



## Review

## Deep learning in neural networks: An overview



Jürgen Schmidhuber

The Swiss AI Lab IDSIA, Istituto Dalle Molle di Studi sull'Intelligenza Artificiale, University of Lugano &amp; SUPSI, Galleria 2, 6928 Manno-Lugano, Switzerland

## ARTICLE INFO

## Article history:

Received 2 May 2014

Received in revised form 12 September 2014

Accepted 14 September 2014

Available online 13 October 2014

## Keywords:

Deep learning

Supervised learning

Unsupervised learning

Reinforcement learning

Evolutionary computation

## ABSTRACT

In recent years, deep artificial neural networks (including recurrent ones) have won numerous contests in pattern recognition and machine learning. This historical survey compactly summarizes relevant work, much of it from the previous millennium. Shallow and Deep Learners are distinguished by the depth of their *credit assignment paths*, which are chains of possibly learnable, causal links between actions and effects. I review deep supervised learning (also recapitulating the history of backpropagation), unsupervised learning, reinforcement learning & evolutionary computation, and indirect search for short programs encoding deep and large networks.

© 2014 Published by Elsevier Ltd.

## Contents

1. Introduction to Deep Learning (DL) in Neural Networks (NNs).....	86
2. Event-oriented notation for activation spreading in NNs .....	87
3. Depth of Credit Assignment Paths (CAPs) and of problems .....	88
4. Recurring themes of Deep Learning.....	88
4.1. Dynamic programming for Supervised/Reinforcement Learning (SL/RL).....	88
4.2. Unsupervised Learning (UL) facilitating SL and RL.....	89
4.3. Learning hierarchical representations through deep SL, UL, RL .....	89
4.4. Occam's razor: compression and Minimum Description Length (MDL) .....	89
4.5. Fast Graphics Processing Units (GPUs) for DL in NNs.....	89
5. Supervised NNs, some helped by unsupervised NNs.....	89
5.1. Early NNs since the 1940s (and the 1800s).....	90
5.2. Around 1960: visual cortex provides inspiration for DL (Sections 5.4, 5.11) .....	90
5.3. 1965: deep networks based on the Group Method of Data Handling.....	90
5.4. 1979: convolution + weight replication + subsampling (Neocognitron).....	90
5.5. 1960–1981 and beyond: development of backpropagation (BP) for NNs .....	90
5.5.1. BP for weight-sharing feedforward NNs (FNNs) and recurrent NNs (RNNs).....	91
5.6. Late 1980s–2000 and beyond: numerous improvements of NNs .....	91
5.6.1. Ideas for dealing with long time lags and deep CAPs.....	91
5.6.2. Better BP through advanced gradient descent (compare Section 5.24).....	92
5.6.3. Searching for simple, low-complexity, problem-solving NNs (Section 5.24).....	92
5.6.4. Potential benefits of UL for SL (compare Sections 5.7, 5.10, 5.15).....	92
5.7. 1987: UL through Autoencoder (AE) hierarchies (compare Section 5.15).....	93
5.8. 1989: BP for convolutional NNs (CNNs, Section 5.4).....	93
5.9. 1991: Fundamental Deep Learning Problem of gradient descent .....	93
5.10. 1991: UL-based history compression through a deep stack of RNNs.....	94
5.11. 1992: Max-Pooling (MP): towards MPCNNs (compare Sections 5.16, 5.19) .....	94

E-mail address: [juergen@idsia.ch](mailto:juergen@idsia.ch).

5.12.	1994: early contest-winning NNs.....	95
5.13.	1995: supervised recurrent very Deep Learner (LSTM RNN).....	95
5.14.	2003: more contest-winning/record-setting NNs; successful deep NNs.....	96
5.15.	2006/7: UL for deep belief networks/AE stacks fine-tuned by BP.....	96
5.16.	2006/7: improved CNNs/GPU-CNNs/BP for MPCNNs/LSTM stacks.....	96
5.17.	2009: first official competitions won by RNNs, and with MPCNNs.....	97
5.18.	2010: plain backprop (+ distortions) on GPU breaks MNIST record.....	97
5.19.	2011: MPCNNs on GPU achieve superhuman vision performance.....	97
5.20.	2011: Hessian-free optimization for RNNs.....	98
5.21.	2012: first contests won on ImageNet, object detection, segmentation.....	98
5.22.	2013–: more contests and benchmark records.....	98
5.23.	Currently successful techniques: LSTM RNNs and GPU-MPCNNs.....	99
5.24.	Recent tricks for improving SL deep NNs (compare Sections 5.6.2, 5.6.3).....	99
5.25.	Consequences for neuroscience.....	100
5.26.	DL with spiking neurons?.....	100
6.	DL in FNNs and RNNs for Reinforcement Learning (RL).....	100
6.1.	RL through NN world models yields RNNs with deep CAPs.....	100
6.2.	Deep FNNs for traditional RL and Markov Decision Processes (MDPs).....	101
6.3.	Deep RL RNNs for partially observable MDPs (POMDPs).....	101
6.4.	RL facilitated by deep UL in FNNs and RNNs.....	102
6.5.	Deep hierarchical RL (HRL) and subgoal learning with FNNs and RNNs.....	102
6.6.	Deep RL by direct NN search/policy gradients/evolution.....	102
6.7.	Deep RL by indirect policy search/compressed NN search.....	103
6.8.	Universal RL.....	103
7.	Conclusion and outlook.....	103
	Acknowledgments.....	104
	References.....	104

