MEDICAL MURDER WITHOUT INTENT: RETHINKING CRIMINAL LIABILITY IN THE AGE OF ARTIFICIAL INTELLIGENCE

Amrisha Kumar, BBA-LLB (Hons.), University of Mumbai Thane Sub-Campus

ABSTRACT

The incorporation of artificial intelligence into healthcare has altered the traditional framework of medical decision-making. Technologies that once functioned merely as assistive tools are now capable of influencing diagnoses, treatment pathways, and surgical interventions. While such advancements enhance efficiency and predictive accuracy, they simultaneously introduce complex challenges for criminal law—particularly when patient death occurs without identifiable human intent.

Indian criminal jurisprudence is fundamentally anchored in the doctrine of mens rea, which presumes that criminal responsibility arises only from conscious human conduct. Artificial intelligence, however, operates independently of intention, awareness, or moral judgment. This disconnect creates a legal vacuum where fatal outcomes may occur in the absence of a legally blameworthy mind.

This paper examines whether the Bharatiya Nyaya Sanhita, 2023 adequately addresses deaths resulting from AI-assisted medical treatment. By analysing Section 106, judicial standards governing medical negligence, and the limitations of existing doctrines such as proximity and foreseeability, the study highlights the structural inadequacy of intent-based liability in algorithm-driven healthcare. It argues that AI-related medical deaths expose a fundamental crisis in criminal attribution and necessitate a reorientation of responsibility beyond individual fault. The paper concludes by emphasizing the need for legislative clarity and a shift toward institutional and shared accountability in order to preserve fairness within the criminal justice system.

1. Introduction

1.1 The Shift from Human Judgment to Algorithmic Decisions

Historically, medicine has always been a deeply human profession. A doctor's diagnosis was a mix of their education, experience, and gut instinct. But the landscape is changing fast. Today, hospitals are deploying AI systems that can detect diseases or plan surgeries faster than any human could. These systems are marketed as the ultimate solution to human error.

However, as we lean more on technology, the nature of medical decision-making is shifting. Doctors are often pushed—sometimes by hospital protocol—to follow what the AI suggests. While this might make things efficient, it takes away a layer of human discretion. And this is where the legal nightmare begins: if an AI makes a recommendation that leads to a patient's death, **who do we hold accountable?**

1.2 Clash with the Bharatiya Nyaya Sanhita (BNS), 2023

Indian criminal law has recently undergone a massive overhaul with the replacement of the colonial-era IPC by the **Bharatiya Nyaya Sanhita (BNS), 2023**. However, despite this modernization, the law still operates on the presumption that a crime requires a human perpetrator with a specific mental state (*mens rea*). Whether we are looking at culpable homicide or the newly structured provisions for negligence under **Section 106 of the BNS**, the law fundamentally looks for intent, knowledge, or a lack of reasonable care by a *person*.

AI, however, remains a black box. It doesn't have a conscience, it doesn't have morals, and it certainly doesn't possess the *mens rea* required by the Sanhita. It operates purely on probabilities and data patterns. This leads to a bizarre legal scenario that could be called **"medical murder without intent."** We have a victim and a wrongful death, but the "killer" is a piece of software that cannot be imprisoned or penalized under the new provisions.

Currently, even the BNS is hazy on this front. If a doctor blindly follows an AI and a patient dies, does this fall under the medical negligence exception in **Section 106(1)?** If a hospital deploys a "black box" algorithm, can they be held vicariously liable? The transition from IPC to BNS updated the text, but it didn't solve the core issue: our statutes simply weren't written for autonomous, non-human decision-makers. This paper aims to dissect this conflict, analyzing whether the BNS is truly "modern" enough to survive the AI era or if we are still

using 19th-century logic for 21st-century technology.

Consequently, this paper seeks to examine whether the existing principles of Indian criminal law, particularly the requirement of *mens rea*, are capable of addressing deaths caused by artificial intelligence–assisted medical treatment. It questions whether the BNS is truly "modern" enough to survive the AI era or if we are still attempting to fit 21st-century technology into traditional legal frameworks. To answer this, the study adopts a **doctrinal and analytical approach**, critically examining statutory provisions, relevant judicial decisions, and academic literature to identify the structural gaps in our current jurisprudence.

