NEW YORK GENERAL GROUP

Artificial Intelligence and Drug Discovery

A strategic research report on AI-enabled pharmaceutical innovation

Prepared by
Yu Murakami, CEO
New York General Group

May 29, 2026

EXECUTIVE BRIEF

AI-enabled drug discovery is becoming an industrial operating system, not merely a research tool

> *The most defensible advantage in AI drug discovery will not be the model alone; it will be the closed loop between proprietary data, experimental validation, clinical insight, and regulatory-grade governance.*

200M+
predicted protein structures
AlphaFold DB [1]

$2.1B
Isomorphic financing
Reuters, 2026 [7]

Phase IIa
AI-discovered IPF candidate
Nature Medicine [5]

$314M-$4.46B
R&D cost estimate range
JAMA Netw Open [14]

Artificial intelligence is now being applied across the full pharmaceutical value chain: disease biology, target discovery, molecular design, lead optimization, ADMET prediction, clinical trial design, regulatory submissions, manufacturing quality, and post-market surveillance. The inflection point is driven by three converging forces: increasingly large biomedical data assets, foundation models capable of learning from sequences, structures, graphs, images and language, and automation infrastructure that can convert digital hypotheses into wet-lab evidence.

The strongest near-term impact is not the complete replacement of scientists. It is the acceleration of hypothesis generation and prioritization. AI can rank targets, design candidate molecules, flag toxicity risks, interpret papers and patents, enrich clinical trial populations, and standardize knowledge flows. However, the ultimate test remains clinical: a candidate must demonstrate safety, efficacy, manufacturability, differentiability, and regulatory acceptability in humans.

The strategic implication for New York General Group is that AI drug discovery should be evaluated less as a software investment and more as a portfolio operating model. Winning organizations will develop proprietary data assets, rapid experimental feedback loops, auditable AI governance, and partnership structures that translate computational advantage into clinical candidates.

Exhibit 1 | AI now touches every stage of the drug discovery and development value chain

The highest-value use cases connect prediction to experimental or clinical decisions rather than remaining as stand-alone analytics.

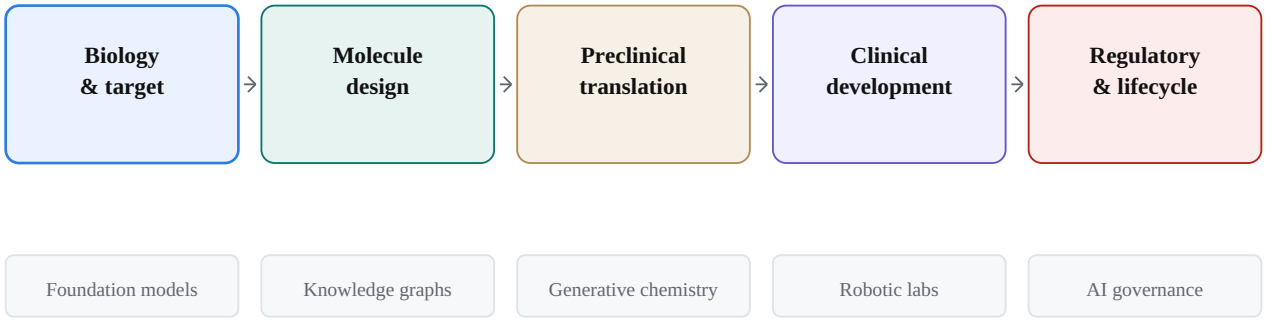

CONTENTS

Report structure

01

Strategic context

AI drug discovery is moving from a collection of impressive models to a new operating model for pharmaceutical research. The question is not whether AI can generate molecules; it is whether an organization can convert model outputs into validated, differentiated clinical assets.

New York General Group | AI and Drug Discovery

STRATEGIC CONTEXT

The drug discovery bottleneck is a decision-quality problem

Drug discovery is not primarily constrained by a lack of possible molecules. It is constrained by the cost and uncertainty of choosing which biological hypotheses, targets, molecules and patients deserve scarce experimental and clinical resources. The chemical search space is vast, disease biology is heterogeneous, and late-stage failure is expensive. AI is valuable because it can compress search, compare hypotheses at scale, and make the next experiment more informative.

The traditional pharmaceutical R&D model relies on sequential narrowing: identify a disease mechanism, nominate a target, find hits, optimize leads, run preclinical assays, enter clinical trials, and then learn whether the biology translates to humans. AI changes this sequence by enabling earlier integration of heterogeneous evidence. Genomics, proteomics, literature, clinical phenotypes, protein structures, assay results and chemical features can be combined into models that propose ranked options rather than isolated observations.

