Biological faithfulness is not necessary for machine learning

An investigation of backpropagation as a deep learning network with no biological correlate

Meagan Wiederman

ABSTRACT

Artificial intelligence (AI) is an area of research where machines are programmed to produce ‘learned’ outputs rather than explicitly programmed outputs. Some advancements in AI have focused on fully replicating the thinking human brain in machinery, which would create a model that ‘thinks’. However, this is inefficient given the complexity of the brain. AI programs that are not modelled after the brain (non-biologically faithful) may yield more promising advancements towards reproducing thought in AI. Backpropagation is an error minimizing algorithm used in AI that has no biological counterpart. Backpropagation is successfully used in medicine for diagnoses, to optimize image compression, and in pharmacological drug concentration modelling. The recent success of backpropagation demonstrates that biological faithfulness is not required for machine thought.

NATURAL INTELLIGENCE, ARTIFICIAL INTELLIGENCE, AND BIOLOGICAL FAITHFULNESS

The brain, or natural intelligence (NI), is considered the gold standard model of intelligence. NI systems are able to ‘think’, defined as the judgment of mental concepts, and ‘learn’, defined as accommodating new thoughts. In contrast to NI, artificial intelligence (AI) are machines that are programmed to produce outputs that are “learned” rather than being explicitly programmed.¹ Progress inspired by NI advances the field of AI design under the assumption that making a machine like the thinking brain would translate into a thinking machine. However, this assumption does not exclude that machine thinking that is different from NI (non-biologically faithful) may be considered thinking. Since there are many exhaustive intricacies inherent in replicating the brain, it may be more productive to pursue machine learning that doesn’t replicate an identical architecture to the brain.

A thinking machine must solve problems in the environment of a biological thing.² This criterion of biologically faithful adaptations is not restricted to biologically faithful machines. The Turing test (imitation game) is one example of a setting where machines solve biologically relevant problems to test a machine’s ability to think.³ In the imitation game, machine intelligence is evaluated by a judge who compares the machine’s response to questions to a human’s response. The Turing test proposes that successful imitation is evidence that machines “think”. Since the machine and human cannot be seen and it’s only their output to the judge’s questions

that indicates their intelligence, the Turing test can allow for non-biologically faithful models to be considered “thinking”.

BIOLOGICALLY FAITHFUL NEURAL AI

McClelland et al argue that models should be informed and constrained by our knowledge of the neural processes that underpin cognition, since learning is inherently anthropomorphic, and that machines will require humanistic qualities to be capable of such feats.⁴ This group advocates for use of parallel distributed processing, a brain-inspired neural network that processes information in many places at once, as the standard model in AI. In this thought camp, termed “connectionism”, the brain (connectome) is considered as an exhaustive collection of cells (nodes) with synapses (connections) that produces thought and is viewed as the gold-standard for AI. Mapping the connectome is currently being piloted in the fruit fly visual cortex.⁵ Continuing this exhaustive mapping of all the connections would eventually present a viable model for an accurate, biologically faithful adaptation of the brain.⁵ One could imagine that if one were to exactly replicate all the nodes and their connections from a human brain into a machine, then one should have an intelligent machine. Indeed, both Moravec and Kurzweil have argued that it is technologically feasible to copy the brain directly into hardware and software, and that such a simulation will be essentially identical to the original.^{6,7} If one accepts the brain as a thinking thing, an exact replica of the brain using connectomics and intricate mechanical design may present one case of a thinking machine.

THE SUCCESS OF BACKPROPAGATION IN AI, A NON-BIOLOGICALLY FAITHFUL PROCESS

To date, neural realism has not produced most of the current progress in AI. The greatest current advancement is backpropagation, which does not reflect any biological phenomena. Backpropagation is the algorithm for supervised training of modern AI.⁸ During training, the network is shown sample inputs and the correct output classifications and it executes two phases: propagation and weight updating. During propagation, the networks first approximate the non-linear relationship between the input and the output by assigning random value weights to the connections of input to output through any number of hidden layers and calculating an output (Figure 1, black). During weight updating, the received output is compared to the desired or “target” output using a loss function to produce an error term for each node. This error is propagated backwards from the output layer towards the

input to adjust the weight of the connections between the nodes in the opposite direction of their contribution to the error (Figure 1, red). Repeated iteration of propagation and weight updating causes the network to converge upon values that minimize the loss function and produce the desired output (Figure 2). During training, input nodes organize themselves to recognize different characteristics so that these nodes will respond to new patterns input with these features. Thus, backpropagation uses a sort of “squashing” function to extract the relevant characteristics. AI using backpropagation deal with decision-like tasks such as pattern recognition and automatic control by recognizing these features. The backpropagation network can correctly generate an output for a novel input not included in its training patterns, demonstrating “learning” that can be considered machine thinking.

Figure 1. (Top) A backpropagation network starts with the random assignment of weights to generate output (black). Weight updating using error term (red).

Figure 2. (Bottom) New weights are calculated from error term. The backpropagation network now reliably generates the desired output and successfully characterizes novel inputs.

There are currently no known theories of the brain using the output to calculate an error term that alters individual synaptic strength, unlike the process that occurs with backpropagation. Rather, synaptic plasticity exists in only 2 ways: Hebbian and homeostatic.^{9,10} Hebbian synaptic plasticity involves synapses deteriorating when failing to fire synchronously or propagating if it successfully fires synchronously with its post-synaptic partner. This form of plasticity reinforces and builds networks and prunes unnecessary connections. Homeostatic plasticity reduces the total firing of the brain back to homeostasis. Neither of these two methods of plasticity is dependent on the output of the neural network since only the activity of local synapses is considered. Nor do either of these two forms of neural learning plasticity aim to reach a desired goal. Hebbian plasticity is used non-specifically for reinforcement while homeostatic plasticity is restorative. Therefore, based on our current understanding of synaptic plasticity, the brain doesn't complete backpropagation. Despite backpropagation being non-biologically faithful and a movement away from neural inspiration, it has been very successful.

