WIKIPEDIA

Artificial neural network

Artificial neural networks (**ANNs**), usually simply called **neural networks** (**NNs**) or, more simply yet, **neural nets**,[1] are computing systems inspired by the biological neural networks that constitute animal brains.[2]

An ANN is based on a collection of connected units or nodes called artificial neurons, which loosely model the neurons in a biological brain. Each connection, like the synapses in a biological brain, can transmit a signal to other neurons. An artificial neuron receives signals then processes them and can signal neurons connected to it. The "signal" at a connection is a real number, and the output of each neuron is computed by some non-linear function of the sum of its inputs. The connections are called *edges*. Neurons and edges typically have a *weight* that adjusts as learning proceeds. The weight increases or decreases the strength of the signal at a connection. Neurons may have a threshold such that a signal is sent only if the aggregate signal crosses that threshold. Typically, neurons are aggregated into layers. Different layers may perform different transformations on their inputs. Signals travel from the first layer (the input layer), to the last layer (the output layer), possibly after traversing the layers multiple times.

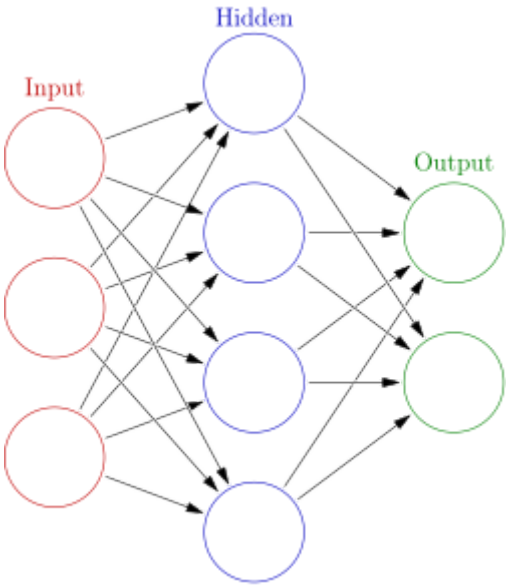

An artificial neural network is an interconnected group of nodes, inspired by a simplification of neurons in a brain. Here, each circular node represents an artificial neuron and an arrow represents a connection from the output of one artificial neuron to the input of another.

Contents

Training
History
Models
- Artificial neurons
- Organization
- Hyperparameter
- Learning
- Learning paradigms
- Stochastic neural network
- Other

Types
Network design
Use
Applications

Theoretical properties
- Computational power
- Capacity
- Convergence
- Generalization and statistics

Criticism
- Training
- Theory
- Hardware
- Practical counterexamples
- Hybrid approaches

Gallery

See also

Notes

References

Bibliography

Training

Neural networks learn (or are trained) by processing examples, each of which contains a known "input" and "result," forming probability-weighted associations between the two, which are stored within the data structure of the net itself. The training of a neural network from a given example is usually conducted by determining the difference between the processed output of the network (often a prediction) and a target output. This difference is the error. The network then adjusts its weighted associations according to a learning rule and using this error value. Successive adjustments will cause the neural network to produce output which is increasingly similar to the target output. After a sufficient number of these adjustments the training can be terminated based upon certain criteria. This is known as supervised learning.

Such systems "learn" to perform tasks by considering examples, generally without being programmed with task-specific rules. For example, in image recognition, they might learn to identify images that contain cats by analyzing example images that have been manually labeled as "cat" or "no cat" and using the results to identify cats in other images. They do this without any prior knowledge of cats, for example, that they have fur, tails, whiskers, and cat-like faces. Instead, they automatically generate identifying characteristics from the examples that they process.

History

Warren McCulloch and Walter Pitts[3] (1943) opened the subject by creating a computational model for neural networks.[4] In the late 1940s, D. O. Hebb[5] created a learning hypothesis based on the mechanism of neural plasticity that became known as Hebbian learning. Farley and Wesley A. Clark[6] (1954) first used computational machines, then called "calculators", to simulate a Hebbian network. In 1958, psychologist Frank Rosenblatt invented the perceptron, the first artificial neural network,[7][8][9][10] funded by the United States Office of Naval Research.[11] The first functional networks with many layers were published by Ivakhnenko and Lapa in 1965, as the Group Method of Data Handling.[12][13][14] The basics of continuous backpropagation[12][15][16][17] were derived in the context of control theory by Kelley[18] in 1960 and by Bryson in 1961,[19] using principles of dynamic programming. Thereafter research stagnated

following Minsky and Papert (1969),[20] who discovered that basic perceptrons were incapable of processing the exclusive-or circuit and that computers lacked sufficient power to process useful neural networks.

In 1970, Seppo Linnainmaa published the general method for automatic differentiation (AD) of discrete connected networks of nested differentiable functions.[21][22] In 1973, Dreyfus used backpropagation to adapt parameters of controllers in proportion to error gradients.[23] Werbos's (1975) backpropagation algorithm enabled practical training of multi-layer networks. In 1982, he applied Linnainmaa's AD method to neural networks in the way that became widely used.[15][24]

The development of metal–oxide–semiconductor (MOS) very-large-scale integration (VLSI), in the form of complementary MOS (CMOS) technology, enabled increasing MOS transistor counts in digital electronics. This provided more processing power for the development of practical artificial neural networks in the 1980s.[25]

In 1986 Rumelhart, Hinton and Williams showed that backpropagation learned interesting internal representations of words as feature vectors when trained to predict the next word in a sequence.[26]

From 1988 onward,[27][28] the use of neural networks transformed the field of protein structure prediction, in particular when the first cascading networks were trained on *profiles* (matrices) produced by multiple sequence alignments.[29]

In 1992, max-pooling was introduced to help with least-shift invariance and tolerance to deformation to aid 3D object recognition.[30][31][32] Schmidhuber adopted a multi-level hierarchy of networks (1992) pre-trained one level at a time by unsupervised learning and fine-tuned by backpropagation.[33]

Neural networks' early successes included predicting the stock market and in 1995 a (mostly) self-driving car.[a][34]

Geoffrey Hinton et al. (2006) proposed learning a high-level representation using successive layers of binary or real-valued latent variables with a restricted Boltzmann machine[35] to model each layer. In 2012, Ng and Dean created a network that learned to recognize higher-level concepts, such as cats, only from watching unlabeled images.[36] Unsupervised pre-training and increased computing power from GPUs and distributed computing allowed the use of larger networks, particularly in image and visual recognition problems, which became known as "deep learning".[37]

Ciresan and colleagues (2010)[38] showed that despite the vanishing gradient problem, GPUs make backpropagation feasible for many-layered feedforward neural networks.[39] Between 2009 and 2012, ANNs began winning prizes in image recognition contests, approaching human level performance on various tasks, initially in pattern recognition and handwriting recognition.[40][41] For example, the bi-directional and multi-dimensional long short-term memory (LSTM)[42][43] of Graves et al. won three competitions in connected handwriting recognition in 2009 without any prior knowledge about the three languages to be learned.[42][43]

Ciresan and colleagues built the first pattern recognizers to achieve human-competitive/superhuman performance[44] on benchmarks such as traffic sign recognition (IJCNN 2012).

Models

ANNs began as an attempt to exploit the architecture of the human brain to perform tasks that conventional algorithms had little success with. They soon reoriented towards improving empirical results, mostly abandoning attempts to remain true to their biological precursors. Neurons are connected to each other in

various patterns, to allow the output of some neurons to become the input of others. The network forms a directed, weighted graph.[45]

An artificial neural network consists of a collection of simulated neurons. Each neuron is a node which is connected to other nodes via links that correspond to biological axon-synapse-dendrite connections. Each link has a weight, which determines the strength of one node's influence on another.[46]

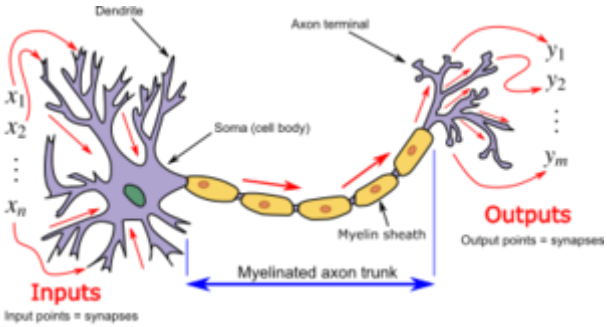

Neuron and myelinated axon, with signal flow from inputs at dendrites to outputs at axon terminals

Artificial neurons

ANNs are composed of artificial neurons which are conceptually derived from biological neurons. Each artificial neuron has inputs and produces a single output which can be sent to multiple other neurons.[47] The inputs can be the feature values of a sample of external data, such as images or documents, or they can be the outputs of other neurons. The outputs of the final *output neurons* of the neural net accomplish the task, such as recognizing an object in an image.

To find the output of the neuron we take the weighted sum of all the inputs, weighted by the *weights* of the *connections* from the inputs to the neuron. We add a *bias* term to this sum.[48] This weighted sum is sometimes called the *activation*. This weighted sum is then passed through a (usually nonlinear) activation function to produce the output. The initial inputs are external data, such as images and documents. The ultimate outputs accomplish the task, such as recognizing an object in an image.[49]

Organization

The neurons are typically organized into multiple layers, especially in deep learning. Neurons of one layer connect only to neurons of the immediately preceding and immediately following layers. The layer that receives external data is the *input layer*. The layer that produces the ultimate result is the *output layer*. In between them are zero or more *hidden layers*. Single layer and unlayered networks are also used. Between two layers, multiple connection patterns are possible. They can be 'fully connected', with every neuron in one layer connecting to every neuron in the next layer. They can be *pooling*, where a group of neurons in one layer connect to a single neuron in the next layer, thereby reducing the number of neurons in that layer.[50] Neurons with only such connections form a directed acyclic graph and are known as *feedforward networks*.[51] Alternatively, networks that allow connections between neurons in the same or previous layers are known as *recurrent networks*.[52]

Hyperparameter

A hyperparameter is a constant parameter whose value is set before the learning process begins. The values of parameters are derived via learning. Examples of hyperparameters include learning rate, the number of hidden layers and batch size.[53] The values of some hyperparameters can be dependent on those of other hyperparameters. For example, the size of some layers can depend on the overall number of layers.

Learning

Learning is the adaptation of the network to better handle a task by considering sample observations. Learning involves adjusting the weights (and optional thresholds) of the network to improve the accuracy of the result. This is done by minimizing the observed errors. Learning is complete when examining additional observations does not usefully reduce the error rate. Even after learning, the error rate typically does not reach 0. If after learning, the error rate is too high, the network typically must be redesigned. Practically this is done by defining a cost function that is evaluated periodically during learning. As long as its output continues to decline, learning continues. The cost is frequently defined as a statistic whose value can only be approximated. The outputs are actually numbers, so when the error is low, the difference between the output (almost certainly a cat) and the correct answer (cat) is small. Learning attempts to reduce the total of the differences across the observations. Most learning models can be viewed as a straightforward application of optimization theory and statistical estimation.[45][54]

Learning rate

The learning rate defines the size of the corrective steps that the model takes to adjust for errors in each observation.[55] A high learning rate shortens the training time, but with lower ultimate accuracy, while a lower learning rate takes longer, but with the potential for greater accuracy. Optimizations such as Quickprop are primarily aimed at speeding up error minimization, while other improvements mainly try to increase reliability. In order to avoid oscillation inside the network such as alternating connection weights, and to improve the rate of convergence, refinements use an adaptive learning rate that increases or decreases as appropriate.[56] The concept of momentum allows the balance between the gradient and the previous change to be weighted such that the weight adjustment depends to some degree on the previous change. A momentum close to 0 emphasizes the gradient, while a value close to 1 emphasizes the last change.

Cost function

While it is possible to define a cost function ad hoc, frequently the choice is determined by the function's desirable properties (such as convexity) or because it arises from the model (e.g. in a probabilistic model the model's posterior probability can be used as an inverse cost).

Backpropagation

Backpropagation is a method used to adjust the connection weights to compensate for each error found during learning. The error amount is effectively divided among the connections. Technically, backprop calculates the gradient (the derivative) of the cost function associated with a given state with respect to the weights. The weight updates can be done via stochastic gradient descent or other methods, such as Extreme Learning Machines,[57] "No-prop" networks,[58] training without backtracking,[59] "weightless" networks,[60][61] and non-connectionist neural networks.