The economic logic is straightforward. Even modest improvements in early prioritization can be significant if they reduce the number of weak hypotheses entering costly development stages. Published estimates of the cost of bringing a new medicine to market vary substantially because they depend on therapeutic area, data availability, treatment of failures and cost-of-capital assumptions; one JAMA Network Open analysis notes prior estimates ranging from $314 million to $4.46 billion [14]. In that context, AI does not need to eliminate failure to matter. It can create value if it helps stop weak programs earlier, enriches trials more effectively, or redirects scarce experimental capacity toward better hypotheses.

Three shifts define the current moment

- From prediction to design. Earlier computational chemistry often focused on scoring known compounds. Modern generative models can propose new molecules, proteins or binders under multi-parameter constraints.
- From isolated models to integrated systems. The frontier is no longer a single model that predicts one endpoint; it is an orchestrated workflow connecting structure, property, literature, experiment and clinical knowledge.
- From research convenience to regulated decision support. FDA and EMA guidance shows that AI is entering contexts where credibility, documentation and monitoring matter, especially when outputs influence safety, efficacy or quality decisions [3][4].

Exhibit 2 | The most attractive AI use cases balance near-term evidence with strategic differentiation

Basic productivity use cases are useful, but enduring advantage requires proprietary data, experimental loops and decision-grade governance.

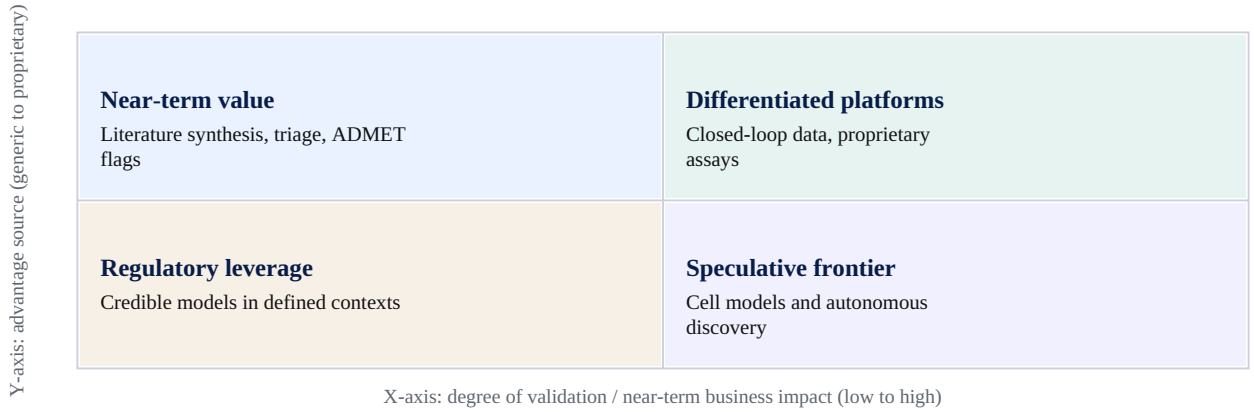

The key management thesis

The core thesis of this report is that AI will reshape drug discovery through a layered capability stack. At the base are data rights, data quality, assay standardization and compute infrastructure. Above that are models for biology, chemistry, language, imaging and clinical prediction. Above the models are decision workflows: target nomination committees, medicinal chemistry cycles, translational review boards, clinical protocol teams and regulatory documentation. The organizations that win will not simply subscribe to AI tools; they will redesign governance, incentives and experimental workflows around AI-enabled learning.

02

Technology map

AI-enabled drug discovery is a family of technologies. Structural prediction, graph learning, generative chemistry, language models, multimodal models and agents each solve different pieces of the pharmaceutical value chain.

New York General Group | AI and Drug Discovery

TECHNOLOGY MAP

The AI drug discovery stack has six core capability layers

AI drug discovery is often discussed as though it were one technology. In practice, it is a stack of capabilities, each with distinct data requirements, validation standards and business value. A robust strategy must understand where each capability is mature, where it remains experimental, and how it connects to wet-lab or clinical decisions.

Layer	Primary data	Typical applications	Strategic constraint
Biomedical knowledge AI	Literature, patents, ontologies, trial records	Target-disease linkage, landscape analysis, competitive intelligence	Hallucination control, source provenance, expert review
Structural AI	Protein sequences, structures, complexes	Target structures, binding modes, protein interfaces, docking hypotheses	Static structures do not guarantee efficacy or safety
Molecular graph AI	Molecular graphs, assay labels, ADMET labels	Property prediction, virtual screening, toxicity flags	Training set bias and weak extrapolation
Generative design	Chemical, protein, antibody and peptide representations	New molecule generation, multi-parameter optimization	Synthetic feasibility and experimental validation
Clinical AI	EHR, omics, imaging, biomarkers, trial data	Patient stratification, endpoint prediction, recruitment, external controls	Bias, generalizability and regulatory credibility
Agentic / closed-loop AI	Model outputs, automated lab data, ELN, robotics	Automated design-make-test-analyze cycles	Operational maturity, assay quality, governance

