Hopfield Networks, Boltzmann Machines and Feed-forward Networks

Michael Wollowski

Summary of the 2024 Physics Nobel Prize Documents:

https://www.nobelprize.org/uploads/2024/09/advanced-physicsprize2024.pdf

https://www.nobelprize.org/uploads/2024/10/popular-physicsprize2024-2.pdf

Thinking

- "Although computers cannot think, machines can now mimic functions such as memory and learning."

Difference Between Rule-Based and NN AI

- Traditional AI works like a type of recipe.
  - The software receives data, which is processed according to a clear description and produces the results.
  - This is similar to how one might collect ingredients and processes them by following a recipe, producing a cake.
- In machine learning, the computer learns by example,
  - Enabling it to tackle problems that are too vague and complicated to be managed by step by step instructions.
  - One example is interpreting a picture to identify the objects in it.
- In summary: ANN are trained to perform certain tasks rather than asked to execute a predetermined set of instructions.

Historical Background

- The first electronic-based computers appeared in the 1940s, and were invented for military and scientific purposes.
- They were intended to carry out computations that were cumbersome and time-consuming for humans.
- In the 1950s, the opposite need emerged, namely to get computers to do what humans and other mammals are good at – pattern recognition.

Characteristics of NN

- "An artificial neural network processes information using the entire network structure."

Natural and artificial neurons

The brain's neural network is built from living cells, neurons, with advanced internal machinery. They can send signals to each other through the synapses. When we learn things, the connections between some neurons get stronger, while others get weaker.

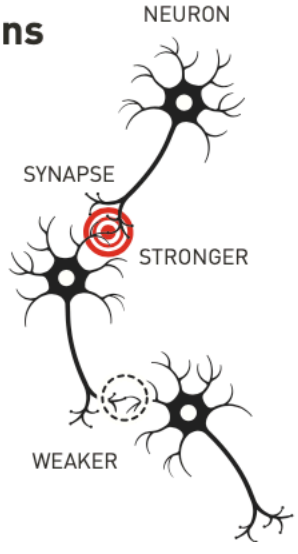

Artificial neural networks are built from nodes that are coded with a value. The nodes are connected to each other and, when the network is trained, the connections between nodes that are active at the same time get stronger, otherwise they get weaker.

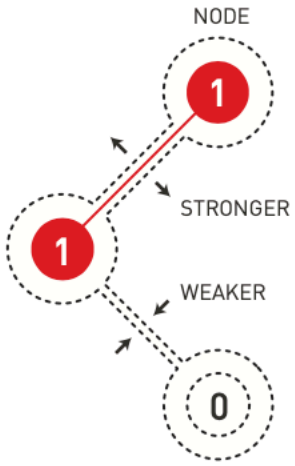

Associative Memory

- Imagine that you are trying to remember a fairly unusual word that you rarely use, such as one for that sloping floor often found in cinemas and lecture halls.
- You search your memory. It's something like ramp… perhaps rad…ial? No, not that. Rake, that's it!
- This process of searching through similar words to find the right one is reminiscent of the associative memory that the physicist John Hopfield discovered in 1982.
- The Hopfield network can store patterns and has a method for recreating them.
- When the network is given an incomplete or slightly distorted pattern, the method can find the stored pattern that is most similar.

Science and Sticking to your Guns

- When Hopfield was invited to a meeting about neuroscience he encountered research into the structure of the brain.
- He was fascinated by what he learned and started to think about the dynamics of simple neural networks.
- When neurons act together, they can give rise to new and powerful characteristics that are not apparent to someone who only looks at the network's separate components.
- In 1980, Hopfield left his position at Princeton University, where his research interests had taken him outside the areas in which his colleagues in physics worked, and moved across the continent.
- He had accepted the offer of a professorship in chemistry and biology at Caltech.
- There, he had access to computer resources that he could use for free experimentation and to develop his ideas about neural networks.

Intuition about Hopfield Networks

- Hopfield networks are pattern “re-constructors.”
- Give it the pattern on the left, it will reproduce the pattern on the right.

- Why not just save the image itself and compare it to another image being tested?
- Hopfield’s method is special because several pictures can be saved at the same time and the network can usually differentiate between them.

Intuition about Hopfield Networks

- Hopfield likened searching the network for a saved state to rolling a ball through a landscape of peaks and valleys, with friction that slows its movement.
- If the ball is dropped in a particular location, it will roll into the nearest valley and stop there.
- If the network is given a pattern that is close to one of the saved patterns it will, in the same way, keep moving forward until it ends up at the bottom of a valley in the energy landscape, thus finding the closest pattern in its memory.

