Approaches to Artificial Intelligence as a Subject in School Education

Peter Micheuz

University Klagenfurt, Institute of Informatics Didactics, Austria,
peter.micheuz@aau.at

Abstract. Due to recent developments in the field of artificial intelligence (AI) and its impact on many areas of life, this paper provides an overview of that field, focussing on current approaches, especially in schools. After a clarification of the particular terminology in a wider context, and after a short journey into the past of AI in schools, current initiatives and AI-related approaches on a school level are described. In the following section, the disciplinary aspect of AI is highlighted. This paper concludes with some implications for the practice of AI education.

Keywords. Artificial Intelligence, Machine Learning, Deep Learning, Data Science, School Education

1 Introduction

Writing a contribution about artificial intelligence in schools is a challenging task. One reason is the abundance of relevant existing on- and off-line resources, and the other is the difficulty to offer an overview of the already many studies and initiatives in that field. However, in this paper, an attempt is made to give a comprehensive overview of this multifaceted and broad field.

The pace of recent developments in AI has surprised not only insiders, but also the public and schools. If anyone thought that the subject of computing in schools is already completed and curricula do not need any major revisions, then they have recently been shown a different picture.

Although there is a plethora of insightful and valuable books, papers and a vivid blogosphere about AI terminology (the particular AI glossary of Wikipedia consists already of more than 300 terms), a compact overview is given in section one.

About thirty years ago, when the subject of computing was in its infancy, AI already played a certain role in school education. Section three gives some insight into this historical period. The following section provides an exemplary overview of recent initiatives and AI-related projects, and is concluded by an illustrative rapid run-through of some approaches to impart AI at various levels of education. In section five, the interdisciplinary nature of AI is explored, followed by a short conclusion in which some practical aspects are addressed.

The challenge for providing good practices in AI education and to convey a complete picture of this field is open. It can be expected that AI is not a fad and is here to enter school education, and to stay. From to the title, this paper is concerned with

"learning about AI", and does not elaborate on "learning through AI" in the context of educational technologies, except for the following notes on this issue. AI has the potential to play an important role in educational technology, with many potential applications, from inspiring ones like personalised learning, (automatic) assessment facilitation, assisted language learning and translation to less favourable ones like advanced cheating. However, we still do not know how the digitalisation of education and the adoption of AI will shape learning in this decade [1].

2 What is it all about?

Let us start with a term sometimes used in an educational context: "Deep learning". It stands for meaningful learning, in contrast to human surface and rote learning [2]. "Deep learning" with regard to AI is a method that mimics the workings of the human brain in processing big data for use in predictions and decision making [3]. Its results affect our lives in a way which could not have been foreseen some years ago. It is very likely that most of us have unknowingly been using deep learning models already on a daily basis. A deep learning model is almost certainly used every time we use an internet search engine, a face recognition system on a social media website, a translation system or a speech interface to a smart device. Accordingly, deep learning can be regarded as one of the most powerful and fastest growing applications of artificial intelligence within the sub-field of machine learning.

2.1 From Deep Learning to Artificial Intelligence

All three areas, deep learning (DL), machine learning (ML) and artificial intelligence stand in a hierarchical relationship to each other (Fig. 1), although the concept of what defines AI has changed over time. But, at the core, there has always been the idea of building machines (computers) which are capable of "thinking" like humans.

Artificial Intelligence

Machine Learning

Deep Learning
The subset of machine learning composed of algorithms that permit software to train itself to perform tasks, like speech and image recognition, by exposing multilayered neural networks to vast amounts of data.

A subset of AI that includes abstruse statistical techniques that enable machines to improve at tasks with experience. The category includes deep learning

Any technique that enables computers to mimic human intelligence, using logic, if-then rules, decision trees, and machine learning (including deep learning)
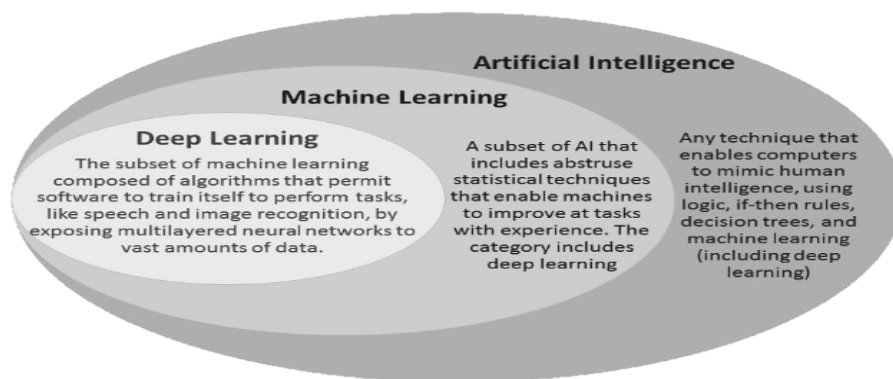
Fig. 1: Relationship between DL, ML and AI [4]

