ClinicalKey AI

Free, open and unaccountable? The growing risk of generalist AI in clinical care

Exploring how Europe's evolving AI and data regulations are forcing a long overdue reset in how clinical AI is designed, governed and trusted — and why healthcare can no longer treat AI as just another productivity tool.

By Dr. Rahul Goyal, MD
Clinical executive and practising clinician

ELSEVIER

Advancing human progress together

A new line has been drawn for clinical AI in Europe

Across Europe, a new and more consequential conversation is taking place around the role of artificial intelligence (AI) in healthcare. In many settings, it is already part of the day-to-day clinical workflow. Every day we're seeing new use cases, from drafting patient notes to supporting decisions where time is limited.

In the years I've practiced medicine, adoption of new technologies has always lagged behind other industries. And yet, much of the early wave of AI slipped into practice without the same expectations placed on new treatments, medical devices or even clinical guidelines. That period is beginning to give way to a more structured regulatory environment, as policymakers seek to define how AI can be used safely and responsibly in healthcare.

This regulatory shift is also one that Elsevier's parent company, RELX, is engaged in directly. RELX has been appointed to the **EU AI Act's Advisory Forum**, the official expert body advising the European Commission on implementation of the AI Act. Its appointment, from more than 700 applications — and one of only 32 companies — reflects the group's standing in the AI policy landscape and the importance of ensuring that science, research, healthcare and trusted information perspectives are represented as Europe defines responsible AI in practice.

This moment reflects a broader recognition that healthcare is not a neutral testing ground for innovation. Clinical AI operates in high-risk environments, where patient safety, professional accountability and evidence integrity are non-negotiable. While innovation has moved quickly, Europe is now formalizing what acceptable and responsible AI looks like in clinical settings. When AI tools influence clinical decisions, the margin for error narrows and the consequences of poor governance become far more serious.

Why patient care depends on holding healthcare AI to a higher standard

Despite AI adoption in healthcare accelerating faster than anyone could have predicted, governance has lagged behind use. This imbalance is not sustainable, and it places an unreasonable burden on clinicians to compensate for gaps they did not create.

> **Availability and speed may drive adoption, but trust, regulatory readiness and accountability will increasingly determine which AI tools can operate at scale in clinical settings.**

Patient care is becoming more complex, with aging populations and increasing multimorbidity both contributing to rising clinical workloads, leaving many of us looking for support that genuinely helps, rather than creating additional steps or administrative burden. When tools emerge that promise to save time or simplify decisions, it is only natural to see interest in how they might help relieve these pain points.

According to Elsevier's 2026 Clinician of the Future report, nearly half of clinicians report using an AI tool for work purposes, and among those users, more than half rely frequently or always on generalist AI tools. These tools are often the very same tools our patients are using and referencing

when they come in for an appointment. This widespread uptake underscores both the perceived value of AI and the reality that many tools being used were not designed or validated for regulated clinical environments — in other words, not fit for purpose for clinical workflows.

What's emerging is a clear gap. While institutions and governments are still working through how these tools should be introduced, those of us practicing medicine on the frontline are already using what is readily available, often with limited guidance.

In healthcare, the gap between what is governed versus what is being used in practice has a much greater impact compared to other industries. As clinicians, we spend years training and honing our judgment, asking the right questions, assessing evidence, understanding where uncertainty sits. There is also a professional obligation underpinning it all, a clear expectation to do no harm and act in the patient's best interest. That responsibility is closely tied to the trust patients place in us, often when they are most vulnerable.

> **If that gap is left unaddressed, reliance on tools that cannot be explained or examined will steadily erode that trust.**

In healthcare, patient safety cannot be an afterthought or a reaction to widespread implementation. We should not be expected to act as the final safeguard for tools that lack clear governance, validation or accountability. When responsibility is pushed down to the individual user, risk is unevenly distributed and patient safety becomes harder to guarantee.

The Clinician of the Future report also highlights where institutions are struggling most, with AI governance and training consistently ranking among the lowest-performing areas. Fewer than half of healthcare organizations report having effective AI governance in place, and even fewer provide adequate training on the use of AI tools. This gap points to a systemic issue: responsible AI adoption cannot be achieved through individual vigilance alone. It must be embedded at the level of design, deployment and organizational oversight. For clinical and hospital leaders, this is now an urgent responsibility rather than a future ambition. When clinicians are already reaching for whatever is available, the task is to put credible, fit-for-purpose AI in their hands and actively steer adoption towards it so that the safe, governed choice becomes the easy one.

As regulation tightens across Europe, proactive governance is no longer optional. It is the only sustainable way to align AI use with clinical responsibility and patient trust.

Some clinical AIs are ready for regulation. Many are not.

Not all AI used in healthcare is created equal. Yet too often, clinical and generalist AI tools are discussed as if they carry the same level of risk, accountability and readiness for use in healthcare settings.

This is especially important as many widely used, free generalist tools were not developed for clinical use. They are typically trained on broad internet content rather than curated medical evidence and often lack the clinical oversight, governance and accountability expected in healthcare. As a result, they can generate inaccurate, incomplete or outdated information, perpetuate biases and unsupported generalizations, and provide outputs with limited transparency around how conclusions have been reached.

What should separate clinical AI from general-purpose tools is trust. **When it comes to using AI tools, the first thing I ask myself is, where is the answer coming from, and can I trust it?** The stronger, more mature systems are usually built on curated clinical evidence, with sources that can be checked and reviewed. I am also more comfortable using tools where there is clear editorial or clinical oversight behind the content. At the point of care, that is the test that matters: every answer should be explainable and verifiable, grounded in credible content I can trace, check and challenge before I act on it.

Systems that generate opaque responses from mixed or unverified sources are much harder to validate and much harder to trust in clinical settings.

