

Ingo Bothe

Managing Consultant, Strategy Consulting

Alexis Anderson

Principal, Strategy Consulting

Alex Fink

Senior Vice President, Strategy Consulting

Anthony Guethert

Chief Technology Officer



Pharma's AI Dilemma: Balancing Risks and Rewards

**A strategic approach to overcoming
challenges and harnessing AI's potential**

Introduction

Artificial Intelligence (AI) is poised to revolutionize the pharmaceutical industry, offering unprecedented opportunities in drug discovery, clinical trials, and patient care. However, this potentially rapid transformation presents significant challenges and risks. This report, the second¹ in our series on AI in healthcare, provides pharmaceutical leaders with a comprehensive analysis of six very real dilemmas associated with AI adoption – **ethical considerations** (privacy & confidentiality, and bias & fairness), **operational factors** (safety, and transparency & explainability), and **human factors** (autonomy & decision-making, and training). We explore each dilemma in depth, offering industry-specific examples, risk assessments, and practical mitigation strategies.

While AI presents formidable challenges, we firmly believe that proactive management and strategic implementation can significantly mitigate risks. We conclude with three actionable recommendations for pharmaceutical leaders and decision-makers to help navigate the AI landscape responsibly and effectively.

The promise of AI in pharmaceuticals

AI is on the brink of revolutionizing various industries, including healthcare, in part by leveraging computational models and algorithms to analyze and interpret vast amounts of data.

It's already clear that Generative AI (GenAI), a branch of AI that focuses on creating new data or content based on existing information, holds significant

potential in healthcare. Unlike traditional AI systems that mainly work by classifying or predicting, GenAI aims to generate new, original data by learning from patterns and structures within existing datasets. This technology utilizes complex neural networks and deep learning techniques to synthesize information and create content that is similar to the original data it was trained on.

Recent examples

Sanofi's partnership with Exscientia leverages AI for accelerated drug discovery, potentially reducing time and costs in the development process.²

Novartis has employed IBM Watson Health's AI capabilities to enhance breast cancer diagnosis and treatment, providing evidence-based options to oncologists to enhance decision-making.³

Johnson & Johnson's collaboration with Microsoft has led to AI-driven educational tools for surgeons, creating virtual reality simulations for surgical training.⁴



The need for responsible AI adoption

While the pharmaceutical industry increasingly leverages GenAI, it is critical to balance innovation with responsibility.

Our aim is to guide pharmaceutical leaders in harnessing AI's transformative potential while navigating its inherent risks. We advocate for an approach of intelligent and safe AI adoption, one that maximizes benefits while rigorously addressing key considerations such as patient safety, data privacy, ethical implications, regulatory compliance, AI interpretability, and liability concerns.

By carefully managing these aspects, pharmaceutical companies can position themselves to reap the rewards of AI innovation while maintaining the trust

of patients, healthcare providers, and regulators alike. In the following sections, we will comprehensively examine the risks associated with AI adoption in the pharmaceutical industry.

Our analysis will cover a range of 6 critical concerns, from data security and algorithmic bias to regulatory challenges and ethical considerations. It's important to clarify that we don't believe these risks should deter innovation in any way. Rather, they should inform a strategic approach to AI implementation. To that end, each section also offers practical strategies for mitigating these risks. We also conclude with actionable recommendations and an overview of how Lumanity can help to empower pharmaceutical leaders to harness AI's potential safely and effectively.



Six ongoing dilemmas in AI for pharmaceutical leaders

Ethical considerations

01



Privacy & confidentiality

AI systems process sensitive patient data. Ensuring robust privacy protection is crucial to maintain patient trust and comply with regulations, while inadequate privacy safeguards could lead to unauthorized access or data breaches.

02



Bias & fairness

AI models learn from historical data, which may contain biases. Ensuring fairness across diverse patient populations is vital, yet biased algorithms could perpetuate health disparities.

Operational challenges

03



Safety & harm mitigation

AI can improve patient safety but also introduce new risks. Rigorous testing and monitoring are necessary; however, errors or unintended consequences could harm patients.

04



Transparency & 'explainability'

AI systems often operate as "black boxes." Understanding their decision-making process is challenging and a lack of transparency undermines trust and accountability.

Human factors

05



Autonomy & decision-making

AI algorithms influence clinical decisions. Balancing the roles of physicians and machines is essential, though overreliance on AI might diminish human autonomy and agency.

06



Education & training

Physicians need to understand AI's capabilities and limitations. Integrating AI education into medical training is essential, as uninformed use of AI may lead to suboptimal patient care.

