

# An Introduction to Deep Learning: Part III

LASSE AMUNDSEN, HONGBO ZHOU and MARTIN LANDRØ

---

*There's this idea that ideas in science are a bit like epidemics of viruses. There are apparently five or six basic strains of flu viruses, and apparently each one comes back with a period of around 25 years. People get infected, and they develop an immune response, and so they don't get infected for the next 25 years. And then there is a new generation that is ready to be infected by the same strain of virus. In science, people fall in love with an idea, get excited about it, hammer it to death, and then get immunized – they get tired of it. So ideas should have the same kind of periodicity.*

*Tomaso Armando Poggio (1947–), Professor of MIT Brain and Cognitive Sciences, and Director of The Center for Brains, Minds and Machines.*

---

We learn by doing, Aristotle (384–322 BC) told us. Today, experiential learning does not apply just to children; computers and robots are increasingly able to act and adapt based on experience. Most successes so far in deep learning have been based on supervised learning. The holy grail, however, is unsupervised learning, as illustrated here by five-year-old girls figuring out how to row a boat. In the unsupervised approach, the computer program figures out what the data means on its own. It may detect characteristics that humans cannot see.

Lasse Amundsen



Previously in this series we showed that deep learning is nothing other than neural networks – an approach to artificial intelligence (AI) which has been going in and out of fashion since the 1950s. Deep learning is a ‘strain of virus’, and since 2012 it has returned with great excitement and speed, accelerated by the increased processing power of graphics chips – as well as Big Data (all the images, videos, emails, driving patterns, phrases, objects and so on that are used to train the networks).

Examples of AI systems utilising deep learning range from speech recognisers on smartphones and automatic translators, to the technology behind driverless cars, enabling them to identify a stop sign or to distinguish a pedestrian from a lamppost. Current automated driving development requires millions of images and thousands of hours of video to train the neural network. Many automakers have announced plans to offer fully automated cars in the future. The driverless car is coming, but will we actually want to cede control?

Part II of this series mentioned a deep learning computer model that learns to perform classification tasks directly from images. The one that started it all was the 2012 publication ‘ImageNet Classification with Deep Convolutional Networks’ by Krizhevsky, Sutskever and Hinton, which discussed the architecture of the network, called AlexNet. They used a relatively simple layout, compared to modern architectures. The network was made up of only eight layers, where the first five were convolutional layers followed by fully connected layers.

### Deep Learning with MATLAB

Here, we show that you can apply deep learning with MATLAB using AlexNet: it is not perfect, but it is fun. MATLAB is a numerical computing environment with over 2 million users across industry and academia. One of MATLAB’s developers has made a demonstration of how to apply deep learning in just 11 lines of code, which sets up your web camera to perform object recognition. First, the camera acquires images, then AlexNet takes the image as input and provides a label for the object in the image. All you need to do is to download the webcam support package and the AlexNet support package into MATLAB.

Let us experiment with objects in our surroundings to see how accurate AlexNet is. The first object that we let the webcam see is a pile of bananas. It is easily recognised by AlexNet. In the second example, we put a Powerpoint-made page of the cover of the book *Introduction to Exploration Geophysics with Recent Advances* in front of the webcam. AlexNet now suggests “comic book”. The authors of the book are amazed that AlexNet recognises this is a book. Perhaps it is comic, too? It is up to the readers to judge.



```
clear
camera = webcam; % Connect to the camera
nnet = alexnet; % Load the neural net

while true
    picture = camera.snapshot; % Take a picture
    picture = imresize(picture, [227,227]); % Resize the picture

    label = classify(nnet,picture); % Classify the picture

    image(picture); % Show the picture
    title(char(label)); % Show the label
    drawnow;
end
```

