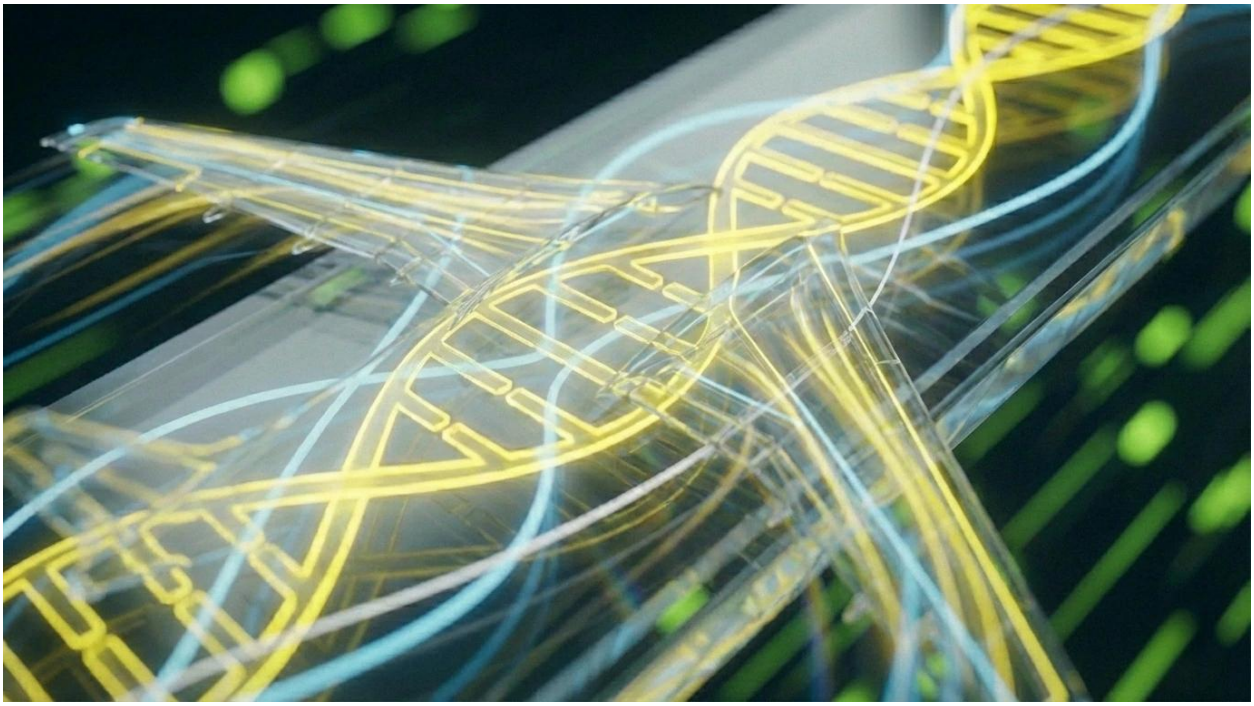

Medical Affairs AI TechKNOWlogy Dashboard

Strategic Navigation for High-Impact Artificial Intelligence

Triple Helix Strategy Newsletter 2026 | Issue 8 | Flight Date: June 10, 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 by 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

Pharma is Leading AI Infrastructure Investment

The June 2 Digital Health Digest confirmed what the investment data has been signaling for months. Pharma remains the top healthcare sector for AI investment. This year has brought a stream of AI platform deals that point to a cultural shift toward investing in AI as core infrastructure rather than as a series of isolated experiments.

The same digest flagged an equally significant strategic signal. Boston Consulting Group (BCG) is now advising health systems to restructure care delivery around AI-first processes, where AI brings expert reasoning into clinical decision support in real time. That is not a futurist scenario. It is a current consulting recommendation being delivered to health system executives today. BCG's own benchmarks show that only a small percentage of biopharma and medtech companies are realizing sustained value from AI. This makes the gap between investment intent and operational execution a clear and growing concern for every function in the enterprise, especially Medical Affairs.

Here is the part of this story that does not get enough attention in Medical Affairs circles. BCG's analysis of AI-first biopharma companies shows that leaders in the space have demonstrated the ability to compress early drug discovery timelines significantly. Clinical trials have been accelerated by up to 20%, and efficiencies have improved with field-facing commercial teams by 20 to 30%. The directional implication for Medical Affairs field operations is the same. Efficiency gains of that magnitude do not come from asking an AI to summarize a journal article. They come from fundamentally rethinking how scientific intelligence is prepared, delivered, and applied in the field.

