1 Artificial Intelligence Strategy

Digital Agenda Vienna

30 July 2019

Commissioned by:

City of Vienna

2 Management summary

As a cutting-edge technology, artificial intelligence is bound to play a role in the public sector as well. The Vienna City Administration has accorded this technology a pivotal role in the ICT strategy and intends to put forward specific ideas and recommendations for using artificial intelligence to improve quality of life in the City of Vienna.

This AI strategy was mapped out between September 2018 and January 2019. Chapter 2 and 3 set out the scope and define the strategic objective as regards artificial intelligence in the Vienna City Administration.

Chapters 4 through 6 describe the technical requirements for artificial intelligence and probe the challenges expected to arise from the use of artificial intelligence in the public sector, including current trends and progress achieved. With a view to developing an understanding for the broad use of artificial intelligence in various areas of work and life, chapter 7 provides a detailed description of its areas of application. To ensure smooth implementation of this technology, chapter 8 lays out a conceivable AI framework which is meant to provide terms of reference for the sustained use of artificial intelligence.

The AI strategy's central focus is to identify, develop and implement specific use cases for AI application which were further refined during preparation of the AI strategy using the assortment of generic use cases described in chapter 9.1 . These generic use cases should also be understood as modules for future use cases and are thus meant to serve the City of Vienna as tools for the independent identification of potential AI use cases.

A rundown of the results and recommended action are provided in chapter 10 .

3 Scope

The AI strategy of the Vienna City Administration is intended for all organisational units that fall into the remit of the Vienna City Administration and carry out AI projects in Vienna. In particular, these include the offices of the City of Vienna and their enterprises.

4 Strategic goals at the sub-strategy level

4.1 Role of the AI strategy in the Digital Agenda

The AI strategy of the Vienna City Administration should be understood as part of the overall ICT strategy and is intended to provide the strategic terms of reference for implementation of AI projects in the City of Vienna. This document was drawn up in a participatory process involving experts, key city administration stakeholders and interested citizens.

4.2 Vision and mission

Over recent years, artificial intelligence has entered a new stage of maturity and is becoming the driver for innovation and the digitalisation of systems in all areas of life. In 2018, the European Commission developed a Strategy on Artificial Intelligence[1] and therein proposed an approach that put the human at the centre of AI development ('human-centred AI'). All areas of city administration are apt to benefit from artificial intelligence. Its use must always take due consideration of moral and ethical principles. The principles of efficiency, effectiveness, security and utility for citizens must be upheld to ensure the successful implementation of AI projects.

Using methods of artificial intelligence, we want to create new services in the Vienna City Administration and make all existing services in all areas more efficient and convenient for citizens. Efficiency is enhanced by using the completely new opportunities that artificial intelligence has to offer in automation and process support.

Also, collaborating with the companies and startups in Vienna when implementing AI methods in the Vienna City Administration will further strengthen Vienna as a business hub. The use of AI methods can contribute significantly to establishing Vienna as Europe's digital capital.

Given the pace of developments when it comes to AI methods and applications, only those who begin to actively use artificial intelligence now will cut it in a digital society. Society, business, administration and science are encouraged to recognise the potential and risks of AI.

[1] ec.europa.eu/newsroom/dae/document.cfm?doc_id=51625

4.3 Strategic goals

1. Our intention is to harness the potential of artificial intelligence in an effort to improve services for citizens and, furthermore, to create new services that take into account the citizens' interests.
2. The application of artificial intelligence improves user-friendliness in citizen services and makes it more convenient to use such services.
3. We have opted for sustainable, reusable and universally applicable core AI infrastructure components. Generally, all services and applications should get to access these AI components and integrate them in their processing.
4. Generic application scenarios and specific case studies provide terms of reference for the meaningful and expedient implementation of artificial intelligence and, in so doing, offer a baseline for the assessment of AI projects.
5. Taking into account the marketability of AI technologies, the security, availability, stability and functionality of AI systems, the City of Vienna will identify and implement appropriate use cases.
6. AI skills will gradually be developed over time in the ICT Department of the City of Vienna and the specialised departments.
7. The responsible use of artificial intelligence will ensure that decisions are not left entirely to a computer system. Human beings will continue to be responsible for decisions.
8. In our efforts to further develop AI applications, we rely on collaboration with companies, startups and research and development institutions.

5 Artificial intelligence

Artificial intelligence is an area of computer science that aims to enable machines to perform tasks 'intelligently' (Fraunhofer-Gesellschaft, 2018). In its broad sense, artificial intelligence consists of computer systems that can act intelligently. On that account, such systems usually do not simply implement instructions that have been programmed into it. Rather, they adapt their behaviour to suit the situation on hand using the information available. Machine learning frequently plays an important role: for example, based on a set of historic data that reveals a system's ideal behaviour in a given situation, a system learns which behaviour to adopt in the future when a similar situation arises.

Only AI systems specialised to perform specific tasks are currently of any practical relevance. In some cases, they may actually be able to do the specific job better than human experts.

5.1 The history of artificial intelligence

Artificial neural networks were first conceived as early as in the 1940s and initial implementations were developed a decade later. After some initial success in the 1950s and 1960s, research and development slowed in the 1970s on account of the high complexity and the slow computers, precipitating what is also known as the first 'AI winter'. It was not until the 1980s, that AI experienced a second boom triggered by the development of what are known as expert systems, which use a knowledge base consisting of manually entered logical rules. However, it soon became clear that it would become increasingly difficult to consistently build up knowledge bases over time and that it was practically always impossible to define all conceivable prerequisites for a given task. This led to the second 'AI winter' in the late 1980s. At the turn of the century, the progress achieved in computer technology and the increased availability of large data volumes set the stage for complex artificial neural network learning, initiating the success of today's artificial intelligence (Döbel, et al., 2010, p. 9f).

5.2 Definition of terms

5.2.1 Artificial intelligence

Artificial intelligence has been defined and interpreted in many different ways. John McCarthy, an AI pioneer, was the first to define the term in 1955 as the goal 'to develop machines that behave as though they were intelligent'. However, seemingly complex behaviour can be produced using simple electrical circuits, which renders this definition inadequate given that artificial intelligence has set itself the goal of addressing many serious practical problems that require more than simple circuits.

Encyclopaedia Britannica defines the term as follows: 'Artificial intelligence [is] the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings'. This definition has its shortcomings, too, since a computer that can store a long text and retrieve it anytime or is capable of performing complex multiplication, for example, may be considered to have intelligent abilities.

