BEYOND CLINICIAN LIABILITY

AI in Healthcare: A Medico-Legal and Ethical Perspective

SAMLA Conference • Dr Ife Ofunne
MBBS| MPhil (Stellenbosch) | LLM (MedLaw) UKZN| MSc (UK)| PGD (Applied Ethics)| PGD (Family Medicine) Adv Dip (Aesthetic Medicine) cum Laude (SA)

From “Who is liable?” → “Where did the system fail?”

The problem with the “AI is only a tool” narrative

The traditional model places liability on the clinician: AI is treated as a tool, not a juristic person.

But what if the tool is imposed by the organisation?

What if the clinician is still learning to use it?

What if the clinician had no meaningful role in selecting, validating or developing the system?

A bad workman may quarrel with his tools — but what if the tool is defective, opaque, mandated or poorly governed?

Unpacking liability: tool questions vs policy questions

TOOL QUESTIONS

Was AI used appropriately?

Was it validated for this clinical context?

Was the clinician trained?

Was human oversight meaningful?

Was the output explainable / contestable?

POLICY QUESTIONS

Was there an organisational AI policy?

Who approved and reviewed it?

What law/regulation governed deployment?

Was risk assessment performed?

Who bears responsibility for governance failures?

AI error ≠ automatically clinical negligence

Negligence remains a question of duty, breach, causation, reasonableness and foreseeability.

Wrong AI output does not, by itself, establish negligence.

Wrong clinician input does not, by itself, establish negligence.

An adverse patient outcome does not, by itself, establish negligence.

A correct diagnosis may still be followed by harm from implementation failure or an idiosyncratic/unforeseen event.

A wrong diagnosis that causes no harm may be an error — but not necessarily actionable negligence.

Standard of care: the unit of analysis is changing

Traditional lens: Was the clinician’s conduct reasonable for a practitioner in the circumstances?

AI-supported practice adds a second systemic layer: Was the AI-enabled clinical system reasonably designed, deployed, monitored and governed?

Relevant organisational factors may include: AI policy, procurement, validation, training, oversight, audit, incident reporting and quality assurance.

The standard of care may therefore have both a clinician component and a system / organisational component.

Clinical competence + technological competence + organisational governance

Causation is not the same as liability

Causation is integral to liability — and to meaningful reform.

The AI-enabled causal chain may include: clinician → algorithm → data → developer → vendor → deployer → institution → patient.

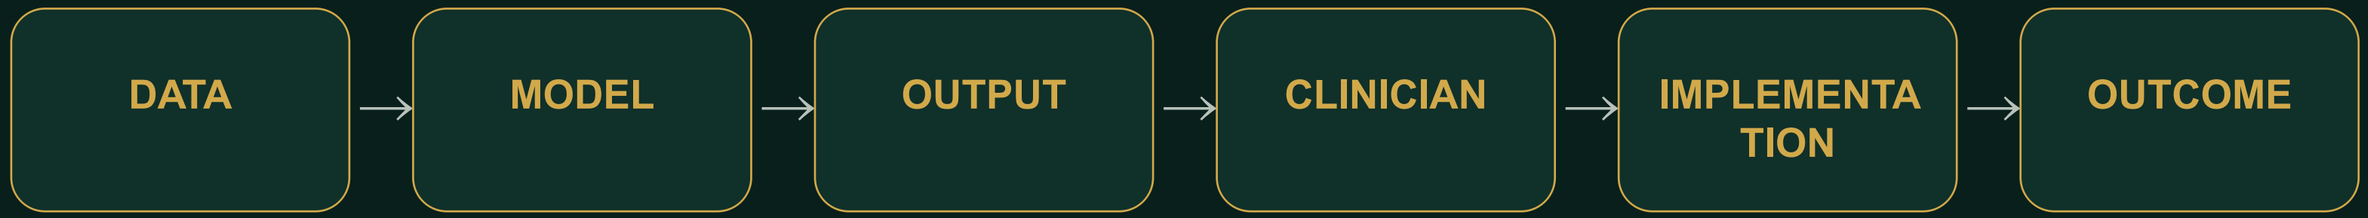
The clinically relevant question is not simply “Who is liable?”

It is: “Which failure contributed to the harm, and was that failure legally attributable?”

Who/what caused the harm? ≠ Who can we blame?

The AI–clinician incident review

At what point did the legally relevant failure occur?



Who controlled that point?

What duty applied?

Was the failure foreseeable?

Did it cause patient harm?

What system change prevents recurrence?

Duty of care: who actually owed the duty?



