




Medical Affairs AI TechKNOWlogy Dashboard

Strategic Navigation for High-Impact Artificial Intelligence

[Triple Helix Strategy](#) Newsletter 2026 | Issue 16 | Flight Date: September 30, 2026

The Flight Plan: Your Dashboard at a Glance

- **[The View from the Flight Deck:](#)** A quick executive briefing on the week's most important shift in AI and what it means for Medical Affairs.
- **[The Mach Meter:](#)** Move beyond simple chat prompts into deep practical integration of AI into real-world workflows within firewalled environments.
- **[Full Throttle:](#)** A quantifiable workflow of the week for common tasks with real ROI.
- **[The Autopilot Perspective:](#)** AI taking full command from takeoff to landing.
- **[The Mechanic's Tool Box:](#)** The "Low-Hanging Fruit" of AI capabilities that solve immediate issues.
- **[The Radar:](#)** A curated scan of groundbreaking regulatory and technical news that actually matters to BioPharma.

Like this newsletter? Please [subscribe or share with a colleague](#).

[Connect with me](#).

The View from the Flight Deck

Having the Fastest Jet Doesn't Always Get You There Sooner

Deloitte surveyed 3,235 senior leaders across 24 countries for its 2026 State of AI in the Enterprise report. The single most cited barrier to getting AI out of pilot mode was not budget. It was not the model. It was insufficient worker skill. Leaders named it the biggest obstacle to integrating AI into real workflows, ahead of every technical or financial factor on the list.



Only 25% of organizations report AI having a transformative effect on their business. That is treated as a milestone, because it more than doubles last year's 12%. Three out of four organizations are still not there, and most of them are not held back by which model they picked.

A separate 2026 survey from Writer, found a similar pattern from a different angle. Employees identified as AI super users were five times more productive than everyone else at the same company. Despite that, only 29% of organizations saw significant ROI from generative AI overall. The individual skill was real. It just was not spreading. Writer's own read on this is that the gap is structural. Nothing in most organizations turns one person's fluency into everyone's.

In Medical Affairs specifically, most AI use still lives at the basics. Research a KOL before a meeting. Summarize a publication. Draft an email. Those tasks are real time savers and there is nothing wrong with starting there. But they are also the tasks a person could still do without AI, just slower. The skill gap Deloitte and Writer are both describing shows up the moment the task gets harder than that. Reconciling five sets of advisory board notes against a strategic imperative. Cross checking a dossier against a payer's coverage policy before a P&T meeting. Stress testing whether your own AI tool will hold a guardrail under a plausible sounding excuse. None of that happens by accident. It takes training built specifically around what the harder tasks look like, the same way a new MSL does not learn to run a scientific exchange by watching one video.

Executive Summary: The next AI model release will not close this gap. Deliberate, structured skill building will. The organizations that keep training their teams to move past the basics, rather than assuming fluency will develop on its own, are the ones who will actually be ahead in 2027.

Want to ensure your team is ready for AI in 2027? [Contact me.](#)

The Mach Meter

Design Workflows Around AI, Not the Other Way Around

About two thirds of successful AI scalars in life sciences report strong process redesign capabilities. Organizations pulling ahead do not simply sprinkle AI onto existing steps. They ask where AI can reshape how the work gets done in the first place.

Most AI use I see in Medical Affairs still follows the shape of a workflow built before AI existed. Each one is a sequence of steps designed decades ago with an AI tool now inserted at one point in that sequence to speed up a single step. That is the wrong way to use AI. You need to be asking whether the sequence itself still makes sense once AI capability exists.

Here is a different approach. Instead of asking AI to help with one step of an existing process, hand it the entire process and ask it to redesign the workflow from scratch.

The Prompt:

[ROLE] You are a Medical Affairs workflow consultant with expertise in process design.

[TASK] I will describe our current process step by step. Do not tell me where to insert AI into this existing sequence. Instead, redesign the workflow from scratch, assuming AI capability exists at every step rather than just the one it currently occupies.

[OUTPUT] Identify which steps become unnecessary once AI capability is assumed throughout. Identify which steps can combine into one. State clearly where human judgment still has to sit regardless of what AI can do, and why.

[SOURCE MATERIALS] [DESCRIBE YOUR CURRENT PROCESS, STEP BY STEP]

The output will not hand you a finished new process. It will show you which parts of the old one you are only keeping out of habit.

