

# Machine Learning with Python


## For PC, Raspberry Pi, and Maixduino



```

t sensor
t image
t lcd
rt KPU as kpu

init()
or.reset()
or.set_pixformat(sensor.RGB565)
or.set_framesize(sensor.QVGA)
or.run(1)
k = kpu.load(0x300000)
hor = (1.889, 2.5245, 2.9465, 3.94056, 3.99987, 5.3658, 5.155437)
kpu.init_yolo2(task, 0.5, 0.3, 5, anchor)
le(True):
img = sensor.snapshot()
code = kpu.run_yolo2(task, img)
if code:
for i in code:
print(i)
a = img.draw_rectangle(i.rect())
a = lcd.display(img)
= kpu.deinit(task)
    
```



```

"w":135, "h":181, "value":0.949942, "classid":0, "index":0, "objnum":1}
"w":135, "h":181, "value":0.938012, "classid":0, "index":0, "objnum":1}
"w":135, "h":181, "value":0.938012, "classid":0, "index":0, "objnum":1}
"w":135, "h":181, "value":0.938012, "classid":0, "index":0, "objnum":1}
"w":135, "h":181, "value":0.938012, "classid":0, "index":0, "objnum":1}
"w":135, "h":181, "value":0.923467, "classid":0, "index":0, "objnum":1}
"w":135, "h":181, "value":0.923467, "classid":0, "index":0, "objnum":1}
"w":135, "h":181, "value":0.938012, "classid":0, "index":0, "objnum":1}
"w":135, "h":181, "value":0.938012, "classid":0, "index":0, "objnum":1}
"w":135, "h":181, "value":0.938012, "classid":0, "index":0, "objnum":1}
"w":135, "h":181, "value":0.923467, "classid":0, "index":0, "objnum":1}
    
```



Günter Spinner



---

# Machine Learning with Python

For PC, Raspberry Pi, and MaixDuino



Dr. Gunter Spanner

---

● This is an Elektor Publication. Elektor is the media brand of Elektor International Media B.V.  
PO Box 11, NL-6114-ZG Susteren, The Netherlands  
Phone: +31 46 4389444

● All rights reserved. No part of this book may be reproduced in any material form, including photocopying, or storing in any medium by electronic means and whether or not transiently or incidentally to some other use of this publication, without the written permission of the copyright holder except in accordance with the provisions of the Copyright Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licencing Agency Ltd., 90 Tottenham Court Road, London, England W1P 9HE. Applications for the copyright holder's permission to reproduce any part of the publication should be addressed to the publishers.

● **Declaration**

The Author and Publisher have used their best efforts in ensuring the correctness of the information contained in this book. They do not assume, and hereby disclaim, any liability to any party for any loss or damage caused by errors or omissions in this book, whether such errors or omissions result from negligence, accident, or any other cause.

All the programs given in the book are Copyright of the Author and Elektor International Media. These programs may only be used for educational purposes. Written permission from the Author or Elektor must be obtained before any of these programs can be used for commercial purposes.

● British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

● **ISBN 978-3-89576-502-5** Print

**ISBN 978-3-89576-503-2** eBook

● © Copyright 2022: Elektor International Media B.V.

Editor: Jan Buiting

Prepress Production: D-Vision, Julian van den Berg

Elektor is part of EIM, the world's leading source of essential technical information and electronics products for pro engineers, electronics designers, and the companies seeking to engage them. Each day, our international team develops and delivers high-quality content - via a variety of media channels (including magazines, video, digital media, and social media) in several languages - relating to electronics design and DIY electronics. [www.elektormagazine.com](http://www.elektormagazine.com)