Traditional clinical practice → clinician-centred duty

AI-assisted practice → potentially distributed duties across the socio-technical system

Evidence gap: AI is moving faster than clinical validation

2024 Lancet Digital Health scoping review

86 RCTs • 70/86 (81%) reported positive primary endpoints

63% were single-centre trials; 92% were conducted in a single country.

The authors raised concerns about generalisability and practicality.

The review also highlighted limited demographic reporting and potential publication bias.

Most trials focused on diagnostic yield/performance rather than broader patient-relevant outcomes.

Source: Han et al., The Lancet Digital Health, 2024.

Whose AI is this?

Clinicians? health organizations? technology companies? — or a genuine co-production among 'stakeholders'?

AI is being deployed into clinical decision-making, but clinicians are not always material participants in development, validation or procurement.

A multi-billion-dollar AI industry creates powerful incentives for rapid adoption.

If clinical AI is allowed to reshape the standard of care, clinicians must be part of the governance conversation.

Otherwise, we risk creating a new form of technological paternalism: technology dictates practice, while clinicians absorb the liability.

Ethics I — Autonomy & disclosure

Should patients be informed that AI contributed to their care or decision?

Does disclosure depend on material risk, context and patient expectations?

Where does therapeutic privilege fit — and can it justify non-disclosure?

Can meaningful consent exist when the patient (and or clinician) does not fully understand the role of AI ?

Transparency should facilitate autonomy not become a technical data dump.

Patient autonomy is not simply a demand for technical transparency; it is meaningful participation in decisions affecting the patient.

AI disclosure: transparency without turning consent into a technical manual.

Ethics II — Beneficence, non-maleficence & justice

BENEFICENCE

AI must demonstrably serve the patient's best interests. Risk–benefit analysis remains clinical and moral, not merely technical.

NON-MALEFICENCE

Do not introduce avoidable technological harms — including automation bias, alert fatigue, hidden error and unsafe implementation.

JUSTICE

Who benefits? Who is represented in the data? Who bears the risks? African patients should not simply become consumers of models developed elsewhere.

Ethics III — From principles to moral philosophy

Kantianism: duties to respect persons, avoid instrumentalising patients, and take obligations seriously.

Virtue ethics: phronesis — practical wisdom — is essential when rules cannot anticipate every AI-enabled clinical scenario.

A clinician cannot outsource moral agency to an algorithm.

AI may inform judgment; it cannot replace professional responsibility, contextual reasoning or ethical discernment.

The algorithm can calculate. The clinician must still deliberate.

The governance response

Every health organisation deploying AI should have an AI governance framework — not merely a software licence.

Prospectively allocate risks through procurement, contracts, indemnities, audit rights, validation requirements and incident-reporting mechanisms.

Define who may use AI, for what purpose, under what supervision, with what documentation and escalation pathway.

Train clinicians before deployment — and involve clinicians in selection, validation and review.

Create AI-specific incident review processes that examine system failure, not only individual blame.

A new professional pathway: the Clinician–AI Scientist

Clinical AI needs professionals who understand medicine, data, ethics, law and implementation.

A Clinician–AI Scientist pathway could bridge the development–deployment divide.

The objective is not to turn every doctor into a programmer.

It is to ensure clinicians have meaningful agency in AI design, validation, adoption, oversight and incident review.

This can become both a medico-legal reform and a form of bioethics revisionism.

Inclusive AI governance is safer than retrospective liability.

The questions we should be asking

At what point in the AI-enabled decision chain did the legally relevant failure occur?

Who controlled that point?

What duty applied?

Was the failure foreseeable and unreasonable?

Did it cause the patient’s harm?

What governance failure allowed it to occur?

THE CENTRAL PROPOSITION

We must reject the story that AI liability is simply clinician liability by another name.

AI changes the clinical decision architecture.

Therefore, law, ethics and governance must change with it.

The question is no longer only “Who is liable?”

It is also: “Where did the system fail?”

Selected references

Han R, Acosta JN, Shakeri Z, et al. Randomised controlled trials evaluating artificial intelligence in clinical practice: a scoping review. The Lancet Digital Health. 2024;6(5):e367–e373. doi:10.1016/S2589-7500(24)00047-5.

WHO. Ethics and governance of artificial intelligence for health. World Health Organization.

UNESCO. Recommendation on the Ethics of Artificial Intelligence. 2021.

EU. Artificial Intelligence Act (Regulation (EU) 2024/1689).

Additional legal and market-growth references should be inserted from the presenter’s final literature set before submission.