Philosophy and the Sciences

Introduction to the Philosophy of the Cognitive Sciences

EIDYN

Intelligent machines and the human brain

How does one make a clever adaptive machine that can recognise speech, control an aircraft, and detect credit card fraud?

Machine learning and computation

Recent years have seen a revolution in the kinds of tasks computers can do. Underlying these advances is the burgeoning field of machine learning and computational neuroscience.

Machine learning aims to create machines that can learn from their experiences.

A **computation** is a way of solving a problem by following a set of instructions, or a recipe, called an algorithm.

The father of our modern idea of computation is **Alan Turing**. Alan Turing was an English mathematician who discovered what's now known as a universal computing machine.

The universal computing machine that Turing described is known today as the **Universal Turing Machine**. It consists of a long paper tape and a head that can scan along the tape and read and write symbols, guided by a simple set of instructions.

Computational description of the brain

In the 1970s, a cognitive scientist called **David Marr** said that computation could help us to answer three different questions about the brain. First, which task does the brain solve? Second, how does the brain solve that task? Third, why is that task important for the brain to solve? Marr grouped these three questions into three different levels of **computational description**.

Computational level: description of which mathematical function the device computes.

Algorithmic level: concerns how the device solves its task. How does it compute the function?

Implementation level: it describes the role each one of the physical parts of the device plays in implementing the device's algorithm.

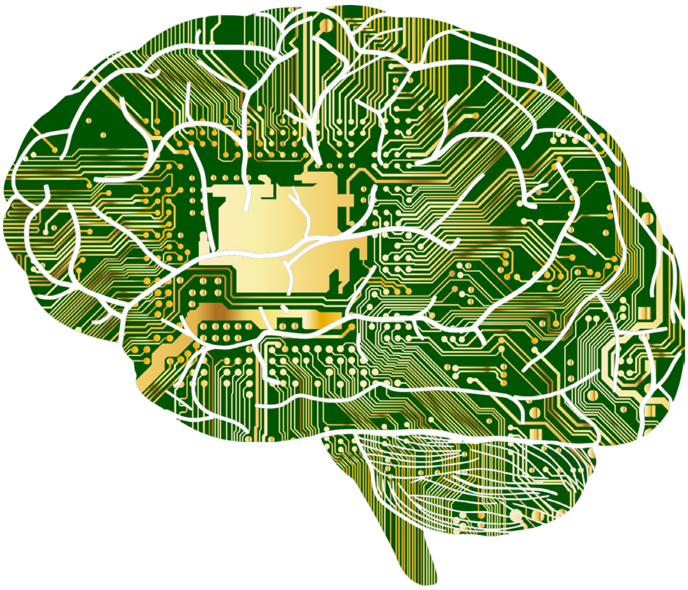

The brain as a probabilistic machine

The human brain is one of the most complex objects in the universe. It has over 100 billion neurons and a **complex web** of approximately a quadrillion connections. The brain performs not one, but many different computations simultaneously.

Recently a new hypothesis has emerged regarding the types of computations that the brain might perform. The idea is that the brain might work like a **probabilistic machine.**

Hermann Von Helmholtz proposed that visual perception was the result of what he called a process of unconscious inference. Through this process the brain would complete missing information using past knowledge and construct a hypothesis about our environment.

The brain works by constantly forming **hypotheses**, or beliefs, about the environment, and the actions to take. Those hypotheses can be described mathematically as conditional probabilities.

Statisticians have shown that the best way to compute this probability is to use **Bayes' Rule**, named after Thomas Baye.

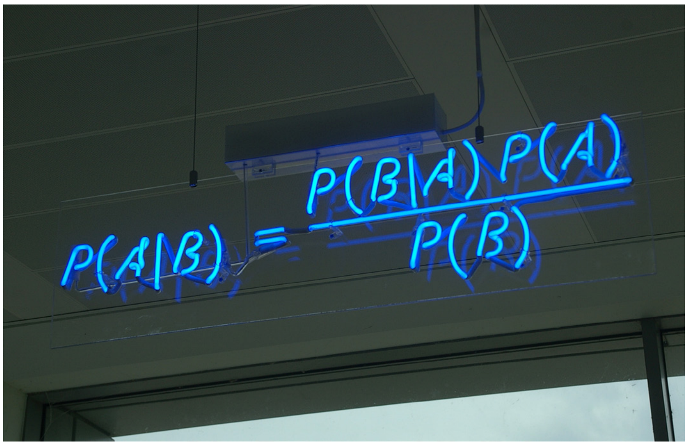

Bayes' theorem spelt out in blue neon at the offices of Autonomy in Cambridge, by mattbuck, CC-BY SA 3.0

Bayes' rule

Bayes' rule states that we can get the probability P of the hypothesis A given the data B by multiplying two other probabilities. First P of the data given the hypothesis (the likelihood) and second P of the hypothesis, which we call the prior probability.

The **likelihood** and the **prior probability** form an internal modern of the world inside the brain.

How can we test if the brain is doing something like Bayesian inference?

The **Bayesian model** predicts not only how to combine different pieces of information but also how to incorporate prior knowledge.

A lot of very important questions remain. Where do those prior beliefs come from? How do we learn them? Are they the same for everybody? How do they depend on our previous experience?

The reality that we perceive is consistent with our expectations but not with the actual physical world. We make such assumptions all the time. For example we expect objects to be symmetrical and to change smoothly in space and time, orientations to be more frequently horizontal or vertical and angles to look like perpendicular corners.

In situation of strong **uncertainty**, or when the objects don't conform to the average statistics, using expectations can lead to **illusions**.

This is illustrated, for example, by a classic illusion known as the hollow mask illusion (right).

We tend to perceive reality as being more similar to what we expect than it really is.

Objects will be seen as being slower, more symmetrical, or maybe smoother in space and time, etcetera. The Bayesian approach helps formalizing these ideas in a **quantitative way.**

Experimental work shows that our brains create prior expectations automatically and unconsciously all the time. We collect information about our environment and try to use it to predict what could come next.

Our brain constantly revises its assumptions and updates its internal model of the world.

The idea that the brain would work like a **probabilistic machine** is not restricted to perception, but has been applied to all domains of cognition. In particular, Bayesian models could be very useful for **psychiatry**.

Bayesian models are very useful for describing perception and behaviour at the computational level as Mark explained. However, how those algorithms are implemented in the brain is still very unclear. In fact it is still quite debated whether Bayesian models can make predictions for **neurobiology**.