---

<b>Cautionary Notices</b> . . . . .	<b>10</b>
<b>Program Downloads</b> . . . . .	<b>11</b>
<b>Chapter 1 • Introduction</b> . . . . .	<b>12</b>
1.1 "Super Intelligence" in three steps? . . . . .	13
1.2 How machines can learn. . . . .	14
<b>Chapter 2 • A Brief History of ML and AI</b> . . . . .	<b>16</b>
<b>Chapter 3 • Learning from "Big Data"</b> . . . . .	<b>19</b>
3.1 Machine Learning and Artificial Intelligence. . . . .	19
<b>Chapter 4 • The Hardware Base</b> . . . . .	<b>23</b>
<b>Chapter 5 • The PC as Universal AI Machine</b> . . . . .	<b>24</b>
5.1 The computer as a programming center. . . . .	24
<b>Chapter 6 • The Raspberry Pi</b> . . . . .	<b>27</b>
6.1 The Remote Desktop . . . . .	28
6.2 Using smartphones and tablets as displays. . . . .	30
6.3 FileZilla . . . . .	30
6.4 Pimp my Pi . . . . .	31
<b>Chapter 7 • Sipeed Maix, aka "MaixDuino"</b> . . . . .	<b>34</b>
7.1 Small but mighty: the performance figures of the MaixDuino . . . . .	34
7.2 A wealth of applications . . . . .	37
7.3 Initial start-up and functional test . . . . .	38
7.4 Power supply and stand-alone operation . . . . .	40
<b>Chapter 8 • Programming and Development Environments</b> . . . . .	<b>41</b>
8.1 Thonny — a Python IDE for beginners and intermediates . . . . .	41
8.2 Thonny as a universal IDE for RPi and MaixDuino. . . . .	44
8.3 Working with files . . . . .	45
8.4 Thonny on the Raspberry Pi . . . . .	48
8.5 Tips for troubleshooting the Thonny IDE. . . . .	50
8.6 The MaixPy IDE. . . . .	53
8.7 A MicroPython interpreter for MaixDuino. . . . .	54
8.8 The Flash tool in action . . . . .	56
8.9 Machine Learning and interactive Python . . . . .	57
8.10 Anaconda . . . . .	58

8.11 Jupyter . . . . .	60
8.12 Installation and Start-Up . . . . .	61
8.13 Using MicroPython Kernels in Jupyter . . . . .	64
8.14 Communication setup to the MaixDuino . . . . .	65
8.15 Kernels . . . . .	66
8.16 Working with Notebooks . . . . .	66
8.17 All libraries available?. . . . .	67
8.18 Using Spyder for Python Programming . . . . .	68
8.19 Who's programming who?. . . . .	70
<b>Chapter 9 • Python in a Nutshell . . . . .</b>	<b>72</b>
9.1 Comments make your life easier . . . . .	74
9.2 The print() statement. . . . .	76
9.3 Output to the display . . . . .	76
9.4 Indentations and Blocks . . . . .	78
9.5 Time Control and Sleep . . . . .	78
9.6 Hardware under control: digital inputs and outputs . . . . .	80
9.7 For vital values: variables and constants . . . . .	82
9.8 Numbers and variable types . . . . .	83
9.9 Converting number types . . . . .	84
9.10 Arrays as a basis for neural networks. . . . .	85
9.11 Operators . . . . .	86
9.12 Conditions, branches and loops . . . . .	87
9.13 Trial and error — try and except . . . . .	89
<b>Chapter 10 • Useful Assistants: Libraries! . . . . .</b>	<b>91</b>
10.1 Matplotlib as a graphics artist. . . . .	92
10.2 The math genius: Numpy . . . . .	96
10.3 Data-mining using Pandas . . . . .	100
10.4 Learning and visualization using Scikit, Scipy, SkImage & Co. . . . .	102
10.5 Machine Vision using OpenCV . . . . .	103
10.6 Brainiacs: KERAS and TensorFlow . . . . .	109
10.7 Knowledge transfer: sharing the learning achievements . . . . .	113
10.8 Graphical representation of network structures . . . . .	114

---

10.9 Solution of the XOR problem using KERAS . . . . .	115
10.10 Virtual environments . . . . .	116
<b>Chapter 11 • Practical Machine Learning Applications . . . . .</b>	<b>119</b>
11.1 Transfer functions and multilayer networks . . . . .	119
11.2 Flowers and data . . . . .	120
11.3 Graphical representations of data sets . . . . .	122
11.4 A net for iris flowers. . . . .	124
11.5 Training and testing . . . . .	126
11.6 What's blossoming here? . . . . .	128
11.7 Test and learning behavior . . . . .	130
<b>Chapter 12 • Recognition of Handwritten Numbers . . . . .</b>	<b>133</b>
12.1 "Hello ML" — the MNIST data set . . . . .	134
12.2 A neural network reads digits . . . . .	136
12.3 Training, tests and predictions. . . . .	137
12.4 Live recognition of digits. . . . .	139
12.5 KERAS can do even better! . . . . .	140
12.6 Convolutional networks . . . . .	142
12.7 Power training. . . . .	146
12.8 Quality control — an absolute must! . . . . .	146
12.9 Recognizing live images . . . . .	147
12.10 Batch sizes and epochs. . . . .	150
12.11 MaixDuino also reads digits . . . . .	151
<b>Chapter 13 • How Machines Learn to See: Object Recognition . . . . .</b>	<b>154</b>
13.1 TensorFlow for Raspberry Pi . . . . .	154
13.2 Virtual environments in action. . . . .	155
13.3 Using a Universal TFlite Model. . . . .	157
13.4 Ideal for sloths: clothes-sorting . . . . .	160
13.5 Construction and training of the model. . . . .	163
13.6 MaixDuino recognizes 20 objects . . . . .	166
13.7 Recognizing, counting and sorting objects. . . . .	169
<b>Chapter 14 • Machines Learn to Listen and Speak. . . . .</b>	<b>172</b>

