Connected by Design: How AI and Automation are Transforming Drug Discovery at BMS

When I think about our vision for AI and automation at Bristol Myers Squibb, I think about the music video by the band OK Go for their song "This Too Shall Pass."

Stick with me – I promise this will make sense.

The video begins with the band's bassist donning a pair of protective goggles and then rolling a toy truck into a line of upright dominos.

What follows is one of the most complex – and spectacular – Rube Goldberg machines I've ever seen. The dominoes cascade into a tethered string that propels a Hot Wheels car down a ramp, which knocks a billiard ball down another ramp, which knocks over a book connected to a string, and by then, the system is off to the races.

Over the course of about four minutes, tires roll, fans blow, and balloons and umbrellas fall from the ceiling. It ends with all four band members being sprayed with paintballs (hence the protective eyewear at the start), an outcome that began when that first domino fell.

What strikes me most about this video is that each isolated event – the dominoes falling, the tire rolling – is connected with every other step of this Rube Goldberg machine. The dominoes must fall not only for the Hot Wheels car to hit the billiard ball, but for the payoff at the end. It is a fully integrated system.

That is our vision for AI and automation at BMS: an integrated learning ecosystem, where AI helps every experiment, clinical readout, and partnership compound into higher-conviction scientific decisions, faster.

We aren't there yet – as an industry or at BMS – but we are taking deliberate steps to build the components and integrate them into a system that scales our ability to discover transformational medicines for patients.

Building a system that learns

Moving from discrete applications of AI to an integrated system is the difference between a line of dominoes falling in isolation…and having them trigger a paintball detonation a half mile away.

When I wrote about AI on this blog early last year, I discussed how BMS scientists were harnessing AI tools to solve complex biological problems and support our R&D principles to achieve sustained, top-tier R&D productivity.

The challenge I articulated then is still very real today: Only about 1 in 10 medicines in clinical development will become an approved medicine that reaches patients, and only about 1 in 3 of those are truly meaningful in the lives of patients. Ultimately, the greatest return on AI is improving those odds – moving from a system that is expensive, inefficient, and failure prone to one that is designed to improve our ability to deliver transformational medicines to the patients who need them.

In early 2025, we were still largely thinking of using AI for separate tasks: using phenotypic and pooled optical screens to identify **causal human biology**; applying our "Predict First" strategy to design and optimize molecules using AI to **match modality to mechanism**; and leveraging AI to identify biologically relevant disease segments and **accelerate the path to clinical proof of concept**.

That work is ongoing. But in a sign of just how quickly this technology has advanced in a year, we are now thinking of those applications as not individual steps but rather parts of an integrated learning ecosystem. This step-change in thinking is key to overcoming some of the hurdles we face in drug discovery and development.

Indeed, one of the greatest challenges in biopharma R&D is the complexity of research workflows. As OpenAI recently put it in their blog about their latest life sciences model:

Scientists must work across large volumes of literature, specialized databases, experimental data, and evolving hypotheses in order to generate and evaluate new ideas. These workflows are often time-intensive, fragmented, and difficult to scale.

If we are truly going to increase the probability of success in drug discovery, we cannot think in discrete steps – we cannot build a line of dominoes in isolation. Rather than relying on the siloed workflows described above, we need to build a system that is fast, integrated, and scalable.

A GPS navigation for human biology

Building such a system doesn't happen overnight. As impressive as the OK Go Rube Goldberg machine is, they did 89 takes – and only three were successful from beginning to end. The half-mile course took six months to plan, involving input from 60 experts including NASA engineers, and it took 30 minutes to reset after each attempt.

All this to say: building a seamlessly integrated system looks easy once it works, but getting there is hard!

We already have a strong foundation at BMS when it comes to applying AI to our first three R&D principles, as described above. But our vision goes further. It requires a "dream" state, which is enabled by a "build" phase.

Our dream is a future where we have **complete prediction** for the safety and efficacy profile of all targets nominated in early discovery. In this future state, we generate development candidates through **a fully automated lab-in-the-loop model:** AI-guided design–make–test–analyze cycles that shorten iteration time while maintaining scientific accountability. We dynamically manage every clinical experiment using **predictive modeling** and an investor mindset. And we use AI to analyze clinical feedback and external data to update mechanistic understanding and enhance model fidelity.

