

Simple Rules for a Complex World with Artificial Intelligence

Jesús Fernández-Villaverde¹

April 28, 2020

¹University of Pennsylvania

A complex world with AI and ML

- It is hard to open the internet or browse a bookstore without finding an article or a book discussing artificial intelligence (AI) and, in particular, machine learning (ML).
- I will define some of these concepts more carefully later on.
- Many concerns: effects of AI and ML on jobs, wealth and income inequality, ...
- Some observers even claim that gone is the era of large armies and powerful navies. The nations with the best algorithms, as the Roman said in the Aeneid, “will rule mankind, and make the world obey.”


**AI
SUPER-
POWERS**
CHINA,
SILICON VALLEY,
AND THE
NEW WORLD ORDER
KAI-FU LEE

AI, ML, and public policy, I

- One thread that increasingly appears in these works, and which is bound to become a prominent concern, is how AI can address policy questions.
- Some proposals are mild.
- For example, thanks to ML:
 1. We might be able to detect early strains in the financial markets and allow regulatory agencies to respond before damage occurs.
 2. We might be able to fight electronic fraud and money laundering.
 3. We might be able to improve public health.
- I do not have any problem with these applications. Natural extension of models used in econometrics for decades.

- Some proposals are radical: “digital socialism” (Saros, 2014, Phillips and Rozworski, 2019, and Morozov, 2019).
- The defenders of this view argue that some form of AI or ML can finally achieve the promise of central planning: we will feed a computer (or a system of computers) all the relevant information and, thanks to ML (the details are often fuzzy), we will get an optimal allocation of goods and income.
- I am not arguing against a straw man: gathering attention from “establishment” figures.

Leigh Phillips & Michal Rozworski



PEOPLE'S REPUBLIC OF WALMART

HOW *the* WORLD'S BIGGEST
CORPORATIONS ARE LAYING *the*
FOUNDATION *for* SOCIALISM

Jack Ma, founder of Alibaba

“Over the past 100 years, we have come to believe that the market economy is the best system, but in my opinion, there will be a significant change in the next three decades, and the planned economy will become increasingly big. Why? Because with access to all kinds of data, we may be able to find the invisible hand of the market.

The planned economy I am talking about is not the same as the one used by the Soviet Union or at the beginning of the founding of the People’s Republic of China...

With the help of artificial intelligence or multiple intelligence, our perception of the world will be elevated to a new level. As such, big data will make the market smarter and make it possible to plan and predict market forces so as to allow us to finally achieve a planned economy.”



Benedetto Cristofani

Christmas Specials

Dec 18th 2019 edition >

Essay

Can technology plan economies and destroy democracy?

How algorithms could someday be used to optimise the ballot box

BEIJING AND SAN FRANCISCO

ABOUT A CENTURY ago, engineers created a new sort of space: the control room. Before then, things that needed control were controlled by people on the spot. But as district heating systems, railway networks, electric grids and the like grew more complex, it began to make sense to put the controls all in one place. Dials and light bulbs brought the way the world was working into the

Old idea in new clothes

- Many economists have made similar arguments in the past: Oskar Lange and Edmond Malinvaud.
- The System for Optimal Functioning of the Economy and the Central Economic-Mathematical Institute in the Soviet Union led by Nikolay Fedorenko.
- The Project Cybersyn in Chile led by Stafford Beer ([Medina, 2011](#)).

The Central Economic Mathematical Institute by Leonid Pavlov



Project Cybersyn



My argument, I

- While “digital socialism” might be a hipster thing to talk about in Williamsburg or Shoreditch, it is as much of a chimera as “analog socialism.”
- ML does not fix the problems of central planning.
- Why? Three increasingly more serious barriers:
 1. ML requires enormous datasets that are unlikely to exist in most cases of interest.
 2. Except in a few cases, ML suffers from the Lucas critique ([Lucas, 1976](#)).
 3. The allocation problem in a society is not how to solve an optimization problem, but to generate the right information when information is dispersed and agents do not have incentives (or even the capabilities) to disclose such information to a central planner ([Hayek, 1945](#)).

My argument, II

- Historically, we have organized our economic life around a set of simple rules:
 - Private ownership of most goods.
 - Freedom of contract.
 - Uniform and impersonal norms.
- These simple rules are subject to many caveats and limitations, but they create a market economy that allocates goods and income.
- Markets (and related mechanisms) are likely to remain the best course of action.
- I will argue that, in fact, these simple rules have evolved over time in a related manner to the ways a particular class of ML algorithms (reinforcement learning) works.