14.1 Talk to me! . . . . .	172
14.2 RPi Learns to talk. . . . .	174
14.3 Talking instruments . . . . .	176
14.4 Sorry, didn't get you ... . . . .	179
14.5 RPi as a ChatBot . . . . .	183
14.6 From ELIZA to ChatterBots . . . . .	186
14.7 The Talking Eye . . . . .	188
14.8 An AI Bat . . . . .	189
<b>Chapter 15 • Facial Recognition and Identification . . . . .</b>	<b>192</b>
15.1 The right to your own image . . . . .	193
15.2 Machines recognize people and faces . . . . .	194
15.3 MaixDuino as a Door Viewer . . . . .	198
15.4 How many guest were at the party?. . . . .	200
15.5 Person-detection alarm. . . . .	202
15.6 Social minefields? — face identification. . . . .	203
15.7 Big Brother RPi: face identification in practice . . . . .	204
15.8 Smile, please ;-) . . . . .	206
15.9 Photo Training . . . . .	213
15.10 "Know thyself!" ... and others . . . . .	215
15.11 A Biometric scanner as a door opener . . . . .	216
15.12 Recognizing gender and age . . . . .	217
<b>Chapter 16 • Train Your Own Models . . . . .</b>	<b>220</b>
16.1 Creation of a model for the MaixDuino . . . . .	220
16.2 Electronic parts recognition with the MaixDuino. . . . .	222
16.3 Performance of the trained network . . . . .	225
16.4 Field test . . . . .	226
16.5 Outlook: Multi-object detectors . . . . .	228
<b>Chapter 17 • Dreams of the Future: from KPU to Neuromorphic Chips . . . . .</b>	<b>230</b>
<b>Chapter 18 • Electronic Components . . . . .</b>	<b>236</b>
18.1 Breadboards . . . . .	236
18.2 Wires and jumpers. . . . .	238



---

18.3 Resistors . . . . .	238
18.4 Light-emitting diodes (LEDs) . . . . .	239
18.5 Transistors . . . . .	240
18.6 Sensors . . . . .	241
18.7 Ultrasound range finder . . . . .	242
<b>Chapter 19 • Troubleshooting . . . . .</b>	<b>243</b>
<b>Chapter 20 • Buyers Guide. . . . .</b>	<b>244</b>
<b>Chapter 21 • References; Bibliography . . . . .</b>	<b>245</b>
<b>Index . . . . .</b>	<b>246</b>

## Cautionary Notices

1. The circuits and boards (Raspberry Pi, MaixDuino) in this book may only be operated with approved, double-insulated AC power supplies. Insulation faults in a simple power supply unit can result in life-threatening voltages on non-insulated components.
2. Powerful LEDs can cause eye damage. Never look directly into an LED!
3. Neither the Author nor the Publisher assume any liability for any damage, loss, or injury caused by, or as a result of, setting up the projects described.
4. Electronic circuits can emit electromagnetic interference. Neither the Publisher nor the Author have any control over the technical implementations of the user. Consequently, the user is responsible for compliance with the relevant emission limit values.

## Program Downloads

The programs discussed in this book can be downloaded free of charge from the Elektor Store website:

[www.elektor.com](http://www.elektor.com)

Once at the Elektor Store, enter the book title in the Search box. Once at the book resources and info page, click on the **Downloads** tab. Download the .zip archive file, save it locally on your computer, then extract it.

If a program is not identical to the one described in the book, the version from the download package should be used, as this is the most current one.

## Chapter 1 • Introduction

We are increasingly confronted with the applications of "Artificial Intelligence" (AI): recommendations for music or video, navigation systems, shopping suggestions and so on are based on procedures that can be more or less assigned to this area. However, there is often a lack of clarity regarding the terms

- Artificial Intelligence (AI)
- Machine Learning (ML)
- Deep Learning (DL)
- Neural Networks (NN)

One of the most frequent queries to popular search engines is therefore:

"What is the actual difference between AI and Machine Learning?"

Although the terms are closely related, they are by no means synonymous:

**Artificial Intelligence** is, in principle, the general knowledge base. It forms the *foundation*, just as mathematics is the basis for physics or electrical engineering. AI tries to solve problems without rigid algorithms using highly developed machines.

**Machine Learning** is a branch of artificial intelligence. The focus is on automatic learning from experience, i.e., without a special program being created for it. There are various methods for this, for example neural networks or general statistical methods.