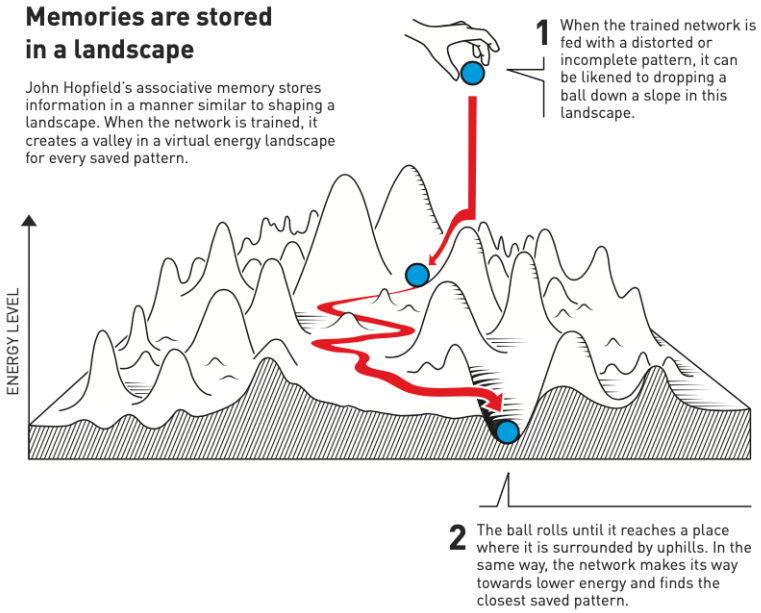

Recurrent and Feedforward Networks

- Two architectures for systems of interconnected nodes were explored, "recurrent" and "feedforward" networks.
- The former allows for feedback interactions

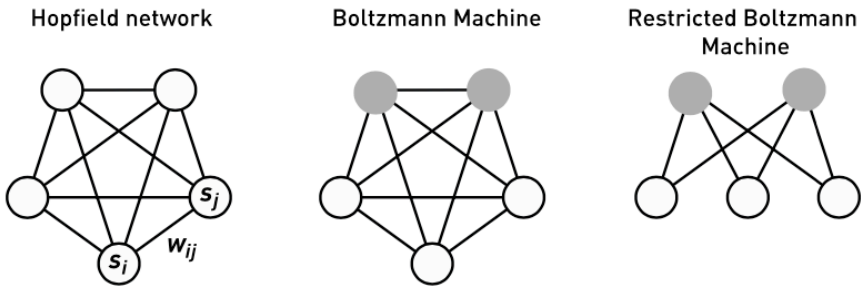

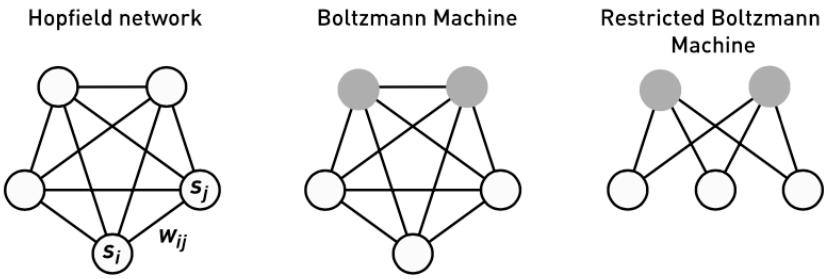

Figure 1. *Recurrent networks of N binary nodes s_i (0 or 1), with connection weights w_{ij}. (Left) The Hopfield model. (Centre) Boltzmann machine. The nodes are divided into two groups, visible (open circles) and hidden (grey) nodes. The network is trained to approximate the probability distribution of a given set of visible patterns. Once trained, the network can be used to generate new instances from the learned distribution. (Right) Restricted Boltzmann Machine (RBM). Same as the Boltzmann machine, but without any couplings within the visible layer or between hidden nodes. This variant can be used for layer-by-layer pre-training of deep networks.*

Boltzmann Machines

- The Boltzmann machine is commonly used with two different types of nodes.
- Information is fed to one group, which are called visible nodes.
- The other nodes form a hidden layer.
- The hidden nodes' values and connections also contribute to the energy of the network as a whole.