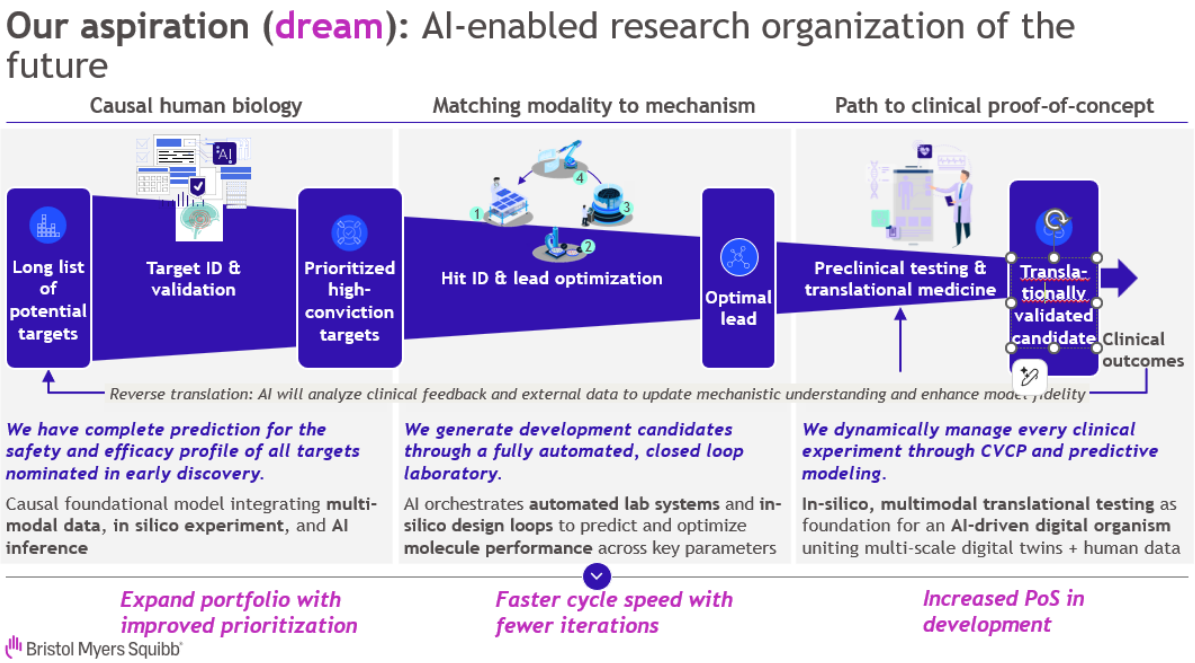

Bringing this vision to life are two foundational enablers: **a shared data and architecture backbone** and **AI co-scientists and agents**.

The backbone provides a unified, AI-ready data infrastructure spanning Research and Drug Development, eliminating the fragmented systems, siloed knowledge, and manual data assembly that slow decisions today, and enabling real-time access to integrated information across every layer of the system.

The co-scientists and agents help our scientists navigate, synthesize, and analyze diverse data and knowledge sources, removing the burden of manual literature review so that our scientists can focus on decision making and hypothesis testing.

Together, these capabilities are designed to build what my colleague Sally John refers to as a “GPS navigation” for human biology and disease. Just as GPS doesn't simply show you where you are but reroutes you when conditions change, our system would not only help our scientists identify the most promising path to a medicine, but dynamically guide them

there. While it will take time to build this system, doing so will bring us closer to realizing our ultimate dream of an AI-enabled lab of the future.

