A Brief Introduction to AI and Its Use in the Legal Practice

*Georges Gielen**

1 Introduction

Beyond any doubt, artificial intelligence (AI) is here to stay! It is taking our world by storm, disruptively affecting all aspects of our society. The field of law and the legal practice are no exception to this. This contribution[1] will give a brief introduction to AI and describe its huge potential and current limitations, followed by an overview of possible applications in the legal practice, such as in smart contracting, the focus of this book.

2 AI: Definition and Basic Techniques

Long before the introduction of ChatGPT, our society has been characterised by a gradual but steady process of digitisation and digitalisation (see the figure below), which is fundamentally transforming our daily way of living and working. This goes from digitally processing all data and documents (instead of using paper) over digitising all processes in companies and organisations (tax declarations, bank transactions, ordering concert or train tickets, etc.) to reshaping entire organisations.

Source: www.channelinsider.com

* Prof. dr. ir. Georges Gielen. ESAT-MICAS, Department of Electrical Engineering (ESAT), KU Leuven, Belgium.

1 *This introductory text by no means pretends to be complete in any way, and – considering the fast evolutions in the field – is likely to be outdated before it is even completed.*

Apart from the socio-cultural problems that this may create (the *digital divide* between different social classes within a population or between younger and older generations or between developed and developing countries), the benefits of this digital transformation are deemed to outweigh by far its drawbacks.

Technologically, this digital transformation is, of course, enabled by the ever-continuing progress in the semiconductor technology used to fabricate electronic chips for computing. This relentless progression over many decades now, in line with the so-called 'Moore's law',[2] made it possible to continuously build ever more powerful computer chips, running ever more powerful software applications in acceptable time and/or at affordable cost.

Originally, computers were developed to compute, i.e. to perform large calculations ('number crunching'). Illustrative examples are calculating the orbit of the Apollo rockets on their journey to the moon, calculating the weather forecasts for the coming days/weeks, calculating the profit/loss account of a company, etc. Over the years, the software applications have diversified and have supported other tasks such as text processing in editors, audio filtering in hearing aids, data streaming in WiFi base stations, movie editing for social media posting, radar-based distance control in cars, to name a few. The relentless progression of hardware and software over many decades has created a wide range of products and applications we use in our daily lives today.

Although extremely powerful at computing tasks, computers, as we knew them, were not very effective at cognitive tasks, such as recognising faces on photos, creating a new poem or a new painting, summarising a discussion or debate, or understanding the nuances in a speech. Humans were believed to be superior at such cognitive tasks, and this was attributed to the way their brain functions – different from classical computers. All this has changed, however, with the recent re-emergence of techniques of AI. Originally coined in 1955 by John McCarthy, the progress in AI techniques over the years has experienced several ups and downs. Only since the mid-2010s did the right ingredients of progress in computing power, progress in artificial neural network algorithms and the growing availability of huge amounts of digital data (through digitisation, the internet, the cloud, the internet of things and the many wireless sensor networks) come together to make practical and efficient AI systems possible. The

2 Moore's law is the prediction by Gordon Moore, co-founder of Intel, in 1965 that the semiconductor technology would develop further in such a way that the number of transistors on a (computer) chip and hence its computing power would double every 1.5 years (later on modified to every 2 years). Repeating this over many decades led to an exponential increase in computing capabilities (for the same price) over time [Moore65].

capabilities seem to be limited only by the availability of proper data and the energy efficiency of the computing devices.

Roughly based on the EC JRC Flagship report on AI[3] and on the publication of Rémy Bonnaffé,[4] the following definition describes the field of AI pretty well:

> *Definition:* ***Artificial Intelligence (AI)*** is a generic term that refers to any machine or algorithm that is capable of observing its environment, learning, analysing and reasoning and, based on the knowledge and experience gained, of taking intelligent action or proposing decisions, to solve a problem at hand, comparable to human intelligence.

The AI Act of the EU[5] uses a slightly different definition:

> *Definition:* An ***AI system*** means a machine-based system designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs, such as predictions, content, recommendations or decisions that can influence physical or virtual environments.

Essential properties of an AI system are as follows:
1. its *autonomy*: it processes the observed inputs entering the system to come to a conclusion by itself. This may initiate proper action automatically in an autonomous system, or in many practical cases the conclusion is presented to the user who has the final say.
2. Its *adaptivity*: it evolves over time by learning from its prior successes and failures. This learning can be static (i.e. the learning is once, in the beginning, based on past available data) or continual over time (based on new additional inputs-outputs being processed).

Given these basic properties of autonomy and learning, the similarity with human intelligence is evident. It is no surprise that AI techniques today are used in almost any application domain. Some examples of tasks addressed by AI systems are the recognition

3 Annoni, A., Benczur, P., Bertoldi, P., Delipetrev, B., De Prato, G., Feijoo, C., Fernandez Macias, E., Gomez Gutierrez, E., Iglesias Portela, M., Junklewitz, H., Lopez Cobo, M., Martens, B., Figueiredo Do Nascimento, S., Nativi, S., Polvora, A., Sanchez Martin, J.I., Tolan, S., Tuomi, I. and Vesnic Alujevic, L., "Artificial Intelligence: A European Perspective," Publications Office of the European Union, Luxemburg, 2018 (DOI:10.2760/936974).

