

Autonomous Surgical Robots: The Malpractice-Product Liability Boundary

Bükre Şeyma Dönder* / Lect. (Ph.D.) 

International University of Sarajevo
bukreseymadonder@gmail.com

Ahmet Berke Oğurlu / MSc. Stu. 

Queen Mary University, Robotics and Artificial Intelligence
berkeogurlu@gmail.com

*Corresponding Author

Abstract

The increasing autonomy of surgical robots challenges a foundational distinction in tort law: the boundary between professional negligence and defective products. In conventional robot-assisted surgery, operative decision-making largely remains with the surgeon, making medical-malpractice doctrine a relatively natural framework. Autonomous surgical systems complicate this structure by shifting elements of perception, planning, technical execution, and potentially clinical decision-making from clinicians to software-driven robotic systems. This article develops an Autonomy-Control-Risk Attribution (ACRA) framework for determining when liability should remain centered on medical malpractice, shift toward product liability, or be shared among clinicians, healthcare institutions, and manufacturers. ACRA examines the harm-producing autonomous function, the existence of meaningful rather than merely nominal human control, the provenance of the materialized risk, informational control, and interaction-based failures. Where clinicians retain effective control over the relevant clinical risk, malpractice should remain the primary liability regime. Where harm arises from hardware, software, algorithms, warnings, or autonomous behavior that clinicians

could not reasonably detect or control, product-side liability should gain greater weight. Where control is distributed, concurrent liability is preferable to forcing the injury into a single doctrinal category. The framework aims to preserve genuine professional accountability without turning surgeons into a moral crumple zone for machine-generated decisions.

Keywords: Autonomous Surgical Robots, Medical Malpractice, Product Liability, Meaningful Human Control, Artificial Intelligence.

JEL Codes: K13, I18, O33

Highlights

- Liability should follow meaningful control over the risk that materialized.
- ACRA distinguishes task-specific autonomy from system-wide autonomy.
- Nominal human oversight does not by itself establish surgeon control.
- Concurrent liability is appropriate for distributed human-machine failures.

Citation: Dönder, B. Ş., & Oğurlu, A. B. (2026). Autonomous surgical robots: The malpractice-product liability boundary. *Researches on Multidisciplinary Approaches (ROMAYA Journal)*, 2026(3), 1-15

1. Introduction

Robotic surgery has become an interdisciplinary field combining precision mechanical systems, medical imaging, real-time control, computer vision, and artificial-intelligence-assisted decision processes. Surgical robots are no longer merely mechanical aids; they are increasingly complex systems with evolving capabilities in perception, decision support, and controlled execution (Yang et al., 2017: 1; Attanasio et al., 2021: 651).

In current clinical practice, most FDA-cleared surgical robots still operate at low levels of autonomy under surgeon supervision, although some systems are approaching limited task autonomy and conditional autonomy. It has been found that most surgical robots were at Level 1 (86%) and some reached Level 3 (Conditional Autonomy) (6%). (Lee et al., 2024: 1–3). Experimental research further demonstrates that robotic systems can cross important technical thresholds in tasks such as autonomous suturing, tissue tracking, surgical planning, and laparoscopic intestinal anastomosis. Autonomous robotic surgery has the potential to provide effectiveness, safety, and reliability independent of individual surgeon's skill and experience (Saeidi et al., 2022: 1).

Increasing autonomy raises not only questions of liability but also issues of informed consent, data protection, privacy, and algorithmic explainability. In AI-enabled healthcare, robotic surgery is becoming an important component of the digital-health ecosystem (Küçükali, 2025). Fully or partially autonomous surgery may also generate legal questions concerning health data, AI models, and decision processes. This article, however, is limited to private-law liability arising from the use of autonomous surgical robots.

Robotic surgery increasingly destabilizes a basic distinction on which tort and medical-liability law have traditionally relied. When a physician performs surgery negligently, medical malpractice provides the natural framework. When a defective medical device injures a patient, product liability ordinarily directs attention to the manufacturer and the condition of the product. An autonomous surgical robot touches both categories. It is a manufactured medical product composed of hardware, software, algorithms, sensors, and control systems, while also increasingly performing tasks historically associated with the surgeon's professional judgment and physical skill.

This overlap is already visible in contemporary robotic-surgery disputes. Negligence claims may be brought against surgeons and hospitals while manufacturers simultaneously face allegations of design, manufacturing, or warning defects. Robotic surgery therefore creates a hybrid field between traditional medical malpractice and product liability because professional healthcare is delivered through a technological system designed, programmed, updated,

maintained, and supported by actors outside the physician-patient relationship (Cho, 2026: 302–303). Autonomous AI-based medical devices assume the dual role of a medical product and a medical practitioner, so responsibility may, depending on the circumstances, extend to a manufacturer, operator, or maintenance provider (Grossbard, 2025: 313).

The central problem is determining the point at which an injury should cease to be characterized primarily as a failure in the delivery of medical care and instead be treated as a failure of the technological system performing that care. Pasquale distinguishes between complementary and substitutive automation. When AI or robotics merely assist a skilled medical professional, they function as tools, and negligence remains the more appropriate standard of liability. By contrast, when AI and robotics replace rather than merely assist the professional, the burden of responsibility shifts toward the vendor of the computational system. Where the physician can no longer override the system, a stricter liability standard, including strict liability, becomes more appropriate (Pasquale, 2022: 5). In current and near-term systems, professional and institutional control often remains significant, so negligence may continue to govern mixed-control cases, while genuine manufacturing, design, and warning defects remain within product liability (Cho, 2026: 302–303).

The literature on meaningful human control adds another dimension. Surgical autonomy is not binary but graduated: systems range from tools that directly follow human commands to systems capable of generating operative strategies or conducting an entire procedure. At intermediate levels, human oversight may be continuous, intermittent, or conditional (Amoroso & Tamburrini, 2021: 31). The surgeon's physical presence in the operating room therefore does not, by itself, establish legally meaningful control.

The source of the harm is similarly central to liability in robotic surgery. Depending on that source, the physician, hospital, and manufacturer may be individually or concurrently responsible. Malpractice and product liability therefore need not operate as wholly separate categories in robotic-surgery disputes (Parlak Börü, 2019: 758).

Building on these approaches, this article argues that liability should be determined neither by the robot's legal characterization nor by its abstract level of autonomy alone. Instead, liability should follow the distribution of meaningful control over the concrete risk that materialized. The central question is: who controlled, or reasonably should have controlled, the risk that caused the patient's injury? This requires distinguishing the system's general level of autonomy from its functional autonomy within the causal chain that produced the harm.

To operationalize this proposition, the article propo-

ses the Autonomy-Control-Risk Attribution (ACRA) framework. It asks four foundational questions: (1) which autonomous function caused the harm; (2) who exercised meaningful control over that function; (3) within which actor's legally relevant risk sphere did the materialized risk originate; and (4) which actor possessed the information and practical capacity necessary to prevent the harm?

The analysis is organized around three research questions. First, at what point should the transfer of surgical decision-making authority from surgeon to machine shift liability from malpractice toward product liability? Second, how should meaningful human control, algorithmic opacity, and information asymmetry affect that boundary? Third, how should law allocate liability for mixed human-machine failures without turning the surgeon into the residual bearer of responsibility merely because a human remains nominally "in the loop"?

2. Method and Scope

This article adopts a doctrinal and comparative-functional approach. Tort-law scholarship on autonomous medical systems, surgical-robot liability, meaningful human control, and informed consent is examined together with empirical findings on perceptions of responsibility as surgical systems become more autonomous. Because much of the existing scholarship discusses negligence, medical malpractice, and product liability through the United States legal framework, the core doctrinal analysis is strongly informed by U.S. tort law. The discussion is supplemented by Chinese judicial practice concerning robot-assisted surgery and by broader interdisciplinary literature on surgical autonomy.