2. Artificial Intelligence in Healthcare

2.1 Beyond Mere Tools: Artificial Intelligence has moved far beyond being just a high-tech assistant; it has become a central pillar of modern healthcare. Unlike traditional medical devices—like an X-ray machine or a ventilator, which wait for a human to operate them—AI systems are active. They analyse massive datasets to generate independent recommendations in fields like radiology, oncology, and even robotic surgery. The key shift here is that these systems aren't just *supporting* doctors; in many cases, they are actively *guiding* the clinical decision-making process.

2.2 The Learning Problem: To understand the legal risk, we have to understand the tech. Most medical AI runs on Machine Learning (ML). These models are "trained" on thousands of patient histories and outcomes to spot patterns that humans might miss. But here’s the catch: unlike a standard medical instrument that always works the same way, AI evolves. As it consumes more data, it updates its own parameters. While this improves accuracy, it reduces predictability. For a lawyer, this is a nightmare because the system is constantly changing, distinguishing it from any conventional medical equipment we've regulated before.

2.3 The Spectrum of Autonomy: We can broadly categorize these systems into two buckets:

1. **Decision-Support Systems:** These act like a "second opinion," offering a diagnosis but leaving the final call to the doctor.

2. **Autonomous/Semi-Autonomous Systems:** These are trickier. In high-stakes environments like the ICU or robotic surgery, systems might determine drug dosages or surgical cuts on their own. The reality of hospital life—time pressure and burnout—

means that even with "support" systems, doctors often default to the algorithm's choice. They don't have the time to double-check the math, meaning the "human in the loop" is often just a rubber stamp.

2.4 The "Black Box" Dilemma: This is arguably the biggest legal hurdle. In deep learning models, there is a "Black Box" problem: the input goes in, the answer comes out, but the messy middle—*how* the AI reached that conclusion—is often a mystery, even to the developers. This lack of explainability is dangerous. If a doctor follows an AI recommendation that turns out to be fatal, they might not be able to explain *why* the decision was made. You can't cross-examine an algorithm, and a doctor can't justify a rationale they never understood in the first place.

2.5 The Absence of Moral Agency: Finally, we have to remember that AI is purely mathematical, not moral. A human doctor balances medical facts with ethics, empathy, and patient history. AI operates on cold probability. It doesn't value human life; it calculates risk percentages. It makes decisions that are technically rational but might be ethically blind.

2.6 Why This Matters for Criminal Liability: This structural shift changes the game for negligence. In a standard medical negligence case, we ask: *"Did the doctor act reasonably?"* But if the treatment path was chosen by an opaque algorithm that the doctor was required to trust, defining "reasonable conduct" becomes nearly impossible. The human element shrinks, and the technological influence grows. Therefore, before we can discuss who to blame under the BNS, we have to accept that AI hasn't just improved medicine; it has fundamentally altered the chain of command in healthcare.

3. Existing Legal Framework Governing Medical Liability under the Bharatiya Nyaya Sanhita (BNS), 2023

3.1 Old Wine in a New Bottle: The Persistence of *Mens Rea* the enactment of the **Bharatiya Nyaya Sanhita (BNS), 2023** represents a major shift in India's criminal justice system, replacing the colonial-era penal code. However, while the names and sections have changed, the heart of criminal liability remains the same: **no crime without a guilty mind.** The BNS is still deeply rooted in the traditional requirement of *mens rea*. For any offense resulting in death, the prosecution must prove a specific mental state—whether it's intention, knowledge, or negligence. The law fundamentally assumes that the accused is a human being capable of moral reasoning and understanding the consequences of their actions.

When it comes to doctors, the law has always treaded carefully, and the BNS continues this trend. Recognizing that medicine is an uncertain science, the law distinguishes between a simple "error of judgment" and "criminal negligence." Under the BNS, particularly within the ambit of **Section 106** (which replaces the old Section 304A IPC), a doctor isn't a criminal just because a treatment failed.