**Deep Learning** or **Deep Neural Learning** is again a subset of machine learning. Here, with the help of neural networks, methods and systems get developed that are modeled on biological brains.

The term Artificial Intelligence was coined in 1956 at an international conference called *The Dartmouth Summer Research Project*. A fundamental idea was to model the human brain and to construct more advanced computer systems based on this knowledge. It was believed that it would soon transpire how the human mind works. The transfer to a machine should then only be a small step.

This turned out to be slightly overoptimistic. Nevertheless, some important findings have been gathered at the conference. It was stated, for example, that the key factors for an intelligent machine should be

- independent learning;
- natural language processing;
- creativity.

Wisdom, creativity and social skills are terms that hardly apply to machines up to now. Rather, the tasks that AI scientists are currently concentrating on are:

- speech recognition and machine language translation;
- medical diagnosis such as evaluation of x-rays or computer tomography scans;
- Internet search engines;
- optimization of sales strategies;
- scientific and technical applications, in astronomy and geophysics etc.;
- forecast of stock prices and foreign exchange rates;
- handwriting recognition;
- personal identification using face and fingerprint recognition.

All of these applications require a certain number of skills originally reserved to humans. But whether this is already "intelligence" is often questioned. It therefore makes sense to take a look at what the term intelligence should actually include.

### 1.1 "Super Intelligence" in three steps?

First of all, artificial intelligence can be divided into two areas: weak and strong AI. The "strong" or general AI is essentially the attempt to create artificial human beings. The goal would be machines or robots that have all the mental powers of humans, including a consciousness [12].

A "weak" AI, on the other hand, can perform only certain tasks very well, but does not even come close to human capabilities in other areas. Examples of weak AI include computers that can beat people at chess — not only average players but also world chess champions. Weak artificial intelligence is nowadays widespread in science, business, and healthcare. For example, AI enterprise systems contain powerful analytical tools and can provide recommendations and insights for business development. They are used for risk management and planning.

Strong AI begins as soon as machines become "human", make their own decisions, or learn new skills without human help. These systems are not only able to solve classical mathematical or logical problems, but also have certain emotional abilities. It is quite easy to program machines to perform some "emotional" or "verbal" actions in response to certain stimuli. ChatBots and virtual assistants are for example already quite good at having a conversation. Some experiments in which robots learn to read human emotions have also been carried out successfully. However, these efforts were more concerned with producing emotional responses than real mental competences.

The next step would be a so-called "super intelligence". This is a name given to systems or machines that are superior to humans in many or even all areas of intelligence. It includes all, or at least many, areas of the intellect, in particular, creative and problem-solving skills or emotional and social skills. Self-confidence or a comprehensive memory are also frequently considered prerequisites for real super intelligence. It does not matter whether these functions are based on biological or technical principles.

However, considerations about a Super Intelligence quickly led to dystopian scenarios. Once an artificial intelligence has been created that can continue to learn completely independently, the step to a superhuman level of intelligence is ultimately only a matter of

time. The controllability of such an AI could ultimately become an existential question for all of humanity.

Scientists who work in this area and make a good living from it, are downplaying this danger. This is certainly all too human. It is therefore important that as many people as possible deal with these topics. Only in this way, a realistic assessment of the potentials and risks will be possible in the end.

## 1.2 How machines can learn

Machine learning is part of the broad field of artificial intelligence that focuses on how to teach systems to learn without being programmed to perform specific tasks. The key idea behind ML is that it is possible to develop machines that independently learn from data, generate knowledge and make useful predictions. To get closer to this goal, three requirements must be met:

### 1. Availability of extensive data sets

These can consist of numbers, images, text, audio samples or other information. Creating the data sets is often the most extensive part of creating an AI system.

### 2. Formulation of a learning objective

In so-called supervised learning, the training data contains "correct" and "incorrect" solutions. These are marked accordingly. During the learning process, the system learns how to make correct assessments.

In the case of unsupervised learning, on the other hand, the system must decide by itself in which categories a data record can be divided. A typical example is a data set with pictures of people. Various systems have already succeeded in independently recognizing that they can be divided into men and women.

### 3. An adaptive system

In order to achieve the best performance, different procedures are often combined. At the beginning of the learning phase, the ML system only contains a large number of default parameters. With an appropriate algorithm, the parameters are varied until the system delivers results that are as correct as possible. Although neural networks are a widespread variant here, they are not the only possibility for the development of adaptive systems.

Deep learning is a machine-learning method that was inspired by the structure of a human brain. With deep learning algorithms, complex, multilayered neural networks can be trained efficiently. In particular, the "back propagation" algorithm is considered to be one of the most promising approaches.

