

Whitepaper – AI Prioritisation for Life Sciences

Authors: Diana Gamez, Head of Data & AI Engineering

Category: Innovation & Technology

Estimated read time: ~15 min

Basel, Switzerland – June 2, 2025

As the adoption of large language models (LLMs) accelerates across nearly every industry, Life Sciences organisations are increasingly exploring how to integrate these technologies into their operations¹. Generative AI holds immense promise—ranging from reducing human error to significantly lowering operational costs. Yet, despite this potential, the Life Sciences sector remains notably cautious² when it comes to technology adoption³, particularly due to its high regulatory burden and low risk tolerance⁴.

This conservative approach has left many legacy systems firmly in place, and it presents a unique challenge for Chief Information Officers (CIOs): where should they focus their AI investments to ensure successful adoption? Unlike in other industries, failed initiatives in Life Sciences—especially those involving AI—are rarely acceptable, making it difficult to move past the proof-of-concept (POC) stage⁵.

At [MIGx](#), we bring deep cross-functional experience at the intersection of AI, data standards understanding, and enterprise-scale Life Sciences IT. Having supported global pharmaceutical and biotech organisations through digital transformations, we understand the unique cultural and operational constraints of this industry.

This whitepaper introduces a strategic, risk-aware framework for prioritising AI initiatives within Life Sciences. Instead of beginning with a broad exploration of use cases, we propose flipping the script: start by identifying processes where a tolerance for AI-driven decisions already exists. This approach not only accelerates time to production but also aligns with the sector’s compliance-driven mindset.

¹ <https://www.businesswire.com/news/home/20240607372298/en/Large-Language-Model-LLM-Markets-2024-2034-with-OpenAI-Google-Meta-Microsoft-Tencent-Yandex-Set-to-Dominate-the-%2485-Billion-Industry---ResearchAndMarkets.com>

² <https://www.scilife.io/blog/ai-pharma-innovation-challenges>

³ <https://www.coherentsolutions.com/insights/artificial-intelligence-in-pharmaceuticals-and-biotechnology-current-trends-and-innovations>

⁴ <https://www.hoover.org/research/fdas-risky-risk-aversion>

⁵ <https://www.bvp.com/atlas/the-healthcare-ai-adoption-index>

The overarching approach

Traditionally, many Life Sciences organisations have taken a bottom-up approach to AI implementation. This typically begins with the broad generation of use cases across departments, followed by a selection process for proof-of-concept (POC) trials, and—if successful—eventual progression to production deployment. While this may seem logical on the surface, it has frequently resulted in significant inefficiencies and missed opportunities.

Here's how the conventional model usually unfolds:

- A multitude of business cases are generated using standardised templates.
- Each department selects the cases it deems promising for a POC.
- Cross-functional teams implement AI-driven changes to the selected processes.
- A polished presentation showcases the outcomes of the pilot.

Just as the project gains traction, compliance or legal reviews raise concerns about risk, leading to project delays—or outright termination. This pattern illustrates a recurring issue: risk considerations are introduced too late in the process, often derailing projects that are otherwise technically and operationally sound.

At MIGx, we advocate for inverting the traditional model. Rather than beginning with a wide net of use cases, organisations should first identify business processes that are AI-compatible, tolerant of errors or uncertainty and would highly benefit from employing AI. Only then should they explore which tasks within those processes can benefit most from AI augmentation.

By aligning AI initiatives with processes that have a built-in resilience to failure, Life Sciences companies can:

- Accelerate time-to-production
- Reduce wasted investment in impractical use cases
- Navigate compliance and risk management earlier in the process
- Build momentum and confidence across the organisation

In our view, this risk-first, process-led strategy is the only viable path to achieving scalable, production-level AI adoption in an industry where failure is often unacceptable.

The POC-to-Production Bottleneck in Life Sciences

Why haven't self-driving cars gone mainstream, even though they are statistically safer than human drivers⁶? The answer lies not in performance, but in societal expectations and risk tolerance. While human error is broadly accepted, even rare mistakes by autonomous systems are seen as unacceptable. Society demands near-perfection from machines, especially when safety is at stake.

This same dynamic is playing out in Life Sciences. Although AI-powered solutions have shown promise during proof-of-concept (POC) stages, few make it to production. The reason? The industry's low tolerance for failure—particularly in regulated and high-stakes domains—becomes a roadblock when it's time to scale.

Despite impressive results in sandbox environments, projects often grind to a halt once they face:

- Compliance audits
- Legal and ethical reviews
- Risk committees
- Concerns over data integrity or explainability
- Technical capacity restrictions, such as context window limitations in LLMs.

This reality leads to a frustrating cycle where organisations invest heavily in exploratory AI projects, only to see them stall before delivering tangible business value.

To break this cycle, companies must reframe their selection criteria from "What's possible?" to "What's permissible?". The goal should be to start with business processes where some degree of imperfection is tolerated, enabling the organisation to build real-world experience, generate value, and gradually increase organisational confidence.

By focusing on risk-tolerant processes, Life Sciences companies can finally move beyond perpetual pilots and begin to realise AI's potential at scale.