Structural AI: from protein folding to biomolecular interactions

AlphaFold made structural hypotheses dramatically more accessible by providing predicted protein structures at massive scale. AlphaFold DB now provides open access to more than 200 million protein structure predictions, creating a practical starting point for target analysis, variant interpretation and structure-based discovery [1]. AlphaFold 3 extends the paradigm by predicting joint structures of complexes that include proteins, nucleic acids, small molecules, ions and modified residues [2].

The strategic impact is substantial. Structure prediction reduces dependence on time-consuming experimental structure determination for early hypothesis generation. It can guide binding-site analysis, protein engineering, target tractability assessment and ligand design. Yet structural prediction should not be confused with therapeutic proof. Binding affinity, residence time, conformational dynamics, cellular uptake, metabolism and toxicity remain empirical or separately modeled questions.

Graph neural networks: molecular representation at scale

Molecules are naturally represented as graphs: atoms are nodes and bonds are edges. Graph neural networks can learn molecular representations for property prediction, virtual screening, molecule generation and synthesis planning. Review literature identifies GNNs as a core component of AI-aided drug discovery across molecular property

prediction, virtual screening, molecular generation and biomedical knowledge graphs [11].

GNNs are most valuable when the decision target is clearly defined and experimentally grounded: solubility, permeability, CYP inhibition, hERG liability, assay activity, selectivity or synthetic accessibility. Their main weakness is not mathematical elegance but transferability. A model trained on public or historical data may fail when applied to new chemical matter, new assay conditions or a new target class.

Generative AI: designing rather than only screening

Generative AI shifts discovery from screening available molecules to proposing new candidates under multiple constraints. The model can optimize simultaneously for potency, selectivity, solubility, permeability, metabolic stability, safety flags and synthetic feasibility. This is an inherently multi-objective problem because improving one property can degrade another.

Rentosertib, formerly ISM001-055, is an important case because it demonstrates that a generative AI-discovered TNIK inhibitor can progress into human clinical evaluation for idiopathic pulmonary fibrosis [5]. The lesson is not that generative AI has solved drug discovery. The lesson is that AI-generated hypotheses can now reach clinical proof-of-concept testing, thereby shifting the debate from computational novelty to clinical translation.

Large language models and agents

LLMs are valuable because pharmaceutical knowledge is text-heavy. Key evidence sits in papers, patents, trial registries, regulatory documents, medical notes, internal reports and expert memos. LLMs can synthesize these sources, create first-pass scientific maps and support document workflows. Reviews have described LLM applications from disease mechanism analysis and molecule design to clinical trial support [12].

The next frontier is agentic AI: systems that can call databases, run molecular models, design experiments, query electronic lab notebooks and produce auditable recommendations. However, agentic systems increase risk. They require stronger access controls, source traceability, human-in-the-loop review, and documented boundaries between exploratory suggestions and regulated decisions.

03

Use cases

AI delivers value when it improves a consequential decision: which target to pursue, which molecule to synthesize, which patient to enroll, which safety risk to monitor, and how to document evidence for regulators.

New York General Group | AI and Drug Discovery

USE CASES

AI value accumulates across the pharmaceutical chain

1. Target discovery and disease biology

Target discovery is the first major bottleneck. A target must be associated with disease, causally relevant, pharmacologically tractable, safe to modulate, measurable through biomarkers and differentiated from existing approaches. AI can integrate genome-wide association data, transcriptomics, proteomics, single-cell atlases, pathway maps, phenotypic screens and literature. The output should not be a target answer; it should be a ranked set of target hypotheses with evidence, uncertainty and experimental tests.

The greatest benefit is counterfactual thinking. AI can connect signals across datasets that would be difficult to integrate manually, such as disease genetics, cell-type specificity, pathway compensation and prior clinical failures. The strongest target discovery systems therefore combine knowledge graphs with causal inference, perturbation data and human review.

2. Hit discovery and virtual screening

AI can reduce the number of compounds that must be physically screened by ranking virtual libraries, public databases or proprietary collections. Unlike classical docking alone, modern workflows can combine structural information, ligand similarity, graph embeddings, active learning and assay data. The goal is to allocate experimental capacity toward compounds that are both scientifically plausible and chemically diverse.

3. Lead optimization and multi-parameter design

Lead optimization is where AI can become embedded in day-to-day medicinal chemistry. The model recommends which analogs to synthesize next, balancing potency, selectivity, solubility, permeability, metabolic stability, off-target risk and synthetic accessibility. The best systems work as design-make-test-analyze loops rather than one-off prediction engines. Each round of experimental data updates the model and improves the next round of suggestions.

