frontiers | Frontiers in Physics

TYPE Opinion
PUBLISHED 14 October 2024
DOI 10.3389/fphy.2024.1484701

OPEN ACCESS

EDITED BY
Alex Hansen,
NTNU, Norway

REVIEWED BY
Chuandong Lin,
Sun Yat-sen University, China

*CORRESPONDENCE
Sauro Succi,
sauro.succi@iit.it

RECEIVED 22 August 2024
ACCEPTED 11 September 2024
PUBLISHED 14 October 2024

CITATION
Succi S (2024) Chatbots and zero sales resistence.
Front. Phys. 12:1484701.
doi: 10.3389/fphy.2024.1484701

Chatbots and zero sales resistence

Sauro Succi[1,2,3]*

[1]Italian Institute of Technology (IIT), Rome, Italy, [2]Physics Department, Harvard University, Cambridge, MA, United States, [3]Mechanical Engineering Department, University College London, London, United Kingdom

KEYWORDS

artificial neural network-ANN, artificial intelligence, machine learning, energy saving, complex systems, modeling, computer simulation

Not a day goes by without we hear of the latest AI breakthroughs, such as chatbots that write up texts or generate images increasingly harder to tell apart from their human-made counterparts. These headlines come with a heavy load of hype, but even with hype factored out, a highly seductive promise stands tall, the promise to capture levels of complexity largely out of grasp for our best theories, models and simulations. Briefly, AI would supplant the time-honored Scientific Method, as we know it since Galileo's time [1, 2].

While heavily pumped up, this promise is not empty, addressing as it does, among others, one of the most vexing Achille's heels of the scientific method, the infamous Curse of Dimensionality (CoD) [3]. Indeed, CoD compounds with a profound hallmark of Complexity, namely, the fact that complex systems are *sneaky*: they inhabit ultra-dimensional spaces but do not fill them up [4–6]. To the contrary, "interesting things" take place in ultrathin and often highly scattered portions of the huge state space available to them. Nature likes to play hide and seek and big time so. An illuminating example can be found in the book of Frenkel and Smit [7], where we learn that the chance of making a sensible Monte Carlo move in the state space of hundred hard-spheres (please note, hundred, not Avogadro's) is about 10^{-260}! The golden nuggets are well hidden indeed.

Computational science has devised a number of clever techniques to visit the regions hosting these preciously rare golden nuggets without waiting many ages of the Universe [8]. Yet, the CoD still remains a very tough cookie for the scientific method to the present day.

Artificial Intelligence, mostly powered by Machine Learning (ML), promises a new and unprecedently powerful angle of attack to Complexity in science and society. And again, the promise is largely overblown but not empty, as witnessed by a number of success stories: chess and GO winnings, self-driving cars, DeepFold mapping of protein structure, stand out as some the most spectular (ized) cases in point [9].

It is worth discussing where this "magic" comes from in some little more detail.

The basic idea of ML is to represent a given D-dimensional output y (target) through the recursive application of a simple non linear map [10]. For a neural network (NN) consisting of an input layer x, L hidden layers $z_1 \dots z_L$, each containing N neurons, and an output layer y, the update chain $x \to z_1 \cdots \to z_L \to y$ reads symbolically as follows:

$$x = input \tag{1}$$

$$z_1 = f(W_1 x - b_1), \dots z_L = f(W_L z_{L-1} - b_L) \tag{2}$$

$$y = f(W_{L+1} z_L - b_{L+1}) \tag{3}$$

Where W_l are $N \times N$ matrices of weights, b_l are N-dimensional arrays of biases and f is a nonlinear activation function, to be chosen out of a large palette of options. The output y is then compared with a given training target y_T ("Truth") and the weights are recursively updated in such a way as to minimize the discrepancy between y and y_T (Loss function), up to the desired tolerance. This latter task is pursued by changing the weights along the direction of maximum change of the Loss function. In equations

$$W'_{ij} = W_{ij} - \gamma \frac{\partial \mathcal{L}}{\partial W_{ij}} \tag{4}$$

where $\mathcal{L}[W]$ is the loss function, which depends on the full set of weights W, and γ is a numerical parameter in control of the convergence of the overall process.

The idea is that with enough (big) data for training, the combination of Equations 1–3 and Equation 4 can reach *any* target, with no need of any model/theory aimed at capturing the causal structure of the problem at hand. Whence the alleged demise of the scientific method [1, 2].

