

ARTIFICIAL SCIENCE

How *in-silico* minds will accelerate science

By Paul Cohen

The technology of science is as old as science, itself. Astrolabes and microscopes, image-displaying tachistoscopes and gene sequencers, mass spectrometers and atomic clocks, molecular tweezers, and the gene-editing tool CRISPR have all accelerated science by making the invisible visible and the uncontrolled do our bidding. Compared with these physical technologies, information technologies for science do different work: They store and manipulate data, they automate tedious tasks, they facilitate replicability and sharing, they learn, and one day they will assist scientists as ably as a good graduate student.


Science has been a frontier for artificial intelligence (AI) since the 1970s. The prevailing view was that AI should do what our most notable scientists do, so autonomous scientific discovery was the highest prize, and the prevailing methods of AI – various forms of generate-and-test – were presented as theories of the scientific discovery process¹. For example, the DARPA-sponsored DENDRAL project generated graph models of molecules in organic chemistry and tested them against mass spectrometer data, and a later program called CONGEN was able to discover all chemical graphs that satisfy empirical constraints². At roughly the same time, other researchers applied rudimentary data-mining heuristics to re-discover Kepler's Law and other physical laws in data³.

If these early efforts failed to convince us that AI could “do science,” it was probably because they did so little of what scientists do: They didn't read the literature, go to seminars, discuss theories with colleagues, prepare samples, design and run experiments, clean noisy data, or test hypotheses. They focused on the “aha moment” of discovery, not on the daily work of science.


Even now, the heroic theory of science, which holds that scientific discovery is the product of individual genius, influences discussions of AI approaches to science. For example, in 2014 Hiroaki Kitano, whose several research affiliations include president and CEO of Sony Computer Science Laboratories, proposed “a new grand challenge for AI: to develop an AI system that can make major scientific discoveries in biomedical sciences and that is worthy of a Nobel Prize.”⁴ In fact, to date no AI system can recognize a significant scientific result as such, and all “scientific discovery” systems are carefully managed by humans.

Instead of asking how humans can build AI machines capable of making scientific discoveries, it might be more productive to ask how machines can facilitate scientific discovery by humans⁵. This was the premise of DARPA's Big Mechanism and World Modelers programs, which I had designed in my recent stint as a DARPA program manager to investigate if and how AI can accelerate science as well as create new opportunities for AI research.

The Big Mechanism program was designed to develop technology to help humans build causal models of complicated systems. The program



The Big Mechanism program aims to develop technology to read research abstracts and papers to extract pieces of causal mechanisms, assemble these pieces into more complete causal models, and reason over these models to produce explanations. The domain of the program is cancer biology with an emphasis on signaling pathways.



Joshua Lederberg inaugurates the Stanford University Medical Experimental Computer (SUMEX), a computer designed to encourage the application of artificial intelligence in medicine and to provide a computer facility for sharing and interaction among researchers nationwide. Lederberg was one of those who worked on the DENDRAL project, an early effort to use artificial intelligence to further the scientific process. INSET: Cover sheet of the Proposal for Continuation of the Stanford Artificial Intelligence Project and the Heuristic DENDRAL Project (June 1969).

focused on the complicated molecular interactions in cells that, when they go wrong, result in cancer. Cell-signaling pathways are sequences of protein-protein interactions that transmit information to the cell nucleus and determine cell fate. The literature on cell signaling is vast, and each paper describes just a few signaling interactions. So the Big Mechanism program developed technology to help scientists explain the effects of drugs on entire pathways.

In an experiment in 2016, machines were able to explain 25 known drug-pathway interactions. Given a previously published model of 336 relevant genes (each of which encodes a protein) the machines used natural-

PROPOSAL FOR CONTINUATION OF
THE STANFORD ARTIFICIAL INTELLIGENCE PROJECT

JOHN MCCARTHY, Professor of Computer Science
Principal Investigator

ARTHUR SAMUEL, Senior Research Associate in Computer Science
Associate Investigator

and

THE HEURISTIC DENDRAL PROJECT

EDWARD FEIGENBAUM, Professor of Computer Science
Co-Principal Investigator

JOSHUA LEDERBERG, Professor of Genetics
Co-Principal Investigator

ARPA SD-183
JUNE 1969



Instead of asking how humans can build AI machines capable of making scientific discoveries, it might be more productive to ask how machines can facilitate scientific discovery by humans.

