Towards a Governance Framework for Generative AI in Drug Discovery: Ethical, Regulatory, and Practical Challenges

Imran Nasim[1,3] , **Adam Nasim**[2,3] ,
[1]IBM UK
[2]Quantitative Pharmacology, Merck KGaA, Lausanne, Switzerland
[3]Department of Mathematics, University of Surrey
imran.nasim@ibm.com, adam.nasim@merckgroup.com

Abstract

Generative AI technologies are revolutionizing drug discovery by accelerating the design of novel compounds, optimizing molecular interactions, and simulating clinical trials. These advancements promise to reduce the time and cost of drug development significantly. However, generative AI also introduces unique governance challenges, including ethical risks, regulatory compliance, and operational transparency. Evolving governance policies present challenges for deploying these technologies effectively. This work examines these developments and proposes a comprehensive governance framework tailored for the responsible deployment of generative AI frameworks in drug discovery.

1 Introduction

The drug discovery process is notoriously expensive and time-intensive, often costing over USD 2.8 billion and taking more than 12 years to develop a single novel drug [DiMasi *et al.*, 2016]. To address these challenges and rising costs, adopting more efficient strategies is imperative. Generative AI (GenAI) is revolutionizing drug discovery by streamlining the traditionally complex and costly process of identifying viable compounds. By leveraging advancements in algorithmic design and computational hardware, GenAI facilitates the design of novel molecular structures and accurately predicts their biological impacts. This innovation accelerates drug discovery, complements traditional methods, and offers significant savings in time and cost [Mak *et al.*, 2024]. Tools such as DeepMind's AlphaFold and Nvidia's BioNemo platforms exemplify the transformative potential of these technologies [Jumper *et al.*, 2021] [John *et al.*, 2024]. However, despite the promise of GenAI in drug discovery, its widespread adoption raises critical governance challenges.

- How can GenAI outputs align with ethical principles?
- What governance mechanisms ensure compliance with evolving policies?
- What unique considerations are required for deploying GenAI in pharmaceutical workflows?

This paper explores recent advancements in AI governance and proposes a framework to ensure responsible and innovative applications of GenAI in drug discovery.

2 Generative AI in Drug Discovery: Current Landscape and Challenges

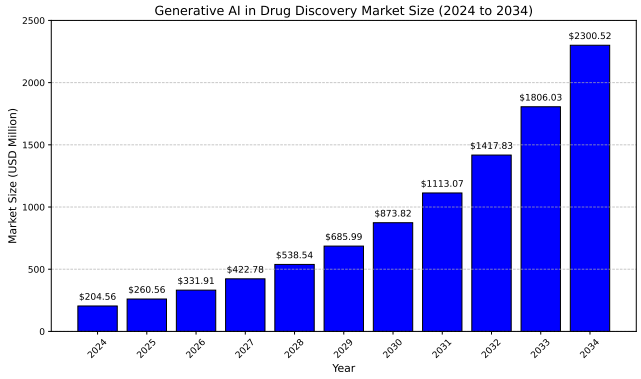

Figure 1: Projected Growth of Generative AI in Drug Discovery Market (2024-2034).

The global market for GenAI in drug discovery is estimated to be USD 204.56 million in 2024 and is projected to reach approximately USD 2.3 billion by 2034, reflecting a compound annual growth rate (CAGR) of 27.38% over the forecast period [1]. Such a rapidly developing field presents unique challenges from an ethical and regulatory point of view.

2.1 Ethical Risks

- **Bias in Training Data:** AI systems risk perpetuating biases in drug discovery [Mittermaier *et al.*, 2023], leading to inequitable healthcare solutions.
- **Dual-Use Concerns:** Generative models could inadvertently create harmful substances, posing security risks [Urbina *et al.*, 2022].

[1]Precedence Research https://www.precedenceresearch.com/generative-ai-in-drug-discovery-market

2.2 Regulatory and Compliance

Global AI governance regulations are critical for ensuring the ethical and transparent use of AI in drug discovery. The EU AI Act, for instance, classifies healthcare AI as high-risk, mandating strict transparency, risk management, and compliance protocols to prevent misuse, such as unsafe molecule generation. Similarly, the U.S. SR-11-7, while primarily for banking, emphasizes model validation and monitoring to ensure AI systems perform as intended without drifting from their purpose. In China and across Asia-Pacific, emerging frameworks prioritize privacy and ethical safeguards, while Canada's Directive on Automated Decision-Making highlights the importance of human oversight in high-impact systems. In the pharmaceutical domain, the FDA and EMA play pivotal roles in regulating drug development processes, including AI-driven innovations [Agency, 2024; Food and Administration, 2024]. Both agencies emphasize the need for transparency, reproducibility, and alignment with safety and efficacy standards for regulatory approval. These frameworks, alongside guidance from the FDA and EMA, provide insights to inspire a holistic governance process for addressing the unique challenges of generative AI in drug discovery, ensuring ethical and responsible innovation.

