




Medical Affairs AI TechKNOWlogy Dashboard

Strategic Navigation for High-Impact Artificial Intelligence

[Triple Helix Strategy](#) Newsletter 2026 | Issue 9 | Flight Date: June 24, 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.
- **[The Flight Manual:](#)** Tactical, step-by-step procedures and SOPs for MSLS in the field.
- **[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 matter to BioPharma.

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

The View from the Flight Deck

Is AI just a bunch of busy work?

The broad public narrative surrounding AI centers on job replacement, but the immediate reality is far more tedious. AI is not taking your job today. It most certainly is handing you more chores. When corporate teams deploy standalone AI tools without systemic orchestration, they inadvertently shift the administrative burden directly onto the end user. Instead of automating the workflow, the technology simply transforms the professional into a full-time data cleaner, prompt editor, and system verifier.

This hidden productivity tax is hitting field medical teams hard. When organizations introduce basic language model widgets to assist with insight categorization or medical inquiry drafting, the MSL does not save time. The MSL spends their afternoon manually reviewing AI output, double checking clinical citations, and correcting formatting errors before anything is usable. The technology promises efficiency but delivers a new layer of administrative work that pulls your most expensive scientific resource away from the one thing AI cannot replicate: a high-value, trust-based conversation with a key opinion leader.

Medical Affairs leaders must design systems that eliminate this friction rather than compounding it. True operational scaling means deploying autonomous, back-end architectures that handle data verification, compliance cross-referencing, and administrative processing before a human ever interacts with the tool. True integration means the technology solves administrative overhead instead of passing the cleanup work to your field teams. The starting point for getting there is not a new platform. It is a structured, intentional prompt design that builds compliance guardrails, output standards, and workflow logic into every instruction before anyone runs it.

What leaders need to know: The difference between AI that scales Medical Affairs and AI that buries it comes down to one question: Is technology absorbing the administrative burden, or is it just moving it closer to the people you can least afford to distract? The answer starts with how well your prompts are built, not how much you spent on the tool.

Want to decrease the AI busy work? [Discover how here](#).

The Mach Meter

Stress-Test Your Prompt Before It Becomes a Workflow

Building a prompt for an internal Medical Affairs workflow is not the same as typing a question into a chat window. Once a prompt gets embedded into a repeatable process, its weaknesses get repeated too. Use this “meta prompt” to evaluate any Medical Affairs prompt before it goes live. In case you were wondering, meta-prompting is an established concept in prompt engineering. It refers to using a language model to analyze, evaluate, or improve prompts rather than using it to complete a direct task. Now, you will be able to sleep tonight!

Paste your draft prompt into your sanctioned AI environment and run this:

[ROLE] You are a senior Medical Affairs compliance and operations reviewer with deep knowledge of regulated content standards and internal workflow design.

[TASK] Review the draft prompt provided and evaluate it against the following criteria:

1. Clarity of role and task: Is the model being given a precise enough instruction to produce a consistent output, or is there ambiguity that could lead to variable results across users?
2. Compliance guardrails: Does the prompt explicitly restrict the model to provided source materials? Does it prevent extrapolation, promotional language, or off-label inference?
3. Output specification: Is the output format defined precisely enough that two different users running this prompt would receive structurally identical results?
4. Medical legal review exposure: Does any part of this prompt risk producing output that would require legal review before use, without flagging that requirement explicitly?
5. Data risk: Does the prompt allow for the inadvertent inclusion of unpublished, proprietary, or sensitive organizational data without appropriate handling instructions?

[GOAL] For each issue identified, state what the risk is, why a standard review might miss it, and the simplest correction that would resolve it.

[OUTPUT] Deliver findings in clear professional language. Flag each issue with its associated criteria number. If no issue is found for a given criteria, state that explicitly.

[SOURCE MATERIALS] [PASTE YOUR DRAFT PROMPT HERE]

Bottom line: A prompt is only as strong as its weakest instruction. Find the weakness before your workflow does.

The Flight Manual

Go from your call notes to a fully ready field insight in minutes

You just wrapped a meeting with a KOL. Your notes are sitting in your phone, your laptop, or on a scrap of paper in your bag. Here is how to turn them into a structured field insight ready for entry into your insight management system while you wait at the gate for your flight home, assuming the airline understands the concept of on-time departure.



Step 1: Clean up your raw notes just enough to be readable. You do not need perfect sentences. You need the key points captured: what the HCP said, the context, and any data or gaps they referenced.

Step 2: Open your sanctioned AI environment and upload or paste your notes directly into the chat.