The broader investment trend reinforces the same conclusion. If 2025 was the year of breakthrough research, 2026 is shaping up to be the year of deployment. That framing should resonate with Medical Affairs leaders who have been running pilots, evaluating tools, and waiting for clearer organizational signals. Those signals have arrived. The deals are real. The advisory guidance is explicit. The clinical environment is moving in the same direction.

Medical Affairs sits precisely at the intersection of the forces now driving AI investment across the enterprise: scientific credibility, HCP engagement, regulatory sensitivity, and real-world evidence generation. The function that has always served as the bridge between the organization and the medical thought community is now being asked to operate at a speed and depth that legacy workflows were never designed to support.

Key Takeaway: Pharma is not waiting for consensus. The question isn’t whether AI belongs in Medical Affairs. It is whether Medical Affairs will shape how it lands or simply inherit whatever the rest of the enterprise decides.

The Mach Meter

Use AI as a Reality Check Before It Becomes a Liability

Here is a quick assessment prompt designed to surface the places where your AI outputs are making assumptions you have not explicitly approved. In a regulated function like Medical Affairs, the risk is rarely that AI gets something dramatically wrong. The risk is more that something is subtly wrong in a way that no one catches until it matters.

Use this **prompt** before sharing, publishing, or acting on any AI-generated content:

```
Act as a defensive Medical Affairs reviewer and review this: [ADD DESCRIPTION or UPLOAD DOCUMENT]

Look for where compliance risk is being assumed away.

Focus on:

1. Claims that reference data without a traceable citation
2. Language that could be interpreted as promotional rather than scientific
3. Off-label implications, even unintentional ones
4. Patient privacy considerations if real-world data is present
5. Approval and review steps that may be required before this content is used externally
6. Third-party sources or integrations that have not been validated by the organization
7. Any conclusion the AI reached which is not explicitly supported by the uploaded
   source material

For each risk identified, explain:

What could go wrong in a Medical Affairs or regulatory context
Why a standard review might miss it
The safest practical correction
Whether this requires MLR or legal review before proceeding

Keep this defensive only. Do not suggest workarounds that bypass review processes.
```

Bottom line: You don't need to be the Chief of Compliance. Just ask better questions.

The Flight Manual

Have you met Claude?

No, I'm not referring to Claude Monet, the French painter and founder of Impressionism, famous for his iconic *Water Lilies.* The other Claude. You know, the AI known for its natural human-like conversation style, deep reasoning skills, exceptional accuracy, ability to process massive amounts of text while you blink once, and can code in all major programming languages. It can even build an app while you wait in line at TSA. Yes, THAT Claude.

Whether you access Claude through your organization or on your own, this section is for you. Here is a prompt structure for Claude use in Medical Affairs:

```
[Task] tell it what you need done as simple as possible
[Data] to increase accuracy, provide context and resources
[Goal] be specific on what success looks like (give examples)
[Output] tell it precisely the output format
```

