Carsten Schulte (Paderborn University)
Yannik Fleischer (Paderborn University)
Lukas Höper (Paderborn University)
Rolf Biehler, Daniel Frischemeier, Sven Hüsing, and
Susanne Podworny (Paderborn University)

Exploring the data-driven world: teaching AI and ML from a data-centric perspective

Schulte, C., Fleischer, Y., Höper, L., Biehler, R., Frischemeier, D., Hüsing, S., & Podworny, S. (2022). Exploring the data-driven world: teaching AI and ML from a data-centric perspective. In AI, data science, and young people. Understanding computing education (Vol 3). Proceedings of the Raspberry Pi Foundation Research Seminars.

Available at: rpf.io/seminar-proceedings-vol-3-schulte-et-al

Exploring the data-driven world: teaching AI and ML from a data-centric perspective

Carsten Schulte (Paderborn University)
Yannik Fleischer (Paderborn University)
Lukas Höper (Paderborn University)
Rolf Biehler, Daniel Frischemeier, Sven Hüsing, and
Susanne Podworny (Paderborn University)

Abstract

In our talk, we discussed the question of whether artificial intelligence (AI) and machine learning (ML) should be taught differently from other themes in the computer science curriculum, and if so, how to teach them. The tentative answer is that these topics require a paradigm shift for some teachers, and that this shift has to do with the changing role of algorithms, data, and the societal context. The talk presented three teaching examples from the beginning of secondary school (11-13 year old students) to illuminate the possible differences in teaching. The first example drew upon the Matchbox Computer and successors like the Sweet Learning Computer to teach the machine learning process, the second was about enactive teaching of decision trees, and the third was about analysing location data.

Introduction

Teaching artificial intelligence (AI) and machine learning (ML) from what we have called a 'data centric perspective' is an idea that originated from our project *ProDaBi (Project Data Science and Big Data at school)*. We started with a symposium in 2017 to collect some ideas on the topic of data science at school and how to incorporate it in school curricula. Following that symposium, the project started in 2018 with the aim to develop a curriculum for data science at schools, including AI.

We roughly oriented the work around the curricular spider web (see Figure 1); this helped us identify that questions to answer when designing a curriculum should be related to, for example, the content, suitable learning activities, and so on.

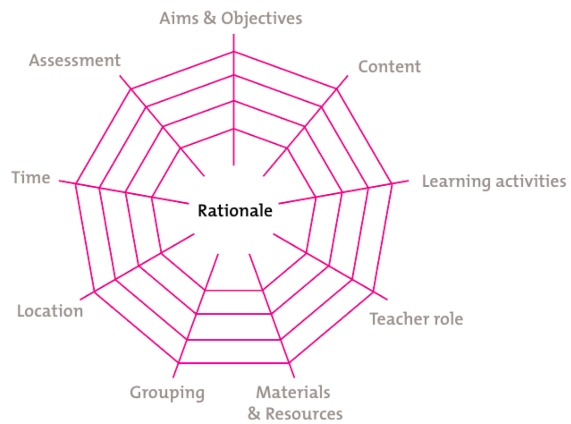

Figure 1. Curricular spider web. From Curriculum in development *by Thijs, A., & Van Den Akker, J., 2009, p. 11, Netherlands Institute for Curriculum Development (SLO). Copyright: 2009, Netherlands Institute for Curriculum Development (SLO).*

At first, we focused on developing materials and resources, and tried to find a place or time slot at school. We initially started with a special elective course at the end of secondary school, and from there gradually adapted resources and teaching ideas to also be useful in lower secondary grades. However, the focus of this presentation is not on the materials, but on the rationale, that is, the middle of the curricular spider web (see Figure 1). So, the question is why we have done it the way we have, and what is behind it. We will present this framework first, or at least some insights into the framework, and then make it more concrete with three examples.

Paradigm shift needed

We tried to answer the question how new and probably changing and complex topics like data science and AI can be included in a curriculum. The motivation to include these topics was — besides future job opportunities — the increasing number of AI and data-driven applications people can use in their daily lives. This is a typical issue for computer science, and one answer is rather popular, which we do not completely agree with. It's formulated as the slogan: "ideas not artefacts" (e.g., Wing, 2006), and is based on experiences as captured in this paper's title: *Computer science in English high schools: We lost the S, now the C is going* (Clark & Boyle, 2006), referring to a misconceived focus on (using) computers instead of teaching the science behind computers. We think it is important to not only use applications, but to also look under the hood. However, this doesn't just refer to ideas. Modern AI applications are mostly discussed in terms of the increased role and progress of ML, hence data-driven applications. And these can only be understood and evaluated by also taking into account the data they are processing, for example, for training their models — this is discussed and argued for in the paper called *Machine behaviour* by Rahwan et al. (2019). The authors argue that these data-driven applications can only be understood by including their behaviour 'in the wild', so to speak.