Taking a
closer look: How
can we mitigate
the risks?



01



Privacy & confidentiality

Ethical considerations

Definition:

GenAI systems can process and analyze vast amounts of sensitive patient data with unparalleled efficiency. This capability can significantly enhance the quality of healthcare services by providing personalized treatment plans, predicting patient outcomes, and identifying potential health risks before they become critical. Ensuring robust privacy protection is crucial for maintaining patient trust and complying with stringent regulations like HIPAA (Health Insurance Portability and Accountability Act) in the United States and GDPR (General Data Protection Regulation) in Europe.

Example:

IBM's Watson Health recently demonstrated the effective implementation of privacy-preserving AI in healthcare. The system successfully balances innovation with data protection by implementing robust role-based access controls and maintaining data sovereignty

for healthcare providers. This approach ensured that while Watson Health can process and analyze health records to support clinical decision-making, the original data remains secure and controlled by the healthcare organizations that own it. Healthcare providers maintain full control over their patient data while benefiting from AI-driven insights for diagnosis, treatment recommendations, and predictive analytics.³

Risks:

The primary risk associated with using GenAI to handle patient data is the potential for inadequate privacy safeguards. If these systems are not designed with stringent security measures, they could become targets for unauthorized access or data breaches. Such incidents can expose sensitive patient information, resulting in identity theft, financial fraud, and a significant loss of patient trust. Additionally, healthcare organizations could face severe legal and financial repercussions for failing to protect patient data.

Recommended risk mitigation measures:

01

Implement data minimization principles by clearly defining what patient data is necessary for specific use cases and why. Organizations must ensure they collect and process only the minimum viable data needed for their intended purpose, with proper consent and documentation.

02

Establish comprehensive data governance frameworks that include clear policies on data collection, consent management, and purpose limitation. Ensure that data usage aligns with the original purpose for which consent was obtained, and there is clear ownership of data used to train an AI model.

03

Deploy role-based access control (RBAC) systems that restrict data access to those who need it for direct patient care or specific authorized purposes. This should include regular access reviews and robust authentication mechanisms.

04

Apply appropriate de-identification and anonymization techniques before processing data through AI systems. This should be done according to established standards and regulatory requirements.

05

Implement encryption for all patient data, both at rest and in transit, as part of a comprehensive security strategy. However, encryption should never be relied upon as the sole protection mechanism, especially when considering the use of external AI platforms or services.

By implementing these mitigation strategies, healthcare organizations can leverage the benefits of GenAI while significantly reducing the risks associated with privacy and confidentiality breaches.

02



Bias & fairness

Ethical considerations

Definition:

AI models, particularly those used in healthcare, learn from vast amounts of historical data. This capability allows them to identify patterns and make predictions that can enhance patient care. However, the data these models learn from may contain inherent biases, reflecting historical inequalities and disparities. Ensuring fairness across diverse patient populations is therefore vital. By carefully addressing these biases, AI can become a powerful tool to promote equity in healthcare, offering more accurate diagnoses and personalized treatments to a wider range of patient demographics.

Example:

A positive example here is the **UCSF Health AI program**, which aims to develop machine learning models to predict patient

outcomes and optimize treatment plans. The program emphasizes the importance of fairness by incorporating diverse datasets and employing bias mitigation techniques to ensure that their AI models serve all patient populations equitably.⁵

Risks:

The primary risk associated with AI in this context is that biased algorithms could perpetuate or even exacerbate existing health disparities. If the historical data used to train AI models contain biases, the AI systems might generate recommendations that favor certain groups over others, leading to unequal treatment and outcomes.

Recommended risk mitigation measures:

01

Train AI models using diverse and representative datasets that encompass various demographics, ensuring that they learn patterns that apply broadly across different populations. And, be clear about where your gaps are when publishing insights.

02

Regularly conduct bias audits and impact assessments to identify and address any biases in AI algorithms. These audits can help detect unintended disparities in AI predictions.

03

Implement algorithmic fairness techniques, such as reweighting or adversarial debiasing, to reduce biases in AI models and promote equitable outcomes.

04

Continuously monitor AI systems in real-world applications to identify and correct any emerging biases, ensuring that the models remain fair over time.

05

Engage with diverse stakeholders, including patients and advocacy groups, to understand their perspectives and incorporate their feedback when developing and deploying AI systems.

By taking these steps, healthcare organizations can leverage AI to promote fairness and equity in patient care while minimizing the risk of perpetuating health disparities.

