Human+AI and Pharma:

Defining the Balance Between Innovation and Oversight

DHCGROUP
NETWORK • INSIGHTS • EVENTS

A DHC Group Point of View

Based on Executive Interviews and Industry Roundtable

December 2025

Executive Summary

The pharmaceutical industry allocates significant resources to AI—**44% of life sciences organizations now dedicate 20% or more of commercial budgets to AI initiatives**. Yet only 11% report having AI-ready data systems, and scalability remains the critical bottleneck separating ambition from execution (Source: IQVIA).

Through DHC Group interviews with pharmaceutical and technology executives, plus a DHCG roundtable featuring leaders from IQVIA, Eversana, and Healthline Media, a clear framework emerges: successful AI implementation requires a 'human-at-the-core' approach—where human judgment establishes strategic parameters and AI operates within those boundaries.

Three critical variables determine appropriate oversight: product lifecycle stage (inline vs. launch), audience type (HCP vs. consumer), and therapeutic area severity (oncology vs. dermatology). Organizations that master risk-tiered governance capture AI's efficiency gains without compromising patient safety or regulatory compliance. The competitive advantage belongs to those who view oversight not as friction but as an enabler—building sophisticated governance frameworks that allow AI to scale responsibly.

About the Research

AI has moved from experimental to operational necessity in pharmaceutical sales and marketing, creating a paradox: the technology that promises unprecedented efficiency also introduces new risks where mistakes carry severe consequences—from regulatory sanctions to patient harm.

This Point of View synthesizes insights from in-depth interviews with industry leaders spanning pharmaceutical companies, technology providers, and publishers, combined with a DHC Group roundtable discussion featuring:

Andrew Burkus

Senior Director
IQVIA Digital

Abid Rahman

SVP, Innovation
EVERSANA

Josh Moffett

VP of Data Healthline Media

Mark Bard

Co-Founder
The DHC Group

The Current State: Momentum Meets Infrastructure Reality

IQVIA's AI in Lifesciences survey of over 100 life sciences executives reveals significant momentum—44% allocate 20% or more of commercial budgets to AI, and over 80% have advanced beyond pilot phases. Andrew Burkus noted that one-third now describe themselves as AI-advanced, having streamlined processes and created feedback loops for continuous optimization.

However, infrastructure gaps constrain progress. Only 11% report AI-ready data systems, with three primary barriers:

1

Data Quality Issues

Siloed data and legacy systems inhibit personalization at scale

2

Regulatory Constraints

Cited by 51% as the top concern, creating strategic and operational hurdles

3

Investment Misalignment

Sales and marketing prioritized, but disconnects remain between investment and outcomes

Four key application areas emerged: procedural automation, patient experience enhancement, supply chain resilience, and intelligent targeting. Yet as Andrew Burkus emphasized, the regulatory environment creates urgency:

> **"The administration saying that they're going to double down on enforcement, regulation and oversight in these areas really only further validates how important (integrating AI into our workflows) is for 2026 and beyond." —Andrew Burkus, IQVIA**

From Human-in-the-Loop to Human-at-the-Core

The evolution from 'human-in-the-loop' to 'human-at-the-core' reflects a fundamental shift. Rather than humans checking AI outputs after generation, humans establish strategic parameters, define brand voice, and set compliance guardrails within which AI operates autonomously. During the interviews, one pharmaceutical executive captured this: "We're not setting AI loose. We have a foundation of 40-50 years of messaging. AI helps us on the fringes."

Abid Rahman articulated Eversana's 80/20 framework:

> **"We want to automate things that should be automated, which is about 80% of the work. The other part, there is human-in-the-loop in various different ways. The level of involvement varies based on the type of content we're dealing with and the type of task we're executing."**
> —Abid Rahman, Eversana

Automated Work

80% - Tasks that should be automated

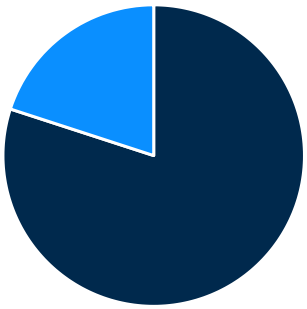

Human Contribution

20% - Verification, brand voice, compliance, empathy

Critically, the 20% human contribution focuses on qualitatively different work—verification, brand voice alignment, regulatory compliance, and empathy validation rather than content creation from scratch.

