



Neuromorphic Cognitive Learning Systems: The Future of Artificial Intelligence?

Vassilis Cutsuridis^{1,2}

Received: 20 March 2024 / Accepted: 12 May 2024 / Published online: 19 May 2024

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2024

Abstract

In this position paper, I outline the caveats of the current artificial intelligence (AI) field driven by deep learning (DL) and large data volumes. Although AI/DL has demonstrated huge potential and attracted huge investments globally, it encounters big problems – it not only need to collect huge datasets and spend enormous time and resources to be trained on them, but also the trained system cannot deal effectively with any never encountered before (novel) data. From a human perspective, any current AI/DL system is completely unintelligent. It is only able to represent information but have no awareness of what this information means. I propose as an alternative the Neuromorphic Cognitive Learning Systems (NCLS), intimate imitations of animal and human brains, able to address the AI/DL limitations and achieve true artificial general intelligence. Similar to human and animal brains NCLS are unparalleled in their ability to rapidly, and on their own, adapt and learn from changing and unexpected environmental contingencies with very limited resources. I describe how NCLS driven AI inspired by human/animal brains can pave the way to new computing technologies with the potential to revolutionize the industry, economy and society. It is my strong belief that NCLS investigations will have major impact to real-time autonomous systems to achieve human-like intelligence capabilities.

Keywords Neuromorphic computing · Brain-inspired AI · Artificial general intelligence · Deep learning · Spiking neural networks · Cognitive learning systems

AI driven by DL has demonstrated huge potential and attracted huge investments globally. The DL revolution has been successful in a wide variety of applications in nearly every scientific field including medical imaging, e-commerce, drug discovery, and agriculture. The basis of any DL system is artificial neural networks, crude approximations of biological neural networks with nodes connected via weighted edges that could be strengthened or weakened according to a prescribed learning rule. They are end-to-end systems with an opacity problem [1]. Despite intense research and huge monetary investments, it is still unclear why AI/DL works so well [2]. To reach such seemingly impressive results, they require huge datasets and

enormous time and resources. They have a voracious appetite for computing power [3], are extremely data-dependent, require massive amounts of labelled data to harvest from, and have long off-line training via backpropagation to adequately learn from them. They are notoriously unable to learn quickly (require many epochs) and stably without catastrophically forgetting their past knowledge [4]. When trained for a new class, these networks lose the previously trained information, often leading to catastrophic forgetting [5], and thus, they need to be trained anew. They also need to be trained with almost every possible input/output pair and be tested on data similar to the training ones, thus allowing answers to be interpolated in between old ones [6]. If the trained system ever encounters any before unseen (novel) data, then its performance crumbles. In reality, this means a well-trained autonomous vehicle cannot cope with unfamiliar roads, or a well-trained robot cannot work again if the environment is slightly modified. They are unable to generalize well. A notable example, DeepMind's Atari deep neural networks (DNNs) trained day and night for weeks, on more than a thousand processors and GPUs,

✉ Vassilis Cutsuridis
vcutsuridis@lincoln.ac.uk; vcutsuridis@gmail.com

¹ Lincoln Group of Excellence for Neuromorphic Cognitive Learning Systems, College of Health and Science, University of Lincoln, Lincoln, UK

² School of Computer Science, University of Lincoln, Lincoln LN6 7TS, UK

using petabytes and petabytes of training data, while playing games against each other to improve their skills, eventually achieving incredible results, and then a minor perceptible pixel perturbation, and their performances crumbled [7]. As if the DNNs never really understood anything about what they were doing, but simply memorized data, overfitted them, and superficially acted upon them without awareness of what these pixels represented. From a human perspective, they were monumentally unintelligent [8]. They cannot understand the hierarchical structure, causal relations, abstractions, symbols, and ontological inference. They are only able to represent information but have no awareness of what this information means. They have been shown to misclassify TV sets for orangutans, generate dirty words and racist remarks when asked to generate children's stories, and give ridiculous answers to simple common-sense reasoning questions [8]. Perceptible to human eyes, alterations in images, originally correctly classified (e.g. as a lion), caused a DL to label the image as something else entirely (e.g. mislabeling a lion for a library) [9].

