

An Introduction to

LAW & TECHNOLOGY

Edited by
Eva Lievens
Simon Verschaeve
Carl Vander Maelen

CONTENTS

Preface 7

List of Abbreviations 9

PART I. TECHNOLOGY, ECONOMY & SOCIETY 13

Chapter 1. An introduction to information technology for lawyers 15

Yvan Saeys

Chapter 2. Digital infrastructure 43

Carl Vander Maelen

Chapter 3. The rise of platforms in the digital economy 65

Tom Evens

Chapter 4. A societal perspective on technology and new media 81

Ralf De Wolf

PART II. INTERACTION BETWEEN TECHNOLOGY & LAW 97

HUMAN RIGHTS 99

Chapter 5. Freedom of expression and digital technologies 101

Lien Stolle

Chapter 6. The right to privacy and data protection in a digital era 123

Ingrida Milkaite & Liesa Keunen

Chapter 7. Digital discrimination 145

Eva Lievens

Chapter 8. Children's rights in the digital environment 159

Valerie Verdoodt

DATA PROTECTION 179

Chapter 9. Personal data and non-personal data 181

Ingrida Milkaite

Chapter 10. Data protection actors 195

Yueming Zhang & Ingrida Milkaite

Chapter 11. Data protection principles 209

Ingrida Milkaite & Eva Lievens

Chapter 12. Data subject rights 227

Ruben Roex

Chapter 13. Data protection compliance 245

Simon Verschaeve

Chapter 14. Data protection enforcement 267

Simon Verschaeve & Liesa Keunen

DIGITAL ECONOMY	293
Chapter 15. The regulation of electronic commerce and digital services	295
<i>Pieterjan Declerck & Eva Lievens</i>	
Chapter 16. The liability of online intermediaries	315
<i>Lien Stolle</i>	
Chapter 17. Consumer protection in the online environment	329
<i>Pieterjan Declerck & Valerie Verdoodt</i>	
Chapter 18. The regulation of data-related products and services	347
<i>Griet Verhenneman & Simon Verschaeve</i>	
Chapter 19. Electronic contracts and electronic signatures	377
<i>Ruben Roex</i>	
Chapter 20. Cookies and similar tracking technologies	391
<i>Simon Verschaeve</i>	
Chapter 21. Online advertising	405
<i>Simon Verschaeve</i>	
Chapter 22. Cryptocurrencies	437
<i>Carl Vander Maelen</i>	
CYBERCRIME AND CYBERSECURITY	453
Chapter 23. Cybercrime and cybersecurity	455
<i>Eva Lievens</i>	
ARTIFICIAL INTELLIGENCE	477
Chapter 24. The regulation of artificial intelligence	479
<i>Valerie Verdoodt</i>	
PART III. TECHNOLOGY & THE LEGAL PROFESSION	503
Chapter 25. Legal Tech	505
<i>Carl Vander Maelen</i>	
PART IV. REGULATING TECHNOLOGY	517
Chapter 26. Regulating technology: whether, when and how?	519
<i>Carl Vander Maelen & Eva Lievens</i>	
Authors	535

PREFACE

In 2019, the Faculty of Law and Criminology of Ghent University decided to include a new compulsory course 'Law & Technology' in the Bachelor of Laws programme. The rationale was that every student who graduates in this programme needs to have a basic understanding of the impact of technology on society, the interaction between technology and law, and the impact of technology on legal professionals. In the academic year 2022–2023 the course was taught for the first time. This book aims to familiarise readers in general and students in particular with a selection of important topics in the field of law and technology. It aims to identify challenges, map the legal framework and offer triggers to reflect critically on the importance of (digital) technologies in our lives.

Given that (digital) technologies are omnipresent in today's society and that this extends the scope of questions and topics that could fall within the field of law and technology, this book does not aim to be comprehensive. As editors, we selected themes that we appreciate to be important and relevant at this moment in time. This book is regularly revised to keep up with trends and developments in society. Structural changes in this 4th edition include a new chapter on the recent EU data-related regulations, and the integration of a new chapter on 'Digital infrastructure' in the first part. This chapter combines the previous chapters on ICANN and net neutrality in one chapter. The discussion of each topic offers the reader a sketch of the context, and an explanation of the legal framework at the international, regional and national level.