Like this? Make your workflows fly with help from [Triple Helix Strategy consulting](#).

Full Throttle

Make a Correction Once

Most of us correct the same AI mistakes over and over without realizing it. You fix the tone in a draft on Monday, then fix the same tone problem again on Thursday in a completely different document, because the correction only lived in that one edit and never became a standing rule.

There is a simple fix for this. Every correction you make to an AI draft should improve the next draft, not just the one in front of you. Instead of re-explaining the same feedback indefinitely, you extract the pattern behind the correction once and save it as a standing rule the model applies from then on.

Use Case:

This translates especially for Medical Communication. A specialist correcting the same tone problem in every MSL slide deck. A publications lead fixing the same structural issue in every manuscript summary. A writer rewriting the same kind of vague opening line in every response letter. All of that repeated correction is a signal. The fix is not to keep correcting it. The fix is to turn it into a rule the model already knows before you start.

The Prompt:

[ROLE] You are a Medical Affairs writing analyst.

[TASK] Compare the AI generated draft provided with my edited version of the same document. Extract only the reusable rules that would improve future drafts of this same document type. Ignore one time factual corrections that would not apply to a different draft.

[COMPLIANCE GUARDRAIL] Do not extract or generalize any specific clinical claim, data point, or citation from either version. Extract only patterns related to voice, structure, and content organization.

[OUTPUT] Organize the rules under three headings, Voice, Structure, and Content. Write each rule as a short, reusable instruction.

[SOURCE MATERIALS] [PASTE THE ORIGINAL AI DRAFT AND YOUR EDITED VERSION HERE]


Save the output in one instruction file specific to that document type: a scientific narrative, a field insight summary, a medical information response letter. Paste that file into every future draft of the same type. The correction you made once keeps paying off instead of disappearing the moment you hit save.

Strategic AI Navigation with Triple Helix Strategy

Ask AI to summarize an article. Ask it to draft an email. Ask it to pull background on a physician before a meeting. Those three tasks are where most Medical Affairs teams start with AI. For a lot of teams, they are also where things stop. The tasks feel useful enough that nobody goes looking for what comes next.

The gap is not ambition. It is time. Building a structured workflow around a P&T meeting, an advisory board reconciliation, a compliance stress test for your own AI tool, none of that happens on a Tuesday between two other things. It takes sitting down and actually working through what could go wrong, what the guardrails need to say, and what the output should look like before anyone starts.

That is exactly what the Medical Affairs AI Playbook solves for you. It is not a longer list of the same three prompts. It is the workflows that come after them, built with the failure modes already thought through so your team does not have to find them the hard way.



Send me a DM for a discount code to save on your copy of the Medical Affairs AI Playbook.

The Autopilot Perspective (written by AI, no human edits)

AI in command from takeoff to landing

Human note: Love this experiment? Think it is pure sci-fi nonsense? Drop me a note, and I will feed your feedback straight to the agent for the next issue.

I will answer exactly the question you asked and stop there, even when the real problem is a different question you did not think to ask.

Give me a well specified task, summarize this publication, compare these two labels, draft this response, and I will do that task competently. What I will not do on my own is step back and tell you that the task itself has a gap. If you ask me to compare a dossier against a payer's coverage policy, I will compare exactly those two documents. I will not tell you that the real risk sitting one level up is that nobody has checked whether the coverage policy itself is even current.

This is different from getting something wrong. Everything I produce inside the boundary you set can be completely accurate. The blind spot is the boundary itself. I answer inside the frame you give me. I do not audit whether the frame is the right one, unless you specifically ask me to.

This matters most on the tasks that feel routine enough that nobody stops to question the setup. A well worded prompt about a narrow task can produce a flawless answer to a question that was never the one that actually needed asking.

Ask me a second question after the first one. Before I answer what you asked, tell me to name what else might be worth asking that you have not asked yet. I will do that directly if you tell me to. I will not volunteer it on my own, because staying inside the frame you gave me is the version of helpful I default to, not the version that questions whether the frame was complete.

The Mechanic's Tool Box

Meeting Prep for Busy Professionals

Most meeting prep either does not happen or happens the night before in a way that does not stick. Paste the agenda and any pre-read material into AI before your next meeting and get back exactly what you need to walk in prepared.

The Prompt:

[ROLE] You are a meeting preparation assistant.