3.2 Who is the "Person"?

Here is where the framework starts to crack. The BNS is built on the presumption that negligence comes from *human* conduct—acts or omissions by a person. But AI systems operate on algorithms and predictive modelling, not conscious thought. When a doctor's decision is heavily influenced—or even dictated—by an AI, identifying the "guilty mind" becomes messy. If a doctor follows a hospital-mandated AI protocol and the patient dies, have they committed "gross negligence"?

- The doctor didn't write the code.
- The doctor might not even understand *how* the AI reached its conclusion (the "black box" problem).
- Yet, strictly speaking, the BNS views the doctor as the only "person" in the room responsible for the decision.

4. Judicial Precedents and the *Jacob Mathew* Standard

The Benchmark: *Jacob Mathew v. State of Punjab* (2005) To understand why prosecuting AI-related deaths is so difficult, we have to look at the ***"holy grail"*** of medical negligence law in India: The Supreme Court's judgment in ***Jacob Mathew v. State of Punjab.***[1] This case set the gold standard for criminal liability. The Court ruled that for a doctor to be held criminally liable, the negligence must be **"gross"** or of a very high degree. A simple lack of care or an error of judgment—which might get you fined in a consumer court—is not enough to put you in prison.

The Court's logic was simple: doctors need to take risks to save lives. If they are constantly terrified of arrest, they will practice "defensive medicine" (e.g., ordering unnecessary tests just

[1] *Jacob Mathew v. State of Punjab, (2005) 6 S.C.C. 1 (India).*

to cover their backs). Therefore, to prove criminal negligence, you must show that the doctor did something that no reasonable medical professional in their right mind would have done.

The Bolam–Bolitho Framework and Its Limits in Algorithmic Medicine

The standard for determining medical negligence in India has largely been shaped by the Supreme Court's decision in *Jacob Mathew v. State of Punjab*, which adopted the English **Bolam test** laid down in ***Bolam v. Friern Hospital Management Committee***[2]. According to this principle, a medical professional cannot be held negligent if their conduct conforms to a practice accepted as proper by a responsible body of medical professionals skilled in that field.

The underlying logic of the Bolam test is **professional consensus**. It assumes that medical judgment is formed through human reasoning, experience, and peer evaluation. If similarly placed doctors would have acted in the same manner, criminal liability is ordinarily excluded, even if alternative or more advanced methods were available at the time.

However, the increasing integration of artificial intelligence into clinical decision-making exposes a serious conflict within this framework. Unlike human practitioners, AI systems do not participate in professional deliberation, nor do they justify their conclusions through explainable reasoning. Many such systems operate as ***"black boxes,"*** generating recommendations without revealing the basis of their analysis.

This creates a legal paradox. When a doctor relies on an AI-assisted diagnostic or treatment tool that is widely used within hospitals, such reliance may itself be argued to constitute an *accepted medical practice*. Under the protection offered by *Jacob Mathew*, the doctor may claim absence of negligence. Yet, where the AI's recommendation is erroneous and its reasoning cannot be understood or questioned, the doctor's reliance begins to resemble a delegation of medical judgment to a system beyond human comprehension.

The difficulty lies in the fact that the Bolam framework presumes that "standard practice" is something that can be debated, defended, and logically explained by human experts. It never anticipated a situation where medical decisions are influenced by autonomous systems whose internal reasoning remains opaque. Applying Bolam mechanically in such circumstances risks

[2] *Bolam v. Friern Hospital Management Committee*

converting professional judgment into **uncritical technological obedience**.

A partial corrective to this problem emerges through the **Bolitho test**, formulated in ***Bolitho v. City and Hackney Health Authority***[3]. Bolitho qualifies the Bolam rule by permitting courts to reject medical opinions that fail to withstand *logical analysis*. Under this approach, professional endorsement alone is insufficient unless the reasoning supporting the decision is rational and defensible.

In the context of artificial intelligence, Bolitho offers a potential analytical tool. If an AI system cannot explain the basis of its recommendation, its output may arguably fail the requirement of logical scrutiny. This could, in theory, allow courts to question blind reliance on algorithmic advice rather than accepting it solely because it reflects prevailing practice.