Three caveats

1. This is a discussion where I have “skin-in-the-game.” I do ML “for a living.”
2. AI and ML are enormously useful tools in many contexts. Nothing of what I say should be interpreted as implying that AI and ML cannot be applied to many problems of interest. In fact, as a European, I am quite worried the European Union is falling behind the U.S. and China in this field.
3. But it is crucial to frame the promises of AI realistically and avoid the disappointment that could come from overpromising. Neither AI nor public policy will be well served by hyperbole or careless vows.

Classical artificial intelligence

- Modern AI started in a famed two-month summer workshop at Dartmouth College in 1956: John McCarthy, Herbert Simon, and Marvin Minsky.
- Focused on symbolic reasoning, expert systems, and cognitive model approaches.
- Loosely speaking, we were trying to “replicate” human intelligence.
- We aimed at writing code that could pass the Turing test.
- Some sense of disappointment, with several “winters of discontent.”

Herbert Simon (1965). The Shape of Automation for Men and Management (p. 96)

“Machines will be capable, within twenty years, of doing any work a man can do.”





What is machine learning?

- Machine learning comes from a very different perspective.
- Set of algorithms to detect and learn from patterns in the data and use them for decision making or to forecast future realizations of random variables.
- Both French (“apprentissage automatique”) or Spanish (“aprendizaje automático”) capture this idea much better than English (although “statistical learning” can also be used).
- Deep learning and reinforcement learning are subsets of this more general class of algorithms.
- Focus on recursive processing of information to improve performance over time.
- Operational definition of learning that does not aim at passing Turing test.

Two-person games

- Think about the example of how to program a computer to play a two-person game: checkers, chess, go.
- Checkers (*le jeu de dames*) and chess are relatively easy to code: Chinook has “solved” checkers and Deep Blue defeated Garry Kasparov in 1996.
- Go is a completely different beast:
 1. Board with a 19x19 grid of lines, containing 361 points to place stones.
 2. The number of legal board positions is $\approx 2 \times 10^{170}$ (in comparison, checkers legal positions are $\approx 5 \times 10^{20}$, chess legal positions are $\approx 10^{43} - 10^{50}$, and the numbers of atoms in the university is $\approx 2 \times 10^{80}$).
- Around 2012, consensus was that it would take at least a decade to come up with software that could be competitive at the top level in Go.

Jonathan Schaeffer

One Jump Ahead

Computer Perfection at Checkers

Revised Edition


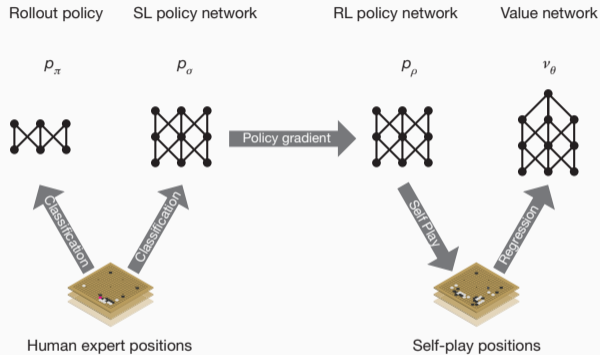
 Springer

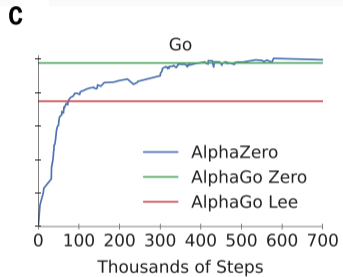
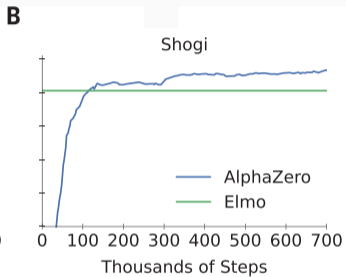
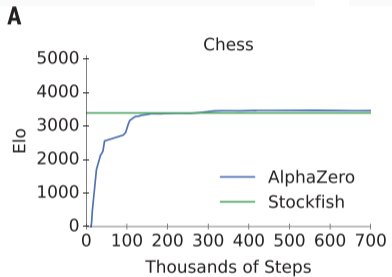




AlphaGo

- AlphaGo vs. Lee Sedol in March 2016.
- *Silver et al. (2018)*: now applied to chess, shogi, Go, and StarCraft II.
- Check also:
 1. <https://deepmind.com/research/alphago/>.
 2. <https://www.alphagomovie.com/>
 3. <https://deepmind.com/blog/article/alphastar-mastering-real-time-strategy-game-starcraft-ii>
- Very different than Chinook and Deep Blue.
- New and surprising strategies.
- However, you need to keep this accomplishment in perspective.





