

Case Report

The AI-Atlas: Didactics for Teaching AI and Machine Learning On-Site, Online, and Hybrid

Thilo Stadelmann ^{1,*},[†] , Julian Keuzenkamp ², Helmut Grabner ³ and Christoph Würsch ⁴ 

¹ CAI Centre for Artificial Intelligence, ZHAW Zurich University of Applied Sciences, Obere Kirchgasse 2, 8400 Winterthur, Switzerland

² ZHAW Digital, ZHAW Zurich University of Applied Sciences, Gertrudstrasse 15, 8400 Winterthur, Switzerland; julian.keuzenkamp@zhaw.ch

³ IDP Institute of Data Analysis and Process Design, ZHAW Zurich University of Applied Sciences, Rosenstrasse 3, 8400 Winterthur, Switzerland; helmut.grabner@zhaw.ch

⁴ ICE Institute for Computational Engineering, OST Eastern Switzerland University of Applied Sciences, Werdenbergstrasse 4, 9471 Buchs, Switzerland; christoph.wuersch@ost.ch

* Correspondence: thilo.stadelmann@zhaw.ch

† ECLT European Centre for Living Technology, 30123 Venice, Italy.

Abstract: We present the “AI-Atlas” didactic concept as a coherent set of best practices for teaching Artificial Intelligence (AI) and Machine Learning (ML) to a technical audience in tertiary education, and report on its implementation and evaluation within a design-based research framework and two actual courses: an introduction to AI within the final year of an undergraduate computer science program, as well as an introduction to ML within an interdisciplinary graduate program in engineering. The concept was developed in reaction to the recent AI surge and corresponding demand for foundational teaching on the subject to a broad and diverse audience, with on-site teaching of small classes in mind and designed to build on the specific strengths in motivational public speaking of the lecturers. The research question and focus of our evaluation is to what extent the concept serves this purpose, specifically taking into account the necessary but unforeseen transfer to ongoing hybrid and fully online teaching since March 2020 due to the COVID-19 pandemic. Our contribution is two-fold: besides (i) presenting a general didactic concept for tertiary engineering education in AI and ML, ready for adoption, we (ii) draw conclusions from the comparison of qualitative student evaluations ($n = 24\text{--}30$) and quantitative exam results ($n = 62\text{--}113$) of two full semesters under pandemic conditions with the result of previous years (participants from Zurich, Switzerland). This yields specific recommendations for the adoption of any technical curriculum under flexible teaching conditions—be it on-site, hybrid, or online.

Keywords: flexible educational design; e-learning; constructivism; design-based research; COVID-19; post-pandemic tertiary engineering education; artificial intelligence



Citation: Stadelmann, T.; Keuzenkamp, J.; Grabner, H.; Würsch, C. The AI-Atlas: Didactics for Teaching AI and Machine Learning On-Site, Online, and Hybrid. *Educ. Sci.* **2021**, *11*, 318. <https://doi.org/10.3390/educsci11070318>

Academic Editor: Georgios N. Kouziokas

Received: 31 March 2021

Accepted: 22 June 2021

Published: 25 June 2021

Publisher’s Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Analysis and Exploration

1.1. The Problem of Teaching Artificial Intelligence as a Foundational Subject

Students enter their careers in a time that awaits nothing short of a digital disruption [1]. The core of the disruptive potential of the digital transformation is provided by technological developments, foremost by Artificial Intelligence (AI) and its currently most prominent branch, Machine Learning (ML) [2]. While technical in nature, AI and ML have the potential to disrupt society on all levels, including business, public service, justice, art, science and health [3]. Numerous scientific articles are published daily [4], and coverage in popular media contributes to these technological developments not just a scientific, but also a mainstream hype [5]. The hype leads to misconceptions about what the technology is capable of, or designed to do, including that it could be concerned with human intelligence or creating conscious machines [6] (while the reality boils down to

mere “complex problem solving” [7]). This exemplary problem highlights the necessity of a foundational understanding of AI for engineers in two respects: (i) the understanding needs to be solid, i.e., rooted in the foundations of the discipline rather than purely in its applications, tools or business value, to manage expectations; (ii) like algebra and other foundational subjects, the understanding of AI needs to serve as an underlying methodological framework for an increasing number of future engineering efforts, i.e., it needs to serve as a foundation for future careers in engineering rather than merely as an interesting fashionable specialization.