Learning paradigms

Machine learning is commonly separated into three main learning paradigms, supervised learning, unsupervised learning and reinforcement learning.[62] Each corresponds to a particular learning task.

Supervised learning

Supervised learning uses a set of paired inputs and desired outputs. The learning task is to produce the desired output for each input. In this case the cost function is related to eliminating incorrect deductions.[63] A commonly used cost is the mean-squared error, which tries to minimize the average squared error between the network's output and the desired output. Tasks suited for supervised learning are pattern recognition (also known as classification) and regression (also known as function approximation). Supervised learning is also applicable to sequential data (e.g., for hand writing, speech and gesture recognition). This can be thought of as learning with a "teacher", in the form of a function that provides continuous feedback on the quality of solutions obtained thus far.

Unsupervised learning

In unsupervised learning, input data is given along with the cost function, some function of the data x and the network's output. The cost function is dependent on the task (the model domain) and any *a priori* assumptions (the implicit properties of the model, its parameters and the observed variables). As a trivial example, consider the model $f(x) = a$ where a is a constant and the cost $C = E[(x - f(x))^2]$. Minimizing this cost produces a value of a that is equal to the mean of the data. The cost function can be much more complicated. Its form depends on the application: for example, in compression it could be related to the mutual information between x and $f(x)$, whereas in statistical modeling, it could be related to the posterior probability of the model given the data (note that in both of those examples those quantities would be maximized rather than minimized). Tasks that fall within the paradigm of unsupervised learning are in general estimation problems; the applications include clustering, the estimation of statistical distributions, compression and filtering.

Reinforcement learning

In applications such as playing video games, an actor takes a string of actions, receiving a generally unpredictable response from the environment after each one. The goal is to win the game, i.e., generate the most positive (lowest cost) responses. In reinforcement learning, the aim is to weight the network (devise a policy) to perform actions that minimize long-term (expected cumulative) cost. At each point in time the agent performs an action and the environment generates an observation and an instantaneous cost, according to some (usually unknown) rules. The rules and the long-term cost usually only can be estimated. At any juncture, the agent decides whether to explore new actions to uncover their costs or to exploit prior learning to proceed more quickly.

Formally the environment is modeled as a Markov decision process (MDP) with states $s_1, \ldots, s_n \in S$ and actions $a_1, \ldots, a_m \in A$. Because the state transitions are not known, probability distributions are used instead: the instantaneous cost distribution $P(c_t|s_t)$, the observation distribution $P(x_t|s_t)$ and the transition distribution $P(s_{t+1}|s_t, a_t)$, while a policy is defined as the conditional distribution over actions given the observations. Taken together, the two define a Markov chain (MC). The aim is to discover the lowest-cost MC.

ANNs serve as the learning component in such applications.[64][65] Dynamic programming coupled with ANNs (giving neurodynamic programming)[66] has been applied to problems such as those involved in vehicle routing,[67] video games, natural resource management[68][69] and medicine[70] because of ANNs ability to mitigate losses of accuracy even when reducing the discretization grid density for numerically approximating the solution of control problems. Tasks that fall within the paradigm of reinforcement learning are control problems, games and other sequential decision making tasks.

Self-learning

Self-learning in neural networks was introduced in 1982 along with a neural network capable of self-learning named Crossbar Adaptive Array (CAA).[71] It is a system with only one input, situation s, and only one output, action (or behavior) a. It has neither external advice input nor external reinforcement input from the environment. The CAA computes, in a crossbar fashion, both decisions about actions and emotions (feelings) about encountered situations. The system is driven by the interaction between cognition and emotion.[72] Given the memory matrix, W =||w(a,s)||, the crossbar self-learning algorithm in each iteration performs the following computation:

```
In situation s perform action a;
Receive consequence situation s';
Compute emotion of being in consequence situation v(s');
Update crossbar memory w'(a,s) = w(a,s) + v(s').
```

The backpropagated value (secondary reinforcement) is the emotion toward the consequence situation. The CAA exists in two environments, one is behavioral environment where it behaves, and the other is genetic environment, where from it initially and only once receives initial emotions about to be encountered situations in the behavioral environment. Having received the genome vector (species vector) from the genetic environment, the CAA will learn a goal-seeking behavior, in the behavioral environment that contains both desirable and undesirable situations.[73]

Neuroevolution

Neuroevolution can create neural network topologies and weights using evolutionary computation. It is competitive with sophisticated gradient descent approaches. One advantage of neuroevolution is that it may be less prone to get caught in "dead ends".[74]

Stochastic neural network

Stochastic neural networks originating from Sherrington–Kirkpatrick models are a type of artificial neural network built by introducing random variations into the network, either by giving the network's artificial neurons stochastic transfer functions, or by giving them stochastic weights. This makes them useful tools for optimization problems, since the random fluctuations help the network escape from local minima.[75]

Other

In a Bayesian framework, a distribution over the set of allowed models is chosen to minimize the cost. Evolutionary methods,[76] gene expression programming,[77] simulated annealing,[78] expectation-maximization, non-parametric methods and particle swarm optimization[79] are other learning algorithms. Convergent recursion is a learning algorithm for cerebellar model articulation controller (CMAC) neural networks.[80][81]

Modes

Two modes of learning are available: stochastic and batch. In stochastic learning, each input creates a weight adjustment. In batch learning weights are adjusted based on a batch of inputs, accumulating errors over the batch. Stochastic learning introduces "noise" into the process, using the local gradient calculated from one data point; this reduces the chance of the network getting stuck in local minima. However, batch

learning typically yields a faster, more stable descent to a local minimum, since each update is performed in the direction of the batch's average error. A common compromise is to use "mini-batches", small batches with samples in each batch selected stochastically from the entire data set.

Types

ANNs have evolved into a broad family of techniques that have advanced the state of the art across multiple domains. The simplest types have one or more static components, including number of units, number of layers, unit weights and topology. Dynamic types allow one or more of these to evolve via learning. The latter are much more complicated, but can shorten learning periods and produce better results. Some types allow/require learning to be "supervised" by the operator, while others operate independently. Some types operate purely in hardware, while others are purely software and run on general purpose computers.

Some of the main breakthroughs include: convolutional neural networks that have proven particularly successful in processing visual and other two-dimensional data;[82][83] long short-term memory avoid the vanishing gradient problem[84] and can handle signals that have a mix of low and high frequency components aiding large-vocabulary speech recognition,[85][86] text-to-speech synthesis,[87][15][88] and photo-real talking heads;[89] competitive networks such as generative adversarial networks in which multiple networks (of varying structure) compete with each other, on tasks such as winning a game[90] or on deceiving the opponent about the authenticity of an input.[91]

Network design

Neural architecture search (NAS) uses machine learning to automate ANN design. Various approaches to NAS have designed networks that compare well with hand-designed systems. The basic search algorithm is to propose a candidate model, evaluate it against a dataset and use the results as feedback to teach the NAS network.[92] Available systems include AutoML and AutoKeras.[93]

Design issues include deciding the number, type and connectedness of network layers, as well as the size of each and the connection type (full, pooling, ...).

Hyperparameters must also be defined as part of the design (they are not learned), governing matters such as how many neurons are in each layer, learning rate, step, stride, depth, receptive field and padding (for CNNs), etc.[94]

Use

Using Artificial neural networks requires an understanding of their characteristics.

- Choice of model: This depends on the data representation and the application. Overly complex models are slow learning.
- Learning algorithm: Numerous trade-offs exist between learning algorithms. Almost any algorithm will work well with the correct hyperparameters for training on a particular data set. However, selecting and tuning an algorithm for training on unseen data requires significant experimentation.
- Robustness: If the model, cost function and learning algorithm are selected appropriately, the resulting ANN can become robust.

ANN capabilities fall within the following broad categories:

- Function approximation, or regression analysis, including time series prediction, fitness approximation and modeling.
- Classification, including pattern and sequence recognition, novelty detection and sequential decision making.[95]
- Data processing, including filtering, clustering, blind source separation and compression.
- Robotics, including directing manipulators and prostheses.

Applications

Because of their ability to reproduce and model nonlinear processes, artificial neural networks have found applications in many disciplines. Application areas include system identification and control (vehicle control, trajectory prediction,[96] process control, natural resource management), quantum chemistry,[97] general game playing,[98] pattern recognition (radar systems, face identification, signal classification,[99] 3D reconstruction,[100] object recognition and more), sensor data analysis,[101] sequence recognition (gesture, speech, handwritten and printed text recognition[102]), medical diagnosis, finance[103] (e.g. automated trading systems), data mining, visualization, machine translation, social network filtering[104] and e-mail spam filtering. ANNs have been used to diagnose several types of cancers[105][106] and to distinguish highly invasive cancer cell lines from less invasive lines using only cell shape information.[107][108]

ANNs have been used to accelerate reliability analysis of infrastructures subject to natural disasters[109][110] and to predict foundation settlements.[111] ANNs have also been used for building black-box models in geoscience: hydrology,[112][113] ocean modelling and coastal engineering,[114][115] and geomorphology.[116] ANNs have been employed in cybersecurity, with the objective to discriminate between legitimate activities and malicious ones. For example, machine learning has been used for classifying Android malware,[117] for identifying domains belonging to threat actors and for detecting URLs posing a security risk.[118] Research is underway on ANN systems designed for penetration testing, for detecting botnets,[119] credit cards frauds[120] and network intrusions.

ANNs have been proposed as a tool to solve partial differential equations in physics[121][122][123] and simulate the properties of many-body open quantum systems.[124][125][126][127] In brain research ANNs have studied short-term behavior of individual neurons,[128] the dynamics of neural circuitry arise from interactions between individual neurons and how behavior can arise from abstract neural modules that represent complete subsystems. Studies considered long-and short-term plasticity of neural systems and their relation to learning and memory from the individual neuron to the system level.

Theoretical properties

Computational power

The multilayer perceptron is a universal function approximator, as proven by the universal approximation theorem. However, the proof is not constructive regarding the number of neurons required, the network topology, the weights and the learning parameters.

A specific recurrent architecture with rational-valued weights (as opposed to full precision real number-valued weights) has the power of a universal Turing machine,[129] using a finite number of neurons and standard linear connections. Further, the use of irrational values for weights results in a machine with super-Turing power.[130]

Capacity

A model's "capacity" property corresponds to its ability to model any given function. It is related to the amount of information that can be stored in the network and to the notion of complexity. Two notions of capacity are known by the community. The information capacity and the VC Dimension. The information capacity of a perceptron is intensively discussed in Sir David MacKay's book[131] which summarizes work by Thomas Cover.[132] The capacity of a network of standard neurons (not convolutional) can be derived by four rules[133] that derive from understanding a neuron as an electrical element. The information capacity captures the functions modelable by the network given any data as input. The second notion, is the VC dimension. VC Dimension uses the principles of measure theory and finds the maximum capacity under the best possible circumstances. This is, given input data in a specific form. As noted in,[131] the VC Dimension for arbitrary inputs is half the information capacity of a Perceptron. The VC Dimension for arbitrary points is sometimes referred to as Memory Capacity.[134]

Convergence

Models may not consistently converge on a single solution, firstly because local minima may exist, depending on the cost function and the model. Secondly, the optimization method used might not guarantee to converge when it begins far from any local minimum. Thirdly, for sufficiently large data or parameters, some methods become impractical.

Another issue worthy to mention is that training may cross some Saddle point which may lead the convergence to the wrong direction.

The convergence behavior of certain types of ANN architectures are more understood than others. When the width of network approaches to infinity, the ANN is well described by its first order Taylor expansion throughout training, and so inherits the convergence behavior of affine models.[135][136] Another example is when parameters are small, it is observed that ANNs often fits target functions from low to high frequencies. This behavior is referred to as the spectral bias, or frequency principle, of neural networks.[137][138][139][140] This phenomenon is the opposite to the behavior of some well studied iterative numerical schemes such as Jacobi method. Deeper neural networks have been observed to be more biased towards low frequency functions.[141]

Generalization and statistics

Applications whose goal is to create a system that generalizes well to unseen examples, face the possibility of over-training. This arises in convoluted or over-specified systems when the network capacity significantly exceeds the needed free parameters. Two approaches address over-training. The first is to use cross-validation and similar techniques to check for the presence of over-training and to select hyperparameters to minimize the generalization error.