MACHINE LEARNING

ETHEM ALPAYDIN

010001010111010001101000
011001010110110100100000
010000010110110001110000
011000010111100101100100
011010010110111000001101
0000101001001100110101100001
011000110110100001101001
011011100110010100100000
010011000110010101100001
011100100110111001101001
011011100110011100001101
00001010010101010001101000
011001010010000001001110
011001010111011100100000
010000010100100100001101
00001010

THE MIT PRESS ESSENTIAL KNOWLEDGE SERIES

DEEP LEARNING

JOHN D. KELLEHER

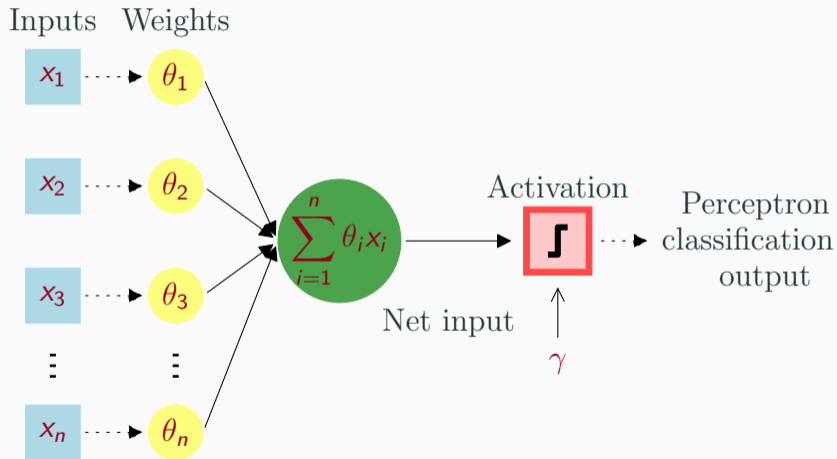


THE MIT PRESS ESSENTIAL KNOWLEDGE SERIES


Why now?

- Many of the ideas of machine learning (e.g., basic neural network by [McCulloch and Pitts, 1943](#), and perceptron by [Rosenblatt, 1958](#)) are decades old.
- Previous waves of excitement: late 1980s and early 1990s. Those decades were followed by a backlash.
- Three forces behind the revival:
 1. Big data.
 2. Cheap computational power.
 3. Algorithmic advances (influence of Google).
- Likely that these forces will become stronger over time.
- Exponential growth \Rightarrow plenty of packages for Python, R, and Julia.

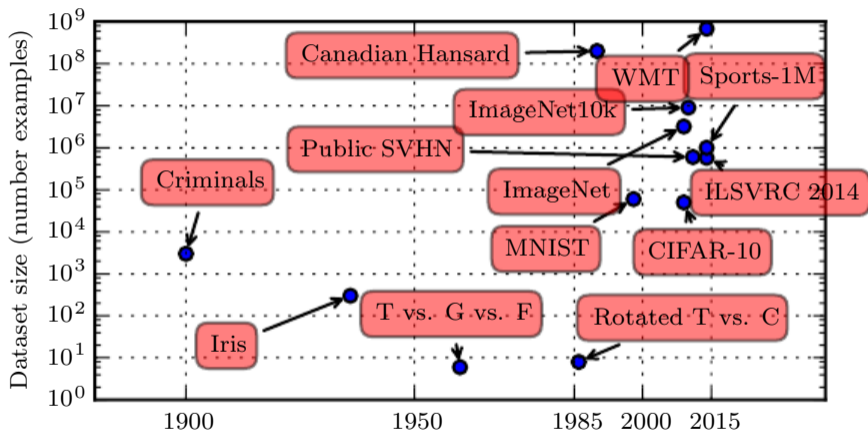
Flow representation




The biological analog



Data sizes



Number of transistors





TensorFlow

The data problem, I

- ML requires enormous datasets that are unlikely to exist in most cases of interest.
- For example, the rule-of-thumb in the industry is that one needs around 10^7 labeled observations to train a complex neural network properly.
- Walmart, Amazon, or Netflix, with access to millions of observations from their customers, have access to such datasets (but only enough for their problems!).
- In some situations, you can create your “own” data by simulation or experimentation.