The argument is twofold. First, data-driven applications rely on vast amounts of data; their performance thus can only be understood by knowing about this — and also using such applications unavoidably impacts the role of data and the need to collect data, often including, for example, user or interaction data. Secondly, such applications behave differently from traditional algorithms. A sorting algorithm also relies on data, but its internal mechanics are independent from the used data. It is a mathematical function transforming an unsorted input to a sorted output — and the algorithm itself (e.g., a bubble sort) is not being changed, no matter what data is fed into the sorting application using bubble sort.

In contrast, data-driven applications like ML applications rely on a model that is derived from the data used to train the model, and this model is not independent from the data but a direct result of the data: different data usually leads to different models.

When such an application is used, the input data is processed based on the trained model — hence based on the prior fed data — and if there are any biases or other issues with the data or the training process, the model can produce unexpected or unwanted results. Moreover, the input data can also be stored and used to further train and change the model. Therefore, unlike traditional algorithms, such data-driven applications cannot be examined and evaluated by using some test cases before being employed and confronted with real data. Hence, discussions surrounding research areas like machine behaviour are needed to examine these technologies when in use.

For education, this means we cannot only consider the ideas, algorithms, and training processes when we want to explain data-driven applications, we have to deal with the data too. And this data, or rather its properties and meanings, are bound to a societal context. That context can be situational: for instance, where, when, and under what conditions data for self-driving cars or data for predicting success as employee or learner in some specific domain and institution is collected (a range of examples is discussed by O'Neil, 2017, and Rahwan et al., 2019).

While Rahwan et al. (2019) suggest discussing and analysing machine behaviour rather abstractly on different levels and within a framework from science, namely biology, we suggest for education to differ in two aspects (see Schulte & Budde, 2018). To demonstrate the everyday application and transferability of ideas to artefacts, we think it can be useful to include the notion of interaction between humans and machines. During such interactions, humans in different roles shape and are being shaped by digital systems. For a discussion on shaping and being shaped see, for example, Rushkoff (2010), or the debate on hybrid interaction systems, man-machine or human-in-the-loop ideas.

Bell and Duncan (2018, p. 141) argued that "[a] complaint about older curricula is that they focus primarily on the applications and the data, algorithms, programs and infrastructure are treated as a black box, while the human is expected to conform to the system, rather than viewing the interface critically and considering what is good about it and what might be improved". Instead, "the big picture of an interaction with artefacts should be at the centre of attention: If we can explicitly confront students with all elements of digital systems in a form that makes sense in their world, we can give them a better understanding of how everything works and enable them to be creators, not consumers" (Bell & Duncan, 2018, p. 142).

In these debates, a common theme is to reflect on which roles and responsibilities should be reserved for humans, and what aspects can and/or should be automated to be processed by the digital artefact. Note, this viewpoint makes it important to not only focus on the artefact itself, but to also include the societal context in which it is used — and in which possible different interaction roles it can unfold. This shift can also be seen regarding the role of data: without context, data is just transformed from input to output; but with context, issues like bias, fairness, completeness, or the need to change the data can occur and be included as topics for education.

In the following section, we present three examples in which we explore and develop approaches to balance the role of the ML mechanisms (e.g., the algorithms used for training a model), the role of data and its contexts, and also the role of artefacts in contrast to abstract ideas.

Examples

Man machine computer

This example is based on the idea of the Sweet Learning Computer (Curzon & McOwan, 2015), as referred to in the report on AI and teaching it at school from The Royal Society (2017). Originally, it was part of a set of teaching ideas for demystifying machine learning (Curzon et al., 2008). The origin of this example, however, dates back to the Matchbox Computer.

The example roughly works as follows. It presents a very simplified chess game with only three figures and a 3 × 3 playing field. This simplified game has a limited number of possible moves overall, and the second player can always win when choosing the right ones. Here, the human has to make the first move. The machine has a list of all possible and correct answering moves and randomly chooses one of them. In the beginning, it is likely that the machine will lose. However, as a kind of machine learning system, every move that leads to the machine losing is removed, so that eventually the only moves that are left as choices are those that let the machine win.