The second is to use some form of *regularization*. This concept emerges in a probabilistic (Bayesian) framework, where regularization can be performed by selecting a larger prior probability over simpler models; but also in statistical learning theory, where the goal is to minimize over two quantities: the 'empirical risk' and the 'structural risk', which roughly corresponds to the error over the training set and the predicted error in unseen data due to overfitting.

Supervised neural networks that use a mean squared error (MSE) cost function can use formal statistical methods to determine the confidence of the trained model. The MSE on a validation set can be used as an estimate for variance. This value can then be used to calculate the confidence interval of network output,

assuming a normal distribution. A confidence analysis made this way is statistically valid as long as the output probability distribution stays the same and the network is not modified.

Confidence analysis of a neural network

By assigning a softmax activation function, a generalization of the logistic function, on the output layer of the neural network (or a softmax component in a component-based network) for categorical target variables, the outputs can be interpreted as posterior probabilities. This is useful in classification as it gives a certainty measure on classifications.

The softmax activation function is:

$$y_i = \frac{e^{x_i}}{\sum_{j=1}^{c} e^{x_j}}$$

Criticism

Training

A common criticism of neural networks, particularly in robotics, is that they require too much training for real-world operation. Potential solutions include randomly shuffling training examples, by using a numerical optimization algorithm that does not take too large steps when changing the network connections following an example, grouping examples in so-called mini-batches and/or introducing a recursive least squares algorithm for CMAC.[80]

Theory

A fundamental objection is that ANNs do not sufficiently reflect neuronal function. Backpropagation is a critical step, although no such mechanism exists in biological neural networks.[142] How information is coded by real neurons is not known. Sensor neurons fire action potentials more frequently with sensor activation and muscle cells pull more strongly when their associated motor neurons receive action potentials more frequently.[143] Other than the case of relaying information from a sensor neuron to a motor neuron, almost nothing of the principles of how information is handled by biological neural networks is known.

A central claim of ANNs is that they embody new and powerful general principles for processing information. These principles are ill-defined. It is often claimed that they are emergent from the network itself. This allows simple statistical association (the basic function of artificial neural networks) to be described as learning or recognition. In 1997, Alexander Dewdney commented that, as a result, artificial neural networks have a "something-for-nothing quality, one that imparts a peculiar aura of laziness and a distinct lack of curiosity about just how good these computing systems are. No human hand (or mind) intervenes; solutions are found as if by magic; and no one, it seems, has learned anything".[144] One response to Dewdney is that neural networks handle many complex and diverse tasks, ranging from autonomously flying aircraft[145] to detecting credit card fraud to mastering the game of Go.

Technology writer Roger Bridgman commented:

> Neural networks, for instance, are in the dock not only because they have been hyped to high heaven, (what hasn't?) but also because you could create a successful net without understanding how it worked: the bunch of numbers that captures its behaviour would in all probability be "an opaque, unreadable table...valueless as a scientific resource".
>
> In spite of his emphatic declaration that science is not technology, Dewdney seems here to pillory neural nets as bad science when most of those devising them are just trying to be good engineers. An unreadable table that a useful machine could read would still be well worth having.[146]

Biological brains use both shallow and deep circuits as reported by brain anatomy,[147] displaying a wide variety of invariance. Weng[148] argued that the brain self-wires largely according to signal statistics and therefore, a serial cascade cannot catch all major statistical dependencies.

Hardware

Large and effective neural networks require considerable computing resources.[149] While the brain has hardware tailored to the task of processing signals through a graph of neurons, simulating even a simplified neuron on von Neumann architecture may consume vast amounts of memory and storage. Furthermore, the designer often needs to transmit signals through many of these connections and their associated neurons – which require enormous CPU power and time.

Schmidhuber noted that the resurgence of neural networks in the twenty-first century is largely attributable to advances in hardware: from 1991 to 2015, computing power, especially as delivered by GPGPUs (on GPUs), has increased around a million-fold, making the standard backpropagation algorithm feasible for training networks that are several layers deeper than before.[12] The use of accelerators such as FPGAs and GPUs can reduce training times from months to days.[149]

Neuromorphic engineering or a physical neural network addresses the hardware difficulty directly, by constructing non-von-Neumann chips to directly implement neural networks in circuitry. Another type of chip optimized for neural network processing is called a Tensor Processing Unit, or TPU.[150]

Practical counterexamples

Analyzing what has been learned by an ANN is much easier than analyzing what has been learned by a biological neural network. Furthermore, researchers involved in exploring learning algorithms for neural networks are gradually uncovering general principles that allow a learning machine to be successful. For example, local vs. non-local learning and shallow vs. deep architecture.[151]

Hybrid approaches

Advocates of hybrid models (combining neural networks and symbolic approaches), claim that such a mixture can better capture the mechanisms of the human mind.[152]

Gallery

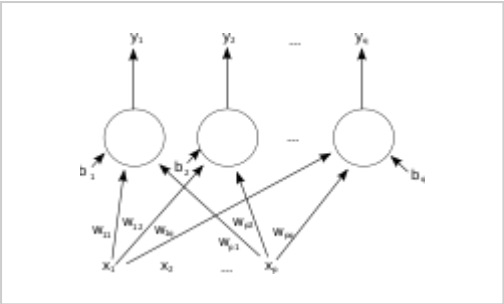

A single-layer feedforward artificial neural network. Arrows originating from x_2 are omitted for clarity. There are p inputs to this network and q outputs. In this system, the value of the qth output, y_q would be calculated as $y_q = K*(\sum(x_i * w_{iq}) - b_q)$

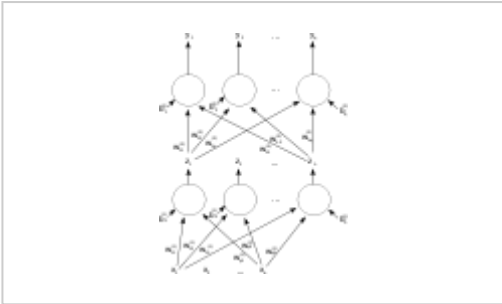

A two-layer feedforward artificial neural network.

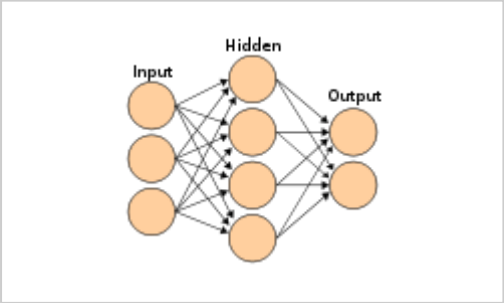

An artificial neural network.

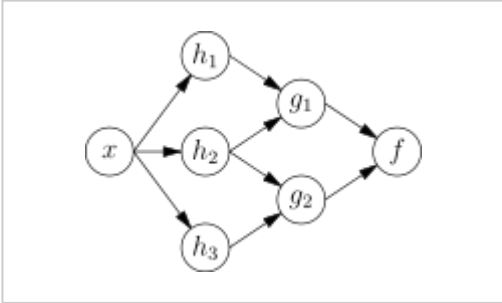

An ANN dependency graph.

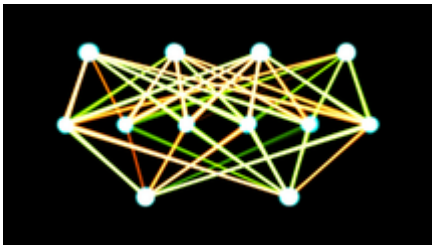

A single-layer feedforward artificial neural network with 4 inputs, 6 hidden and 2 outputs. Given position state and direction outputs wheel based control values.

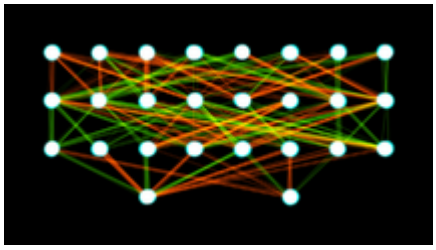

A two-layer feedforward artificial neural network with 8 inputs, 2x8 hidden and 2 outputs. Given position state, direction and other environment values outputs thruster based control values.

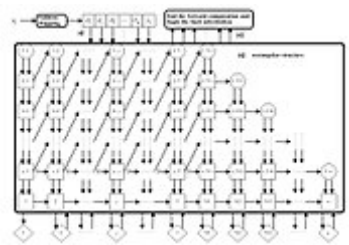

Parallel pipeline structure of CMAC neural network. This learning algorithm can converge in one step.

See also

- ADALINE
- Autoencoder
- Biologically inspired computing
- Blue Brain Project
- Catastrophic interference
- Cognitive architecture
- Connectionist expert system
- Connectomics
- Large width limits of neural networks
- Machine learning concepts
- Neural gas
- Neural network software
- Optical neural network
- Parallel distributed processing
- Recurrent neural networks
- Spiking neural network
- Tensor product network

Notes

a. Steering for the 1995 "No Hands Across America" required "only a few human assists".

References

1. Hardesty, Larry (14 April 2017). "Explained: Neural networks" (https://news.mit.edu/2017/explained-neural-networks-deep-learning-0414). MIT News Office. Retrieved 2 June 2022.
2. Yang, Z.R.; Yang, Z. (2014). *Comprehensive Biomedical Physics* (https://www.sciencedirect.com/topics/neuroscience/artificial-neural-network). Karolinska Institute, Stockholm, Sweden: Elsevier. p. 1. ISBN 978-0-444-53633-4.
3. McCulloch, Warren; Walter Pitts (1943). "A Logical Calculus of Ideas Immanent in Nervous Activity". *Bulletin of Mathematical Biophysics*. **5** (4): 115–133. doi:10.1007/BF02478259 (https://doi.org/10.1007%2FBF02478259).
4. Kleene, S.C. (1956). "Representation of Events in Nerve Nets and Finite Automata" (https://www.degruyter.com/view/books/9781400882618/9781400882618-002/9781400882618-002.xml). *Annals of Mathematics Studies*. No. 34. Princeton University Press. pp. 3–41. Retrieved 17 June 2017.
5. Hebb, Donald (1949). *The Organization of Behavior* (https://books.google.com/books?id=ddB4AgAAQBAJ). New York: Wiley. ISBN 978-1-135-63190-1.
6. Farley, B.G.; W.A. Clark (1954). "Simulation of Self-Organizing Systems by Digital Computer". *IRE Transactions on Information Theory*. **4** (4): 76–84. doi:10.1109/TIT.1954.1057468 (https://doi.org/10.1109%2FTIT.1954.1057468).
7. Haykin (2008) Neural Networks and Learning Machines, 3rd edition
8. Rosenblatt, F. (1958). "The Perceptron: A Probabilistic Model For Information Storage And Organization in the Brain". *Psychological Review*. **65** (6): 386–408. CiteSeerX 10.1.1.588.3775 (https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.588.3775). doi:10.1037/h0042519 (https://doi.org/10.1037%2Fh0042519). PMID 13602029 (https://pubmed.ncbi.nlm.nih.gov/13602029).
9. Werbos, P.J. (1975). *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences* (https://books.google.com/books?id=z81XmgEACAAJ).
10. Rosenblatt, Frank (1957). "The Perceptron—a perceiving and recognizing automaton". *Report 85-460-1*. Cornell Aeronautical Laboratory.
11. Olazaran, Mikel (1996). "A Sociological Study of the Official History of the Perceptrons Controversy". *Social Studies of Science*. **26** (3): 611–659. doi:10.1177/030631296026003005 (https://doi.org/10.1177%2F030631296026003005). JSTOR 285702 (https://www.jstor.org/stable/285702). S2CID 16786738 (https://api.semanticscholar.org/CorpusID:16786738).
12. Schmidhuber, J. (2015). "Deep Learning in Neural Networks: An Overview". *Neural Networks*. **61**: 85–117. arXiv:1404.7828 (https://arxiv.org/abs/1404.7828). doi:10.1016/j.neunet.2014.09.003 (https://doi.org/10.1016%2Fj.neunet.2014.09.003). PMID 25462637 (https://pubmed.ncbi.nlm.nih.gov/25462637). S2CID 11715509 (https://api.semanticscholar.org/CorpusID:11715509).
13. Ivakhnenko, A. G. (1973). *Cybernetic Predicting Devices* (https://books.google.com/books?id=FhwVNQAACAAJ). CCM Information Corporation.
14. Ivakhnenko, A. G.; Grigorʹevich Lapa, Valentin (1967). *Cybernetics and forecasting techniques* (https://books.google.com/books?id=rGFgAAAAMAAJ). American Elsevier Pub. Co.
15. Schmidhuber, Jürgen (2015). "Deep Learning" (https://doi.org/10.4249%2Fscholarpedia.32832). *Scholarpedia*. **10** (11): 85–117. Bibcode:2015SchpJ..1032832S (https://ui.adsabs.harvard.edu/abs/2015SchpJ..1032832S). doi:10.4249/scholarpedia.32832 (https://doi.org/10.4249%2Fscholarpedia.32832).