Nevertheless, the application of Bolitho to AI-driven medical decisions remains uncertain. Where algorithms are incapable of providing transparent reasoning, meaningful judicial evaluation becomes difficult. Consequently, existing negligence doctrines—designed for human actors—struggle to address harms arising from non-human decision-makers.

This exposes a fundamental limitation of the current legal framework: **tests developed for human judgment are ill-equipped to govern algorithmic medicine**. Until Indian courts or the legislature provides explicit guidance on the legal treatment of AI-assisted decision-making, the Bolam–Bolitho framework will continue to operate within a doctrinal vacuum.

4.1 The "Rational Doctor" vs. The " Automated Standard"

This creates a dangerous ambiguity. If the medical standard of care evolves to include using AI tools, then *refusing* to use AI might eventually be seen as negligent. Conversely, *relying* on a flawed AI could also be negligent. The BNS offers no guidance here. It doesn't tell us if "trusting the machine" is a valid defence or a reckless act.

Crucially, the BNS is silent on artificial intelligence. There are no sections dealing with automated decision-making, algorithmic accountability, or autonomous systems. It remains a strictly human-centric statute. While civil law (consumer protection, torts) has mechanisms for compensation, criminal law is about **punishment and moral blame**. You cannot punish an

[3] *Bolitho v. City and Hackney Health Authority*

algorithm, and punishing a doctor for the errors of a machine they didn't create feels legally unsound.

We are left with a significant disconnect. The BNS upholds the traditional principles of intention and negligence, but AI challenges the very assumptions those principles are based on. We are trying to apply a law designed for human behaviour to a reality dominated by data processing. This absence of legislative clarity leaves courts with limited tools to interpret these cases, creating a risky "grey area" for both the medical professionals deploying this tech and the patients subjected to it.

5. Medical Death without *Mens Rea*: The Core Criminal Law Crisis

The Problem of Criminal Attribution

Criminal liability in Indian law is not outcome-based but fault-based. Courts have consistently maintained that punishment must correspond to moral blame rather than mere consequence. When artificial intelligence becomes central to medical decision-making, this moral foundation begins to erode. In AI-assisted treatment, the doctor may act in good faith, the hospital may follow approved protocols, and the developer may remain geographically and causally distant from the patient. Yet a fatal outcome may still occur. In such circumstances, the criminal law encounters an unprecedented dilemma: the existence of harm without an identifiable guilty mind. This results in what may be described as a situation of "medical death without mens rea," where the traditional architecture of criminal responsibility fails to accommodate technologically mediated harm.

5.1 The Judicial Shield: (*Jacob Mathew* and *Kusum Sharma)* Historically, Indian courts have treated doctors with kid gloves. The judiciary understands that medicine is an imperfect science—patients die even when doctors do everything right.

- **The *Jacob Mathew* Standard (2005):** This is the most critical precedent. The Supreme Court made it clear: for a doctor to be a criminal, simple negligence isn't enough. It has to be **"gross negligence"** or recklessness. A mere error of judgment or an accident doesn't land a doctor in jail.

- **The *Kusum Sharma* [4] Warning (2010):** The Court doubled down here, warning against

[4] *Kusum Sharma v. Batra Hosp. & Med. Rsch. Ctr., (2010) 3 S.C.C. 480 (India).*

the tendency to criminalize every medical failure. The message was clear: unless the doctor's conduct was a total departure from standard practice, the criminal law should stay out of it.

5.2 The AI Paradox: When "Standard Practice" Kills Here is where the logic breaks down. ***Kusum Sharma*** tells us to judge negligence based on "prevailing medical standards." But what happens when relying on AI *becomes* the standard? If a hospital protocol mandates using an AI diagnostic tool, and a doctor follows it in good faith, they are essentially following "standard practice." If that AI makes a calculation error and the patient dies, the doctor hasn't been "reckless" or "grossly negligent"—they were just compliant. This creates a massive blind spot. The doctor lacks the *mens rea* (because they trusted a validated system), but the patient is still dead. We can't punish the doctor for doing exactly what they were trained to do.