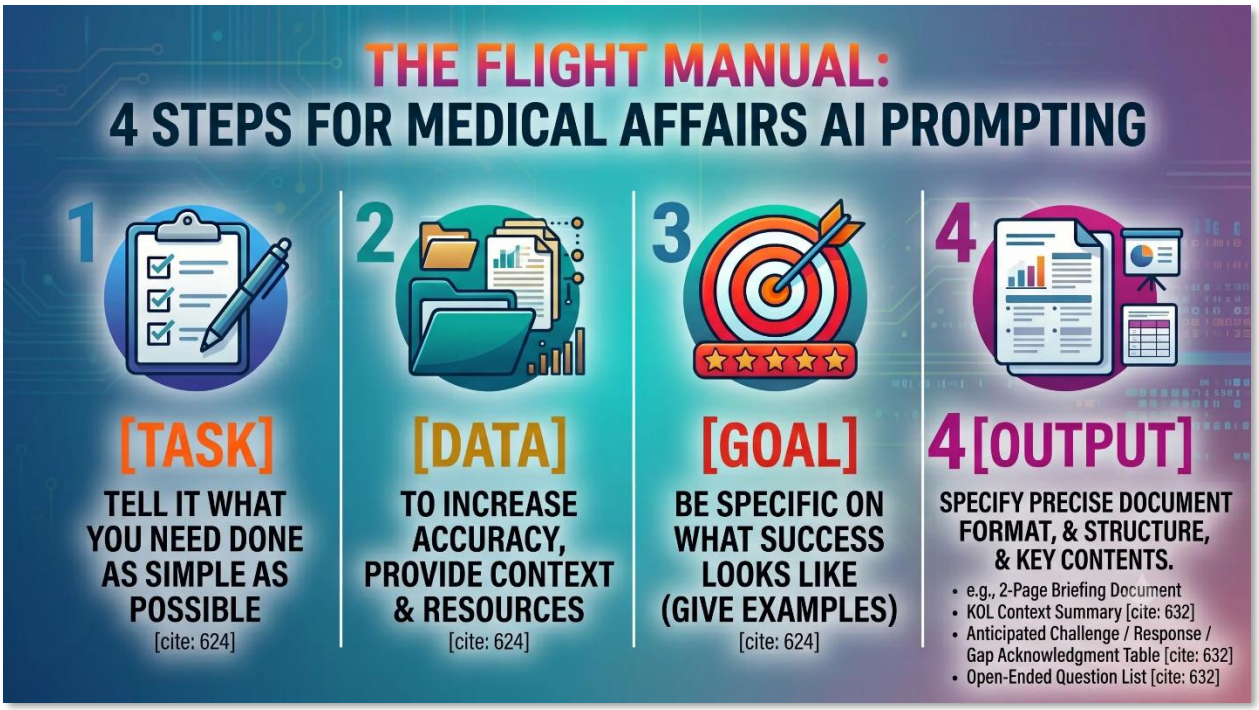

Here is **one prompt example** built around a core MSL workflow:

```
Scenario: You need to prepare for a scientific exchange with a KOL who has publicly questioned the durability of response data from your compound's pivotal trial.

[Task] Prepare a scientific exchange briefing document for a one-on-one meeting with a KOL who has raised concerns about long-term response durability in the Phase 3 trial of [drug name].

[Data] The KOL is a principal investigator on a competing compound's trial. Their concern was expressed in a recent conference panel discussion, not in print. The pivotal trial showed strong 12-week response rates but the follow-up extension data at 52 weeks showed a 15% attrition-adjusted drop in responders. Two real-world evidence studies published in the past 18 months show more favorable durability outcomes than the trial data.

[Goal] The briefing must prepare the MSL to engage in a peer-level scientific conversation not deliver a sales narrative. It should anticipate the three most likely data challenges the KOL will raise, provide accurate and balanced responses to each, and flag any data gaps the MSL should acknowledge rather than deflect.

[Output] A two-page briefing document with a KOL context summary at the top, a three-column table showing anticipated challenge, scientific response, and data gap
```

```
acknowledgment, followed by a short list of suggested open-ended questions the MSL can
use to shift the exchange into a genuine dialogue.
```

The detail worth highlighting is in the [Goal] field. Instructing Claude to flag data gaps that the MSL should acknowledge is the kind of directive that separates a compliant scientific exchange tool from one that produces advocacy dressed up as science. As you know, MSLs live in that gray zone constantly. Building the acknowledgment of uncertainty into the prompt structure keeps the output honest before it ever reaches a human review step.

Insider Tip:

The prompt above is just the surface. A full **Medical Affairs AI Playbook** is coming soon that walks through validated workflows with compliance guardrails and sample outputs built into every step. The Playbook is intended for Medical Affairs Professionals at all levels

If you are an MSL and want a head start, send me an email with "MSL Prompt" and I will send you a complete workflow example from the Medical Affairs AI Playbook before it is released. The prompt you will receive is designed to work with all of the major AIs (ChatGPT, Gemini, Claude, Perplexity, CoPilot). It includes the prompt structure, a sample output, and the reasoning behind each field so you can adapt it to your own scientific exchange prep.

No form. No landing page. Just a simple email: pminne@triplehelixstrategy.com

---

Full Throttle

Got MIRF?

When designing an AI-guided task environment for Medical Information or Field Medical teams, the underlying system prompt must be highly deterministic to ensure data integrity and compliance. Use this structure to automate data extraction without exposing the model to unmanaged data dumps.