By playing this game repeatedly (we suggest at least ten times), students can experience this gradual and data-based 'learning' process. To make this process more visible, we divided the steps the machine has to take into several sub-roles, each played by a student. This way, the complete mechanism becomes apparent, and students can literally see that the machine is just following an algorithm — there is no human intelligence needed or involved in letting this process of machine learning unfold. The 'intelligence', so to speak, lies in the setup of the machine. The machine does not really understand or learn to play — it just has fewer and fewer possible moves to choose from during the training phase.

It is interesting to take a close look at the data, and the role of the human player or trainer. If the human tries to win and actually wins, a move causing the machine to lose can be removed. But if the human does not make winning moves, the machine cannot remove its own bad moves and does not learn. We can see these moves of the human as input and training data, and students can experience that the result of ML depends on the training data. Regarding the role of the human, if the human chooses to lose or play badly during the training phase, and to later play well, in this way, the human can affect the machine's learning to prevent it from becoming (too) smart.

This way, the man machine game can teach some basic insights into AI and ML. It is, however, an interesting question whether these insights really become conscious to the learners and whether they can relate these insights to real AI applications, e.g., to autonomous cars. This is discussed in the paper by Große-Bölting and Mühling (2020), where students were asked about their understanding of the inner workings of ML systems after having played the game outlined above (in a somewhat simplified version). Interestingly, the authors conclude that there was no real transfer and interpret the internalisation of this concept as being. The role of a verbalisation and reflection phase in addition to playing the game thus seems important, and such a phase should probably include some explicit transfer to real-world applications. Just teaching ideas without making the relationship to artefacts explicit seems not to guarantee the desired learning outcome.

Teaching the systematic creation of decision trees with data cards

This series of lessons aims to give students in grades 5 and 6 an idea of supervised machine learning and artificial intelligence by learning about data-based decision trees. The series is mainly based on unplugged materials that enable action-oriented learning on an enactive level. Additionally, a digital learning environment (for instance, menu-based Jupyter Notebook) can be used flexibly at the end of the series. The selected context of food is relevant for all students and especially suitable for younger students.

Food can be classified as 'rather recommendable' or 'rather not recommendable' based on nutritional information. Several characteristics, such as the amount of fat, sugar, and calories, can be taken into account. Multi-level rule systems that can perform such classifications are called decision trees. Such decision trees can be created based on data. In this case, data means a set of foods for which nutritional information is given and the target attribute (rather recommendable vs. rather not recommendable) is known. Based on this, users can manually create decision trees step by step that classify the food items with a decreasing misclassification rate for every added step. This creation process can also be automated to find optimal decision rules according to specific criteria. Automation requires representing each food item digitally as a 'data card' — that is, a list of numerical values related to the various nutritional characteristics. A machine learning algorithm then develops a decision tree for this data. In practice, other types of classifiers, e.g. neural networks, are used in addition to decision trees, with machine learning methods adapted to them.

Decision trees have the advantage that they can be understood by students as a system of rules, and the procedures for creating a tree can first be worked out manually with unplugged material and then automated on the computer. In class, food items are represented as physical data cards (see Figure 2) and students can sort and classify the cards to understand the process of creating data-based decision trees on an enactive level. The goal is to gain behind-the-scenes insight into a machine learning algorithm and not just to train classifiers with given

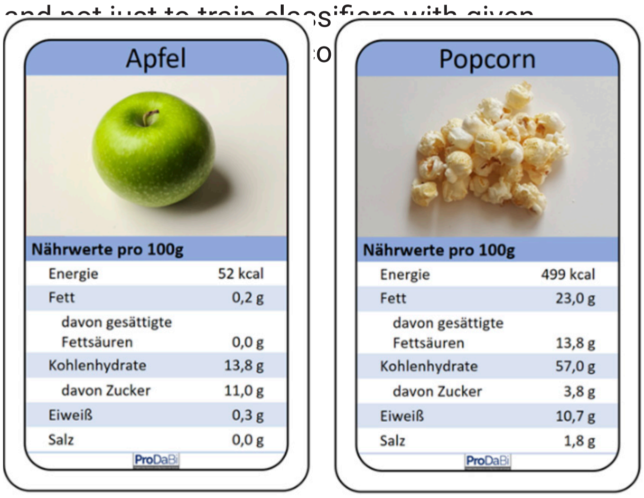

Figure 2. Examples of data cards.