We think of AI and its sub-discipline ML as one of the five pillars of the discipline of computer science (besides theoretical, practical, technical, and applied computer science [8]). As such, it requires similar treatment to other foundational subjects (e.g., thorough establishment of basic ideas and theory), but holds a peculiarity: different from foundations, such as algebra, AI and ML already build upon a body of knowledge from computer science. This implies a later slot in respective programs, and hence more experienced students who can better judge the impact of AI on other aspects of their profession or society as a whole. However, teaching AI as a foundational aspect of a technical education was until recently largely neglected in curricula according to surveys [9], and novel data science courses [10] do not cover the same terrain.

We, as educators and researchers in this context, thus, are faced with the problem of having to prepare experienced students with diverse professional backgrounds for the reality of AI and ML applications and their implications. For this, we need university courses that provide a solid theoretical foundation while providing opportunities and a mind-set for applied life-long-learning, especially in tertiary engineering education at universities of applied science. We need to teach solid AI and ML foundations to this diverse, technical audience while the field is rapidly evolving, students may have initial misconceptions, and the outcome has to be practice-oriented and relevant for a broad range of careers. In short: we need a didactic concept to solve for this problem by providing a powerful heuristic—a set of coherent best practices—for curriculum and didactic planning and implementation. While many courses exist, and new ones are rapidly created to keep up with the current surge, to our knowledge, no such didactic concept exists explicitly. Earlier AI didactics, like References [11,12], focused only on special teaching situations or one of the sub-discipline of AI.

In this paper, we report on the iterative design of two courses based on the “AI-Atlas” didactic concept. We discuss how the concept itself was developed in specific response to the aforementioned challenges induced by the current surge of AI, and how the resulting course design was adapted based on student feedback and the rapidly changing context of COVID-19. Specifically, our case report is based on the structure of Design-Based Research (DBR) [13] to present the AI-Atlas (didactic concept) and two exemplary implementations (concrete courses), and evaluate the merits of the concept based on student results and feedback. The courses, “Artificial Intelligence 1” (referred to as the “AI course” below) for undergraduate students of computer science, and “Machine Learning” (“ML course”) within an interdisciplinary graduate program in engineering, are described in detail in Appendix A. The atlas (or map) metaphor itself that gave rise to the AI-Atlas concept and its underlying design choices is introduced in Section 2.1. We chose to present our findings as a case-report rather than a research article because there are certain limitations regarding our evaluation data as explained in Section 3.2. Our intention is, therefore, to introduce our didactic concept and to motivate educators to test it in order to gather more experience and evidence for the suitability of the AI-Atlas as a flexible basis for designing a curriculum for AI and ML courses in tertiary education.

1.2. Existing AI Curricula and Their Relation to Our Didactic Concept

It is beyond the scope of this case report to present a comprehensive survey of related literature on AI and ML teaching methodology. Nevertheless, the AI-Atlas fits well within the recent discussion on AI curricula: like Williams et al. [14], constructionist ideas are

central to our didactic design as discussed in Section 2.3. Moreover, the origins of the AI-Atlas [15,16] evolved from teaching best practices of the involved lecturers, putting this case report in line with similar post-hoc analyses, for example, References [17–19].

Chiu and Chai [20] reflect AI curriculum development for K-12 [21] by teachers with and without prior exposure to AI training. They use the psychological framework of Self-Determination Theory (SDT) to understand how teachers' motivation controls curriculum design. Further, the authors use the four planning paradigms of curriculum as either content, product, process, or practice to distinguish fundamentally different ways to think about course development and respective didactic designs. While SDT can also be used to explain our motivation for the specific design of the AI-Atlas, our concept can be classified as being a blend of content (the syllabus is important to us), product (we care about achieved educational objective), and practice (we frequently connect the learning experience to real-world applications); see Section 2.5.