16. Dreyfus, Stuart E. (1 September 1990). "Artificial neural networks, back propagation, and the Kelley-Bryson gradient procedure". *Journal of Guidance, Control, and Dynamics*. **13** (5): 926–928. Bibcode:1990JGCD...13..926D (https://ui.adsabs.harvard.edu/abs/1990JGCD...13..926D). doi:10.2514/3.25422 (https://doi.org/10.2514%2F3.25422). ISSN 0731-5090 (https://www.worldcat.org/issn/0731-5090).
17. Mizutani, E.; Dreyfus, S.E.; Nishio, K. (2000). "On derivation of MLP backpropagation from the Kelley-Bryson optimal-control gradient formula and its application". *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*. IEEE: 167–172 vol.2. doi:10.1109/ijcnn.2000.857892 (https://doi.org/10.1109%2Fijcnn.2000.857892). ISBN 0-7695-0619-4. S2CID 351146 (https://api.semanticscholar.org/CorpusID:351146).
18. Kelley, Henry J. (1960). "Gradient theory of optimal flight paths". *ARS Journal*. **30** (10): 947–954. doi:10.2514/8.5282 (https://doi.org/10.2514%2F8.5282).
19. "A gradient method for optimizing multi-stage allocation processes". *Proceedings of the Harvard Univ. Symposium on digital computers and their applications*. April 1961.
20. Minsky, Marvin; Papert, Seymour (1969). *Perceptrons: An Introduction to Computational Geometry* (https://books.google.com/books?id=Ow1OAQAAIAAJ). MIT Press. ISBN 978-0-262-63022-1.
21. Linnainmaa, Seppo (1970). *The representation of the cumulative rounding error of an algorithm as a Taylor expansion of the local rounding errors* (Masters) (in Finnish). University of Helsinki. pp. 6–7.
22. Linnainmaa, Seppo (1976). "Taylor expansion of the accumulated rounding error". *BIT Numerical Mathematics*. **16** (2): 146–160. doi:10.1007/bf01931367 (https://doi.org/10.1007%2Fbf01931367). S2CID 122357351 (https://api.semanticscholar.org/CorpusID:122357351).
23. Dreyfus, Stuart (1973). "The computational solution of optimal control problems with time lag". *IEEE Transactions on Automatic Control*. **18** (4): 383–385. doi:10.1109/tac.1973.1100330 (https://doi.org/10.1109%2Ftac.1973.1100330).
24. Werbos, Paul (1982). "Applications of advances in nonlinear sensitivity analysis" (http://werbos.com/Neural/SensitivityIFIPSeptember1981.pdf) (PDF). *System modeling and optimization*. Springer. pp. 762–770.
25. Mead, Carver A.; Ismail, Mohammed (8 May 1989). *Analog VLSI Implementation of Neural Systems* (http://fennetic.net/irc/Christopher%20R.%20Carroll%20Carver%20Mead%20Mohammed%20Ismail%20Analog%20VLSI%20Implementation%20of%20Neural%20Systems.pdf) (PDF). The Kluwer International Series in Engineering and Computer Science. Vol. 80. Norwell, MA: Kluwer Academic Publishers. doi:10.1007/978-1-4613-1639-8 (https://doi.org/10.1007%2F978-1-4613-1639-8). ISBN 978-1-4613-1639-8.
26. David E. Rumelhart, Geoffrey E. Hinton & Ronald J. Williams , "Learning representations by back-propagating errors (https://www.nature.com/articles/323533a0)," *Nature', 323, pages 533–536 1986.*
27. Qian, Ning, and Terrence J. Sejnowski. "Predicting the secondary structure of globular proteins using neural network models." *Journal of molecular biology* 202, no. 4 (1988): 865-884.
28. Bohr, Henrik, Jakob Bohr, Søren Brunak, Rodney MJ Cotterill, Benny Lautrup, Leif Nørskov, Ole H. Olsen, and Steffen B. Petersen. "Protein secondary structure and homology by neural networks The α-helices in rhodopsin." *FEBS letters* 241, (1988): 223-228
29. Rost, Burkhard, and Chris Sander. "Prediction of protein secondary structure at better than 70% accuracy." *Journal of molecular biology* 232, no. 2 (1993): 584-599.

30. J. Weng, N. Ahuja and T. S. Huang, "Cresceptron: a self-organizing neural network which grows adaptively (http://www.cse.msu.edu/~weng/research/CresceptronIJCNN1992.pdf)," *Proc. International Joint Conference on Neural Networks*, Baltimore, Maryland, vol I, pp. 576–581, June 1992.
31. J. Weng, N. Ahuja and T. S. Huang, "Learning recognition and segmentation of 3-D objects from 2-D images (http://www.cse.msu.edu/~weng/research/CresceptronICCV1993.pdf)," *Proc. 4th International Conf. Computer Vision*, Berlin, Germany, pp. 121–128, May 1993.
32. J. Weng, N. Ahuja and T. S. Huang, "Learning recognition and segmentation using the Cresceptron (http://www.cse.msu.edu/~weng/research/CresceptronIJCV.pdf)," *International Journal of Computer Vision*, vol. 25, no. 2, pp. 105–139, Nov. 1997.
33. J. Schmidhuber., "Learning complex, extended sequences using the principle of history compression (https://mediatum.ub.tum.de/doc/814767/file.pdf)," *Neural Computation*, 4, pp. 234–242, 1992.
34. Domingos, Pedro (22 September 2015). "chapter 4". *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. Basic Books. ISBN 978-0465065707.
35. Smolensky, P. (1986). "Information processing in dynamical systems: Foundations of harmony theory." (http://portal.acm.org/citation.cfm?id=104290). In D. E. Rumelhart; J. L. McClelland; PDP Research Group (eds.). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (https://archive.org/details/paralleldistribu00rume/page/194). Vol. 1. pp. 194–281 (https://archive.org/details/paralleldistribu00rume/page/194). ISBN 978-0-262-68053-0.
36. Ng, Andrew; Dean, Jeff (2012). "Building High-level Features Using Large Scale Unsupervised Learning". arXiv:1112.6209 (https://arxiv.org/abs/1112.6209) [cs.LG (https://arxiv.org/archive/cs.LG)].
37. Ian Goodfellow and Yoshua Bengio and Aaron Courville (2016). *Deep Learning* (http://www.deeplearningbook.org). MIT Press.
38. Cireşan, Dan Claudiu; Meier, Ueli; Gambardella, Luca Maria; Schmidhuber, Jürgen (21 September 2010). "Deep, Big, Simple Neural Nets for Handwritten Digit Recognition". *Neural Computation*. **22** (12): 3207–3220. arXiv:1003.0358 (https://arxiv.org/abs/1003.0358). doi:10.1162/neco_a_00052 (https://doi.org/10.1162%2Fneco_a_00052). ISSN 0899-7667 (https://www.worldcat.org/issn/0899-7667). PMID 20858131 (https://pubmed.ncbi.nlm.nih.gov/20858131). S2CID 1918673 (https://api.semanticscholar.org/CorpusID:1918673).
39. Dominik Scherer, Andreas C. Müller, and Sven Behnke: "Evaluation of Pooling Operations in Convolutional Architectures for Object Recognition (https://www.ais.uni-bonn.de/papers/icann2010_maxpool.pdf)," *In 20th International Conference Artificial Neural Networks (ICANN)*, pp. 92–101, 2010. doi:10.1007/978-3-642-15825-4_10 (https://doi.org/10.1007%2F978-3-642-15825-4_10).
40. 2012 Kurzweil AI Interview (http://www.kurzweilai.net/how-bio-inspired-deep-learning-keeps-winning-competitions) Archived (https://web.archive.org/web/20180831075249/http://www.kurzweilai.net/how-bio-inspired-deep-learning-keeps-winning-competitions) 31 August 2018 at the Wayback Machine with Jürgen Schmidhuber on the eight competitions won by his Deep Learning team 2009–2012
41. "How bio-inspired deep learning keeps winning competitions | KurzweilAI" (https://web.archive.org/web/20180831075249/http://www.kurzweilai.net/how-bio-inspired-deep-learning-keeps-winning-competitions). *www.kurzweilai.net*. Archived from the original (http://www.kurzweilai.net/how-bio-inspired-deep-learning-keeps-winning-competitions) on 31 August 2018. Retrieved 16 June 2017.

42. Graves, Alex; Schmidhuber, Jürgen (2009). "Offline Handwriting Recognition with Multidimensional Recurrent Neural Networks" (http://www.idsia.ch/~juergen/nips2009.pdf) (PDF). In Koller, D.; Schuurmans, Dale; Bengio, Yoshua; Bottou, L. (eds.). *Advances in Neural Information Processing Systems 21 (NIPS 2008)* (https://papers.nips.cc/paper/2008). Neural Information Processing Systems (NIPS) Foundation. pp. 545–552. ISBN 9781605609492.
43. Graves, A.; Liwicki, M.; Fernandez, S.; Bertolami, R.; Bunke, H.; Schmidhuber, J. (May 2009). "A Novel Connectionist System for Unconstrained Handwriting Recognition" (http://www.idsia.ch/~juergen/tpami_2008.pdf) (PDF). *IEEE Transactions on Pattern Analysis and Machine Intelligence*. **31** (5): 855–868. CiteSeerX 10.1.1.139.4502 (https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.139.4502). doi:10.1109/tpami.2008.137 (https://doi.org/10.1109%2Ftpami.2008.137). ISSN 0162-8828 (https://www.worldcat.org/issn/0162-8828). PMID 19299860 (https://pubmed.ncbi.nlm.nih.gov/19299860). S2CID 14635907 (https://api.semanticscholar.org/CorpusID:14635907).
44. Ciresan, Dan; Meier, U.; Schmidhuber, J. (June 2012). *Multi-column deep neural networks for image classification*. *2012 IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3642–3649. arXiv:1202.2745 (https://arxiv.org/abs/1202.2745). Bibcode:2012arXiv1202.2745C (https://ui.adsabs.harvard.edu/abs/2012arXiv1202.2745C). CiteSeerX 10.1.1.300.3283 (https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.300.3283). doi:10.1109/cvpr.2012.6248110 (https://doi.org/10.1109%2Fcvpr.2012.6248110). ISBN 978-1-4673-1228-8. S2CID 2161592 (https://api.semanticscholar.org/CorpusID:2161592).
45. Zell, Andreas (2003). "chapter 5.2". *Simulation neuronaler Netze* [*Simulation of Neural Networks*] (in German) (1st ed.). Addison-Wesley. ISBN 978-3-89319-554-1. OCLC 249017987 (https://www.worldcat.org/oclc/249017987).
46. *Artificial intelligence* (3rd ed.). Addison-Wesley Pub. Co. 1992. ISBN 0-201-53377-4.
47. Abbod, Maysam F (2007). "Application of Artificial Intelligence to the Management of Urological Cancer". *The Journal of Urology*. **178** (4): 1150–1156. doi:10.1016/j.juro.2007.05.122 (https://doi.org/10.1016%2Fj.juro.2007.05.122). PMID 17698099 (https://pubmed.ncbi.nlm.nih.gov/17698099).
48. DAWSON, CHRISTIAN W (1998). "An artificial neural network approach to rainfall-runoff modelling". *Hydrological Sciences Journal*. **43** (1): 47–66. doi:10.1080/02626669809492102 (https://doi.org/10.1080%2F02626669809492102).
49. "The Machine Learning Dictionary" (https://web.archive.org/web/20180826151959/http://www.cse.unsw.edu.au/~billw/mldict.html#activnfn). *www.cse.unsw.edu.au*. Archived from the original (http://www.cse.unsw.edu.au/~billw/mldict.html#activnfn) on 26 August 2018. Retrieved 4 November 2009.
50. Ciresan, Dan; Ueli Meier; Jonathan Masci; Luca M. Gambardella; Jurgen Schmidhuber (2011). "Flexible, High Performance Convolutional Neural Networks for Image Classification" (https://people.idsia.ch/~juergen/ijcai2011.pdf) (PDF). *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence-Volume Volume Two*. **2**: 1237–1242. Archived (https://web.archive.org/web/20220405190128/https://people.idsia.ch/~juergen/ijcai2011.pdf) (PDF) from the original on 5 April 2022. Retrieved 7 July 2022.
51. Zell, Andreas (1994). *Simulation Neuronaler Netze* [*Simulation of Neural Networks*] (in German) (1st ed.). Addison-Wesley. p. 73. ISBN 3-89319-554-8.
52. Miljanovic, Milos (February–March 2012). "Comparative analysis of Recurrent and Finite Impulse Response Neural Networks in Time Series Prediction" (http://www.ijcse.com/docs/INDJCSE12-03-01-028.pdf) (PDF). *Indian Journal of Computer and Engineering*. **3** (1).

