Sandro Skansi

Introduction to Deep Learning

From Logical Calculus to Artificial Intelligence

Springer International Publishing, 2018

This book is published as a part of "Undergraduate Topics in Computer Science" series, published by Springer Publishing Company, which is a series dedicated to "delivering instructional content for undergraduates studying in all areas of computing and information science. (…) The texts are all authored by established experts in their fields and reviewed by an international advisory board". This textbook is intended as an introduction to deep learning topics for graduate and advanced undergraduate students of computer science, cognitive science and mathematics, as well as scholars and scientist from linguistics, logic, philosophy, and psychology. The book is divided into eleven paragraphs and a preface, discussing about the very basic elements of machine learning as well as mathematical and computational prerequisites for deep learning.

In the preface of the book, Skansi states that this book "does not contain any new scientific results, that his only contribution is to compile existing knowledge and try to explain it with his own examples and intuition" (p. V). And that is one of many strong points of this book. It does not lose its focus trying to explicate and introduce new ideas or deconstruct old ones under the pretence of contextual digressions, complicating the very thing it is trying to simplify. The author uses intuitive and easy to understand examples that can be very helpful to a reader that comes from a socio-humanistic scientific background, in grasping the concepts such as convolutional neural networks, the notion of distributed representation, neural Turing machine, memory networks, autoencoders, etc. For example, the author uses Kendo, a Japanese martial art similar to fencing, as a guideline on how to learn Deep learning.

The book starts with a chapter called "From Logic to Cognitive Science", on historical background and development of AI and it's undeniable ties to philosophy. It analyses the influence of Gottfried Leibniz, John Stuart Mill, Alan Turing and George Boole, then the work and exploits of McCulloch, Pitts, Lettvin and Weiner, and their connections with Russell and Carnap. Then the reader is introduced to a scholar responsible for the first true learning rule for neural networks, Frank Rosenblatt, and two people who discovered one of the great problems of for neural networks, the XOR problems. It shows the gradual development of deep learning or neural networks, the problems it encountered on the path to acknowledged, respected discipline, with exponentially growing impact and significance, not just for scientific purposes, but industrial, everyday appliance.

Chapter two, titled "Mathematical and Computational Prerequisites", gives the essential mathematical preliminaries, so that any reader could understand the chapters that will ensue. Deep learning cannot be understood without understanding its main engine called *backpropagation*, with its intrinsic featurette, *gradient descent*. The first part of this chapter is about derivations, which would serve as a helping tool for a reader as he tackles gradient and gradient descent. It also gives a brief overview of logic and Turing machines aimed at better understanding some of the concepts the reader will encounter in the next chapters. The last section is a brief, but sufficient introduction to Python, a language that will be used in the examples in the book.

The author starts with *machine learning* in chapter three, called "Machine Learning Basics". As the title suggests, this chapter deals with the basics of machine learning, nonetheless basics that are not part of deep learning but are carefully placed so that eventually a reader's better understanding of deep learning could be deriving from these basics. In a nutshell, this chapter addresses supervised learning, its terminology, and unsupervised learning. Classifiers such as naive Bayes, simplest neural networks and logistic regression as the essence of deep learning are explained in a detailed manner. Among other things mentioned in this chapter, the reader is presented with two techniques of unsupervised learning, K-means and principal component analysis (PCA).

In chapter four, titled "Feedforward Neural Networks", the author focuses on how shallow neural network works (or "simple feedforward neural networks", as the author is calling them). This is all an introduction to backpropagation as the core method of learning 'deep learning'. The shallow neural network is in a sense a core of deep learning, and what deep learning is all about, is fixing problems that are encountered when there is an attempt to add more layers to a shallow neural network. There is a detailed account

of Rosenblatt's perception rule and the Delta rule as its alternative. It tackles the problem of representation abstract and graphical objects as mathematical objects. The error function is explained and their role in the whole system. In this chapter, as well as throughout the entire book, the author refers the reader to many books that can help the reader with refining his knowledge on neural networks.

