InnoStrat

Medical Affairs' AI Survival Guide

Lessons from the legal sector for Medical Affairs AI adoption

June 2026

InnoStrat

Same Risk - Different Robe

By
Stanislav Smirnov, Medical Director
Mikkel Nielsen, Senior Partner

What Law's AI experience tells Medical Affairs

Two regulated professions, one shared problem: Interpreting governed content, contextualising the content to a specific situation, and standing behind an output that must survive formal review. Law is roughly 18–24 months ahead of Medical Affairs on this curve. That makes Law's experience a usable forecast for Medical Affairs.

InnoStrat

Executive Summary

Medical Affairs (MA) and law firms share a structural job:
Take regulated, authoritative source material, interpret it, tailor it to a specific situation, and produce content that a formal review process must clear.

For Law that gatekeeper is the court and the bar; for MA it is medical-legal-regulatory (MLR) review and the named medical signatory.

Because the legal sector reached majority AI adoption ahead of MA, its missteps now show MA what is coming.

The legal experience points to six lessons:

1

Individual adoption outruns institutional readiness. The gap, not the Technology, is the core risk.

2

Bans backfire, pushing staff toward unsanctioned consumer tools (“Shadow AI”) and destroying data visibility.

3

The durable use cases are document- and research-adjacent, covering the same drafting, summarisation, and search tasks MA is already automating.

4

The verification burden is the true cost centre, and it is where professionals get hurt.

5

Liability lands on the human who failed to check, never on the tool.

6

Oversight bodies move late and inconsistently, so internal governance cannot wait for external rules.

InnoStrat

Direct translation

For MA, the translation is direct:

- Expect fast bottom-up uptake of literature, drafting, and MLR-support tools;
- Pair any restriction with approved, healthcare-grade alternatives; treat MLR and the medical signatory as more critical as content volume scales;
- Invest in verification discipline and training now, because the legal record shows warnings alone do not prevent the failure.

What we see in Medical Affairs organisations today is consistent with this forecast: Adoption is still uneven and is held back less by reluctance than by uncertainty over what is sanctioned. It is the same gap, just a few steps behind Law.

InnoStrat

Why the analogy holds

The comparison is more than rhetorical. Both functions perform regulated-content interpretation and approval, and both must contextualise general authority to a specific case.

A lawyer adapts case law and statute to a particular client and jurisdiction; an MSL or medical advisor adapts trial data, guidelines, and real-world evidence to a particular KOL, therapy area, and product life-cycle stage.

In both, the deliverable, whether a brief, a response letter, a slide deck, or an advisory-board follow-up, must pass a formal gate before it carries weight.

The two fields also share the same failure surface. Generative models produce fluent, confident text that can misstate a source, cite something that does not exist, or attach a real reference to a claim it does not support.

In Law, this surfaces as fabricated citations in court filings; in MA it would surface as unsupported or mis-referenced scientific claims—precisely the category MLR exists to catch. Because the underlying risk is identical, the legal sector's two-year head start is predictive, not just illustrative.

InnoStrat

What MA can learn from law about AI use

1

Adoption races ahead of governance

This is the clearest signal. Among legal professionals, generative-AI use roughly doubled in a single year, from about 31% to nearly 70%, while more than half of firms provided no AI training and over 40% had no formal policy.[1]

The defining tension was repeatedly framed as individuals sprinting while their institutions barely kept pace. MA should expect the same pattern: MSLs, medical information teams, and writers will adopt faster than their organisations can govern.

The same surveys show adoption and apprehension rising together: Even as a majority of firms formally adopted generative AI, a large majority of leaders reported concern about its reliability and risk — and that unease did not slow uptake, which makes governance, not caution, the differentiator between organisations that get value from AI and those that get burned by it.[2]

What we observe in practice

Industry surveys already show MA with high planned AI investment but most organisations still in an “ideation or foundational” phase — the same gap as Law, one step earlier.

This is the picture we see across Medical Affairs organisations: real day-to-day use is still uneven, concentrated in a minority of staff and lightest in the most heavily governed functions, the closest internal analogue to Law’s slow-moving, high-stakes work.

The uneven, function-by-function curve the legal sector went through is recognisable inside pharma today, and broader workforce research points the same way, with most organisations reporting employees already using AI well ahead of the policies meant to govern it.

[1] 8am, 2026 Legal Industry Report (Nicole Black, March 2026), based on a survey of more than 1,300 legal professionals conducted September–October 2025. The report found generative-AI use among legal professionals rose to 69% (from 31% a year earlier), with 54% of firms providing no AI training and 43% lacking a formal AI policy.
[2] SurePoint Technologies, 2025 State of the Legal Industry report. The report found that 63% of mid-sized law firms had formally adopted generative AI while 81% of firm leaders reported internal concern about its reliability and risk.