Put down in such simplistic and bombastic terms, the idea is readily debunked, based on well-known properties of complex systems [11, 12]. Yet, it is true that neural nets prove sometimes capable of representing "sneaky" functions in hyperdimensional spaces which would be extremely hard to attain by any other method.

So, where does such magic come from?

The key point is that for a DNN (deep neural network) of depth L (number of layers) and width N (numbers of neurons per layer), there are $P = N^L$ possible paths connecting any single item x_i in the input layer to any another single item y_j in the output layer. Hence a DNN with $N = 10^3$ and $L = 10^2$ features $N^2 L = 10^8$ weights and $P = 10^{30}$ paths. Moreover, the search for the target can proceed in parallel across all of these paths. If you think that this is sci-fi, please think again, as current leading edge ML applications, such as DeepFold or Large Language Models motoring the most powerful "ask-me-anything" chatbots are using up to 100 billions weights, basically the number of neurons in our brain. Except that our brain works at 20 Watt while the largest ML models are now sucking up at least ten million times more energy, a point to which we shall return shortly.

These numbers unveil the magic behind ML: DNN duel the CoD face up, by unleashing an exponential number of paths, and adjusting them in such a way as to sensibly populate the sneaky regions where the golden nuggets are to be found.

This strategy is an opinion splitter: pragmatists are enthusiastic at the conceptual simplicity of this black-box and, speaking of weights, to them "too much is not enough". Scientists fond of Insight, are horrified at the diverging number of parameters, their "prejudice" being that parameters are fudge factors concealing lack of understanding, so, to them the motto is rather "the least the best". There are of course many nuances to be considered between these two opposite fronts, but here we shall focus on just two interconnected ones: *Explainability and Sustainability*.

Besides the exponential number of weights, in order to converge, the ML procedure needs correspondingly huge training datasets. Alas, training does not come for free: it is estimated that new generation chatbots will come near to the Gigawatt power demand, in excess of most existing power plants. The comparison with the 20 some Watts of our brain is embarassing. Given the energy-devouring nature of the large-scale ML procedure, the obvious question is: is it really worth sacrificing a substantial share of the total energy budget worldwide to the totem of chatbots?

This may make Jensen Huang, the founder of Nvidia, the richest man on Earth, but still it hardly looks like the way to go for the rest of us, apart the super-elite who may be able to afford extra-terrestrial life for the decades to come. The point is that, even when it works, ML is hardly Explainable, it offers little clue on the physical meaning of the parameters: Control but little Insight. And with Insight out of the game, there's no guarantee that what works for seen data will keep working again for the unseen ones (extrapolation).

Of course, one can close eyes and keep going on steroids with weights, but, given the power bill discussed above, this sounds reckless at best. There must be better ways. An interesting clue in this direction is that the overwhelming majority of the weights, the experts tell me, are close to zero, meaning that they do not contribute significantly to the success of the ordeal. This is scientifically *very* interesting and it begs for understanding, not just because this is what science is all about, but also because understanding here means saving oceans of Gigawatts.

Recently Elon Musk advocated the need of putting AI on a rigorous scientific basis, in his own words "Join xAI (Explainable AI) if you believe in our mission of understanding the universe, which requires maximally rigorous pursuit of the truth, without regard to popularity or political correctness." Yann LeCun, one of the most respected computer scientists worldwide and Turing awardee, promptly countered that "Musk wants a maximally rigorous pursuit of the truth but spews crazy-ass conspiracy theories on his own social platform." And, upon being questioned by Musk about his recent science, LeCun goes on by quoting his some 80 papers, as opposed to Musk's zero entries in this ballgame.

Now, it is ironic enough to hear one of the most muscular and hungriest AI energy consumers on the planet to advocate the rigorous pursuit of scientific truth. And despite his towering status in computer science, it is only slightly less ironic to see LeCun taking up the role of the guardian of science. Indeed, LeCun, incidentally also Chief AI Scientist at Zuckerberg's Facebook, is a champion of that kind of computer science where the dismissal of Insight in favor of Control is largely tolerated, see [13, 14].

But let's give Musk the benefit of doubt and assume he's genuinely interested in understanding the Universe. The tip is fairly simple: stop leveraging the muscular power of ML with legions of GPU's and *surprise us with more understanding and less weights*. The name of the game being causal AI or explainable AI, acknowledging the fact that whenever Correlation can replace Causality, what we are talking about is not Science but Control.