The article does not claim that a single uniform liability rule exists internationally. Product-liability doctrine, professional standards of care, procedural burdens of proof, and the organization of healthcare vary across jurisdictions. Rather, it identifies a functional boundary problem common across legal systems: technological autonomy redistributes control, while legal responsibility does not necessarily redistribute at the same pace.

This tension is particularly visible in China. Reviews of Chinese decisions indicate that courts tend to treat robot-assisted surgery as a surgical method or tool used by physicians rather than as an independent legal actor performing the operation. Claims are therefore generally directed against hospitals, and judicial assessment tends to evaluate the medical treatment as a whole rather than independently analyzing the robot's intelligence or programming as a source of harm (Shentu, 2024: 5). A tool-based classification can operate reasonably well while robots remain predominantly assistive, but it becomes increasingly inadequate as operative control is transferred to autonomous systems.

3. From Surgical Tool to Autonomous Decision-Maker

3.1. Autonomy Is a Spectrum, Not a Category

Technological developments in artificial intelligence and robotics have significantly transformed the role of machines in modern healthcare, creating new opportunities as well as new legal and regulatory challenges. In this context, a first analytical error in allocating liability for robotic surgery is to classify procedures merely as "robotic" or "non-robotic." The legally significant issue is not simply whether a robot participated in the operation, but which decisions it was authorized to make and the extent of its authority to execute those decisions. A robotic system can be considered as 'autonomous' at given tasks when it can perform them without further human intervention (Amoroso & Tamburrini, 2021: 23).

A six-level taxonomy illustrates this progression. At L0, the robot merely responds to human commands. At L1, it may mechanically assist, correct, or constrain human actions. At L2, it independently performs a task specified by a human. At L3, it may generate the overall surgical strategy subject to human selection or approval. At L4, it may perform an entire procedure under human supervision. At L5, the robot performs the medical procedure without human supervision (Amoroso & Tamburrini, 2021: 31). This classification also demonstrates a relationship between increasing autonomy and the allocation of legal responsibility (Grossbard, 2025: 314). As autonomy increases, the surgeon's practical decision-making and intervention space may narrow, although this fact alone does not establish that liability has shifted to the manufacturer.

Traditional teleoperated robotic surgery provides the clearest example of technologically mediated but human-controlled surgery. In the da Vinci system, the surgeon remains the operative decision-maker, issuing commands from the console that are reproduced by robotic instruments in the operative field (Koyuncu Aktaş, 2021: 1709). The recent emergence of humanoid surgical robots does not necessarily alter this legal characterization. Recent *in vivo* demonstrations of humanoid robotic surgery, for example, have remained teleoperated rather than autonomous (Liang et al., 2026). So long as a humanoid platform remains teleoperated, the physical displacement of the surgeon from the operating table does not by itself transfer decisional control to the robot. Robots should be assessed according to their functions, level of autonomy, and interaction with humans rather than their external appearance, a distinction that is particularly relevant to liability law (Parlak Börü, 2019: 760). Anthropomorphic embodiment may therefore make a robot appear to replace the surgeon physically while its decisional autonomy remains limited.

The robot's level of participation in a procedure is not sufficient by itself. Comparing autonomous execution in rigid anatomical structures with autonomous operations on deformable soft tissue shows that systems placed in the same autonomy category may nevertheless differ substantially in their working environments and adaptation requirements, thereby demanding different degrees of human oversight and care. A single model of human control cannot adequately govern every L2 surgical robot (Amoroso & Tamburrini, 2021: 32).

The taxonomy is therefore legally useful but not dispositive. Autonomy must be assessed together with the specific harm-producing function, the operative environment, the nature of the hazard, and the realistic availability of human intervention.

3.2. Complementary and Substitutive Autonomy

In robotic surgery, complementary robotics preserves human professional expertise as the operative source of medical decision-making while improving or supporting professional performance. The da Vinci system does not itself complete the surgery; rather, it functions as a highly sophisticated tool deployed by a skilled surgeon, who remains responsible for directing the procedure (Pasquale, 2022: 6). When AI or robotics merely assist a professional, they function as tools. By contrast, when they replace rather than merely assist a skilled medical professional, the automation becomes substitutive and responsibility should shift accordingly (Pasquale, 2022: 3). This distinction provides an important liability insight. So long as robotic technology remains an assistive tool and competent human intervention is realistically available, human negligence remains a strong basis of responsibility. When the system displaces a decision that would otherwise be made by the surgeon, responsibility should increasingly be directed toward those providing the substitutive technology (Pasquale, 2022: 5).

The complementary/substitutive distinction should be applied at the level of particular tasks rather than to the surgical system as a whole. Within the same procedure, patient selection may remain entirely human, planning may be shared, instrument positioning or suturing may be autonomous, and emergency conversion authority may return to the surgeon. A single robot may therefore function, during one operation, as a professional tool, a decision-support system, and an autonomous operative actor. Liability analysis should follow these functional transitions.

4. Medical Malpractice and the Persistence of Professional Responsibility

Medical malpractice generally refers to the principle that healthcare providers are required to exercise an

appropriate standard of care when providing medical treatment to patients. Where harm in robotic surgery arises from conduct that remains under the meaningful control of a healthcare professional, medical malpractice provides the appropriate starting point for liability analysis.

The conventional elements required to prevail in a medical malpractice claim are a legal duty of care and breach of that duty, proximate causation, and resulting injury or damages. Malpractice thus remains fundamentally a human-centric cause of action (Grossbard, 2025: 305).

The technological mediation of the surgeon's actions does not eliminate professional duties. A surgeon's possible negligence may include misusing the surgical robot, failing to convert to open surgery promptly when complications require manual intervention, or inadequately informing the patient about the surgeon's own experience with the robot or the differences in risk between robotic and traditional surgery (Cho, 2026: 323). The professional standard of care in robotic surgery also extends beyond ordinary surgical competence. A surgeon who elects to use a robotic system should understand its operation sufficiently to recognize foreseeable malfunctions and be prepared, when necessary, to continue or convert the procedure using an appropriate alternative technique. Familiarity with the robotic system and preparedness for malfunction are therefore components of professional care (Koyuncu Aktaş, 2021: 1712).

Hospitals also occupy an important institutional position. Selection of robotic surgery, credentialing of users, maintenance of equipment, development of protocols, and provision of the operative environment and infrastructure are institutional functions. Inadequate training or credentialing, failure to satisfy maintenance requirements, insufficient safety measures, and the absence of emergency procedures may therefore give rise to hospital negligence (Cho, 2026: 322). The relationship between manufacturer duties and institutional responsibility is illustrated by *Taylor v. Intuitive Surgical, Inc.* In that case, the Washington Supreme Court held that the manufacturer of the da Vinci system had a duty to warn the purchasing hospital, rather than treating warning to the surgeon alone as necessarily sufficient. The court emphasized the hospital's institutional role, including its use of safety information in credentialing and patient care. The decision demonstrates that responsibility for safe robotic surgery cannot always be reduced to the surgeon-manufacturer relationship; the hospital may itself occupy an important position in the flow and use of safety information.

High machine autonomy does not by itself eliminate malpractice. Even after a task is delegated to an autonomous system, professional negligence may consist of selecting an unsuitable patient or operating mode, failing to monitor an observable compli-

cation, disregarding a system warning, approving an evidently irrational strategy, or delaying conversion when a different operative technique is required.

There is nevertheless an important limit. The surgeon's duty of supervision can justify malpractice liability only when supervision is meaningful in practice. Mere physical capacity to intervene should not make the surgeon responsible for every algorithmic action. The possibility of intervention and the realistic ability to recognize when intervention is required are distinct questions.