The field of research - already with impacting and fruitful applications in recent years - has become known as "machine learning". Even moreso, it has become so integral to contemporary AI that the terms "artificial intelligence" and "machine learning" are sometimes used interchangeably.

Machine learning is one of the primary approaches to artificial intelligence, but by far not the only one, as will be seen later. There are many similar definitions around, varying just in the wording, but with the same semantics. Wikipedia's definition is: "Machine learning (ML) is the scientific study of algorithms and statistical models that computer systems use to perform a specific task without using explicit instructions, relying on patterns and inference instead. It is seen as a subset of artificial intelligence. Machine learning algorithms build a mathematical model based on sample data, known as "training data", in order to make predictions or decisions without being explicitly programmed to perform the task. Machine learning algorithms are used in a wide variety of applications, such as email filtering and computer vision, where it is difficult or infeasible to develop a conventional algorithm for effectively performing the task." [5]

Used to solve problems which were previously considered too complex, and using the model of neural networks involving large amounts of data and rapidly growing computational power, AI has revolutionised the quality of speech recognition, language processing and computer vision, vehicle identification, driver assistance and other domains.

Generally, AI comprises advanced algorithms based on advanced mathematics, which can handle higher processes similar to humans. Examples include visual perception, speech recognition, decision-making, and translations between languages. Among other trends in information technology, such as the internet of things (IoT), robotics, 3-dimensional (3D) printing, big data, blockchain technology, virtual and augmented reality, AI is one of the leading topics of our digitally penetrated world. AI is often accompanied by misleading stories and thus leads to diverging feelings in the general public, ranging from utopian enthusiasm to dystopian fear. Accordingly, there is a huge challenge for education and schools to provide all pupils with a solid understanding in that field.

As with every new technological achievement, deep learning as the main application of machine learning has its dystopian implications. It is potentially worrying that the trail of data and metadata we are leaving and delivering voluntarily, and largely unnoticed when moving through the online world, is also being processed and analysed using deep learning models. This is why it is so important to at least understand what artificial intelligence is, how smart technologies work, and what they are capable of, and what their current limitations are [6].

It is important to recognise also that AI is a constantly moving target. Things that were once considered within the domain of artificial intelligence - optical character recognition and computer chess, for example - are now perceived as routine computing. Today, robotics, image recognition, natural language processing, real-time analytics tools and various connected systems within the IoT are all increasingly using AI in order to be augmented with more advanced features and capabilities.

Deep learning and, or with, neural networks, meanwhile, gain the most attention because they are particularly well-suited for tasks involving image, video, and audio data

[7]. For text and numerical information, though, the older methods can still be more suitable.

AI's transformative effects on technology will increase over the coming decades, with the development and adoption of deep learning continuing to be driven by rapidly growing datasets, the development of new algorithms, and improved hardware. These trends are not stopping.

2.2 Artificial Intelligence in a Wider Context

We cannot discuss AI without considering the highly related and overlapping wide area of data science (Fig. 2). Simply put, data science is the study of data, involving developing methods of recording, storing, and analysing data to effectively extract useful information. The goal of data science is to gain insights and knowledge from structured and unstructured data. It is the science which brings the saying "data is the new oil" to life.

