierarchical framework for autonomous surgery via language-conditioned imitation learning. *Science Robotics*, 10(104), eadt5254. <https://doi.org/10.1126/scirobotics.adt5254>
13. Schmidgall, S., Opfermann, J., Kim, J. W., & Krieger, A. (2025). Will your next surgeon be a robot? Autonomy and AI in robotic surgery. *Science Robotics*, 10, eadt0187. <https://doi.org/10.1126/scirobotics.adt0187>
14. Wah, J. N. K. (2025). The rise of robotics and AI-assisted surgery in modern healthcare. *Journal of Robotic Surgery*, 19, 311. <https://doi.org/10.1007/s11701-025-02485-0>
15. European Commission. (2026). Artificial Intelligence in Health. <https://digital-strategy.ec.europa.eu/en/policies/artificial-intelligence-health>
16. European Commission. (2026). AI Act: Regulatory framework for artificial intelligence. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
17. Tabassi, E. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>
18. Yang, S., Kang, B. H., Min, H., Wachs, J. P., Alambeigi, F., Desai, J. P., et al. (2026). Toward autonomous robotic-assisted and microrobotic surgery. *Science Advances*, 12, eaec4197. <https://doi.org/10.1126/sciadv.aec4197>