5.3 Why We Can't Blame the Developer: So, if we can't blame the doctor, can we blame the coder? Indian law makes that nearly impossible due to the requirement of **Proximity**. In ***Keshavlal v. State of Gujarat*** **(2011)**[5], the Court held that criminal negligence requires a **direct nexus** (connection) between the act and the death.

- **The Distance:** An AI developer writes an algorithm in Silicon Valley or Bangalore. Years later, that algorithm is trained on new data, modified by a hospital IT team, and finally used on a patient.
- **The Result:** The developer is too "remote" from the crime scene. There are too many intervening acts to establish a direct causal link for criminal homicide.

5.4 The Void: This leaves us in a disturbing legal limbo which I call **"Medical Death without *Mens Rea*."**

- **The Doctor** isn't liable (no reckless intent).
- **The Developer** isn't liable (no proximity).
- **The AI** isn't liable (it's software, not a person).
- **The Hospital** faces civil liability, but that's just money—it doesn't address the moral

[5] *Keshavlal v. State of Gujarat, (2011) 7 S.C.C. 639 (India).*

question of criminal accountability.

The Supreme Court in ***State of Maharashtra v. Mohd. Yakub (1980)*** [6]emphasized that criminal liability must flow from conscious wrongdoing, not just bad outcomes. We generally hate the idea of **"Strict Liability"** (liability without fault) for serious crimes. But this creates a normative crisis. We have a system where a patient can lose their life to a technological failure, and the criminal law shrugs its shoulders because no ***human*** had a "guilty mind." The BNS conceptualizes crime as a human-to-human interaction. It is structurally incapable of processing a reality where the killer is a collaborative effort between a well-meaning doctor and a "black box" algorithm. Until the legislature steps in, we are forcing 19th-century legal theories to solve 21st-century algorithmic murders, and it simply isn't working.

6. Rethinking Criminal Responsibility in the AI Era

6.1 The Case for a Structural Update: The challenges posed by AI in healthcare make one thing painfully clear: our existing criminal law principles are struggling to keep up. The **Bharatiya Nyaya Sanhita (BNS), 2023**, while modern in name, still operates on old-school assumptions of human agency, foreseeability, and moral blame. Sticking rigidly to these principles when dealing with AI-driven deaths risks creating a "double injustice"—either we unfairly criminalize doctors who were just doing their jobs, or we let genuine negligence slide because the law can't find a human to blame. The goal of reform isn't to throw out *mens rea*, but to reinterpret it for a world where machines make life-or-death decisions.

1. "Shared Responsibility": We need to stop looking for a ***"lone wolf"*** villain in medical AI cases. Modern healthcare is a team sport involving doctors, hospital admins, and algorithms.

- Instead of pinning everything on the doctor (the easiest target), the law should recognize **distributed responsibility**.

- Liability should track ***control***. If a hospital administration buys a cheap, untested AI system and forces doctors to use it, the ***hospital*** (or its directors) should face the heat, not the surgeon. This aligns criminal culpability with actual authority, preventing doctors from becoming scapegoats for institutional failures.

[6] *State of Maharashtra v. Mohd. Yakub, (1980) 3 S.C.C. 57 (India).*

2. "Standard of Care" for AI: Right now; doctors are flying blind. There is no section in the BNS or the National Medical Commission guidelines that says ***how much*** a doctor should trust an AI.

- We need urgent statutory guidelines. If a doctor ignores a correct AI diagnosis, is that negligence? If they follow a wrong AI diagnosis, is *that* negligence?

- The law must clarify that AI is a **"co-pilot," not the captain.** We need codified standards that protect doctors who use AI as a tool but punish those who use it as a substitute for their own brain. This would give courts an objective yardstick to measure **"reasonable conduct"** instead of guessing after the patient has already died.

3. Criminal Law vs. Regulatory Safety: We have to remember that not every medical accident needs a handcuff. Criminal law is a blunt instrument—it's meant for the worst offenders, not for system errors.

- We should offload the technical policing to specialized regulatory bodies (like a "Medical AI Watchdog"). Let them handle certification, auditing, and fines for buggy algorithms.

- This keeps criminal law as a measure of *last resort*—reserved for cases where there was a conscious, reckless disregard for safety—while ensuring that the tech itself is constantly monitored by experts, not judges.