Risk-Based Tiering: The Core Framework

Effective governance requires categorizing use cases by risk and applying proportionate oversight based on three variables:

Therapeutic Area Severity

Multiple executives noted dramatic oversight differences based on disease state. A pharmaceutical executive explained: "We have one process for dermatology versus oncology. When someone says 'I think I may have prostate cancer,' we're extremely careful about content and empathy." Oncology content undergoes specialized medical review while routine dermatology queries proceed with algorithmic oversight.

Audience Type

Companies exhibit a significantly higher risk tolerance for HCP than for consumer content. Physicians can critically evaluate information; consumers demand heightened scrutiny due to the potential for misunderstanding and legal exposure. Andrew Burkus highlighted that external communications now require robust human-in-the-loop oversight, ensuring that AI-driven messaging aligns with strategic objectives while maintaining accuracy, relevance, and compliance amid growing regulatory scrutiny and heightened visibility.

Product Lifecycle Stage

Established products benefit from years of approved content serving as training data. Launch brands maintain heavy human involvement. Interestingly, mature brands with minimal budgets represent opportunities for aggressive AI experimentation—limited downside justifies higher risk tolerance for efficiency gains.

These variables create three possible oversight tiers: Low-risk (periodic review), Medium-risk (human validation required), and High-risk (human leadership with AI support). Pharmaceutical organizations should map every AI initiative to these tiers architecturally rather than retrofitting later.

Critical Intervention Points

Understanding where human intervention delivers maximum value separates effective from performative governance. For content generation, automated checkpoints escalate for human review when adverse events, off-label use, or inappropriate tone appears.

MLR review's primary bottleneck isn't the review itself—it's submission quality. Abid Rahman described how AI addresses this:

> "We have created AI-based guidelines. For promotional content, we have an AI solution that generates evidence for why things are being said—not just evidence, but also where it comes from, in annotations. From a human standpoint, it becomes very simple. I don't have to spend hours double-checking." **—Abid Rahman, Eversana**

For agentic systems, pharmaceutical companies face intense constraints around medical advice. Josh Moffett outlined a de-risking strategy:

> **"If it's high risk, can you pull forward one step in the health journey and do something slightly lower-risk? *'Here are providers suited for those treatment plans.' 'Let me connect you to that provider.'* How do you pull forward to de-risk what you're offering?"** —Josh Moffett, Healthline Media