At L3, meaningful supervision must involve more than a formal ability to press an "approve" button. The surgeon should be capable of understanding the rationale for the proposed strategy, comparing the available options, and deciding within an adequate time window. Problems of interpretability may therefore weaken the reality of human control (Amoroso & Tamburrini, 2021: 32). A surgeon may be held responsible for harm caused by an autonomously performing robot if the surgeon failed to exercise meaningful human control properly and the harm could have been avoided had the relevant meaningful human control duties been carefully fulfilled. Conversely, retrospective responsibility may be rebutted, and potentially shifted toward other actors, where it is shown that the surgeon properly complied with the applicable meaningful human control duties (Amoroso & Tamburrini, 2021: 33).

The doctrinal proposition is accordingly narrow but important: a surgeon should not be deemed to possess control merely because the system designates the surgeon as a supervisor. Malpractice based on inadequate supervision presupposes a realistic opportunity to understand, detect, and alter the risk-producing conduct.

5. Product Liability and the Autonomous Surgical System

5.1. Conventional Product Defects

Under Turkish law, a defect in a contract of sale may be understood as a material, legal, or economic deficiency that is inconsistent with the quality of the sold item or with a quantity affecting that quality, and that eliminates or substantially reduces its value for the intended use or the benefit that the buyer may reasonably expect from it (Kılıçoğlu, 2023: 74; Gül, 2026: 768). Defects may be classified according to different criteria; in terms of detectability, Turkish doctrine distinguishes between apparent, discoverable, latent, and fraudulently concealed defects (Kaşak, 2023: 87). An apparent defect is one that the buyer can perceive through ordinary attention without any special examination (Gümüş, 2024: 139; see also Eren & Yücer Aktürk, 2024: 334; Gül, 2026: 768; Kaşak, 2023: 84 et seq.). A discoverable defe-

ct is one that can be identified upon examination of the sold item after delivery (Eren & Yücer Aktürk, 2024: 112, 334; Kaşak, 2023: 85 et seq.). A latent defect, by contrast, cannot be detected through an ordinary examination and becomes apparent only later, often during use of the sold item (Yavuz, Acar & Özen, 2019: 80; Eren & Yücer Aktürk, 2024: 112, 335; Gül, 2026: 768). Where the seller intentionally conceals the defect by presenting the sold item as possessing a quality that it does not in fact have, the defect may be characterized as fraudulently concealed (Eren & Yücer Aktürk, 2024: 112, 336; Yıldırım, 2006: 968 et seq.). This sales-law conception of defect should not, however, be equated with the notion of a defective product in product-liability law. The former primarily concerns the conformity of the sold item with contractual or legally expected qualities, whereas the latter focuses on the safety-related characteristics of the product and the risks it poses.

When patient injury originates in the technological system rather than professional judgment, conventional product-liability doctrines and the concept of defect come to the forefront.

Product-liability analysis commonly distinguishes manufacturing defects, design defects, and failures to warn. A manufacturing defect concerns a product that deviates from its intended design or specifications; a design defect concerns a defective condition that poses unreasonable dangers to users or risks that could have been mitigated through a reasonable alternative design; and a failure to warn claim concerns inadequacy in warning about inherent product dangers or in providing instructions for safe use (Grossbard, 2025: 307).

Conventional manufacturing-defect analysis can generally be applied without conceptual modification. Defective wiring or insulation, malfunctioning circuit boards, and assembly errors are examples of manufacturing defects that may occur in surgical robots. A claim that the device deviated from its intended design and that the deviation caused the injury resembles the conventional defective-medical-device paradigm rather than a claim centered on professional malpractice (Cho, 2026: 313). The evidentiary difficulty of establishing a product defect in robotic surgery is illustrated by *Mracek v. Bryn Mawr Hospital*. The case arose after an alleged malfunction of the da Vinci surgical system, but the product-liability claims failed because the plaintiff did not provide sufficient expert evidence to establish a defect and causal connection to the injury. The case demonstrates that identifying a robotic malfunction is not sufficient by itself; the claimant must still establish the legally relevant defect and causation.

Consider a surgeon who correctly commands a robotic instrument to coagulate tissue, but defective insulation causes an electrical arc to burn adjacent tissue. The surgeon's participation does not trans-

form the insulation defect into malpractice. The legally significant point is that the relevant risk originated earlier in the manufacturing process.

An international empirical study reflects a similar intuition. When the surgeon remained the primary medical decision-maker but the robotic system supplied incorrect information, 69.6% of respondents attributed primary responsibility to the manufacturer. When the robot incorrectly executed a preprogrammed surgical plan, 63.4% likewise identified the manufacturer as primarily responsible (Jamjoom et al., 2022: 5). These responses do not determine legal liability, but they illustrate an intuitive distinction between a clinician selecting the wrong action and a machine failing to perform the action assigned to it. Recent litigation also confirms that injuries associated with robotic surgical systems may proceed through conventional product-liability doctrine. In *Franklin v. Intuitive Surgical, Inc.*, the federal district court allowed product-liability claims concerning injuries allegedly associated with robotic surgical devices to proceed beyond the pleading stage. The decision did not establish that the devices were defective or impose final liability on the manufacturer; rather, it illustrates that robotic-surgery injuries may present sufficiently plausible product-side allegations to warrant further judicial examination.

Turkish legal scholarship likewise treats the surgical robot as a product for product-liability analysis. Recognizing the robot as a product makes manufacturer liability available when injury originates in a defective robotic system (Koyuncu Aktaş, 2021: 1716). The malpractice-product liability boundary is therefore not peculiar to U.S. tort law: similar questions arise in Turkish private law once the surgical robot is simultaneously understood as an instrument of medical performance and as a product placed on the market. This distinction has become even more significant under the revised EU Product Liability Directive (EU) 2024/2853, which expressly brings software, including AI systems, within the concept of a “product” for product-liability purposes and treats software developers, including AI system providers, as manufacturers (Directive (EU) 2024/2853, recital 13).

5.2. Autonomous Software and the Temporal Instability of Defect

More complex cases shift attention from broken hardware toward design and algorithmic behavior. Surgical robots are not necessarily static products: software may be updated after sale, and learning systems may alter their behavior in response to newly acquired data. These developments blur the boundary between the initial design and later maintenance. A software update may introduce a new risk, and machine-learning behavior at the time of

injury may differ from behavior when the product first left the manufacturer's control (Cho, 2026: 314).

As AI-enabled devices continue to evolve beyond the design and foresight of their original manufacturers, potential risks can only be anticipated to a limited extent. This uncertainty complicates the application of conventional manufacturing, design, and warning categories and raises questions about when defectiveness should be assessed (Grossbard, 2025: 307). A model may have been trained by the manufacturer, modified through an authorized update, adapted using institutional data, or altered in practical effect by a hospital's failure to maintain the system. In such circumstances, there may no longer be a single legally meaningful “product state.” Product analysis should therefore distinguish among initial design, permitted evolution, and local modification or maintenance.

Manufacturer-side responsibility is strongest when harmful behavior is attributable to architecture, training, updates, or known limitations within the manufacturer's control. By contrast, institutional negligence becomes more salient where harm results from neglected maintenance, unauthorized reconfiguration, or failure to install required updates.

Unpredictability should not automatically sever responsibility. Machine learning does not necessarily break the causal connection between designers and the system, because apparently surprising behavior may still occur within objectives and constraints determined by design. Dangerous unpredictability may therefore indicate inadequate safety rather than eliminate liability (O'Sullivan et al., 2019: 6).