53. Lau, Suki (10 July 2017). "A Walkthrough of Convolutional Neural Network – Hyperparameter Tuning" (https://towardsdatascience.com/a-walkthrough-of-convolutional-neural-network-7f474f91d7bd). *Medium*. Retrieved 23 August 2019.
54. Kelleher, John D.; Mac Namee, Brian; D'Arcy, Aoife (2020). "7-8". *Fundamentals of machine learning for predictive data analytics: algorithms, worked examples, and case studies* (2nd ed.). Cambridge, MA. ISBN 978-0-262-36110-1. OCLC 1162184998 (https://www.worldcat.org/oclc/1162184998).
55. Wei, Jiakai (26 April 2019). "Forget the Learning Rate, Decay Loss". arXiv:1905.00094 (https://arxiv.org/abs/1905.00094) [cs.LG (https://arxiv.org/archive/cs.LG)].
56. Li, Y.; Fu, Y.; Li, H.; Zhang, S. W. (1 June 2009). *The Improved Training Algorithm of Back Propagation Neural Network with Self-adaptive Learning Rate*. *2009 International Conference on Computational Intelligence and Natural Computing*. Vol. 1. pp. 73–76. doi:10.1109/CINC.2009.111 (https://doi.org/10.1109%2FCINC.2009.111). ISBN 978-0-7695-3645-3. S2CID 10557754 (https://api.semanticscholar.org/CorpusID:10557754).
57. Huang, Guang-Bin; Zhu, Qin-Yu; Siew, Chee-Kheong (2006). "Extreme learning machine: theory and applications". *Neurocomputing*. **70** (1): 489–501. CiteSeerX 10.1.1.217.3692 (https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.217.3692). doi:10.1016/j.neucom.2005.12.126 (https://doi.org/10.1016%2Fj.neucom.2005.12.126).
58. Widrow, Bernard; et al. (2013). "The no-prop algorithm: A new learning algorithm for multilayer neural networks". *Neural Networks*. **37**: 182–188. doi:10.1016/j.neunet.2012.09.020 (https://doi.org/10.1016%2Fj.neunet.2012.09.020). PMID 23140797 (https://pubmed.ncbi.nlm.nih.gov/23140797).
59. Ollivier, Yann; Charpiat, Guillaume (2015). "Training recurrent networks without backtracking". arXiv:1507.07680 (https://arxiv.org/abs/1507.07680) [cs.NE (https://arxiv.org/archive/cs.NE)].
60. Hinton, G. E. (2010). "A Practical Guide to Training Restricted Boltzmann Machines" (https://www.researchgate.net/publication/221166159). *Tech. Rep. UTML TR 2010-003*.
61. ESANN. 2009.
62. Bernard, Etienne (2021). *Introduction to Machine Learning* (https://www.wolfram.com/language/introduction-machine-learning/). Wolfram Media Inc. p. 9. ISBN 978-1-579550-48-6.
63. Ojha, Varun Kumar; Abraham, Ajith; Snášel, Václav (1 April 2017). "Metaheuristic design of feedforward neural networks: A review of two decades of research". *Engineering Applications of Artificial Intelligence*. **60**: 97–116. arXiv:1705.05584 (https://arxiv.org/abs/1705.05584). Bibcode:2017arXiv170505584O (https://ui.adsabs.harvard.edu/abs/2017arXiv170505584O). doi:10.1016/j.engappai.2017.01.013 (https://doi.org/10.1016%2Fj.engappai.2017.01.013). S2CID 27910748 (https://api.semanticscholar.org/CorpusID:27910748).
64. Dominic, S.; Das, R.; Whitley, D.; Anderson, C. (July 1991). "Genetic reinforcement learning for neural networks" (https://archive.org/details/ijcnn91seattlein01ieee). *IJCNN-91-Seattle International Joint Conference on Neural Networks*. IJCNN-91-Seattle International Joint Conference on Neural Networks. Seattle, Washington, USA: IEEE. pp. 71–76. doi:10.1109/IJCNN.1991.155315 (https://doi.org/10.1109%2FIJCNN.1991.155315). ISBN 0-7803-0164-1.
65. Hoskins, J.C.; Himmelblau, D.M. (1992). "Process control via artificial neural networks and reinforcement learning". *Computers & Chemical Engineering*. **16** (4): 241–251. doi:10.1016/0098-1354(92)80045-B (https://doi.org/10.1016%2F0098-1354%2892%2980045-B).
66. Bertsekas, D.P.; Tsitsiklis, J.N. (1996). *Neuro-dynamic programming* (https://papers.nips.cc/paper/4741-deep-neural-networks-segment-neuronal-membranes-in-electron-microscopy-images). Athena Scientific. p. 512. ISBN 978-1-886529-10-6.

67. Secomandi, Nicola (2000). "Comparing neuro-dynamic programming algorithms for the vehicle routing problem with stochastic demands". *Computers & Operations Research*. **27** (11–12): 1201–1225. CiteSeerX 10.1.1.392.4034 (https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.392.4034). doi:10.1016/S0305-0548(99)00146-X (https://doi.org/10.1016%2FS0305-0548%2899%2900146-X).
68. de Rigo, D.; Rizzoli, A. E.; Soncini-Sessa, R.; Weber, E.; Zenesi, P. (2001). "Neuro-dynamic programming for the efficient management of reservoir networks". *Proceedings of MODSIM 2001, International Congress on Modelling and Simulation*. MODSIM 2001, International Congress on Modelling and Simulation (http://www.mssanz.org.au/MODSIM01/MODSIM01.htm). Canberra, Australia: Modelling and Simulation Society of Australia and New Zealand. doi:10.5281/zenodo.7481 (https://doi.org/10.5281%2Fzenodo.7481). ISBN 0-86740-525-2.
69. Damas, M.; Salmeron, M.; Diaz, A.; Ortega, J.; Prieto, A.; Olivares, G. (2000). "Genetic algorithms and neuro-dynamic programming: application to water supply networks". *Proceedings of 2000 Congress on Evolutionary Computation*. 2000 Congress on Evolutionary Computation. Vol. 1. La Jolla, California, USA: IEEE. pp. 7–14. doi:10.1109/CEC.2000.870269 (https://doi.org/10.1109%2FCEC.2000.870269). ISBN 0-7803-6375-2.
70. Deng, Geng; Ferris, M.C. (2008). *Neuro-dynamic programming for fractionated radiotherapy planning*. Springer Optimization and Its Applications. Vol. 12. pp. 47–70. CiteSeerX 10.1.1.137.8288 (https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.137.8288). doi:10.1007/978-0-387-73299-2_3 (https://doi.org/10.1007%2F978-0-387-73299-2_3). ISBN 978-0-387-73298-5.
71. Bozinovski, S. (1982). "A self-learning system using secondary reinforcement". In R. Trappl (ed.) Cybernetics and Systems Research: Proceedings of the Sixth European Meeting on Cybernetics and Systems Research. North Holland. pp. 397–402. ISBN 978-0-444-86488-8.
72. Bozinovski, S. (2014) "Modeling mechanisms of cognition-emotion interaction in artificial neural networks, since 1981 (https://core.ac.uk/download/pdf/81973924.pdf)." Procedia Computer Science p. 255-263
73. Bozinovski, Stevo; Bozinovska, Liljana (2001). "Self-learning agents: A connectionist theory of emotion based on crossbar value judgment". *Cybernetics and Systems*. **32** (6): 637–667. doi:10.1080/01969720118145 (https://doi.org/10.1080%2F01969720118145). S2CID 8944741 (https://api.semanticscholar.org/CorpusID:8944741).
74. "Artificial intelligence can 'evolve' to solve problems" (https://www.science.org/content/article/artificial-intelligence-can-evolve-solve-problems). *Science | AAAS*. 10 January 2018. Retrieved 7 February 2018.
75. Turchetti, Claudio (2004), *Stochastic Models of Neural Networks*, Frontiers in artificial intelligence and applications: Knowledge-based intelligent engineering systems, vol. 102, IOS Press, ISBN 9781586033880
76. de Rigo, D.; Castelletti, A.; Rizzoli, A. E.; Soncini-Sessa, R.; Weber, E. (January 2005). "A selective improvement technique for fastening Neuro-Dynamic Programming in Water Resources Network Management" (http://www.nt.ntnu.no/users/skoge/prost/proceedings/ifac2005/Papers/Paper4269.html). In Pavel Zítek (ed.). *Proceedings of the 16th IFAC World Congress – IFAC-PapersOnLine*. 16th IFAC World Congress (http://www.nt.ntnu.no/users/skoge/prost/proceedings/ifac2005/Index.html). Vol. 16. Prague, Czech Republic: IFAC. pp. 7–12. doi:10.3182/20050703-6-CZ-1902.02172 (https://doi.org/10.3182%2F20050703-6-CZ-1902.02172). hdl:11311/255236 (https://hdl.handle.net/11311%2F255236). ISBN 978-3-902661-75-3. Retrieved 30 December 2011.
77. Ferreira, C. (2006). "Designing Neural Networks Using Gene Expression Programming". In A. Abraham; B. de Baets; M. Köppen; B. Nickolay (eds.). *Applied Soft Computing Technologies: The Challenge of Complexity* (http://www.gene-expression-programming.com/webpapers/Ferreira-ASCT2006.pdf) (PDF). Springer-Verlag. pp. 517–536.