In neural networks, information is transmitted via virtual connection channels. During a learning or "training" session the connection strengths of these channels are optimized. Using a huge amount of data, a large number of parameters are determined, which ultimately form a ready-to-use system.

Not only numerical or visual data can be processed. Speech recognition systems can also be implemented in this way. The soundwaves can be displayed as spectrograms to which a neural network can assign certain words.

However, there is currently no perfect, universal algorithm that works equally well for all tasks. There are actually four main groups of ML algorithms. In addition to monitored and unsupervised learning, a distinction is made between semi-monitored and "reinforcement learning". Each variant has its own strengths and weaknesses.

Supervised learning is commonly used for classification and regressions. Application examples are spam filtering, speech recognition, computer vision, image recognition and classification.

In the case of unsupervised learning, initially the program is not given any pre-classification. The system automatically divides the samples into groups. Unsupervised learning is often used to break down data based on certain similarities. The method is therefore well suited for complex data analysis. This enables machines to find patterns that humans cannot recognize due to their inability to process huge amounts of data. This variant is suitable, among other things, for discovering fraudulent financial transactions, forecasting sales or analyzing customer preferences. The classic applications are data segmentation, detection of anomalies, customer recommendation systems, risk management and the detection of fake data or images.

In semi-supervised learning, there is a mixture of labeled and unlabeled data. In principle, certain prediction results are known, but the model can and may also find its own patterns in order to find a certain data structure and optimize predictions.

Reinforcement learning comes closest to its human counterpart. People don't need constant supervision to learn effectively. Rather, new insights are drawn from positive or negative experiences. Touching a hot saucepan or riding a bicycle over a pin or nail led to a corresponding learning success. Training with ready-made data sets is no longer necessary. Instead, the system can learn in dynamic, real-world environments.

In a first step, simulated worlds are often used in intensified learning research. These offer ideal, data-rich environments. The achievable scores in video games are ideally suited to train reward-motivated behaviors. The main applications are autopilot systems, self-driving cars and trains or autonomous robots [12].

In some areas these methods have already been used to develop systems that are superior to humans. For example, when diagnosing X-ray images, certain types of cancer are already better recognized by ML systems than by trained X-ray specialists. In addition, unlike the doctors, these machines can carry out routine evaluations 24 hours a day, 7 days a week.

## Chapter 2 • A Brief History of ML and AI

Although the term artificial intelligence first appeared at Dartmouth, the question of whether machines can really think is much older. In a famous paper entitled *As We May Think*, a system was proposed as early as 1945 that would model knowledge processing on human ways of thinking. A little later, Alan Turing published a technical article about how machines could take on tasks that require a high degree of "intelligence". In this context, chess was also explicitly mentioned.

Nowadays, the exorbitant computing power of modern computers is undisputed. Nevertheless, it is doubtful whether a machine can really "think intelligently". Even the lack of an exact definition of the term "intelligence" is problematic. Moreover, the most important advances in the field of AI are not even recognizable to many people. The new methods are often used in very subtle ways. For example, when it comes to investigating customer behavior and, subsequently, influencing purchasing decisions, there is no interest in public education in this regard.

Furthermore, there is a certain tendency to constantly redefine what "intelligent" means. When machines have solved a problem, afterwards it is quickly regarded as trivial or as plain computing power. For example, chess — the "Game of Kings" — was considered for centuries to be a strategy game requiring massive intelligence. After the chess world champion was defeated, the game was considered nothing more than "computing power".

The future of AI becomes easier to understand if we look at some important points in its history. After the Dartmouth Conference, one of the next milestones was a computer program capable of communication with humans through a keyboard and a screen. ELIZA surprised many people with the illusion of a human conversation partner.

Expert systems were finally used for the first medical applications. Highly-developed computer programs supported medics in diagnoses and routine analyses. From 1986 onward, computers slowly learned to speak. Using predefined phonetic sequences, machines were able to pronounce whole words and sentences intelligibly for the first time.

Finally, further technological developments paved the way for artificial intelligence in everyday life. Powerful processors in smartphones and tablets offer consumers extensive AI applications. Apple's Siri, Microsoft's Cortana software, and Amazon's Alexa are conquering the markets. Since 2018, AI systems have been discussing space travel and arranging hairdresser's appointments.

Finally, a system called AlphaGo managed to beat the reigning world champion in the game of Go. This Asian board game has significantly more variants than chess. With the pure pre-calculation of possible moves, the chances of success are therefore extremely low. With the help of AI methods, however, the seemingly impossible was achieved. A very special move attracted particular attention. Much later, it became clear how brilliant it was. Many experts were already talking about creativity and intuition on the part of the machine.