Elaine Rieh and Kevin Knight resolve this dilemma by defining artificial intelligence as 'the study of how computers do things at which, at the moment, people are better'. This definition will continue to be applicable even in 2050 (Ertl, 2016, p. 1ff).

5.2.2 Machine learning

Machine learning emerged after the convergence of statistics and artificial intelligence and has the objective of 'learning' from data and making as accurate predictions as possible. The resulting output (the function that has been learned) is called a *model*. In the final analysis, nearly all models that ML creates from sample data are statistical models, which is why statistics constitutes the backbone of the machine learning theory (Fraunhofer-Gesellschaft, 2018). Model training requires lots of calculations and is thus expensive. Although pre-trained models are available for some purposes, they are often only suitable for specific and thus limited applications such as the simple identification of objects in pictures.

5.2.3 Artificial neural networks and deep learning

Artificial neural networks are part of what constitutes machine learning, and enormous progress has been achieved in this field since 2006, especially in the analysis of image and video as well as voice and text data. Machines using artificial neural networks can translate texts, answer questions, compress messages and even generate images or texts themselves (Döbel, et al., 2010). In these artificial neural networks, an entered value passes through multiple calculation levels to generate an output value. At least three levels are needed so that a suitable pattern can be extracted from the raw data (Reder, 2018). 'Deep' networks comprising lots of different levels have proven to be particularly successfully, establishing the notion of 'deep learning'.

5.3 Machine learning algorithms

Machine learning uses a large number of model types and algorithms, each of which is suitable for a different set of tasks. This section provides an overview of the classification and regression methods, instance-based methods, decision trees, Bayesian models, clustering and artificial neural networks along with information on how each of them work and which tasks they can sensibly be used for. When selecting suitable algorithms, it is essential to logically understand the added value they provide in the specific case at hand and which data is required.

5.3.1 Classification methods

Classification is the separation of objects into different classes based on matching features. Images can be sorted by their content, for example. The number of classes may vary. In handwriting recognition, for example, a separate class needs to be created for numbers. Classification methods usually do not determine the specific class of an object but predict the class with a certain degree of probability instead (Zeigermann & Nguyen, 2018, p. 24f).

5.3.2 Regression methods

With regression problems the goal is to predict a value, provided there is continuity in the output. For example, this method can be used to predict a person's annual income based on level of education, age and place of residence. Mathematically, regression is not very different from classifications. However, in terms of the purpose of application it is: the goal of regression is to predict continuous values (a machine's temperature, for example), whereas the point of classification is to distinguish classes (the machine overheats or does not overheat, for example) (Müller & Guido, 2017, p. 32f).

5.3.3 Instance-based methods

Instance-based algorithms compare new data with pre-trained data that is deemed important. Typically, new data is compared with sample data previously deposited in a database, and the dataset with the closest similarity is used for classification. Therefore, these algorithms are also called 'winner takes all' methods, since only the most similar dataset can be used for classification. The algorithms vary in their similarity measurement, which gauges the distances between events (Kropp & Töbel, 2017, S. 16).

5.3.4 Decision trees

A decision tree is used to organise objects into groups based on suitable features with respect to a given target variable. The flowchart has a tree-like structure with a root, several internal nodes and terminal nodes. The classification of an object using a decision tree starts at the root and assigns the object to the internal nodes of a new layer until the object has been clearly classified. In practice, this algorithm is frequently used to analyse the response to advertising, verifying the effectiveness of medicinal products or checking the credit standing of bank customers (Bankhofer & Vogel, 2008, p. 273).

5.3.5 Bayesian classifiers

Some algorithms and methods use Bayes' theorem for classification and regression. This theorem describes the probability of an event based on the interplay of independent events (Kropp & Töbel, 2017). Bayesian classifiers are used to identify unsolicited emails, for example. Using training data, the algorithm learns to define the classes *solicited* and *unsolicited* and then determines the relative frequency of terms and phrases that often appear in what is called email spam. If the words *money* and *rich* appear several times in a message, for example, chances are high that the message is spam. However, if the message contains the word *money* and other words that are not typical of spam, such as *safe deposit*, it is very likely that the message is not an unsolicited message.

5.3.6 Clustering

Clustering is an example of unsupervised learning, i.e. a system produces information via a dataset for which no 'correct' answers are available (the assignment of individual data points to categories, for example). In clustering, data is automatically categorised and sorted into related groups or clusters without first defining these groups. Human intervention is only needed to interpret the groups (Zeigermann & Nguyen, 2018, p. 28f).

5.3.7 Artificial neural networks

Artificial neural networks (ANN) are made up of layers of nodes called neurons. The connections between the neurons are referred to as weights and are continuously adjusted during neural network learning until the output is satisfactory (Döbel, et al., 2010, S. 11). Each node (neuron) has an activation function that determines whether and how the node sends on the signals (numbers) it received from the previous nodes. ANNs may consist of several hundred layers and billions of weights to be learned and are thus capable of approximating complicated mathematical functions. Compared to other algorithms, ANNs require huge data volumes to ensure that they are worthwhile (Hecker, Dr., et al., 2017, p. 8f).

5.4 Learning styles in machine learning

Machine learning involves various learning styles which can be used to different ends. These learning styles are *supervised learning*, *unsupervised learning*, *semi-supervised learning* and *reinforcement learning*.

5.4.1 Supervised learning

In supervised learning, the algorithm is trained using known data in order to generate a mapping function between the input and output data. The training dataset includes both known datasets and the related outputs to be predicted. The trained model is then applied to new datasets in order to make a prediction.

Two types of analysis are used in supervised learning: classification and regression. Classification produces suitable and freely selectable categories whereas regression turns out a continuous and real value (Kropp & Töbel, 2017).

5.4.2 Unsupervised learning

In unsupervised learning, the system is obliged to learn without any instructions, because the training data is not labelled or has an unknown structure. The aim is to identify the unknown structure in order to discover unknown patterns (Dorn, 2018, p. 18). The best-known use case for this method is clustering, where persons can be segmented according to their known characteristics, for example, without clearly defining the segments first (small, medium and large in clothes sizes, for example). Another case would be the clustering of documents by their content without specifying the categories beforehand.

5.4.3 Semi-supervised learning

Frequently, the volume of labelled data in the total data pool is too small and it is impossible to create a representative and balanced model and ensure training via supervised learning. Semi-supervised learning can be of service in this case, since this learning method can be implemented for training with a small volume of labelled data while the unlabelled data can be used to improve the model's accuracy (Bonaccorso, 2018, p. 19f). For example, semi-supervised learning is used in sentiment analysis to determine a text's emotional polarity (e.g. positive or negative) As a result, unlabelled sentences can still help to improve the trained model on account of their similarity to labelled sentences.