Risk Tolerance as a Compass for AI Prioritisation

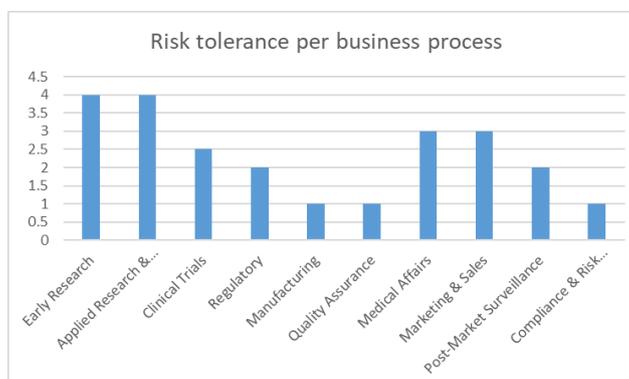
In the Life Sciences industry, risk is not just a factor—it's a framework. Regulatory agencies such as the FDA define explicit thresholds for acceptable risks⁷ in areas like clinical trials, pharmacovigilance, and manufacturing quality⁸. These thresholds provide a valuable guidepost for identifying where AI applications can realistically be deployed.

⁶ <https://waymo.com/blog/2024/12/new-swiss-re-study-waymo>, <https://www.euronews.com/next/2024/06/20/self-driving-cars-are-generally-safer-than-human-driven-ones-research-shows>, <https://www.nature.com/articles/s41467-024-48526-4>

⁷ <https://www.fda.gov/media/170461/download>

⁸ <https://pmc.ncbi.nlm.nih.gov/articles/PMC7785646/>

For instance, in clinical trials, the tolerance for adverse events is extremely low—often measured in occurrences per thousands or even millions of patient exposures. This stringent requirement implies that any AI solution deployed in such contexts must demonstrate near-flawless performance and be fully validated, traceable, and explainable. As a result, clinical operations are among the least permissive environments for early-stage AI integration.



Note: computed by MIGx as illustration based on discussion with customers

Conversely, earlier-stage research processes—such as target discovery, compound screening, or protein structure modeling—offer a greater degree of experimentation and flexibility. These functions are more insulated from patient-facing consequences and regulatory scrutiny, making them ideal candidates for piloting AI in production environments.

Take, for example, DeepMind's AlphaFold, which has revolutionised protein structure prediction. This application of AI is not only tolerated but actively celebrated in the research community because the downstream risk is relatively low, while the potential scientific and commercial value is immense.

To support prioritisation, organisations should consider creating a risk map of their high-level business processes—ranging from R&D and clinical trials to manufacturing and commercial operations. By evaluating the regulatory expectations, error tolerance, and existing controls within each area, CIOs can identify where AI is most likely to succeed in the short term.

This process-centered, risk-informed view transforms AI prioritisation from a speculative wish list into a focused, executable roadmap that aligns with the industry's core operating principles.

Strategic Implications for Life Sciences CIOs

For CIOs in the Life Sciences sector, the implications of a risk-first approach to AI adoption are clear: prioritisation must begin not with the most exciting technologies, but with the business processes most tolerant to AI-supported decisions.

The first step is to assess the organisation's risk appetite across departments. Each functional area—whether research, clinical, regulatory, manufacturing, or commercial—carries a distinct set of regulatory requirements and operational expectations. CIOs should work closely with compliance, legal, and business leaders to evaluate:

- What level of error is acceptable?
- Where can AI outputs be safely validated or overridden?
- What internal controls already exist to mitigate risk?

Once high-potential areas are identified, the next step is to drill down into specific tasks. Focus on those with:

- High volume or manual repetition
- Moderate-to-low impact of occasional failure
- Clear business ownership and measurable KPIs

These tasks are ideal entry points for automation or augmentation via AI. Importantly, successful deployment in these contexts builds both organisational confidence and proof points that can be used to guide more ambitious AI efforts over time.

Ultimately, CIOs must shift from a technology-driven roadmap to a risk-informed prioritisation model, positioning AI not just as a tool for innovation—but as a responsible, scalable component of digital transformation. This ensures to leave enough time for the technology to continue to progress, thus continuously decreasing the risk of failure, hence having more and more applicable business process.

Conclusion: From Exploration to Execution

The Life Sciences industry stands at the threshold of a new era—one where artificial intelligence offers transformational value, but only if deployed thoughtfully, strategically, and in alignment with regulatory realities.

The conventional bottom-up approach—starting with a broad list of use cases and filtering them through layers of review—has proven too slow, too costly, and too often unsuccessful. Instead, organisations should reverse their approach: start with risk tolerance, identify processes that can absorb AI-driven decisions, and then target specific, high-leverage tasks where automation or augmentation will be both impactful and permissible.

This risk-informed methodology allows companies to:

- Prioritise AI initiatives that have a real chance of reaching production
- Accelerate time-to-value by avoiding costly dead-ends
- Build internal momentum and regulatory trust in AI capabilities
- Create a scalable foundation for future, higher-risk deployments

6

For CIOs, the mandate is clear: don't aim to explore everything—aim to deploy what works. Focus less on speculative ideation and more on surgical implementation within domains that can truly accept it.

By doing so, Life Sciences organisations can finally unlock AI's promise—not just as a tool for innovation, but as a dependable force for productivity, quality, and long-term transformation.

Need guidance on where to start?

Our experts help Life Sciences leaders move from ideation to execution. [Book a strategy session](#) to assess your AI readiness and identify your first deployment opportunity.