A little later, a successor system learned the game without any human support. This means that self-learning machines have reached a new dimension. The extensive training with human partners was not required anymore. The new system learned the basics and intricacies of the game, and even developed its own strategies.

Both in chess and Go, all information is available to everyone involved. Poker, on the other hand, is a game with incomplete or hidden information. This gives "intuition" and the assessment of competing players a special meaning. Nevertheless, the machines achieved decisive success here as well, by winning important poker tournaments. AI systems are also able now to learn the game of poker independently from human training partners. Some even discovered the ability to bluff by themselves.

The knowledge gained in poker is also of interest in other areas such as medicine or professional negotiation techniques. Here too, intuition often plays an important role. Therefore, medical students are usually inferior to experienced medics in diagnosing diseases. Experience and intuition are also required in finance and investment. This is why there are increasing applications in these areas for the results gained from the AI's poker victories.

The development of AI was by no means free from setbacks and crises. Among other topics, this was demonstrated by the so-called XOR problem. The XOR or "exclusive or" gate is a classic problem in digital technology. The task is to construct a network element that emulates an XOR logic gate. This gate delivers a logic one at the output if the two inputs are not equal. If, on the other hand, the inputs are equal, a logic zero is occurs at the output. In tabular form:

Input 1	Input 2	Output
0	0	0
0	1	1
1	0	1
1	1	0

At first glance, the XOR gate seems to be a very simple problem. However, it turned out that this functionality poses a major problem for neural network architectures. So-called "perceptrons" are elementary units that correspond to biological neurons. They therefore form the basis of artificial neural networks. Each preceptron can receive an input from other units and calculate a certain output. Typically, the weighted sum of all received values is formed and a decision is made as to whether a signal is to be passed on to other units. However, the XOR problem cannot be linearly separated. This can best be demonstrated when the XOR input values are plotted on a graph. As can be seen from Figure 2.1, there is no way to separate the predictions by a single straight classification line. Therefore, a perceptron can only separate classes that can be divided by straight lines.

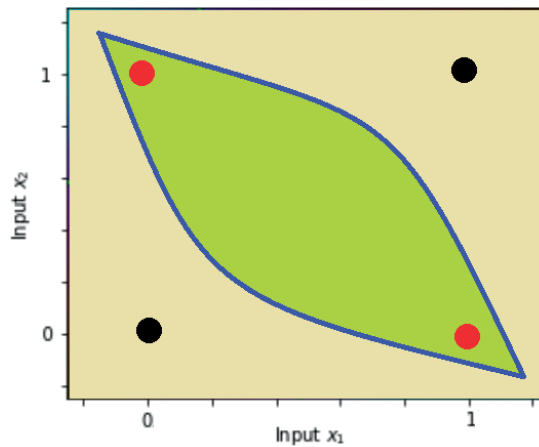


Figure 2.1: The XOR problem.

Hence, the XOR problem cannot be solved with a perceptron. The other logic operators AND, OR and NOT could be emulated without any problems, but not the XOR function! This realization plunged AI development into a deep crisis. It took several years until research groups seriously addressed the topic of neural networks and machine learning again.

The solution to the problem is to go beyond the single-layer architecture. Additional layers of "neurons", in particular so-called hidden layers, finally made it possible to solve the XOR problem. How this solution looks in detail is explained in section 10.9.

Despite its long, colorful history and considerable success rate, artificial intelligence is still not mature in many areas. Many applications need to become much more dependable or fault-tolerant. Sensitive areas such as autonomous driving or medical diagnoses require the highest degree of reliability. In addition, AI systems must become increasingly transparent. All decisions must remain comprehensible to human beings. Otherwise, the AI-based decision on granting a loan or applying a certain a medical therapy will never really be accepted in everyday life.

## Chapter 3 • Learning from "Big Data"

One of the important tasks in the research area of Machine Learning is to analyze, recognize or interpret all kind of data. In most cases, a large accumulation of randomly distributed data is more or less useless. Only the categorization and division into meaningful groups, statistical evaluations or trend analyses etc., give the data a certain practical use. One goal of machine learning is therefore to transform data into knowledge and draw useful conclusions from it.


For a long time, the application of ML methods was reserved to a small group of experts. The algorithms used were complex and hardly understandable for laymen. It was only through the development of open-source libraries that non-experts were given the opportunity to deal with AI topics and machine learning in more detail. This made AI procedures and data structures accessible to a wide range of users. The Python programming language in particular made it possible to embed extensive functions in simple "wrappers". With the help of uncomplicated Python instructions, today nearly anyone can use the powerful algorithms. No supercomputers are required. Rather, even small systems can recognize patterns in large amounts of data or derive predictions about future events and developments [2].