3 Governance Framework Considerations

i. Societal and Stakeholder Impact

- **Stakeholder Inclusion:** Engage diverse stakeholders, including patients, healthcare professionals, and regulators, to ensure solutions align with societal needs.
- **Anticipate Risks:** Evaluate potential negative outcomes, such as inequitable access to AI-generated drugs or misuse of AI for harmful applications.

ii. Bias Control

- **Data Audits:** Regularly review training datasets for biases, especially in underrepresented populations, to prevent skewed results.
- **Clinical Inclusivity:** Incorporate diverse demographic and genomic data to ensure generated drug candidates are effective across varied populations.

iii. Transparency & Explainability

- **Model Explainability:** Develop tools to provide insights into how molecules are generated and why specific candidates are selected for further testing.
- **Regulatory Readiness:** Maintain clear documentation of AI models, datasets, and algorithms to meet the transparency standards set by regulatory bodies such as the FDA and EMA.

iv. Accountability: Lifecycle Management

- **Defined Roles:** Establish clear accountability structures across teams managing data curation, model development, and deployment.
- **Ethical Oversight:** Create an AI ethics board to review potential risks, monitor compliance with regulatory frameworks, and manage ethical dilemmas.

v. Safety and Security

- **Pre-deployment Testing:** Conduct rigorous safety tests to validate the reliability of AI models in generating viable drug candidates.
- **Continuous Monitoring:** Deploy automated monitoring systems to identify performance drift, anomalies, or safety risks in real-time.
- **Model Security:** Safeguard AI models against cyber threats that could compromise sensitive pharmaceutical data.

vi. Privacy Protection and Data Security

- **Privacy by Design:** Implement encryption and anonymisation techniques when handling patient data.
- **Compliance Assurance:** Ensure all AI processes comply with global privacy laws, such as GDPR and HIPAA.

vii. Implementation Strategies

1. **Bias Monitoring and Equity Assurance:** Regularly assess AI outputs for biases that may exclude underrepresented populations and collaborate with clinical researchers to test the efficacy of AI-generated compounds.
2. **Consumer and Patient Protection:** Use GenAI to design safer drugs by predicting and reducing potential toxicity during molecule generation and ensure all outputs meet regulatory standards before clinical trials.
3. **Worker Support and Collaboration:** Equip researchers and developers with tools to integrate AI insights into traditional workflows and provide ongoing training on AI advancements and governance requirements.
4. **Global Leadership and Innovation:** Align GenAI practices with international standards, promote cross-border collaborations, and encourage open-source innovation while maintaining governance protocols.

viii. Measuring Governance Effectiveness

Organizations must implement metrics to monitor and evaluate the effectiveness of governance strategies:

- **Data Quality Metrics:** Track the diversity, completeness, and accuracy of training datasets.
- **Model Reliability Indicators:** Monitor success rates of AI-generated candidates advancing to clinical trials.
- **Compliance Audits:** Conduct periodic reviews to ensure adherence to regulatory requirements and internal governance policies.
- **Cost-Benefit Analysis:** Measure the value AI adds to drug discovery processes, such as reduced timelines and costs.

Conclusion

This proposed governance framework ensures that GenAI is utilised responsibly in drug discovery by addressing ethical, technical, and regulatory challenges. By adhering to these principals, organizations can drive innovation while safeguarding patients, stakeholders, and broader society.

References

[Agency, 2024] European Medicines Agency. Reflection paper on the use of artificial intelligence (ai) in the medicinal product lifecycle, 2024. Accessed: 2024-11-21.

[DiMasi *et al.*, 2016] Joseph A DiMasi, Henry G Grabowski, and Ronald W Hansen. Innovation in the pharmaceutical industry: new estimates of r&d costs. *Journal of health economics*, 47:20–33, 2016.

[Food and Administration, 2024] U.S. Food and Drug Administration. Artificial intelligence and machine learning (ai/ml) for drug development, 2024. Accessed: 2024-11-21.