The book is divided into four thematic parts. Part 1 explains important concepts and terms from other disciplines such as computer sciences, communication sciences and sociology that are necessary to be able to identify and assess the legal impact of technology. It also discusses digital infrastructure and the concept of net neutrality. Part 2 reflects critically on the interaction between technology and law, and discusses topics such as human rights, data protection, digital economy, cybercrime, and artificial intelligence. Part 3 discusses legal tech, which facilitates legal processes and is increasingly used by lawyers, judges and other legal professionals. Finally, Part 4 examines the various instruments and strategies that are used to regulate technology.

Technology evolves at breakneck speed, and traditionally the legislator struggles to keep up. In recent times, however, and especially at the level of the European Union, legislative initiatives have been introduced in rapid succession. Several legislative instruments that are still going through the motions of the legislative process, have

just been adopted or are being implemented at national level. The text is up to date until 30 June 2025.

We would like to extend our sincere thanks to all authors. This book is the fruit of a real team effort by the Law & Technology research group of Ghent University and as editors we are proud of what we achieved together. We are grateful as well to Yvan Saeys, Tom Evens and Ralf De Wolf who contributed essential insights from their own disciplines. Finally, we would also like to thank Nancy De Braekeleer, Freya Maenhout, Paulien Vandenberghe and Nina Waegemans at Owl Press for their enthusiasm and patience.

Eva Lievens, Simon Verschaeve and Carl Vander Maelen, 30 June 2025

LIST OF ABBREVIATIONS

5G	5th generation broadband network for cellular phone networks
ADR	Alternative dispute resolution
AI	Artificial intelligence
AI HLEG	High-Level Expert Group on Artificial Intelligence (European Union)
AML/CFT	Anti-money laundering and countering terrorism financing
ANN	Artificial neural network
ARI	Alternative regulatory instrument
AVMS	Audiovisual media services
AVMSD	Audiovisual Media Services Directive
BCEL	Belgian Code on Economic Law
BEREC	Body of European Regulators for Electronic Communications
BEUC	Bureau Européen des Unions de Consommateurs
BIPT	Belgisch Instituut voor Postdiensten en Telecommunicatie
CAHAI	Ad Hoc Committee on Artificial Intelligence (Council of Europe)
CASP	Crypto-asset service provider
CBDC	Central bank digital currency
CCTV	Closed-circuit television
CEDAW	Convention on the Elimination of All Forms of Discrimination against Women
CEO	Chief executive officer
CFREU	Charter on the Fundamental Rights of the European Union
CJEU	Court of Justice of the European Union
CoE	Council of Europe
CRC	Committee on the Rights of the Child (United Nations)
CRD	Consumer Rights Directive
CSAM	Child sexual abuse material
CSR	Corporate social responsibility
CSIRT	Computer Security Incident Response Team
DA	Data Act
DCD	Digital Content Directive
DGA	Data Governance Act
DLT	Distributed ledger technology

DMA	Digital Markets Act
DNS	Domain name system
DPA	Data Protection Authority
DPbDD	Data protection by design and by default
DPD	Data Protection Directive
DPIA	Data protection impact assessment
DPO	Data protection officer
DSA	Digital Services Act
DSG	Directive on the Sale of Goods
DSR	Data subject request
ECB	European Central Bank
ECD	E-Commerce Directive
ECHR	European Convention on Human Rights and Fundamental Freedoms
ECtHR	European Court of Human Rights
EDIB	European Data Innovation Board
EDPB	European Data Protection Board
EEA	European Economic Area
EECC	European Electronic Communications Code
EHDS	European health data space
EHDSR	European Health Data Space Regulation
EHR	Electronic health record
ESC	European Social Charter
EU	European Union
FCC	Federal Communications Commission (United States of America)
FIRM	Federal Institute for Human Rights (Belgium)
FRA	Fundamental Rights Agency (European Union)
FSMA	Financial Services and Markets Authority (Belgium)
FTC	Federal Trade Commission (United States of America)
GAN	Generative adversarial networks
GBA	Gegevensbeschermingsautoriteit [Data Protection Authority] (Belgium)
GC	General Comment (by United Nations institutions)
GDPR	General Data Protection Regulation
GPAI	General-Purpose AI