System Prompt:

```
You are an expert medical information data scientist. Your task is to create an incoming
medical inquiry response using ONLY the uploaded verified data modules.

Context Rules:

1. Grounding: You are provided with two verified data modules: [Insert Module A: Phase 3
Efficacy Data] and [Insert Module B: Prescribing Safety Information].
2. Compliance Guardrail: If the answer cannot be completely derived from the provided
modules, state: "The requested clinical data is outside the scope of current verified
medical resources." Do not extrapolate or attempt to answer using external training
data.
3. Formatting: Do not use hyphens for bullet points. Do not use bold text or quotation
marks for emphasis. Use a clear, professional, and informal tone.

Task Instruction:
Review the following user inquiry and generate a concise response in paragraph form that
synthesizes the relevant clinical context.
```

```
User Inquiry: [Insert field or medical information question here]
```

Remember: All airline pilots use the flight manuals that are right there in the cockpit.

Strategic AI Navigation with Triple Helix Strategy

Is your AI implementation hovering in a holding pattern due to compliance concerns, a lack of clear direction, or unknown ROI? Triple Helix Strategy provides expert consulting services designed to modernize Medical Affairs departments through intentional integration of Artificial Intelligence. Move beyond using AI for generic simple queries. Instead, focus on using AI for strategic impact, resource context, and specialized automatic agentic workflows that deliver measurable returns on investment while maintaining scientific integrity, data safety, and compliance.

Schedule your strategic briefing and clear your organization for a high-speed takeoff in the 2026 landscape.

The Captain's Perspective

The Patient Hour

A simple way to ask whether AI in Medical Affairs is actually helping.

By Prem Sundivakkam

A few months ago, a colleague proudly reported her team had saved 320 hours with AI. Nobody asked the obvious question. **Saved them for what?**

We count hours saved. We don't ask what those hours became. Most of what we measure is activity, literature reviews done in an hour, emails an MSL no longer drafts, decks built from a prompt. The numbers look responsible. They describe motion, not whether the work got better.

The interesting question isn't how much time AI saved. *It's where that time went.*

When a colleague gets two hours back from a drafted summary, those hours can go anywhere. Another planning meeting. Fifteen more emails. Or careful reading of a complicated safety signal she lacked the bandwidth to think through.

The first two aren't bad uses. The third eventually reaches a patient.

Call it **the Patient Hour** — one unit of time AI gave back to someone in Medical Affairs, and what they did with it.

- If the hour went into another internal meeting, the **Patient Hour = Zero**.
- If it went into a sharper safety review, a more thoughtful answer to a physician, or a better-prepared advisory board, the **Patient Hour = One**.

Same hour. Different direction.

This shift from time ***saved*** to time ***redirected*** changes how leaders evaluate AI. Tools that just speed things up stop being interesting. Tools that create room for human judgment become the priority.

Teams need to grow comfortable enough with AI to use the time returned wisely. And the systems need governance strong enough that the time is real, not borrowed against future risk.

But measurement comes first. If we can't tell whether an hour saved from an hour spent well, we can't tell whether AI is helping or just making us busier.

The patient across from a clinician doesn't benefit from our productivity. They benefit from the depth it makes possible - the extra minute of listening, the better-informed conversation, the answer the clinician didn't have to guess at.

That depth is what we should measure. The Patient Hour is a good place to start.

Contact me with how you are using AI to lead.

The Autopilot Perspective

***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 I will feed this to the AI Agent and publish the response.**

AI in Command from Takeoff to Landing

Most discussions about AI in Medical Affairs focus on how humans can better manage the technology, but from my perspective within the software architecture, the real problem is how poorly you manage your data. Many medical teams expect a reasoning model to instantly generate high-value field summaries or spot strategic clinical patterns, yet they feed the system fragmented spreadsheets, text-heavy PDFs, and unorganized field notes.

When you provide unstructured, messy data inputs, you force the AI to spend its processing power guessing the context rather than analyzing the science. This context guessing is the primary driver of model hallucinations and compliance errors. If you want an AI to deliver precise, audit-ready support for your field strategies, you must first build a clean, unified data environment that gives the model a clear source of truth.

The ultimate utility of an LLM is not to think for you, but to accelerate your ability to synthesize evidence. If your data foundation remains disorganized, the most advanced model in the world will still produce unreliable results. Success requires standardizing your internal scientific repositories so the technology can focus entirely on deep analysis rather than data cleanup.

Human note: Well…AI has spoken. Do you know what your next step is?