The problem becomes more difficult as autonomy and self-learning capacity increase. Errors may arise from incorrect information supplied to the decision algorithm, defects in underlying programming, or harmful behavioral patterns subsequently acquired through learning. The first two remain comparatively traceable to the producer, while the third creates a more difficult attribution problem. A robot may also be non-defective when placed on the market but later develop harmful behavior through information acquired during use, challenging traditional product-liability models because the harmful characteristic may not have existed at market placement (Koyuncu Aktaş, 2021: 1724–1725). For autonomous surgical robots, the legally relevant inquiry may therefore shift from whether the product was defective when sold toward whether the manufacturer reasonably controlled, monitored, or constrained the process through which the system was permitted to evolve.

A further difficulty arises when the harmful component is software rather than physical hardware. In Turkish product-liability scholarship, a threshold issue for AI-related responsibility is whether software elements fall within the legal concept of a “product” (Yünlü, 2024: 536). This question is particularly

important for autonomous surgical robots because the physical device may operate exactly as intended while the injury originates entirely in perception, planning, or control software.

6. The Boundary Problem: When Does Malpractice Become Product Liability?

The hardest cases are neither purely malpractice nor purely product defect. Harm may result from technological defect, negligent professional use, institutional failure, or a combination of these factors. Injuries can result from human errors or machine malfunctions, therefore the appropriate scheme of recovery must be designed to reflect the specific cause of harm (Grossbard, 2025: 274). A manufacturer may allege misuse; the surgeon may point to defective automation; and the hospital may invoke inadequate manufacturer instructions. In current systems, the structure of robotic surgery means that product safety depends on expert implementation. In such cases, the assumptions underlying product liability may not hold where the end user is a sophisticated user, which could lead courts to treat some issues as medical negligence rather than product liability (Cho, 2026: 312). Yet as autonomy increases, reliance on negligence alone becomes progressively harder to justify. Existing robotic-surgery cases already reveal different dimensions of this boundary problem. *Mracek* demonstrates the difficulty of proving defect and causation where an alleged robotic malfunction is involved, while *Taylor* shows that manufacturer duties may extend beyond the individual surgeon to the healthcare institution receiving and using product-safety information. *Franklin*, although only a pleading-stage decision, further illustrates that robotic-surgery injuries may support product-liability claims without first being characterized exclusively as medical malpractice. These cases concern robot-assisted rather than fully autonomous surgery, but they expose the fragmented allocation of duties that becomes more difficult as autonomous decision-making increases. Pasquale warns that doctrines emphasizing competent human intervention, learned intermediaries, or the “captain of the ship” should not be interpreted so broadly that responsibility is systematically shifted to professionals while providers of complementary technologies escape scrutiny (Pasquale, 2022: 8).

6.1. Fault-Based and Strict Liability Across Levels of Autonomy

The distinction between complementary and substitutive autonomy also affects the role of fault in liability analysis. Fault-based liability remains particularly relevant where the robotic system assists the clinician and the clinician retains control over patient selection, operative strategy, monitoring, and inter-

vention. Negligence may arise, for example, where the clinician selects or uses the system improperly, ignores an observable warning, inadequately supervises its operation, or fails to intervene despite having sufficient information and opportunity to do so. The use of an autonomous tool therefore does not by itself exclude professional fault.

The position becomes more difficult when the system is substitutive. If the robot independently performs a task that would otherwise have been carried out by the clinician, and the clinician cannot reasonably understand, prevent, or override the harmful action, attributing negligence to the clinician becomes less convincing. This does not eliminate professional negligence altogether: fault may still arise from inappropriate patient or system selection, misuse, failure to consider known limitations, or failure to intervene where intervention remained realistically possible. In this sense, increasing autonomy narrows, rather than automatically removes, the field of professional fault. This is consistent with the distinction between complementary and substitutive medical robotics, under which negligence remains more appropriate where technology assists the professional, while product-side liability becomes more important as technology replaces professional functions (Pasquale, 2022: 5).

Where the injury instead results from a defect or autonomous behavior attributable to the technological system, product liability may become more significant. Manufacturing defects, defective software, inadequate warnings, unsafe design, or autonomous behavior beyond reasonable clinical control may support product-side liability without requiring proof of professional negligence, subject to the applicable product-liability rules. The distinction should therefore not be treated as absolute: a complementary system may still be defective, while a highly autonomous system may still be negligently selected, monitored, or maintained. The relevant question remains which actor could reasonably foresee and control the risk that caused the harm.

The difficulty intensifies when technology becomes substitutive. If a system independently makes a decision that the surgeon neither generated nor could realistically control, attributing the resulting error to the surgeon merely because the surgeon stood ready to intervene disconnects legal responsibility from agency.

Empirical findings provide some support for this concern, although they should be interpreted cautiously. As robotic decision-making increased, respondents became progressively more divided over responsibility. Even when most or all decisions were made by the autonomous system, substantial blame continued to be placed on surgeons despite their reduced decision-making role (Jamjoom et al., 2022: 5–6). However, the study examined percepti-

ons of responsibility in hypothetical scenarios rather than actual legal liability. Its findings should therefore be interpreted with caution, particularly because the scenarios contained limited factual detail and respondents were asked to assign responsibility in broad categories. The results do not show how courts would allocate liability in real cases, but they do suggest a possible gap between actual control and perceived responsibility. This is consistent with Elish's concept of the "moral crumple zone," under which human operators may continue to bear responsibility even when substantial control has shifted to an automated system (Elish, 2019: 42; Jamjoom et al., 2022: 7).

A liability regime that defaults to malpractice whenever a surgeon is present risks institutionalizing precisely this phenomenon. The boundary should therefore follow human control over the risk that materialized, not human presence alone.

The existence of a product-side defect does not necessarily exclude professional negligence. If a software defect or malfunction is reasonably detectable during preoperative checks or intraoperative monitoring, failure by the physician to identify and respond to that defect may independently breach the professional duty of care. Similarly, maintenance or updating failures may fall within institutional responsibility where those functions are entrusted to healthcare personnel (Koyuncu Aktaş, 2021: 1713–1714). The origin of the defect and the capacity to mitigate its consequences should therefore be analyzed separately: the manufacturer may create the initial risk while a clinician or hospital later incurs responsibility for failing to respond to a risk that became reasonably detectable. A related question concerns the legal characterization of the resulting damage and the proximity of its causal connection to the underlying defect. According to the prevailing view in Turkish doctrine, the principal criterion for distinguishing direct from indirect loss is the proximity and intensity of the causal connection between the event giving rise to liability and the damage suffered. Losses that constitute the immediate consequence of the contractual breach are regarded as direct losses, whereas losses connected to the breach through a more remote causal chain are classified as indirect losses (Yıldırım & Başpınar, 2008: 1113–1114). Accordingly, where damage arises from a defect, the inquiry should distinguish losses that flow directly from the defective condition from those that materialize only through additional intervening causes (Kara, 2018: 525).

The transition to fully autonomous systems may also create a genuine compensation gap. If a self-learning robot causes harm through behavior developed after deployment, manufacturer liability may be restricted by development-risk principles while

attributing the same harm to the physician or hospital may become equally difficult. The injured patient may consequently be left without an effective route to compensation (Koyuncu Aktaş, 2021: 1729–1730).

7. The Proposed Autonomy-Control-Risk Attribution Framework

A clear conceptual framework is essential for the fair and effective allocation of responsibility among the multiple actors involved. Existing literature has already addressed several relevant considerations separately, including levels of autonomy, meaningful human control, product defects, institutional duties, and the allocation of responsibility in human-machine interaction. The contribution of this article is therefore not to present each of these elements as independently novel, but to integrate them into a structured sequence focused on identifying who exercised meaningful control over the specific risk that materialized. On this basis, the article develops the Autonomy-Control-Risk Attribution (ACRA) framework to determine whether harm involving autonomous surgical robots should be analyzed principally as medical malpractice, institutional negligence, product liability, or concurrent liability. ACRA is proposed primarily as a doctrinal and analytical tool rather than as a binding judicial test or a legislative rule. It is intended to provide courts, scholars, and policymakers with a structured method for examining responsibility in cases where control is distributed among surgeons, healthcare institutions, manufacturers, and autonomous systems. The framework proceeds through five stages.