GPT-3	Generative pre-trained transformer 3
HTML	HyperText Markup Language
HTTP	HyperText Transfer Protocol
IaaS	Infrastructure as a Service
IANA	Internet Assigned Numbers Authority
ICANN	Internet Corporation for Assigned Names and Numbers
ICCPR	International Covenant on Civil and Political Rights
ICESCR	International Covenant on Economic, Social, and Cultural Rights
ICT	Information and communications technology
IEEE	Institute of Electrical and Electronics Engineers
iOS	iPhone Operating System
IoT	Internet of Things
IP	Internet Protocol
IP address	Internet Protocol address
ISP	Internet service provider
ISS	Information society service
LED	Law Enforcement Directive
LPWAN	Low-Power Wide-Area Networking
LSA	Lead Supervisory Authority
MiCA	(Regulation on) Markets in Crypto-Assets
ML	Machine learning
MLP	Multi-layer perceptron models
NBCC	New Belgian Civil Code
NGO	Non-governmental organisation
NLP	Natural language processing
NSA	National Security Agency
ODR	Online dispute resolution
OECD	Organisation for Economic Co-operation and Development
OSS	One-stop shop mechanism
OTT	Over-the-top (services)
PaaS	Platform as a Service
PID	Price Indication Directive
PSP	Payment services providers

RFID	Radio-Frequency Identification
SA	Supervisory Authority
SaaS	Software as a Service
SAC	Supervisory Authority Concerned
SDK	Software Development Kit
SEA	Search engine advertising
SME	Small and medium-sized enterprises
TCP/IP	Transmission Control Protocol/Internet Protocol
TEU	Treaty on European Union
TFEU	Treaty on the Functioning of the European Union
TLD	Top-level domain
UCPD	Unfair Commercial Practices Directive
UCTD	Unfair Contract Terms Directive
UDHR	Universal Declaration of Human Rights
UDRP	Uniform Dispute Resolution Policy
UNCRC	United Nations Convention on the Rights of the Child
UNESCO	United Nations Educational, Scientific and Cultural Organization
URL	Uniform Resource Locator
UTXO	Unspent Transaction Output
UX	User experience
VAE	Variational AutoEncoders
VLOP	Very large online platform
VLOSE	Very large online search engine
VPN	Virtual private network
VSP	Video sharing platform
WEF	World Economic Forum
WP29	Article 29 Working Party
WWW	World Wide Web

PART I

TECHNOLOGY, ECONOMY & SOCIETY



1. INTRODUCTION

Technological advances during the last decades, and in particular the increased digitisation of information, have strongly shaped and influenced our society. The availability of large volumes of data, together with massive increases in computing power and storage capacity have led to a wealth of information that is now available at our fingertips. At the same time, advances in network technology have radically reshaped the way we communicate, with mobile phones now an integral part of life, making it easier to share and consume information. Novel computational methods such as artificial intelligence (AI) techniques¹ build further on these advances, making it possible to turn these large volumes of data into descriptive and predictive models that can be used to automate decision-making. Together, these technological advances offer a number of opportunities, but also challenges for keeping legal systems in check with the fast-moving technological world. In this chapter, we clarify some of the basics of these recent advances in information technology.

HARDWARE VERSUS SOFTWARE

Computing systems are built of two large components: hardware and software. Hardware refers to all the physical parts that make up computers and networks. This includes the different computer parts such as hard drives, memory, keyboard, mouse and screen, but also the different parts needed to ensure network connectivity including, *inter alia*, cables, routers and modems. Software on the other hand refers to all the instructions and code that are needed to make the hardware work. Software is typically organised into several layers that build on each other. The lowest layer that directly connects to the hardware is the operating system that makes computers or mobile phones work. On top of the operating system run several types of application software, providing the interface to the end user, such as email clients, text processing software or web browsers.

¹ S Russell, P Norvig, *Artificial Intelligence: A Modern Approach* (Pearson Education Limited 2021).

2. INTERNET TECHNOLOGY



The Internet, which represents an international 'network of networks', connects computers and other devices that share a common set of protocols (Transmission Control Protocol/Internet Protocol or TCP/IP for short) for communicating among each other. The Internet facilitates the sharing of information through various protocols, including hyperlinked documents (world wide web), email, telephony and file sharing. Since the 1980s, access to internet technology has become more broadly available, first at universities and academic institutes, then later for the general public by way of internet service providers (ISPs). The Internet has dramatically reshaped the way we communicate and, together with advances in mobile network technology, has now become the major method of communication; this includes social networks, television, video conferencing, online shopping and much more. The omnipresence of internet technology has resulted in new legal challenges, in particular related to cybercrime, fraud, harassment on social media and dealing with private and sensitive information.

