

Responsible AI in healthcare analytics: The MedInsight approach

Jake Linenthal, Chas Busenburg, Robert Richards, Jake Stone,
and Iyibo Jack

Milliman MedInsight

Introduction

The recent explosion of artificial intelligence (AI) capability and deployment at scale has dramatically altered conversations and expectations for technology solutions. Healthcare and healthcare analytics are no exceptions to this. AI as a technological innovation presents significant opportunities for greater efficiency and sophistication of healthcare analytics tools. To fully harness the power of AI in delivering reliable analysis and actionable insights to MedInsight clients, we feel it is important to describe the pillars of “good analysis” and the principles guiding our application of AI to healthcare analytics.

Analysis in the AI era

AI represents an economic shift at least as much as a technological one. Tasks like query drafting, content generation, pattern recognition, and translation between technical and non-technical representations are now accomplished much more quickly. We primarily view AI as a force multiplier that accelerates value delivery, enables more customization at a given cost, and expands who can interact with data to derive insights from it. While these economic shifts are fundamentally transforming our industry and many others, it is also critically important to lay out what has not changed.

First, our clients and all users of analytics technology, AI-enabled or otherwise, still require results that are sound, useful, explainable, and trustworthy. The standards that have always governed healthcare analytics remain unchanged. Crucially, this means the force multiplier of AI is only valuable to the extent that what is being multiplied is sound.

Second, the process of doing good analysis has not changed. In our experience, “good analysis” rests on four pillars:

1. A clear definition of the problem and the decision context
2. Data that adequately represents the domain being analyzed
3. Methodologies and validation strategies that are logically sound and suitable for the analytical question
4. Effective delivery of results that support decision-making and action

Indeed, clarity and precision in defining clear questions, identifying suitable data and methodologies, and presenting results with a useful and defensible interpretation are more imperative with AI than ever before. Users of AI-powered analytics tools may not be trained analysts with deep understanding of any of the four pillars. AI excels at navigating the pathways between these pillars, accelerating processes within each, and driving towards an answer, but without significant direction from humans it will fail to uphold these pillars of good analysis.

For example, consider the data pillar. MedInsight has long recognized the need for robust data validation, especially in the complex world of healthcare. Our Data Confidence Model, which has been refined over our almost three decades of experience in healthcare analytics, is a testament to this commitment. AI responses can only be as good as the data provided to them. AI improves nothing if data confidence is low. AI could silently incorporate low-quality data into final responses with no warning or signs of danger while seeming to satisfy the requests made by the user or system.

Similar consideration ought to be given to each of the four pillars. As we integrate AI into our products and platforms, we adhere to the principles laid out below.

Principles of AI adoption at MedInsight

Existing expertise, validated methodologies, and institutional knowledge are the inputs that AI amplifies, whether those qualities are well-developed or not. The value of getting those right is higher now, not lower. The platform, methodologies, and domain knowledge MedInsight has built over the three decades become more valuable in an AI-enabled world, not less relevant. As we seek to leverage the power of AI in furthering our mission and amplifying the value we provide to our clients, we apply the following principles in our development of AI-powered healthcare analytics solutions.

Principle #1: Humans direct, AI executes — and workflows must be designed accordingly

Humans and AI are not interchangeable labor inputs. AI excels at executing well-specified tasks — logistics, legwork, translating requirements into actions, producing content against defined criteria.

However, AI models are naively agreeable by design and therefore ill-suited to direct analysis or final decision-making. Humans must retain ownership of vision, objectives, and the definition of success. AI systems that are given latitude to define the problem they are solving, rather than being given the problem, produce confident-looking results that can look right and be materially wrong in subtle, hard-to-detect ways.

The practical implications of this principle are twofold. First, AI workflows ought to be iterative and conversational, not single-shot. Systems that engage the user by surfacing assumptions, clarifying the analytical question, presenting intermediate outputs for review, and offering to adjust plans are better able to distribute risk across a process where humans can intervene. Second, and relatedly, AI-driven analytics platforms ought to be designed around the kinds of users they are meant to serve, allowing and encouraging those users to inject their expertise into the system's context (more on how that context should be set up below). The most reliable path to a sound analytical result is a progressive dialogue between director and worker, not a vending machine for answers of unpredictable quality.