03



Safety & harm mitigation

Operational challenges

Definition:

AI is set to significantly improve patient safety by providing early warnings of deteriorating conditions, automating routine tasks, and reducing human error. However, the introduction of AI also brings new risks, including potential errors and unintended consequences. Rigorous testing and continuous monitoring are necessary to ensure that AI systems enhance patient safety without introducing new hazards.

Example:

The **AI-powered Sepsis Watch system developed by Duke University Health System** is designed to identify early signs

of sepsis in patients. By continuously monitoring patient data and alerting clinicians to potential sepsis cases, the system aims to improve patient outcomes and reduce mortality rates associated with this critical condition.⁶

Risks:

Despite the potential benefits, AI can introduce new risks to patient safety. Errors in AI algorithms or unintended consequences of AI-driven decisions could harm patients. For example, an incorrect diagnosis or inappropriate treatment recommendation from an AI system could lead to adverse outcomes.

Recommended risk mitigation measures:

01

Conduct extensive testing and validation of AI systems before deployment to ensure their accuracy, reliability, and safety in real-world clinical settings.

02

Implement continuous monitoring and real-time auditing of AI systems to detect and address any errors or unintended consequences as they arise.

03

Design AI systems with redundancy safeguards, such as human-in-the-loop mechanisms, to ensure that clinicians can review and override AI recommendations if necessary.

04

Continuously update and improve AI models based on feedback and new data, ensuring that they remain accurate and effective over time.

05

Establish ethical guidelines and best practices for the use of AI in healthcare, emphasizing patient safety and the responsible deployment of AI technologies.

By implementing these measures, healthcare organizations can harness the potential of AI to improve patient safety while minimizing the risks associated with its use.



04



Transparency & explainability

Operational challenges

Definition:

GenAI systems often operate as “black boxes,” making it challenging to understand their decision-making processes. This lack of transparency can be problematic, particularly in healthcare, where trust and accountability are paramount. It is essential to check the accuracy of data used to train an AI model, minimizing the risk of transparency issues. By enhancing the transparency and explainability of AI systems, healthcare providers can better understand how AI-derived insights are generated, fostering trust among patients and enabling clinicians to make more informed decisions.

Example:

Google’s DeepMind has been working on making their AI models more interpretable, particularly in healthcare applications. For instance, their AI system for predicting acute kidney injury (AKI) provides clinicians with not only the prediction but also the contributing factors and confidence levels, helping doctors understand the rationale

behind the AI’s recommendations. However, it is important to note that Google has faced scrutiny for not fully adhering to GDPR rules in the UK, highlighting the necessity of compliance with data protection regulations to maintain trust. Using newer AI methods such as Retrieval-Augmented Generation (RAG) within Generative AI on deidentified data may be used to avoid these issues. This technique combines the strengths of retrieval-based systems and generative models, thereby enhancing transparency and reducing the reliance on personal data.⁷

Risks:

The lack of transparency in AI systems can undermine trust and accountability. If healthcare providers cannot understand or explain the decisions made by AI, they may be less likely to adopt these technologies. Additionally, patients may be reluctant to accept AI-driven recommendations without clear explanations, potentially compromising the effectiveness of AI in healthcare. Furthermore, non-compliance with data protection regulations like GDPR can erode trust and lead to legal repercussions.

Recommended risk mitigation measures:

01

Curate source data, develop and use AI models that prioritize interpretability, allowing clinicians to understand and explain the decision-making process.

02

Implement visualization tools that help clinicians see how AI models arrive at their conclusions, providing insights into the factors influencing predictions.

03

Provide thorough documentation and explanations of AI algorithms, including their limitations and potential biases, to help healthcare professionals make informed decisions.

04

Advocate for regulatory standards that mandate transparency and explainability in AI systems used in healthcare, ensuring that these technologies meet rigorous criteria for trustworthiness.

05

Ensure compliance with data protection regulations such as GDPR to maintain trust and avoid legal issues.

06

Minimize the use of personal data in AI models to reduce the risk of transparency issues and enhance privacy.

07

Educate healthcare professionals about AI technologies, including how to interpret and communicate AI-derived insights, to enhance their ability to use these tools effectively.

By prioritizing transparency and explainability, minimizing the use of personal data, and ensuring compliance with data protection regulations, healthcare organizations can foster trust and accountability in AI systems, ultimately enhancing their adoption and effectiveness in clinical practice.

05



Autonomy & decision making

Human factors

Definition:

GenAI algorithms have the potential to significantly influence clinical decisions by providing healthcare professionals with data-driven insights, diagnostic suggestions, and treatment recommendations. These AI systems can analyze vast amounts of medical data, including patient histories, imaging studies, and clinical trials, to offer evidence-based guidance. The integration of AI into clinical decision-making can enhance the accuracy of diagnoses, personalize treatment plans, and improve patient outcomes. However, it is essential to balance the roles of physicians and machines to ensure that human expertise and judgment remain central to patient care.

