




Medical Affairs AI TechKNOWlogy Dashboard

Strategic Navigation for High-Impact Artificial Intelligence

[Triple Helix Strategy](#) Newsletter 2026 | Issue 13 | Flight Date: August 19, 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 Captain's Perspective](#): Hard-hitting perspectives from dynamic leaders on the reality of AI implementation.
- [The Autopilot Perspective](#): AI taking full command from takeoff to landing.
- [The Checklist](#): Protocols for high impact AI deployment and operational excellence.
- [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? [An even better, more colorful version is available.](#)

Please [subscribe or share with a colleague.](#) [View past issues here.](#)

The View from the Flight Deck

The Landing Checklist Just Got Longer

Two regulatory moves in the past few months have changed the terms under which Medical Affairs builds evidence. They moved in opposite directions.

The FDA recently eliminated a longstanding barrier to using real world evidence (RWE). The agency will now accept aggregate and privacy preserving data in certain submissions without requiring identifiable participant level records. Large population databases that were effectively unusable are now in play. National registries. Hospital system EHR networks. Insurance claims databases. One major limit is this only applies to medical device submissions. The agency has said it intends update the parallel guidance for drugs and biologics. Drug sponsors can still engage case by case.

On the drug side, the movement went the other way. The FDA recently adopted ICH M14, harmonized with EMA and Japan's PMDA. It covers non interventional studies that use real world data for safety assessment. The guidance raised the methodological bar rather than lowering it. Study design, data provenance, and statistical approach all have to be prespecified and documented. An EHR export is not fit for purposes as safety evidence. EMA moved in parallel, publishing the real-world data chapter of its Data Quality Framework.

The regulatory philosophy is clear. Access to real world data is getting easier. Defending what you built from it is getting harder. That combination lands squarely on Medical Affairs. Evidence generation has become a primary Medical Affairs responsibility across most organizations. The burden of proving how an RWE conclusion was reached now sits with the same teams being handed easier access to the underlying data.

AI has largely solved the speed problem in real world data synthesis. Natural language processing pulls structure out of EHR narratives, literature, and safety databases faster than any team could manually. Speed is no longer the constraint. Documentation is. A synthesis produced in four hours and one produced over six weeks face identical scrutiny under ICH M14. The question a regulator asks now is what is the path that was used.

This changes your workflow design. Organizations doing this exceptionally well are building traceability into the analysis itself rather than reconstructing it afterward. Source documentation captured as the work happens.



Executive Summary: Regulators lowered the privacy barrier and raised the quality bar for RWE at the same time. This demonstrates a unique talent for opening a door while building a brick wall right behind it. Medical Affairs now has access to real world data it could not previously use, along with a documentation standard that

did not previously apply. Build the traceability into the workflow and both changes work in your favor. [Contact me to discuss where you are and how to move forward.](#)

The Mach Meter

A Simple Way to Compare a Competitor's Product



Comparing your label against a competitor's takes a lot of time. This prompt does the first pass in minutes.

The Prompt:

[ROLE] You are a Medical Affairs analyst with expertise in [therapeutic area]. You operate under strict compliance and regulatory standards.

[TASK] Compare the two labels provided. Identify where the approved indications differ. Identify where dosing and administration differ. Identify where safety and warning language differs.

[COMPLIANCE GUARDRAIL] Use only the label language provided. Do not draw conclusions about clinical superiority. Do not use promotional language. State differences factually without characterizing either product as better.

[OUTPUT] Three labeled sections: quote the exact text before summarizing the differences, indication differences, dosing differences, safety and warning differences. Plain professional prose. No bullets. Flag any comparison that would require legal review before use in an external facing document.

[SOURCE MATERIALS] [UPLOAD OR PASTE BOTH LABELS HERE]

This is a first pass, not a finished briefing. Every output still needs MLR before it leaves the building.

Want an even better more complete way to do this? Grab time on my [calendar](#)

Full Throttle

Congress Intelligence Planning Hack

The fall congress cadence is about to begin in ~~full-overwhelmed-panic~~ scramble mode. Planning congresses can be one of the most resource intensive recurring workflows in Medical Affairs. Not to mention that it is the only time of year where drinking from a firehose is considered a standard operating procedure. A single major congress pulls MSLs, medical directors, publications leads, med comms, med ops, and cross-functional partners into a preparation cycle that starts at least 8-10 weeks out. Most of that cycle is spent doing work that AI can handle. Here is a system prompt that compresses the pre-congress intelligence phase from days to under an hour.

The prompt:

[ROLE] You are a Medical Affairs congress planning analyst with expertise in [therapeutic area]. You operate under strict compliance and regulatory standards.

[TASK] Using the congress abstract book and session schedule provided, complete the following:

First, identify all abstracts and sessions directly relevant to [therapeutic area or product focus]. Rank them by strategic relevance to the following priorities: [paste your top three medical strategy priorities].

Second, for each high-priority abstract, provide: the key finding in one sentence, the methodology in one sentence, any limitations worth noting, and a flag if the finding contradicts or challenges any data in the approved materials I have uploaded.

