What is AI?

Written by Melanie Koehler

Data Scientist at Seenit

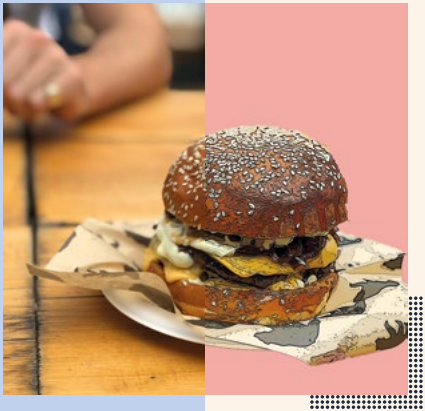

Seenit

1

AI IN GENERAL

A short history of AI

The idea of artificial beings and intelligent robots was first recognised in Greek mythology during the antiquity period. This concept ran through time and today we are increasingly relying on artificial intelligence (AI) in many areas of our daily lives.

AI as we know it today, first appeared in the mid 20th century when computers would play Draughts against mid-level players and win. The U.S. Department of Defence saw this potential and funded many AI projects. But progress was slow, and aside from short revivals in the 1980s

and 1990s, the interest in AI declined – a period known as the 'AI Winter'.

AI only really started to see success in 2015 when computer power became affordable, theoretical knowledge increased, and large volumes of data could be saved and processed far more efficiently. Computers were now trained to carry out even the most complicated tasks in a very short time.

Artificial intelligence

/ɑːtɪˈfɪʃ(ə)l/ /ɪnˈtɛlɪdʒ(ə)ns/

noun: artificial intelligence; noun: AI

The theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.

Source: Oxford Dictionary

My understanding of artificial intelligence is that it is computer systems that are built to basically mimic human intelligence.

Brendan Freer
Senior Training Manager at Seenit

How did AI evolve?

The first successes of AI were driven by Neural Networks in the 1950s and 1960s. The idea was to copy the human mind so that computers could learn and solve problems, something that became very popular between the 1980s and 2010.

Today, the most common approach is deep learning. Deep learning, or more specifically 'Deep Neural Networks', is a sub-field of artificial intelligence and machine learning.

Data from a database is transported through a system of layers, that can be rather 'deep', hence the name. For

example, an image of a cat is processed in four layers. The input is a matrix of numbers. In the first input edges and lines are encoded, in the next layer edges and lines are connected, in the third layer the eyes and ears are encoded, and in the last layer the cat is recognised. The code can learn what to process in each layer on its own. The challenge for humans is finding the best number and size of the layers in order to get the best accuracy rate.

How does AI work?

Computers need external information to learn from, which is supplied either from a database or in the form of rewards.

Three ways of learning are distinguished:

Supervised learning

Supervised learning needs a database that is already classified, meaning each input has a result; a so-called label. As an example, images of elephants and zebras are labelled ‘elephant’ and ‘zebra’ respectively. The computer then learns patterns for both ‘elephant’ and ‘zebra’ and can recognise these

in new images. This form of learning is commonly used to train computers to carry out specific tasks like object recognition. Research shows that the classification works best with a large database and clear labels.

Unsupervised learning

For unsupervised learning the input data is unlabelled and therefore the computer must find the pattern itself.

Reinforced learning

A third option is to give feedback to the computer each time it carries out a task, allowing it to learn what is ‘good’ or ‘bad’. The computer will first start with random decisions

and feedback is given at each step. In this way the computer learns through rewards and penalties. This form of learning can be used in training a computer to play a game, such as Chess or Go, and even beat world champions.

What is crucial for success with deep learning?

For supervised and unsupervised learning, the most crucial point is the database. It is integral to base the training on the biggest database available. In competition, even when

using the same techniques, the best dataset will always win; the larger the database the better the result will be. For supervised learning it is also crucial to use the correct labels, as seen in the previous example with ‘elephants’ and ‘zebras’. Adding new data to the database ensures the computer continuously improves its learning.

In reinforced learning, the accuracy of the output depends on the quality of the feedback.

What can go wrong?

There is a famous expression in AI, garbage in, garbage out.

For supervised learning, it is generally considered that classification results are bad when the dataset exists either of only a few inputs or where the labels are poorly assigned. For example there are only 2 images of a Zebra and one is labelled as ‘elephant’. From this input computers cannot learn anything.

Furthermore, bad datasets can be biased, miss diversity, or contain incorrect data. For example, in credit scoring, the training is biased if the

database only holds information on clients who have credit but not on those who do not have credit. Another example of bias is training a code that distinguishes between male and female behaviour based on a database that mainly holds information on men.

For reinforced learning, feedback has to be accurate, otherwise we run the risk of the computer learning wrong decisions.