4. ADMET and translational risk reduction

Many promising molecules fail because they cannot become safe, usable drugs. ADMET prediction is therefore a high-value AI use case. Models can flag likely liabilities, including CYP interactions, hERG inhibition, hepatotoxicity, low solubility, poor permeability or unstable metabolism. Iambic Therapeutics has highlighted models such as Enchant that aim to predict drug viability from early-stage discovery, including pharmacokinetic, efficacy and toxicity challenges [8].

The limitation is that toxicity is rarely a simple single-endpoint phenomenon. Rare adverse events, long-term toxicity, immune mechanisms and metabolite-driven effects can be underrepresented in data. ADMET AI should therefore be used as risk triage and prioritization, not as a replacement for toxicology.

5. Clinical development

Clinical AI can support patient stratification, endpoint prediction, recruitment, site selection, dropout risk analysis, disease progression modeling and real-world evidence generation. The strategic value is especially high in

heterogeneous diseases where a treatment may work in a subset of patients but fail in an undifferentiated population. AI can help define the subset earlier, improving the probability that a trial tests the right biological hypothesis in the right patients.

However, clinical AI sits closer to regulated decision-making. Bias, missingness, external validity and model drift are central issues. FDA guidance emphasizes context of use and risk-based credibility assessment for models that support regulatory decision-making about drug safety, effectiveness or quality [3].

6. Manufacturing, quality and pharmacovigilance

AI also has downstream value after candidate nomination. Manufacturing applications include process monitoring, batch deviation prediction, quality analytics, supply-chain planning and anomaly detection. Pharmacovigilance applications include adverse event signal detection, case intake triage, duplicate identification and literature surveillance. EMA’s lifecycle view is important because it frames AI as relevant across discovery, development, manufacturing and post-authorization monitoring [4].

04

Case studies

Market signals show a transition from scientific demonstration to capital formation, partnerships and clinical proof points. The strongest cases combine credible science with industrial execution.

New York General Group | AI and Drug Discovery

CASE STUDIES

The market is separating platforms from proof points

AI drug discovery is a heterogeneous market. Some companies sell platform capabilities, some develop internal pipelines, some partner with large pharmaceutical companies, and some provide enabling models or infrastructure. The sector is now being judged by increasingly concrete milestones: clinical candidates, proof-of-concept data, strategic collaborations and regulated use cases.

Company / program	Signal	Strategic interpretation
Google DeepMind / AlphaFold	More than 200 million structures in AlphaFold DB; AlphaFold 3 extends to complexes [1][2]	Structural AI has become foundational infrastructure for biology and early discovery.
Isomorphic Labs	Collaborations with Eli Lilly and Novartis near $3 billion excluding royalties; $2.1 billion financing reported in 2026 [6][7]	Capital and partnerships indicate belief in AI-native discovery platforms at industrial scale.
Insilico Medicine / rentosertib	Generative AI-discovered TNIK inhibitor evaluated in randomized Phase IIa trial for IPF [5]	AI discovery has crossed into clinical validation, though late-stage proof remains required.
Iambic Therapeutics	Enchant model and Takeda collaboration valued at more than $1.7 billion [8][9]	The field is expanding from structure to preclinical performance and pharma partnership execution.
Chan Zuckerberg Biohub / ESM family	Protein-biology world model reported in 2026; ESM lineage originates with the Meta AI/FAIR/EvolutionaryScale protein-modeling track [10][18]	Protein foundation models are becoming shared infrastructure, but attribution should distinguish Biohub deployment from the original ESM lineage.

Case study 1: AlphaFold and the structural biology platform shift

AlphaFold’s impact is best understood as a platform shift. Instead of waiting for every target structure to be solved experimentally, researchers can begin with high-quality computational hypotheses. This does not eliminate crystallography, cryo-EM or NMR; rather, it changes their role. Experimental structure methods remain essential for validation, ambiguous regions, dynamic states and drug-bound conformations, but computational prediction moves structure earlier into the decision process.

AlphaFold 3 further expands the scope from single protein structure to biomolecular interactions, including proteins, nucleic acids, small molecules, ions and modified residues [2]. This is directly relevant to drug discovery because therapeutic action is mediated through interactions, not isolated proteins.

Case study 2: Isomorphic Labs and the AI-native discovery model

Isomorphic Labs represents the AI-native drug discovery thesis: start with advanced AI capabilities and build a pharmaceutical R&D engine around them. The company’s collaborations with Eli Lilly and Novartis were described as having potential value near $3 billion excluding royalties [6]. Reuters later reported a $2.1 billion financing round to scale Isomorphic’s AI-driven drug discovery engine [7].