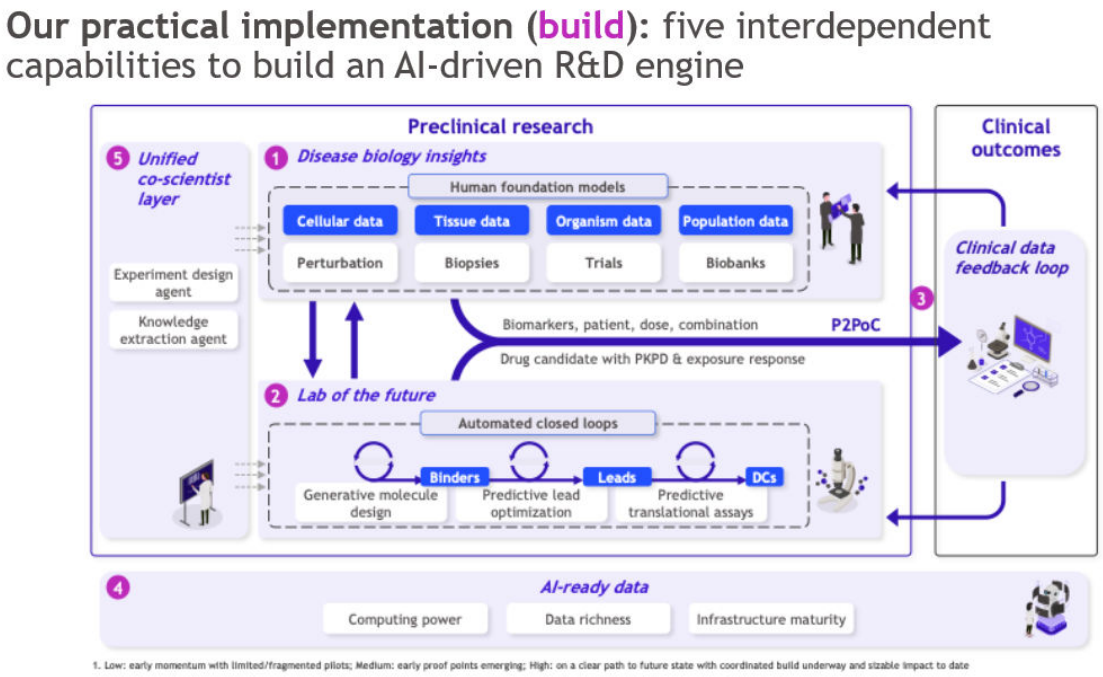

What transformation looks like in practice

So what might this look like in practice?

I recently provided some examples at the Stanford Drug Discovery Symposium, where I first reviewed our R&D principles and then described the Dream-Build progression of AI in drug discovery. Further, I provided concrete use cases of dream-build in action at BMS. (You can find the full presentation here.)

To start, I'll focus on our TYK2 inhibitor because – as many readers of this blog will know – this is a favorite target of mine (see this 2014 publication), and I think it provides a great example of what we are doing today. Moreover, it provides a roadmap of what the future of drug discovery could look like in an AI-enabled lab of the future.

Our TYK2 inhibitor is a highly selective small-molecule drug that binds to an allosteric site on the TYK2 molecule. We discovered this molecule using our R&D principles – but long before AI was integrated into the fabric of our Research organization. It was classic empirical discovery: a phenotypic pathway screen that led to an allosteric mechanism-of-action finding, an allosteric site screen and lead optimization. The time from screening to a drug candidate? **5 to 7 years**.

We have since partnered with Octant to analyze more than 20,000 TYK2 variants using AI-powered deep mutational scans (DMS) to reveal therapeutic insights about TYK2 signaling and disease. This allowed us to produce high-resolution structure-function mapping and identify novel allosteric sites. By coupling DMS with inhibitor treatment, we uncovered TYK2 variants that alter treatment potency and showed that DMS can prospectively reveal novel druggable sites.

Now dream with me: a lab of the future, fully enabled by AI. Starting from the beginning, that same technology could be used to help identify more novel targets like TYK2 and clarify structure-activity relationships. From there, we could use physics and machine learning-enabled tools for structural prediction and informed design, and then leverage our Predict First strategy to optimize our lead molecule, powered by an automated lab-in-the-loop model. The timeline for that process? **About 3.5 years**, based on current estimates – an approximately 30–50% potential improvement in time to the clinic.

The story would not end there. We would use AI to leverage data across potential autoimmune indications, link clinical and translational data with disease biomarkers, and then accelerate our decision making going into full development. Further, clinical and translational data from any clinical trials would be fed back into the discovery phase to inform models for other autoimmune therapies.

The TYK2 example also highlights a current limitation: sometimes we don't have the data we need to train the AI models. We performed DMS to create a function-phenotype map of a single protein, which was used to discover, retroactively, the allosteric binding site. Today, we have insufficient function-phenotype data on most proteins to train AI models to uncover allosteric sites.