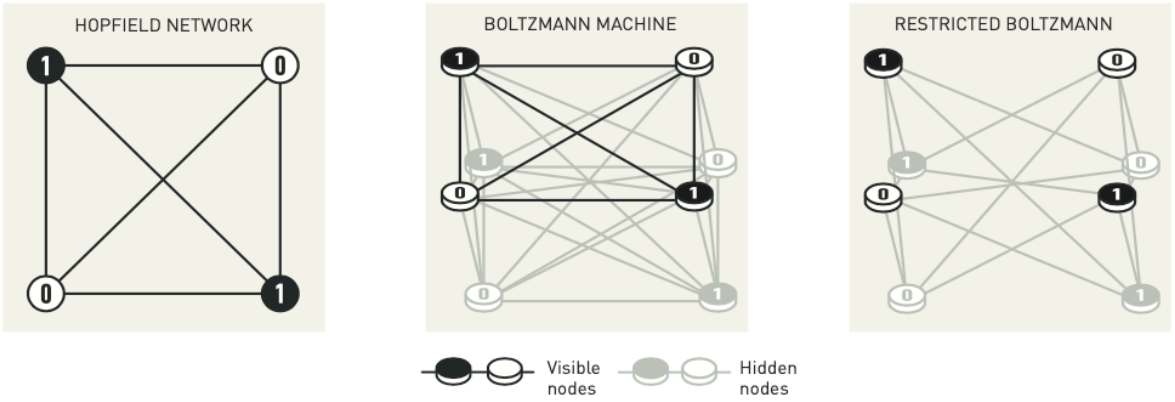

John Hopfield's associative memory is built so that all the nodes are connected to each other. Information is fed in and read out from all the nodes.

Geoffrey Hinton's Boltzmann machine is often constructed in two layers, where information is fed in and read out using a layer of *visible* nodes. They are connected to *hidden* nodes, which affect how the network functions in its entirety.

In a restricted Boltzmann machine, there are no connections between nodes in the same layer. The machines are frequently used in a chain, one after the other. After training the first restricted Boltzmann machine, the content of the hidden nodes is used to train the next machine, and so on.

Boltzmann Machines

- The Boltzmann machine can learn – not from instructions, but from being given examples.
- It is trained by updating the values in the network's connections so that the example patterns, which were fed to the visible nodes when it was trained, have the highest possible probability of occurring when the machine is run.
- If the same pattern is repeated several times during this training, the probability for this pattern is even higher.
- Training also affects the probability of outputting new patterns that resemble the examples on which the machine was trained.

Boltzmann Machines

- A trained Boltzmann machine can recognise familiar traits in information it has not previously seen.
- Imagine meeting a friend's sibling, and you can immediately see that they must be related.
- In a similar way, the Boltzmann machine can recognise an entirely new example if it belongs to a category found in the training material, and differentiate it from material that is dissimilar.

Boltzmann Machines

- Unlike the Hopfield model, it focuses on statistical distributions of patterns rather than individual patterns.
- It contains visible nodes that correspond to the patterns to be learned as well as additional hidden nodes, where the latter are included to enable modelling of more general probability distributions.

Restricted Boltzmann Machines

- While theoretically interesting, in practice, the Boltzmann machine was initially of limited use.
- It was fairly inefficient and takes a long time to find solutions.
- However, a slimmed-down version of it with fewer weights, called the restricted Boltzmann machine, developed into a versatile tool.

Feedforward Networks

- Both the Hopfield model and the Boltzmann machine are recurrent neural networks.
- The 1980s also saw important progress on feedforward networks.
- A key advance was the demonstration by David Rumelhart, Hinton and Ronald Williams in 1986 of how architectures with one or more hidden layers could be trained for classification using an algorithm known as backpropagation.
- Here, the objective is to minimize the mean square deviation, *D*, between output from the network and training data, by gradient descent.
- This requires computing the partial derivatives of *D* with respect to all weights in the network.

Feedforward Network

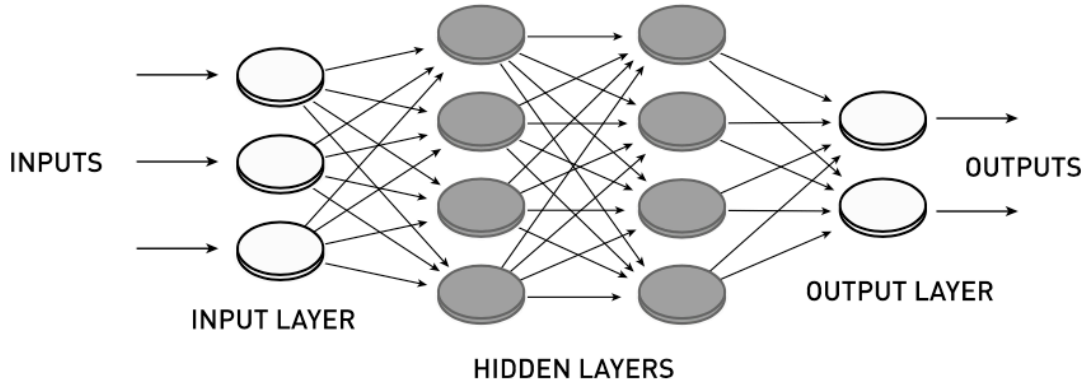

1982: Still not enough Resources to do NN work

- In Hopfield's article on associative memory, from 1982:
- He used a network with 30 nodes.
- If all the nodes are connected to each other, there are 435 connections.
- The nodes have their values, the connections have different strengths and, in total, there are fewer than 500 parameters to keep track of.
- He also tried a network with 100 nodes, but this was too complicated, given the computer he was using at the time.