InnoStrat

What MA can learn from law about AI use

2

Bans backfire into “Shadow AI”

Where firms prohibited AI without offering a sanctioned alternative, staff under efficiency pressure turned to free, consumer-grade tools on personal devices, and the firm lost all visibility into where confidential data was going.

For MA, the equivalent of “client data” is unpublished trial results, proprietary data on file, and HCP or patient information. A lockdown without approved tooling does not reduce exposure; it relocates it to channels no one can see.

This is the same data-privacy risk consistently named as MA’s top AI concern, made worse by policy.

What we observe in practice

In our experience the dominant barrier in Medical Affairs is the same one Law ran into: not reluctance, but uncertainty over what is permitted. People want to use AI and simply cannot tell which tools are sanctioned.

Broader workplace surveys describe the same dynamic, with a large share of employees admitting to using unapproved tools and pasting sensitive material into them when no clear, approved option is provided. When that uncertainty meets efficiency pressure, the consumer-tool workaround, Shadow AI, is the predictable result.

InnoStrat

What MA can learn from law about AI use

3

The use cases that stick are document- and research-adjacent

Law's durable applications, namely legal research, drafting, summarisation, and document review, map almost one-to-one onto MA's emerging core uses: Literature monitoring and summarisation, first-draft response letters and slide decks, insights synthesis, and MLR support.

Both sectors are now entering the same second wave, "agentic" AI that executes multi-step workflows autonomously.

In legal, agentic use in early 2026 looked like generative AI did in 2024, early experimentation pointing toward broader implementation, accompanied by warnings that greater autonomy demands stronger oversight.

MA's agentic-MSL-workflow narrative is on the same trajectory, roughly a step behind.

What we observe in practice

We see the same prioritisation play out in practice: the use cases that gain traction first are the low-risk, document-adjacent ones, such as meeting and content summarisation, first-draft correspondence, and competitive-intelligence briefings.

Medical-content generation is approached most cautiously, precisely because it sits in the MLR-sensitive territory where Law's hallucination failures concentrated.

What Law predicts about the challenges MA will face

1

The verification burden is the real cost

The most important warning from Law is that AI does not remove work so much as shift it toward checking. Hallucinated citations moved from a 2023 novelty to a systemic 2025 problem, with the large majority of documented court decisions on the issue arising in a single year and the tracked rate climbing from a couple of incidents a week to several a day.[1]

Crucially, this kept happening even though everyone knew the risk; studies found that warnings alone produced only modest improvements in verification behaviour. A sharper finding still: the harder a claim is to support, the more the model tends to fabricate support for it, because it is optimising to satisfy the request.[2]

For MA, that is exactly the danger zone. An MSL or writer seeking backing for a weakly evidenced claim will be handed plausible-looking but invalid support, and that is the content MLR must intercept.

Practitioners in Law have given this cost shift a name worth borrowing: **verification debt**. As the cost of generating output collapses, the cost of verifying it rises. Now the verification cost is dominant.

A solution has two parts, generation and verification, and generation has little value on its own because expert humans remain better at validating output than the model is at validating itself.

Carelessly generated work creates a debt that is always paid eventually, in one of three ways: senior reviewers spending their time checking it, the organisation silently absorbing elevated risk, or paying outside experts to do the checking.

For MA, the named medical reviewer is the one who pays, so the mitigation is to **generate more intelligently rather than simply faster.**

[1] Damien Charlotin, AI Hallucination Cases Database (damiencharlotin.com/hallucinations), updated continuously. Charlotin reports that documented filings citing AI-fabricated content rose from roughly two a week in early 2025 to two or three a day by late 2025; the SurePoint 2025 report, drawing on the same database, counts 487 U.S. court instances in 2025—more than ten times the 2024 total.
[2] As Damien Charlotin put it: "The harder your legal argument is to make, the more the model will tend to hallucinate, because they will try to please you." (Quoted in The Daily Record, 13 October 2025.)

What Law predicts about the challenges MA will face

2

Liability attaches to the human, not the tool

Every sanction in the legal record turned on a failure to verify, not on the model's error itself.

Courts also calibrated their response to candour and scale: harsher when AI use was evident yet denied, lighter when professionals were upfront; a few stray errors drew less than many errors across multiple filings.

The MA lesson is direct: the named medical reviewer or signatory owns the output, "the AI produced it" is no defence, and a culture of disclosure plus low error tolerance is protective.

