Modeling and Managing the World's Complicated Systems: A Challenge for AI
Paul Cohen

National Academy President's Circle Meeting, April 17, 2018

Close your eyes for a moment and think about yourself: Are you sleepy, hungry, grumpy…which of the seven dwarves are you right now? And why? Was your flight delayed? Did your lunch not agree with you? Are you distracted by your phone? Think about how you got here, how your lunch got here, how your email got here. Ponder for a moment how your experience at this instant depends on hundreds of interacting systems: your microbiome, the auditorium lighting, food supply chains, the weather, and so on.

Back in the old days, the sun would come up, you'd work your farm, talk with your neighbors, sell some crops…the web of systems was loose and slow to change. Today you are embedded in thousands of interacting systems that are changing very quickly, with very significant consequences. Climate is a major driver, but for a more immediate example consider the consequences of ill-considered shocks to our systems such as trade tariffs.[1]

I'm frankly worried that our systems are so big and their interactions are so consequential that we are flying blind in virtually all policy decisions. So here's a challenge for Artificial Intelligence: *Help humans to model and manage the world's complicated, interacting systems.* Let me repeat this, as it has become an urgent mission for me: AI can help humans to model and manage the systems on which we depend for our survival. This challenge will drive AI in directions it needs to go, it will accelerate science and engineering, and it might save our bacon. Let me elaborate.

The Australian government released a study of land use in 2015[2]. The work lashed together nine computational models of energy systems, water systems, industries and markets, under four different climate scenarios. The effort took one hundred person years: 40 people for 2.5 years. What does this tell me? You *can* model the interactions between complicated systems, it's just expensive. This is good news to an engineer like me.

So at DARPA I started two programs to develop technology to help humans model complicated systems. The first, called Big Mechanism, focused on the hellishly complicated molecular processes 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 we developed natural language processing systems to read journal articles, extract these interactions, gather them and assemble them into signaling pathways. Compared with machine-curated pathway models, *human*-curated ones are about 95% incomplete. In an experiment last year, machines were challenged to explain 25 drug-protein interactions. They started with 336 relevant genes, found and read 95K articles, extracted nearly a million assertions, filtered them and assembled them into signaling models sufficient to simulate the cellular dynamics *and* explained all 25 empirical results accurately. Using Amazon clusters the whole thing took less than a day.

So machine reading and modeling can accelerate science in the crude sense that no human can do what the machines did. None of us can read 95K articles or process a million assertions.

Buoyed by this success I started another DARPA program called World Modelers. Here the goal is not to build models de novo, as in Big Mechanism, but to assemble dozens of extant models into super-models, the way the Australians did when looking at land use. We have terrific models of crop yields, commodities markets, soil degradation, climate and weather, but we don't have the technology to easily assemble these into, say, super-models of food insecurity in South Sudan. The World Modelers program will integrate work on ontologies, natural language, workflow systems, and simulation to make it very easy for machines and people to assemble super-models.

I promised that modeling and managing the world's complicated, interacting systems would drive AI and science and engineering, generally, so let me say why.

First, modern science is ultraspecialized, but none of the consequential problems of this century fit neatly within an academic discipline. We must get specialists working together on consequential problems, and modeling the world's interacting systems is one way to do it. Second, we are trained to look downward at pieces of systems, not upward to systems themselves, much less to the interactions of systems. If we had technological support for modeling interacting systems, some people would do it, and new science and engineering would emerge. Third, from the get-go, my programs were about modeling causal influences and mechanisms. It's hard to manage systems if you can't be sure what's causal and what's not, and besides, scientists and engineers alike lust after explanations. But algorithms to extract causal models from Big Data aren't up to the task of modeling really complicated systems, nor are our ontologies, natural language capabilities, and many other AI technologies, so modeling and managing the world's complicated systems will keep AI busy for years to come.

Of course, we don't have to. We can continue on the current path, which is to find exploitable associations in data. Associations between drugs and symptoms, between what you bought last week and what you'll buy this week. There's knowledge to be mined from data and money to be made, but the pressing problems of the day are all *systemic*, so for my money the grand challenge to AI isn't to find the next exploitable correlation, it's to help us build causal models of the complicated systems that we depend on to survive.

---

[1] At the National Academy meeting, where there was much discussion of self-driving cars, I used a different example: A few weeks ago, crossing a long bridge on the way to work, I saw cars swerving into the other lane despite snow and sleet. The cause? A self-driving car doing the speed limit. However wonderful these cars are, we have to ask whether we'll be safer when 10% of the cars on the road are law-abiding and the rest have drivers like me. The point, of course, is that the effects of this technology are systemic, going far beyond the cars themselves.

[2] Australian National Outlook, https://www.csiro.au/nationaloutlook/

I left DARPA last summer to start the School of Computing and Information at the University of Pittsburgh. Not surprisingly, our mission is to develop the kinds technology I've been discussing. We're holding an inaugural conference on May 21-22, to which you are invited along with government agencies, foundations, NGOs, industry and other stakeholders. Check our website or email me. One of my friends called it a convening of geeks, wonks, quants and do-gooders, and that's right, because ultimately our greatest and most worthy *collective* challenge -- and I include AI in the collective -- is to *manage* this complicated world *for good*.