In a chapter called "Modifications and Extensions to a Feed-Forward Neural Network", Skansi deals with the problem of local minima as well as with modifications and extensions to simple feed-forward neural networks. Furthermore, a reader can learn about the basic concepts of regularisation, learning rate and momentum as a techniques against local minima, stochastic gradient, etc. We could say that this chapter, with its concepts and ideas, is a prelude to deep learning itself.

In chapter six, "Convolutional Neural Networks", the reader is introduced to convolutional neural networks, first presented by Yann LeCun and others in 1998. This concept stems from ideas of David H. Hubel and Torsten Weisel, presented in their 1968 seminal paper which won them the Nobel prize in Physiology and Medicine in 1981. The chapter starts with defining 1D and 2D convolutional layers as the natural extension of logistic regression, and among other things, details receptive field as a core concept of any convolutional architecture, and the ideas of padding in the visual setting and pooling in the general setting. Featured is also simple to use but powerful deep learning library for Python code called Keras, commented and presented in detail, so that even a complete novice to a programming world be able to understand it.

Chapter seven, "Recurrent Neural Networks", tackles recurrent connections to a simple feed-forward neural network. A recurrent neural network (RNN) is a class of artificial neural networks where connections between nodes form a direct graph along a temporal sequence. This allows it to exhibit temporal dynamic behaviour. In the same chapter, reader is presented with lots of illustrations on the roles of hidden states in the setting of a simple recurrent network, both with Elman and Jordan units, long short-term memory (LSTM), example of a recurrent neural network for learning text without labels on a world level with very detailed explanation of code.

In the next two chapters, the author turns his attention to unsupervised deep learning or learning distributed representations. The author starts by referencing the reader to chapter three and the discussion on PCA as a form of learning distributed representation. However, in this chapter, the author tries to fill the part that was left out in chapter three, after which the reader is introduced with the main idea of this chapter, the unsupervised learning technique for deep learning, the autoencoder. The chapter ends with explaining the autoencoder that can learn how to draw cats from watching unlabelled videos.

Chapter nine – "Neural Language Models" – explains that a Language Model assigns probability values to a sequence of words and that neural language models are distributed representations of words and sentences, in other words, they are numerical vectors. The word embedding is any method which converts word in numbers, and this chapter describes it with Word2vec algorithm, in both Skip-gram and CBOW variant whereas the latter is explored in Python code. In this chapter, there is also a simple reasoning calculus that is done with word analogies, encapsulating neural networking without symbolic manipulation.

In chapter ten, titled "An Overview of Different Neural Network Architectures", the focus is on a specific class of neural networks, an energy-based model. The simplest one is the Hopfield Network, which is made of neurons which are connected with weights. Each neuron has a threshold and a denotation. Once a network is assembled, the training can start. After the energy-based model, attention is shifted to memory-based models such as neural Turing-machines (NTM) and the memory networks (MemNN). The NTM is similar to the Turing-machine, but in the NTM, the objective is to do soft computation based on trainable components, so those components should learn how to do it well. Memory networks are all neural networks and their components are a memory, input feature map, updater, output feature map, and a responder. All of these connections are well illustrated in a detailed fashion in this chapter. Also, the question regarding surpassing benchmarks made by memory networks is asked in this chapter, more specifically, how to get a neural network to reason.

The last chapter of this book is a short reiteration of concepts, notions, and ideas discussed in this book as well as incomplete but more importantly, interdisciplinary list of interesting open research problems in deep learning and connectionism, which goes to show how diverse the deep learning research field can be. Some of the problems featured in this list are can we find new and better activation functions, can reasoning be learned or can we formalise necessary analogies?

The thing amiss in this book is the fact that it does not address Generative Adversarial

Networks (GAN) at all. GAN are deep neural networks architectures compromised of two nets, pitting one against the other. GAN's potential is huge because they can learn to mimic any distribution of data, which means that GANs can be taught to create worlds similar to our own. Nonetheless, the thing that sets apart this book from the other that discuss the same topic is that this book is fully accessible and easy to understand to those students that come from a social science and humanities background.

