Deep Learning and Computational Neuroscience

Erik De Schutter¹

Published online: 20 January 2018

© Springer Science+Business Media, LLC, part of Springer Nature 2018

Especially young colleagues are fascinated by the potential of deep learning for neuroscience. This was obvious at the recent Society for Neuroscience meeting in Washington DC, where the few posters that had the magical words in their title attracted large crowds of attendees who seemed almost exclusively in their twenties. The success of deep learning of data representation has led to impressive applications in image, video and speech processing.¹ Compared to these, recent advances in applying reinforcement learning to playing games are outright mind blowing, with AlphaGo Zero achieving superhuman performance in just three days of training on a single machine with specialized hardware.² It is, therefore, easy to predict that the interest in deep learning among young computational neuroscientists will only increase, but the reality may be more complex than they surmise. In this Editorial, I will focus on the question of correspondence between deep learning and how the brain works.³ I will not consider the many opportunities of applying deep learning as a supporting technology.

The original breakthrough leading to the success of deep learning tested the method on an image recognition task, classifying handwritten digits.⁴ Correspondingly, most of the

applications of deep learning to computational neuroscience are about understanding the visual system (including the posters at the recent Society for Neuroscience meeting). As pointed out in a recent review⁵ one category of deep learning models, goal-driven hierarchical convolutional neural networks, has been very successful at predicting neural responses in several layers of primate visual cortex, including V1, V2, V4 and inferior temporal cortex (IT). But the authors also point out that this success is probably due to convolutional neural networks closely mimicking the overall architecture of cortex⁵, in particular implementing features similar to receptive fields of increasing size across the hierarchy. This leads to a warning that any correspondences between deep learning methods and the brain may not generalize to all deep learning. In fact, though the field of machine learning has clearly been inspired by neuroscience³, it has never seen this as a limitation on the methods it can use. For example, the breakthrough referred to earlier⁴ was a method to teach layers in a multilayer network one at a time, something that is hard to imagine occurring in a real brain. Deep learning networks have typically also many more layers than corresponding brain systems and one of the current hypes are “very deep” models with tens of layers.⁶ A recent breakthrough, also used in AlphaGo Zero, are residual networks where shortcut connections are used that connect units in lower layers directly with units in higher layers⁶. Residual networks are an example of deep learning methods that do not reflect real neural systems, this would be like V1 densely projecting directly to V4 or IT. Conversely, there are well known brain circuits that have clearly quite different architectures than visual cortex, like for example the olfactory system.

Another difference between deep learning and human brains is the number of training examples required, with

¹ LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436–444. <https://doi.org/10.1007/s10994-013-5335-x>.

² Silver, D., et al. (2017). Mastering the game of go without human knowledge. *Nature*, 550, 354–359. <https://doi.org/10.1038/nature24270>.

³ A more extensive analysis emphasizing the brain to deep learning connection can be found in Demis Hassabis et al. (2017). Neuroscience-inspired artificial intelligence. *Neuron*, 95, 245–258. <https://doi.org/10.1016/j.neuron.2017.06.011>.

⁴ Hinton, G. E., Osindero, S., & Teh, Y.-W. (2006). A fast learning algorithm for deep belief nets. *Neural Computation*, 18, 1527–1554. <https://doi.org/10.1162/neco.2006.18.7.1527>.

✉ Erik De Schutter
erik@oist.jp

¹ Computational Neuroscience Unit, Okinawa Institute of Science and Technology Graduate University, Okinawa, Japan

⁵ Yamins, D. L. K., & DiCarlo, J. J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience*, 19, 356–365. <https://doi.org/10.1017/S0140525X0001863X>.

⁶ He, K., et al. (2016). Deep residual learning for image recognition. *Conference on Computer Vision and Pattern Recognition*.

millions of labeled images needed to learn simple categorization tasks⁵. In fact, deep learning would not exist if the digital revolution hadn't made big data available. Fortunately, the human brain is better at generalizing from smaller sets of experiences, but recent machine learning approaches try to mimic this³. Conversely, we may not learn to recognize some features because we do not routinely train ourselves on labeled data. An example is the recent controversial study claiming that a deep network learned to recognize sexual preferences of people by analyzing pictures on a dating website.⁷ In newspapers, this was reported as a demonstration of how artificial intelligence can now beat the human mind. In real life, however, people often don't know who is gay or not (we lack the label) and I doubt any rational person would consider training themselves by going through all the profiles on a dating site.