This series of lessons consists of about nine lessons. First, students prepare the data by labelling data cards as 'rather recommendable' or 'rather not recommendable'. The goal within the lesson series is to create a multi-level rule system for classifying food items. The students first learn to derive decision rules (single-level decision trees) from the data. This is done with the concept of data split, where the data cards are split into two subgroups based on a characteristic and a so-called threshold value (e.g., food with up to 10g of fat or over 10g of fat). In both groups, the majority value is used as the choice of class for food items with similar conditions. The students first learn this concept in a setting of statistics with embodied activities, and then the students use it in small groups with their own set of data cards. The students work out how to systematically search for good decision rules. It becomes obvious that a multi-level rule system is needed. Therefore, based on the first rule, more features are included to create decision rules in the second level of the tree. Depending on how fast the students work, they can create two-level or multi-level decision trees. After different groups of students have created different decision trees, these trees are applied to new food items that the students themselves have created on blank cards. These new food items are classified using all the trees. This makes it obvious that there are also uncertainties in the decision trees, as some food items are classified differently by different trees. In order to systematically investigate the uncertainties in decision trees, each group tests its decision tree with the 15 test cards that are marked as yellow cards. This makes it possible to compare the performance of the decision trees. After carrying out the whole process manually, students can use a prepared menu-based environment in a Jupyter Notebook, for example, to automatically create a decision tree on food data using a computer. They can also change the data in the process and observe the effects on the decision tree. Finally, students reflect on how decision trees are created from data with the help of the computer, what advantages and disadvantages this has, and where the students find such decision models in their everyday lives.

Data awareness in the mobile phone network

This example aims to foster students' data awareness, which means to be aware of the collection of personal data, and its usage and processing for various purposes, during interaction with data-driven digital artefacts. Students should be enabled to be aware of the role of personal data during interactions with data-driven digital artefacts in their everyday lives. This should help students to assess the possibilities, implications, and mutual influences of interaction with a data-driven digital artefact.

This teaching unit for middle school students lasts for about four lessons (45 minutes each), consists of three parts, and addresses the mobile phone network as an example system that collects and processes location data during interaction with it. The unit also connects the students' insights to further examples from students' everyday lives.

In the first part, the context of the mobile phone network is introduced and its composition and inner workings are examined using the example of making a mobile phone call. Thereby, the students also identify which personal data is collected, and what it is primarily used for in this context. For example, the location data of the base station of the mobile phone network to which the user is connected. This location data is necessary to ensure the efficient establishment of a connection between mobile phones — we call this the primary purpose for using and processing the collected data.

In the second part, the students are given some data that was collected by the mobile phone network and published by a German politician (he tried to draw attention to the role of such data). These real-world data include location data collected by using the mobile phone network, e.g., during calls or when texting a mobile phone, while browsing the internet with the phone, or just from the fact that the phone is logged into the mobile phone network. We call this collection of location data by the mobile phone network an implicit data collection. We developed a web application with which the students can explore these location data. They are set the task of finding out as much information about a person as they can. So, they create a profile or characterisation of a person that the students did not know before, simply by exploring the person's location data. By doing so, students gain some interesting insights about the person, for example, about their leisure activities, or finding out where the person lives or works (Höper et al., 2021). While discussing the profiles created by the students, it becomes apparent why such profiling is regulated by laws in many countries (especially in Europe). The students can then argue for such reasons in a more meaningful way because they have experienced an example of what one can conclude from such data.

In the third part, the insights about the collection and processing of location data during interaction with the mobile phone network are transferred and applied to other data-driven digital artefacts in students' everyday lives. Consequently, the students generalise the insights and examine other data-driven digital artefacts that also collect location data, such as various apps on their smartphones, including those that collect GPS data. During an evaluation and assessment of the collection and processing of location data in various contexts, the advantages and disadvantages of the collection and processing of location data can be discussed. This will foster students' skills for reflective decisions regarding the release of their personal data during everyday interactions with data-driven digital artefacts.

Conclusion

AI education requires developing an adequate picture of the hybrid interaction system — a kind of data-driven, emergent ecosystem that needs to be made explicit to understand its transformative role as well as the technological basics of these AI tools and how they are related to data science.

Interacting with digital artefacts, especially data-driven applications, is often done within a social context, with aims or tasks a human has to or wants to reach and complete. While the technical system, by automated processing, helps the human to do so, the question also arises as to which aspects are automated, and what range of possibilities to act and to decide are transferred to the machine, and which are still within the direct control and responsibility of the human. One can use the terms shaping or being shaped, or program or be programmed, to refer to this fundamental issue. It can also be related to the role of the machine. Is it to form a human–machine symbiosis? Is the machinery's purpose to replace humans, or to augment?