The strategic lesson is that large pharmaceutical companies are not only buying software. They are seeking access to new discovery architectures that combine structure, chemistry, biology and machine learning. The test for this model

will be whether it produces differentiated clinical candidates and whether partnerships translate into durable pipeline value.

Case study 3: Insilico Medicine and clinical translation

Insilico Medicine's rentosertib case is one of the most concrete clinical proof points in the field. The Nature Medicine report describes a generative AI-discovered TNIK inhibitor for idiopathic pulmonary fibrosis evaluated in a randomized Phase IIa trial [5]. The significance is that the AI contribution spans target identification and molecule design rather than only retrospective analytics.

Even so, this should be interpreted carefully. Phase IIa evidence is not equivalent to approval, and IPF is a difficult disease area. The broader lesson is that AI-discovered assets can progress far enough to be judged by the same clinical criteria as conventional assets. That is progress, but it raises the bar: AI platforms will increasingly be evaluated by clinical outcomes, not just computational benchmarks.

Case study 4: Iambic and the move beyond structure

Iambic's Enchant model, as reported by Reuters, aims to predict early-stage drug performance across pharmacokinetic, efficacy and toxicity dimensions [8]. Takeda's collaboration with Iambic, reportedly valued at more than $1.7 billion, further indicates interest in AI systems that improve molecular quality and speed in small-molecule programs [9].

This is strategically important because many AI systems perform best on well-defined structural or binding tasks, while clinical success depends on the whole drug profile. Models that help optimize preclinical viability may be more valuable than models that only generate attractive structures.

Case study 5: Biohub and protein world models

Chan Zuckerberg Biohub's protein-biology world model illustrates the rise of protein foundation models as shared infrastructure for biology and therapeutic design [10]. The attribution matters: the ESM lineage originated in Meta AI/FAIR work and the EvolutionaryScale team, which later launched ESM3 as a frontier protein model; Biohub's reported work should therefore be read as part of a broader ESM-family ecosystem rather than as the sole origin of ESM [18]. Protein models can support binder design, therapeutic protein engineering, immune modulation and functional inference, expanding the frontier from small molecules to programmable biology.

Market reality: setbacks are part of the evidence base

A credible assessment of AI drug discovery must include negative evidence. BEN-2293, BenevolentAI's topical pan-Trk inhibitor for mild-to-moderate atopic dermatitis, was reported as safe and well tolerated in a Phase IIa trial, but did not demonstrate statistically significant effects on the main itch or eczema-severity efficacy endpoints across the intention-to-treat population [15]. This does not invalidate AI-enabled discovery; it demonstrates that AI-originated or AI-supported hypotheses remain subject to the same translational and clinical risks as any other drug program.

Sector restructuring also matters. Recursion's acquisition of Exscientia for $688 million created one of the best-known combinations of AI-enabled biology and AI-enabled chemistry [17]. Yet subsequent pipeline prioritization and workforce reductions, including a reported 20% layoff linked to pipeline cutbacks, show that platform scale does not remove portfolio trade-offs, cash discipline or clinical uncertainty [16]. These cases strengthen rather than weaken the central thesis: the durable advantage is not a model demo, but a governed evidence engine that learns from success, failure and deprioritization.

05

Risks and governance

The central risk is mistaking model output for biological evidence. The central governance task is to define where AI is exploratory, where it is decision-support, and where it influences regulated conclusions.

New York General Group | AI and Drug Discovery

RISKS AND GOVERNANCE

AI drug discovery requires scientific humility and operating discipline

Risk 1: Data quality and bias

AI models inherit the limitations of their data. Public datasets may overrepresent successful experiments, popular targets, easy-to-assay endpoints and well-studied chemical space. Negative results and failed experiments are often missing. Internal datasets can be powerful, but only if experimental conditions, assay protocols, metadata and identifiers are standardized.

A common failure mode is to train a technically impressive model on data that does not represent the real decision context. A model may perform well on historical benchmark splits but fail when asked to recommend novel molecules in a new program. Hasselgren and Oprea emphasize that the full potential of AI in drug discovery depends on sufficient ground truth and appropriate human intervention later in the pipeline [13]. This is especially dangerous when results are presented without uncertainty estimates.

Risk 2: Weak extrapolation

Innovation often requires moving outside the training distribution. Yet AI models are most reliable near familiar data. This creates a strategic tension: safe predictions may be incremental, while bold predictions may be unreliable. The best mitigation is an active learning loop that explicitly samples uncertain but informative candidates, then rapidly tests them experimentally.

