








AI for Life Sciences: What is it Good For?

AI visionaries have spent years making ambitious claims about the potential transformative impact of AI on life sciences, but most of the real successes have resulted from using AI to improve less glamorous operational processes.

CONTENTS

-  We Are Nearing the End of the Current AI Hype Cycle
-  What AI Is Good For
-  Doing the Right Thing
-  Doing Things Right
-  Conclusion

ABOUT THE AUTHORS

Jim Manzi • Partner, Foundry.ai

Jim is a Partner and co-founder of Foundry.ai. He was co-founder, CEO and Chairman of Applied Predictive Technologies, which became the world's largest cloud-based AI software company and the dominant platform for rapid, iterative randomized trials of business programs. Previously, Jim developed pattern recognition software at AT&T Laboratories, and worked as a corporate strategy consultant.

Jim is the author of several software patents for the automation of randomized experiments, as well as the 2014 Harvard Business Review article "The Discipline of Business Experimentation." His 2012 book *Uncontrolled* on the development and application of randomized trials was widely reviewed in the New York Times, Wall Street Journal and other national publications.

Jim received an SB in mathematics from MIT, and was subsequently awarded a Dean's Fellowship in statistics to the doctoral program at the Wharton School of the University of Pennsylvania. He serves on the Board of Directors of Aledade, an innovative value-based care company backed by Venrock and Google Ventures.

Scott Setrakian • Vice-Chairman, Foundry.ai

Scott leads Foundry.ai's San Francisco office. Prior to joining Foundry.ai, Scott was co-founder and Managing Director of Applied Predictive Technologies, and led the company's work for life sciences companies in applying randomized experiments to improving commercial activities such as DTC advertising, targeting of speaker program and detailing efforts, patient adherence programs and RWE for therapeutic benefits vs. costs. Previously, he sat on the Board of Directors of Mercer Management Consulting, and ran the firm's global pharmaceuticals practice.

Scott received an AB in Human Biology and an MBA from Stanford University. He sits on the Board of Directors of the Buena Vista Funds, the William Saroyan Foundation, and the San Francisco Zoo.



WE ARE NEARING THE END OF THE CURRENT AI HYPE CYCLE

The term ‘AI’ (Artificial Intelligence) has gathered momentum over the last few years and is arguably today’s most over-hyped business buzz phrase. The vast majority of Global 2000 CEOs are being challenged by their Boards to demonstrate operational and financial benefits from applying AI to their businesses. Almost every large company has established some kind of working team to brainstorm and prioritize AI applications, and many companies have funded some specific pilot projects that have emerged from this process.

In our experience, most of these initiatives will end up disappointing their sponsors. The problem is *not* that experienced senior executives will somehow be misled by all this buzz. They know a hype cycle when they see one. They also understand the eternal verities of successfully introducing an important new digital technology:

1. Pilot quickly at low-cost
2. Demand some measurable short-term impact, and reinforce success
3. Maintain an unwavering focus on the bottom-line

This has been the playbook for the large companies that have most successfully introduced major new technologies, from enterprise data warehouses to CRM systems to the Web, and it is a requirement for successfully introducing AI to create shareholder value.

The problem, rather, is that most executives do not have a sufficiently granular understanding of AI to allocate pilot-stage resources well. As with all technological advances, it is unnecessary for CXOs to understand the detailed inner workings of a piece of AI software, just as it is unnecessary for them to understand the detailed engineering of their mobile phones' operating systems. But what they *do* need to know about AI is the answer to one question: *What is it good for?*

Right now, most senior executives don't have a clear answer to this question, and thus many companies are undertaking pilots addressing the wrong business problems, measuring the wrong short-term metrics, and trying to build platforms and roadmaps which – at this stage of AI's maturity – will often do more harm than good.

Nowhere has this been truer than in the attempt to bring AI to the life sciences industry. Capital has been pouring into the task. Healthcare AI start-ups raised almost \$1 billion of venture capital¹ just in the second quarter of 2019. Large incumbents are also making enormous investments in internal AI capabilities. Novartis, as an example, has roughly 250 full-time data scientists², has built a global AI hub, and has created major AI partnerships with IBM Watson, MIT, Intel and Quantum Black.

Unsurprisingly, much of this investment has been focused on trying to unlock huge, previously unsolved problems. There are more than 150 significant venture-backed AI companies³ *just* focused on the topic of drug development. In fact, there are so many now that they appear to be running out of name ideas – for instance, there are three separate AI-driven companies called, respectively, Helix, HelixAI and Healx.