Finally, Li et al. [22] explore the fit between university AI curricula and the demands of the job market. Our design with its focus on professional applicability can be seen as a solution to the problems they identify within their analyzed courses: What distinguishes our concept from the many other adoptions of, e.g., the AIMA textbook (cf. Reference [23]), is an end-to-end focus on the usefulness of the taught foundations in daily application that should lead to successful transfer into personal problem-solving skills and professional practice. This implies to be less focused on covering a large amount of different AI/ML algorithms, but rather teach the underlying relationships. Of course, practical relevance is a major concern for any course at a university of applied sciences. While we cannot compare the AI-Atlas to a previous AI course at our institutions (it was developed for the initial design of the first AI and ML courses here), we notice that our approach takes an atypical route to this destination: most courses are made more "applied" by stripping them of theoretical concepts. However, we put theoretical foundations at the center. The reasoning for this seemingly uncommon choice follows below.

2. Design and Construction

In this section, we explain why we named our didactic concept the AI-Atlas and what the core mentality is it means to convey. We also look at three didactic principles that guided us in its design and three aspects about AI and ML teaching we believe it needs to cover in order to solve the above stated problem before breaking them down into the suggestion of specific didactic settings. Together, the principles and ideas layed out in the following five subsections constitute the AI-Atlas didactic concept.

2.1. The Atlas Metaphor

In the late 16th century, Gerhard Mercator's "Atlas Sive Cosmographicae Meditationes de Fabrica Mundi et Fabricati Figura" combined maps and associated explanations of the known world [24]. They were used by generations to explore, push boundaries, and further trade and development [25]. For an atoll, for example, they would show all individual islands with their border, list their characteristics, and show their relation to each other. This aids travelers by allowing them to plan the most effective or efficient route based on their current needs or interest. However, it does not set an pre-conceived path and, if a new island is discovered, it can be added to the atlas without disrupting existing knowledge.

Within the AI-Atlas, we think of the sub-disciplines of AI—such as, e.g., search, planning, and machine learning—as individual islands of an atoll: well developed in themselves and somewhat related to each other, but missing a direct connection. Hence, to solve the problem we have identified, future generations of professionals need the analogue to what Mercator gave to his contemporaries: an atlas to the world of AI (cf. Figure 1). An explorer can profit from the help of this atlas to get an overview and find the best path given a specific journey (i.e., application). It highlights main routes (i.e., baseline approaches), special landmarks (e.g., important algorithms, killer applications), and borders (i.e., limits

of the approaches, future work) but never restricts learners to knowing or using only a single path.



Figure 1. Illustration of the AI-Atlas metaphor: the sub-disciplines of AI are like an atoll that is best navigated with the help of an atlas. The AI-Atlas didactic concept provides the means, to be employed by educators, to let such atlases emerge in their learner’s minds. Image credit: Colorbox, used with permission.

Our AI-Atlas didactic concept contains *the means* to let the actual atlas emerge only in a learner’s mind. It is, thus, created by analog means and stored in analog form (in natural neural networks) and is not manifested in some digital format on the learner’s computer. This aspect of the metaphor underlines the AI philosophy underlying the didactic concept and derived courses: artificial intelligence is not primarily replacing human intelligence, and machine learning does not render human learning unnecessary, just like digital does not primarily replace analog, but augments it [26]. AI, thus, finds an optimal environment for application where human and machine complement each other with their respective strengths and weaknesses.

2.2. *The Core Mentality: AI Professionals Are Explorers*

The discipline of AI and its major tool of ML do not have a single goal (“creating intelligence”), but rather offer a methodological toolbox to approach multiple targets (“solving complex problems”) [7]. Thus, at their core, they are not constituted by technology as much as by a specific mentality: since AI’s inception as a discipline in the 1950s, AI researchers notoriously approached the kind of problems with creativity and pragmatism that had been laid aside by fellow researchers from other disciplines as “too hard” [27]. In other words, AI researchers explored previously unknown territories. They did so by employing an interdisciplinary “let’s do it” mentality. Today, this mentality distinguishes the work of the AI professional from other modeling approaches used by software engineers, database designers or statisticians, although skills in all these areas are relevant for success in and with AI or ML. The AI-Atlas not only acknowledges but actively hones this explorer mentality [28]. It does so by building on a set of corresponding didactic principles.