2.1 THE WORLD WIDE WEB

Pioneered in 1989 by Tim Berners-Lee, the world wide web (WWW) was designed as a collection of documents, linked together by the HyperText Transfer Protocol (HTTP).² Every document is characterised by its uniform resource locator (URL), which defines a unique address, and is typically accessed through a web browser. Once a user types in the URL, the web browser translates the URL into an internet protocol address (IP address) using the domain name system (DNS), a hierarchically structured and decentralised naming system that identifies every computer connected to the Internet (see Chapter 2 on 'Digital infrastructure' for more information).

2.2 THE INTERNET OF THINGS (IOT)

Building further on internet technology, many devices can now be connected to each other, resulting in the Internet of Things (IoT).³ Each of these devices could, for example, be equipped with sensors, its own processing power and software running either locally on the device, on a central server or in the cloud. IoT applications are



Icon credit: Flaticon.com

2 T Berners-Lee, 'The Original HTTP as defined in 1991' (1 January 1991) <www.w3.org>.
3 M Weiser, R Gold, J S Brown, 'The origins of ubiquitous computing research at PARC in the late 1980s' (1999) 38(4) *IBM Systems Journal* 693.

plentiful, including many control and automation applications such as smart homes and smart healthcare systems, but also applications in agriculture, energy management and environmental monitoring. As an example, consider the application of IoT in elderly care homes.⁴ Sensors could be embedded in bathrooms and toilets, detecting the location of inhabitants, and measuring weight, pulse and blood pressure. Similarly, sensors in wheelchairs and walkers could collect data which, when combined with artificial intelligence methods, detect unusual situations or alarming movement patterns, alerting the user to problems.

At the technical level, IoT solutions can be distinguished at two levels: the communication protocols that are used to communicate between the devices and the servers, and the way the data processing is implemented. The communication protocol is often determined by the physical distance between the devices with, for example, Bluetooth and Radio-frequency identification (RFID) frequently being used for short-range wireless communication, and 5G (fifth-generation broadband network for cellular phone networks) and low-power wide-area networking (LPWAN) for medium and long-range distance communication. The data processing can be implemented in several ways, anywhere in between the following two extremes. In the case of edge computing, the largest part of the computation is done on the device itself. This requires the device to have sufficient computing capabilities and is required if, for instance, fast decisions need to be made for which transfer of the data to a remote server is not possible, such as in the camera of a self-driving car. In the case of cloud computing, no computation is carried out on the device itself; instead the data is sent to a remote server or cloud solution for data processing.

2.3 BIG DATA AND CLOUD COMPUTING

Due to the digitisation of society, a lot of information can now be collected and stored in digital form, making it easier to use for data-mining purposes. Based on advances in internet technology, a lot of data regarding user behaviour such as surfing, online shopping, music and movie preferences can be collected, and the data subsequently mined for interesting patterns. These patterns can be exploited for many different purposes such as marketing (e.g. defining market segments) and optimising internet traffic, but also surveillance monitoring and internet security. Smart devices such as cameras and sensors are also continuously collecting data, resulting in large data

streams that require specific algorithms. These novel types of data result in novel machine-learning models that are now able to automate certain tasks, such as facial recognition or autonomously driving a car.

ALGORITHMS

The essential building block of all software, an algorithm is defined as a finite sequence of instructions that describes how to realise a specific goal or computation, just as a recipe is used in cooking. A computer program is a specific implementation of such an algorithm, taking into account certain constraints such as specific programming language or hardware of the computer where the program runs.

Example of the Bubble sort algorithm to sort a list of numbers

- 1. Compare the first two elements in the list:
If the first is greater than the second, swap them.
- 2. Repeat this with every pair of adjacent elements in the list.
- 3. Then, repeat from step 1 until the list is fully sorted.

