



## Imperfect patterns

### Description

Training a neural network (NN) involves automatically adjusting numerical weights and thresholds (biases) to account for all input and output pairs presented to the NN. After training on these input-output patterns the NN should reproduce the appropriate output pattern when presented with any one of the input patterns.

Note that it is not patterns that are stored, but the method of constructing them. In their chapter "The appeal of Parallel distributed processing" • Humelhart, McClelland and Hinton, reinforce the point that

"the patterns themselves are not stored. Rather, what is stored is the connection strengths between units that allow these patterns to be re-created" • [31]

## PDP

A neural network is also called a "parallel distributed processing" (PDP) system, the term commonly used when we wrote our 1990 articles on the subject. As the name implies, PDP systems lend themselves to parallel computer architectures, though the whole process is generally simulated with iterative computer programs using standard architectures.

As pointed out by McClelland and Rumelhart in their seminal books on the subject, this leads to interesting and useful analogies between cognition and computing, and interesting explanations of cognitive phenomena. In summary, once they have been trained on a set of input and output patterns NNs can be regarded as operating in a distributed way to process input patterns. Input patterns are not reduced to primitive elements for analysis but are processed via the properties of the whole network. Where it occurs, the process of discovering similarities and categorizing input features is a by-product of the PDP process rather than part of a sequence designed into the network.

Grant Sanderson demonstrates this point in his animated video of a NN trained to recognise hand drawn integers. He explains that the hidden layers are there to detect the sorts of features a human being might recognise, particularly the relative sizes, angles and relationships between lines in the hand

drawn integers. However, he shows that any features in the weights and biases in the hidden layers after training are opaque to human inspection. Without the hidden layers the NN would be unable to generalise about line and shapes. So we conclude that feature detection that is at least legible to the NN must be occurring.

## Degrading gracefully

As I showed in previous posts, NNs are a good way of simulating the human capacity for pattern completion. Furthermore, successful matches between input and output patterns can still be accomplished with imperfect patterns. Where there is poor or ambiguous input the system will always produce some form of output. This may be an amalgamation of outputs or it could be the output corresponding to the most closely matching stored input pattern.

If a few nodes and arcs in the network have less than optimal values, or fail to perform as predicted then that may introduce noise into the performance of the network, but it is unlikely to fail entirely. Humelhart, McClelland and Hinton assert that the

pattern retrieval performance of the model degrades gracefully both under degraded input and under damage [36].

Destruction or modification to one part of the network results in uniform degradation of the performance of the network rather than the loss of any particular stored pattern. Unlike the effects of a bug in a procedural computer program, a neural network accommodates a degree of error and noise.

## Bibliography

- Clark, Andy. *Microcognition: Philosophy, Cognitive Science, and Parallel Distributed Processing*. Cambridge, MA: MIT Press, 1989.
- Coyne, Richard, and Arthur Postmus. "Spatial applications of neural networks in computer-aided design." *Artificial Intelligence in Engineering* 5, no. 1 (1990): 9-22.
- Hinton, G. E., and T. J. Sejnowski. "Learning and relearning in Boltzmann machines." In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 1, Foundations*, edited by D. E. Rumelhart, and J. L. McClelland, 282-314. Cambridge, MA: MIT Press, 1986.
- Rumelhart, D. E., and J. L. (eds) McClelland. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Volume 1, Fundamentals*. Cambridge, Mass.: MIT Press, 1987.
- McClelland, J. L., D. E. Rumelhart, and G. E. Hinton. "The appeal of parallel distributed processing." In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 1, Foundations*, edited by D. E. Rumelhart, and J. L. McClelland, 3-44. Cambridge, MA: MIT Press, 1986.
- Sanderson, Grant. "Neural Networks: The basics of neural networks, and the math behind how they learn." *3Blue1Brown*, 2023. Accessed 12 February 2023. <https://www.3blue1brown.com/topics/neural-networks>

## Note

- Featured image is from MidJourney prompted by "networks that degrade gracefully."

### Category

1. Artificial Intelligence

### Tags

1. neural networks
2. parallel distributed processing
3. PDP

### Date Created

April 1, 2023

### Author

rcoyne99

*default watermark*