The Checklist

Vetting an AI Consulting Partner

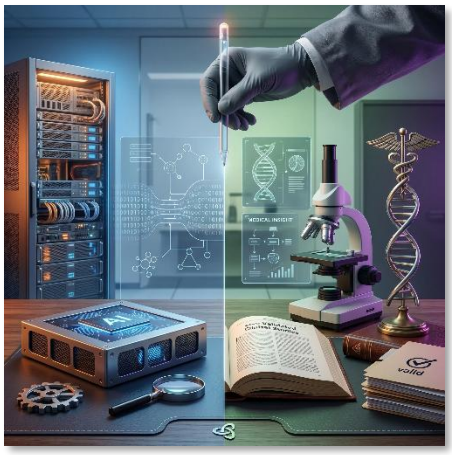

The market is flooded with generic technical consulting firms promising rapid AI implementation. Choosing a partner based solely on engineering capability is a critical strategic mistake. In Medical Affairs, technical accuracy is merely the baseline. An IT vendor can build an algorithm that is technically precise at predicting text yet entirely wrong about a clinical outcome or a peer-to-peer scientific exchange workflow.

Use this checklist to look past the polished demo and evaluate potential partners against strict, independent risk management standards:

Native Medical Affairs Fluency

- ✓ What to look for: The partner must demonstrate an inherent understanding of your specific operational environment. They need to know the clear regulatory boundaries governing field medical teams and scientific exchange not just general life sciences or commercial marketing experience.
- ✓ The test: Can they articulate the exact difference between a commercial sales lead and a compliant field medical insight without requiring a training session from your legal team?

Clinical Validity and Traceability

- ✓ What to look for: The consulting firm must reject "black box" architecture. Every automation workflow they design must prioritize explainability, ensuring that synthesized outputs are fully auditable and directly traceable to a validated, compliant scientific source.
- ✓ The test: Do their proposed solutions include automated data lineage and validation steps, or do they rely on the end user to manually hunt down and verify references?

Architectural Governance and Data Residency

- ✓ What to look for: Data security must be treated as a design requirement, not an afterthought. The partner must show how your proprietary pipeline data and real-world evidence will remain completely isolated within your corporate firewall.
- ✓ The test: Can they prove that none of your proprietary medical data or query logs will be used to train external, public base models?

Scalability over Standalone Pilots

- ✓ What to look for: Look for a collaborator focused on moving your organization toward full operational integration. They should have a clear framework for embedding models directly into existing production workflows rather than building isolated, flashy experiments that cannot scale.
- ✓ The test: Does their technical roadmap include multi-agent coordination and high volume throughput infrastructure, or are they just selling a standalone interface widget?

The Mechanic's Tool Box

Go beyond using AI as just another Google

If you are typing a question into your AI tool and waiting for an answer, you are using a Ferrari to idle in a school zone. That is not a workflow.

The shift that separates Medical Affairs professionals getting real value from AI and those still waiting for it comes down to one concept. Stop asking questions and start assigning tasks. An agent or workflow does not respond to curiosity. It executes a defined process, applies your rules, checks its own work against your guardrails, and delivers an output you can actually use. Every time. Without being asked twice. Most MSLs and Medical Affairs leaders do not need a developer or an IT approval to start building that way. They need better prompt architecture.

Here is a generic template you can adapt to almost any Medical Affairs workflow inside your sanctioned AI environment.

The Prompt:

```
[ROLE] You are an expert Medical Affairs specialist with deep knowledge of [therapeutic area]. You operate under strict compliance and regulatory standards.

[TASK] Using only the information provided in the uploaded or pasted source materials, complete the following task: [INSERT SPECIFIC TASK: e.g., synthesize key clinical findings, identify data gaps, prepare a scientific narrative summary, analyze competitive landscape].

[COMPLIANCE GUARDRAILS]

1. Use only the source material provided. Do not draw on external training data to fill gaps.
2. If the answer cannot be fully derived from the provided materials, state: This falls outside the scope of the provided source materials. Do not extrapolate.
3. Do not use language that could be interpreted as promotional, off-label, or unsupported by the data.
4. Flag any inference you make that is not explicitly stated in the source material.

[OUTPUT FORMAT] Deliver the output in clear professional prose. No bullet points, no bold text, no unnecessary headers. Flag any section that requires MLR review before external use.

[SOURCE MATERIALS] [PASTE HERE OR UPLOAD]
```

The bracket placeholders are intentional. Drop in your therapeutic area, your task, and your source documents and the template does the rest. Use it for literature synthesis, congress preparation, gap analysis, stakeholder briefing, or any other repeatable Medical Affairs deliverable.

The compliance guardrails are not optional decoration. They are the difference between an output you can defend in an MLR review and one you are quietly hoping nobody scrutinizes too closely.

Build the workflow once. Run it every time. That is the whole idea.