References

Bell, T., & Duncan, C. (2018). Teaching computing in primary schools. In S. Sentance, E. Barendsen, & C. Schulte (eds.). *Computer Science Education: Perspectives on teaching and learning in school* (pp. 131–150). Bloomsbury Academic. http://dx.doi.org/10.5040/9781350057142.ch-010

Clark, M. A. C., & Boyle, R. D. (2006, November). Computer science in English high schools: We lost the S, now the C is going. In *International Conference on Informatics in Secondary Schools-Evolution and Perspectives* (pp. 83–93). Springer. https://doi.org/10.1007/11915355_8

Curzon, P., & McOwan, P. W. (2015). *The Sweet Learning Computer*. cs4fn. Retrieved March 26, 2019, from http://www.cs4fn.org/machinelearning/sweetlearningcomputer.php

Curzon, P., McOwan, P. W., & Black, J. (2008). *Artificial intelligence ... but where is the intelligence?* cs4fn. www.cs4fn.org/ai/downloads/aiwhereistheintelligence.pdf

Große-Bolting, G., & Muhling, A. (2020, April). Students perception of the inner workings of learning machines. In *LATICE Seventh International Conference on Learning and Teaching in Computing and Engineering*. https://www.researchgate.net/profile/Gregor-Grosse-Boelting/publication/356253786_Students_Perception_of_the_Inner_Workings_of_Learning_Machines/links/61938fca3068c54fa5eebfac/Students-Perception-of-the-Inner-Workings-of-Learning-Machines.pdf

Höper, L., Podworny, S., Schulte, C., & Frischemeier, D. (2021). Exploration of location data: Real data in the context of interaction with a cellular network. In R. Helenius, & E. Falck (Eds.). *Statistics education in the era of Data Science. Proceedings of the satellite conference of the International Association for Statistical Education.* https://doi.org/10.52041/iase.nkppy

O'Neil, C. (2017). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Penguin.

Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bonnefon, J. F., Breazeal, C., Crandall, J. W., Christakis, N. A., Couzin, I. D., Jackson, M. O., Jennings, N. R., Kamar, E., Kloumann, I. M., Larochelle, H., Lazer, D., McElreath, R., Mislove, A., Parkes, D. C., Pentland, A. S., Roberts, M. E., Shariff, A., Tenenbaum, J. B., & Wellman, M. (2019). Machine behaviour. *Nature, 568*(7753), 477–486. https://doi.org/10.1038/s41586-019-1138-y

The Royal Society. (2017). *Machine learning: the power and promise of computers that learn by example*. https://royalsociety.org/-/media/policy/projects/machine-learning/publications/machine-learning-report.pdf

Rushkoff, D. (2010). *Program or be programmed: Ten commands for a digital age*. Or Books. https://doi.org/10.2307/j.ctt207g7rj

Schulte, C., & Budde, L. (2018, November). A Framework for Computing Education: Hybrid Interaction System: The need for a bigger picture in computing education. In *Proceedings of the 18th Koli Calling International Conference on Computing Education Research* (pp. 1–10). https://doi.org/10.1145/3279720.3279733

Thijs, A., & Van Den Akker, J. (2009). *Curriculum in development*. Netherlands Institute for Curriculum Development (SLO). https://www.academia.edu/5132063/Curriculum_in_development_1_?email_work_card=title

Wing, J. M. (2006). Computational thinking. *Communications of the ACM, 49*(3), 33–35. https://doi.org/10.1145/1118178.1118215

Carsten Schulte (Paderborn University, Germany)

Dr Carsten Schulte is a professor of computing education research at Paderborn University, Germany. His work and research interests are the philosophy of computing education and empirical research into teaching-learning processes (including eye movement research). Since 2017, he has been working together with Didactics of Mathematics (Paderborn University) in the ProDaBi project, in which Data Science and Artificial Intelligence are prepared as teaching topics. He is also PI in the collaborative research centre 'Constructing Explainability' on explainable AI.

Yannik Fleischer (Paderborn University, Germany)

Yannik Fleischer is a PhD student in mathematics education research at Paderborn University, Germany. His main research interest is to develop a concept to teach machine learning methods in school with a focus on decision trees, and to evaluate this by developing and examining teaching materials in practice. Since 2019, he has been supervising year-long project courses on data science in upper secondary and developing, implementing, and evaluating teaching modules for different levels in secondary school, mainly about machine learning with decision trees.

Lukas Höper (Paderborn University, Germany)

Lukas Höper is a PhD student in computing education research at Paderborn University, Germany. His main research interest is to develop the concept of data awareness for computing education and evaluate this by developing and examining teaching materials in practice. Since 2020, he has been working on data awareness in the ProDaBi project, among other topics on AI and Data Science in schools.

Raspberry Pi
Foundation