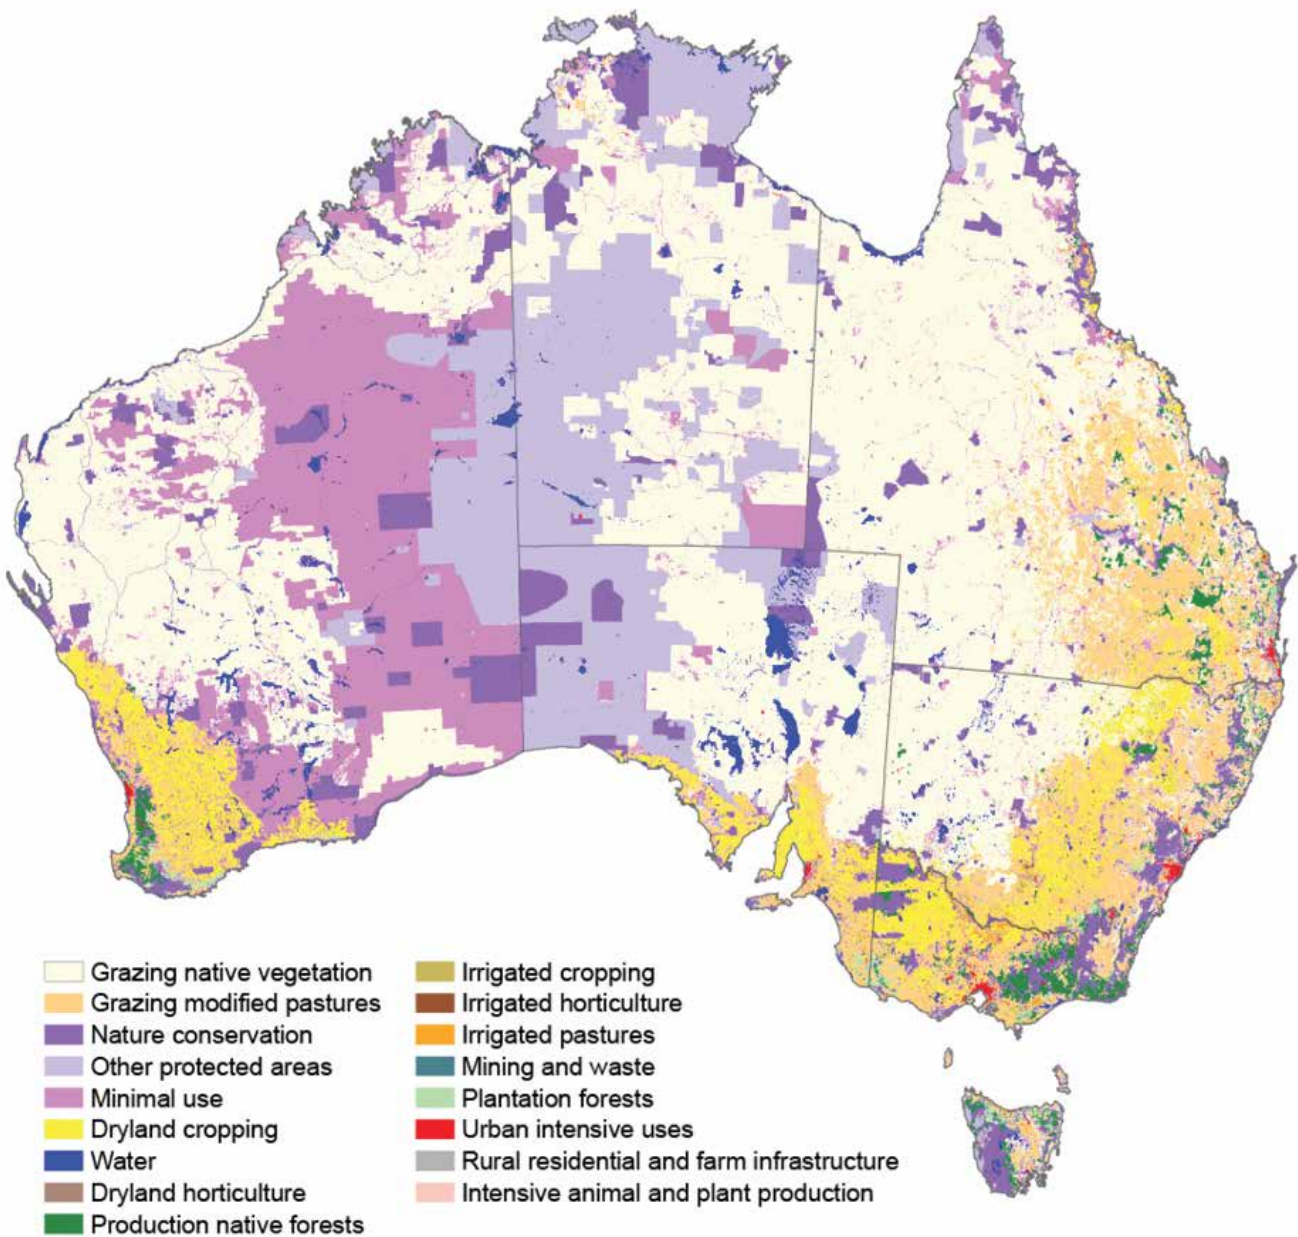


IMAGE COURTESY OF AUSTRALIAN GOVERNMENT DEPARTMENT OF AGRICULTURE AND WATER RESOURCES

language-understanding technologies to discover and read 95,000 journal articles, from which they extracted nearly a million causal assertions about protein-protein interactions. These were filtered and assembled into a single plausible signaling model that not only simulated the dynamics of protein concentrations but also explained all the drug-protein interactions. Using cloud-computing clusters, the whole process took less than a day.