4 Rémy Bonnafé, "New Technologies and the Law: The Impact of Artificial Intelligence on the Practice of Law," https://medium.com/@remybonnaffe/new-technologies-and-the-law-the-impact-of-artificial-intelligence-on-the-practice-of-law-c456904688d1.

5 https://artificialintelligenceact.eu/.

of faces or objects in images, the interpretation and/or automatic translation of speech or texts, the answering of questions by chatbots, the creation of drawings or videos, the medical diagnostics on scan images, the steering of vehicles or drones, the online monitoring of production machines in factories, the screening of contracts, and the making of summary reports of meetings.

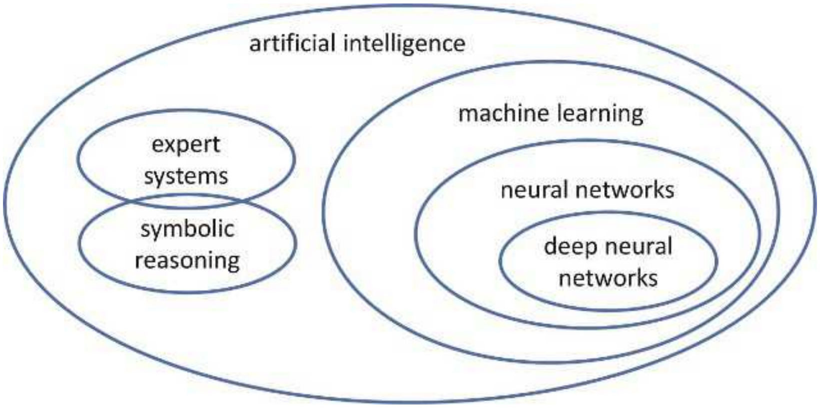

So, besides the progress in computing power, what has made this huge progress possible? The major innovation has come from the introduction of artificial neural networks and, in particular, deep neural networks. Before diving more deeply into this, it is worthwhile to note that – as indicated in the figure below – artificial neural networks are merely a subset of the broader field of machine intelligence, i.e. machines or man-made systems aiming at carrying out 'intelligent' tasks equivalent to humans, and that other techniques, such as expert systems and symbolic reasoning, exist to try and implement artificially intelligent systems. (Deep) neural networks have, however, caused the recent boost.

3 Artificial Neural Networks

Roughly trying to mimic how our human brain works as an interaction of huge numbers of connected neurons, artificial neural networks (see the figure of the standard feed-forward network below) consist of several 'hidden' layers of interconnected 'neurons' (i.e. simple computing units) between the input layer and the output layer. Each connection has a certain strength (possibly zero), indicated by the corresponding weighting coefficient. Each neuron then combines the weighted inputs and generates its output according to its mathematical model (e.g. a thresholding function). The weighting coefficients of the interconnections need to be tuned towards the targeted application for the artificial neural network to adequately perform its desired task (e.g. recognising dogs on images). Deep neural networks are similar in structure but are characterised by a huge number of hidden-neuron layers between the input and output layers. State-of-the-art large-language models (LLMs) such as GPT-4 are based on models, with each more than a few hundred billion parameters that need to be determined during network training.

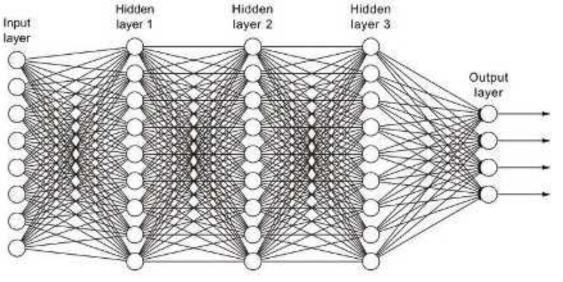

Source: freecontent.manning.com

Note that besides the values of all the weighting coefficients of the interconnections and besides the number of hidden neural layers, many alternative neural network architectures different from the above feed-forward structure have been developed. Essentially, they vary in regard to how the neurons are interconnected, and each of these architectures is better suited for some applications than for others, depending on the task to be executed. A non-exhaustive list of some frequently used alternative neural network structures and illustrative applications is the following:[6]

- convolutional neural networks: these are suited for, among other things, image processing;
- recurrent neural networks: these are suited for, among other things, processing of audio and speech streams;
- transformer systems: these are suited for, among other things, speech/text generation;
- recommender systems: these are suited for, among other things, predicting future preferences of a user based on past choices.

3.1 Training the Neural Network

This brings us to the important concept of *network training*. For artificial neural networks to be successful for a certain task, the network architecture needs to be well chosen and the values of the weights need to be well determined towards the task at hand. This is typically done by training the weights with many example inputs through some sort

6 See for instance https://towardsai.net/p/machine-learning/main-types-of-neural-networks-and-its-applications-tutorial-734480d7ec8e for a more complete overview.