The data problem, II

- Unfortunately, in many problems of interest in public policy, we do not have access to such a wealth of data and, most likely, we never will
- For instance, in macroeconomics, we usually have at most 291 data points (1947:Q1-2019:Q3).
- Often, not even that!
- Micro data does not help that much.
- In fact, the problems of Walmart, Amazon, and Netflix are, in some sense, radically simpler than the problems of running an economy.
- Think about the fundamental problem of trading-off the present vs. the future.

The Lucas critique

- Observed behavior comes from the decisions of agents under a set of rules.
- If you change the rules, agents decisions will change.
- Thus, lessons learned from (reduced-form) statistical models, such as ML, based on observed behavior are useless to forecast behavior under a new set of rules (there are a few exceptions).
- In general, this means that you cannot use ML to forecast the effect of many policies of interest (or that your ability to do so is limited).

A simple example

- In the U.S. 2016 presidential election, Hillary Clinton got 65,853,514 votes against Trump's 62,984,828.
- However Trump won the electoral college.
 1. Naive conclusion: a direct voting system would have elected Clinton.
 2. Lucas critique's answer: we do not know. Under an alternative set of rules, both candidates would have campaigned differently (location of rallies, TV spending, stump speech content, debate strategies, ...) and voters would have behaved differently (did you bother to vote if you were a Republican in California or a Democrat in Wyoming?).
- Or closer to home: will you vote in the same way if France changed the two-round electoral system for the first-past-the-post of the UK or the U.S.? Will France have the same political party system?
- My own experience on both sides of the electoral process (campaigning and voting).

- The problem of economics was never about computing an optimal solution to an allocation problem given some data.
- The problem is, and will always be, determining the preferences, abilities, and effort of the agents in a world where everyone has an incentive to misrepresent those preferences, abilities, and effort.
- ML can do very little to alleviate this problem.

Example: teaching matrix at Penn Economics, I

- Every year, we need to determine the teaching matrix at Penn Economics.
 1. Who should teach what?
 2. Who will develop new courses?
 3. At what level of quality?
- If I had access to the data, determining Penn Economics' teaching matrix is a simple problem. I can write a computer program to get to the optimum in a few seconds.
- The challenge is gathering the required information.

Example: teaching matrix at Penn Economics, II

- Every member of the faculty is asked about his/her preferences.
- Classes are allocated and there are some teaching evaluations filled by students.
- Every step is filled with lack of information:
 1. I do not have incentives to tell my true preferences (innovative, large class vs. old, small class).
 2. I do not have incentives to innovate in teaching.
 3. I do not have incentives to perform at the optimal level.
- While some mechanisms can be designed to fix some of the problems, they are probably not very robust in real life and, even if they were, they are difficult to scale to a whole society.

A lesson from economic history

- The trouble with the Soviet Union was not that their computers were not powerful enough (although they were not) or that its planning algorithms were poor (they were terrible).
- The problem with the Soviet Union was that central planning is inefficient *in essentia sua*.
- The evidence from economic history is overwhelming.
- *The Truth About Soviet Whaling: A Memoir* by Alfred Berzin: Soviet whalers killed at least 180,000 more whales than they reported between 1948 and 1973. Why?



A coda, I

- We did not come up with simple rules thanks to an enlightened legislator (or a blue-ribbon committee of academics “with a plan”).
- Simple rules were the product of an evolutionary process that appeared over centuries thanks to the decisions of thousands and thousands of agents:
 1. Roman law and its rediscovery in the Middle Ages in continental Europe and Scotland.
 2. Lex mercatoria across all Europe.
 3. Common law in England.

A coda, II

- Good law is nothing more than applied optimal mechanism design. The forces of evolution, through trial and error, took us over the optimal solution of such a mechanism design problem.
- This process is surprisingly similar to another area of AI, reinforcement learning (RL).
- The history of Western law and the simple rules that emerged from it is an example of decentralized RL. Jurists, through reasoning and experience, saw what worked and what did not. Those rules that led to Pareto improvements survived and thrived. Those who did not, dwindled.
- There is a sharp lesson from AI: our trust in simple rules has deeper roots than our high positivist era of the administrative state recognizes.