In the next section, the basic structure of artificial intelligence and machine learning will be presented. In addition, the associated terminology will be explained in more detail. This will provide the basics for solving practical applications. In particular, the following topics will play a major role:

- Different variants of machine learning and AI.
- Learning systems, neural networks and training methods.
- Installation of a development environment for the programming language "Python" or its interactive variant IPython.
- Creating a software basis for ML methods and procedures.

### 3.1 Machine Learning and Artificial Intelligence

The research fields of artificial intelligence are closely related. It is often difficult to separate the individual sub-areas from each other. In addition to machine learning methods, there are other areas such as classical expert systems or so-called evolutionary algorithms. However, many of these different regions have lost some of their importance in recent years, as current developments are increasingly focusing on neural networks and deep learning.



*Figure 3.1: Different areas of artificial intelligence.*

Machine learning as a subregion of artificial intelligence is increasingly becoming the central core of current research. Independent "learning" from extensive amounts of data replaces personnel-intensive and problem-specific, explicit programming. The processing of data with known correlations and the "learning" of structures through so-called training play a central role.

Later, the extracted rules can also be applied to new data and initially unknown contexts. The new methods analyze known data, establish certain rules, and finally make decisions based on them. Based on various mathematical methods, data sets are divided into hierarchical levels. This can eventually lead to practical applications such as image recognition or categorization. With the help of Python, the complex mathematics of neural networks can be reduced to simple functions. This allows users to quickly implement practical projects and yields relevant results in a short period of time.

So-called deep learning can in turn be considered a special field of machine-learning methods. The goal is to imitate human learning behavior as effectively as possible. The basis for this is the use of large amounts of data as "learning material". Deep-learning algorithms for neural networks currently deliver the best results in the area of image or pattern recognition, etc. The neural networks used for this purpose usually consist of several intermediate layers that serve as a connection between the input and output nodes.

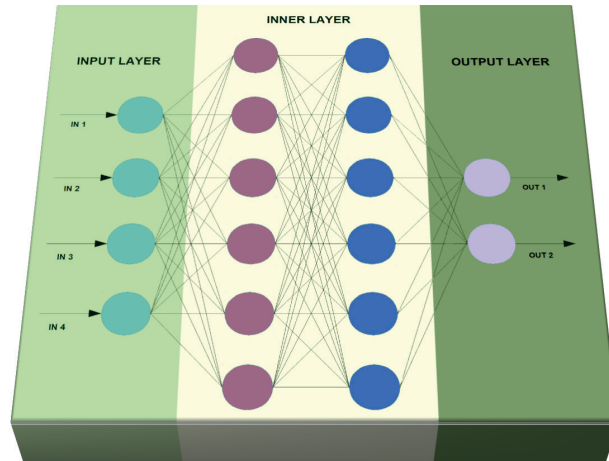


Figure 3.2: A neuronal network.

The layers contain so-called artificial neurons. The function of these neurons is emulated by mathematical operations (see Chapter 11). The aim of this development is to simulate the structure of the human brain with its network of neurons, axons, dendrites and synapses as efficiently as possible. Deep learning is characterized by the use of multiple layers in the neural networks. The input layer processes either the direct raw data, for example, the individual pixels of an image, or preprocessed or filtered data streams.

The inner or "hidden" layers take on the further processing and data reduction. The output layer finally delivers the results. The typical structure of such a neural network is shown in Figure 3.2. The deep learning methods can also be used to model complex relationships. The large number of network levels also allows to map highly non-linear functions. Without this structure, even simple logical gate functions would be unsolvable, as can be seen for example with the XOR gate (see Chapter 2).

In contrast, deep learning algorithms are able to solve comparatively complex tasks. Through multiple "training runs", the in-depth learning can be improved with each calculation step. It has thus developed into one of the central development drivers in the field of artificial intelligence [11].

The main differences between machine learning and deep learning are conveniently summarized in a table.

	<b>Learning data</b>	<b>Size of the data sets required for learning</b>	<b>Necessary Hardware structure</b>	<b>Required Training times</b>	<b>Interpretation of decisions</b>
<b>Machine Learning</b>	structured	All Sizes	Simple	Short	Simple
<b>Deep Learning</b>	No special requirements	Extensive	Challenging	Days to weeks	Almost impossible

An important advantage of deep learning compared to general machine learning is the ability to process unstructured data. Real information such as images, sounds and videos etc. can be used here more or less directly as input data. Other machine learning algorithms, such as decision tree methods do not have this capability. If, for example, images are used as input data, complex and special program adjustments must always be carried out by real people.