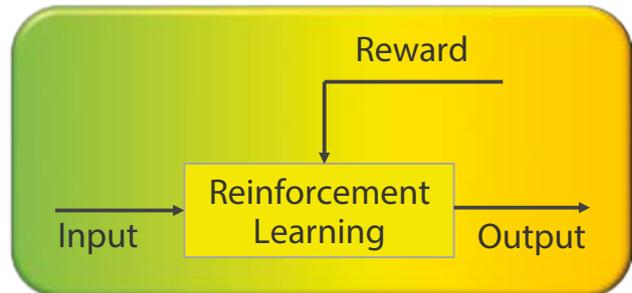
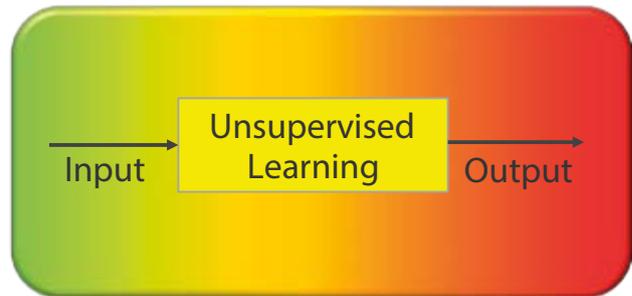
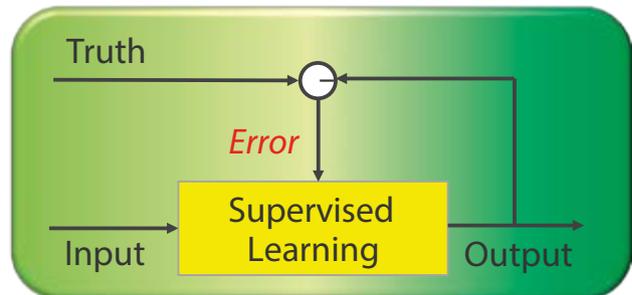
*Deep learning in 11 lines of MATLAB code. All you need is MATLAB, a simple webcam, and a deep neural network – AlexNet – to identify objects in your surroundings (Hicklin, 2017). A few years ago, this example would have been considered science fiction.*

### Supervised Learning

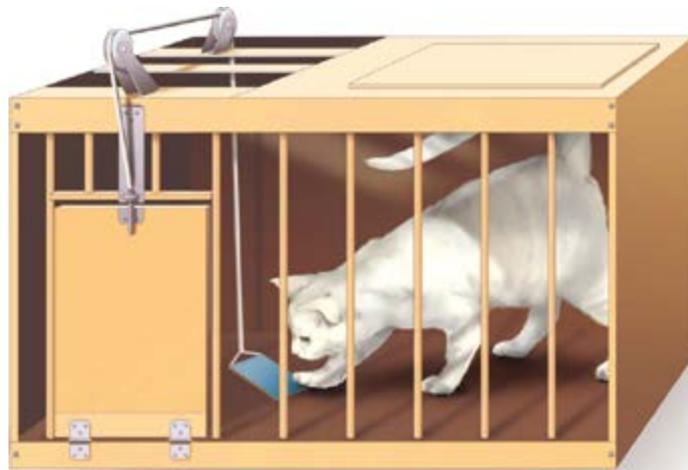
In Part II, we briefly discussed how deep learning can be used in supervised learning, in reinforcement learning, and in unsupervised learning.

In supervised learning, during training, data are fed to the network but have been painstakingly labelled in advance. This process can be thought of as a ‘teacher’ supervising the learning process. The teacher knows the correct answers, the algorithm iteratively makes predictions on the training

*Supervised learning relies on data where the true class of the data (truth) is given. Unsupervised learning means that the deep learning algorithm does not have any labels attached to supervise the learning. Reinforcement learning focuses on the end outcome to learn. A reward function is provided that tells how ‘good’ certain states are.*



## Recent Advances in Technology



*The American behaviourist psychologist, Edward Lee Thorndike (1874–1949), placed cats inside boxes from which they could escape only by pressing a lever, to examine the phenomenon of reinforcement in learning.*

data and is corrected by the teacher. Learning stops when the algorithm achieves an acceptable level of performance. For example, to train a neural network to identify pictures of bananas, it needs to be fed images that are labelled bananas. The network finds what all pictures labelled banana have in common, so they can eventually use those recognised patterns to more accurately predict what they are seeing in new images. The more labelled pictures the network sees, the better it can refine the accuracy of its predictions.

Supervised learning has shown impressive results for image recognition. For instance, Facebook's deep learning software is able to match two images of an unfamiliar person at the same level of accuracy as a human. Researchers from Stanford University in January 2017 published a deep learning algorithm that recognises skin cancer in photos as well as dermatologists do. It was built on the architecture of the GoogleNet Inception v3, a convolutional neural network algorithm, and trained on a set of nearly 130,000 images of skin lesions from more than 2,000 diseases. In the future, might you test skin anomalies with your smart phone?

