

A Paradigm Shift in Computer Science?

Moshe Y. Vardi

Rice University

Is Computer Science Fundamentally Changing?

Model-Driven Science vs Data-Driven Science

- We are at peak hype about Artificial Intelligence and Machine Learning!
- *Common perception*: A Kuhnian paradigm shift!
- **Paradigm Shift**: a fundamental change in the accepted concepts and practices of a scientific discipline – “Throw out the old, bring in the new!”
- **Example**: Geocentric model → Heliocentric model

Highly accurate protein-structure prediction with AlphaFold

- Proteins are essential to life, and understanding their structure can facilitate a mechanistic understanding of their function.
- Through an enormous experimental effort, the structures of around 100,000 unique proteins have been determined, but this represents a small fraction of the billions of known protein sequences.
- **AlphaFold**: a novel machine-learning approach that incorporates physical and biological knowledge about protein structure into the design of the deep-learning algorithm.
- AlphaFold demonstrated accuracy competitive with experimental structures in a majority of cases and greatly outperforming other methods.
- Nature – 2021, Nobel Prize in Chemistry – 2024

Vasant Dhar: The Paradigm Shifts in AI

CACM, Oct. 2024: The current general AI has breached a major barrier where machines can learn through self-supervision, transforming AI from an application to a general-purpose technology like electricity that will enable the rapid development of transformative applications across the economy.

- “For the first time, we can converse with an entity, however imperfectly, about anything, as we do with other humans.”

The Scaling Hypothesis

Rich Sutton, *The Bitter Lesson*, 2019: Applying large amounts of compute has been a more successful than leveraging human knowledge to design algorithms specific to these tasks.

- **GPT-2**: 17M parameters
- **GPT-3**: 175B parameters
- Around 5Q parameters: AGI!

But:

- **Nov. 13, 2024** – “Bloomberg reports that OpenAI, Google, and Anthropic are having trouble making better models, seeing less progress despite higher costs.”
- **Nov. 15, 2024** – Reuters reports that OpenAI and others seek new path to smarter AI as current methods hit limitations.

The End of Scaling Age

Ilya Sutskever, Nov. 2024 (ousted OpenAI co-founder): results from scaling up pre-training - the phase of training an AI model that use s a vast amount of unlabeled data to understand language patterns and structures - have plateaued.

- “The 2010s were the age of scaling, now we’re back in the age of wonder and discovery once again. Everyone is looking for the next thing,” Sutskever said. “Scaling the *right* thing matters more now than ever.”

Paradigm Shift or Paradigm Glide?

- The progress in ML is revolutionary indeed!
- **But:** The paradigm-shift paradigm has been exaggerated!
- *In reality:* New scientific theories *refine* old ones.
 - After all, we went to the moon with Newtonian Mechanics!
- *My Thesis:* Data-driven science *refines* model-driven science, but does not replace it!

“Understanding the Limitations of Mathematical Reasoning in Large Language Models”

Oct. 7, 2024: Mehrdad Farajtabar et al.: “We found no evidence of formal reasoning in language models . Their behavior is better explained by sophisticated pattern matching – so fragile, in fact, that changing names can alter results by 10%!”

Conversation with ChatGPT o1-preview:

- Q: How many times in my life can I subtract 5 from 35?
- A: You can subtract 5 from 35 only once in your life. ... After you subtract 5 from 35 the number is no longer 35 but 30.

Logic vs. Machine Learning

Daniel Kahneman, *Thinking, Fast and Slow*, 2011:

- **Machine Learning**: fast thinking, e.g., “Is this a stop sign?”
- **Logic**: slow thinking, e.g., “Should you stop at a stop sign?”
 - Logic is about language and meaning!

Example—Autonomous Vehicles: how to establish safety? [Shashua, '17]

- *Data Driven*: Drive 1B miles!
- *Data+Model Driven*: Combine data (1M miles) with reasoning.

Regulating Automated Decisions Systems

James Larus and Chris Hankin, CACM, Aug. 2018: “The widespread adoption of AD systems will be economically disruptive and will raise new and complex societal challenges. ... Disdain for regulation is pervasive throughout the tech industry. In the case of automated decision making, this attitude is mistaken.”