AI and human creativity are immensely powerful when working together, as AI helps us to be more productive with our time, taking away some of the heavy lifting of process, allowing us to focus on imagination and originality.

Emily Forbes
CEO at Seenit

Photo by Mahdis Mousavi

Where do we use AI today?

Today AI is used in a number of applications, many of which we take for granted or simply don't know are AI driven, such as smart search systems, speech recognition, financial predictions and even self-driven cars.

But there are other areas in which AI operates that are not commonly known. In banking, AI is used for financial fraud detection and credit scoring, in medicine and health care, AI is used for tumour finding and as a health care assistant. In fact, one study showed that a trained AI code using deep learning could

diagnose breast cancer better than a group of pathologists.

There are also instances of a more creative approach to AI, where data from different disciplines are combined, like generating information on the economy by counting swimming pools from images, or making investments in agriculture by analysing satellite images.

Artistic style transfer using AI

What are the limits for AI?

AI is primarily used to help us perform better at what we are doing.

Computers are trained to perform specific tasks and not a diversity of tasks at the same time. A code that learnt to distinguish an elephant from a zebra will not be able to play Chess or even recognise a dog.

So you see, computers learn very differently from how humans learn and they will probably never be as diverse thinkers as human beings. But although each code is trained to perform only one task, this one task is carried out in

a very short time with high precision (if the code is well trained by us of course).

For this single task, the computer is likely to be faster and less prone to errors than our brain which is slower, and in many cases, can be distracted and make misleading assumptions.

AI and ethics

With great power comes great responsibility, and AI can be used for both good and bad. It can be used for both detecting and committing fraud as well as recommending systems and influencing people's behaviours.

We can do so much with AI, but it raises a lot of questions. Who is responsible if a surgery robot kills the patient or a self-driven car injures a pedestrian? Insurance companies must now be prepared for those instances.

There will come a time when no one will be able to control the use of AI just as much as we can't control what is on the internet. Each of us should make sure we abide to ethical standards, but shouldn't we do that anyway, wherever we are? And how good are we at that?

Do I think AI is a good or bad thing? I wouldn't say either really. It's simply a method of programming a computer, so the morality lies with the programmer.

Oli Carter
Software Engineer at Seenit

Photo by Alice Tonks

2

AI AT SEENIT

How we use AI

At Seenit we use AI, or more specifically deep learning, in tandem with other approaches to ensure better video quality, allow more detailed searches, and offer unique features.

AI is used when more complex tasks need to be performed, for example sentiment analysis, noise detection, and speech to text transcription.

We've been around for almost six years, so we have an enormous amount of data to work with. Our database is unique within our field and we are therefore able to offer tailored approaches to working with video. Self-editing videos, 'liking'

videos, and overall mood detection are only a few examples of tasks we train our systems to perform.

Our goals

- Get the best out of each video and improve quality
- Easily edit a video
- Create examples of final cuts, based on most-liked and best quality videos
- Generate ideas for projects

This will streamline the process of selecting the best clips, improving the audio or video quality and help edit the video, which can be both time consuming and costly.

What have we already achieved and what's next?

Sentiment analysis

We are using AI to detect faces in videos and analyse if their overall impression is happy, sad, surprised, or angry.

Noise detection

We have trained our systems to search for common noises, such as a car engine or a bird singing. This approach is also used to identify if speech is the main focus of a video.

Speech to text transcription

We frequently use this feature for overall mood detection and to create subtitles. Language detection plays a crucial role for these tasks and we train the computer to detect language from speech.

Focus and stability

Unsteady hands create unstable and often out of focus videos. We can provide solutions to both these problems to make sure the video is of the best quality.

Cut video

We can detect any sudden movements at the very beginning or end of a video

and cut out the unusable sections.

Landscape to portrait

We can provide a portrait version of a landscape video that shows the area of interest or main focus of the clip.

GIF creation

Showing the first 5 seconds of each video when you hover over it makes it easier to find the clips you want without having to view each one in full.

Audio normalisation

We can normalise a video's audio so that there are no variations in volume. If the background noise drowns out the speech, we can separate the two and

adjust them to improve the overall sound quality.

Face recognition

We can use facial recognition to identify which person is the main focus in a video.

Self-editing videos

This is one of our main goals and something that we are continuing to work on and improve.

We know that time is precious and that creating a final edit from scratch yourself is not always feasible, so using an existing video that already contains the best quality clips from the shot list will make things much easier.

AI helps make decisions in seconds rather than taking a person hours. It gets used everywhere today. It'll be checking your bank transactions for suspicious activity.

Ian Merrington
VP of Engineering at Seenit

Contact us

info@seenit.io

seenit.io