This is the biggest lie which has been served to us by AAI, where AAI stands for Aggressive AI, in order to distinguish it from the many important contributions of AI to technology and society. But the king is now hopefully naked, since the pursuit of Control regardless of Insight comes with a energy price tag that planet Earth just cannot sustain. Incidentally, many scientists (most physicists and mathematicians) much less visible than Musk and LeCun, are already doing this out of the limelight [15–17].

Musk's glorious statement about xAI calls for another comment, mostly related to a strategy that, back in 1944 in his prophetic book "The abolition of man", CS Lewis dubbed *Zero Sales Resistence*, ZSR for short [18]. Pursuing financial/economic interests under the glorified veil of world-saving intentions is a well-known strategy since long. What is new, though, is the unprecedented power of the modern AAI lords. Here comes the point. The speed of modern technology has basically collapsed spacetime: a single message can reach any (connected) individual on this planet in virtually no time. The result is that a well-crafted message, suitably conveyed by your most seductive influencer of choice (maybe a chatbot?), can win billions of brains in a single swoop, a process called *brain condensation* [19]. The associated profits go with some square root of this condensation ratio, as reflected by the four-five orders of magnitude gap between the salary of top executives *versus* their least paid employees [20]. Under the ZSR light it is hardly a surprise to hear Zuckerberg proclaiming that "connectivity is a human right" (to which Bill Gates allegedly returned "Have you ever heard of things likewater and bread?).

The next interesting step is to realize that steering the "sentiments" of human beings amounts to controlling a comparatively small number of high-level "psychological variables". This is piece of cake for the most powerful ML algorithms, no point of Insight, Control is all that matters [19].

Trying to cure Alzheimer with ML shows a very different movie, one where CoD hits hard, Correlation cannot replace Causation and Control cannot replace Insight [21], on pain of providing to society a cure far worst than the disease it is meant to relieve.

Yet, the tools are the same, it is basically the very same machinery described by equations (1 + 2+3 + 4) above! That's why Science is so easily served as a strawman for Control.

This is where AAI becomes as dangerous as never before in the history of science: it just trickles into our habits, step by step, pretending to ease all our pains without asking anything in return, "other than" data on all sorts of penchants of ours, food, movies or mating patterns alike. A silent but relentless conquer of our brain towards the ZSR goal. And please, make no mistake, the main item on sale with ZSR is not the commercial product but our brain instead.

The motto *Why learn if you can look it up?* speaks loud for the above. A perfect echo to the comment we find in Domingo's book, where we are informed that the Master Algorithm will do "anything we want before we even ask"!

So, while science occasionally gets significant contributions from ML, hardly ever in proportion to the bombastic headlines, what the Master Algorithm does for sure is secure sky-rocketing sales.

CS Lewis says it best, quoting verbatim: *"What we call Man's power over Nature turns out to be a power exercised by some men over other men with Nature as its instrument"*. In more mundane and actual terms, "If something is free, you are the product" (anonymous).

What to do then?

In the recent years there have been increasing calls to "Algorethics", the idea being of injecting "ethical constraints" into the ML algorithms [22, 23]. This is certainly a commendable goal, one which makes the object of high-level agreements between Institutions and AI companies.

The question is: will it really fix the ZSR issue? Will it stop AAI from promising master algorithms which do "anything we want before we even ask"?

I sincerely doubt it, no matter how good and well meant the law, history shows that the smart villain always manages to find the next loophole. More effective, I think, is to pursue the inherent spiritual drive of Science (and capital R Religion, of course), namely, the pursuit of Insight for the pleasure of finding things out, period. The rest will follow, as it always did. Control without Insight, on the contrary, is a sure recipe for major failure of society.

Author contributions

SS: Conceptualization, Writing–original draft.

Funding

The author(s) declare that financial support was received for the research, authorship, and/or publication of this article. The author is grateful to SISSA for financial support under the "Collaborations of Excellence" initiative as well as to the Simons Foundation for supporting several enriching visits of his.

Acknowledgments

The author wishes to acknowledge many enlightening discussions over the years with P. V. Coveney, G. Galli, A. Laio and D. Spergel. I am also grateful to M. Durve for critical reading of the manuscript.