Data is worth very little if there are no highly-skilled professionals who can derive actionable insights from it. Undoubtedly, the competence to understand, use, process and intepret data has become indispensable and a requirement for an expanding range of jobs and careers. (Big) data are ubiquitous. It is estimated that about ninety percent of the world's data has been created in the last two years [8].
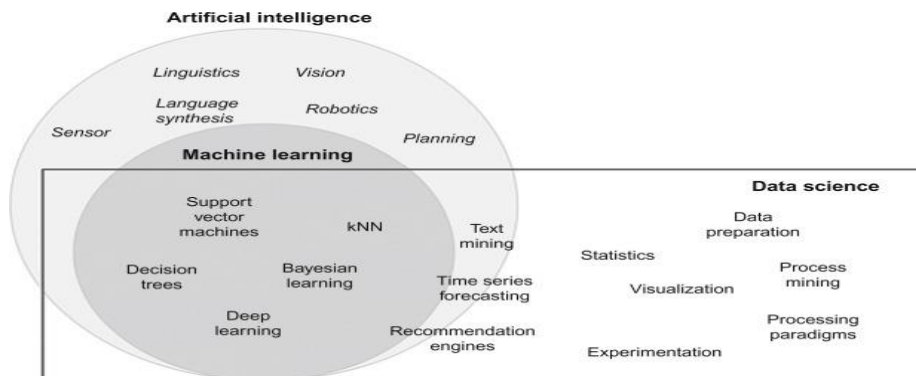

Fig. 2: Relationship between Data Science and AI [9]

Mastering data science and harnessing data requires a solid grounding in: mathematics for processing and structuring (big) data; statistics; programming not just in one specific programming language; and last but not least analytical and computational thinking, including problem solving and logical reasoning.

The new data-driven world requires individuals to be constantly separating fact from fiction. In short, the need to analyse and interpret data is no longer confined to engineering or computer programming; it has become an essential life skill. Yet, the K-12 education landscape is lagging behind. Schools have not recognized today the changes the data explosion has made to society. Curricula we teach currently should be revised and provide more practical utility for the 21st century. There is a widening gap between

competencies students need in life and what is taught in schools. It can be argued that data science including AI should be building blocks of a modern school education.

Machine learning is one of the primary approaches to artificial intelligence, but it has to be seen in the wider context of data science, which encompasses important areas such as the often underestimated and arduous work of data preparation on the one hand, and the fascinating field of data visualisation on the other.

Before learning machines and machine learning provided us with suggestions and predictions in (still) particular situations, (contrary to "general artificial intelligence" with super intelligent robots exceeding the abilities of human beings), they had to be trained with (big) data sets which could be accomplished through "supervised", "unsupervised" or through "reinforced learning". In short, supervised learning requires the supply of training data and correct answers; unsupervised learning occurs when machines learn from a dataset on their own, and reinforcement learning is based on permanent feedback from the environment. Machine learning uses algorithms to learn from data and data patterns, and the knowledge acquired can be used to make predictions and decisions.

Whereas AI is just at the beginning of being included in curricula and lessons, data have been already for a long time a building block of computing education. In some countries, the term "Electronic Data Processing" was the predecessor for the later subject "Informatics". Accordingly, AI is naturally embedded in all aspects and fields around data; that is, data literacy and data management. Recently, a comprehensive model of data key concepts and a competence model of data literacy have been published [10, 11]. This holistic view on data shows convincingly that data driven computing education is very broad, with AI playing an increasingly important role.
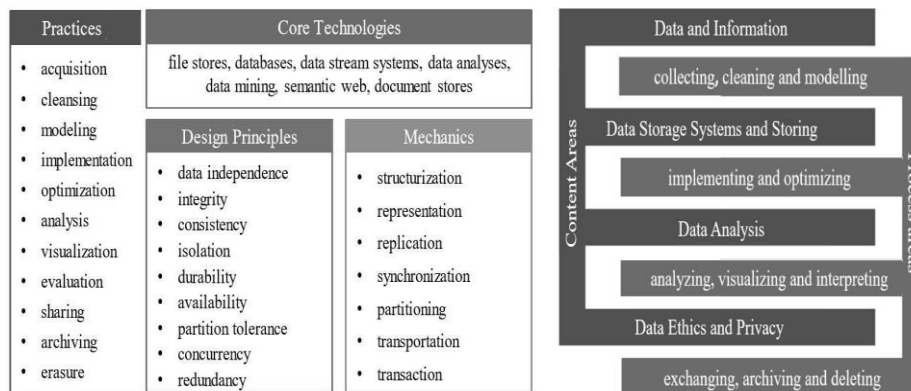

Fig. 3. Key Concepts of Data Management, Data Literacy Competence Model

3 Historical Context

Although AI may be regarded as the 'hot topic' of the moment in (digital) technologies, and the driving force behind many technological breakthroughs of recent years, at least the term is not all that new. During the last decades, AI has moved out of the domain

of science fiction and into the real world, and the theory and the fundamental computer science which makes it possible has been around for decades.

There are many very useful resources on the web which describe this timeline in general and AI milestones in particular [12]. Such content on the web can be harnessed and discussed by pupils in a historical, interdisciplinary context.