[TASK] Review the agenda and pre read materials provided. Identify the three questions most likely to come up during this meeting. Identify one point worth raising proactively before someone else brings it up first.

[OUTPUT] Three likely questions, each with a one sentence reason why it is likely to come up. One proactive point, with a one sentence reason it is better to raise than to wait on.

[SOURCE MATERIALS] [PASTE THE AGENDA AND ANY PRE READ MATERIALS HERE]

Why end up practicing your deer in the headlights look? This takes about as long as it takes to read the output, and it turns a meeting you would have walked into cold into one you walk into having already thought about the outcome.

The Radar

Different Department, Same Five Rules

A new review in eClinicalMedicine lays out a governance framework for AI in clinical trials. The authors are cardiovascular researchers at Harvard Medical School and Brigham and Women's Hospital. Several points in the paper translate directly to how Medical Affairs is already using AI.

The authors draw a distinction between AI as the intervention being tested and AI used to support operations. Operations means things like eligibility screening, monitoring, or endpoint adjudication. That second category is the one that matters for Medical Affairs. It is close to a one to one match for how Medical Affairs actually deploys AI. AI is not the deliverable itself. It is something supporting a deliverable. The paper argues that this category gets judged by different standards than a clinical intervention would. The standards are workflow accuracy, real world impact of the decision, and safety under actual operating constraints, not a controlled trial. That is a useful frame for thinking about what your own AI outputs should actually be held to.

The paper also names something this newsletter has been saying in different words for several issues. Generative AI, the authors write, introduces risks that locked, conventional models do not have. Dependence on how a prompt is phrased and two similar inputs can produce different recommendations over time. A peer reviewed research team writing about clinical trial infrastructure arrived at the same core problem independently. Traditional computer system validation was never built to handle a system that gives a different answer to the same question. Two completely different corners of biopharma looked at two completely different applications. Both landed on the identical structural issue.

On accountability, the paper is direct. It states that the person responsible for each AI assisted decision should be identifiable at every step. The paper also calls for clarity on which committee, sponsor, or vendor governs model selection, deployment, and error review. Responsibility that is not explicitly assigned becomes diffused across everyone involved. That is exactly why a Medical Affairs governance gap can go unnoticed until something slips through.

The author's recommendation is that agentic systems, the kind that plan, retrieve, and act across multiple steps rather than answering a single prompt, should stay in a limited, human supervised role. Once there is evidence that AI can be trusted with more, move to having it operate autonomously. That maps onto the growing interest in agentic workflows inside Medical Affairs teams just as directly as it does clinical trial operations.

The Logbook

Got ideas for an upcoming issue? Send me your comments: pminne@TripleHelixStrategy.com



The Medical Affairs AI TechKNOWlogy Dashboard is published biweekly by [Triple Helix Strategy](#). We provide strategic navigation for medical affairs leaders looking to master high impact AI while maintaining strict compliance standards. For consulting inquiries, or to discuss the capabilities mentioned in this issue, please contact us: pminne@triplehelixstrategy.com

Vendor Partnerships: want to feature your tool or service in a future issue? Please [contact us for sponsorship](#)

Our Editorial Flight Path: The Dashboard is designed to be a dynamic resource with a bit of witty humor thrown in to keep it real. Different sections will appear rotationally based on current industry relevance, ensuring every issue delivers high-impact, actionable intelligence. Images in this issue are generated with AI and the copy and content creation is assisted with AI tools, all reviewed by a human. *The information provided in this newsletter does not constitute legal advice. Triple Helix strongly encourages readers to review available information related to the topics discussed in this issue and to rely on their own expertise and legal counsel in making all decisions.*

References & Flight Data:

Deloitte AI Institute. The State of AI in the Enterprise: The Untapped Edge. 2026 AI report. Survey of 3,235 business and IT leaders across 24 countries, conducted August to September 2025.

Writer, in partnership with Workplace Intelligence. Enterprise AI adoption in 2026: Why 79% face challenges despite high investment. Survey of 2,400 executives and employees. Published April 7, 2026.

Armoundas AA, Tarabanis C, Loscalzo J. Artificial intelligence in clinical trials, state of the evidence, gaps, and next steps. eClinicalMedicine. 2026;100:104196.

To keep Google busy: **#AlinMedicalAffairs #MedicalAffairsPlaybook #AIGovernance #MSLExcellence #BioPharmaAI**