Returning to the AlphaGo Zero example², the success of deep learning may soon be surpassed by reinforcement learning, which is again directly based on neuroscience concepts.^{3,8} In this case the big data challenge was overcome by having AlphaGo play games against itself, so no prior data was required. This may be conceptually similar to many forms of learning in infancy, where trial and error clearly play a very big role. But while being successful at playing Go and chess seems a big achievement because of the close to infinite number of possible game states, the important limiting metric for reinforcement learning is the number of possible actions and that number is, clearly, limited for any board game. Nevertheless, it is probably worthwhile to carefully investigate lessons that can be drawn from AlphaGo Zero.

Finally, I want to report on an interesting recent report⁹ that shows a fairly realistic way to solve one of the most vexing problems in mapping machine learning to the brain: the credit assignment problem. The first machine learning revolution in

the 80ies was based on the discovery of the back-propagation algorithm¹⁰ and it is well known that real brains have no back-propagation³, i.e. transmission of weight updates from higher layers to lower layers. In more modern terms, this is called the credit assignment problem: which neurons in lower layers directly contributed to the final behavioral outcome? Interestingly, leading researchers in deep learning are quite concerned about finding solutions to this conundrum^{11,12} but till now only partial answers were proposed. Guerguiev et al.⁹ propose to use the dendritic structure of neurons and this is a remarkably complete scheme, although there are still a few unresolved issues. Specifically, sparse feedback connections carrying higher-order feedback to the apical dendrites are used to drive changes in synaptic weights in basal dendrites that receive sensory input. The dendrites are essential because they provide for physical separation of the two inputs onto the same neuron. I will not describe the results in detail, but encourage reading of the paper.

From my perspective, this is quite an ironic result. As an avid modeler of dendrites^{13,14} I have been frustrated by the swing of the field of computational neuroscience towards point neuron network modeling. I have seen a gradual loss of interest among students of summer schools in computational neuroscience in simulating single neuron models and the shift is also noticeable in recent textbooks.¹⁵ If what Guerguiev and colleagues (including a scientist working at DeepMind) propose⁹ is true, then understanding how neural networks learn complex tasks will require models that include both apical and basal dendrites, though the models themselves do not have to be very complex. This nicely matches the recent demonstration in in vivo experiments of the behavioral importance of dendritic spikes¹⁶ and suggests a bright future for modeling dendrites.

⁷ www.economist.com/news/science-and-technology/21728614-machines-read-faces-are-coming-advances-ai-are-used-spot-signs.

⁸ Littman, M. L. (2015). Reinforcement learning improves behaviour from evaluative feedback. *Nature*, 521, 445–451. <https://doi.org/10.1177/0278364913495721>.

⁹ Guerguiev, J., Lillicrap, T. P., & Richards, B. A. (2017). Towards deep learning with segregated dendrites. *eLife*, 6. <https://doi.org/10.7554/eLife.22901>.

¹⁰ Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. In T. A. Polk and C. M. Seifert (eds), *Cognitive modeling*. MIT Press.

¹¹ Lillicrap, T. P. et al. (2016). Random synaptic feedback weights support error backpropagation for deep learning. *Nature Communications*, 7, 13276. <https://doi.org/10.1038/ncomms13276>.

¹² Bengio, Y., Lee, D.-H., Bomschein, J., & Lin, Z. (2015). Towards biologically plausible deep learning. arXiv:1502.04156.

¹³ Steuber, V. et al. (2007). Cerebellar LTD and pattern recognition by Purkinje cells. *Neuron*, 54, 121–136. <https://doi.org/10.1016/j.neuron.2007.03.015>.

¹⁴ Anwar, H., et al. (2013). Stochastic calcium mechanisms cause dendritic calcium spike variability. *The Journal of Neuroscience*, 33, 15848–15867. <https://doi.org/10.1523/JNEUROSCI.1722-13.2013>.

¹⁵ Gerstner, W., et al. (2009). *Neuronal dynamics*. Cambridge University Press. <https://doi.org/10.1017/CBO9781107447615>.

¹⁶ Takahashi, N., et al. (2016). Active cortical dendrites modulate perception. *Science*, 354, 1587–1590. <https://doi.org/10.1126/science.aah6066>.