But senior executives with P&L responsibility in publicly-traded corporations cannot exist forever on press releases, long term projections and big dreams. Eventually they have to produce results. And these leaders are now starting to say out loud what has been whispered all along by life sciences veterans: there is a huge disconnect between rhetoric and reality.

In one three-month period in mid-2019:

Pfizer Executive Chairman Ian Read said⁴:

“Using AI in drug discovery is extremely difficult and unlikely to be productive in the near term because our understanding of biology is not as deep as we’d like to believe it to be.”

Novartis CEO Vas Narasimhan said⁵:

“Opportunities presented by artificial intelligence are on the margin ... [I’m] skeptical that AI can keep up with the rapid advances in complex diseases like cancer, or predict which drugs will work better than humans can.”

Even Chris Gibson, CEO of AI-intensive Recursion Pharmaceuticals, said of AI applied to drug discovery⁶:

“We are just coming off a peak hype cycle and about to go through the valley of disillusionment.”

IBM halted sales of Watson for Drug Discovery, which had been the most promoted piece of software on the planet for applying AI to this area⁷.

A dark, moody photograph of a laboratory setup. In the foreground, there are two glass beakers, one on the left and one on the right, both containing a clear liquid. A metal stand with a circular top and a base is visible in the background. The lighting is low, creating a somber and scientific atmosphere.

“

These leaders are now starting to say out loud what has been whispered all along by life sciences veterans: there is a huge disconnect between rhetoric and reality”

Novartis's Narasimhan, a physician and former McKinsey consultant, has been as sophisticated and aggressive as any Big Pharma CEO in the world in driving the AI revolution within his company. He recently outlined the hard lessons he has learned in a revealing discussion with venture giant Andreessen Horowitz⁸:

“As we've gotten quite scaled and working on digital health and data science, we've learned there's a lot of talk and very little in terms of actual delivery of impact.”

He then pointed out topic after topic where results have not lived up to the hype:

TOPIC	QUOTE
Undirected development of novel insights	“we've not been able to crack”
Predicting in vivo therapy impacts without randomized trials	“puts a lot on the statistics that I don't think we have”
Using sensors to radically change the economics of clinical trials	“a lot of hype about expectations”

Interestingly, however, Narasimhan did drop a couple of breadcrumbs about two very boring-sounding areas where AI appears to be driving actionable results in Novartis: **forecasting corporate financial results** and **clinical trial enrollments**. Neither is as transformational as automated discovery of a new blockbuster molecule, but when AI is properly understood, these are exactly the kinds of successes we should expect.

WHAT AI IS GOOD FOR

AI on television is robots playing Jeopardy. But in our experience, AI that makes money for large corporations almost always follows a specific pattern: it is software that uses data + math to create statistical improvement in a repetitive business decision process. Operational processes like the details of executing clinical trials. Or back-office processes like finance or procurement. Or go-to-market processes like targeting detailing efforts, DTC advertising and digital outreach. Or even certain sub-processes within drug discovery where there have been significant successes, such as synthesis prediction⁹, automated compound synthesis¹⁰, bioactivity modeling¹¹ or using image recognition¹² to analyze phenotypic screening data.

In fact, once you know what to look for, profitable AI opportunities are hiding in plain sight everywhere inside of a large life sciences company. “Isn’t this just process automation,” one might ask? Yes, in a sense, but it is automation (and to be more accurate, often semi-automation) of a specific kind of process: *cognitive* business processes that require decision-making under conditions of uncertainty.

DOING THE RIGHT THING

Choosing the right business application areas for initial focus is the starting point for effective executive management of AI. In theory, starting with a list of new technologies and then rigorously determining the value of applying them to various business challenges should get to roughly the same place as starting from business problems and quantifying the business case for applying new technologies against them. However, in practice we have found companies are much better off starting with the business problems, primarily because evaluating technical feasibility is a far more delegatable task than judging where the profit opportunities sit in a business.

As a general guideline, senior executives should identify a short list of core repeated decision processes with high-profit leverage that would be improved with better data utilization. We have rarely found them to be wrong about this. The work of the staff is then to estimate the value-at-stake for each process that is addressable with the AI technology of today, and not the potential technology of five years from now.