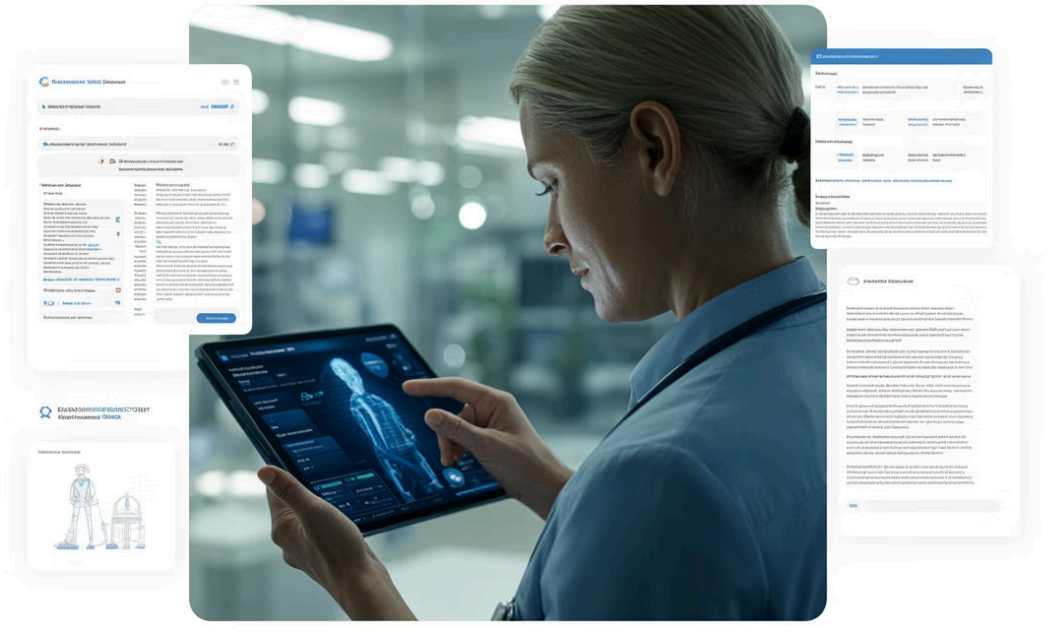

Empathy Engineering and Scalable Oversight

Beyond Prompts to Systematic Design

Empathy requires more than instructing AI to "be empathetic." Sophisticated systems analyze input language to infer appropriate response tone. Josh Moffett described Healthline Media's approach:

> "Someone asking a very scientific question deserves a response that matches. Or someone saying, 'I cannot get my three-year-old to sleep. Please help'—that's deeply personal. If we can infer that, we can ensure the response demonstrates more empathy, matching that one-to-one engagement level." —**Josh Moffett, Healthline Media**

One consumer health AI executive emphasized continuous human review: "We have human reviewers constantly checking—are we too clinical? Does it lack emotional nuance? Is it missing context for someone in that state of mind?"

AI-Validates-AI Systems

Scalable oversight involves deploying AI to validate AI—using one model to check another's outputs before human review. An AI platform executive explained: "The highest ROI comes from disagreement rates. When AI judges disagree with each other or humans, that's where human expertise becomes most valuable. You can't review thousands of daily predictions, but you can focus on disagreements."

01

Primary AI generates content

02

Validation AI evaluates outputs

03

Human experts review disagreements and conduct spot audits

Andrew Burkus also advocated simulation environments for sandbox testing before production deployment, giving domain experts opportunities to identify failure modes safely and in controlled environments.

Josh Moffett articulated a counterintuitive finding: "There's always this trade-off—if you bring humans in the loop, do you slow yourself down? What we actually found is that AI evaluation frameworks are a competitive advantage, because by capturing that feedback loop early and often, you develop faster on the back end and bring higher quality products to market."

Redefining Roles and Data Governance

The Age of AI Conductors

AI transforms marketing roles rather than eliminating them. Marketers shift from creating campaign assets to directing AI to generate variants, reviewing outputs for brand alignment, and approving final selections. Andrew Burkus used a musical metaphor: "A successful marketing professional has to become that AI conductor, adding domain expertise to the skill sets required."

Critically, AI conductors must be domain experts. Josh Moffett emphasized:

> **"Humans in the loop need to be domain experts. That's where you get an understanding of whether an AI output is good, bad, right, or wrong. That's where you create a moat different from previous technologies—you need domain expertise to really excel here."**
> —Josh Moffett, Healthline Media

Cultural transformation proves essential. Josh Moffett described Healthline Media's milestone: "We got to a point where instead of our engineering teams pushing AI applications on business teams, we have the opposite—they bring them to us. Creating that has been a huge benefit."

Data Governance and Compliance

When AI-generated content violates regulations, who bears responsibility? Andrew Burkus addressed this directly: "Ownership and liability ultimately start at the brand level, at the pharma level..... That said, tech platforms and partners need to be as transparent as possible with audit trails and robust documentation."

Josh Moffett added: "Ultimately, the entity triggering that AI output bears primary compliance liability. Partner with legal and compliance from the start. Don't take the approach of 'there'll be someone to point a finger at.'"

Abid Rahman outlined Eversana's foundational principle: "We decided we are not going to train AI models on client data. Being able to put that in place automatically takes away a lot of compliance challenges." This requires investing in alternative training data but provides clients assurance that their proprietary information won't inform competitors' applications.

Key Takeaways and Actionable Steps

Strategic Imperatives

1. **Embrace Human-at-the-Core:** Humans define strategy and compliance parameters; AI executes within boundaries.
2. **Implement Risk-Based Tiering:** Map initiatives to oversight tiers based on therapeutic area, audience, and lifecycle stage.
3. **Start with Inline Brands:** Leverage approved content libraries for initial experimentation on lower-risk elements.
4. **Respect Practice-of-Medicine Boundaries:** For consumer AI, facilitate connections rather than making recommendations.
5. **Invest in Empathy Engineering:** Build systematic input analysis and continuous human review for emotional appropriateness.