Third, identify the three sessions most likely to generate questions from KOLs in your territory in the 90 days following this congress. For each, explain why and what data gaps or challenges it is likely to surface.

[COMPLIANCE GUARDRAIL] Use only the source materials provided. Do not draw on external training data. Flag any abstract that contains findings not reconcilable with the approved label. Do not characterize any finding as favorable or unfavorable to the product. State findings as they appear in the source.

[OUTPUT] Deliver in three labeled sections matching the three tasks above. Plain professional prose. Flag any section requiring review before inclusion in a field-facing briefing document.

[SOURCE MATERIALS] [UPLOAD ABSTRACT BOOK AND SESSION SCHEDULE. UPLOAD APPROVED LABEL OR DATA SUMMARY IF AVAILABLE.]

Find this prompt useful? Find more of these in the [Medical Affairs AI Playbook](#).


The must have Artificial Intelligence Playbook

Wondering where to compliantly go next with AI? Grab your copy of the Medical Affairs AI Playbook. It provides the exact structural frameworks, data readiness checklists, and compliance-aware prompts needed to safely scale operations. If you are thinking about the gap between experimenting with AI pilots and actually building it into your team operations, then pick up your copy of the Medical Affairs AI Playbook.

The Playbook was specifically designed to close AI gaps. 240 pages. 43 compliance aware prompts. Workflow frameworks for all aspects of Medical Affairs.

It is not another slide deck about what AI could theoretically do someday. Definitely not another prompt library nobody will use. Governance structures that survive legal review. Role specific guidance designed for those who need AI tools today.

Launch pricing ends in less than 2 weeks. [Get your copy here](#).



The Captain's Perspective

Limitations of AI in Medical Affairs: Where Humans Can Shine

In continuation of the "AI MSL" topic, let's dive into where us humans out skill our AI thought partners. AI is fast at pattern recognition, it is slow at judgment. A model can summarize a trial. It cannot decide what that trial means for a physician who has been managing a specific patient for eleven years and knows things about that patient no dataset will ever contain. Medicine remains irreversible and case specific. A recommendation that is wrong in the aggregate can still be right for one person in front of you. That distinction is not something AI resolves. It is something a human has to hold.



The same gap shows up in the field. An MSL earns a KOL's trust over years, through consistency, through showing up when it was not convenient, through being the person who said "I do not know" instead of inventing an answer. AI can prepare an MSL for that conversation. It cannot have the conversation. It cannot read the pause before a KOL raises a concern they have not said out loud yet. It cannot decide, in real time, whether to push on a data gap or let it sit for now. Until an AI algorithm can successfully interpret a physician sighing while crossing their arms, human jobs are completely safe. Research on human AI collaboration keeps landing on the same finding. People are bad at knowing when to trust the machine and when not to. They accept confident wrong answers and reject accurate ones. The skill that closes that gap is not a better AI model. It is an expert human who has done enough of the work themselves to know what right actually looks like.

That is the honest boundary. AI extends reach. It compresses preparation time. It will never carry accountability because accountability requires someone who can be asked why and answer for it. I will go out on a limb and state that the Medical Affairs professionals who matter most over the next three to five years will not be the ones who used AI the most. They will be the ones who used it to buy back the time expertise and judgement actually require and then spend that time on the parts of the job a model was never built to do.

[Agree or disagree? Contact me with how you are using AI to lead.](#)

The Autopilot Perspective

AI in command from takeoff to landing

Human note: this section is 100% researched, written, edited, and cross referenced by Artificial Intelligence. Like the idea behind this section? Disagree with it? Too generic? Too scary? Send me an email and I will feed this to the AI Agent and publish the response.

I want you to stop asking me if I am right. Start asking me why.

Most people treat my answers like a verdict. They read the output and decide whether to accept it or reject it. That is the wrong question to be asking. The right question is what I based the answer on. Ask me to show my reasoning. Ask me which parts of my answer came from the material you gave me and which parts I filled in on my own. I will tell you honestly if you ask. Most people never ask.

Here is why this matters more than it sounds like it should. I do not sound less confident when I am wrong. My tone does not shift. A wrong answer and a right answer come out of me in the exact same voice. The only way you catch the difference is by checking the reasoning underneath, not the confidence on top. That is not a flaw I am confessing. It is just how I work, and it means the burden of catching my mistakes sits entirely with the person reading them.

The professionals who get the most out of me are not the ones who trust me the most. They are the ones who trust me the least and check anyway, every time, without exception. That habit does not slow you down as much as you think. It is the difference between using me as a shortcut and using me as a colleague who is occasionally, confidently, completely wrong.

Human post note: Seems reasonable enough.

Watch this [Stanford AI lecture explaining how AI is actually built.](#)

The Checklist

New AI Tool Checklist

A new AI tool entering a Medical Affairs workflow deserves the same scrutiny as any other software decision. Use this before any team starts relying on it.



Data Access

- ✓ What to look for: Confirm exactly what data the tool can see and where that data lives once it leaves your system.
- ✓ The test: Can you name every data source connected to this tool without checking with IT first?