[John *et al.*, 2024] Peter St. John, Dejun Lin, Polina Binder, Malcolm Greaves, Vega Shah, John St. John, Adrian Lange, Patrick Hsu, Rajesh Illango, Arvind Ramanathan, Anima Anandkumar, David H Brookes, Akosua Busia, Abhishaike Mahajan, Stephen Malina, Neha Prasad, Sam Sinai, Lindsay Edwards, Thomas Gaudelet, Cristian Regep, Martin Steinegger, Burkhard Rost, Alexander Brace, Kyle Hippe, Luca Naef, Keisuke Kamata, George Armstrong, Kevin Boyd, Zhonglin Cao, Han-Yi Chou, Simon Chu, Allan dos Santos Costa, Sajad Darabi, Eric Dawson, Kieran Didi, Cong Fu, Mario Geiger, Michelle Gill, Darren Hsu, Gagan Kaushik, Maria Korshunova, Steven Kothen-Hill, Youhan Lee, Meng Liu, Micha Livne, Zachary McClure, Jonathan Mitchell, Alireza Moradzadeh, Ohad Mosafi, Youssef Nashed, Saee Paliwal, Yuxing Peng, Sara Rabhi, Farhad Ramezanghorbani, Danny Reidenbach, Camir Ricketts, Brian Roland, Kushal Shah, Tyler Shimko, Hassan Sirelkhatim, Savitha Srinivasan, Abraham C Stern, Dorota Toczydlowska, Srimukh Prasad Veccham, Niccolò Alberto Elia Venanzi, Anton Vorontsov, Jared Wilber, Isabel Wilkinson, Wei Jing Wong, Eva Xue, Cory Ye, Xin Yu, Yang Zhang, Guoqing Zhou, Becca Zandstein, Christian Dallago, Bruno Trentini, Emine Kucukbenli, Saee Paliwal, Timur Rvachov, Eddie Calleja, Johnny Israeli, Harry Clifford, Risto Haukioja, Nicholas Haemel, Kyle Tretina, Neha Tadimeti, and Anthony B Costa. Bionemo framework: a modular, high-performance library for ai model development in drug discovery, 2024.

[Jumper *et al.*, 2021] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *nature*, 596(7873):583–589, 2021.

[Mak *et al.*, 2024] Kit-Kay Mak, Yi-Hang Wong, and Mallikarjuna Rao Pichika. Artificial intelligence in drug discovery and development. *Drug discovery and evaluation: safety and pharmacokinetic assays*, pages 1461–1498, 2024.

[Mittermaier *et al.*, 2023] Mirja Mittermaier, Marium M Raza, and Joseph C Kvedar. Bias in ai-based models for medical applications: challenges and mitigation strategies. *NPJ Digital Medicine*, 6(1):113, 2023.

[Urbina *et al.*, 2022] Fabio Urbina, Filippa Lentzos, Cédric Invernizzi, and Sean Ekins. Dual use of artificial-intelligence-powered drug discovery. *Nature machine intelligence*, 4(3):189–191, 2022.

A Generative AI Powered vs Classical Drug Discovery

Table 1: Comparison Between Traditional and Generative AI-Powered Drug Discovery Approaches

Aspect	**Traditional Drug Discovery**	**Generative AI-Powered Drug Discovery**
Process	Sequential and linear; progress depends on step-by-step experimentation.	Iterative and adaptive; models refine outputs based on feedback.
Effort	Labor-intensive; researchers manually design experiments and test compounds.	Data-driven and automated; algorithms generate molecules and predict outcomes.
Timeline	Time-consuming; often requires years to develop a viable candidate.	Faster; can reduce timelines by two-thirds compared to traditional methods.
Cost	Extremely expensive; costs often exceed billions of dollars.	Cost-efficient; delivers results at a fraction (as low as one-tenth) of traditional costs.
Data Integration	Limited to experimental results and predefined compound libraries.	Leverages diverse datasets, including genomic, chemical, clinical, and literature data.
Target Selection	Limited to known and predefined biological targets.	Identifies multiple alternative targets for experimentation using AI-driven insights.
Personalization	Focused on developing treatments for broader populations.	Highly personalized; leverages patient-specific biomarkers to tailor drug candidates.
Scalability	Challenging to scale; each project requires significant manual effort.	Highly scalable; AI models can analyze and design for multiple projects simultaneously.
Risk Assessment	Relies heavily on trial-and-error with limited predictive tools.	Uses predictive analytics to assess safety and efficacy earlier in the pipeline.
Innovation Potential	Incremental; heavily reliant on prior knowledge and standard approaches.	Disruptive; capable of discovering entirely novel compounds and mechanisms of action.
Regulatory Pathway	Relies on standardized frameworks but faces challenges with new mechanisms.	Requires new regulatory guidelines to address AI-generated outputs and explainability.