Long before robots were dominating utopian and dystopian arenas in science fiction, in the 17th century, the scientist and philosopher Rene Descartes thought of thinking and decision-making machines. While he was wrong in stating that they would never be able to talk like humans, he already distinguished between machines which might one day learn about performing one specific task, and those which might be able to adapt to any job. Today, these two fields are known as specialised and general AI.

The origin of the term "artificial intelligence" goes back to a conference at Dartmouth College (in the United States of America) in 1956, years before the subject computing (computer science, informatics, information technology) came into play into the curricula of timetables in some countries, and accordingly into textbooks about computing/informatics. In their seminal and modern introduction to computer science, Goldschlager and Lister [13] locate artificial intelligence in the chapter "Algorithms in action: some computer applications", subdivided into "Can machines think?", computer games, understanding language, visual perception, knowledge representation and expert systems. The book ends with the (philosophical) question "Superfluous human?" Their comforting assertion that human beings with their creative capabilities, innovation power and originality will never be superfluous can be judged to be correct, but were they right with their statement that "Computers carry out (only) repetitive tasks"?

About ten years later, the first German comprehensive textbook on "Didaktik der Informatik" contained a full chapter about "artificial intelligence" [14], referring to AI as an academic discipline with its subsections of natural language systems, expert systems, robotics, computer vision, followed by a short historical summary and general explanations about cognitive science, including a sceptical view on neural networks. Herein, the author Baumann dismissed statements (from his students) such as "neuronal networks can learn" as an improper use of language.

A similar misjudgement can be read in Rechenberg's "Was ist Informatik" [15], where neuronal networks are denoted as an offspring, and not really belonging to artificial intelligence. "It looks like this idea is doomed to failure. […] Neuronal networks have proven some applicability, but its performance should be estimated very carefully. There are many examples for which much more efficient solutions have been found with other mathematically trustworthy processes."

What followed was an "AI winter", also in schools, until recently. Currently, we are witnesses of an "AI springtime", not only in research and business, but slowly in schools too. However, there is much evidence that AI in schools is still in its infancy.

4 Current AI Initiatives and Approaches in K-12 Education

Following the hype about AI in the 1980s that used a classical rule-based approach, expert systems and the programming language Prolog [15], we currently observe a regaining of momentum in this field, fuelled by promising approaches as shown below.

From this work, it appears that AI can be introduced in curricula and computing lessons in an appropriate and sustainable manner.

Perhaps the most advanced development in implementing AI in K-12 education can be found in China, where a textbook on AI based on computational thinking has been rewritten and published [16]. The old textbook in that field mainly included knowledge representation, reasoning, expert systems, search, etc., whereas the core concepts of the new textbook encompass intelligent systems, artificial neural networks, and machine learning. The shift has been from focussing on expert systems to the analysis and design of intelligent systems incorporating state of the art AI concepts.

In England, the initiative "Computer Science for Fun", an (online) magazine where "the digital world meets the real world" offers many ideas and teaching materials for machine learning (http://www.cs4fn.org/machinelearning), with the so called "Sweet learning computer (a simplified chess game)" as one example [17].

In Germany, there are various initiatives to pilot AI-related projects and studies in schools. These range from unplugged activities in that field [18], an activity-based explanation of how neurons work and learn in robot controlling [19], to machine learning in the context of data science [20].

Computer vision, which perhaps for pupils is the most striking and impressive aspect of AI, can already be treated and discussed in a phenomenological way for a young age-group, using Google's Autodraw as an example. The question if and how a computer recognises animals (such as the fat giraffe in Fig. 4) can be a starting point for stimulating interest in lessons. Another fertile question can address the future of so called "Captchas" (see Fig. 4, an applied Turing test), including a discussion on how much longer it will be before DL will solve also these puzzles. The right face in Fig. 4. is fake; it does not exist in reality and is AI-generated.

Fig. 4. From Phenomenological AI to "Deep Learning of Deep Learning" [22]

Due to deep learning, great progress has been made with character recognition. This aspect of AI is used almost as a daily practice by all internet users and is therefore suited for computing lessons, addressing (un)supervised and reinforcement learning through training with data, underpinned by theoretical foundations.