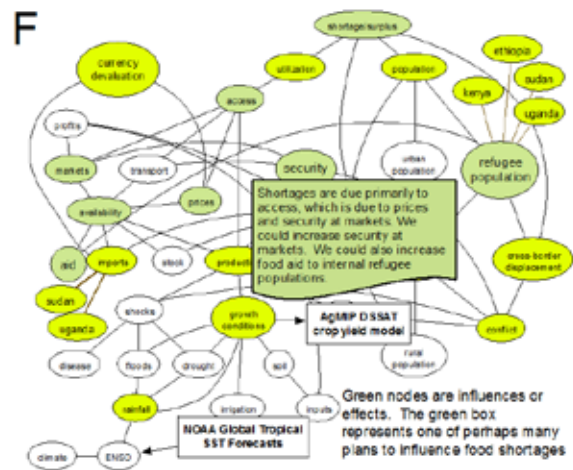
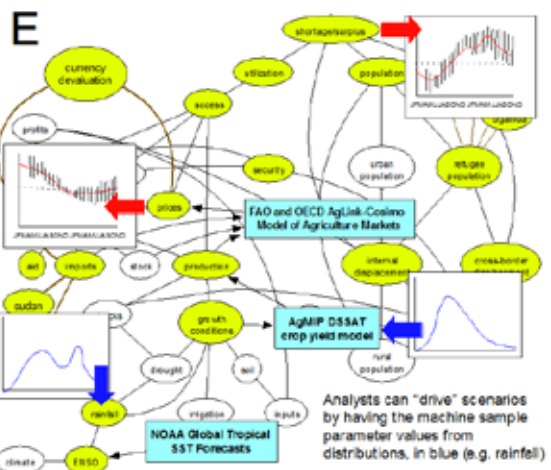
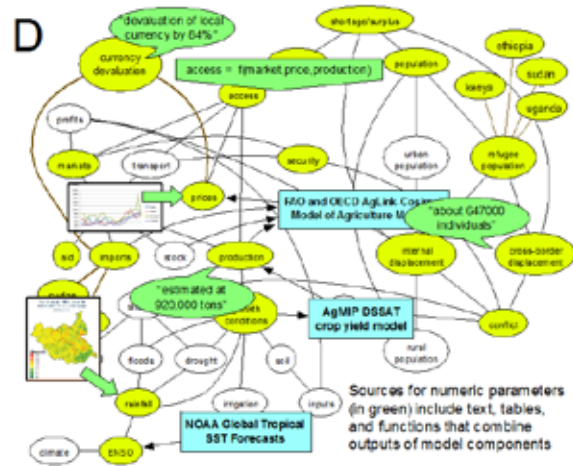
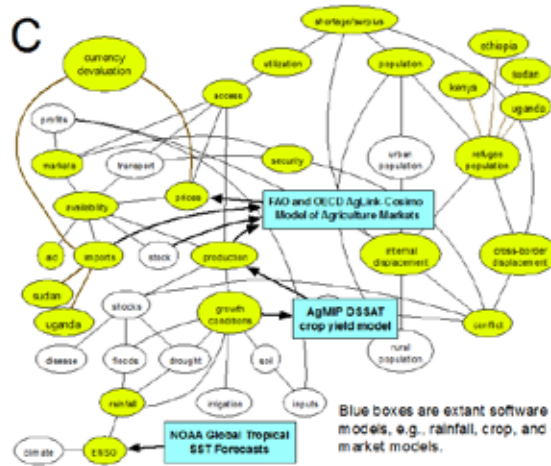
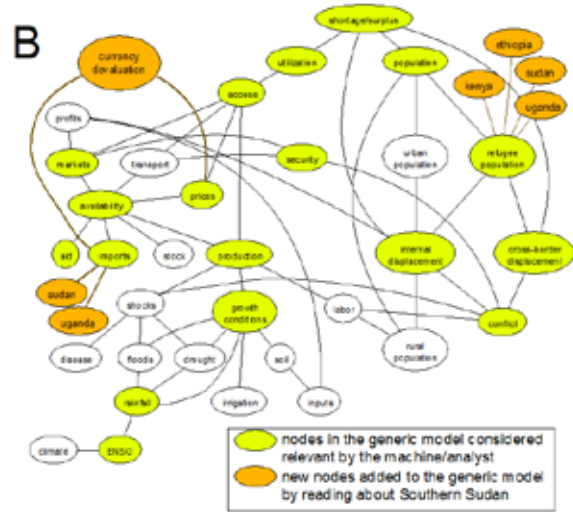
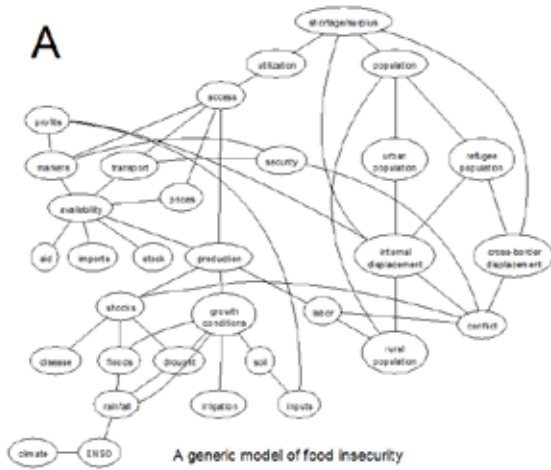
These results showed that machine reading and model-building can accelerate science in the sense that no human can do what the machines did: No one can read 95,000 journal articles or process a million assertions into a causal model that explains empirical results, even if they devoted years to the task. The machines worked alone, but researchers already are demonstrating more interactive versions of the technologies⁶.