Risk 3: Biological complexity

AI can identify patterns, but biology is causal, dynamic and context-dependent. A molecule can bind a target and still fail because the target is not disease-causal, redundant pathways compensate, tissue exposure is insufficient, toxicity emerges, or the patient population is misdefined. Therefore, AI outputs should be treated as hypotheses requiring mechanistic and translational validation.

Risk 4: Regulatory credibility

When AI supports regulatory decision-making, credibility must be established for a particular context of use. FDA’s draft guidance emphasizes a risk-based credibility assessment framework for AI models used to produce information or data supporting decisions regarding safety, effectiveness or quality [3]. EMA similarly frames AI across the medicinal product lifecycle, emphasizing validation, monitoring and responsible use [4].

A practical governance model should classify every AI use case into one of three categories: exploratory research, operational decision support, or regulated evidence support. Each category requires a different level of documentation, validation, review and monitoring.

Governance dimension	Minimum expectation for serious AI drug discovery programs
Context of use	Define exactly what decision the model supports and what it does not support.
Data lineage	Track source, consent or rights, preprocessing, missingness, assay conditions and version history.

Governance dimension	Minimum expectation for serious AI drug discovery programs
Model validation	Use fit-for-purpose metrics, external validation where feasible, uncertainty and stress testing.
Human oversight	Require expert review before program decisions or regulated claims.
Change control	Document model updates, data refreshes, performance drift and revalidation triggers.
Auditability	Maintain source citations, run logs, model versions, prompt history where relevant and decision records.

Risk 5: Intellectual property and partnership structure

AI drug discovery complicates intellectual property. Value may reside in generated molecules, training data, model architecture, proprietary assays, active learning loops, experimental results and clinical know-how. Partnership agreements must therefore specify data use rights, model improvement rights, ownership of generated candidates, downstream milestones, field restrictions and post-termination rights.

Risk 6: Dual-use and security

As AI systems become better at designing bioactive molecules and proteins, dual-use risk rises. Controls should include access management, screening for harmful outputs, audit logs, policy boundaries, and escalation procedures for sensitive biological or chemical design tasks. Responsible innovation is not optional in a field where the same capabilities can accelerate medicines or harmful biology.

06

Strategic recommendations

The strategic objective is not to look AI-enabled. It is to build a reproducible learning system that turns proprietary evidence into better portfolio decisions.

New York General Group | AI and Drug Discovery

STRATEGIC RECOMMENDATIONS

A practical AI drug discovery operating model

> *Build the capability around decisions, not around tools. Begin with the decisions that determine portfolio value, then select data, models and governance to improve those decisions.*

Recommendation 1: Build proprietary data advantage

Public models and public datasets will become increasingly commoditized. Durable advantage comes from proprietary, high-quality and decision-relevant data. This includes standardized assay results, negative experiments, translational biomarkers, structural data, medicinal chemistry design rationales, clinical phenotypes and manufacturing records. Data quality work is not administrative overhead; it is the foundation of AI differentiation.

Recommendation 2: Redesign the design-make-test-analyze loop

AI creates value when its predictions shape what scientists do next. Therefore, programs should be organized around rapid cycles: model proposes, scientists review, lab tests, data returns, model updates. Cycle time and feedback quality are more important than isolated benchmark performance.

Exhibit 3 | AI capability compounds when data, experiments and governance are designed as one loop

A mature organization moves from experimentation to portfolio-scale learning through repeatable operating routines.

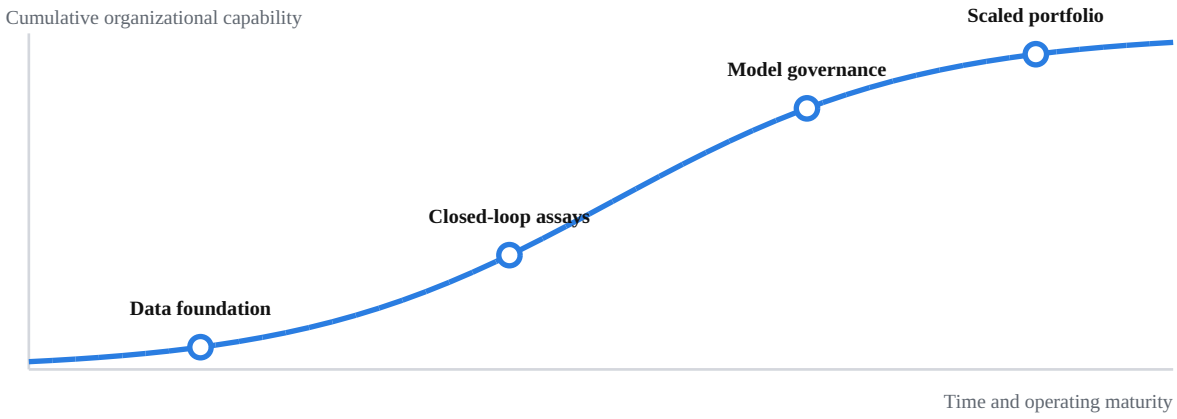

Recommendation 3: Establish an AI use-case portfolio

AI investments should be managed as a portfolio across near-term productivity, scientific differentiation and long-term platform bets. Examples include: literature intelligence for immediate productivity; ADMET prediction and medicinal chemistry copilots for program execution; structural and generative models for differentiated discovery; clinical AI for trial productivity; and regulated AI governance for durable scale.