5 AI is interdisciplinary

Like hardly any other science, AI is interdisciplinary. It uses results from such diverse fields as mathematics, logic, operations research, statistics, control engineering, image processing, linguistics, philosophy, psychology and neurobiology. In addition, in many AI projects, the field of the respective application has to be taken into account. To successfully work on AI projects is therefore not always easy, but almost always exciting and challenging [21]. The challenge is at least as great when thinking of introducing AI into school curricula, as interdisciplinary aspects in this context are rather the exception than the rule.

In Germany, a nationwide initiative about imparting a holistic view on AI within the "Science Year" has been launched [20]. Its ambitious goals, aiming at the target group of 12-18-year-old pupils, comprise:

- a sound explanation of how AI works.
- stimulation of a social discourse on AI.
- a reduction of existing misconceptions.

The course consists of (up to) six modules, containing teaching material, arranged around six modules: "Introduction - students' everyday experiences with AI"; "How does machine learning work?"; "What's the difference between man and machine?"; "Historical overview of the development of artificial intelligence"; "The distribution of roles of man and machine - ethical and societal aspects"; and finally "In which AI world do we want to live?"

Contemporary computer lessons (should, it can be argued) make use of contextualised teaching concepts such as IniK which means "Informatik in Kontext" [23]. IniK is based on the assumption that solely technical computing competences do not suffice to understand the digital world of information technology (IT) systems and digital media. Pupils should be able to use them in a self-determined way. To this end, questions are placed at the centre of computing that go beyond technical issues, and include the social context, aspects of economics, culture, politics or law [24]. Answering these questions makes it possible to develop IT content in a cross-disciplinary way and can lead to sustainable computing competences. Aspects of AI are almost predestined to supplement this approach of IniK with already recommended and elaborated topics, such as "E-mail only(?) for you", "My computer is talking to me!", "Smart and rich through apps", "Social networks", and "Don't trust a picture!"

Lessons according to "Computing in Context" are per se interdisciplinary and subject-linking. Assuming a real-life context, there are manifold references to different subjects. But, most likely, there are only a few teachers who are able to deal with multidisciplinary challenges with professional competence. Maybe appropriate interdisciplinary teacher training in which colleagues from relevant subjects join together to form a team may be the key to remedy this situation. However, the challenge of a sound teacher education in AI is still ahead of us.

6 Implications and Final Remarks

From a theoretical point of view, some studies indicate that many seemingly "intelligent" systems and AI can be demystified in computing lessons, and sooner or later, this much-discussed area of digital technologies will reach school informatics on a broader basis. However, we have to be honest and realise that a constructive and meaningful approach to teach the (currently) 'hot topic' of AI in depth requires a deep understanding of the field, represents a big pedagogical challenge for a teachers and teacher trainers, and of course, a cognitive one for many pupils.

Obviously, it makes a big difference to approach this topic in schools (comprehensively):

- from a social and philosophical viewpoint (talking and reasoning about AI).
- by conscious awareness of AI applications on a phenomenological level (knowing about AI and using AI applications in a reflective way).
- by applying a grey box model of AI, requiring a basic knowledge of its key concepts and programming languages and environments (applying AI).
- through putting the mathematical and computational perspective into the foreground (understanding the foundations of AI and constructing AI).

From a practical point of view, and with the focus on learning outcomes, it is useful to consider the seminal (revised) Bloom's taxonomy [25] in mind, which starts from mere recalling and basic understanding to the creation of AI applications. Referring to the "Tale of Three Learning Outcomes" [26] with three categories "No learning", "Rote Learning" and "Meaningful Learning", the question arises as to how much time can be allocated to AI in (an always) overcrowded curriculum. "Meaningful Learning" is recognised as an important general educational goal and occurs when pupils build their knowledge and cognitive processes that are needed for successful problem solving. This begins with an appropriate mental representation of the problem and ends with the problem solution, in which the pupils devise and carry out a plan for solving the problem.

With regard to didactical aspects of AI education, it is a key question how to deal with the plethora of its possible approaches. The spectrum begins with a discipline of AI in its own right, providing a holistic picture of the field with sequenced und structured lessons, ranging across fragmented approaches within subjects such as computing, mathematics or philosophy. These include singular bottom-up initiatives such as the simulation of a neuron with the programming environment Scratch, and unplugged activities in the form of a role-playing simplified chess-game to demonstrate reinforcement learning. All these approaches in schools are still in their infancy worldwide and at an experimental stage lacking empirical results.

However, recent developments and the obvious progress of machine learning and its impact on all of us suggests the implementation of AI education in school education on a broader and deeper basis. This would have the potential to extend and enrich not only the subject of computing in schools, but education in general.