How to Regulate?

- Fairness, accountability, transparency, *explainability*

Explainable AI: AI in which the results of the solution can be understood by humans – explainable AI is a must for *human-centered AI*.

Corollary: Explanation must be high level, rather than low level – *logic!*

Boole's Symbolic Logic

Boole's insight: Aristotle's syllogisms are about *classes* of objects, which can be treated *algebraically*.

“If an adjective, as ‘good’, is employed as a term of description, let us represent by a letter, as y , all things to which the description ‘good’ is applicable, i.e., ‘all good things’, or the class of ‘good things’. Let it further be agreed that by the combination xy shall be represented that class of things to which the name or description represented by x and y are simultaneously applicable. Thus, if x alone stands for ‘white’ things and y for ‘sheep’, let xy stand for ‘white sheep’.

Boolean Satisfiability

Boolean Satisfiability (SAT); Given a Boolean expression, using “and” ($\wedge$) “or”, ($\vee$) and “not” ($\neg$), *is there a satisfying solution* (an assignment of 0’s and 1’s to the variables that makes the expression equal 1)?

Example:

$$(\neg x_1 \vee x_2 \vee x_3) \wedge (\neg x_2 \vee \neg x_3 \vee x_4) \wedge (x_3 \vee x_1 \vee x_4)$$

Solution: $x_1 = 0, x_2 = 0, x_3 = 1, x_4 = 1$

Complexity of Boolean Reasoning

History:

- William Stanley Jevons, 1835-1882: “I have given much attention, therefore, to lessening both the manual and mental labour of the process, and I shall describe several devices which may be adopted for saving trouble and risk of mistake.”
- Ernst Schröder, 1841-1902: “Getting a handle on the consequences of any premises, or at least the fastest method for obtaining these consequences, seems to me to be one of the noblest, if not the ultimate goal of mathematics and logic.”
- Cook, 1971, Levin, 1973: Boolean Satisfiability is NP-complete.

Algorithmic Boolean Reasoning: Early History

- Newell, Shaw, and Simon, 1955: “Logic Theorist”
- Davis and Putnam, 1958: “Computational Methods in The Propositional calculus”, unpublished report to the NSA
- Davis and Putnam, JACM 1960: “A Computing procedure for quantification theory”
- Davis, Logemman, and Loveland, CACM 1962: “A machine program for theorem proving”

DPLL Method: Propositional Satisfiability Test

- Convert formula to conjunctive normal form (CNF)
- Backtracking search for satisfying truth assignment
- Unit-clause preference

The SAT Revolution

CDCL = conflict-driven clause learning


- Backjumping
- Smart unit-clause preference
- Conflict-driven clause learning
- Smart choice heuristic (brainiac vs speed demon)
- Restarts

Key Tools: GRASP, 1996; Chaff, 2001

Current capacity: *millions* of variables

Some Experience with SAT Solving

Speed-up of 2012 solver over other solvers



Knuth Gets His Satisfaction

SIAM News, July 26, 2016: “Knuth Gives Satisfaction in SIAM von Neumann Lecture”

Donald Knuth gave the 2016 John von Neumann lecture at the SIAM Annual Meeting. The von Neumann lecture is SIAM's most prestigious prize.

Knuth based the lecture, titled “Satisfiability and Combinatorics”, on the latest part (Volume 4, Fascicle 6) of his The Art of Computer Programming book series. He showed us the first page of the fascicle, aptly illustrated with the quote “I can't get no satisfaction,” from the Rolling Stones. In the preface of the fascicle Knuth says “The story of satisfiability is the tale of a triumph of software engineering, blended with rich doses of beautiful mathematics”.

Applications of SAT Solving in SW Engineering

Leonardo De Moura+Nikolaj Björner, 2012: Applications of Z3 at Microsoft

- Symbolic execution
- Model checking
- Static analysis
- Model-based design
- ...

Fact: SAT solving is today an industrial reality!

Grand Challenge: Combine logic with machine learning – *neurosymbolic reasoning*.