MEDICAL APPLICATIONS

Backpropagation networks have been able to successfully distinguish between an infected and non-infected person, effectively presenting a diagnosis.¹¹ This task completed by backpropagation networks accurately replicates a diagnosis by a physician, suggesting that backpropagation networks may have, in this task, successfully paralleled what is seen in a biologically faithful model. Backpropagation has also been used for image compression and storage of X-ray and MRI data.¹² To store medical images, one needs to compress the images by transforming them into simpler data. This can be completed using a backpropagation network to decide how an image's contents relate to their optimum compression ratio for digital storage of X-ray images.¹² This decision task effectively completed by backpropagation exceeds the capabilities of biologically faithful networks. Backpropagation networks are also used to model pharmacological interactions. One example of this is of a backpropagation network that was designed to predict the concentration of intravenous and intragastric losartan following administration.¹³ This network simplifies the calculation of blood drug concentration when compared to what is achieved with NI.

LIMITATION OF BACKPROPAGATION

Backpropagation allows novel inputs to give rise to an output that a network is trained for. However, backpropagation is incapable of “innovation” (novel outputs from given inputs) since it is restricted to judgement tasks. Innovation is critical to refining thought, an aspect that has led to systems like society, governance, and business to be developed by NI. Innovation may represent one aspect of NI that can be achieved by the natural brain but not by non-biologically faithful AIs like backpropagation networks. In addition, biologically faithful AI neural networks have been able to create paintings,¹⁴ suggesting their relative “creativity”.

CONCLUSION AND FUTURE DIRECTIONS

Backpropagation is regarded as one of the most significant advancement in AI currently in use.¹⁵ It has been argued that attempts to build brain-like AI are limited by our understanding of the brain. If we assume success is an indication of progress towards ‘thought’, then the success of backpropagation demonstrates that biologically faithful adaptations of the brain are unnecessary to produce machine thought. The current use of backpropagation in different medical situations is evidence of how a non-biologically faithful machine can achieve accurate results on judgement tasks that are comparable to that of NI, illustrating non-biologically faithful machine success. Both biologically faithful and non-biologically faithful advancements in AI are striving to achieve more effective machines capable of machine learning.

REFERENCES

1. Dreyfus H, Haugeland J. From micro-worlds to knowledge representation: AI at an impasse. MIT Press; 1981. 154 p.
2. Turing A. Computing Machinery and Intelligence. *Mind*. 1950; 49:433-60.
3. Gomes L. Machine-Learning Maestro Michael Jordan on the Delusions of Big Data and Other Huge Engineering Efforts. [Internet]. *IEEE Spectrum*; 2014 [modified 2014 Oct 20; cited 2018 Mar 15]. Available from: <https://spectrum.ieee.org/robotics/artificial-intelligence/machinelearning-maestro-michael-jordan-on-the-delusions-of-big-data-and-other-huge-engineering-efforts>.
4. McClelland JL, Cleeremans A. Oxford Companion to Consciousness: Connectionist Models. New York: Oxford University Press; 2009. 21-7 p.
5. Horne JA, Langille C, McLin S, *et al*. A resource for the Drosophila antennal lobe provided by the connectome of glomerulus VA1v. *eLife*. 2018 Nov;7, e37550. <https://doi.org/10.7554/eLife.37550>
6. Moravec H. Mind Children: The Future of Robot and Human Intelligence. New York: Harvard University Press; 1990. 27 p.
7. Kurzweil R. The Singularity is Near: Achieving the Computational Capacity of the Human Brain. New York: Viking Press; 2005. 9-17p.
8. Hecht-Nielsen R. Theory of the backpropagation neural network. Neural networks for perception. HNC, Inc. and University of California, San Diego; 1992. 65-93 p.
9. Magee JC, Johnston D. A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. *Science*. 1996 Jan;275(5297),209-13. <https://doi.org/10.1126/science.275.5297.209>
10. Turrigiano GG. Homeostatic plasticity in neuronal networks: the more things change, the more they stay the same. *Trends Neurosci*. 1999 May;22(5), 221-7. [https://doi.org/10.1016/S0166-2236\(98\)01341-1](https://doi.org/10.1016/S0166-2236(98)01341-1)
11. Gupta A, Shreevastava M. Medical Diagnosis using Back propagation Algorithm. *Int J of Emerg Tech and Adv Eng*. 2011 Nov;1(1): 2250-2459. ISSN: 2250-2459.
12. Dimililer K. Backpropagation Neural Network Implementation for Medical Image Compression. *J Appl Math*. 2013 Nov;2013:1-8. <https://doi.org/10.1155/2013/453098>
13. Lin B, Lin G, Liu X, *et al*. Application of back- propagation artificial neural network and curve estimation in pharmacokinetics of losartan in rabbit. *Int J Clin Exp Med*. 2015;8(12):22352–8. ISSN:1940-5901/ IJCEM0015908.
14. Guillemin A. La Comtesse De Belamy. [Internet]. OBVIOUS; 2018 [modified 2018; cited 2018 Oct 7]. Available from: <http://obvious-art.com/la-comtesse-de-belamy.html>.
15. Frasconi P, Gori M, Tesi A, *et al*. Successes and Failures of Backpropagation: A theoretical investigation. Greenwich: Ablex Publishing; 1993. 235 p.