Astrophysicists search the galactic sky for exoplanets (planets outside our own solar system) by analysing large amounts of data from NASA's Kepler Space Telescope with both automated software and manual analysis. Kepler observed about 200,000 stars for four years, taking one picture every 30 minutes, creating about 14 billion data points. The data points essentially are light readings – the minuscule change in brightness captured when a planet passed in front of a star.

Researchers from Google and University of Texas at Austin announced in December 2017 that a deep learning algorithm had discovered two new planets, Kepler 80g and Kepler 90i. The algorithm was trained on 15,000 signals from potential planets in the Kepler database. When the scientists tested their model on signals it had never seen before, it correctly identified which signals were planets and which signals were not planets 96% of the time. Then, the scientists took the trained system and set it to hunt for new planets on data from 670 stars that were already known to have multiple planets, as they considered those to be the most likely hiding places. Interestingly, the planet Kepler 90i is the eighth planet discovered orbiting the Kepler-90 star, making it the first known eight-planet system outside our own solar system.

### Maximising Rewards

Reinforcement learning is the problem of getting the computer program to act so as to maximise its rewards. This learning copies a simple principle from nature. The American psychologist Edward Lee Thorndike (1874–1949) placed cats inside boxes from which they could escape only by pressing a lever. After wandering around restlessly and meowing, the cat would eventually step on the lever by chance. After it learned to associate this behaviour with the desired outcome – open door and food outside – it eventually escaped with increasing speed. Thorndike put forward the 'law of effect', stating that rewards for appropriate behaviour always substantially strengthened associations (connections), whereas punishments for inappropriate responses only slightly weakened the association between the stimulus and the wrong response. We can use a similar method to train computers to do many tasks.

DeepMind's AlphaGo computer program that astonishingly beat the world Go champion in 2016 (see Part II) was largely based on reinforcement learning and Monte Carlo tree search. Recently, a descendant of AlphaGo, dubbed AlphaGo Zero (now AlphaZero) because it didn't need any human input, achieved tabula rasa superhuman performance in the game of Go, by representing Go knowledge using deep convolutional neural networks, trained

*AlphaGo Zero can beat all opponents in the game of Go, trained solely by reinforcement learning from games of self-play.*



solely by reinforcement learning from games of self-play. In December 2017, starting from random play, and given no domain knowledge except the game rules, researchers from DeepMind reported that AlphaZero crushed the current world champion chess-playing program, Stockfish 8. In 100 games AlphaZero scored 25 wins and 25 draws with White, while with Black it scored 3 wins and 47 draws. The chess world was shaken.

Reinforcement learning will soon be applicable to more than games. In addition to improving self-driving cars, the technology can now get a robot to grasp objects it has never seen before. But there are challenges to overcome. This learning approach requires huge amounts of data. Successes have come when the computer was able to practise relentlessly in simulations.

### **Unsupervised Learning**

In unsupervised learning, there are no correct answers and no teacher. Algorithms are left to their own devices to dive into the data to discover alone, looking for patterns and connections, and structures in the data. An example of unsupervised learning is: we have lots of photos of N people. There is no information who is in which photo. We want to divide this dataset into N piles, each with photos of one individual.

In 2012, Google demonstrated a deep learning network that was able to decipher cats, faces and other objects from a giant pile of unlabelled images. Unfortunately, so far, unsupervised learning cannot compete or match the accuracy and effectiveness of supervised training. The new buzz word that seems to replace unsupervised learning is 'predictive learning'. Predictive learning is like the astrophysicist's dark matter. Dark matter does not emit or absorb light. We know it is there, but we just do not know how to see it directly.

### **The Future**

Deep learning has had many impressive successes, primarily in image and text recognition, and board gaming. But deep learning is only a small part of machine learning, which is a small part of artificial intelligence. The future of AI may explore ways beyond deep learning. Most of the important learning we do as humans is experiential – and unsupervised. Just think about it. Or look at the picture of the two girls rowing. You teach your children plenty, but after all, their most important learning is unsupervised. Just like humans, we expect that computers and robots will continue to improve as they learn... unsupervised... by doing. When we say that neural networks mimic the brain, that is true at the level of loose inspiration. To build true AI, computer scientists need a better model of how the biological brain works.

*References available online.*

*Download the webcam support package from here:*

<https://www.mathworks.com/matlabcentral/fileexchange/45182-matlab-support-package-for-usb-webcams>

*Download the AlexNet support package from here:*

<https://www.mathworks.com/matlabcentral/fileexchange/59133-neural-network-toolbox-tm--model-for-alexnet-network> ■