5.4.4 Reinforcement learning

Because it is often unclear what will help and what will not help to achieve the desired outcome, it is often impossible to label the data accordingly. Reinforcement learning methods, which always use agent-based approaches, are implemented for such a problem configuration. The agent (software that carries out specific operations with a certain degree of independence) receives continuous feedback in the form of rewards and/or punishments, allowing it to learn the best (possible) strategy over time to solve the problem at hand (Frochte, 2018, p. 23f).

The agent could use this method to independently learn how to play chess by having each move it makes rated as *good*, *bad* or *neutral* based on pre-defined criteria and thereby continuously improve its abilities.

6 Challenges in the use of artificial intelligence

Compared to standard programming, machine learning is a data-based technology with its own set of challenges. The more training data is available, the better the model works and the lower the error rate. At the same time, the initial data must be of high quality and unbiased.

Another difficulty is that the model needs to be kept general enough so that it can also deal with any unknown data emerging after the training stage. Likewise, the models need to be robust, i.e. they need to respond similarly to similar input (Döbel, et al., 2010).

6.1 Data handling

The ongoing digitalisation of work and private life has ensured greater data availability in terms of quantity, quality, data types, access and the fact that mass data can be linked up. Under its Data Excellence Strategy, the City of Vienna is looking to provide reliable information and data as the central asset in an open administration of the future. The main pillars are enterprise data management, data quality management and a system of data governance that spans the entire organisation.

The handling of such data volumes requires a broad subset of skills ranging from data acquisition and processing to data analysis and visualisation and, in current practice, is guided by the CRISP-DM (cross-industry standard process for data mining).

CRISP-DM phases:

- **Business understanding:** Definition of objectives and requirements; inference of specific tasks and a rudimentary procedure

- **Data understanding:** Data collection and first-time triage of the data available; identification of possible data quality issues

- **Data preparation:** Creation of the final dataset to be used for modelling

- **Modelling:** Application of appropriate data mining methods, parameter optimisation, usually identification of several models

- **Evaluation:** Selection of the model that best fulfils the task. Carefully alignment with task at hand.

- **Deployment:** Processing and presentation of results; where applicable, model integration into one of the client's decision-making process (Seiter, 2017, p. 3ff).

Figure1: CRISP-DM process

6.2 Data quality

The quality of the training data largely determines the quality of the model. If the algorithm is shown too many wrong examples, for example, it will be unable to learn the right answers (Döbel, et al., 2010). High quality of the training data is guaranteed by the following:

- Variations in spelling or the misspelling of names or terms constitute a potential source of errors that should be avoided.
- Well structured and standardised data form the basis for the expedient use of artificial intelligence.
- Depending on the use case at hand, varying data volumes are needed to identify relevant patterns and draw valid conclusions.
- Duplicate datasets lead to a skewed output; therefore, the data needs to be screened first.
- Data is usually in some way time-related. Therefore, outdated data should not be used if predictions for future events are to be made, for example.
- Usually, data requires preliminary manual annotation and/or categorisation by domain experts to enhance the quality of the output.
- Where possible, master and movement data from distributed systems (CRM, ticketing or ERP systems, for example) should be pooled.

6.3 Unbiased initial data

The performance of contemporary AI systems hinges on the quantity and quality of the data available. The fact that AI systems reproduce any discrimination existing in the datasets is an important factor that must be taken into account. For example, when historical data shows that men used to obtain more senior positions than women, an AI system may be induced to draw the false conclusion that men are generally more suitable for such positions. However, the systems themselves cannot provide any information on the bias inherent in the underlying data. Even inadequate data will produce an output, although it will be inconclusive. What complicates the matter is that subtle forms of data-based bias are only difficult to identify, e.g. when other factors correlate with sex, and thus affect the outcome. The less apparent the correlation, the more difficult it is to make out the correlation based solely on the results (Dickow & Jacob, 2018, p. 5ff). Therefore, the following aspects of biaslessness should be ensured:

- AI systems should be designed in a way that clearly shows the extent to which the training data is representative for real data applications and in a manner that discloses the error rate of the analysis

- Especially when working with customer data, their objectiveness must be guaranteed, and the source data should be free of any geographical or cultural bias and correspond to the target group at hand.

- Frequently, data serves a specific agenda. What must be ensured is that the data source is reputable, trustworthy, credible and not biased.

- Domain knowledge: It must be clear how the data was collected and which methods (in market research, for example) were implemented. When a questionnaire is poorly structured, the quality of the data will be inferior too.

6.4 The traceability of decisions

The traceability of a model's decisions can be a major challenge. Traceability is easily established using decision trees, whereas neural networks offer only little traceability (Döbel, et al., 2010). If, for example, an AI system is supposed to learn how to distinguish cats from dogs using pictures, the motifs of the training dataset are usually clearly identified as cats or dogs. The AI system 'knows' the right answer, so to speak, but not how to arrive at this answer. Even if the AI system were able to allocate the pictures correctly and reliably, the basis of any such allocation would still be unclear. While humans recognise dogs by their typical appearance, an AI system may base its allocation on specific colour contrasts or a picture's edges, which humans are not even aware of. One way to make this image recognition process more transparent is to draw attention to the image elements that led to a specific classification.

Although perhaps of relatively little consequence when identifying cats and dogs, understanding and critically analysing the reasons for a specific output in AI is vitally important in lots of use cases. Developing transparent and traceable AI processes is essential for the safe use of AI not only in image recognition, which has been cited as an example, but also in many other areas (Otto & Graf, 2018).

Frequently, a model's traceability and its complexity are caught up in a time conflict. A linear model depicts simple relationships and is thus easily explained, but this same simplicity puts a limit on its area of application. Many problem issues follow complex rules and are modelled in a highly abstract manner using neural networks. As a result, complicated problems can be successfully solved, but the transparency of the solution finding process diminishes.

6.5 Ethics and risks of artificial intelligence

Society is in the midst of a transformation heavily influenced by digitalisation and, in particular, by the opportunities that artificial intelligence presents, and the pace and end of this transformation is unknown. From today's vantage point, it can be assumed that human work will not be entirely supplanted by computers. However, some activities and professions will increasingly be replaced by machine automation. It is thus becoming ever more important to reflect on ethical issues in this context.

