Modeling and Managing Complicated Systems

Paul Cohen

March 5, 2020

My name is Paul Cohen and I'm the founding dean of the School of Computing and Information at the University of Pittsburgh. Prior to joining Pitt in 2017, I spent four wonderful years at DARPA. As an AI researcher, I expected to work on common AI problems such as understanding human language but I discovered a new problem that has sustained me since – and, for what it's worth, is a killer app for AI:

> Our world is made of highly connected, complicated systems that work at very different spatial and temporal scales. But people specialize and literatures are huge, and data are incomplete and models are in various ways incompatible, so it's hard to get a complete, causal understanding of how complicated, interacting systems behave.

One day it occurred to me that humans might not be *able* to understand complicated, interacting systems without help. We read pretty slowly and we can only hold so many things in mind. Three ideas quickly followed:

- If humans cannot, machines must. This is stated as an imperative because the world's systems are under stress and I'm terrified that we don't know how to act.
- Humans *can* build models of complicated, interacting systems, it's just hugely expensive. Building AI machines to lower these costs is not easy, but it is an engineering problem whose cost can be amortized over countless large-scale modeling projects.
- For centuries, humans have pulled research results into our heads and synthesized understanding there, but this is expensive. Push science means *machines* pull research results, synthesize understanding, and push it out for humans to explore.

With these ideas I was able to design and implement a DARPA program called Big Mechanism. It brought together systems biologists, biomedical researchers and AI researchers to produce, for the first time, machines that can read journal articles and build systems biology models of cancer processes. My subsequent programs focused on human-machine collaboration (Communicating with Computer) and the general problem of integrating extant models into large ecosystems of models because most interesting problems arise from interactions of the world's complicated systems (World Modelers, managed by Joshua Elliott).

Big Mechanism and World Modelers both address the problem that scientific knowledge is fragmented and distributed. The reasons for this range from cognitive – no human can read 100,000 papers, or dig into the code of a dozen computational models – to sociological – we reward narrow expertise, not synthesis and comprehensive understanding.

Here I should address a posturing but pretty lightweight elephant in the room, big data and machine learning. Some of you are might be thinking that machine learning will somehow build comprehensive models of, say, cancer processes, Alzheimer's etiology, the effects of climate on food insecurity, the effects of new tax codes on small businesses, etc. I don't think so: First, with few exceptions, machine learning is associative – correlational – it doesn't establish causality. Second, machine learning generally associates *patterns*, it doesn't discover *mechanisms*. The job of science is to establish causal mechanisms, and machine learning generally doesn't do that.

Third, machine learning generally solves only two problems: classification and prediction. You can do a lot with machines that can reliably classify or predict, but you generally can't answer a short but humbling question: Why?

The inability of computers to answer "why" questions drove Big Mechanism and my other DARPA programs. Machine learning contributed almost nothing to these programs; in fact, I designed the programs to force AI researchers to look beyond data and correlation and start addressing causality and mechanism.

If we're serious about machines building causal, mechanistic models of the world, then machines should go to where causes and mechanisms are documented, and that's the scientific literature and existing models. Big Mechanism, World Modelers, and Joshua Elliott's ASKE program have all demonstrated that machines can read text, tables, equations and even FORTRAN code from legacy models, and build comprehensive models of the world's complicated, interacting systems. The technology must improve and it will. As it does, it enables a new kind of science, which I call Push Scholarship.

Today we practice Pull Scholarship: We sit in our monastic chambers, reading slowly, one document every few minutes or hours, and slowly construct in our minds coherent pictures of how the world works. It's a bit like doing a jigsaw puzzle where each journal article or dataset provides one piece. When we think pieces are missing we may go to the lab or the field and do the science and add the piece to those being pulled in by other scholars in other monastic cells. Occasionally, someone completes a region of the jigsaw puzzle and writes about it in a review article.

Pull science is incredibly inefficient. It works less and less well because there are simply too many results to synthesize in our heads. So we specialize. We don't read outside our areas, don't do interdisciplinary science, and potentially important results – which you pay for with your taxes and other costs – are lost. Truly lost: Most papers are read by just a handful of people. Papers in lower-ranked journals aren't read at all.

Push science, in contrast, is a new kind of scholarship where instead of pulling results into our heads, machines like those developed in Big Mechanism and World Modelers and ASKE create and maintain causal, mechanistic models, designed to answer "why" questions, and push them out to scientific communities. The quantum of scholarship becomes models, not individual research papers or datasets. These will be computational models, not textual descriptions in review articles, so scholars can test ideas in silico. This could change scholarship profoundly.

So now you've heard a bit of detail about two of the ideas I introduced earlier: Modeling is hugely expensive but machines can help, and the future of science is push science. I conclude with my mantra: If humans cannot, then machines must.

Here is the famous Drake equation, proposed by Frank Drake in 1961 to estimate the number of intelligent species in the universe:

$$N = R_* \cdot f_p \cdot n_e \cdot f_l \cdot f_i \cdot f_c \cdot L$$

Most of the terms concern star formation, planetary orbits, the ability to sustain life, and so on. The one that interest me is the last parameter, L, which denotes the lifespan of intelligent species. Drake recognized that we might not find intelligent species in the universe because they die out so quickly. They exist, but only for brief moments. It seems to me that a species that has the technological capability to mess up the systems upon which it depends for survival, but lacks models of those systems, survives only by extraordinary resilience and good luck. Perhaps we have the resilience, but in the immortal words of Dirty Harry, "You've got to ask yourself one question. Do I feel lucky? Well, do ya, punk?"