Perhaps the most significant contribution of the Big Mechanism program has been to subvert the dominant paradigm of Big Data with its emphasis on exploitable correlations. All other things being equal, scientists prefer causal, explanatory models to opaque, predictive models based on correlations. The Big Mechanism program – its name a poke in the eye of Big Data – showed that machines could build causal, mechanistic models

The Australian government analyzed land use by integrating nine component models, but this effort took approximately one person-century to accomplish. Intelligent machine assistance could radically shorten the timeline for such a project.

of cellular processes such as tumorigenesis. One day, perhaps, these technologies will lead to machine-assisted hypotheses of how to interrupt or enhance cellular processes.

While causal knowledge is the highest prize in science, it is quite difficult to extract from data. Contrary to popular belief, it is ~~probably~~ possible to find causal relationships in correlational data, but the algorithms for doing so are computationally expensive and leave residual uncertainty about whether one thing truly causes another. In contrast, it is not very difficult to find assertions of causality in text. The linguistic constructions can be arcane (e.g., “mitogens stimulate cell division by inhibiting intracellular negative controls”), but researchers have been making steady progress toward extracting causality from text⁷. Eventually machines will read all of the causal assertions in the scientific literature and check them



against each other and against available data. Imagine machines that read millions of papers and find odd results that don't fit the zeitgeist or don't accord with data. Are they due to fraud or incompetence, or are they showing us something new and unexpected? Human scientists ask themselves these questions as they plod slowly through vast literatures; imagine how machines might accelerate the process.

DARPA's World Modelers program intends to support people in assembling large, complicated workflows of many models in order to more quickly gain insight into problems such as food and energy insecurity.

Food insecurity has many causes, from poor soil to political instability, from the El Niño cycle to economic migration. Scientists have been developing models of these individual elements for decades, but superhuman effort is required to link them together in a common software environment to study something like famine in sub-Saharan Africa. For example, the Australian government analyzed land use by integrating nine component models – of energy, water, and markets, among others – but this splendid effort took roughly one person-century of work⁸. The World Modelers program intends to accelerate this kind of analysis through intelligent machine assistance.

These examples suggest that AI will accelerate science in several ways, such as reading and assembling fragmentary results spread widely over literatures, integrating legacy models in common computational frameworks, automating *in silico* experiments, and even designing experiments and controlling the robots that carry them out. Much of this is “good old-fashioned AI,” not contemporary data science. At present, big data and machine learning play roles such as finding associations that might be causal (e.g., associations between genes and phenotypes) and learning computationally efficient approximations to expensive legacy models. But science depends on theories and data, and, importantly, on what people assert about theories and data in published literature. This suggests that future data-science technologies should expand their scope to embrace the interplay of theories, data, and literature.

For all its promise, AI has yet to recreate even the intellectual functions of a good research assistant. Nor is it likely to unless it tackles some problems that are getting in the way. As we review some of these, it might seem incredible that they haven’t been solved, but AI is a field in which seemingly easy things can be very difficult.

What are all those things and processes? Scientists refer to things and processes by names such as “p53” and “desertification.” In general, one thing can have many names and many things can have the same name, and even if things and names were in one-to-one correspondence, machines won’t necessarily know anything about what a name denotes. For example, “EBF” and “breast feeding” are names for very similar processes (“EBF” stands for exclusive breast feeding), but machines can’t know this unless they have access to dictionaries or ontologies that map names to formal descriptions of things. One might hope that someone has specified that “breast feeding” means, well, breast feeding, while EBF means breast feeding exclusive of other kinds of feeding. Ontologies record this kind of information and some fields, such as biology, have excellent ontologies for genes, proteins, drugs, and so on, but in general, scientific fields are poorly ontologized. In short, machines don’t know what scientists are talking about.