Thus, a new shift in AI is required to overcome these caveats and achieve true artificial general intelligence. Neuromorphic cognitive learning systems (NCLS) constitute attractive alternatives, because they are intimate imitations of animal and human brains. Brains are marvels of evolutionary engineering able to process vast amounts of information extremely fast, accurately, and with far less energy than AI/DL. They are unparalleled in their ability to rapidly, and on their own, adapt and learn from changing and unexpected environmental contingencies with very limited resources. Their basic information processing units are spiking neurons. In contrast to ANNs/DL, which are continuously active, spiking neurons are active only when they spike. Because they communicate with spikes, they are energy efficient, because they consume energy only when a spike is generated. Because neurons in conventional DNNs are continuously active, they require high computational and energy costs, especially in models with billions of parameters. NCLS are thus ideal for low-power devices, such as mobile phones or IoT sensors. NCLS also utilize event-based processing in which neurons spike only in response to specific stimuli. This sparsity of activity means that at any given time, only a small fraction of neurons is active, drastically reducing energy consumption. NCLS process information based on the timing of spikes (temporal coding), which allows them to encode temporal information and be temporally precise. Temporal coding has been found to carry more information than the rate-based code and is extremely fast. Neurons can be made to react to single spikes, allowing for extremely fast binary calculations. The human brain can recognize faces in as little as 100 ms [10]. Their ability to process sequential data makes NCLS useful in tasks such as speech recognition or gesture recognition. To avert

catastrophic forgetting, NCLS use temporal learning, a form of biological continual learning that depends on the timing of spikes. Spike timing-dependent plasticity (STDP) [11], an unsupervised learning rule that requires no data labeling, adjusts the strength of the connections between spiking neurons based on the timing of their spikes. STDP as a lifelong learning rule in AI models could allow them to gradually improve their performance as they are exposed to more data. Other temporal learning rules include the tempotron rule [12], the SpikeProp rule [13], the SPAN rule [14], the Chronotron rule [15], and the ReSuMe rule [16]. These rules modify synaptic weights according to a gradient descent approach in an error landscape where the error function depends on the distance between two spike trains. Finally, NCLS are robust to noise and can effectively filter out irrelevant information.

NCLS are currently supported by computing platforms and hardware. SpiNNaker [17] can run real-time billions of spiking neurons to simulate the human brain. TrueNorth [18], Loihi [19], and Tianji [20] neuromorphic chips can implement NCLS with 10,000 spiking neurons. NCLS have shown mastery in visual [21] and olfactory [22] recognition, multi-sensory concept learning [23], knowledge representation and reasoning [24], decision making and motor control [25], learning and memory [26, 27], and natural language processing [28].

Overall, NCLS-driven AI inspired by human/animal brains will open roads to new computing technologies which will revolutionize the industry, economy, and society. NCLS will have a significant impact on real-time energy-efficient autonomous systems to achieve human-like intelligence capabilities.

Author Contributions V.C. wrote the manuscript.

Funding This work was supported by the EU HORIZON 2020 Project ULTRACEPT under Grant 778062.

Data Availability No datasets were generated or analysed during the current study.

Declarations

Competing Interests The authors declare no competing interests.

References

1. Waldrop MM. What are the limits of deep learning? *PNAS*. 2019;116(4):1074–7.
2. Tishby N, Zaslavsky N. Deep learning and the information bottleneck principle. 2015. [arXiv:1503.02406](https://arxiv.org/abs/1503.02406).
3. Thompson NC, Greenewald K, Lee K, Manso GF. The computational limits of deep learning. 2022. [arXiv:2007.05558](https://arxiv.org/abs/2007.05558).

