
11.1 An Incomplete Overview of Open Research Questions

We conclude this book with a list of open research questions. A similar list, from which we have borrowed some of the problems we present here, can be found in [1]. We were hoping to compile a diverse list to show how rich and diverse research in deep learning can be. The problems we find most intriguing are:

1. Can we find something else than gradient descent as a basis for backpropagation? Can we find something as an alternative to backpropagation as a whole for weight updates?
2. Can we find new and better activation functions?
3. Can reasoning be learned? If so, how? If not, how can we approximate symbolic processes in connectionist architectures? How can we incorporate planning, spatial reasoning and knowledge in artificial neural networks? There is more here than meets the eye, since symbolic computation can be approximated with solutions to purely numerical expressions (which can then be optimized). A good nontrivial example is to represent $A \rightarrow B$, $A \vdash B$ with $\frac{B}{A} \cdot A = B$. Since it seems that a numerical representation of logical connectives can be found quite easily, can a neural network find and implement it by itself?
4. There is a basic belief that deep learning approaches consisting of many layers of nonlinear operations correspond to the idea of re-using many subformulas in symbolic systems. Can this analogy be formalized?
5. Why are convolutional networks easy to train? This is of course connected with the number of parameters, but they are still easier to train than other networks with the same number of parameters.
6. Can we make a good strategy for self-taught learning, where training samples are found among unlabelled samples, or even actively sought by an autonomous agent?

7. The approximation of the gradient is good enough for neural networks, but it is currently computationally less efficient than symbolic derivation. For humans, it is much easier to guess a number that is close to a value (e.g. a minimum) than to compute the exact number. Can we find better algorithms for computing approximate gradients?
8. An agent will be faced with an unknown future task. Can we develop a strategy so that it is expecting it and can start learning right away (without forgetting the previous tasks)?
9. Can we prove theoretical results for deep learning which use more than just formalized simple networks with linear activations (threshold gates)?
10. Is there a depth of deep neural networks which is sufficient to reproduce all human behaviour? If so, what would we get by producing a list of human actions ordered by the number of hidden layers a deep neural network needs to reproduce the given action? How would it relate to the Moravec paradox?
11. Do we have a better alternative than simply randomly initializing weights? Since in neural networks everything is in the weights, this is a fundamental problem.
12. Are local minima a fact of life or only an inherent limitation of the presently used architectures? It is known that by adding hand-crafted features helps, and that deep neural networks are capable of extracting features themselves, but why do they get stuck? Curriculum learning helps a lot in some cases, and we can ask whether the curriculum is necessary for some tasks?
13. Are models that are hard to interpret probabilistically (such as stacked autoencoders, transfer learning, multi-task learning) interpretable in other formalisms? Perhaps fuzzy logic?
14. Can deep networks be adapted to learn from trees and graphs, not just vectors?
15. The human cortex is not always feed-forward, it is inherently recurrent, and there is recurrence in most cognitive tasks. Are there cognitive tasks which are learnable only by feed-forward or only by recurrent networks?

11.2 The Spirit of Connectionism and Philosophical Ties

Connectionism today is more alive and vibrant than ever. For the first time in the history of AI, connectionism, under its present name of ‘deep learning’, is trying to take over GOFAI’s central position, and reasoning is the only major cognitive ability that remains largely unconquered. Whether this is a final wall which can never be breached, or just a matter of months, is hard to tell. Artificial neural networks as a research area almost died out a couple of times during similar quests. They were always the underdog, and perhaps this is the most fascinating part. They finally became an important part of AI and Cognitive Science, and today (in part thanks to marketing) they have an almost magical appeal.

A sculptor has to have two things to make a masterpiece: a clear and precise idea what to make, and the skill and tools to make it. Philosophy and mathematics are the two oldest branches of science, old as civilization itself, and most of science

can be seen as a gradual transition from philosophy to mathematics. This can chart one's way in any scientific discipline, and this is especially true of connectionism: whenever you feel without ideas, reach for philosophy, and when you feel you do not have the tools, reach for mathematics. A little research in both can build an astounding career in any branch of science, and neural networks are no exception here.

This book ends here, and if you feel it has been a fantastic journey, then I am happy. This is only the beginning of your path to deep learning. I strongly encourage you to seek out knowledge¹ and never settle for the status quo. Always dismiss when someone says 'why are you doing this, this does not work' or 'you are not qualified to do this' or 'this is not relevant to your field' and continue to research and do your very best. A proverb I like very much² goes: *Every day, write something new. If you do not have anything new, write something old. If you do not have anything old, read something.* At one point, someone with a new brilliant mind will make a breakthrough. It will be hard, and there will be a lot of resistance, and the resistance will take weird forms. But try to find solace in this: neural networks are a symbol of struggle, the struggle of pulling yourself up from rock-bottom, falling again and again, and finally reaching the stars against all odds. The life of the father of neural networks was an omen of all the future struggles. So, remember the story of Walter Pitts, the philosophical logician, the teenager who hid in the library to read the *Principia*, the student who tried to learn from the best, the person who walked out of life straight into the annals of history, the inconsolable man who tried to redeem the world with logic. Let his story be an inspiration.

Reference

1. Y. Bengio, Learning deep architectures for AI. *Found. Trends Mach. Learn.* **2**(1), 1–127 (2009)

¹Books, journal articles, Arxiv, Coursera, Udacity, Udemy, etc—there is a vast universe of resources out there.

²I do not know whose proverb it is, but I do know it was someone's, and I would be very grateful if a reader who knows the author contacts me.