Starting out from deliberations on what *morality* is understood to encompass and how this notion is defined within the realm of ethics, two areas are examined where ethical considerations are relevant for *artificial intelligence*, i.e. the human consequences of working with computer systems deemed *intelligent*, on the one hand, and responsibility issues that arise when such systems are implemented, on the other.

Crucial pertinent issues have already been addressed by Joseph Weizenbaum in *Once More The Computer Revolution* (1979).

I. Who are the beneficiaries of the much-touted technical progress and who are the victims?

II. Which limitations should humans, in general, and researchers as well as engineers, in particular, impose on the use of computers in consideration of human concerns?

Notwithstanding the first point, Weizenbaum puts forward his demand while at the same time cautioning against leaving specific cognitive processes (especially processes concerning the analysis and evaluation of complex situations and decisions with far-reaching consequences that are based on them) entirely to computer systems no matter how perfect these may be, insisting that computers should only be used as tools by humans so as to be able to better, more comprehensively and more rapidly analyse a complex system and maintaining that the actual assessment of the situation as well as all the decisions and action should remain in human hands and/or with the human mind (Neumaier, 1994, pp. 41-64).

The strategic objectives of this AI strategy, which are set out in chapter 3.3 , and the recommended action proposed in chapter 10 duly take into account Weizenbaum's observations on responsibility and ethics.

7 Trends and progress in artificial intelligence

According to a study conducted by Gärtner (Andrews & Herschel, 2017), forty percent of all the applications newly implemented in enterprises will be using AI technology by 2020. These applications will primarily be used in interactive chatbots and/or modern analytics, i.e. wherever large data volumes are available and the use of AI in automation, recommendation systems and decision-making processes promises to deliver major improvements. AI will be implemented along existing business processes, taking on individual steps of the process one at a time. This development will essentially be fuelled by three anticipated trends.

1. Advancements in natural language processing will considerably simplify the use of AI-based applications.

2. The broad integration of AI in existing applications and the simultaneous increase in data availability from the *Internet of Things* will conduce to the improvement of business processes.

3. AI-based applications and systems will interact and collaborate to achieve the desired objectives (Andrews & Herschel, 2017, p. 2ff).

Whereas the entire knowledge base had to be put together manually in expert systems, careful selection of the features suffices in machine learning. Neural networks automatically learn abstract representations in their hidden layers and take care of the feature selection themselves (Hecker, Dr., et al., 2017, p. 10). These advancements are powered by the development of improved algorithms, the growing availability of large data volumes and the provision of adequate computing power.

Also, the *Internet of Things* (IoT) should be considered in close connection with developments in the field of artificial intelligence. The networked things have been providing data in previously unknown volumes and, conversely, IoT will only be able to achieve a breakthrough after data processes can be intelligently analysed and controlled using AI. Ever since the costs for AI and sensors have been plummeting and their application readiness increasing, the correlation between AI and IoT has become more and more apparent (Tiedemann, 2018).

The systems used today are considered to belong to *weak AI*, since *weak AI* is usually able to resolve only distinct problems and has been optimised specifically for a given task. *Strong AI*, by contrast, promises ubiquitous application and would approach and resolve tasks similar to humans. However, whether this will ever be the case is a matter of controversy (Cearley, Burke, Searle, & Walker, 2017, p. 4f). Research is currently focused on new generation AI systems, which are able to use and abstract expert knowledge to become more robust, more applicable in a broader domain and self-explanatory (Hecker, Dr., et al., 2017, p. 6).

8 Areas of AI application

Artificial intelligence is used in a wide range of areas. The systems can either be rather autonomous in its action or rather supportive and collaborative and are generally embedded in processes of the physical world or tied into digital environments. AI-based services can automate knowledge or substantive work, make predictions, monitor or also propose action (Hecker, Dr., et al., 2017, p. 8). This chapter provides a general overview of the most important areas where artificial intelligence is used today.

8.1 Natural language processing

Natural language processing (NLP) is a technology that aims to process human language using computers. It is based on the premise that every form of language, be it spoken or written, must be 'understood' by AI on a level suitable for the use case at hand. Not only the individual word is important in NLP but also its connection with other words, entire sentences or facts. To enable this, the system recognises general patterns which it then uses to understand complex relationships and find answers to problems.

These are some of the problems that NLP-based programmes address (Ryte Wiki, 2018):

- Categorisation of text documents
- Extraction of information from text documents
- Text translation into other languages
- Sentiment analysis of texts
- Simplification of texts
- Text conversion to spoken language (text-to-speech, speech synthesis)
- Conversion of spoken language to text (speech-to-text, speech recognition)
- Understanding search queries in natural language
- Detection of advanced and follow-up questions
- Checking plausibility of answers
- Provision of answers in text form

8.2 Cognitive assistants

Cognitive assistants are systems that communicate with users through text or voice and are meant to support and/or replace humans in the performance of cognitive tasks and decision making. These systems are learning, cooperative and primarily digital in nature and typically use one or several of the NLP technologies dealt with in the previous chapter. The term *cognitive assistant* is used synonymously for virtual, intelligent, cognitive assistants, intelligent and virtual agents (chat)bots, smart assistance, cognitive computing or companion systems.

Cognitive assistants are used both privately and commercially and can be divided into two groups: digital assistants (e.g. Apple Siri, Amazon Alexa or IPSoft Amelia) and assistance systems for decision support (intelligent applications that offer real-time decision support in specific tasks). The following speech technologies are particularly relevant:

- Speech-to-text
- Text-to-speech
- Speech interpretation and speech processing (they respond to spoken or written language)
- Knowledge representation (organisation, categorisation and contextualisation of linguistic information)
- Dialogue (exchange of complex information in a conservation)

- Speaker recognition (identification of the speaker)

Learning assistants are trained and continuously improved using historic dialogues. The availability of learning data in the various languages may constitute a technical challenge if the data available is not sufficiently structured and/or may not be used for legal reasons. Along with technological advances, user acceptance is a prerequisite for the dissemination of these systems (Hecker, Dr., et al., 2017, p. 36f).

8.3 Computer vision

Computer vision (machine vision) describes the ability of systems to identify objects, scenes and activities in images and videos. In the process, complex images analysis tasks are divided up into individual edges, lines and textures of objects with a view to determining whether, with a certain degree of probability, the features identified in an image represent an object. Machine vision is used in various applications, for instance in the diagnosis of illnesses, in facial recognition or for security and surveillance purposes (Gentsch, 2017, p. 39). Computer vision is one of the areas of application that benefit in particular from current developments in artificial neural networks (deep learning).