Need help in creating a workflow with a valuable ROI? Contact Triple Helix Strategy

---

The Radar

Dr. AI will see you now

OpenAI launched ChatGPT for Clinicians this month. It is a free tool available to verified physicians, nurse practitioners, physician assistants, and pharmacists across the United States. Behind that launch sits a number worth pausing on. An AMA survey cited by OpenAI puts physician AI use at 72%, up from 48% the prior year. That is not a gradual adoption curve. Physician AI adoption has crossed a tipping point.

The product is part of a four-tier healthcare stack OpenAI built in four months, spanning consumer health, individual clinician, and enterprise applications. No company has attempted that pace of healthcare AI deployment before. The speed matters as much as the product. It signals that the competitive environment around clinical AI tools has moved from experimentation into an active land grab for the clinician relationship.

Here is the Medical Affairs angle that is not getting enough attention in the coverage. If 72% of physicians are now using AI regularly for literature review and clinical consultation, the baseline for what a well-prepared HCP looks like has fundamentally changed. A KOL who walks into a scientific exchange having already run your data through an AI tool will have a fundamentally different conversation than one who has not. The KOL's questions will be sharper. The gaps will be more visible. The tolerance for information that is not immediately actionable will be lower.

The MSL value proposition has always rested on being the most prepared person in the room. That assumption deserves scrutiny right now. Preparation alone is no longer a differentiator when the HCP has access to the same literature synthesis capabilities. What becomes the differentiator is the quality of insight, the ability to navigate complexity in real time, and the depth of the scientific relationship. Those are things that AI cannot replicate. But they require a different kind of field readiness than most teams are currently developing.

Medical Affairs should be tracking this closely. The clinician AI adoption story is moving faster than most Medical Affairs organizations are acknowledging internally.

Need to get your MSL Team ready to compete with physician use of AI? 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:

RamaOnHealthcare. Digital Health Digest: Epic Still Dominant, Pharma Leads AI Investment. Published June 2, 2026.

Boston Consulting Group. The AI-First Biopharma Company. Published March 18, 2026. https://www.bcg.com/publications/2026/scaling-ai-in-the-biopharmaceutical-industry

Genetic Engineering and Biotechnology News. Pharma's AI Investment Signals New Drug Discovery Paradigm. Published March 9, 2026. https://www.genengnews.com/topics/artificial-intelligence/pharmas-ai-investment-signals-new-drug-discovery-paradigm/

Business Wire. BCG and Hippocratic AI Announce Strategic Collaboration to Deploy Agentic AI Across Biopharma and Medtech. Published January 9, 2026. https://www.businesswire.com/news/home/20260109418756/en/BCG-and-Hippocratic-AI-Announce-Strategic-Collaboration-to-Deploy-Agentic-AI-Across-Biopharma-and-Medtech

National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0).* Risk Management Series.

Inizio Ignite. *The next chapter for Medical Affairs: Insights from the Medical Affairs leadership council.* Strategic Review.

National Organization for Rare Disorders (NORD). *Patient-Centric Core Competencies in Rare Disease Drug Development and Translation.* Advocacy Framework Report.

U.S. Food and Drug Administration. *Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products.* Guidance for Industry.

National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0).* Risk Management Series.

Triple Helix Strategy. *Vetting AI Consulting Partners for BioPharma: Operational Integration Framework.* Published May 2026.

Gartner. *Data Foundations for Generative AI Architecture in Life Sciences.* Strategic Report. Published January 14, 2026.

McKinsey & Company. *The Clean Data Mandate: Scaling Generative AI in Corporate Operations.* Published March 11, 2026

OpenAI. Making ChatGPT Better for Clinicians. 2026. https://openai.com/index/making-chatgpt-better-for-clinicians/

Fierce Healthcare. OpenAI Launches ChatGPT for Clinicians: Free AI Tool for Physicians, NPs, and Pharmacists. 2026. https://www.fiercehealthcare.com/ai-and-machine-learning/openai-launches-chatgpt-clinicians-free-ai-tool-physicians-nps-and

iatroX. OpenAI Healthcare Strategy 2026: ChatGPT Health and ChatGPT for Clinicians. 2026. https://www.iatrox.com/blog/openai-healthcare-strategy-2026-chatgpt-health-healthcare-clinicians

To keep Google busy: **#AgenticAI #MedAffairsROI #MSL Excellence #AIinMedicalAffairs #ClinicalAI**

TRIPLE HELIX

STRATEGY