However, there is still a long way to go to find appropriate approaches for particular age-groups with reasonable levels and requirements. Above all, it needs curious and engaged educators, teachers and teacher trainers able to incorporate this important and prospective field into general, specific and vocational education.

References

1. Editorial Team: https://edtechreview.in /trends-insights/trends/3856-top-five-use-cases-of-ai-in-education (31.10.2020)
2. Entwistle , N. (1988). Styles of Learning and Teaching, David Fulton: London
3. LeCun Y, Bengio Y, Hinton G. Deep learning. Nature. 2015;521(7553):436-444
4. Dhande, M.: What is the difference between AI, machine learning and deep learning. Blog on: https://www.geospatialworld.net/blogs/difference-between-ai-machine-learning-and-deep-learning (31.10.2019)
5. Machine Learning (Definition): https://en.wikipedia.org/wiki/Machine_learning (Page last visited: 31.10.2019)
6. Kelleher J.: Deep Learning. MIT Press. Cambridge, USA (2019)
7. Vincent J.: The biggest headache in machine learning? Cleaning dirty data off the spreadsheets. http://www.data-analysts.org/view/236.html (31.10.2019)
8. Marr B.: How Much Data Do We Create Every Day? The Mind-Blowing Stats Everyone Should Read. https://www.forbes.com/sites/bernardmarr/2018/05/21/how-much-data-do-we-create-every-day-the-mind-blowing-stats-everyone-should-read (31.10.2019)
9. Kotu V., Deshpande B.: Learn more about Artificial Intelligence. In Data Science. Concept and Practice. Elsevier, Amsterdam (2019)
10. Grillenberger, A., Romeike, R.: Key Concepts of Data Management: An Empirical Approach. In Proceedings Koli Calling (2017).
11. Grillenberger, A., Romeike, R.: Developing a theoretically founded data literacy competency model. In Proceedings of WiPSCE (2018).
12. Marr B.: The Most Amazing Artificial Intelligence Milestones So Far. https://www.forbes.com/sites/bernardmarr/2018/12/31/the-most-amazing-artificial-intelligence-milestones-so-far (31.10.2019)
13. Goldschlager L., Lister A.: Computer Science A Modern Introduction Prentice Hall, London (1988)
14. Baumann R.: Didaktik der Informatik. Klett Verlag. Stuttgart (1996)
15. Rechenberg P.: Was ist Informatik? Eine allgemeinverständliche Einführung (3. Aufl.). Hanser Verlag, Munich (2000)
16. Yu, Y., Chen, Y.: Design and development of high school artificial intelligence textbook based on computational thinking. Open Access Library Journal 5(09), 1 (2018)
17. Curzon, P., McOwan, P.W.: Computer science for fun - cs4fn: The sweet learning computer: www.cs4fn.org/machinelearning/sweetlearningcomputer.php (31.10.2019)
18. Seegerer S. et al.: AI Unplugged – Wir ziehen Künstlicher Intelligenz den Stecker. In: Informatik für alle. Proceedings INFOS 2019. Lecture Notes in Informatics. Dortmund (2019)
19. Strecker K., Modrow E.: Eine Unterrichtssequenz zum Einstieg in Konzepte des maschinellen Lernens. In: Informatik für alle. Proceedings INFOS 2019. Lecture Notes in Informatics. Dortmund (2019)
20. Schlichtig M., et. al. Understanding Artificial Intelligence - A Project for the Development of Comprehensive Teaching, To appear in Proceedings ISSEP 2019, Cyprus (2019)
21. Ertl W., Grundkurs Künstliche Intelligenz, Computational Intelligence, Springer Vieweg, Wiesbaden, p. 12 (2016)
22. Ng A., Soo K.: Numsense! Data Science for the Laymen. No Math added. Springer (2017)
23. Diethelm, I.; Koubek, J.; Witten. H.: IniK – Informatik im Kontext, Entwicklungen, Merkmale und Perspektiven. In: LOG IN Heft Nr. 169/170. p. 97-105 (2011)
24. Coy, W.: Informatik im Großen und Ganzen. In: LOG IN,. Heft 136/137, p. 17– 23 (2005)
25. Anderson, L., Bloom, B. and Krathwohl, D., 2000. A Taxonomy For Learning, Teaching, And Assessing. London: Longman.
26. Mayer R.: Rote versus Meaningful Learning. In: Theory into Practice. Volume 41, Ohio State University, USA (2002)