8.4 Intelligent decision support

Expert systems, the predecessors of intelligent support systems, were developed as early as in the 1980s and 1990s. However, their potential was extremely limited because the necessary data volumes and technical means were lacking. The algorithms available today enable decision support in various areas like medicine and law and support humans in making a diagnosis or identifying precedents (Bitkom, 2017, p. 35).

Decision models trained using artificial neural networks, in particular, require transparency and traceability, which cannot be ensured just like that. To avoid unequal treatment and bias, the data used to train learning systems must be representative. What needs to be kept in mind as well is that AI applications lack the sensitivity needed in some specific contexts, especially when unforeseen events occur for which the applications have not been trained (Djeffal, 2018, p. 8).

8.5 Business analytics

The purpose of business analytics is to resolve problem issues using an evidence-based approach in the overall management cycle, from planning, supervision and control. In this context, evidence is understood as any well-founded and objective insight in a subject matter that can be obtained with data from diverse areas within and outside of an organisation using statistical, data mining and machine learning algorithms. The problems that can be solved using business analytics, and their respective database, can be found in a multiplicity of areas within an organisation as demonstrated by the following examples (Seiter, 2017, p. 1ff):

- Provision of predictions about future population trends in individual cities or city districts.
- Evaluation of demand for specific services based on data from social media and/or internal ERP systems.
- Verification to establish whether the qualifications of a job applicant match those of people already working in the organisation in order to hire the best candidate.
- Traffic volume forecasts for various weekdays and/or times of day.

8.6 Robots and sensors

Robots have been used in industrial manufacturing since the 1970s and have been performing repetitive tasks in a controlled and structured environment ever since. Whereas first-generation robots were strictly separated from humans, modern robots are safer, more independent and flexible and, as autonomous robots, are able to screen and adapt their environment and their behaviour accordingly. Artificial intelligence methods are in no small part responsible for the autonomy of robot systems because AI allows these robots to continuously calibrate their environment using known data and patterns and, in so doing, autonomously generate new knowledge and take decision based on this new knowledge.

Generally, the robotics market can be broken down into industry robotics and service robotics. In industry, robots are increasingly used to perform complex tasks together with humans. Outside industrial production, service robots perform semi-automated and fully automated service tasks to improve human well-being (e.g. vacuum cleaning, lawn mowing or pool cleaning).

The progress achieved in language and gesture control is continuously unlocking new ways for humans and robots to work together. Future challenges include the lack of acceptance, the need for developments to be further optimised and safe human-machine interaction (Hecker, Dr., et al., 2017, p. 12ff).

8.7 Robotic desktop, process automation and cognitive computing

Besides carrying out mechanical tasks, robots can now perform administrative work as well. Robotic desktop automation (RDA) is a technology that operates on the basis of pre-defined interaction protocols between the computer and humans and is particularly suitable for decentralised applications on individual desktops, including for the modification of master data or movement data in a single or across several systems. However, RDA software is not capable of autonomously working out solutions to individual problem issues, and the individual working steps need to be initiated by staff members. The use cases can be visually represented, usually occur frequently and exhibit a certain complexity.

Robotic process automation (RPA) goes a step further. What essentially distinguishes this technology from RDA is that the automated processes do not need to be displayed on the screen, it has a higher level of technical integration and no manual trigger is required. Therefore, RPA ensures semi-automated or fully automated processing of structured applications while keeping audit-proof and tamper-safe records of the activities at the same time.

Artificial intelligence only becomes applicable in cognitive computing (CC). CC is able to perform tasks without any predefined setup and can take decisions and generate solutions autonomously based on experience and newly acquired knowledge. This technology does not rely on structured process or data and can also process images or language, for example. Thus, CC is applied to perform specific, complex and sophisticated functions and tasks (Alexander, Haisermann, Schabicki, & Frank, 2018, p. 12f)

8.8 Autonomous means of transport

Autonomous means of transport include vehicles with varying degrees of automation that are able to navigate and/or operate partly and/or entirely without human intervention. A distinction is made between various levels based on the tasks the system is designed to perform and the need of intervention by a human driver. These are:

- Level 0: no automation
- Level 1: driver assistance
- Level 2: partial automation
- Level 3: high automation
- Level 4: full automation
- Level 5: driver-less

Level-1 solutions for individual transport are already available in the market, and current research is strongly focused on levels 2 through 5. It is currently expected that level 5 will be achieved by 2025. The mobility sector holds an enormous potential. Developments revolve especially around autonomous lorries and agricultural machines (CPS.HUB NRW, 2018, p. 15f).

Driverless underground trains are already in operation in many cities of the world and driverless trains will also be used on the future U5 line in Vienna, which is currently under construction. Furthermore, an autonomous bus will start to travel the streets of Seestadt Vienna from 2019.

9 AI framework

The current state of the art does not offer any tool that provides a universal solution for all AI problems. Depending on the task at hand, various tools and methods need to be used. For this reason, it is advisable to use an AI framework as a setup for the comprehensive roll-out of AI in an organisation and/or enterprise. The AI framework helps to prevent the emergence of siloed solutions and individual AI use cases that are not integrated into the overall value chain. In addition, the AI framework offers a game plan for the implementation of AI use cases.

9.1 Gärtner's AI framework

A Framework for Applying AI in the Enterprise

Figure 2: AI framework (Bern & Whit, 2017)

The schema in Figure 2: AI framework (Bern & Whit, 2017) can be read both from left to right and vice versa. It illustrates the extensive use of AI in an organisation. Based on the corporate objectives and the use cases derived from these, one or several corporate divisions are identified (customer service, HR, etc.) that should be tied into an all-encompassing solution. In this case, the use case also determines which general AI solution and application areas (common AI solution areas) will be needed during implementation of the use cases.

This schema also shows that individual or multiple core and key AI technologies (learning, language, vision, analytics, data science, automation, etc.) can be used for the implementation of use cases in all corporate divisions.

9.2 Cloud computing providers

Amazon, Google, IBM and Microsoft are the four largest cloud solution providers and offer an extensive selection of infrastructure components and machine learning tools. These players set themselves apart from one another by the services, operating models and platform services they offer. The portfolios essentially consist of computers (processing power), storage (memory) and networks.

Unlike local architectures, cloud-based operating models have the advantage that specific aspects of the management and server provisioning are automated and thus less time and effort is required for ongoing maintenance and management. At the same time, the service ensures that the application always has sufficient resources available and the server infrastructure is scaled up or down depending on the specific needs. As setting up and operating high-performance and scalable AI systems is expensive, easy and flexible access to cloud computing directly encourages the development of new AI applications.