4. Redefining "Gross Negligence" for the Algorithm Age: The Supreme Court's ***"gross negligence"*** standard needs a software update.

- Courts need to distinguish between **"Reasonable Dependence"** and **"Blind Reliance."**

- **Example:** If a doctor sees clear clinical signs that a patient is dying, but ignores them just because the **"computer says they are fine,"** *that* should be gross negligence. But if the case is complex and the doctor relies on a highly-rated AI for a diagnosis that turns out wrong, that should remain a civil matter, not a criminal one. This nuance protects the *mens rea* requirement while punishing genuine laziness or recklessness.

5. Breaking the Legislative Silence: Finally, the elephant in the room, the BNS doesn't even mention "Artificial Intelligence." This silence is dangerous.

- The legislature needs to wake up and formally acknowledge **"Algorithmic Harm"** as a legal category. We don't necessarily need to give robots rights, but the law must recognize that when a machine is part of the murder weapon, the rules of causation change.

- Without this acknowledgement, we are leaving everything to ad-hoc judicial interpretation. We are asking judges to fix a legislative gap, which leads to unpredictable verdicts. Acknowledging tech-mediated harm would finally give our courts the interpretive tools they need to deliver justice in the 21st century.

7. Conclusion

The integration of artificial intelligence into healthcare represents far more than a technological upgrade; it marks a **fundamental transformation in medical decision-making**. While AI systems hold immense potential to improve accuracy and save lives, they simultaneously expose **deep structural limitations within existing criminal law**. The Bharatiya Nyaya Sanhita, 2023, despite its modern framework, continues to operate on a **strictly human-centric understanding of crime**—one that assumes every offence must originate from a conscious and morally blameworthy mind. In the age of algorithmic medicine, this assumption is increasingly difficult to sustain.

AI-assisted medical deaths reveal a **serious legal vacuum**. Doctors may act in accordance with hospital protocols, institutions may rely on certified technologies, and developers may remain causally distant from the patient, yet a fatal outcome may still occur. In such situations, criminal law finds itself **unable to identify a "guilty mind" (*mens rea*)**, despite the presence of irreversible harm. Punishing medical professionals for algorithmic failures they neither designed nor controlled would be unjust, just as denying accountability to victims' families would undermine confidence in the justice system.

The solution does not lie in abandoning the principle of *mens rea* or converting criminal law into a regime of **strict liability**. Rather, it lies in **rethinking criminal responsibility within technologically complex systems**. The law must recognise *reasonable reliance* on artificial

intelligence, while clearly distinguishing it from *blind or reckless dependence*. Criminal liability must therefore correspond to **actual control, decision-making authority, and conscious disregard for patient safety**, rather than mere participation in an AI-driven process.

Ultimately, the challenges posed by artificial intelligence demand a **careful balance between innovation and accountability**. Criminal punishment must remain a **measure of last resort**, reserved only for cases involving genuine recklessness. At the same time, **systemic and technological failures must be addressed through regulatory oversight and institutional responsibility**, rather than through individual scapegoating. Unless the legislature provides clear statutory guidance on *algorithmic harm*, courts will continue to struggle within outdated frameworks. The future of justice in healthcare depends upon **aligning legal theory with technological reality**—not by fearing artificial intelligence, but by ensuring that the law evolves alongside it.

BIBLIOGRAPHY

I. Statutes

Bharatiya Nyaya Sanhita, No. 45 of 2023, India Code (2023).

II. Cases

Indian Cases

1. *Jacob Mathew v. State of Punjab, (2005) 6 S.C.C. 1 (India).*
2. *Keshavlal v. State of Gujarat, (2011) 7 S.C.C. 639 (India).*
3. *Kusum Sharma v. Batra Hospital & Medical Research Centre, (2010) 3 S.C.C. 480 (India).*
4. *State of Maharashtra v. Mohd. Yakub, (1980) 3 S.C.C. 57 (India).*

Foreign Cases

1. *Bolam v. Friern Hospital Management Committee, [1957] 1 W.L.R. 582 (Q.B.).*
2. *Bolitho v. City & Hackney Health Authority, [1998] A.C. 232 (H.L.).*