The three hallucination categories Law identified translate cleanly: fictitious sources become fake references; fabricated citations to real documents become mis-cited real studies; and real quotes that fail to support the proposition become accurate quotes wrapped around an unsupported medical claim.

3

Formal oversight arrives late and unevenly

For much of the adoption surge, legal had no explicit overarching guidance; courts applied existing rules and individual judges responded inconsistently before standing orders, bar opinions, and court rules gradually formed.

MA's analogous oversight, spanning MLR committees, evolving FDA and EMA expectations, and the EU AI Act, follows the same lagging-then-formalising path.

The practical implication is to build internal governance ahead of external mandates; the firms that waited for the rules to settle absorbed the sanctions in the interim.

Translating the lessons into MA controls

The table below maps each legal lesson to a concrete Medical Affairs control point.

Lesson from Law	Medical Affairs control
Adoption outruns governance	Issue an AI policy and role-based training before tools proliferate; assume bottom-up uptake is already under way.
Bans create shadow AI	Provide an approved, healthcare-grade tool so staff never need consumer apps for proprietary or HCP/patient data.
Document/research uses dominate	Prioritise sanctioned workflows for literature summarisation, draft response letters, and MLR support—the high-value, lower-risk core.
Verification is the true cost	Mandate documented human verification of every AI-assisted claim and citation; resource MLR for higher draft volume, not less.
Liability sits with the human	Make the named reviewer/signatory explicitly accountable; require disclosure of AI assistance and keep an audit trail.
Oversight lags adoption	Stand up internal governance now—don’t wait for FDA/EMA or AI-Act specifics to finalise.

What good looks like

The legal record says what to avoid; the more useful question is what to do instead. The organisations that get this right tend to follow the same shape, regardless of size or therapy area: **governance first, then enablement.**

Lead with clarity

A short, plain-language policy and an approved-tool list remove the ambiguity that pushes people toward consumer tools in the first place. People adopt safely when they simply know what they are allowed to use.

Prove value on low-risk work

Start with the document- and research-adjacent tasks where the upside is high and the data risk is low, building confidence and verification habits before touching MLR-sensitive content.

Deepen into the regulated core last

Bring the most governed, highest-stakes work into scope only once verification discipline, training, and oversight are established, so that the human review owning the output is strengthened rather than bypassed as volume grows.

InnoStrat

The synthesized forecast

Putting the lessons together, the legal experience forecasts a recognisable path for Medical Affairs:

- Rapid, largely bottom-up adoption of drafting, summarisation, and research tools;
- A widening and dangerous gap between that adoption and organisational policy and training;
- Shadow use of consumer tools that amplifies the very data-privacy exposure MA already fears most;
- A steadily rising verification burden that becomes the real cost centre; and accountability landing squarely on the human reviewer, which makes MLR and the medical signatory more central as content volume grows.

The single most transferable mitigation is also the simplest to state:

- Pair any restriction with sanctioned, healthcare-specific tooling;
- Require and document human verification of AI-assisted output;
- Invest in training early, because the legal record demonstrates, repeatedly, that warnings on their own do not prevent the failure.

InnoStrat

How InnoStrat can help

Knowing the shape of a good rollout and running one are different things.

InnoStrat works with Medical Affairs teams to translate these lessons into a plan that fits your organisation—assessing where you are today, defining the governance and sanctioned tooling that lets adoption happen safely, sequencing the use cases that build confidence first, and putting the verification and training disciplines in place before the regulated work scales.

If you'd like to know how this would look in your environment, contact us and ask how we can help you get there.

info@innostrat.net | www.innostrat.net

InnoStrat

A note on the evidence

The figures cited here are drawn from named legal-sector and professional-services sources, footnoted at the relevant points: the 8am 2026 Legal Industry Report, the SurePoint 2025 State of the Legal Industry report, and Damien Charlotin's AI Hallucination Cases Database. Because these sources use differing methodologies and reporting periods, the precise percentages should be read as directional rather than directly comparable.

Observations about how Medical Affairs organisations are adopting AI reflect our own experience in the field rather than a single survey, and are offered as a directional pattern rather than a benchmark. Where broader workplace and professional-services research is referenced, it is included only to show that the same dynamics appear outside Medical Affairs; the central evidence base remains the legal sector's documented experience.

The underlying pattern of adoption outpacing governance, verification as the point of failure, and liability resting on the human is consistent across sources and is the part worth planning around. AI tooling, regulatory expectations, and the EU AI Act's medical-device provisions are all evolving quickly; any compliance or regulatory decision should be confirmed with your legal and regulatory teams.