In order to further facilitate access to machine learning, many cloud computing providers also offer a series of *Machine Learning as a Service* (MLaaS) applications and tools. For specific purposes (in data visualisation, computer vision, NLP, speech recognition and forecasts, for example), pre-trained models can be employed to save time and resources needed for model development and training. There are now many different MLaaS providers with a constantly increasing number of potential use cases.

Cloud services can also be provided as company cloud/on-premises cloud services. To a certain extent, pre-trained models can also be fine-tuned. At the City of Vienna, the use of cloud-based services is governed by the ICT Policy. Any relevant provisions must be observed.

10 The success of AI strategy implementation

Understanding the potential of AI and its uses in one's own organisation is a first step in identifying possible use cases. The identification and implementation of specific AI use cases is essential when it comes to ensuring the successful implementation of the AI strategy. The development and selection of AI use cases will focus on conforming with the strategic goals of the AI strategy mentioned in Chapter 3.

10.1 Generic use cases for the identification of potential use cases

This sub-chapter describes generic AI use cases that can be used as a basis for the identification of specific use cases for the City of Vienna. To ensure a better understanding, each generic AI use case is provided with information on the appropriate input required and/or the output that can be produced. In addition, a flow chart is provided for each use case which describes the technical steps involved and the model training required. The sample solutions should be taken as illustrative examples. In each specific case, all the available technologies and solution architectures should be considered, and ongoing developments should be taken into account.

10.1.1 Evaluation of unstructured data in text form

Business information is not only stored in tables or databases but also in the form of texts and is thus usually unstructured - as documents (e.g. applications, expert opinions, official notices), emails and in the form of machine-generated records. Furthermore, text content can be derived from the internet, in the form of social network postings or posting in discussion forums. The evaluation of unstructured data has great potential since 'old' data sources represent large knowledge repositories, on the one hand, and even modern documentation systems continue to produce lots of *free text*.

On account of the emergence of new analysis tools and the further development of semantic understanding and machine learning, organisations can analyse and use unstructured data. Possible use cases range from simple text extraction and summaries all the way to automatic support in decision-making processes, as in the evaluation of expert opinions.

Figure 3: Input & output in the evaluation of unstructured data in text form

Figure 4: Sample flow chart for the evaluation of unstructured data in text form

1. Text content available in a form that cannot be machine read (e.g. handwritten text or low-quality scans)

2. Conversion into a standard format (e.g. PDF or TIFF)

3. To improve the success rate of OCR, a number of routines are followed in order to remove any image noise, for example

4. Identification of headings, paragraphs and tables in the layout at hand. Depending on the layout, at least 1,000 documents are analysed according to the following criteria: language, font, free text vs. print form, handwriting, etc.

5. Transformation of unstructured data in characters

6. Extraction of information from pre-defined fields

7. Results become available in semi-structured form

8. Further processing and/or decision making

10.1.2 Time series analysis

Generally, time series are collections of data points observed in time order. This includes population or economic data, the demand for goods or supply services, visitor numbers, transaction data, traffic data or even climate data. Time series analysis aims to predict future values and to recognise any structures, such as trends and sudden changes, and can be used in a wide range of applications.

Figure 5: Input & output in time series analysis

Figure 6: Sample flow chart of time series analysis

1. Input in the form of structured time series data
2. Adding data from further sources, which may impact the prediction
3. Data selection and transformation in appropriate time slots (e.g. daily or monthly)
4. Development and training of prediction models using a suitable method. The model is trained and evaluated on an ongoing basis as soon as actual data is generated.
5. Prediction of the desired value

10.1.3 Identification of anomalies

An anomaly is an event with attributes that deviate considerably from the norm and thus constitute an irregularity. In small data volumes, an anomaly can be identified manually. However, in large data volumes, intelligent algorithms are needed to support identification. Accessing internal IT systems from abroad at unusual times of day or changes in a person's purchasing behaviour following the theft of a credit card, for example, may constitute an anomalous signal indicating a departure from the norm. Moreover, irregularities in the operation of a machine or IT devices can be identified, which could serve as a trigger for predictive maintenance (Ott, 2018, s. 26f).

Figure 7: Input & output in anomaly detection

Figure 8: Sample flow chart for anomaly detection

1. Input in the form of master data and, where available, known anomalies
2. Link-up with knowledge base and extension to include peripheral data and/or networks
3. Data transfer into a structure suitable for the model
4. Selection of a suitable model
    a. Anomaly detection: identification of statistical irregularities
    b. Network analysis: analysis of anomalies at correlative level
    c. Prediction: forecast based on known anomalies
5. Identification of anomalies

10.1.4 Computer vision

Images and videos from various sources are suitable for applications in computer vision systems. The objective is always to semantically understand and/or interpret objects or scenes based on the image and video content. The flow chart below describes *photogrammetry, a special case that* determines the spatial location or three-dimensional shape of an object using photographs (e.g. aerial photographs) and exact reference coordinates of the object.

Figure 9: Input & output in computer vision

Figure 10: Sample flow chart for computer vision (photogrammetry)

1. Input in the form of city maps, driving data, image sequences or movement data from GPS devices
2. Sensor fusion and matching of image with the respective location
3. Distance determination by the 3D scanning points
4. Plausibility check of the corresponding points, identification of incorrect attributions
5. Image projection into 3D
6. Development of the 3D model
7. Identification of the relevant areas in the 3D model
8. Classification of the identified areas and further processing and/or decision making

10.1.5 Chatbots & assistants

By using chatbots, organisations get to support their staff and/or clients around the clock and provide them with knowledge. The integration of various systems and data sources allows intelligent bots to recognise the context and - just like a human employee - to make certain assumptions. The use of human language for communication with virtual assistants facilitates interaction and user experience. Instead of presenting a complex user interface with lots of options (that are usually irrelevant), bots focus on the user's specific concern. Any communication in written or spoken form for which feedback (e.g. information) is expected can be used as input for chatbot applications.

Figure 11: Chatbot & assistant input & output

10.2 Recommended approach for the implementation of an AI project

The approach to be taken for the implementation of AI projects must be structured. Unlike regular IT projects, development must focus on the repeated verification of feasibility when it comes to AI applications since feasibility may vary depending on the environment and the data available. In many cases, no empirical data is available because lots of AI methods and algorithms for everyday use are untried and untested or, as the case may be, are very much research-based in nature.