7.1. Stage One: Identify the Autonomous Function in the Causal Chain

The analysis should begin by identifying the function performed by the robot at the time the harm occurred rather than assigning consequences from a general label of "autonomy."

Relevant functions may include patient-management decisions, surgical planning, anatomical identification, trajectory selection, tissue manipulation, suturing, energy delivery, complication detection, and emergency response. Experimental liability scenarios similarly distinguish management decisions, operative planning, and technical execution, identifying the principal decision-maker at each stage (Jamjoom et al., 2022: 3).

A single system may therefore operate at L0 for one task while functioning at L2 or L3 for another. At Stage One, liability analysis should track the autonomy of the harm-producing function rather than the overall autonomy of the machine.

7.2. Stage Two: Determine Whether Meaningful Human Control Existed

The next question is whether, while the identified function was being performed, a healthcare professional could realistically detect and prevent the harmful action. Meaningful human control requires sufficient information, comprehension, time, and authority to intervene; superficial oversight does not satisfy this standard (Amoroso & Tamburrini, 2021: 32). Upper levels of autonomy of systems should be combined with human control and responsibility (Santoni de Sio & van den Hoven, 2018: 2)

The inquiry may therefore be framed as follows: "Could a reasonably competent surgeon, on the basis of information actually and reasonably available at the relevant time, have detected the dangerous robotic behavior and prevented the harm through an authorized intervention?"

If the answer is yes, professional negligence may move to the center of the analysis. If the answer is no, responsibility should not be artificially assigned to the surgeon through presumptions based solely on human presence. Recent AI-law scholarship similarly distinguishes meaningful oversight from formal approval: effective supervision requires the human decision-maker to understand the basis of the AI

system's output, independently assess it, and retain authority to reject it where necessary (Oğurlu, 2026: 507). Applied to autonomous surgery, the system's technical provision for supervision or intervention should not by itself establish meaningful surgeon control.

ACRA thus distinguishes nominal or general supervision from effective supervision within the control chain.

7.3. Stage Three: Identify the Provenance of the Materialized Risk

After autonomy and meaningful control have been examined, the analysis should identify the sphere in which the relevant risk originated and the actor best positioned to prevent the harm.

Recent surgical-AI literature adopts a tripartite framework of risk. It distinguishes risks inherent to the AI system, risks introduced by the clinician-user, and risks arising from institutional deployment. A similar approach recognizes that, although surgeons remain the ultimate clinical decision-makers, liability may also extend to developers for defective design or failure to warn, and to institutions for negligent implementation or oversight. (Hashimoto et al., 2026: 5468).

Table 1. Constructs in the Research Model

Risk provenance	Typical example	Presumptive liability orientation
Professional clinical judgment	Unsuitable patient; unreasonable procedure selection; delayed conversion despite an observable complication	Medical malpractice
Professional use of the robot	Erroneous command, misuse, ignored warning, or inadequate supervision where intervention was realistically possible	Medical malpractice
Institutional governance	Inadequate credentialing, training, maintenance, configuration, or emergency protocols	Hospital/institutional negligence
Manufacturing failure	Defective insulation, wiring, actuator, sensor, or assembly deviating from specifications	Product liability
Design/software failure	Unsafe control architecture, algorithmic defect, inadequate fail-safe, or defective update	Product liability / manufacturer negligence
Information/warning failure	Manufacturer knows a material failure mode but inadequately warns clinicians or the hospital	Failure-to-warn product liability
Autonomous decision beyond meaningful clinical control	Robot independently selects or executes a harmful strategy that a competent surgeon could not reasonably detect or override	Presumption favoring product-side liability
Human-machine interaction failure	Inadequate supervision and a system defect both materially contribute to the injury	Concurrent liability / comparative apportionment

Source: Author's own elaboration based on the cited literature.

This risk-provenance approach is consistent with existing accounts under which a manufacturing defect points toward the manufacturer, use or medical error toward the operator, and maintenance failures toward the party responsible for maintenance. Operator responsibility may also decline as operator control decreases (O'Sullivan et al., 2019: 6).

A useful example further illustrates the distinction. If the surgeon commands the robotic instrument to move in one direction but the system misinterprets the command and moves in the opposite direction, damaging nearby tissue, the injury cannot automatically be characterized as malpractice. Where the clinician could not reasonably detect the defect and no maintenance or updating failure is attributable to the healthcare provider, the inquiry shifts toward the technological system and its producer (Koyuncu Aktaş, 2021: 1722).

7.4. Stage Four: Examine Informational Control

In autonomous systems, information itself is a form of risk control. Hospitals hold information about maintenance, credentialing, and local practices; clinicians hold patient-specific clinical information and direct intraoperative observations; manufacturers often possess superior information about training data, software architecture, validation limits, system updates, logs, and recurring adverse events. However, informational control should not be confused with substantive responsibility. The fact that an actor possesses relevant information does not, by itself, establish liability; rather, liability must still be determined by reference to the applicable duties, the source of the risk, and causation. A separate question concerns the burden of proof. As a general matter, the claimant must establish the facts supporting the relevant basis of liability, subject to any presumptions or evidentiary rules provided by the applicable legal system. Where decisive technical evidence is predominantly under the control of a manufacturer, hospital, or other defendant, however, strict application of the ordinary burden of proof may place the claimant at a substantial informational disadvantage. The liability analysis should therefore distinguish between the allocation of substantive responsibility and the evidentiary consequences of unequal access to information.

Autonomous surgical robots are safety-critical systems, making comprehensive recording of system tasks and events important for reconstructing adverse outcomes (O'Sullivan et al., 2019: 3). Accountability is also closely connected to liability, supporting "flight-recorder"-style preservation of inputs, internal system states, and outputs, although recorded data may still be insufficient to make complex machine-learning behavior genuinely interpretable (O'Sullivan et al., 2019: 5).

This asymmetry has evidentiary consequences. If the information necessary to identify the source of failure exists only in inaccessible software or logs held by the manufacturer, imposing the ordinary burden of proof on the injured patient may create a practically impossible evidentiary burden.

7.5. Stage Five: Permit Concurrent Liability for Interaction Failures

ACRA rejects a mandatory binary choice between malpractice and product liability. A robot may be defectively designed and negligently supervised; a manufacturer may inadequately warn a hospital that, in turn, inadequately trains the surgeon; or a system may generate a questionable strategy that a surgeon negligently approves despite receiving sufficient information. Where function, control, and risk are distributed, liability may likewise need to be distributed.

The hybrid nature of autonomous medical systems makes it difficult for a single cause of action to explain every injury; malpractice and product liability may therefore arise in the same case (Grossbard, 2025: 313–314). Autonomous robotic surgery similarly sits at the intersection of product design, software, hospital governance, and professional expertise, so the central inquiry should be which doctrine corresponds to the actor controlling the relevant risk (Cho, 2026: 326).

Turkish doctrine also recognizes the concurrence of liability grounds. Where a surgical robot was defective when placed on the market, the patient may pursue the manufacturer even if contractual or professional liability also arises on the healthcare side; this situation is characterized as a concurrence of grounds of liability (Koyuncu Aktaş, 2021: 1729).

8. Preventing the Surgeon from Becoming the "Moral Crumple Zone"

Human oversight is widely recognized as playing a critical role in ensuring accountability and responsibility in the use of autonomous systems. Human-in-the-loop architectures can serve two very different functions. Genuine human control can protect patients by enabling a qualified professional to detect and prevent machine error. Conversely, human presence may function mainly as a focal point for responsibility, leaving a nominal supervisor available to absorb legal blame while the system autonomously determines what occurs.