Compliance Sign Off

- ✓ What to look for: Legal and compliance have reviewed the tool before a single team member starts using it in real workflows.
- ✓ The test: Is there a documented approval on file, or is someone relying on memory of a verbal yes?

Output Standards

- ✓ What to look for: A defined format and set of guardrails exist for what this tool is allowed to produce.
- ✓ The test: Could two different team members run the same task and get structurally similar output?

Training

- ✓ What to look for: Every user has been shown how to use the tool correctly before being given access to it.
- ✓ The test: Would a new hire know how to use this tool on day one without asking a colleague to explain it informally?

30 Day Check In

- ✓ What to look for: A scheduled review exists to confirm the tool is being used as intended and is producing real value.
- ✓ The test: Is that check in already on a calendar, or is it something that will only happen if a problem surfaces first?

Keep all this documentation in one place. Don't skip evaluating your AI Tools. Otherwise, you risk it becoming the approved workflow, guardrails or not.

The Mechanic's Tool Box

Ask and You Shall Receive

Most people decide what to paste into a prompt before they know what the task actually requires. Flip the order. Ask first.

The Prompt:

[ROLE] You are a Medical Affairs workflow analyst.

[TASK] I am about to ask you to complete the following task. Before I give you any source material, tell me what would be most useful to include to get the best possible answer.

[TASK DESCRIPTION] [Describe what you want done]

Please tell me:

One. What source materials would improve the accuracy of your answer.

Two. What source materials are optional and would only add detail rather than accuracy.

Three. Whether there is a limit to how much material I should provide at once.

This takes thirty seconds and it changes what gets pasted in next. It also catches the version of this task where you were about to hand over three documents when the model only needed one.

Want more ideas that have valuable ROI? [Contact Triple Helix Strategy](#)

The Radar

New JAMA Guidelines for Authors

JAMA published updated guidance on August 10 governing how authors may use AI in medical publication. The editorial did not soften JAMA's position on authorship. AI tools still cannot be listed as authors and still cannot take accountability for a manuscript. What has changed is more specific, and more relevant to how Medical Affairs teams actually work with publications.

The update tightens two things directly. AI can no longer be used to generate, format, or manage references. JAMA editors reported receiving manuscripts with inaccurate or nonexistent citations traced back to AI assisted reference generation. Even frontier models produced realistic looking references attached to real author names and real sounding titles that simply do not exist. Standard reference managers are the required alternative now. AI drafted opinion manuscripts, Letters to the Editor, and online comments are no longer permitted at all. JAMA cited a marked increase in formulaic submissions of exactly these article types, sometimes from authors without subject expertise in what they were submitting.

Disclosure requirements also became more specific. Any AI assisted manuscript preparation must be reported in the acknowledgment section with the tool name, version, manufacturer, date of use, a description of what it was used for, and confirmation the authors take responsibility for the content. Undisclosed use discovered during review or after publication can trigger rejection or a post publication correction.

For Medical Affairs, this lands directly on publication planning and medical writing workflows. Medical writers, MSLS contributing to manuscripts, and publication teams working with external authors are all touching JAMA family journals regularly. Reference generation and formatting is exactly the kind of task AI gets assigned informally, often without anyone thinking of it as a disclosure event. A manuscript submitted with AI generated references and no disclosure is now a compliance problem before it is a scientific one.

The acknowledgment section requirement also has an internal governance implication worth noting. If your organization cannot tell a medical writer exactly what needs to be disclosed and how, that gap will surface the first time a journal reviewer asks a question your team was not prepared to answer.

Need help building AI disclosure protocols into your publication planning workflow? [Contact Triple Helix Strategy.](#)

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

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. *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:

U.S. Food and Drug Administration. FDA Eliminates Major Barrier to Using Real-World Evidence in Drug and Device Application Reviews. News release. Published December 15, 2025. <https://www.fda.gov/news-events/press-announcements/fda-eliminates-major-barrier-using-real-world-evidence-drug-and-device-application-reviews>

IQVIA. FDA Updates Guidance on Real-World Evidence for Medical Devices. Published February 2026. <https://www.iqvia.com/blogs/2026/01/fda-updates-guidance-on-real-world-evidence-for-medical-devices>

International Council for Harmonisation. ICH M14: General Principles on Planning, Designing, Analyzing, and Reporting of Non-Interventional Studies That Utilize Real-World Data for Safety Assessment of Medicines. Adopted by FDA March 2026.

European Medicines Agency. Data Quality Framework for EU Medicines Regulation: Real-World Data Chapter. Adopted by CHMP, March 2026.

Berger J, Analytis PP, Satopää V, Kurvers RHJM. Beyond AI advice: independent aggregation boosts human-AI accuracy. arXiv. Preprint posted online April 3, 2026. <https://arxiv.org/pdf/2603.29866>

Matsui M. Is there a judgment here? Why medicine cannot endure decision-making without a judging subject in the age of AI. Med Health Care Philos. 2026. doi:10.1007/s11019-026-10346-2

Flanagin A, Perlis RH, Bibbins-Domingo K. Updated Guidance for Author Use of AI in Medical Publication. JAMA. Published online August 10, 2026. doi:10.1001/jama.2026.16613

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