Even though this book is mostly about technical aspects and problems regarding the diverse field of deep learning, it is also in the same sense, about the need for close cooperation between different branches of science and their respective fields. In this case, the reader is referenced to connectionism, an approach in the fields of cognitive science that strives to explain mental phenomena using artificial neural networks. Also, deep learning is known in the philosophy of language for almost 30 years under the same name, *connectionism*. After reading this book, it becomes transparent that the lines between one field and another gets hard to denote when it comes to valid ideas and people behind those ideas in the field of deep learning. It could come from any field that has some relation to cognitive science or deep learning (anthropology, philosophy, linguistics, neuroscience, psychology and computer science) or even so, that ground-breaking, astonishing idea could come from some AI aficionado with no formal connection to any of the previously mentioned fields of science. And that is the subtext and a major strength of this book. It is intended for anyone willing to get started or deepen his knowledge on the matter. Moreover, with the book's concise and unexpectedly entertaining narrative, with intuitive examples and well-placed illustrations, it will serve as an excellent stepping stone for any further explorations and research regarding the topic of deep learning. Additionally, as this book showed the importance of interdisciplinary cooperation, it also tried to evoke the significance of philosophy, especially logic, as one of the pillars not only of deep learning and AI but any worthwhile academic, scientific or industrial endeavour. As the author clearly states:

"… Philosophy and mathematics are the two oldest branches of science, old as civilization itself, and most of the science can be seen as a gradual transition from philosophy to mathematics (…)." (pp. 186–187)

Filip Šoljić

Gottfried Küenzlen

Die Entzauberung der Welt

Studien zu Kultur, Gesellschaft und Religion in der Moderne

Lit Verlag, Berlin 2019

Die Philosophen streben in ihren Betrachtungen und Überlegungen an, die Dinge und Phänomene in der Welt zu verstehen und zu erklären, sie auf Begriffe zu bringen und in einen logischen Zusammenhang zu ordnen. Wie es schon Thomas von Aquin in seinem Kommentar der Aristotelischen *Metaphysik* feststellte, die Hauptaufgabe der Weisheit Liebenden sei, die Ordnung in den Dingen zu finden und alle anderen nach diesem Geordnet-Sein zu leiten: *Nam sapientis est alios ordinare* („Unde necesse est, quod una earum sit aliarum omnium rectrix, quae nomen sapientiae recte vindicat. Nam sapientis est alios ordinare." – Thomas Aquinas, *Commentary on the Metaphysics*, translated by John P. Rowan, Chicago 1961, https://dhspriory.org/thomas/Metaphysics.htm). So versuchen sie ebenfalls, diese logische Ordnung in der modernen Welt zu entdecken. Aber die Moderne lässt sich nur schwer auf die begriffliche Logik einfach zurückführen. Die wirkliche wie auch die ideale Geschichte der Moderne sei „vielschichtig und vielgesichtig", lautet die Ausgangsthese des Verfassers dieses denkerisch anregenden und an subtilen Reflexionen reichen Buches.

Küenzlen geht, in einem bestimmten Anklang an die Denker der Postmoderne, davon aus, dass die Zeit der großen Geschichten auch in der Philosophie vorbei sei. Aber agesehen davon, ob wir über Moderne oder Postmoderne reden, lässt sich unsere Welt in der Tat kaum mehr mit den Großerzählungen erklären. Der Grund dafür liegt in der eindeutigen Tatsache, dass der traditionelle metaphysische Begriffskomplex infrage gestellt worden ist. Die Erklärungskraft der bisherigen wissenschaftlichen Signaturen ist zur Neige gegangen. Die Zeit der Aufklärung und des Fortschrittsoptimismus ist weit hinter uns geblieben. Aber eine Großerzählung hat im gewissen Sinne doch ihre Geltung behalten. Diese ist die Rede von der „Entzauberung der Welt".

Es handelt sich allerdings um eine der Großerzählungen, die uns heute mit großem Nachdruck anspricht. Und gerade davon zeugt das vorliegende Buch. Der Verfasser setzt sie in den Titel des Buches und widmet ihr den Großteil der Aufmerksamkeit. Ihre theo-