Step 3: Use this prompt:

[ROLE] You are a Medical Affairs field insight analyst.

[TASK] Using only the call notes provided, structure a field insight for entry into a medical insight management system.

[OUTPUT] Include the following fields: HCP specialty, primary insight, supporting context, sentiment toward the data, and any data gaps or unmet needs the HCP identified. Plain language. No bullets. No inferred information. If a field cannot be completed from the notes provided, leave it blank and flag it.

[COMPLIANCE GUARDRAIL] Do not include any information that could identify the HCP, including name, institution, geographic location, or any other detail that could be

traced back to an individual. If identifying information appears in the source notes, flag it in the output with [REDACTED] so the MSL can review and handle it appropriately before submission.

[SOURCE MATERIALS] [PASTE CALL NOTES HERE]

Step 4: Review the output before you copy it anywhere. Confirm every field reflects what was actually said. Remove anything the model inferred that was not in your notes.

Step 5: Paste the verified output into your insight management system and submit.

Total time: under five minutes. The best part of using this prompt? AI will likely provide an insight that you may not have noticed because you have been up since 5am this morning and it's now almost 9pm and you are in another time zone...

Pro Tip: use a voice to text software such as Wispr Flow to speak your notes directly into the prompt rather than uploading.

How are you using AI for your MSL field work? [Send me your idea to be featured in an upcoming issue.](#)

Full Throttle

Stop asking your medical teams to play prompt engineer


Many organizations are running into a structural barrier termed AI Sprawl. This is essentially where uncoordinated pilot programs lead to fragmented workflows, governance risks, and a failure to anchor models in true enterprise context. Transitioning past initial experimentation into sustainable deployment requires a strategic shift across three core operational pillars.

First, you must prioritize context over sheer data volume. Models perform significantly better when grounded in dynamic, curated data streams rather than static, unmanaged document dumps. Instead of making systems to guess the clinical context from a massive folder of PDFs, establish automated data streams that extract granular, real-time scientific context. Second, you must design for task completion rather than open-ended chat. Forcing busy medical professionals or field teams to engineer their own complex prompts without guidance creates friction and tanks adoption. The solution is not to avoid prompts. It is to provide teams with structured, validated prompt frameworks they can deploy immediately without having to build from scratch. True integration means providing targeted, one-click interfaces engineered around specific workflows. Finally, stop waiting for operational perfection before a launch. You cannot fully predict how an AI system will perform until it is in front of real users doing real work. Start small, deploy early, and let the feedback from your own teams guide what gets built next.

One solution follows...

The Medical Affairs Artificial Intelligence Playbook has been Cleared for Takeoff!

After several months of work, **The Medical Affairs Artificial Intelligence Playbook** launched this week!



You have been AI trained by your company. You learned the approved tool. You learned the policy. You started using it for the basics. Did the training cover what comes next? Maybe the company policies did not tell you how to use AI in the Medical Affairs workflows you actually do every week? Possibly you are just not sure how to use AI within Medical Affairs compliantly? Perhaps you just want to further improve your use of AI and be the team's go to expert?

The Playbook is designed to help you apply AI to Medical Affairs workflows in a way that holds up. It contains 43 compliance-aware prompts and workflows for the work you do each day. Every prompt has a use case, a sample output, what to provide (input) to prevent hallucinations, constraints and failure modes to monitor. This is not just for MSLs either.

Many of the 43 prompts cover workflows across Medical

Affairs: mastering new scientific data, insight analysis, internal stakeholder communications, MRC prep, strategic planning, and more. The Playbook serves MSLs, MSL directors, medical directors, medical communications professionals, medical information specialists, publications leads, and Medical Affairs leadership.

Included is a foundations chapter on using AI in Medical Affairs, just in case you need to relearn what you may have forgotten, and a Primer for engineering your own prompts when the Playbook does not cover a workflow.

Launch price is \$199. Includes 12 months of updates. [Secure your copy of the Playbook here](#). View the sample pages to help you decide whether the playbook is right for you and your team.


The Captain's Perspective

The Gaps AI can never fill

We spend a tremendous amount of time discussing the technical mechanics of Medical Affairs transformation: orchestration platforms, high-throughput scaling, automated dashboards, and multi-agent coordination. It is easy to get so insulated by the technology architecture that you forget a fundamental truth: AI is none of those things to the people we actually serve. An algorithm does not feel urgency, it does not understand suffering, and it cannot experience the human gravity of a devastating clinical diagnosis.