Every organization is unique, but here are two examples of life sciences business applications that we have often found to be good starting points for investigation:

I Clinical Trials

Over the past two decades, large consumer businesses have pioneered analytical processes designed to maximize the value created by iterative, randomized experiments. Analogous AI-based processes can now be applied to the life sciences domain to drive breakthrough insights, in areas such as:

- Building and applying prediction models to large databases to predict enrollment, expected baseline progression and attrition propensities by potential participant
- Automated extraction of patient context and dosage characteristics from earlier trials and observational data to maximize likelihood of success of subsequent trials
- Integration of numerous external datasets to increase model performance
- Automatic identification of test groups and control populations (where placebo effects can vary predictably) to maximize reliable intervention impact

2 Go-To-Market

New AI-driven prediction models, combined with clinical data and careful extraction and analysis of provider and consumer data, can be used to optimize:

- In-person and digital detailing efforts
- Personalized DTC advertising
- Speaker programs and other customized impact initiatives

DOING THINGS RIGHT

In addition to choosing the right business application areas for AI initiatives, executives need to make sure the conditions are in place for these initiatives to succeed. Beyond the normal challenges of project management, there are specific issues associated with AI projects in life sciences. We believe that a short list of principles will help to navigate them. These are not the principles you will often hear quoted. In fact, we believe that many data analytics truisms can be extremely misleading when applied to AI projects. In each case, we will argue for a replacement to a widely-quoted slogan.

CONVENTIONAL WISDOM

Garbage-In-Garbage-Out

FOUNDRY PRINCIPLE

Data Quality is Defined by
Fitness-for-Use

Everybody knows that data quality, consistency and relevance are enormous issues in the healthcare sector, even more so than in other parts of the economy. When evaluating whether data is good enough in the context of an effort to improve a repeated decision process, the criterion should always be the same: Does the stream

of decisions made with this data create better net financial and operational outcomes than the stream of decisions made without it? If and only if the answer is ‘yes,’ then the data is good enough for use.

And its value is not measured by anything inherent to the data, but instead in the difference in outcomes it yields. One should think of data like a hedge fund does. The key consideration is not whether it creates improvements in outcomes versus no data, but rather whether it provides incremental value to the data one already has. This is a very practical issue in life sciences. For example, the electronic health record (EHR) can provide enormous predictive information about patients, but we have repeatedly found that if an organization already has comprehensive claims data about a patient, then the traditional heart of the EHR – encounters, notes, diagnoses and procedures – often provides surprisingly little incremental predictive power. The benefit created by the EHR in such cases usually arises from a combination of its timeliness (often it will arrive 90 days or more in advance of claims data for some organizations), and in the “other” parts of the record, such as labs and social indicators.

CONVENTIONAL WISDOM

AI Requires Big Data

FOUNDRY PRINCIPLE

Data Diversity Beats Data Volume

It is often claimed that training AI algorithms requires very large datasets. This was true several years ago, but today it is sometimes true and sometimes not. We have repeatedly built patient-level prediction models with state-of-the-art (SOTA) or beyond-SOTA accuracy for topics such as predicting onset of specific diseases, forecasting expected next-year costs, and projecting the impact of a business or therapeutic intervention with just thousands to tens-of-thousands of patient records.

What is almost always true, however, is that integrating new relevant classes of data

dramatically increases predictive accuracy. These can range extremely widely: from healthcare datasets such as EHRs, claims, genomic and medical imagery / sensor data, to other consumer-oriented datasets such as browsing history, neighborhood demographics, travel patterns, writing samples and smartphone app usage.

CONVENTIONAL WISDOM

More Data Beats Better Algorithms

AI practitioners often want to focus on the analytically-fascinating task of finding patterns in numerical data. And pattern-finding models today are much more sophisticated than they were even five years ago. But this is now a broadly commoditized area, and it is very difficult to create material competitive advantage purely by being a better modeler.

FOUNDRY PRINCIPLE

The AI is in the Data Layer

The biggest advantage created by AI is that it allows us to incorporate many more kinds of data. The white-hot center of current AI excitement is around Deep Learning. From an executive perspective, the killer application of Deep Learning is that it lets us create data out of things we did not historically consider data, such as pictures, sounds and sensor data. This is why we can, in practice, now incorporate x-rays, cat scans, EEG output and many other sources of insight into our models. Natural Language Processing allows us to do the same for words such as medical notes and other texts. Related methods allow this for genomic data. Advanced search-and-filter methods enable us to incorporate much more consumer data than ever before. And so on. AI is distributed throughout this whole process, not something we snap on top of our pre-existing view of what we mean by “our data.”