4. Grossberg S. How does a brain build a cognitive code? *Psych Rev.* 1980;87:1–51.
5. McCloskey M, Cohen NJ. Catastrophic interference in connectionist networks: the sequential learning problem. *Psychol Lear Motiv.* 1989;24:109–65.
6. Touretzky DS. Connectionism and compositional semantics. In: Barnden JA, Pollack JB, editors. *High-level connectionist models.* Hillsdale, NJ: Erlbaum; 1991. p. 17–31.
7. Kinsky K, Silver T, Mély DA, Eldawy M et al. Schema networks: zero-shot transfer with a generative causal model of intuitive physics. 2015. [arXiv:1706.04317](https://arxiv.org/abs/1706.04317).
8. Marcus G. Deep learning: a critical appraisal. 2017. [arXiv:1801.00631](https://arxiv.org/abs/1801.00631).
9. Nguyen A, Yosinski J, Clune J. Deep neural networks are easily fooled: high confidence predictions for unrecognizable images. In: *Computer Vision and Pattern Recognition (CVPR '15).* IEEE; 2015.
10. VanRullen R, Guyonneau R, Thorpe SJ. Spike times make sense. *Trends Neurosci.* 2005;28(1):1–4.
11. Bi GQ, Poo MM. Synaptic modification by correlated activity: Hebb's postulate revisited. *Ann Rev Neurosci.* 2001;24:139–66.
12. Güttig R, Sompolinsky H. The tempotron: a neuron that learns spike timing-based decisions. *Nat Neurosci.* 2006;9(3):420–8.
13. Bohte SM, Kok JN, Poutré JAL. Error-backpropagation in temporally encoded networks of spiking neurons. *Neurocomputing.* 2002;48(1–4):17–37.
14. Mohammed A, Schliebs S, Matsuda S, Kasabov N. SPAN: spike pattern association neuron for learning spatio-temporal spike patterns. *Int J Neural Syst.* 2012;22(04):1250.
15. Florian RV. The Chronotron: a neuron that learns to fire temporally precise spike patterns. *PLoS ONE.* 2012;7(8):e40.233.
16. Ponulak R. ReSuMe-new supervised learning method for spiking neural networks. Institute of Control and Information Engineering, Poznań University of Technology, Tech. Rep.; 2005.
17. Painkras E, Plana L, Garside J, Temple S, et al. Spinnaker: a multi-core system-on-chip for massively-parallel neural net simulation. In: *NaIn Proceedings of the IEEE 2012 Custom Integrated Circuits Conference.* San Jose, CA: IEEE; 2012. p. 1–4.
18. Akopyan F, Sawada J, Cassidy A, Alvarez-Icaza R, et al. TrueNorth: design and tool flow of a 65 mw 1 million neuron programmable neurosynaptic chip. *IEEE Trans Comput Aided Des Integr Circ Syst.* 2015;34:1537–57.
19. Davies M, Srinivasa N, Lin T, Chinya G, et al. Loihi: a neuromorphic manycore processor with on-chip learning. *IEEE Micro.* 2018;38:82–99.
20. Pei J, Deng L, Song S, Zhao M, et al. Towards artificial general intelligence with hybrid tianji chip architecture. *Nature.* 2019;572:106–11.
21. Kim S, Park S, Na B, Yoon S. Spiking-YOLO: spiking neural network for energy-efficient object detection. 2019. [arXiv:1903.06530](https://arxiv.org/abs/1903.06530).
22. Imam N, Cleland TA. Rapid online learning and robust recall in a neuromorphic olfactory circuit. *Nat Mach Intell.* 2020;2:181–91.
23. Wang Y, Zeng Y. Multisensory concept learning framework based on spiking neural networks. *Front Syst Neurosci.* 2022;16:845177.
24. Fang H, Zeng Y, Tang J, Wang Y, et al. Brain-inspired graph spiking neural networks for commonsense knowledge representation and reasoning. 2022. <https://doi.org/10.48550/arXiv.2207.05561>.
25. Zeng Y, Zhao D, Zhao F, et al. BrainCog: a spiking neural network based, brain inspired cognitive intelligence engine for brain inspired AI and brain simulation. *Patterns.* 2023;4:100789.
26. Cutsuridis V, Graham BP, Cobb S. Encoding and retrieval in the hippocampal CA1 microcircuit model. *Hippocampus.* 2010;20(3):423–46.
27. Andreanos N, Yue S, Cutsuridis V. Associative memory retrieval evaluation in a brainmorphic microcircuit model. 2024. Under review.
28. Knipper RA, Mishty K, Sadi M, Santu SKK. SNNLP: energy-efficient natural language processing using spiking neural networks. 2024. [arXiv:2401.17911](https://arxiv.org/abs/2401.17911).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.