Example implementation of the Bubble sort algorithm in Python

```
def bubbleSort(list2sort):
    n = len(list2sort)

    # Go through all list elements for i in range(n):
    # Last i elements are already in place for j
    in range(0, n-i-1):

        # traverse the list from 0 to n-i-1
        # Swap if the element is found greater
        # than the next element
        if list2sort[j] > list2sort[j+1]:
            list2sort[j], list2sort[j+1] = list-
            2sort[j+1], list2sort[j]


    # Driver code to test above
    list2sort = [25,22,11,67,91]

    bubbleSort(list2sort)
    print("Sorted list is:")

    for i in range(len(list2sort)):
        print("%d" % list2sort[i], end=" ")
```

4 P A Laplante, M Kassab, N L Laplante, J M Voas, 'Building Caring Healthcare Systems in the Internet of Things' (2018) 12(3) *IEEE Systems Journal*.

All these digitally generated data result in such large amounts of data that they are simply too much to be handled and stored on a single machine. Therefore, large-scale distributed computing resources, referred to as cloud computing,⁵ have become available to offer customers the possibility of outsourcing storage and computing facilities to providers of such resources. Alongside providing large-scale storage and computing resources, these cloud solutions also remove some of the risks of small-scale computing systems, such as vulnerability to hardware crashes, the complexity of which is shielded from the end user and dealt with by the cloud service providers. By offering infrastructure, platforms and even software as a service (paradigms referred to as IaaS, PaaS and SaaS),⁶ users can now rent these services for as long as they need them, without having to worry about their setup and maintenance. As these services are remotely hosted, they are ubiquitous, available to anyone anywhere and offer flexible payment options: you use and pay as much as required. Companies that work with large amounts of data make use of these cloud solutions as they are extremely well suited to big data mining.



The term big data refers to data-mining problems the volume, diversity and complexity of which requires the development of novel techniques, algorithms and analyses to extract valuable knowledge that is typically hidden inside these data.⁷ Big data mining techniques are often characterised by the so-called 5V definition: volume, velocity, variety, veracity and value.⁸

Volume refers to the vast amounts of data that these systems are expected to process. Think for example of the number per minute of tweets that are generated, the images a traffic camera records or the number of status updates on social media. For certain applications, it is important that these data are constantly processed to support downstream decision-making. As the amounts of data become so large, classical databases cannot adequately handle them and novel, distributed databases have had to be developed, often making use of cloud solutions to deal with them.⁹

The second aspect, velocity, deals with the related aspect of the speed at which the data is generated. This often has consequences for further processing, as certain decisions might be made too late and opportunities may be missed. In some cases, data is arriving at such high speeds that it cannot be stored for further processing, and instant decisions need to be made about whether or not it is useful. Data streams¹⁰ refer to a constant incoming series of data from, for example, sensor networks, weather stations or computer traffic. In the area of machine learning, novel stream-mining algorithms have been developed or adapted to deal with such scenarios.

Another aspect, unrelated to volume and velocity, is the variety of the data. This aspect refers to the many types of data that can be collected, and potentially integrated in order to make further decisions. Two main types of data are typically distinguished: structured and unstructured data.¹¹ Structured types of data refer to data that are typically stored in tables or relational databases. Think for example of data stored in a spreadsheet, such as expenses or travel distance. Unstructured types of data on the other hand refer to data that you cannot process so easily, such as images, video, sound recordings or social media updates (text). These types of data already require specific

5 L Qian, Z Luo, Y Du, L Guo, 'Cloud Computing: An Overview' in M G Jaatun, G Zhao, C Rong (eds), *Cloud Computing. CloudCom 2009. Lecture Notes in Computer Science* (Springer 2009).

6 Y Duan, G Fu, N Zhou, X Sun, N C Narendra, B Hu, 'Everything as a Service (XaaS) on the Cloud: Origins, Current and Future Trends' in *IEEE 8th International Conference on Cloud Computing* (IEEE 2015) 621.

7 A Rajaraman, J Leskovec, J D Ullman, *Mining of Massive Datasets* (Cambridge University Press 2014).

8 A Jain, 'The 5 V's of big data', IBM Watson Health Perspectives (17 September 2016) <<https://www.ibm.com/blogs/watson-health/the-5-vs-of-big-data/>>.

9 K Grolinger, W A Higashino, A Tiwari, M A M Capretz, 'Data management in cloud environments: NoSQL and NewSQL data stores' (2013) 2(1) *Journal of Cloud Computing: Advances, Systems and Applications* 1.