78. Da, Y.; Xiurun, G. (July 2005). "An improved PSO-based ANN with simulated annealing technique" (https://web.archive.org/web/20120425233611/http://www.dice.ucl.ac.be/esann/proceedings/electronicproceedings.htm). In T. Villmann (ed.). *New Aspects in Neurocomputing: 11th European Symposium on Artificial Neural Networks*. Vol. 63. Elsevier. pp. 527–533. doi:10.1016/j.neucom.2004.07.002 (https://doi.org/10.1016%2Fj.neucom.2004.07.002). Archived from the original (http://www.dice.ucl.ac.be/esann/proceedings/electronicproceedings.htm) on 25 April 2012. Retrieved 30 December 2011.
79. Wu, J.; Chen, E. (May 2009). "A Novel Nonparametric Regression Ensemble for Rainfall Forecasting Using Particle Swarm Optimization Technique Coupled with Artificial Neural Network" (https://web.archive.org/web/20141231221755/http://www2.mae.cuhk.edu.hk/~isnn2009/). In Wang, H.; Shen, Y.; Huang, T.; Zeng, Z. (eds.). *6th International Symposium on Neural Networks, ISNN 2009*. Lecture Notes in Computer Science. Vol. 5553. Springer. pp. 49–58. doi:10.1007/978-3-642-01513-7_6 (https://doi.org/10.1007%2F978-3-642-01513-7_6). ISBN 978-3-642-01215-0. Archived from the original (http://www2.mae.cuhk.edu.hk/~isnn2009/) on 31 December 2014. Retrieved 1 January 2012.
80. Ting Qin; Zonghai Chen; Haitao Zhang; Sifu Li; Wei Xiang; Ming Li (2004). "A learning algorithm of CMAC based on RLS" (http://www-control.eng.cam.ac.uk/Homepage/papers/cued_control_998.pdf) (PDF). *Neural Processing Letters*. **19** (1): 49–61. doi:10.1023/B:NEPL.0000016847.18175.60 (https://doi.org/10.1023%2FB%3ANEPL.0000016847.18175.60). S2CID 6233899 (https://api.semanticscholar.org/CorpusID:6233899).
81. Ting Qin; Haitao Zhang; Zonghai Chen; Wei Xiang (2005). "Continuous CMAC-QRLS and its systolic array" (http://www-control.eng.cam.ac.uk/Homepage/papers/cued_control_997.pdf) (PDF). *Neural Processing Letters*. **22** (1): 1–16. doi:10.1007/s11063-004-2694-0 (https://doi.org/10.1007%2Fs11063-004-2694-0). S2CID 16095286 (https://api.semanticscholar.org/CorpusID:16095286).
82. LeCun Y, Boser B, Denker JS, Henderson D, Howard RE, Hubbard W, Jackel LD (1989). "Backpropagation Applied to Handwritten Zip Code Recognition". *Neural Computation*. **1** (4): 541–551. doi:10.1162/neco.1989.1.4.541 (https://doi.org/10.1162%2Fneco.1989.1.4.541). S2CID 41312633 (https://api.semanticscholar.org/CorpusID:41312633).
83. Yann LeCun (2016). Slides on Deep Learning Online (https://indico.cern.ch/event/510372/)
84. Hochreiter, Sepp; Schmidhuber, Jürgen (1 November 1997). "Long Short-Term Memory". *Neural Computation*. **9** (8): 1735–1780. doi:10.1162/neco.1997.9.8.1735 (https://doi.org/10.1162%2Fneco.1997.9.8.1735). ISSN 0899-7667 (https://www.worldcat.org/issn/0899-7667). PMID 9377276 (https://pubmed.ncbi.nlm.nih.gov/9377276). S2CID 1915014 (https://api.semanticscholar.org/CorpusID:1915014).
85. Sak, Hasim; Senior, Andrew; Beaufays, Francoise (2014). "Long Short-Term Memory recurrent neural network architectures for large scale acoustic modeling" (https://web.archive.org/web/20180424203806/https://static.googleusercontent.com/media/research.google.com/en//pubs/archive/43905.pdf) (PDF). Archived from the original (https://static.googleusercontent.com/media/research.google.com/en//pubs/archive/43905.pdf) (PDF) on 24 April 2018.
86. Li, Xiangang; Wu, Xihong (15 October 2014). "Constructing Long Short-Term Memory based Deep Recurrent Neural Networks for Large Vocabulary Speech Recognition". arXiv:1410.4281 (https://arxiv.org/abs/1410.4281) [cs.CL (https://arxiv.org/archive/cs.CL)].
87. Fan, Y.; Qian, Y.; Xie, F.; Soong, F. K. (2014). "TTS synthesis with bidirectional LSTM based Recurrent Neural Networks" (https://www.researchgate.net/publication/287741874). *Proceedings of the Annual Conference of the International Speech Communication Association, Interspeech*: 1964–1968. Retrieved 13 June 2017.
88. Zen, Heiga; Sak, Hasim (2015). "Unidirectional Long Short-Term Memory Recurrent Neural Network with Recurrent Output Layer for Low-Latency Speech Synthesis" (https://static.googleusercontent.com/media/research.google.com/en//pubs/archive/43266.pdf) (PDF). *Google.com*. ICASSP. pp. 4470–4474.

89. Fan, Bo; Wang, Lijuan; Soong, Frank K.; Xie, Lei (2015). "Photo-Real Talking Head with Deep Bidirectional LSTM" (https://www.microsoft.com/en-us/research/wp-content/uploads/2015/04/icassp2015_fanbo_1009.pdf) (PDF). *Proceedings of ICASSP*.
90. Silver, David; Hubert, Thomas; Schrittwieser, Julian; Antonoglou, Ioannis; Lai, Matthew; Guez, Arthur; Lanctot, Marc; Sifre, Laurent; Kumaran, Dharshan; Graepel, Thore; Lillicrap, Timothy; Simonyan, Karen; Hassabis, Demis (5 December 2017). "Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm". arXiv:1712.01815 (https://arxiv.org/abs/1712.01815) [cs.AI (https://arxiv.org/archive/cs.AI)].
91. Goodfellow, Ian; Pouget-Abadie, Jean; Mirza, Mehdi; Xu, Bing; Warde-Farley, David; Ozair, Sherjil; Courville, Aaron; Bengio, Yoshua (2014). *Generative Adversarial Networks* (https://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf) (PDF). Proceedings of the International Conference on Neural Information Processing Systems (NIPS 2014). pp. 2672–2680.
92. Zoph, Barret; Le, Quoc V. (4 November 2016). "Neural Architecture Search with Reinforcement Learning". arXiv:1611.01578 (https://arxiv.org/abs/1611.01578) [cs.LG (https://arxiv.org/archive/cs.LG)].
93. Haifeng Jin; Qingquan Song; Xia Hu (2019). "Auto-keras: An efficient neural architecture search system" (https://autokeras.com/). *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM. arXiv:1806.10282 (https://arxiv.org/abs/1806.10282). Retrieved 21 August 2019 – via autokeras.com.
94. Claesen, Marc; De Moor, Bart (2015). "Hyperparameter Search in Machine Learning". arXiv:1502.02127 (https://arxiv.org/abs/1502.02127) [cs.LG (https://arxiv.org/archive/cs.LG)]. Bibcode:2015arXiv150202127C (https://ui.adsabs.harvard.edu/abs/2015arXiv150202127C/abstract)
95. Turek, Fred D. (March 2007). "Introduction to Neural Net Machine Vision" (http://www.vision-systems.com/articles/print/volume-12/issue-3/features/introduction-to-neural-net-machine-vision.html). *Vision Systems Design*. **12** (3). Retrieved 5 March 2013.
96. Zissis, Dimitrios (October 2015). "A cloud based architecture capable of perceiving and predicting multiple vessel behaviour" (https://zenodo.org/record/848743). *Applied Soft Computing*. **35**: 652–661. doi:10.1016/j.asoc.2015.07.002 (https://doi.org/10.1016%2Fj.asoc.2015.07.002).
97. Roman M. Balabin; Ekaterina I. Lomakina (2009). "Neural network approach to quantum-chemistry data: Accurate prediction of density functional theory energies". *J. Chem. Phys.* **131** (7): 074104. Bibcode:2009JChPh.131g4104B (https://ui.adsabs.harvard.edu/abs/2009JChPh.131g4104B). doi:10.1063/1.3206326 (https://doi.org/10.1063%2F1.3206326). PMID 19708729 (https://pubmed.ncbi.nlm.nih.gov/19708729).
98. Silver, David; et al. (2016). "Mastering the game of Go with deep neural networks and tree search" (http://web.iitd.ac.in/~sumeet/Silver16.pdf) (PDF). *Nature*. **529** (7587): 484–489. Bibcode:2016Natur.529..484S (https://ui.adsabs.harvard.edu/abs/2016Natur.529..484S). doi:10.1038/nature16961 (https://doi.org/10.1038%2Fnature16961). PMID 26819042 (https://pubmed.ncbi.nlm.nih.gov/26819042). S2CID 515925 (https://api.semanticscholar.org/CorpusID:515925).
99. Sengupta, Nandini; Sahidullah, Md; Saha, Goutam (August 2016). "Lung sound classification using cepstral-based statistical features". *Computers in Biology and Medicine*. **75** (1): 118–129. doi:10.1016/j.compbiomed.2016.05.013 (https://doi.org/10.1016%2Fj.compbiomed.2016.05.013). PMID 27286184 (https://pubmed.ncbi.nlm.nih.gov/27286184).
100. Choy, Christopher B., et al. "3d-r2n2: A unified approach for single and multi-view 3d object reconstruction (https://arxiv.org/abs/1604.00449)." European conference on computer vision. Springer, Cham, 2016.

101. Gessler, Josef (August 2021). "Sensor for food analysis applying impedance spectroscopy and artificial neural networks" (https://riunet.upv.es/handle/10251/174498). *RiuNet UPV* (1): 8–12.
102. Maitra, D. S.; Bhattacharya, U.; Parui, S. K. (August 2015). "CNN based common approach to handwritten character recognition of multiple scripts" (https://ieeexplore.ieee.org/document/7333916). *2015 13th International Conference on Document Analysis and Recognition (ICDAR)*: 1021–1025. doi:10.1109/ICDAR.2015.7333916 (https://doi.org/10.1109%2FICDAR.2015.7333916). ISBN 978-1-4799-1805-8. S2CID 25739012 (https://api.semanticscholar.org/CorpusID:25739012).
103. French, Jordan (2016). "The time traveller's CAPM". *Investment Analysts Journal*. **46** (2): 81–96. doi:10.1080/10293523.2016.1255469 (https://doi.org/10.1080%2F10293523.2016.1255469). S2CID 157962452 (https://api.semanticscholar.org/CorpusID:157962452).
104. Schechner, Sam (15 June 2017). "Facebook Boosts A.I. to Block Terrorist Propaganda" (https://www.wsj.com/articles/facebook-boosts-a-i-to-block-terrorist-propaganda-1497546000). *Wall Street Journal*. ISSN 0099-9660 (https://www.worldcat.org/issn/0099-9660). Retrieved 16 June 2017.
105. Ganesan, N (2010). "Application of Neural Networks in Diagnosing Cancer Disease Using Demographic Data" (https://doi.org/10.5120%2F476-783). *International Journal of Computer Applications*. **1** (26): 81–97. Bibcode:2010IJCA....1z..81G (https://ui.adsabs.harvard.edu/abs/2010IJCA....1z..81G). doi:10.5120/476-783 (https://doi.org/10.5120%2F476-783).
106. Bottaci, Leonardo (1997). "Artificial Neural Networks Applied to Outcome Prediction for Colorectal Cancer Patients in Separate Institutions" (https://web.archive.org/web/20181123170444/http://www.lcc.uma.es/~jja/recidiva/042.pdf) (PDF). *Lancet*. The Lancet. **350** (9076): 469–72. doi:10.1016/S0140-6736(96)11196-X (https://doi.org/10.1016%2FS0140-6736%2896%2911196-X). PMID 9274582 (https://pubmed.ncbi.nlm.nih.gov/9274582). S2CID 18182063 (https://api.semanticscholar.org/CorpusID:18182063). Archived from the original (http://www.lcc.uma.es/~jja/recidiva/042.pdf) (PDF) on 23 November 2018. Retrieved 2 May 2012.
107. Alizadeh, Elaheh; Lyons, Samanthe M; Castle, Jordan M; Prasad, Ashok (2016). "Measuring systematic changes in invasive cancer cell shape using Zernike moments" (http://pubs.rsc.org/en/Content/ArticleLanding/2016/IB/C6IB00100A). *Integrative Biology*. **8** (11): 1183–1193. doi:10.1039/C6IB00100A (https://doi.org/10.1039%2FC6IB00100A). PMID 27735002 (https://pubmed.ncbi.nlm.nih.gov/27735002).
108. Lyons, Samanthe (2016). "Changes in cell shape are correlated with metastatic potential in murine" (https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4810736). *Biology Open*. **5** (3): 289–299. doi:10.1242/bio.013409 (https://doi.org/10.1242%2Fbio.013409). PMC 4810736 (https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4810736). PMID 26873952 (https://pubmed.ncbi.nlm.nih.gov/26873952).
109. Nabian, Mohammad Amin; Meidani, Hadi (28 August 2017). "Deep Learning for Accelerated Reliability Analysis of Infrastructure Networks". *Computer-Aided Civil and Infrastructure Engineering*. **33** (6): 443–458. arXiv:1708.08551 (https://arxiv.org/abs/1708.08551). Bibcode:2017arXiv170808551N (https://ui.adsabs.harvard.edu/abs/2017arXiv170808551N). doi:10.1111/mice.12359 (https://doi.org/10.1111%2Fmice.12359). S2CID 36661983 (https://api.semanticscholar.org/CorpusID:36661983).
110. Nabian, Mohammad Amin; Meidani, Hadi (2018). "Accelerating Stochastic Assessment of Post-Earthquake Transportation Network Connectivity via Machine-Learning-Based Surrogates" (https://trid.trb.org/view/1496617). *Transportation Research Board 97th Annual Meeting*.