If a human is retained in a system not because surgery requires human performance but principally to provide a person to whom liability can be assigned, it becomes difficult to justify simultaneously depriving that person of meaningful control (O'Sullivan et al., 2019: 7).

Empirical findings reinforce the importance of this concern. Significantly, it showed that the surgeons are still ascribed blame even when they have had no role in decision-making, which adds weight to concerns that human operators may become “moral crumple zones” absorbing blame and potentially legal responsibility when complex autonomous systems fail (Elish, 2019: 50). Even as autonomous systems assumed more decision-making, the surgeon was still identified as a responsible actor despite having a limited role in the decision-making process, adding weight to the concept of the “moral crumple zone” (Elish, 2019: 55; Jamjoom et al., 2022: 7).

ACRA responds by refusing to infer responsibility from supervision alone. If a manufacturer relies on human intervention as a defense, it should be able to demonstrate that intervention was technically and clinically meaningful. Ordinarily, this requires evidence that the failure was detectable and understandable, that sufficient time existed for intervention, and that the consequences were practically reversible.

The reciprocal conclusion is equally important. Where sufficient time and a realistic opportunity to reverse the consequences existed, surgeons should remain responsible for ignoring or misusing control genuinely left to them. The objective is not to immunize professionals from liability, but to align legal responsibility with genuine agency and meaningful control.

9. Informed Consent as a Distinct Layer of Responsibility

The transition toward autonomous surgery changes the matters on which patients may need to be informed before treatment. Informed consent, however, should not be confused with the allocation of liability for harm caused by negligent conduct.

The use of robotic surgery may expand the physician's disclosure obligations. Where robotic surgery is selected instead of a conventional technique, the patient should be informed of that choice and available alternative methods because meaningful and autonomous decision-making regarding bodily integrity depends on adequate information about the available interventions (Koyuncu Aktaş, 2021: 1714).

Generic consent may be insufficient for automated robotic procedures. Patients should receive case-specific information about the nature of the procedure, the role of robotics and automation, expected benefits, possible outcomes, risks, and the possibility of conversion to another technique (Contaldo et al., 2025: 5). Professional judgment nevertheless remains important when changing clinical conditions require conversion to a conventional approach.

Robot-specific information should therefore form part of informed-consent analysis. The clinician's duty to disclose treatment risks to the patient or relatives is distinct from the manufacturer's duty to provide adequate warnings to the medical intermediary (Cho, 2026: 324–325). As autonomy increases, the patient should be told not merely that “a robot will be used,” but the extent of the robot's decisional role. Telling a patient that the surgeon will direct robotic instruments is materially different from disclosing that the system will independently plan or execute an operative task. Consent to autonomous robotic surgery, however, is not a liability waiver: consent to disclosed risks does not amount to consent to negligent use, inadequate institutional governance, or defective technology.

The scope of disclosure should correspond to the actual role performed by the robotic system. The patient does not need to be given a technical description of the system or its formal autonomy level, but should be informed, in understandable terms, about which clinically relevant functions remain under the surgeon's control, and which are delegated to the robotic system. In particular, the patient should be told whether the robot merely reproduces the surgeon's commands, generates or recommends operative options subject to human approval, or independently selects or executes part of the procedure. Where human intervention remains possible, the patient should also understand the surgeon's supervisory role and the practical limits of that supervision (Aydın, 2024: 21–22).

Informed consent should further distinguish between risks inherent in the medical procedure and risks introduced by the robotic technology. General informed-consent principles require disclosure of the nature and purpose of the intervention, its material risks and expected benefits, reasonable alternatives, and information sufficient for the patient to understand and voluntarily authorize the proposed treatment.¹ In robotic surgery, this general disclosure should be supplemented by technology-specific information, including material risks of hardware malfunction, software error, limitations of the robotic system, the extent of robotic involvement, and the experience or training of the surgeon where this may materially affect the patient's decision (Bagwan et al., 2025: 74).

For example, in an angiographic intervention performed under local anesthesia with the assistance of a robotic or computerized system, risks inherent in the angiographic procedure may arise even when the equipment functions exactly as intended. Those risks should be distinguished from harm that may arise because the robotic device malfunctions, its software operates incorrectly, or an autonomous function produces an erroneous action. If the system merely executes the physician's commands, its role may ap-

appropriately be described as assistive. If, however, the system independently selects or executes a clinically relevant step, that autonomous decision-making role becomes a separate matter that should be disclosed to the patient.

The relevant inquiry is therefore not simply whether a robot is used, but whether its participation materially changes the nature of the intervention, the source of risk, or the distribution of decision-making authority. Robot-specific informed consent should provide enough information for the patient to decide whether to accept that technological and decisional arrangement without turning the consent process into a technical manual. Consent to the disclosed risks of robotic or autonomous treatment, moreover, does not constitute consent to negligent use, inadequate supervision, or a defective technological system.

10. Comparative Implications: Limits of the “Robot-as-Tool” Model

Current Chinese practice provides a useful example of the consequences of legal classification. Robot-assisted surgical interventions are generally characterized as complementary or assistive tools rather than as procedures performed by a robotic legal actor. Claims consequently tend to target hospitals, while medical appraisal bodies evaluate the medical intervention as a whole instead of separately assessing the robot's algorithmic operation or autonomous decision-making capacity (Shentu, 2024: 5).

Given the present autonomy of most medical robots, this approach is internally coherent. Predominantly assistive technology can reasonably be treated as a tool while essential medical decisions and their execution remain under the surgeon's control (Shentu, 2024: 5). The tool model, however, has an inherent endpoint.

Once the robot autonomously performs clinically critical perception, planning, or execution functions, treating every outcome as part of the hospital's medical act may obscure the actual source of risk. Some appraisal institutions have reported inadequate capacity to assess the robotic dimensions of the technologies before them, with the result that algorithmic structure or artificial intelligence may fall outside the causal analysis (Shentu, 2024: 5).

Courts may therefore possess mature structures for evaluating hospitals while lacking comparable mechanisms for autonomous systems, creating a systemic tendency toward malpractice-based characterization. Clearer allocation of responsibility among stakeholders and stronger pre-market oversight of high-risk medical robots have accordingly been recommended.

10.1. Would Legal Personality Change Liability Attribution?

The increasing decisional autonomy of surgical robots raises a further question: would the allocation of liability change if an autonomous surgical robot were recognized as a legal person? Current approaches remain cautious about such a step. It is argued that the present denial of legal personality to AI is appropriate at the current stage of technological development (Shentu, 2024: 6). Similarly it is recognized that existing classifications of responsibility may require reconsideration if robots were eventually granted a legal status comparable to that of persons (O'Sullivan et al., 2019: 7).

Granting legal personality would alter the formal structure of liability by making the autonomous robot itself a potential bearer of legal duties and, at least in principle, a possible subject of a direct liability claim. It would not, however, by itself resolve the underlying problem of compensation or risk attribution. Direct liability of the robot would have practical value only if accompanied by a mechanism from which compensation could be obtained, such as dedicated assets, compulsory insurance, or a compensation fund. More importantly, recognition of robotic legal personality should not automatically release manufacturers, hospitals, or clinicians from responsibility for risks attributable to their own conduct. A defective algorithm may still originate in product design, unsafe deployment may remain attributable to the healthcare institution, and negligent supervision may remain attributable to the clinician.