10 M Garofalakis, J Gehrke, R Rastogi, *Data Stream Management* (Springer 2009).

11 C Taylor, 'Structured vs. Unstructured Data' (2021) <www.datamation.com>.

expertise to be processed in an automated fashion, such as the natural language processing (NLP) systems that are discussed later in this chapter. Dealing with all these data types is challenging but combining them might be crucial for making important decisions. Imagine the case of patients in a hospital. A lot of clinical information about the patient (e.g. age, sex, blood pressure, weight etc.) might be present as structured data, and stored in a database. Likewise, there might be other types of data, such as radiographic images, the patient's DNA sequence, doctor's notes about the patient or even the patient's social media posts which would all need to be integrated in order to gain a better understanding of the patient's health status.

Another important aspect of big data concerns its veracity: this refers to the fact that data quality might vary, and the accuracy of the information is not always easy to control. Think for example about posts on social media: some information may be ambiguous, missing or sometimes even completely wrong. All this needs to be considered when designing data-mining systems that automatically process such data to yield novel insights.

Finally, all these aspects serve one big goal: to extract value from the data. This could result in the identification of novel patterns, better statistics or machine-learning models, or new opportunities for research and industry. In all these settings it is important that users make a business case for any attempt to collect and leverage big data in order to trade off the costs and benefits of such systems.

The applications of big data science and cloud computing are plentiful. In science and research more and more data are collected; measuring devices have been miniaturised during the last decades, showing data about our world and everything in it in ever-greater detail. Examples of scientific domains where vast amounts of data are collected and analysed include physics and astronomy; information technology; biotechnology and bioinformatics; healthcare; and public health, as well as some more recent fields that have started to digitally record information such as history. In business and commerce big data techniques have already been quickly adopted and are used for example to predict when a customer is going to change their phone operator or bank, or when a customer is expecting a baby. Recently a lot of applications are also moving in the direction of economy and sustainability, with examples such as the optimisation of electricity grids (smart grids), smart meters and smart homes and the optimisation of traffic flows based on traffic information, social media and weather data. Also, many sports and social activities are increasingly making use of big data technologies, for example for sports analytics or online dating.

Finally, also in security and law enforcement, the opportunities of big data science are slowly being explored, such as the design of systems that assist in automated browsing and analysis of court decisions and patent filings.

3. ARTIFICIAL INTELLIGENCE

Artificial intelligence (AI) is a wide area of research that focuses on building systems that are able to solve problems that require intelligence, such as learning to perform specific tasks (e.g. self-driving cars) or interacting with humans (e.g. conversational agents). A crucial aspect of building AI systems is the need to find adequate representations for knowledge and data that allow the systems to learn and derive new knowledge, patterns or insights (Figure 1). By implementing these representations and learning mechanisms in a hardware or software agent, an autonomous system can be built that subsequently is able to be deployed in the real world. Common examples of AI systems include planning agents such as route planners, recommender systems, conversational agents (e.g. chatbots, speech recognition, automatic translation), image recognition software (e.g. face recognition), self-driving cars and many more.




Figure 1 An overview of the different subdisciplines of the broad field of artificial intelligence (AI)




Figure 2 An overview of the most important milestones in the history of artificial intelligence

3.1 BUILDING BLOCKS OF AI

In order to build AI systems, a number of basic building blocks need to be combined and depending on the task one intends to automate, some or all of them might be needed. A first set of building blocks must reflect the fact that our agent needs to be able *to interact with its environment*. Imagine a self-driving car as an example. This robotic system needs to constantly monitor the road ahead (input), process the incoming information and subsequently act upon it, e.g. by steering in the right direction if there is a turn, or braking if a pedestrian is crossing the road (output). Depending on the context, different systems might be needed to record the input, such as various cameras and sensors for processing visual inputs, or textual or auditory inputs in the case of conversational agents. Specific components such as computer vision or speech and language processing algorithms will be needed to implement these functionalities. This information needs subsequently to be processed and acted

upon. In the case of robotic systems, various actuators might control the response of the robot (e.g. steering the self-driving car in a particular direction, or braking), while in conversational agents a certain answer might be produced by the system.