Nevertheless, the above scenario isn't science fiction – we are already doing this work today. In my presentation, I shared how we used machine learning, physics-based design, and medicinal chemistry to optimize our fetal hemoglobin (HbF) asset for sickle cell disease, which I wrote about previously. In short, AI allowed us to model how modifications to the molecule would affect its ability to reactivate fetal hemoglobin, compressing what would have taken years of iterative lab work into a fraction of the time.

We are also modeling translational and clinical data from our immune reset programs currently in the clinic to optimize target antigen selection, modality, dose, and schedule, and to identify patient segments. Here, the goal is to use real-world clinical signals to sharpen our understanding of which patients are most likely to respond, so that each trial informs the next, and the system gets smarter as it goes. Long term, we are building an atlas of disease subtypes matched to the optimal immune reset therapy – whether that's a cell therapy asset, a T cell engager (TCE), or another modality such as in vivo CAR T.

The challenge, then, is two-fold: optimizing the individual components and connecting all these steps. We have many of the pieces – and we are generating data to optimize the

models of the future. Once we do, we will need to integrate them. We are dreaming of what is possible and building the infrastructure, but it will take time.

Where we are on the journey

Building an integrated AI ecosystem is not a single leap. It is a series of deliberate steps, and intellectual honesty demands we say plainly where we stand. Today, BMS has strong, working capabilities at the component level: we are running AI-guided design-make-test-analyze cycles in medicinal chemistry, deploying co-scientists to synthesize literature and experimental data, and generating the function-phenotype maps that will train the next generation of predictive models. These are not pilots or proofs-of-concept — they are embedded in how our scientists work. But the connections between these components are still being built, and the full feedback loop — where clinical data flows back to sharpen early discovery models in real time — remains a work in progress across the industry, not just at BMS.

How far are we from the dream state? Our honest estimate is that targeted AI-enabled workflows are already compressing timelines meaningfully. We believe 30 to 50% reductions in time-to-candidate are achievable in programs where the underlying data infrastructure is mature, based on modeling of our current TYK2 and HbF experience against historical baselines. But we are clear-eyed that this estimate rests on assumptions about data availability and model fidelity that do not yet hold uniformly across our portfolio. The goal for the next two to three years is to extend those conditions more broadly: building the shared data backbone, scaling the agent layer, and closing the loop between bench and bedside.

Think of it as the OK Go team mid-build: many of the individual mechanisms are working beautifully in isolation. The challenge now – and the opportunity – is the integration.

The people behind the machine

Whether we are in “dream” or “build” mode, however, one piece remains critical: Our people. If you watched the OK Go video until the end, you noticed something special: a throng of people cheering from above all the action. It underscored how many individuals it took to imagine this Rube Goldberg machine and bring it to life.

In R&D, those people are our scientists: individuals who bring human judgment to ambiguous problems, adapt when models fall short, and hold the system accountable to the biology. These scientists bring new skills, like automation engineering, and maintain an agile mindset, willing to learn fast and adjust as new information becomes available. Without them, a dream state will stay just that. If we are going to build a GPS for human biology, we need scientists who know where they want to go – and what tools they need to get them there.

The machine in motion

Remember those protective goggles the bassist puts on at the very start of the OK Go video? Before a single domino has fallen, before a tire has rolled or a balloon has dropped, he is already imagining the paintballs at the end. The whole machine is designed backwards from its payoff.

That is how we think about building our integrated AI ecosystem at BMS. The components we are building today – the shared data backbone, the AI co-scientists and agents, the closed-loop automated laboratory – are each remarkable in isolation. But what makes them transformative is what they are designed to deliver together: higher-conviction scientific decisions, a faster path from hypothesis to medicine, and ultimately, more transformational therapies reaching the patients who need them.

By scaling our R&D principles with AI, we are not just improving efficiency. We are dramatically expanding what our scientists can imagine. They can consider more possibilities, make connections they otherwise might have missed, and pinpoint the most promising hypotheses sooner. Layering automation on top amplifies this creativity even further.

The stakes could not be higher. Every asset that fails in a late-stage clinical trial represents an enormous investment of resources and, more importantly, a patient whose medical need will go unaddressed. Building our integrated system will take time, and the work is hard. But at BMS, we are taking intentional, meaningful steps to make it real – because somewhere at the end of our Rube Goldberg machine, a patient is waiting for a medicine.