Coreference. Humans refer to the same things and processes in different ways, and machines have trouble figuring out coreferences, or repeated references to the same thing or processes. A paragraph might begin “The phosphorylation of ERK” and end with “active ERK,” but how are machines to know that phosphorylated ERK is often referred to as active ERK? How can machines know that “the first horse across the line” is “the winner”? The technology for dealing with coreference is improving, but every missed coreference is a lost opportunity for a machine to extend what it knows about a thing or process.

Semantic depth. Even when machines know about things and processes, they generally don’t know much. Consequently, they can’t answer questions that require semantic depth or nuance. For example, suppose it is important to estimate the per-capita income in a neighborhood, but there is no data

source that provides this information. A good research assistant would think, hmmm, let’s use sales records to find home prices and drone-based imagery to identify the numbers of cars in driveways, and let’s put it all together with public tax records to estimate household income. Machines can’t do this task (i.e., they can’t invent proxies for missing data) unless they know what “income” means and know that other data are proportional to income. The techniques for developing and exploiting meta-data (the things we would like machines to know about data) are improving, but they have yet to incorporate semantically deep enough nuance for machines to invent proxies.

Gritty Engineering. Models often have parameters that represent local conditions. Crop models, for example, need data about soil quality, sunshine, water, and other factors. A good research assistant might integrate a crop model with a soil model, a hydrological model, and weather model. The challenges would include understanding the parameters of the models well enough to use the output of one model as input to another, either directly or following some transformation. As noted, this understanding might require more semantic depth than machines have, but even when semantic issues are solved, gritty engineering issues remain: If the models are linked, meaning that feedback loops exist between the processes they represent, then they should run in a single computational environment. But this can be difficult if they run at different time scales or require very different amounts of computation. The technology of scientific workflows is progressing rapidly, but it isn’t yet possible for machines, rather than humans, to build complicated workflows of many computational models.

This sample of challenges should not discourage anyone from developing AI technology to enhance the work of scientists. Indeed, accelerating science is so important that it should motivate basic AI research on these and other challenges, as happened in the Big Mechanism and World Modelers programs mentioned earlier.

The challenges of our century are systemic, but humans have difficulty modeling and managing systems. Whether we’re modeling the molecular signaling pathways in cancer, the diverse factors contributing to food insecurity, or policies for land use in Australia, we find ourselves struggling with complexity. We have no choice about whether to recruit AI technology to scientific research: We must do it because we can’t understand complicated, interacting systems without help. It is an added benefit that the vision of AI-accelerated science will drive AI research itself for years to come.

1. Langley, P., Simon, H.A., Bradshaw, G.L., & Zytkow, J.M. (1993). *Scientific Discovery*. The MIT Press.
2. <https://profiles.nlm.nih.gov/BB/A/L/A/F/>
3. Op. Cit. Langley, et al,
4. Kitano, H. “Artificial Intelligence to Win the Nobel Prize and Beyond: Creating the Engine for Scientific Discovery.” *AI Magazine*, Spring, 2016.
5. Yolanda Gil, Mark Greaves, James Hendler, Haym Hirsh. “Amplify scientific discovery with artificial intelligence” *SCIENCE* 10 OCT 2014 : 171-172
6. Benjamin M. Gyori, John A. Bachman, Kartik Subramanian, Jeremy L. Muhlich, Lucian Galescu, and Peter K. Sorger. 11/2017. “From word models to executable models of signaling networks using automated assembly.” *Molecular Systems Biology*, 13, 11, Pp. 954. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5731347/>
7. Gus Hahn-Powell, Dane Bell, Marco A Valenzuela-Escárcega, Mihai Surdeanu. “This before that: Causal precedence in the biomedical domain.” <https://arxiv.org/pdf/1606.08089.pdf>
8. Australian National Outlook. https://www.csiro.au/~media/Major-initiatives/Australian-National-Outlook/CSIRO-MAIN_REPORT-National_Outlook_2015.pdf?la=en&hash=B7430AC9FD3C6665A53B9BDD1B2ED5356CB2FA3B