Principle #2: Analytical quality begins with a well-formed question

Analytical quality and AI performance share a common theme: both are driven overwhelmingly by the clarity of the question and the richness of the context, not by architectural sophistication. A well-formed analytical question has clearly defined concepts (metrics, populations, timeframes, comparisons) and a clear sense of what kind of analysis is being performed (trend, cohort identification, causal inference, outcome measurement, etc.). Constructing that question is an area where AI can assist as an early step in an agentic workflow. As mentioned above, the user must be involved in this stage, providing their vision and context. Well-designed AI systems should also validate that the question being answered is the one the user intended to ask.

Principle #3: AI should route to authoritative tools, not generate answers from scratch

While users of AI bring critical expertise or perspective to an analytical exercise, the platform itself must also bring value, specifically in the form of complementary capability the user does not possess. The value of an analytics platform, AI-powered or otherwise, is that it allows analysts with specific questions to leverage deep experience and expertise the analyst does not possess but which is highly relevant to answering the question. As mentioned, AI does not replace this kind of expertise but rather unleashes it at a previously unimaginable speed and scale. In this context, AI functions as an orchestrator, not an oracle.

Rather than generating its own methodology to answer an analytical question, AI should map the question to the best available tools, calculations, and validated data representations available to it, and route through those. Raw query languages like SQL are easy for AI to get wrong in ways that produce misleading or flawed results without obvious errors. On the other hand, validated data, models, and semantic layers expose constrained, well-defined sets of valid states. Routing through those structures keeps results auditable, traceable, and verifiable because the output reflects what the tool is known to do rather than what the AI inferred it should do. The measure of a well-designed AI feature is not how much it generates from scratch, but how reliably it connects the right question to the right validated tools.

Principle #4: Platform coherence and metadata quality are now foundational investments

AI models operate most reliably within coherent, well-structured systems. Inconsistencies that experienced humans can navigate informally — differently named metrics that should be the same, logic scattered across systems, tacit knowledge about implementation decisions — become significant pitfalls for AI. Metadata quality, which was previously a “nice to have,” is now a bottleneck: agents are only as capable as the context they can access, and context that is not explicitly defined cannot be used reliably. That includes the designing organization's own context and expertise as well as client-specific metadata such as preferred definitions, implementation decisions, strategic initiatives, and informal institutional knowledge. The internal work of formalizing business logic, improving metadata, and building coherent platform structures is therefore foundational to AI strategy.

Principle #5: Rigorous evaluation is critical

AI systems that are deployed without structured evaluation quickly drift from their intended behavior, and the drift can be invisible until something goes wrong. Maintaining principled AI use requires deliberate, ongoing evaluation: defined sets of test cases and desired interactions (“given this question, here is how the analysis should proceed and what the result should look like”), review of actual outputs against those criteria, and feedback loops that improve the system without introducing new failure modes. Building and maintaining these evaluation sets takes investment, but it is the best way to determine how AI applications actually perform against the standards we claim.

Conclusion

At Milliman MedInsight, we recognize the enormous opportunity presented by AI. It is clear that data and analytics platforms that leverage this technology thoughtfully can greatly enhance and accelerate insight generation beyond what was previously possible. We are actively incorporating AI into our platforms, analytics, and processes to continue to deliver value to our clients. The principles we have laid out here speak to our commitment to prioritize intentional, responsible, and robust use of AI over deployments without sufficient direction, validation, or oversight. We are excited by the opportunity to leverage AI technology to further our mission of being the most trusted source for healthcare analytics to empower our clients across the healthcare industry.

Milliman MedInsight

Founded in 1998, Milliman MedInsight® is a leading provider of healthcare data and analytics, trusted by over 300 payers, providers/ACOs, employers, and government agencies worldwide. Our comprehensive suite of analytics and data solutions empowers organizations to leverage healthcare data for informed decision-making and improved clinical and financial outcomes. With deep industry expertise and advanced technology, we deliver actionable insights into healthcare utilization, costs, quality, and performance. From risk management to value-based care, Milliman MedInsight empowers stakeholders to navigate the complexities of the healthcare landscape and achieve sustainable success.

Learn more at: medinsight.com

Contact

Jake Linenthal
jake.linenthal@milliman.com

Chas Busenburg
chas.busenburger@milliman.com

Robert Richards
robert.richards@milliman.com

Jake Stone
jake.stone@milliman.com

Iyibo Jack
iyibo.jack@milliman.com