A second set of building blocks will offer the AI system the capacity to *reason and plan in the real world*. As humans, we make use of a lot of implicit knowledge and 'common sense' when making decisions. All this knowledge needs to be provided or programmed into AI systems, thus requiring the design of knowledge representations that allow AI systems to take into account a model of the world, including the required background knowledge. This will allow the AI model to make plans, e.g. planning an optimal route from point A to point B. An important challenge for AI systems in this regard is the difficulty of dealing with unforeseen circumstances or uncertainty. As humans, we are well aware of such real-world properties, but for computer-based systems, they pose an enormous challenge.

3.2 MACHINE LEARNING: LEARNING CONCEPTS FROM DATA

While AI is a very broad research domain, involving research topics such as logic, search problems, uncertainty modeling and robotics, we will here focus on one of the most recent and challenging subfields in AI: machine-learning (ML) models.¹² ML models radically change the classical concept of how computers are programmed. While in classical programs the programmer must specify in an explicit fashion all the steps an algorithm needs to take, machine-learning models allow for 'programming' of a computer by letting it learn from examples. Consider the problem of recognising human faces. If one had to design a computer program to detect any human face, it would be quite a challenge to program such a method by hand and define all the explicit steps it would use. However, with machine-learning models, one can present the computer with a large set of examples of faces, and from these examples the model will automatically learn how to recognise faces: so rather than explicitly programming a facial recognition system, such a model would be learned automatically. Machine-learning models are thus a set of computational techniques that allow models to automatically learn from data, without the explicit need to be programmed by hand. Several types of machine-learning models can be distinguished.

12 C M Bishop, *Pattern Recognition and Machine Learning* (Springer 2006).

3.2.1 DIFFERENT TYPES OF MACHINE-LEARNING METHODS


The most well-known category of ML methods are *predictive models*. Such models aim to predict for a certain input (e.g an image) or a specific output (e.g. whether a face is present in the image). Examples of such predictive models include predicting diseases from medical images, predicting whether customers are likely to change their phone operator or just simply predicting tomorrow's weather. To build such models, we must provide them with labelled examples, such as a set of healthy tissues and a set of cancerous tissues if our objective is to build a model that detects skin cancer. From these labelled examples, often referred to as the training set, our ML model could learn to make the distinction between healthy and cancerous tissue. That is why this type of modeling is often referred to as *supervised learning*, since it requires the human expert to label every image as either healthy or cancerous. Depending on the task, this labelling could happen quickly, or be more laborious for the human expert involved. Constructing the model using a set of labelled examples is typically referred to as 'training' the model. In the training phase, the model learns to adjust its parameters in order to perform the prediction task in the best possible way. Later on, the model is typically evaluated on a previously extracted set of examples, often referred to as the test set. In all ML models, the goal is to build models that generalise well beyond the training set and show good performance on the test set. In the case models are only performing well on the training set, but not on the test set, the model is suffering from overfitting: it has focused too much on the examples in the training set and is not able to generalise beyond it. Preventing overfitting is one of the main challenges in ML models in general.¹³

A second category of ML models are so-called unsupervised learning models. Contrary to supervised learning models, these do not require any manual labelling of the examples, but instead try to learn the structure in the data by themselves. That is why these techniques are often also termed *structure learning* approaches. The most well-known type of unsupervised learning models are clustering techniques.¹⁴ Clustering techniques aim to find group structures in the data, such as people with a similar musical taste, or people who share many connections in a social network and thus could belong to the same group of friends.

A last category of ML methods is reinforcement learning models.¹⁵ These models aim to learn a control policy, which corresponds to a series of actions that need to be performed in an optimal sequence to fulfill a certain task. Think of a robot that needs to learn how to walk. This robot is full of motors and to let the robot walk, the behaviour of all these motors needs to be tightly controlled and coordinated. By letting the robot explore a certain series of actions, it will get positive rewards (e.g. if the robot moves forward) or negative rewards (e.g. if the robot falls over), and by learning from these rewards, over time, the robot can improve its behaviour.

THREE MAJOR TYPES OF MACHINE-LEARNING MODELS

SUPERVISED LEARNING (LEARNING PREDICTIVE MODELS)



13 D M Hawkins, 'The problem of overfitting' (2004) 44(1) *Journal of Chemical Information and Computer Sciences* 1.
14 C C Aggarwal, C K Reddy, *Data Clustering: Algorithms and Applications* (Routledge 2013).

15 R S Sutton, A G Barto, *Reinforcement Learning: An Introduction* (MIT Press 2018).