Examples of AI processes that work without neural networks are the so-called evolutionary algorithms. These programs are also based on a natural learning process. However, it is not the brain that is modeled, but the biological evolution. Here, the algorithms follow one another in different generations. According to Darwin's principle of "survival of the fittest", every possible solution is evaluated and the best version is selected in each case. The "mutated" variants thus lead to an optimized algorithm that is better adapted to the given problem than the previous one.

The last column of the table indicates a possibly serious problem with AI algorithms. In the case of complex neural networks, it is usually no longer possible to understand exactly how decisions are made. In the case of massive wrong decisions by AI systems, this can lead to considerable legal or social problems. Ultimately, the acceptance of such decisions will depend to a large extent on their perceivability. AI applications that do not have a minimum level of transparency will therefore hardly be able to establish themselves in (vital) areas.

## Chapter 4 • The Hardware Base

Just a few years ago, the prevailing opinion was that machine learning or AI could only run on server farms or super computers. The high-performance machines that won chess games became "Jeopardy" masters or beat human GO specialists, shaped the public image of AI. However, with advances in hardware and software, the situation has changed dramatically. Even a mid-range PC now has several gigabytes of RAM and small single-board systems like a Raspberry Pi 4 are available with up to 8 GB of memory.

Even on microcontrollers like the ESP32, ML applications are now performing amazingly. In addition, special chips such as the Kendrite K210 are increasingly seen in the market. The hardware structure of these chips is already specifically designed for neural structures.

Several different systems will therefore be used in this book. In addition to the classic universal PC, both the **Raspberry Pi** and the **MaixDuino** will prove their capabilities in the individual projects. The following table provides an overview of the projects related to the hardware. This should not exclude porting a solution to a different hardware base.

Unfortunately, it is relatively time-consuming, for example, to control peripheral devices directly with a PC, since it often lacks the appropriate interfaces. With the Raspberry Pi or the MaixDuino, on the other hand, LEDs or relays can easily be controlled via the existing input/output lines and pins. Speech recognition and synthesis, on the other hand, would also be feasible on a PC. However, the implementation in Python is much easier on a Raspberry Pi system.

	PC, Laptop	Raspberry Pi	Sipeed MaixDuino
Basics	X	X	
Iris classification	X	X	
Handwriting recognition	X	X	X
Clothes sorting	X	X	
Speech recognition / synthesis		X	
Object recognition		X	X
Face recognition		X	X
Home automation		X	X
Size (approx. cm)	30 × 40 × 15	9 × 6 × 1.5	7 × 6 × 1
Price (approx. €)	500	50	30
Power consumption (W)	50	10	5

## Chapter 5 • The PC as Universal AI Machine

A modern PC or laptop can do a good job in the field of machine learning. However, some requirements should be met for this. The following table provides an overview of the minimum requirements:

<b>CPU</b>	Quadcore with at least 3 GHz
<b>RAM</b>	minimum 16 GB
<b>HDD</b>	1 TB

It is certainly possible to implement some entry-level projects with less powerful equipment. However, in this case the training times for neural networks, for example, become very tedious even for simpler applications.

### 5.1 The computer as a programming center

The PC is not only used as a basis for getting started with ML projects. Rather, it is also required for other purposes. For example, the MaixDuino cannot be programmed directly. The PC is required here as a host computer which programs the Maix board via the USB interface.

An active USB hub can do a good job here. It has the particular advantage in offering a certain protection for the USB port of the PC or laptop used. It is important that it is actually an **active** hub, meaning the device must have its own 5 V power supply via a separate power supply unit. This is the only way to ensure that the computer port is protected in the event of a short-circuit behind the hub.



Figure 5.1: An active USB hub.


Of course, even an active hub does not offer absolute protection. However, it is very unlikely that a short-circuit will break through an active hub and reach the computer's USB port.



If you want to use the Maix board without a PC, i.e., in a stand-alone mode, an external 5 V power supply unit is required too. Ideally, the hub's power supply unit can be used for this. After programming, the board can be used independently of an USB port. It is important for the power supply to provide sufficient current. About 2000 milliamps (2 A) nominal current should be the minimum. If the power supply still has a USB-micro connector, an adapter to USB-C is required (see Figure 5.2). Further information on stand-alone operation can be found in section 7.4.



*Figure 5.2: A USB-C adapter.*



*Figure 5.3: A USB power supply.*

At this point, it should be mentioned that the MaixDuino has an integrated voltage regulator, just as the Arduino. This means that the boards can also be supplied via any power supply with an output voltage of 6 V to 12 V. The unit must be able to supply a current of at least 2 A. For the connection to the respective board, the power supply unit must have a standard coaxial power connector.