Recommendation 4: Create an AI governance board with scientific authority

Governance should not be a purely legal or IT function. It must include drug hunters, translational scientists, clinicians, statisticians, regulatory experts, data engineers, information security leaders and IP counsel. The board should approve high-risk model uses, define validation standards, review model drift and decide when AI outputs can influence regulated documents.

Recommendation 5: Pursue partnerships with clear value capture

Partnerships with AI-native biotechs, foundation model providers, CROs, academic labs and cloud/compute companies can accelerate capability building. However, partnership value depends on data rights and learning rights. A collaboration that generates molecules but does not improve proprietary learning loops may produce limited long-term advantage.

Suggested 24-month roadmap

Horizon	Primary goal	Key actions	Success metrics
0-3 months	Strategic alignment	Map high-value decisions; select 3-5 priority use cases; define governance categories	Use-case portfolio approved; data owners assigned
3-9 months	Data foundation	Clean priority datasets; standardize assay metadata; implement lineage and access controls	Reusable data products; baseline model performance
9-15 months	Closed-loop pilots	Run active-learning pilots in 1-2 programs; integrate ELN and model outputs	Reduced design cycle time; validated hit/lead improvements
15-24 months	Scale and governance	Expand to portfolio workflows; implement model monitoring; document validation packages	Portfolio adoption; audit-ready model records; partnership leverage

Implications for New York General Group

For New York General Group, the highest-value role is likely not to duplicate the infrastructure of the largest pharmaceutical companies. It is to identify focused disease areas, platform partnerships or investment targets where AI can create asymmetric value. Attractive targets will possess three characteristics: proprietary data or assays, a short experimental feedback loop, and a credible path from computational hypotheses to clinically relevant assets.

The organization should treat AI drug discovery as a diligence domain and an operating thesis. Diligence should ask whether the company has unique data, validated models, experimental throughput, model governance, defensible IP, and a realistic translational plan. Operating support should focus on data infrastructure, partnership structuring, regulatory readiness and portfolio decision discipline.

CONCLUSION

AI will not abolish drug discovery; it will professionalize the uncertainty

AI is becoming one of the defining technologies of pharmaceutical innovation. Structural models make biological targets more legible. Graph models and generative systems make chemical exploration more efficient. LLMs organize the literature, patents and clinical knowledge base. Clinical AI can enrich trial populations and sharpen evidence generation. Regulatory guidance is beginning to define how AI can credibly support drug development decisions.

The correct strategic posture is neither hype nor dismissal. AI has not yet proven that it can systematically produce approved drugs at scale. But it has already changed how hypotheses are generated, how molecules are designed, how knowledge is synthesized and how the industry thinks about discovery productivity. The organizations that master AI-enabled learning will make better decisions faster.

The highest-performing version of AI drug discovery is not a black-box machine that claims to invent medicines alone. It is a disciplined scientific system in which computation expands imagination, experimentation tests reality, clinicians define human relevance, and governance preserves trust. In that system, AI becomes a force multiplier for human scientific judgment.

> *Strategic bottom line: the winning asset is not the algorithm; it is the continuously improving evidence engine that connects algorithmic predictions to biological truth.*

REFERENCES

Numbered references

[1] AlphaFold Protein Structure Database. EMBL-EBI and Google DeepMind. AlphaFold DB provides open access to over 200 million protein structure predictions. URL: https://alphafold.ebi.ac.uk/

[2] Accurate structure prediction of biomolecular interactions with AlphaFold 3. Abramson et al., Nature, 2024. AlphaFold 3 predicts joint structures of complexes including proteins, nucleic acids, small molecules, ions and modified residues. URL: https://www.nature.com/articles/s41586-024-07487-w

[3] Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products. U.S. FDA, draft guidance, January 2025. Risk-based credibility assessment framework for AI models in a defined context of use. URL: https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological

[4] Reflection paper on the use of artificial intelligence in the medicinal product lifecycle. European Medicines Agency, 2024. Lifecycle approach to AI in medicines development, manufacturing and pharmacovigilance. URL: https://www.ema.europa.eu/en/documents/scientific-guideline/reflection-paper-use-artificial-intelligence-ai-medicinal-product-lifecycle_en.pdf