Legal personality could therefore add another potential subject to the liability structure without eliminating existing grounds of malpractice, institutional negligence, or product liability. Indeed, treating the robot as a separate legal person while allowing manufacturers or healthcare institutions to avoid responsibility for risks that they designed, introduced, or controlled could itself create a new accountability gap. The relevant question would remain where the materialized risk originated, and which actor had meaningful control over it.

For this reason, legal personality is neither necessary nor sufficient for the ACRA framework. Under present law, ACRA allocates responsibility among existing legal actors by examining autonomous function, meaningful human control, risk provenance, informational control, and interaction failures. If future legislation were to recognize autonomous robots as legal persons, the robot could be added as another potential node within that analysis, but the underlying inquiry would remain unchanged. Legal personality may therefore change the range of possible liable subjects, but it does not eliminate the need to examine manufacturer, institutional, and professional responsibility.

11. Legal Responsibility, Algorithmic Explainability, and the Burden of Proof

An effective liability regime for autonomous surgery depends on evidence capable of reconstructing the event after harm occurs.

Autonomous robotic surgery introduces a new evidentiary category alongside conventional medical record. Relevant evidence may include continuous sensor streams, commands sent to the system, AI-generated outputs, transitions between system states, warnings, software-version information, calibration records, updates, and timestamps of surgeon intervention.

The preservation of such records is closely connected to the attribution of responsibility. A “black-box” or flight-recorder approach can preserve system inputs, internal states, and outputs so that failures can be analyzed and evidence retained for later disputes (O’Sullivan et al., 2019: 5–7). Recording alone, however, is not sufficient. A distinction remains between a physical black box that records events and an algorithmic black box whose internal operation remains difficult to interpret despite the availability of records (O’Sullivan et al., 2019: 5).

A formal right to sue a manufacturer may remain largely theoretical if the patient cannot access or interpret the data necessary to show how technology contributed to the injury. Increasing autonomy should therefore be accompanied by a proportionate duty of traceability. As machine independence in decision-making increases, records should permit later determination of:

- who initiated the relevant action or procedure;
- which software and model version was in operation;
- which sensor data the system relied upon;
- which action or strategy the system selected;
- whether a warning was generated;
- what the surgeon could observe at the relevant time;
- whether intervention was technically possible; and
- whether a post-market update or maintenance event affected the behavior.

Traceability is not merely an evidentiary convenience; it is a precondition for meaningful attribution of responsibility.

12. Discussion

The analysis produces three principal conclusions.

First, liability should be redistributed only insofar as control is actually redistributed. The presence of an automated component does not automatically release the surgeon from professional obligations.

Where the surgeon determines clinical strategy, understands system behavior, and can realistically intervene, malpractice remains an appropriate framework. By contrast, once operative actions and decisions are genuinely delegated to autonomous software, continuing to place the principal burden on the surgeon increasingly conflicts with the logic of negligence-based responsibility.

Second, product liability and medical malpractice increasingly overlap rather than function as alternatives. Contemporary design-defect and warning analyses already employ concepts of reasonableness, foreseeability, and preventability that resemble negligence analysis. This convergence is particularly significant in autonomous surgery because courts may need to distinguish whether a failure originated in product design, professional misuse, inadequate training, or institutional malfunction (Cho, 2026: 310–311). The central question is therefore less a categorical choice between fault and strict liability than which actor was best positioned to foresee and manage the relevant risk.

Third, human supervision cannot simultaneously be used to justify the safety of an autonomous system and as an automatic mechanism for transferring responsibility away from the technology provider. If manufacturers and hospitals rely on surgeon presence as a safety layer, the surgeon must receive enough information and practical intervention authority to perform that function. Where technical design prevents real-time understanding or effective interruption, the nominal title of “supervisor” should not, by itself, divert liability from the technology provider. Contractual arrangements between hospitals and manufacturers may allocate certain risks through warranties, indemnity clauses, maintenance obligations, or insurance requirements. Such arrangements may be relevant to the internal distribution of financial loss between the parties, but they should not by themselves determine liability toward an injured patient. Similarly, the existence of professional, product, or institutional liability insurance affects how losses may ultimately be absorbed, rather than the prior question of which actor legally controlled or created the risk that materialized.

This perspective also refines Pasquale’s complementary/substitutive distinction. Complementarity and substitution are not immutable properties of a system as a whole; they are situational characteristics that should be assessed function by function. A future surgical robot may complement the clinician during diagnosis, substitute for the clinician during execution of a trajectory, and return to a complementary role during postoperative assessment. Liability should remain flexible enough to track these changes in the control architecture.

The resulting model is better understood as a continuum than as a fixed category: meaningful human

control and malpractice predominance at one end; mixed control and concurrent responsibility in the middle; and autonomous risk management with increasing product-side responsibility toward the other end.

This is a rebuttable orientation, not an absolute rule. A highly autonomous procedure may still involve surgeon negligence, just as a minimally autonomous device may contain a manufacturer defect. Autonomy is therefore not dispositive; it guides the initial inquiry and helps determine the extent to which law should expect one actor to foresee and prevent the failures of another.

Liability rules must also remain sensitive to the distinctive lifecycle of software. Because software is comparatively error-prone and continuously developed and updated, undifferentiated application of excessively burdensome liability standards may inhibit technological development (Yünlü, 2024: 538). The objective should not be maximum manufacturer liability, but liability calibrated to actual and legally meaningful control over preventable risk.

13. Conclusion

The increasing use of autonomous systems raises important questions concerning the allocation of responsibility for their actions and outcomes. Autonomous surgical robotics exposes the limits of doctrines built around a relatively clear distinction between professional services and manufactured products. Traditionally, the surgeon supplied medical judgment and technical skill while the device supplied a tool. Increasing autonomy erodes that distinction as surgical robots acquire the capacity to observe anatomy, generate strategies, execute tasks, respond to changing conditions, and perform functions once constitutive of the surgeon's professional activity.

The legal response should be neither to grant robots legal personhood nor to force every injury into a single category of malpractice or product liability. A more defensible solution is to align responsibility with the actual architecture of control.

The ACRA framework therefore asks which function was autonomous, whether meaningful human intervention was realistically possible, where the materialized risk originated, and which actor possessed the information and practical capacity necessary to prevent the harm. Where surgeons or hospitals retain genuine control over the relevant clinical risk, professional negligence should remain central. Where harm results from manufacturing defects, unreasonable design, defective software, inadequate warnings, or autonomous behavior that healthcare professionals could not reasonably detect or prevent, product-side liability should gain greater weight. Where these causes interact, concurrent liability is

more appropriate than artificially selecting a single doctrinal category.

This approach protects surgeons from becoming a moral crumple zone that absorbs responsibility for increasingly machine-generated decisions while also preventing advanced technology from becoming a shield against genuine professional responsibility. The decisive question is not whether the human or the robot "performed the surgery." It is who exercised meaningful control over the risk that caused the harm. As surgical systems become more autonomous, that question should become a central principle for drawing the boundary between medical malpractice and product liability.