Therefore, the following approach is recommended for the implementation of an AI project and should be adjusted accordingly depending on the complexity of the project at hand:

1. Assessment of the project scope and specific problem issue, project planning and initial analysis of technical and financial feasibility.
2. Identification of data sources and their scope or, as the case may be, analysis of data quality. Where applicable, training data must be consolidated if the data volume is insufficient.
3. Implementation of a feasibility study (proof of concept) in order to analyse the project's basic viability.
4. Development of a prototype that performs the required core functionality.
5. Development of a productive system.
6. Integration in the system environment.
7. Go live

Once steps 1 through 4 have been completed, it is often useful to check whether it makes sense to continue with the project based on the insights gained up to that point in time.

11 Summary and recommended action

In order to achieve the strategic goals defined in chapter 3 of the AI strategy drawn up by the City of Vienna, the following action is recommended:

11.1 Process for identifying additional AI use cases

Besides the use cases identified during the development of the AI strategy, additional AI use cases should be created and implemented in a continuous process that has yet to be established. To this end, both the generic use cases described in chapter 9.1 and completely new ideas could be used. Especially when it comes to services for citizens, it is important to clarify whether the use of AI methods will lead to any lasting improvement.

11.2 Implementation of specific AI use cases

Proposals for the implementation of specific AI use cases are drafted on the basis of the generic use cases described in chapter 9.1 . In an initial stage, it is recommended to carry out a feasibility study for 2-3 of these use cases and, in the event of a positive outcome, to have the productive system developed and implemented. During implementation, special attention should be given to the marketability of AI technologies, the security, availability, stability and functionality of the AI systems.

11.3 Development of AI skills in the City Administration

Over the long term, in-depth knowledge and skills relating to artificial intelligence should be developed in the organisation of the City of Vienna in order to ensure the following tasks:

- Ongoing identification of AI use cases within the specialised divisions of the municipal departments and their enterprises.
- Achieving acceptance of and confidence in AI technology through ongoing communication, training courses and workshops with employees.
- Management and ongoing improvement of previously implemented AI use cases.
- Internal project management and coordination of external service providers in the development and implementation of an AI use case.
- Cross connection to the current data excellence strategy.
- Completion of feasibility studies and development of own prototypes for testing.

To carry out these tasks, the following skills profiles are required:

11.3.1 AI / machine learning skills profiles

AI/ML researcher	Makes the basic technology available. Is able to combine and develop state-of-the-art technologies and processes and to devise entirely new methods. Conducts basic research to explore new methods and new areas of applications or applied research to adapt existing technologies for new areas of application.
AI/ML engineer	Uses state-of-the-art basic technologies, combines these and develops machine learning and AI applications. Ensures data selection, pre-processing, processing and follow-on processing. Plans, implements and carries out quality assurance. Plans, implements and carries out quality assurance.
AI/ML assistant	Ensures manual or semi-automatic generation of training data. Has rudimentary basic knowledge of AI/ML and data organisation. Is able to carry out simple quality assurance and monitoring tasks.
KI/ML manager	Is able to set up and manage an AI/ML organisation and organise the financing. Has the ability to assess the technology potential of AI/ML and ensure the general legal and business-related conditions necessary for successful AI/ML projects.

11.3.2 Data science / data engineer skills profiles

Data scientist	Has extensive knowledge of standard machine learning methods and the current state of the art. Is able to independently implement prototypes on the basis of state-of-the-art technologies and processes and draw up detailed technical specifications for productive AI systems.
Analyst	Has an in-depth understanding of various field-tested machine learning methods and implements these in a modular manner using tools or libraries. Has sufficient domain knowledge to answer questions independently and interpret output.
Domain expert	Has in-depth knowledge of the significance of data from the value chain. Also has broad AI and machine learning knowledge. Serves as the interface between the developers of machine learning applications and their users.
Data engineer	Bears responsibility for the provision of data and has in-depth knowledge of database systems, data leaks and data pipelines. Draws up data schemas and ensures data quality and performance. Has basic knowledge of AI/ML.

DevOps engineer	Supervises AI and data systems and ensures a high degree of availability. Plans and implements updates and migrations and ensures system security and system access mechanisms.

11.4 Standardisation

Whoever sets the standards, controls the market. Common norms and standards reduce technical barriers, help to open up markets and enhance competitiveness in business. Common standards can improve the user-friendliness of applications and enable interoperability. On this account, Europe must be empowered to act as a driving force in international standardisation processes. To this end, the City of Vienna, together with experts from science and business, will review the following courses of action:

- Launching an initiative to better represent Viennese and European interests in international standardisation bodies.
- Stronger commitment for the development of open and international standards.
- Involvement of the group *Sofortmaßnahmen der Stadt Wien* (Immediate Action Group).

Competence centres will be developed and networking with other key European players will be ensured without delay. In addition, dedicated competence centres will be set up. The City of Vienna will take the relevant measures in the spirit of these key points within the scope of ongoing programmes and the 2019 budget.

11.5 AI framework

Every time a new AI application is envisaged, commissioned or implemented, a review should be carried out to check whether the AI core and/or key technologies needed for implementation are already available in the City of Vienna and whether these can be integrated in the envisaged solution, even if, for example, an acquired solution brings its own core technologies.

Furthermore, it would be advisable to revert to existing MLaas systems or assess the availability of such systems prior to project start in order to save time and resources when developing of an AI solution.

11.6 Provision of a prototyping environment for developers

A prototyping environment is a dedicated area for testing software (and/or models) where the activities performed do not affect the external environment. On the one hand, this prevents any damage from being done and, on the other hand, the developers can record software action and learn from this action. Providing internal developers at the City of Vienna with a prototyping environment would give them the opportunity to test new AI applications and their suitability without affecting the systems already in place.

11.7 Ensuring human-centred AI

AI systems must assure transparency, traceability and verifiability so that effective protection can be warranted against distortion, discrimination, manipulation or other fraudulent use, especially when using prediction and decision-support systems. This also means that the decision-making authority is not delegated to a computer system but remains the responsibility of humans.

For this purpose, AI system developers and users must assume prime responsibility and carefully check, at the beginning of the process, whether the available data is suitable for the intended analysis. At the same time, clear and understandable information must be provided explaining how an AI system works, why it issues a given recommendation or takes a given decision in a specific situation and which data, if any, is generated. Also, it is especially important to check whether the initial data itself contains any bias in order to preclude self-fulfilling prophecies. In the long run, the success of AI systems can only be guaranteed if broad public acceptance is achieved and users feel they can rely on the technology.