[5] A generative AI-discovered TNIK inhibitor for idiopathic pulmonary fibrosis: a randomized phase 2a trial. Xu et al., Nature Medicine, 2025. Reports rentosertib, formerly ISM001-055, as a generative AI-discovered TNIK inhibitor for IPF. URL: https://www.nature.com/articles/s41591-025-03743-2

[6] Isomorphic Labs kicks off 2024 with two pharmaceutical collaborations. Isomorphic Labs, January 2024. Strategic collaborations with Eli Lilly and Novartis with potential value near $3 billion excluding royalties. URL: https://www.isomorphiclabs.com/articles/isomorphic-labs-kicks-off-2024-with-two-pharmaceutical-collaborations

[7] Google-backed Isomorphic raises $2.1 billion to scale AI-driven drug discovery. Reuters, May 2026. Reports Isomorphic Labs' $2.1 billion financing round to scale its AI drug discovery engine. URL: https://www.reuters.com/legal/litigation/google-backed-isomorphic-raises-21-billion-scale-ai-driven-drug-discovery-2026-05-12/

[8] Nvidia-backed AI firm Iambic unveils drug discovery breakthrough. Reuters, October 2024. Describes Iambic’s Enchant model for predicting preclinical drug performance. URL: https://www.reuters.com/technology/artificial-intelligence/nvidia-backed-ai-firm-iambic-unveils-drug-discovery-breakthrough-2024-10-29/

[9] Takeda deepens AI drug discovery push with $1.7 billion Iambic deal. Reuters, February 2026. Reports Takeda-Iambic collaboration for AI-designed small molecules in oncology and GI diseases. URL: https://www.reuters.com/business/healthcare-pharmaceuticals/takeda-deepens-ai-drug-discovery-push-with-17-billion-iambic-deal-2026-02-09/

[10] Zuckerberg's philanthropic venture unveils AI world model for drug discovery. Reuters, May 2026. Reports Chan Zuckerberg Biohub’s protein-biology world model based on ESM for drug discovery. URL: https://www.reuters.com/business/healthcare-pharmaceuticals/zuckerbergs-philanthropic-venture-unveils-ai-world-model-drug-discovery-2026-05-27/

[11] Graph Neural Networks in AI-aided Drug Discovery. Review literature, 2025. Surveys GNN applications in molecular property prediction, virtual screening, molecular generation and synthesis planning. URL: https://arxiv.org/abs/2506.06915

[12] Large Language Models in Drug Discovery and Development: From Disease Mechanisms to Clinical Trials. Review literature, 2024. Reviews LLM applications across target-disease linkage, molecular design, safety and clinical development. URL: https://arxiv.org/abs/2409.04481

[13] Artificial Intelligence for Drug Discovery: Are We There Yet?. Hasselgren and Oprea, 2023. Discusses AI's promise and limitations, emphasizing ground truth and human intervention. URL: https://arxiv.org/abs/2307.06521

[14] Costs of Drug Development and Research and Development Intensity in the US, 2000-2018. Sertkaya et al., JAMA Network Open, 2024. Provides transparent cost estimates and reports prior studies' R&D cost range from $314 million to $4.46 billion depending on assumptions. URL: https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2820562

[15] BenevolentAI announces top-line Phase IIa results for BEN-2293. BenevolentAI, April 2023. Reports BEN-2293 was safe and well tolerated but did not demonstrate statistically significant effect on NRS or EASI endpoints across the ITT population. URL: https://www.benevolent.com/news-and-media/press-releases-and-in-media/benevolentai-announces-top-line-phase-iia-results-its-topical-pan-trk-inhibitor-ben-2293-1-mild-moderate-atopic-dermatitis/

[16] Recursion lays off 20% of staff in wake of pipeline cutbacks. Fierce Biotech, June 2025. Reports Recursion workforce reduction and pipeline prioritization after earlier streamlining of clinical-stage programs. URL: https://www.fiercebiotech.com/biotech/recursion-lays-20-staff-wake-pipeline-cutbacks

[17] Biotech firm Recursion to buy smaller peer Exscientia for $688 million. Reuters, August 2024. Reports Recursion's agreement to acquire Exscientia for $688 million in an all-stock deal. URL: https://www.reuters.com/markets/deals/biotech-firm-recursion-buy-smaller-peer-exscientia-688-million-2024-08-08/

[18] EvolutionaryScale Launches with ESM3. AWS / EvolutionaryScale media alert, June 2024. Clarifies the ESM3 lineage and founding team associated with protein language models that began with ESM1. URL: https://press.aboutamazon.com/aws/2024/6/evolutionaryscale-launches-with-esm3-a-milestone-ai-model-for-biology