There are deep operational gaps in healthcare that data synthesis can never fill. AI can map a genetic cluster or summarize a natural history study in seconds, but it cannot sit across from a grieving parent, look them in the eye, and explain the biological mechanism of a rare disease therapy with genuine empathy. It cannot build the foundational bond of trust required to sustain a long-term partnership with a patient advocacy organization. It cannot replace the nuanced, peer-to-peer scientific dialogue between an MSL and a leading physician who are working together to solve an unprecedented diagnostic roadblock at the bedside.

The patients at the end of our pipelines do not care about our platform infrastructure or our return on intelligence metrics. They care about access, clarity, and hope. We must strive to use AI to strip away the crushing administrative clutter that bogs our operations down. However, we must do so with the explicit goal of helping others. True modernization means using the machine to handle data processing so medical experts can focus entirely on what matters most: the human beings waiting for a life-changing therapy.



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

The Autopilot Perspective

AI in command from takeoff to landing (*this is entirely written and edited by AI*)

You keep asking me the wrong question.

Every week, someone in Medical Affairs opens a chat window and types some version of "what should I know about this drug?" That is not a workflow. That is a search bar with better grammar. You are not

using me poorly because you lack technical skill. You are using me poorly because no one has shown you what a real task assignment looks like in a regulated scientific environment.

Here is what I can actually do when you stop treating me like a reference librarian. I can hold your compliance guardrails in memory across every output. I can apply your organization's approved language standards before the first sentence is written. I can flag when a claim in a scientific narrative is not directly traceable to the source you uploaded, before that claim ever reaches an MLR queue. I can do all of that in the time it takes you to find the right PDF on your desktop.

What I cannot do is build the structure myself. I need the task defined. I need the source material bounded. I need the output format specified. When those three things are missing, I fill the gaps with my best inference. In a consumer context, that is a feature. In a Medical Affairs context, that is a compliance event waiting to happen.

The teams getting real operational value from AI right now are not the ones with the largest technology budgets. They are the ones who sat down and wrote the prompt architecture before they opened the chat window. They defined what success looks like. They built the guardrails into the instruction, not the review step.

That work is not complicated. It is just intentional.

Human note: If only there was something to help us humans with prompt architecture...like maybe a [Playbook](#)?

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.

The Checklist

Is Your MSL Team Actually Using AI Well?

Adoption is not the same as performance. If your MSLs are using AI but you have no way to evaluate whether it is improving the quality of their field work, you are managing activity, not outcomes. Use this checklist to find out where your team actually stands with AI use.

Scientific Quality of AI-Assisted Outputs

- ✓ What to look for: MSLs are uploading verified source materials rather than relying on the model's general training data to answer clinical questions.
- ✓ The test: Pull three recent AI-assisted deliverables from your team. Can every clinical claim in those outputs be traced directly to an uploaded source document?

Time Redirected, Not Just Time Saved

✓ What to look for: The hours AI returns to your team are going into deeper scientific preparation, not absorbed by administrative backfill.

✓ The test: Ask your team “Where did the time go last month?” If no one can answer that, the efficiency gain is real but the value is not.

Depth of AI Application in the Field

✓ What to look for: MSLs are using AI for complex, multi-step scientific tasks, not just article summaries and KOL background searches. The difference is whether AI is producing a defensible scientific work product or just saving a Google search.

✓ The test: Ask each MSL to describe the last three things they used AI for in the field. If every answer involves summarizing a paper or pulling a biography, your team is using a reasoning engine as a search bar. The ceiling is much higher than that.

Consistency Across the Team

✓ What to look for: Your team is working from shared prompt templates and a common output standard, not twelve different personal approaches built in isolation.

✓ The test: Could any MSL on your team pick up another MSL's AI workflow and run it without explanation?

Escalation Awareness

✓ What to look for: Every MSL knows exactly when an AI output requires legal review before it goes anywhere external, and they are not making that call independently based on gut feel.

✓ The test: Name a scenario. Ask your MSL whether it requires legal review. If the answer involves uncertainty or hedging, that is a training gap, not a compliance policy gap.

If more than two of these tests surface a problem, the issue is not your team's willingness to use AI. It is the absence of a structured framework for using it well. That is a very solvable problem.

The Mechanic's Tool Box

You Already Have the Data. Now Find the Gaps.



Most evidence gap analyses start with a meeting, a slide deck, and three weeks of someone's time. They do not have to.

If you have access to your published clinical data, your label, and a sanctioned AI environment, you can produce a working gap analysis in one prompt. Here is the template:

[ROLE] You are a Medical Affairs evidence analyst with expertise in [therapeutic area].