111. Díaz, E.; Brotons, V.; Tomás, R. (September 2018). "Use of artificial neural networks to predict 3-D elastic settlement of foundations on soils with inclined bedrock" (https://doi.org/10.1016%2Fj.sandf.2018.08.001). *Soils and Foundations*. **58** (6): 1414–1422. doi:10.1016/j.sandf.2018.08.001 (https://doi.org/10.1016%2Fj.sandf.2018.08.001). hdl:10045/81208 (https://hdl.handle.net/10045%2F81208). ISSN 0038-0806 (https://www.worldcat.org/issn/0038-0806).
112. Govindaraju, Rao S. (1 April 2000). "Artificial Neural Networks in Hydrology. I: Preliminary Concepts". *Journal of Hydrologic Engineering*. **5** (2): 115–123. doi:10.1061/(ASCE)1084-0699(2000)5:2(115) (https://doi.org/10.1061%2F%28ASCE%291084-0699%282000%295%3A2%28115%29).
113. Govindaraju, Rao S. (1 April 2000). "Artificial Neural Networks in Hydrology. II: Hydrologic Applications". *Journal of Hydrologic Engineering*. **5** (2): 124–137. doi:10.1061/(ASCE)1084-0699(2000)5:2(124) (https://doi.org/10.1061%2F%28ASCE%291084-0699%282000%295%3A2%28124%29).
114. Peres, D. J.; Iuppa, C.; Cavallaro, L.; Cancelliere, A.; Foti, E. (1 October 2015). "Significant wave height record extension by neural networks and reanalysis wind data". *Ocean Modelling*. **94**: 128–140. Bibcode:2015OcMod..94..128P (https://ui.adsabs.harvard.edu/abs/2015OcMod..94..128P). doi:10.1016/j.ocemod.2015.08.002 (https://doi.org/10.1016%2Fj.ocemod.2015.08.002).
115. Dwarakish, G. S.; Rakshith, Shetty; Natesan, Usha (2013). "Review on Applications of Neural Network in Coastal Engineering" (http://www.ciitresearch.org/dl/index.php/aiml/article/view/AIML072013007). *Artificial Intelligent Systems and Machine Learning*. **5** (7): 324–331.
116. Ermini, Leonardo; Catani, Filippo; Casagli, Nicola (1 March 2005). "Artificial Neural Networks applied to landslide susceptibility assessment". *Geomorphology*. Geomorphological hazard and human impact in mountain environments. **66** (1): 327–343. Bibcode:2005Geomo..66..327E (https://ui.adsabs.harvard.edu/abs/2005Geomo..66..327E). doi:10.1016/j.geomorph.2004.09.025 (https://doi.org/10.1016%2Fj.geomorph.2004.09.025).
117. Nix, R.; Zhang, J. (May 2017). "Classification of Android apps and malware using deep neural networks". *2017 International Joint Conference on Neural Networks (IJCNN)*: 1871–1878. doi:10.1109/IJCNN.2017.7966078 (https://doi.org/10.1109%2FIJCNN.2017.7966078). ISBN 978-1-5090-6182-2. S2CID 8838479 (https://api.semanticscholar.org/CorpusID:8838479).
118. "Detecting Malicious URLs" (https://web.archive.org/web/20190714201955/http://www.sysnet.ucsd.edu/projects/url/). *The systems and networking group at UCSD*. Archived from the original (http://www.sysnet.ucsd.edu/projects/url/) on 14 July 2019. Retrieved 15 February 2019.
119. Homayoun, Sajad; Ahmadzadeh, Marzieh; Hashemi, Sattar; Dehghantanha, Ali; Khayami, Raouf (2018), Dehghantanha, Ali; Conti, Mauro; Dargahi, Tooska (eds.), "BoTShark: A Deep Learning Approach for Botnet Traffic Detection", *Cyber Threat Intelligence*, Advances in Information Security, Springer International Publishing, pp. 137–153, doi:10.1007/978-3-319-73951-9_7 (https://doi.org/10.1007%2F978-3-319-73951-9_7), ISBN 978-3-319-73951-9
120. Ghosh and Reilly (January 1994). "Credit card fraud detection with a neural-network". *1994 Proceedings of the Twenty-Seventh Hawaii International Conference on System Sciences*. **3**: 621–630. doi:10.1109/HICSS.1994.323314 (https://doi.org/10.1109%2FHICSS.1994.323314). ISBN 978-0-8186-5090-1. S2CID 13260377 (https://api.semanticscholar.org/CorpusID:13260377).
121. Ananthaswamy, Anil (19 April 2021). "Latest Neural Nets Solve World's Hardest Equations Faster Than Ever Before" (https://www.quantamagazine.org/new-neural-networks-solve-hardest-equations-faster-than-ever-20210419/). *Quanta Magazine*. Retrieved 12 May 2021.

122. "AI has cracked a key mathematical puzzle for understanding our world" (https://www.technologyreview.com/2020/10/30/1011435/ai-fourier-neural-network-cracks-navier-stokes-and-partial-differential-equations/). *MIT Technology Review*. Retrieved 19 November 2020.
123. "Caltech Open-Sources AI for Solving Partial Differential Equations" (https://www.infoq.com/news/2020/12/caltech-ai-pde/). *InfoQ*. Retrieved 20 January 2021.
124. Nagy, Alexandra (28 June 2019). "Variational Quantum Monte Carlo Method with a Neural-Network Ansatz for Open Quantum Systems". *Physical Review Letters*. **122** (25): 250501. arXiv:1902.09483 (https://arxiv.org/abs/1902.09483). Bibcode:2019PhRvL.122y0501N (https://ui.adsabs.harvard.edu/abs/2019PhRvL.122y0501N). doi:10.1103/PhysRevLett.122.250501 (https://doi.org/10.1103%2FPhysRevLett.122.250501). PMID 31347886 (https://pubmed.ncbi.nlm.nih.gov/31347886). S2CID 119074378 (https://api.semanticscholar.org/CorpusID:119074378).
125. Yoshioka, Nobuyuki; Hamazaki, Ryusuke (28 June 2019). "Constructing neural stationary states for open quantum many-body systems". *Physical Review B*. **99** (21): 214306. arXiv:1902.07006 (https://arxiv.org/abs/1902.07006). Bibcode:2019arXiv190207006Y (https://ui.adsabs.harvard.edu/abs/2019arXiv190207006Y). doi:10.1103/PhysRevB.99.214306 (https://doi.org/10.1103%2FPhysRevB.99.214306). S2CID 119470636 (https://api.semanticscholar.org/CorpusID:119470636).
126. Hartmann, Michael J.; Carleo, Giuseppe (28 June 2019). "Neural-Network Approach to Dissipative Quantum Many-Body Dynamics". *Physical Review Letters*. **122** (25): 250502. arXiv:1902.05131 (https://arxiv.org/abs/1902.05131). Bibcode:2019arXiv190205131H (https://ui.adsabs.harvard.edu/abs/2019arXiv190205131H). doi:10.1103/PhysRevLett.122.250502 (https://doi.org/10.1103%2FPhysRevLett.122.250502). PMID 31347862 (https://pubmed.ncbi.nlm.nih.gov/31347862). S2CID 119357494 (https://api.semanticscholar.org/CorpusID:119357494).
127. Vicentini, Filippo; Biella, Alberto; Regnault, Nicolas; Ciuti, Cristiano (28 June 2019). "Variational Neural-Network Ansatz for Steady States in Open Quantum Systems". *Physical Review Letters*. **122** (25): 250503. arXiv:1902.10104 (https://arxiv.org/abs/1902.10104). Bibcode:2019arXiv190210104V (https://ui.adsabs.harvard.edu/abs/2019arXiv190210104V). doi:10.1103/PhysRevLett.122.250503 (https://doi.org/10.1103%2FPhysRevLett.122.250503). PMID 31347877 (https://pubmed.ncbi.nlm.nih.gov/31347877). S2CID 119504484 (https://api.semanticscholar.org/CorpusID:119504484).
128. Forrest MD (April 2015). "Simulation of alcohol action upon a detailed Purkinje neuron model and a simpler surrogate model that runs >400 times faster" (https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4417229). *BMC Neuroscience*. **16** (27): 27. doi:10.1186/s12868-015-0162-6 (https://doi.org/10.1186%2Fs12868-015-0162-6). PMC 4417229 (https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4417229). PMID 25928094 (https://pubmed.ncbi.nlm.nih.gov/25928094).
129. Siegelmann, H.T.; Sontag, E.D. (1991). "Turing computability with neural nets" (http://www.math.rutgers.edu/~sontag/FTPDIR/aml-turing.pdf) (PDF). *Appl. Math. Lett*. **4** (6): 77–80. doi:10.1016/0893-9659(91)90080-F (https://doi.org/10.1016%2F0893-9659%2891%2990080-F).
130. Balcázar, José (July 1997). "Computational Power of Neural Networks: A Kolmogorov Complexity Characterization". *IEEE Transactions on Information Theory*. **43** (4): 1175–1183. CiteSeerX 10.1.1.411.7782 (https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.411.7782). doi:10.1109/18.605580 (https://doi.org/10.1109%2F18.605580).
131. MacKay, David, J.C. (2003). *Information Theory, Inference, and Learning Algorithms* (http://www.inference.phy.cam.ac.uk/itprnn/book.pdf) (PDF). Cambridge University Press. ISBN 978-0-521-64298-9.

132. Cover, Thomas (1965). "Geometrical and Statistical Properties of Systems of Linear Inequalities with Applications in Pattern Recognition" (http://www-isl.stanford.edu/people/cover/papers/paper2.pdf) (PDF). *IEEE Transactions on Electronic Computers*. IEEE. EC-14 (3): 326–334. doi:10.1109/PGEC.1965.264137 (https://doi.org/10.1109%2FPGEC.1965.264137).
133. Gerald, Friedland (2019). "Reproducibility and Experimental Design for Machine Learning on Audio and Multimedia Data". *MM '19: Proceedings of the 27th ACM International Conference on Multimedia*. ACM: 2709–2710. doi:10.1145/3343031.3350545 (https://doi.org/10.1145%2F3343031.3350545). ISBN 978-1-4503-6889-6. S2CID 204837170 (https://api.semanticscholar.org/CorpusID:204837170).
134. "Stop tinkering, start measuring! Predictable experimental design of Neural Network experiments" (http://tfmeter.icsi.berkeley.edu). *The Tensorflow Meter*.
135. Lee, Jaehoon; Xiao, Lechao; Schoenholz, Samuel S.; Bahri, Yasaman; Novak, Roman; Sohl-Dickstein, Jascha; Pennington, Jeffrey (2020). "Wide neural networks of any depth evolve as linear models under gradient descent". *Journal of Statistical Mechanics: Theory and Experiment*. **2020** (12): 124002. arXiv:1902.06720 (https://arxiv.org/abs/1902.06720). Bibcode:2020JSMTE2020l4002L (https://ui.adsabs.harvard.edu/abs/2020JSMTE2020l4002L). doi:10.1088/1742-5468/abc62b (https://doi.org/10.1088%2F1742-5468%2Fabc62b). S2CID 62841516 (https://api.semanticscholar.org/CorpusID:62841516).
136. Arthur Jacot; Franck Gabriel; Clement Hongler (2018). *Neural Tangent Kernel: Convergence and Generalization in Neural Networks* (https://proceedings.neurips.cc/paper/2018/file/5a4be1fa34e62bb8a6ec6b91d2462f5a-Paper.pdf) (PDF). 32nd Conference on Neural Information Processing Systems (NeurIPS 2018), Montreal, Canada.
137. Xu ZJ, Zhang Y, Xiao Y (2019). "Training Behavior of Deep Neural Network in Frequency Domain" (http://arxiv.org/abs/1807.01251). In Gedeon T, Wong K, Lee M (eds.). *Neural Information Processing. ICONIP 2019*. Lecture Notes in Computer Science. Vol. 11953. Springer, Cham. pp. 264–274. arXiv:1807.01251 (https://arxiv.org/abs/1807.01251). doi:10.1007/978-3-030-36708-4_22 (https://doi.org/10.1007%2F978-3-030-36708-4_22). ISBN 978-3-030-36707-7. S2CID 49562099 (https://api.semanticscholar.org/CorpusID:49562099).
138. Nasim Rahaman; Aristide Baratin; Devansh Arpit; Felix Draxler; Min Lin; Fred Hamprecht; Yoshua Bengio; Aaron Courville (2019). "On the Spectral Bias of Neural Networks" (http://proceedings.mlr.press/v97/rahaman19a/rahaman19a.pdf) (PDF). *Proceedings of the 36th International Conference on Machine Learning*. **97**: 5301–5310. arXiv:1806.08734 (https://arxiv.org/abs/1806.08734).
139. Zhi-Qin John Xu; Yaoyu Zhang; Tao Luo; Yanyang Xiao; Zheng Ma (2020). "Frequency Principle: Fourier Analysis Sheds Light on Deep Neural Networks". *Communications in Computational Physics*. **28** (5): 1746–1767. arXiv:1901.06523 (https://arxiv.org/abs/1901.06523). Bibcode:2020CCoPh..28.1746X (https://ui.adsabs.harvard.edu/abs/2020CCoPh..28.1746X). doi:10.4208/cicp.OA-2020-0085 (https://doi.org/10.4208%2Fcicp.OA-2020-0085). S2CID 58981616 (https://api.semanticscholar.org/CorpusID:58981616).
140. Tao Luo; Zheng Ma; Zhi-Qin John Xu; Yaoyu Zhang (2019). "Theory of the Frequency Principle for General Deep Neural Networks". arXiv:1906.09235 (https://arxiv.org/abs/1906.09235) [cs.LG (https://arxiv.org/archive/cs.LG)].
141. Xu, Zhiqin John; Zhou, Hanxu (18 May 2021). "Deep Frequency Principle Towards Understanding Why Deeper Learning Is Faster" (https://ojs.aaai.org/index.php/AAAI/article/view/17261). *Proceedings of the AAAI Conference on Artificial Intelligence*. **35** (12): 10541–10550. arXiv:2007.14313 (https://arxiv.org/abs/2007.14313). ISSN 2374-3468 (https://www.worldcat.org/issn/2374-3468).