CONCLUSION

AI visionaries have spent years making ambitious claims about the potential transformative impact of AI on life sciences, but most of the real successes have resulted from using AI to improve less glamorous operational processes. We have seen executives who follow the guidelines in this paper successfully capture this opportunity, creating significant value within surprisingly short timelines. We hope these concepts help you in this important domain and more broadly as you consider how to drive focused, practical, and profitable AI initiatives within your organization.

SOURCES

- 1 Nick Paul Taylor, "Healthcare AI funding hits new high as sector matures," MedTech Dive, <https://www.medtechdive.com/news/healthcare-ai-funding-hits-new-high-as-sector-matures/560396>, (August 7, 2019)
- 2 Lydia Ramsey, "'We like to think of ourselves as the lead turtle in the race of the turtles': How Big Pharma is turning to Silicon Valley to supercharge drug development," Insider, <https://www.insider.com/novartis-jay-bradner-on-silicon-valley-intersecting-with-pharma-2018-6>, (June 9, 2018)
- 3 Simon Smith, "150 Startups Using Artificial Intelligence in Drug Discovery," BenchSci, <https://blog.benchsci.com/startups-using-artificial-intelligence-in-drug-discovery>, (September 3, 2019)
- 4 Joseph Walker, Peter Loftus, and Brianna Abbott, "CEOs in Health Care Discuss Challenges of Working With Artificial Intelligence," Wall Street Journal, <https://www.wsj.com/articles/ceos-in-health-care-discuss-challenges-of-working-with-artificial-intelligence-11556724253>, (May 1, 2019)
- 5 Joseph Walker, Peter Loftus, and Brianna Abbott, "CEOs in Health Care Discuss Challenges of Working With Artificial Intelligence," Wall Street Journal, <https://www.wsj.com/articles/ceos-in-health-care-discuss-challenges-of-working-with-artificial-intelligence-11556724253>, (May 1, 2019)
- 6 Emma Court, "A startup using AI to find new treatments for rare diseases just raised \$121 million. Here's why its CEO still has a dire prediction for the industry." Business Insider, <https://www.businessinsider.com/recursion-pharma-ceo-chris-gibson-on-ai-finding-new-drugs-2019-7>, (July 15, 2019)
- 7 Derek Lowe, "Farewell to 'Watson For Drug Discovery'," Science Translational Medicine, <https://blogs.sciencemag.org/pipeline/archives/2019/04/18/farewell-to-watson-for-drug-discovery>, (April 18, 2019)
- 8 Andreessen Horowitz, "The Science and Business of Innovative Medicines," a16z Podcast, <https://a16z.com/2019/01/13/pharma-business-innovation-medicine-next-therapeutics>, (January 13, 2019)
- 9 Marwin H. S. Segler, Mike Preuss, and Mark P. Waller, "Planning chemical syntheses with deep neural networks and symbolic AI," Nature: Volume 555, pages 604–610, <https://www.nature.com/articles/nature25978>, (March 29, 2018)
- 10 Connor W. Coley¹, Dale A. Thomas III, Justin A. M. Lummiss, Jonathan N. Jaworski, Christopher P. Breen, Victor Schultz, Travis Hart, Joshua S. Fishman, Luke Rogers, Hanyu Gao, Robert W. Hicklin, Pieter P. Plehiers, Joshua Byington, John S. Piotti, William H. Green, A. John Hart, Timothy F. Jamison, and Klavs F. Jensen, "A robotic platform for flow synthesis of organic compounds informed by AI planning," Science: Volume 365, Issue 6453, <https://science.sciencemag.org/content/365/6453/eaax1566>, (August 9, 2019)
- 11 Andreas Bender, "Why AI and Drug Discovery are no match made in heaven," DrugDiscovery.NET, <http://www.drugdiscovery.net/2019/08/14/why-ai-and-drug-discovery-are-no-match-made-in-heaven>, (August 14, 2019)
- 12 Ben T. Grys, Dara S. Lo, Nil Sahin, Oren Z. Kraus, Quaid Morris, Charles Boone, and Brenda J. Andrews, "Machine learning and computer vision approaches for phenotypic profiling," JCB Journal of Cell Biology, <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5223612>, (January 2, 2017)



Foundry.ai

1920 L St NW, Suite 800
Washington, DC 20036

www.foundry.ai