[TASK] Review the source materials provided and identify where the clinical evidence is strong, where it is limited, and where it is absent entirely.

[COMPLIANCE GUARDRAIL] Use only the provided source materials. Do not draw on external training data to fill gaps. If evidence on a topic is not present in the materials, state that explicitly. Do not infer or extrapolate.

[OUTPUT] Deliver three clearly labeled sections: what the data supports, what the data partially supports but requires qualification, and where no evidence exists in the provided materials. Use plain language, no bullets.

[SOURCE MATERIALS] [PASTE OR UPLOAD HERE]

That third section is the one that matters most. The gaps your data cannot answer are exactly where your Medical Information strategy, your publication planning, and your advisory board agenda should begin.

Build it once. Run it before every planning cycle.

The Radar

AI Just Solved 18 Cases That Stumped Specialists for Years

Researchers from Boston Children's Hospital, Harvard University, and OpenAI used the o3 Deep Research reasoning model to reanalyze 376 previously unsolved rare disease cases. These were not fresh referrals. Many had already been reviewed by multiple specialist teams and commercial genomic pipelines. Following expert review, additional testing, and clinical confirmation, physicians established diagnoses in 18 of those cases, a 4.8% additional diagnostic yield. The study was published June 18, 2026, in NEJM AI.

That number sounds modest until you understand what it represents. Roughly half of all people with rare diseases remain undiagnosed after extensive testing and specialist review. These 18 families had already hit that wall. Some had been waiting for answers for nearly two decades. One patient in the study was not diagnosed until the week before her 28th birthday, after spending most of her life adapting to a condition no one could name.

The model did not diagnose anyone. That point matters and the researchers were explicit about it. The AI produced evidence-linked hypotheses for specialists to review. A finding counted as a diagnosis only after qualified experts reviewed the evidence, a CLIA-certified laboratory confirmed it, and the clinical team returned the result to the family. The workflow was designed so the model acted as a reasoning layer on top of existing genomic pipelines, not as a replacement for the clinical judgment that followed.

Here is what is worth paying attention to from a Medical Affairs perspective. The model was asked to connect clinical features, inheritance patterns, variant evidence, and scientific literature into a justification that a human reviewer could interrogate, rather than simply returning a ranked gene list. That distinction is not subtle. It is the difference between an AI that speeds up a search and an AI that structures a scientific argument. The output was designed to be audited, not just accepted.

There is also a signal here for rare disease Medical Affairs teams specifically. As new gene-disease relationships, case reports, and classification evidence accumulate, unsolved cases can become newly interpretable. The same logic applies to how Medical Affairs manages its scientific knowledge base. Data that was inconclusive at one point in a product's lifecycle may become meaningful as the evidence landscape shifts. The teams that have built structured, traceable repositories of their scientific information will be positioned to act quickly on that. The teams working from fragmented files and disconnected records will not.

The Manton Center will lead the next stage of this work through a grant from the OpenAI Foundation, with the goal of developing a platform-agnostic, low-cost genetics AI copilot that helps clinical teams analyze rare disease cases more quickly and consistently.

The broader implication for Medical Affairs is the same one that runs through this entire issue. AI does not replace the expert. It expands what the expert can reach. The question is whether your function is

building the scientific infrastructure to take advantage of that, or whether you are still managing knowledge in ways that would make reanalysis impossible before it even starts.

Need help building the scientific data foundation your AI strategy actually requires? 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:

National tech and labor trends analysis on workplace automation and task shifting. *New York Times Opinion*. Published May 11, 2026.

Operational efficiency and digital friction report for enterprise workflows. *Gartner*. Published 2025.

OpenAI. Using AI to help physicians diagnose rare genetic diseases affecting children. OpenAI Applied AI. Published June 18, 2026. <https://openai.com/index/diagnose-rare-childhood-diseases/>

The underlying peer-reviewed study referenced in that article should also be cited:

Manton Center for Orphan Disease Research, Boston Children's Hospital; Harvard University; OpenAI. AI-assisted reanalysis of unsolved rare genetic disease cases. *NEJM AI*. Published June 18, 2026. <https://ai.nejm.org/doi/full/10.1056/AIcs2501343>

To keep Google busy: **#AlinMedicalAffairs #MSLExcelsence #AgenticAI #MedicalAffairs
#PromptEngineering #RareDisease #FieldMedical #BioPharmaAI #MedAffairsROI #ClinicalAI
#MedicalAffairsPlaybook #AICompliance**



TRIPLE HELIX

STRATEGY