Energy Footprint of ML

October 2024:

- “Hungry for energy, Amazon, Google and Microsoft turn to nuclear power” – Large technology companies are investing billions of dollars in nuclear energy as an emissions-free source of electricity for artificial intelligence.
- “Former Google CEO Eric Schmidt says energy demand for AI is infinite and we are never going to meet our climate goals anyway, so we may as well bet on building AI to solve the problem.”

Fundamental Question: Can we reduce the energy footprint of Machine Learning?

Accelerating large-scale neural network training on CPUs

AWS Machine Learning Blog, Feb. 2024:

We have developed a sparse deep-learning engine, known as BOLT, that is specifically designed for training and deploying models on standard CPU hardware as opposed to costly and energy-intensive accelerators like GPUs. Many of our customers have reported strong satisfaction with ThirdAI's ability to train and deploy deep learning models for critical business problems on cost-effective CPU infrastructure.

At ThirdAI, we achieve these breakthroughs in efficient neural-network training on CPUs through proprietary dynamic sparse algorithms that activate only a subset of neurons for a given input, thereby side-stepping the need for full dense computations. ThirdAI uses *locality-sensitive hashing* to dynamically select neurons for a given input. In certain cases, our sparse CPU-based models train faster than the comparable dense architecture on GPUs.

Hashing algorithms for large-scale learning

P. Li, A. Shrivastava, J. Moore, and A. König: Neurips 2011

Minwise hashing is a standard technique in the context of search for efficiently computing set similarities. The recent development of b-bit minwise hashing provides a substantial improvement by storing only the lowest b bits of each hashed value. In this paper, we demonstrate that b-bit minwise hashing can be naturally integrated with linear learning algorithms such as linear SVM and logistic regression, to solve large-scale and high-dimensional statistical learning tasks, especially when the data do not fit in memory.

Crux: It is a classical algorithmic paper!

From Model-Driven Computer Science to Data-Driven Computer Science and Back

In Summary:

- It is a *paradigm glide*, not *paradigm shift*.
- Data-driven CS *refines* model-driven CS, it does *not* replace it.
- Physicists still teach Mechanics, Electromagnetism, and Optics.
- But we must bridge the gap between machine learning and logic to get *human-centered AI*.

The Real Paradigm Shift

Research Budgets

- NSF Computer and Information Science and Engineering Directorate: about \$1.5B in 2024
- NSF funds about 80% of US academic CS research
- April 2024: “Google will spend more than \$100B on AI, exec says”

Crux: The center of gravity of computing research has moved decisively to Big Tech.

Does it matter where research is done?

Purpose of Research:

- AAUP 1940 statement on academic freedom: “Institutions of higher education are conducted for the common good.”
- ACM’s Code of Ethics: “Computing professionals actions change the world. To act responsibly, they should reflect upon the wider impacts of their work, consistently supporting the public good.”

For Big Tech – *Profits > CommonGood*

- MYV, 01/24: “Computing, You Have Blood on Your Hands!”

The Open AI Saga

The plot:

- 2015: Founded as a nonprofit.
- 2019: Established for-profit subsidiary – to attract investment – with nonprofit in control.
- 12/2022: launched a free preview of ChatGPT, a new AI chatbot based on GPT-3.5.
- 17/11/2023: Board fired CEO Sam Altman out of concerns that he was more focused on growth rather than safety.
- Huge pressure from investors – Altman was reinstated five days later.

It goes beyond ethics!

Google Authors: A graph placement methodology for fast chip design, Nature, 2021 – DRL beats simulated annealing

- > 600 citations

Benchmarks: “previous generation of TPUs”

I. Markov, Google’s Reinforcement Learning for IC Macro Placement, CACM 2024: “Crosschecked data indicate that the integrity of the Nature paper is substantially undermined owing to errors in conduct, analysis, and reporting. Nature has been slow to enforce its own policies.”

In Conclusion

- Progress in ML has been revolutionary.
- But so has been progress in automated reasoning.
- The reports of the demise of formal CS are greatly exaggerated.
- But the center of gravity of computing research moved to Big Tech.
- And that is a problem for all of us.