LME (local interpretable model-agnostic explanations) is an extremely comprehensive explanation technique to ensure explicability. LIME provides explanations that are found locally using simple, interpretable models and, in so doing, allows for an approximate understanding of complex machine learning models.

LIME is based on three principles:

- Explanations are found locally and independently for each instance.
  - A simple model is adjusted locally to predictions from the complex model.
  - Explanations are given based on the original variables, even if actual classification is based on abstractions of the original variables, as in a convolutional neural network (CNN).

This means that the decisions can be directly retraced and the image pixels that were decisive for classification can be marked, for example. For standard tabular data, the features that contributed most heavily to classification and the manner in which they contributed to classification can be determined; it becomes clear which words in a text played the most significant role in the prediction of the model.

12 List of references

(16 Nov. 2018). Retrieved from Ryte Wiki: https://de.ryte.com/wiki/Natural_Language_Processing

Alexander, S., Haisermann, A., Schabicki, T., & Frank, S. (March 2018). Robotic Process Automation (RPA) im Rechnungswesen und Controlling - welche Chance ergeben sich? *Controlling - Zeitschrift für erfolgsorientierte Unternehmenssteuerung,* S. 11-19.

Andrews, W., & Herschel, G. (2017). *Develop Your Artificial Intelligence Strategy Expecting These Three Trends to Shape Its Future.* Stamford: Gärtner.

Bankhofer, U., & Vogel, J. (2008). *Datenanalyse und Statistik.* Wiesbaden: Gabler.

Bern, E., & Whit, A. (2017). *A Framework for Applying AI in the Enterprise.* Stamford: Gärtner, Inc.

Bitkom. (2017). *Künstliche Intelligenz: Wirtschaftliche Bedeutung, gesellschaftliche Herausforderungen, menschliche Verantwortung.* Berlin: Bundesverband Informationswirtschaft, Telekommunikation und neue Medien e. V.

Bonaccorso, G. (2018). *Machine Learning Algorithms: Populär algorithms for data science and machine learning.* Birmingham: Packt Publishing.

Cearley, D., Burke, B., Searle, S., & Walker, M. (2017). *Top 10 Strategie Technology Trends for 2018.* Stamford: Gärtner.

CPS.HUB NRW. (2018). *Künstliche Intelligenz und Maschinelles Lernen - Chancen, Potenziale und Herausforderungen in Nordrhein-Westfalen.* Wuppertal: CPS.HUB NRW.

Dickow, M., & Jacob, D. (April 2018). *Das globale Ringen um die Zukunft der künstlichen Intelligenz.* Berlin: Stiftung Wissenschaft und Politik.

Djeffal, C. (2018). *Künstliche Intelligenz: Beitrag zum Handbuch Digitalisierung in Staat und Verwaltung.* Berlin: Alexander von Humboldt Institut für Internet und Gesellschaft.

Döbel, I., Dr. Leis, M., Vogelsang, M., Neustroev, D., Dr. Petzka, H., Riemer, A., et al. (2010). *Maschinelles Lernen - Eine Analyse zu Kompetenzen, Forschung und Anwendung.* München: Fraunhofer-Gesellschaft.

Dorn, S. (2018). *Programmieren für Ingeneure und Naturwissenschaftler.* Berlin: Springer-Verlag.

Ertl, W. (2016). *Grundkurs Künstliche Intelligenz: Eine praxisorientierte Einführung.* Berlin: Springer-Verlag

Fraunhofer-Gesellschaft. (2018). *Maschinelles Lernen - Kompetenzen, Anwendungen und Forschungsbedarf.* München.

Frochte, J. (2018). *Maschinelles Lernen: Grundlagen und Algorithmen in Python.* München: Carl Hanser Verlag.

Gentsch, P. (2017). *Künstliche Intelligenz für Sa/es, Marketing und Service; Mit AI und Bots zu einem Algorithmic Business - Konzepte, Technologien und Best Practices.* Berlin: Springer-Verlag.

Hager, T. (04 2018). Künstliche Intelligenz - Mehr Sicherheit oder neue Sicherheitsrisiken? *Organisator, p*. 40.

Hecker, Dr., D., Döbel, I., Petersen, U., Rauschert, A., Schmitz, V., & Voss, Dr., A. (2017). *Zukunftsmarkt Künstliche Intelligenz: Potenziale und Anwendungen.* München: Fraunhofer-Allianz Big Data.

Kropp, M. D., & Töbel, L. (2017). *Machine Learning.* Troisdorf: Sigs Datacom.

Müller, A. C., & Guido, S. (2017). *Einführung in Machine Learning mit Python: Praxiswissen Data Science.* Sebastopol: O'Reilly.

Neumaier, O. (1994). Was hat "Künstliche Intelligenz" mit Ethik zu tun? *ConceptusXXVII, pp*. 41-76.

Ott, T. (2 2018). Anomalie-Erkennung mit Machine Learning. *BI-Spektrum, pp*. 26-30.

Otto, P., & Graf, E. (2018). Künstliche Intelligenz - Häufig gestellte Fragen. Berlin: Konrad-Adenauer Stiftung e. V.

Reder, B. (2018). *Studie Machine Learning/ Deep Learning 2018.* München: IDG Business Media GmbH.

Seiter, M. (2017). *Business Analytics: Effektive Nutzung fortschrittlicher Algorithmen in der Unternehmenssteuerung.* München: Vahlen-Verlag.

Tiedemann, M. (15 January 2018). Warum Künstliche Intelligenz dem Internet of Things zum Durchbruch verhelfen wird. *Digitale Welt.*

Zeigermann, O., & Nguyen, C. (2018). *Machine Learning kurz & gut.* Sebastopol: O'Reilly.

13 List of figures

14 List of Abbreviations

Abbreviation	Definition
CC	Cognitive computing
CNN	Convolutional neural network
CRF	Conditional random field
CRISP-DM	Cross-industry standard process for data mining
DL	Deep learning
ERP	Enterprise resource planning
FCNN	Fully connected neural network
ICT	Information and communication technologies
IoT	Internet of Things
AI	Artificial intelligence
ANN	Artificial neural network(s)
LIME	Local interpretable model-agnostic explanations
LSTM	Long short-term memory
ML	Machine learning
MLaaS	Machine learning as a service
NLP	Natural language processing
OCR	Optical character recognition
RDA	Robotic desktop automation
RF	Random forest
RPA	Robotic process automation