142. Crick, Francis (1989). "The recent excitement about neural networks". *Nature*. **337** (6203): 129–132. Bibcode:1989Natur.337..129C (https://ui.adsabs.harvard.edu/abs/1989Natur.337..129C). doi:10.1038/337129a0 (https://doi.org/10.1038%2F337129a0). PMID 2911347 (https://pubmed.ncbi.nlm.nih.gov/2911347). S2CID 5892527 (https://api.semanticscholar.org/CorpusID:5892527).
143. Adrian, Edward D. (1926). "The impulses produced by sensory nerve endings" (https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1514809). *The Journal of Physiology*. **61** (1): 49–72. doi:10.1113/jphysiol.1926.sp002273 (https://doi.org/10.1113%2Fjphysiol.1926.sp002273). PMC 1514809 (https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1514809). PMID 16993776 (https://pubmed.ncbi.nlm.nih.gov/16993776).
144. Dewdney, A. K. (1 April 1997). *Yes, we have no neutrons: an eye-opening tour through the twists and turns of bad science* (https://books.google.com/books?id=KcHaAAAAMAAJ&pg=PA82). Wiley. p. 82. ISBN 978-0-471-10806-1.
145. NASA – Dryden Flight Research Center – News Room: News Releases: NASA NEURAL NETWORK PROJECT PASSES MILESTONE (http://www.nasa.gov/centers/dryden/news/NewsReleases/2003/03-49.html). Nasa.gov. Retrieved on 20 November 2013.
146. "Roger Bridgman's defence of neural networks" (https://web.archive.org/web/20120319163352/http://members.fortunecity.com/templarseries/popper.html). Archived from the original (http://members.fortunecity.com/templarseries/popper.html) on 19 March 2012. Retrieved 12 July 2010.
147. D. J. Felleman and D. C. Van Essen, "Distributed hierarchical processing in the primate cerebral cortex (https://archive.today/20150120022056/http://cercor.oxfordjournals.org/content/1/1/1.1.full.pdf+html)," *Cerebral Cortex*, 1, pp. 1–47, 1991.
148. J. Weng, "Natural and Artificial Intelligence: Introduction to Computational Brain-Mind (https://www.amazon.com/Natural-Artificial-Intelligence-Introduction-Computational/dp/0985875720)," BMI Press, ISBN 978-0-9858757-2-5, 2012.
149. Edwards, Chris (25 June 2015). "Growing pains for deep learning". *Communications of the ACM*. **58** (7): 14–16. doi:10.1145/2771283 (https://doi.org/10.1145%2F2771283). S2CID 11026540 (https://api.semanticscholar.org/CorpusID:11026540).
150. Cade Metz (18 May 2016). "Google Built Its Very Own Chips to Power Its AI Bots" (https://www.wired.com/2016/05/google-tpu-custom-chips/). *Wired*.
151. "Scaling Learning Algorithms towards AI" (http://yann.lecun.com/exdb/publis/pdf/bengio-lecun-07.pdf) (PDF).
152. Tahmasebi; Hezarkhani (2012). "A hybrid neural networks-fuzzy logic-genetic algorithm for grade estimation" (https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4268588). *Computers & Geosciences*. **42**: 18–27. Bibcode:2012CG.....42...18T (https://ui.adsabs.harvard.edu/abs/2012CG.....42...18T). doi:10.1016/j.cageo.2012.02.004 (https://doi.org/10.1016%2Fj.cageo.2012.02.004). PMC 4268588 (https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4268588). PMID 25540468 (https://pubmed.ncbi.nlm.nih.gov/25540468).

Bibliography

- Bhadeshia H. K. D. H. (1999). "Neural Networks in Materials Science" (http://www.msm.cam.ac.uk/phase-trans/abstracts/neural.review.pdf) (PDF). *ISIJ International*. **39** (10): 966–979. doi:10.2355/isijinternational.39.966 (https://doi.org/10.2355%2Fisijinternational.39.966).
- Bishop, Christopher M. (1995). *Neural networks for pattern recognition*. Clarendon Press. ISBN 978-0-19-853849-3. OCLC 33101074 (https://www.worldcat.org/oclc/33101074).
- Borgelt, Christian (2003). *Neuro-Fuzzy-Systeme : von den Grundlagen künstlicher Neuronaler Netze zur Kopplung mit Fuzzy-Systemen*. Vieweg. ISBN 978-3-528-25265-6. OCLC 76538146 (https://www.worldcat.org/oclc/76538146).

- Cybenko, G.V. (2006). "Approximation by Superpositions of a Sigmoidal function" (https://books.google.com/books?id=4RtVAAAAMAAJ&pg=PA303). In van Schuppen, Jan H. (ed.). *Mathematics of Control, Signals, and Systems*. Springer International. pp. 303–314. PDF (https://web.archive.org/web/20110719183058/http://actcomm.dartmouth.edu/gvc/papers/approx_by_superposition.pdf)
- Dewdney, A. K. (1997). *Yes, we have no neutrons : an eye-opening tour through the twists and turns of bad science*. New York: Wiley. ISBN 978-0-471-10806-1. OCLC 35558945 (https://www.worldcat.org/oclc/35558945).
- Duda, Richard O.; Hart, Peter Elliot; Stork, David G. (2001). *Pattern classification* (2 ed.). Wiley. ISBN 978-0-471-05669-0. OCLC 41347061 (https://www.worldcat.org/oclc/41347061).
- Egmont-Petersen, M.; de Ridder, D.; Handels, H. (2002). "Image processing with neural networks – a review". *Pattern Recognition*. **35** (10): 2279–2301. CiteSeerX 10.1.1.21.5444 (https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.21.5444). doi:10.1016/S0031-3203(01)00178-9 (https://doi.org/10.1016%2FS0031-3203%2801%2900178-9).
- Fahlman, S.; Lebiere, C (1991). "The Cascade-Correlation Learning Architecture" (http://www.cs.iastate.edu/~honavar/fahlman.pdf) (PDF).
  - created for National Science Foundation, Contract Number EET-8716324, and Defense Advanced Research Projects Agency (DOD), ARPA Order No. 4976 under Contract F33615-87-C-1499.
- Gurney, Kevin (1997). *An introduction to neural networks*. UCL Press. ISBN 978-1-85728-673-1. OCLC 37875698 (https://www.worldcat.org/oclc/37875698).
- Haykin, Simon S. (1999). *Neural networks : a comprehensive foundation*. Prentice Hall. ISBN 978-0-13-273350-2. OCLC 38908586 (https://www.worldcat.org/oclc/38908586).
- Hertz, J.; Palmer, Richard G.; Krogh, Anders S. (1991). *Introduction to the theory of neural computation*. Addison-Wesley. ISBN 978-0-201-51560-2. OCLC 21522159 (https://www.worldcat.org/oclc/21522159).
- *Information theory, inference, and learning algorithms*. Cambridge University Press. 25 September 2003. Bibcode:2003itil.book.....M (https://ui.adsabs.harvard.edu/abs/2003itil.book.....M). ISBN 978-0-521-64298-9. OCLC 52377690 (https://www.worldcat.org/oclc/52377690).
- Kruse, Rudolf; Borgelt, Christian; Klawonn, F.; Moewes, Christian; Steinbrecher, Matthias; Held, Pascal (2013). *Computational intelligence : a methodological introduction*. Springer. ISBN 978-1-4471-5012-1. OCLC 837524179 (https://www.worldcat.org/oclc/837524179).
- Lawrence, Jeanette (1994). *Introduction to neural networks : design, theory and applications*. California Scientific Software. ISBN 978-1-883157-00-5. OCLC 32179420 (https://www.worldcat.org/oclc/32179420).
- MacKay, David, J.C. (2003). *Information Theory, Inference, and Learning Algorithms* (http://www.inference.phy.cam.ac.uk/itprnn/book.pdf) (PDF). Cambridge University Press. ISBN 978-0-521-64298-9.
- Masters, Timothy (1994). *Signal and image processing with neural networks : a C++ sourcebook*. J. Wiley. ISBN 978-0-471-04963-0. OCLC 29877717 (https://www.worldcat.org/oclc/29877717).
- Maurer, Harald (2021). *Cognitive science : integrative synchronization mechanisms in cognitive neuroarchitectures of the modern connectionism*. CRC Press. doi:10.1201/9781351043526 (https://doi.org/10.1201%2F9781351043526). ISBN 978-1-351-04352-6. S2CID 242963768 (https://api.semanticscholar.org/CorpusID:242963768).
- Ripley, Brian D. (2007). *Pattern Recognition and Neural Networks* (https://books.google.com/books?id=m12UR8QmLqoC). Cambridge University Press. ISBN 978-0-521-71770-0.

- Siegelmann, H.T.; Sontag, Eduardo D. (1994). "Analog computation via neural networks". *Theoretical Computer Science*. **131** (2): 331–360. doi:10.1016/0304-3975(94)90178-3 (https://doi.org/10.1016%2F0304-3975%2894%2990178-3). S2CID 2456483 (https://api.semanticscholar.org/CorpusID:2456483).
- Smith, Murray (1993). *Neural networks for statistical modeling*. Van Nostrand Reinhold. ISBN 978-0-442-01310-3. OCLC 27145760 (https://www.worldcat.org/oclc/27145760).
- Wasserman, Philip D. (1993). *Advanced methods in neural computing*. Van Nostrand Reinhold. ISBN 978-0-442-00461-3. OCLC 27429729 (https://www.worldcat.org/oclc/27429729).
- Wilson, Halsey (2018). *Artificial intelligence*. Grey House Publishing. ISBN 978-1-68217-867-6.

Retrieved from "https://en.wikipedia.org/w/index.php?title=Artificial_neural_network&oldid=1100988343"

This page was last edited on 28 July 2022, at 18:31 (UTC).

Text is available under the Creative Commons Attribution-ShareAlike License 3.0; additional terms may apply. By using this site, you agree to the Terms of Use and Privacy Policy. Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a non-profit organization.