Operational Best Practices

1. **Build AI-Validates-AI Systems:** Deploy validation AI before human review; focus experts on disagreements.
2. **Establish Feedback Loops:** Systematically collect HCP and patient feedback; use disagreements as training signals.
3. **Develop Expert Certification:** Maintain reviewer expertise through ongoing certification as AI capabilities improve.
4. **Utilize Simulation Environments:** Test AI systems with edge cases in safe sandboxes before production.

Cultural Transformation

1. **Redefine Roles as AI Conductors:** Train marketers to orchestrate AI rather than manually create content.
2. **Build AI-First Culture:** Embed governance understanding; measure success by business teams requesting AI.
3. **Balance Platform Leverage with Control:** Use centralized platforms for infrastructure while building proprietary governance.

Conclusion

The pharmaceutical industry's cautious AI approach positions organizations for sustainable advantage rather than risky short-term gains. The shift from human-in-the-loop to human-at-the-core represents strategic maturity—recognizing oversight as an essential architecture that enables AI to operate effectively at scale.

Andrew Burkus captured this perspective: "Yes, this impedes speed and affects agility. But compliance can be a differentiator if clients trust you because you show the work, show the compliance. That trust in partners remains paramount."

1. **Risk-Proportionate Design:** Apply oversight based on actual risk rather than uniform processes.

2. **Domain Expert Elevation:** As AI handles execution, expertise becomes the differentiating contribution.

3. **Cultural Priority:** Technology alone doesn't create AI-advanced organizations—culture does.

Organizations with strong governance foundations can adapt as capabilities expand and regulations evolve—confident that their systems will catch problems before patient harm occurs. Those who pursued aggressive automation without oversight will face costly retrofitting when regulations tighten.

Action Steps

Immediate (0-3 Months)

- Conduct an AI oversight audit—map initiatives to risk tiers, identify gaps
- Establish a cross-functional governance team with clear approval hierarchies
- Start an inline brand pilot with low-risk elements while documenting learnings
- Partner with legal on data usage rights and AI-specific contract terms

Near-Term (3-6 Months)

- Implement an AI-validates-AI framework with human review on disagreements
- Build a simulation environment for testing edge cases before production
- Develop empathy engineering with input analysis and continuous review
- Launch AI conductor training, blending technical literacy with domain expertise

Medium-Term (6-12 Months)

- Scale risk-tiered oversight across therapeutic areas with customization
- Establish an expert certification program with ongoing validation requirements
- Build automated monitoring that alerts on anomalies and classifies severity
- Implement structured feedback loops, incorporating learnings systematically

Long-Term (12+ Months)

- Cultivate an AI-first culture where business teams drive AI application requests
- Deploy consumer-facing agentic AI with practice-of-medicine safeguards
- Optimize MLR integration using AI for upfront quality checks
- Evaluate platform strategy, balancing proprietary development with centralized infrastructure

Success requires organizations to simultaneously accelerate and restrain—moving fast while maintaining rigorous oversight. Pharmaceutical companies that master this balance will define industry best practices, capturing AI's transformative benefits without compromising healthcare trust.

About The DHC Group

The DHC Group is known for cutting edge research, expert strategy, and analytics-driven insights.

Building on the leadership team's combined 50+ years of industry experience working with innovative companies, brands, and a wide range of pharmaceutical, biotech, and medical device companies, the DHC Group provides industry thought leaders and innovative organizations a selective opportunity to define the future of digital health and pharmaceutical marketing.

For more information:
Christine Franklin, Executive Director, cfranklin@thedhcgroup.com

Project Partners

EVERSANA INTOUCH

IQVIA DIGITAL

HEALTHLINE MEDIA

More from The DHC Group

PHARMA FORWARD: Join us April 16, 2026 for a Live Summit in NYC, hosted by DHCG x Google

HCP + AI - Latest Research from The DHC Group