Example:

A prominent example of AI in clinical decision-making is the implementation of the AI platform developed by PathAI. This technology assists pathologists in diagnosing diseases by analyzing

pathology slides with high precision. The AI system can detect cancerous cells, measure tumor size, and even predict the likelihood of disease progression. By providing pathologists with detailed and accurate analyses, PathAI enhances their ability to make informed decisions, ultimately improving patient care.⁸

Risks:

One of the significant risks associated with the integration of AI in clinical decision-making is the potential for overreliance on AI systems. This overreliance could diminish human autonomy and agency, as healthcare professionals might begin to defer to AI recommendations without critically evaluating them. Such dependence could lead to a reduction in critical thinking and clinical judgment, undermining the physician's role and potentially leading to suboptimal patient care if the AI system provides incorrect or incomplete recommendations.

Recommended risk mitigation measures:

01

Ensure that AI recommendations are subject to clinical oversight by qualified healthcare professionals. Physicians should validate AI-generated insights and integrate them with their clinical expertise and patient-specific knowledge.

02

Foster a collaborative decision-making process where AI serves as a tool to augment, rather than replace, human judgment. Encourage interdisciplinary discussions to evaluate AI recommendations and make well-informed clinical decisions.

03

Provide comprehensive training for healthcare professionals on the capabilities and limitations of AI systems. Educating clinicians about how AI algorithms work can help them understand when to rely on AI and when to exercise caution.

04

Ensure transparency in AI algorithms by making their decision-making processes interpretable and understandable to clinicians. This transparency can help physicians critically assess AI recommendations and identify potential biases or errors.

05

Implement continuous monitoring and evaluation of AI system performance. Regularly update and validate AI models to ensure they remain accurate, relevant, and aligned with the latest medical research and guidelines.

By implementing these mitigation strategies, healthcare organizations can harness the benefits of AI in clinical decision-making while preserving the autonomy and agency of healthcare professionals. This balanced approach ensures that AI serves as a valuable tool to enhance, rather than replace, human expertise in patient care.

06



Education & training

Human factors

Definition:

For AI to be effectively integrated into healthcare, physicians and other healthcare professionals need to understand its capabilities and limitations. Integrating AI education into medical training is essential to equip clinicians with the knowledge and skills necessary to use AI tools effectively. This education can enhance their ability to critically evaluate AI recommendations, make informed decisions, and improve patient care.

Example:

The Mayo Clinic has developed a comprehensive training program that includes AI education for its medical staff.

This program covers the basics of AI, its

applications in healthcare, and practical training on how to use AI tools in clinical practice. By providing this education, the Mayo Clinic aims to empower its clinicians to leverage AI technologies effectively.⁹

Risks:

The uninformed use of AI in healthcare may lead to suboptimal patient care. If healthcare professionals do not fully understand how AI systems work, their limitations, and potential biases, they may misuse these tools, resulting in incorrect diagnoses, inappropriate treatments, and adverse patient outcomes.

Recommended risk mitigation measures:

01

Develop and implement comprehensive training programs that cover the fundamentals of AI, its applications in healthcare, and practical guidance on using AI tools.

02

Offer continuing education opportunities for healthcare professionals to stay updated on the latest developments in AI and its evolving role in clinical practice.

03

Foster interdisciplinary collaboration between AI experts and healthcare professionals to facilitate knowledge exchange and improve the integration of AI into clinical workflows.

04

Utilize simulation-based training to provide hands-on experience with AI tools, allowing clinicians to practice using these technologies in a controlled environment.

05

Include education on the ethical and legal implications of AI in healthcare, ensuring that clinicians are aware of the responsibilities and challenges associated with AI use.

By integrating AI education into medical training and providing ongoing learning opportunities, healthcare organizations can equip their clinicians with the knowledge and skills necessary to use AI effectively and responsibly.

Charting the path forward: Partnering for responsible AI innovation in pharma

As explored throughout this report, integrating AI in the pharmaceutical industry undoubtedly presents unprecedented opportunities along with significant challenges or dilemmas. To help your organization analyze and navigate

this complex landscape and emerge as a leader in responsible AI adoption, we recommend three critical steps that pharma company leaders should take right now:

01

Conduct a comprehensive AI readiness assessment

Understanding your organization's current AI capabilities and identifying gaps is crucial for strategic AI integration. This assessment evaluates your existing AI infrastructure, analyses data practices, assesses team AI literacy, and benchmarks your AI maturity against industry standards. Lumanity can provide a thorough evaluation, delivering a customized roadmap for strategic AI integration that aligns with your business goals and addresses potential risks.