## Preface

This is the preprint of an invited *Deep Learning* (DL) overview. One of its goals is to assign credit to those who contributed to the present state of the art. I acknowledge the limitations of attempting to achieve this goal. The DL research community itself may be viewed as a continually evolving, deep network of scientists who have influenced each other in complex ways. Starting from recent DL results, I tried to trace back the origins of relevant ideas through the past half century and beyond, sometimes using “local search” to follow citations of citations backwards in time. Since not all DL publications properly acknowledge earlier relevant work, additional global search strategies were employed, aided by consulting numerous neural network experts. As a result, the present preprint mostly consists of references. Nevertheless, through an expert selection bias I may have missed important work. A related bias was surely introduced by my special familiarity with the work of my own DL research group in the past quarter-century. For these reasons, this work should be viewed as merely a snapshot of an ongoing credit assignment process. To help improve it, please do not hesitate to send corrections and suggestions to [juergen@idsia.ch](mailto:juergen@idsia.ch).

## 1. Introduction to Deep Learning (DL) in Neural Networks (NNs)

Which modifiable components of a learning system are responsible for its success or failure? What changes to them improve performance? This has been called the *fundamental credit assignment problem* (Minsky, 1963). There are general credit assignment methods for *universal problem solvers* that are time-optimal in various theoretical senses (Section 6.8). The present survey, however, will focus on the narrower, but now commercially important, subfield of *Deep Learning* (DL) in *Artificial Neural Networks* (NNs).

A standard neural network (NN) consists of many simple, connected processors called neurons, each producing a sequence of real-valued activations. Input neurons get activated through sensors perceiving the environment, other neurons get activated through weighted connections from previously active neurons (details in Section 2). Some neurons may influence the environment by triggering actions. *Learning* or *credit assignment* is about finding

weights that make the NN exhibit *desired* behavior, such as driving a car. Depending on the problem and how the neurons are connected, such behavior may require long causal chains of computational stages (Section 3), where each stage transforms (often in a non-linear way) the aggregate activation of the network. Deep Learning is about accurately assigning credit across *many* such stages.

*Shallow* NN-like models with *few* such stages have been around for many decades if not centuries (Section 5.1). Models with several successive nonlinear layers of neurons date back at least to the 1960s (Section 5.3) and 1970s (Section 5.5). An efficient gradient descent method for teacher-based *Supervised Learning* (SL) in discrete, differentiable networks of arbitrary depth called *back-propagation* (BP) was developed in the 1960s and 1970s, and applied to NNs in 1981 (Section 5.5). BP-based training of *deep* NNs with *many* layers, however, had been found to be difficult in practice by the late 1980s (Section 5.6), and had become an explicit research subject by the early 1990s (Section 5.9). DL became practically feasible to some extent through the help of *Unsupervised Learning* (UL), e.g., Section 5.10 (1991), Section 5.15 (2006). The 1990s and 2000s also saw many improvements of purely supervised DL (Section 5). In the new millennium, deep NNs have finally attracted wide-spread attention, mainly by outperforming alternative machine learning methods such as kernel machines (Schölkopf, Burges, & Smola, 1998; Vapnik, 1995) in numerous important applications. In fact, since 2009, supervised deep NNs have won many official international pattern recognition competitions (e.g., Sections 5.17, 5.19, 5.21 and 5.22), achieving the first superhuman visual pattern recognition results in limited domains (Section 5.19, 2011). Deep NNs also have become relevant for the more general field of *Reinforcement Learning* (RL) where there is no supervising teacher (Section 6).

Both *feedforward* (acyclic) NNs (FNNs) and *recurrent* (cyclic) NNs (RNNs) have won contests (Sections 5.12, 5.14, 5.17, 5.19, 5.21, 5.22). In a sense, RNNs are the deepest of all NNs (Section 3)—they are general computers more powerful than FNNs, and can in principle create and process memories of arbitrary sequences of input patterns (e.g., Schmidhuber, 1990a; Siegelmann & Sontag, 1991). Unlike traditional methods for automatic sequential program synthesis (e.g., Balzer, 1985; Deville & Lau, 1994; Soloway,

Download English Version:

<https://daneshyari.com/en/article/6863319>

Download Persian Version:

<https://daneshyari.com/article/6863319>

[Daneshyari.com](https://daneshyari.com)