Conflict of interest

The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

The author(s) declared that they were an editorial board member of Frontiers, at the time of submission. This had no impact on the peer review process and the final decision.

Publisher's note

References

1. Anderson C. The end of theory: the data deluge makes the scientific method obsolete. *Wired* (2008). Available online at: https://www.wired.com/2008/06/pb-theory/

2. Domingos P. *The master algorithm: how the quest for the ultimate learning machine will remake our world.* New York: Basic Books (2015).

3. Bellman R. *Dynamic programming*. Princeton, NJ: Princeton University Press (1957).

4. Succi S. *Sailing the ocean of complexity: lessons from the physics-biology interface*. Oxford: Oxford University Press (2022).

5. Succi S. Of naturalness and complexity. *The Eur Phys J Plus* (2019) 134(3):97. doi:10.1140/epjp/i2019-12576-3

6. Briscolini M, Santangelo P, Succi S, Benzi R. Extended self-similarity in the numerical simulation of three-dimensional homogeneous flows. *Phys Rev E* (1994) 50(3):R1745–7. doi:10.1103/physreve.50.r1745

7. Frenkel D, Smit B. *Understanding molecular simulation*. New York: Elsevier (2001).

8. Bussi G, Laio A. Using metadynamics to explore complex free-energy landscapes. *Nat Rev Phys* (2020) 2(4):200–12. doi:10.1038/s42254-020-0153-0

9. Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, et al. Highly accurate protein structure prediction with alphafold. *Nature* (2021) 596:583–9. doi:10.1038/s41586-021-03819-2

10. Lecun Y, Bengio J, Hilton G. Deep learning. *Nature* (2015) 521(7553):436–44. doi:10.1038/nature14539

11. Coveney PV, Dougherty ER, Highfield RR. Big data need big theory too. *Philosophical Trans R Soc A: Math Phys Eng Sci* (2016) 374:20160153. doi:10.1098/rsta.2016.0153

12. Succi S, Coveney PV. Big data: the end of the scientific method? *Philosophical Trans R Soc A* (2019) 377:20180145. doi:10.1098/rsta.2018.0145

13. Coveney PV, Highfield RR. Artificial intelligence must be made more scientific. *J Chem Inf Model* (2024) 64:5739–41. doi:10.1021/acs.jcim.4c01091

14. Coveney PV, Highfield RR. Artificial intelligence must be made more scientific. *J Chem Inf Model* (2024) 64:5739–41. doi:10.1021/acs.jcim.4c01091

15. Del Tatto V, Fortunato G, Bueti D, Laio A. Robust inference of causality in high-dimensional dynamical processes from the information imbalance of distance ranks. *Proc Natl Acad Sci* (2024) 121(19):e2317256121. doi:10.1073/pnas.2317256121

16. Gendron G, Witbrock M, Dobbie G. A survey of methods, challenges and perspectives in causality. *arXiv* [Preprint]. arXiv:2302.00293 (2023). Available online at: https://doi.org/10.48550/arXiv.2302.00293

17. Goldt S, Mezard M, Krzakala F, Zdeborová L. Modeling the influence of data structure on learning in neural networks: the hidden manifold model. *Phys Rev X* (2020) 10:041044. doi:10.1103/physrevx.10.041044

18. Lewis CS. *The abolition of man*. New York: Harper-Collins (1944).

19. Riva G, Wiederhold BK, Succi S. Zero sales resistance: the dark side of big data and artificial intelligence. *Cyberpsychology, Behav Social Networking* (2022) 25(3):169–73. doi:10.1089/cyber.2022.0035

20. O'Neil C. *Weapons of math destruction: how big data increases inequality and threatens democracy*. London: Penguin Books (2017).

21. Coveney PV, Highfield RR. *Virtual you*. Princeton, NJ: Princeton University Press (2023).

22. Benanti P, Patel RB, Parikh SM, Prajapati J. "Artificial Intelligence and Ethics". in Pathak YV, Saikia S, Pathak S, Patel JK, Prajapati JB, editors. *Ethical Issues in AI for Bioinformatics and Chemoinformatics*. Boca Raton, FL: CRC Press (2023). p. 1–11. doi:10.1201/9781003353751-1

23. Kearns M, Roth A. *The ethical algorithm: the science of socially aware algorithm design*. Oxford: Oxford University Press (2024).