02

Develop your AI governance framework

Establishing a robust governance structure is essential for responsible AI adoption. This framework encompasses ethical guidelines, risk management protocols, regulatory compliance processes, and mechanisms for maintaining transparency in AI decisions. We can help you design and implement a comprehensive AI governance framework tailored to the unique needs of the pharmaceutical industry, ensuring that your AI initiatives are both innovative and responsible.

03

Implement an AI-specific change management & training program

Successful AI adoption requires a shift in organizational culture and capabilities. This program includes AI literacy training for executives, skill development for technical teams, change management strategies, and communication plans to build trust in AI initiatives. Our experts can develop and deliver customized training and change management programs that will empower your team to leverage AI effectively and responsibly.

By focusing on these three critical steps, your organization can accelerate responsible AI adoption, minimize implementation risks, enhance competitive advantage through ethical AI practices, ensure regulatory compliance, and foster a culture of innovation balanced with responsibility.

Lumanity stands at the forefront in the use and integration of AI across the commercialization lifecycle. Our teams provide expertise to companies looking to leverage the full AI potential, with a clear and informed strategy to consider and mitigate associated risks, and rapidly gain real value from AI. [Contact us](#) to learn more about how Lumanity can support your unique challenge.

Our authors

Lumarity is building upon a history of innovation and existing technology-focused capabilities, having invested in the formation of a cross-functional technology team. This team is dedicated to developing, enhancing, and delivering new technology-driven solutions.

Leading this team is [Greg Haskins](#) as Chief Product Officer and [Anthony Guethert](#) as Chief Technology Officer.



Ingo Bothe

Managing Consultant, Strategy Consulting

Ingo began his career in biomedical sciences and rapidly transitioned to become a seasoned life sciences strategy consultant. With over seven years of consulting experience, he focuses his profound scientific expertise on delivering innovation-led solutions across a range of commercial strategy challenges for biotech and pharmaceutical clients.



Alexis Anderson

Principal, Strategy Consulting

Alexis has worked as a strategy and management consultant for the life sciences and MedTech industries for 12+ years and has deep experience in advising senior leaders and driving transformation in the key subject areas (digital health, patient engagement, and health equity).



Alex Fink

Senior Vice President, Strategy Consulting

Alex has worked as a consultant in the life sciences industry for 25 years in and has deep experience in advising senior leaders, across multiple therapy areas and more than 40 markets.

He leads our commercial consulting team in the EU and focuses on helping client bringing their innovations to patients.



Anthony Guethert

Chief Technology Officer

Anthony brings 25+ years of experience in banking, infrastructure capital, and life sciences sectors across Asia, America, and Europe, and leads the creation, enhancement, and delivery of technology-driven products and services.

References

1. <https://lumanity.com/perspectives/the-ai-will-see-you-now-opportunities-and-challenges-presented-by-the-use-of-generative-artificial-intelligence-ai-across-healthcare-and-the-pharma-value-chain/>
2. Sanofi. (2022). Sanofi and Exscientia collaborate to accelerate drug discovery with AI. [Sanofi Press Release](#)
3. IBM Watson Health. (2018). Transforming health: The impact of Watson Health. [IBM Watson Health Report](#)
4. Johnson & Johnson Medical Devices Companies. (n.d.). Johnson & Johnson Institute: Advancing medical training and education. [Johnson & Johnson Institute Overview](#)
5. UCSF Health. (n.d.). UCSF Health AI program. <https://www.ucsfhealth.org/>
6. Duke University Health System. (n.d.). Sepsis Watch. <https://www.dukehealth.org/>
7. DeepMind. (n.d.). DeepMind's AI system for acute kidney injury. <https://deepmind.google/>
8. PathAI. (n.d.). PathAI. <https://www.pathai.com/>
9. Mayo Clinic. (n.d.). Mayo Clinic AI training program. <https://www.mayoclinic.org/>

Lumanity applies incisive thinking and decisive action to cut through complex situations and deliver transformative outcomes to accelerate and optimize access to medical advances. With deep experience in medical, commercial, and regulatory affairs, Lumanity transforms data and information into real-world insights and evidence that powers successful commercialization and empowers patients, providers, payers, and regulators to take timely and decisive action.

Contact us to learn more about how Lumanity can support your unique challenge.

lumanity.com