That is where AI maturity starts to matter. When governance and validation are weak, we spend more time double-checking outputs, tracing sources and deciding whether a recommendation is reliable enough to act on. What looks fast or low-friction on the surface can create extra cognitive load, duplicated work and additional legal or clinical risk for both clinicians and institutions.

More importantly, without clear governance and oversight, there is a growing risk that healthcare professionals will begin relying on freely available AI tools for tasks they were never designed to support. Whether it's drafting clinical notes, interpreting laboratory results or summarizing patient histories, the privacy, safety and accountability implications are not always fully understood.

Recent examples show that the risks of AI misuse in healthcare are already being felt in practice. In February 2026, a Nature Medicine study on a generalist AI tool found that the system under-triaged 52% of emergency cases in structured testing, including directing patients with diabetic ketoacidosis or impending respiratory failure to 24–48-hour follow-up rather than emergency care, while also showing inconsistent crisis-safeguard activation in suicidal ideation scenarios. Earlier this year, Ontario's auditor general found that AI scribes used by doctors could hallucinate, invent patient details and record incorrect medication names, while results from ECRI's 2026 Health Tech Hazard Report found that the misuse of AI chatbots in health care is the most significant health technology hazard for 2026.

These are not isolated problems; they reflect a longer-standing healthcare challenge around privacy, accountability and unsafe digital adoption, now amplified by more capable and more widely used AI tools.

In healthcare, "free" or open access is not a governance model. Tools that draw heavily on open access or unverified sources may require additional scrutiny precisely because evidence quality, provenance and accountability are central to clinical decision-making. As regulatory expectations rise, only those AI systems designed with these requirements in mind will be able to scale responsibly in European healthcare.

ClinicalKey AI as an example of governed, evidence-led AI fit for use in healthcare

So, what does healthcare-appropriate AI look like in practice? ClinicalKey AI is one example of a system built specifically for healthcare rather than general-purpose use.

From my perspective as a clinician, the difference is often reflected in the evidence base and transparency behind the system. ClinicalKey AI draws on curated medical content, including peer-reviewed journals, clinical guidelines and reference texts, rather than broad, unfiltered internet sources. The underlying content spans a wide range of specialties and includes both primary literature and guidance from professional medical societies.

But the content itself is only part of the picture. Equally important is understanding who is responsible for curating, maintaining and governing that information. Elsevier has been

involved in scientific and medical publishing for more than 140 years, with longstanding relationships across the clinical, research and academic communities. That history is reflected in the development of tools such as ClinicalKey AI through established editorial processes, governance and experience in managing scientific information at scale. At a time when many AI tools are emerging from companies whose primary focus is technology rather than healthcare or science, I see reassurance in a platform that is built on an existing foundation of medical knowledge and publishing expertise.

What I find particularly useful is that the sourcing is transparent. Clinicians can trace recommendations back to the underlying evidence, review citations directly and interrogate where conclusions are coming from. Elsevier has also expanded the platform's "glass box" approach with real-time citation validation, clinician-in-the-loop review processes and daily refreshed content updates across more than 1,000 journals.

That creates a very different experience from using a general-purpose LLM without strong clinical guardrails. With many open models, clinicians are often presented with fluent answers without visibility into source quality, evidence hierarchy or how the model arrived at a recommendation. The issue is not that these models are incapable of being useful. It is that, in clinical environments, unsupported confidence can quickly become a patient safety and governance problem.

This does not remove the need for clinical judgment. When reviewing an answer, we still need to assess the recommendations provided in context, challenge outputs where appropriate and make decisions based on the patient in front of us. But systems designed around evidence quality, editorial oversight, traceability and validated medical content create a much safer foundation for clinical use than models optimized primarily for conversational fluency or internet-scale breadth.

As the EU AI Act takes effect, such characteristics are no longer differentiators, they are prerequisites. Evidence-led, governed AI aligns more naturally with the direction of travel in European regulation and with the long-standing values that underpin clinical practice globally.

Trust and leadership will define the next phase of clinical AI

AI will remain central to the future of healthcare innovation, but Europe is sending a clear message that scale alone is no longer enough. As regulatory expectations evolve, the EU AI Act reinforces principles healthcare has always depended on, such as patient safety, clinical accountability and evidence integrity.

In this new landscape, **leadership in healthcare will not be defined simply by who adopts AI fastest, but by who adopts it responsibly and ethically.** Increasingly, that means giving

clinicians a credible, evidence-led alternative and making it the easier choice, rather than leaving them to default to whatever is free and open. Healthcare systems are under real pressure to innovate, from workforce shortages and rising patient demand to administrative burden and widening gaps in access to care. AI will inevitably become part of how organizations respond to those pressures, but its value will depend on how safely and thoughtfully it is deployed.

At the same time, the relationship between clinicians and patients is already changing. Patients now have access to many of the same general-purpose AI tools as us, and increasingly arrive with AI-generated information, interpretations and recommendations in hand. That shift makes trust even more important. If we are expected to rely on systems that cannot clearly explain where information comes from or how conclusions were reached, there is a risk of further eroding confidence at a moment when healthcare interactions are already becoming more complex and influenced by new technologies.

For healthcare organizations, the challenge they currently face is around selecting technologies that can support clinical decision-making without undermining accountability, transparency or patient trust. The decisions being made now will shape not only regulatory compliance, but how credible and trusted these systems remain once they are embedded into everyday care. The leaders who build on a credible, evidence-led foundation now will help define what trusted, responsible clinical AI looks like. Anything less does not belong at the point of care.

To explore what governed, evidence-led clinical AI looks like in practice, request a free 90-day trial of ClinicalKey AI.*

*Individual subscriptions and trial availability vary by country and are currently offered in selected markets, including the UK and parts of Europe.

26-2600 TM/MR