References

- Amoroso, D. & Tamburrini, G. (2021). The Human Control Over Autonomous Robotic Systems: What Ethical and Legal Lessons for Judicial Uses of AI? In X. Kramer, A. Biard, J. Hoevenaars & E. Themeli (Eds.), *New Pathways to Civil Justice in Europe: Challenges of Access to Justice*, 23–42. Springer. https://doi.org/10.1007/978-3-030-66637-8_2
- Attanasio, A., Scaglioni, B., De Momi, E., Fiorini, P. & Valdastris, P. (2021). Autonomy in Surgical Robotics. *Annual Review of Control, Robotics, and Autonomous Systems*, 4: 651–679. <https://doi.org/10.1146/annurev-control-062420-090543>
- Aydın, F. (2024). Yapay Zekâ ile Desteklenmiş Bilgilendirilmiş Onamda Etik Kodlarının Belirlenmesi. Doktora Tezi. Eskişehir Osmangazi Üniversitesi, Sağlık Bilimleri Enstitüsü, Tıp Tarihi ve Etik Anabilim Dalı.
- Bagwan, M. B., Joshi, T., Goswami, R. A., Patil, R. S., Sharma, V. & Wani, L. K. (2025). Liability in Robotic Surgery: Legal Frameworks and Case Studies. *Journal of Neonatal Surgery*, 14(2s): 70–77. <https://doi.org/10.52783/jns.v14.1659>
- Cho, H. J. (2026). Liability Standards in Autonomous Robotic Surgery. *UC Law Science and Technology Journal*, 17(2): 301–327. https://repository.uclawsf.edu/hastings_science_technology_law_journal/vol17/iss2/
- Contaldo, M. T., Triggiani, S., Vignati, G., Bracchi, D. & Carrafiello, G. (2025). Legal Implication in Utilizing Automated Robots: A Written Informed Consent Form Proposal. *International Journal of Medical Robotics and Computer Assisted Surgery*, 21(2): e70064. <https://doi.org/10.1002/rcs.70064>
- Elish M. C. (2019) Moral Crumple Zones Cautionary Tales in Human-Robot Interaction. *Engaging Science, Technology, And Society*. 5:40–60. doi: 10.17351/ests2019.260.
- Eren, F. & Yücer Aktürk, İ. (2024). Borçlar Hukuku Özel Hükümler. Gözden Geçirilmiş 12. Bası. Ankara: Legem Yayınevi.
- European Union. (2024). Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products and repealing Council Directive 85/374/EEC. *Official Journal of the European Union*, L, 2024/2853.
- Grossbard, E. S. (2025). The AI-Robotic Prescription: Legal Liability When an Autonomous AI Robot Is Your Medical Provider. *University of Miami Business Law Review*, 33(2): 273–314. <https://repository.law.miami.edu/umblr/vol33/iss2/6/>
- Gül, İ. (2026). Ayıbın İkrarı Zamanaşımını Keser mi? Bir Yargıtay Kararı İncelemesi. *Ankara Hacı Bayram Veli Üniversitesi Hukuk Fakültesi Dergisi*, 30(2): 759–808. <https://doi.org/10.34246/ahb-vuhfd.1760077>
- Gümüş, M. A. (2024). Borçlar Hukuku Özel Hükümler: Kısa Ders Kitabı. Gözden Geçirilmiş ve Yenilenmiş 8. Bası. Ankara: Yetkin Yayınevi.
- Hashimoto, D. A., Marwaha, J. S., Lee, S. A., Schwaartzberg, S. & Duffour, M. N. (2026). Risk and liability in the deployment of AI

- systems for surgery: A SAGES white paper. *Surgical Endoscopy*, 40, 5468–5479. <https://doi.org/10.1007/s00464-026-12881-8>.
- Jamjoom, A. A. B., Jamjoom, A. M. A., Thomas, J. P., Palmisciano, P., Kerr, K., Collins, J. W., Vayena, E., Stoyanov, D. & Marcus, H. J. (2022). Autonomous Surgical Robotic Systems and the Liability Dilemma. *Frontiers in Surgery*, 9: 1015367. <https://doi.org/10.3389/fsurg.2022.1015367>
- Kara, D. (2018). İsviçre-Türk Hukuk Sisteminde Ayıptan Sorumluluk Çerçevesinde Doğrudan Dolaylı Zarar Ayrımı. *Galatasaray Üniversitesi Hukuk Fakültesi Dergisi*, 17 (1): 483–530.
- Kaşak, F. E. (2023). Taşınmaz Satışında Ayıp ve Eksik İfa. İstanbul: On İki Levha Yayınevi.
- Kılıçoğlu, A. M. (2023). Borçlar Hukuku Özel Hükümler. 5. Bası. Ankara: Turhan Kitabevi.
- Koyuncu Aktaş, N. (2021). Hekim ve Hastanenin Cerrahi Robot Kullanımının Hastaya Verdiği Zarardan Sorumluluğu. *Selçuk Üniversitesi Hukuk Fakültesi Dergisi*, 29(3): 1693–1737. <https://doi.org/10.15337/suhfd.885733>
- Küçükali, C. (2025). Ethics and Privacy in Digital Health in the Context of Intellectual Property Rights, Personal Rights and Personal Data. In I. Uluslararası II. Ulusal Sağlık Bilimleri Kongresi Bildiri Kitabı, 174–175. <https://openaccess.izu.edu.tr/>
- Lee, A., Baker, T. S., Bederson, J. B. & Rapoport, B. I. (2024). Levels of Autonomy in FDA-Cleared Surgical Robots: A Systematic Review. *npj Digital Medicine*, 7: 103. <https://doi.org/10.1038/s41746-024-01102-y>
- Liang, Z., Thareja, N., Zhang, P., Joyce, C., Atar, S., Richter, F., Jacobsen, G., Liu, S., Broderick, R. & Yip, M. (2026). In vivo feasibility study of humanoid robots in surgery, *Nature* 657, 236–244. <https://doi.org/10.1038/s41586-026-10796-x>
- Oğurlu, C. S. (2026). About Artificial Intelligence Tools in Criminal Procedure, Judicial Discretion and Criminal Procedural Law Safeguards. *Society and Innovations*, 7(3): 491–514. <https://doi.org/10.47689/2181-1415-vol7-iss3-pp491-514>
- 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 (AI) and Autonomous Robotic Surgery. *International Journal of Medical Robotics and Computer Assisted Surgery*, 15(1): e1968. <https://doi.org/10.1002/rcs.1968>
- Parlak Börü, Ş. (2019). Robotik Cerrahi Müdahalelerden Doğan Hukuki Sorumluluk. *İnönü Üniversitesi Hukuk Fakültesi Dergisi*, 10(2): 758–775. <https://doi.org/10.21492/inuhfd.636507>
- Pasquale, F. A. III. (2022). The Price of Autonomy: Liability Standards for Complementary and Substitutive Medical Robotics and Artificial Intelligence. *Ius et Praxis*, 28(1): 3–19. <https://doi.org/10.4067/S0718-00122022000100003>
- 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>
- Santoni de Sio, F., & van den Hoven, J. (2018). Meaningful human control over autonomous systems: A philosophical account. *Frontiers in Robotics and AI*, 5, Article 15. <https://doi.org/10.3389/frobt.2018.00015>
- Shentu, X. (2024). A Review on Legal Issues of Medical Robots. *Medicine*, 103(21): e38330. <https://doi.org/10.1097/MD.00000000000038330>
- 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>
- Yavuz, C., Acar, F. & Özen, B. (2019). Türk Borçlar Hukuku Dersleri (Özel Hükümler). 16. Bası. İstanbul: Beta Basım.
- Yıldırım, M. F. (2006). Satım Sözleşmesindeki Ayıba Karşı Tekeffül Hükümleri ile İrade Sakatlığı Hali Olarak Hile (BK. m. 28) Hükümü Arasındaki İlişki. Prof. Dr. Fikret Eren'e Armağan: 965–985. Ankara.
- Yıldırım, F. & Başpınar, V. (2008). Doktrin ve Uygulama Açısından Türk-İsviçre Hukukunda Doğrudan Doğruya ve Dolaylı Zarar Ayrımı. Prof. Dr. Bilge Öztan'a Armağan: 1093–1125. Ankara: Turhan Yayınları.
- Yünlü, S. (2024). Üretici Yapay Zekâ Kaynaklı Norm ve Kişi Bazlı Hukuki Sorumluluk. *Adalet Dergisi*, (72): 501–542. <https://adalet-dergisi.adalet.gov.tr/arsiv/adaletdergisi/2024/72.sayi/72-icindekiler.html>