2.3. *Didactic Principles*

2.3.1. Principle of “Doing It Yourself”

Since the late 1980s, constructivist ideas have increasingly found their way into pedagogy, as well as into discussions on the design of teaching-learning environments. Constructivism postulates that individuals do not take over information faithfully from external sources, but actively construct it through social negotiation processes based on previously made experiences. Knowledge is situation-specific and must be actively and independently linked by the individual to prior knowledge [29]. It follows that teaching-

learning settings must be designed in such a way that learners are given the opportunity to *actively engage* with the learning content, as well as the associated problems and to relate these to their prior experiences, whereby active engagement can be of both visible and non-visible nature.

2.3.2. Principle of “Intrinsic Motivation”

In their Self-Determination Theory of motivation (SDT), Deci and Ryan [30] explain the relationship between motivation, learning, and the influence of the social environment on the fulfillment of basic needs. Intrinsically motivated behavior is associated with individuals freely engaging with the subject matter and striving of their own accord to understand phenomena and master activities that appear to them to be at least personally highly significant. As a further component of their theory, Deci and Ryan postulate three basic human needs that motivate behavior: (i) the need to *experience competence*, (ii) the need for *social inclusion*, and (iii) the need for *autonomy*. Deci and Ryan now assume that the striving for satisfaction of these three basic needs explains why persons pursue certain goals of action and why certain actions are more likely to be perceived as motivating by themselves.

2.3.3. Principle of “EEE” (Good Explanation, Enthusiasm, and Empathy)

According to Winteler [31], the following characteristics of university teaching promote student learning: the instructor’s preparation, organization of the course, clarity and comprehensibility, perceived efficiency of the teaching, the instructor’s openness to questions and discussion, and openness to other opinions. Helmke and Schrader [32] reduce the state of research on key characteristics of successful university teaching to the short formula “EEE”: (i) good *explanation*, which facilitates information processing and arouses curiosity and factual interest; (ii) commitment and *enthusiasm*, i.e., the infectious enthusiasm of the lecturer about the content; and (iii) *empathy*, by which they include personal appreciation of students, openness to problems, and efforts to obtain feedback to better adapt teaching. The fact that the didactic setup of a course is well planned and fine-tuned is a “*conditio sine qua non*”, i.e., a necessary but not sufficient condition for a successful course. Nevertheless, a committed, authentic (i.e., experienced in the field), and enthusiastic teacher, open for questions and igniting curiosity and factual interest, can bring the majority of the students to engage in the topic and start the learning process in a self-motivated manner.

2.4. The Aspects of Establishing AI Foundations

The above didactic principles need to be combined with the proper mediation of AI and ML foundations, if the AI-Atlas is to successfully guide our explorers-in-training. We suggest that there are three aspects to which the principles need to be applied: *canonization*, *deconstruction*, and *cross-linkage*.

2.4.1. Canonization

The aforementioned hype [5] around AI, and especially deep learning, and the daily growth of scientific literature on the topic [4] make a proper selection of content a key aspect of teaching AI and ML. Hence, a key aspect of the AI-Atlas is to suggest a timely selection of materials that emphasize topics with future relevance alongside their historic development, thereby making the overarching principles that stood the test of time stand out. This is given priority over intriguing detail or formal derivations. Thus, for, e.g., a specific implementation of the AI-Atlas for an introduction to AI course, canonization means that we ensure to teach the full canon of relevant methods (ranging from heuristic search and logical planning to machine learning). We link each of these areas with practice (e.g., controlling a fashionable browser game, building a dragnet investigation system, or decision support for second-hand vehicles). This way, students see for themselves that not only the currently most fashionable methodology, ML, and, within ML not only neural networks, have practical relevance.

2.4.2. Deconstruction

Due to the current extensive media coverage of AI and ML, many misconceptions about the field abound in prospective students (such as the focus of the field being the understanding of human intelligence or the creation of conscious machines [6]). Thus, an important aspect of teaching according to the AI-Atlas is a form of demystification that keeps the original motivation of the students and channels it into more realistic, sustainable paths. We suggest to support students with forming a personal view through critical engagement with scientific texts and programming tasks, which they then present in own write-ups and oral discussions.

2.4.3. Cross-Linkage

Both aspects above—a stable body of knowledge in AI and ML fundamentals and careful treatise of real and misguided excitement—become a firm foundation given the third ingredient: a dense network of cross-references to other subjects in the study program that is compatible with the different occupations of a professional career in engineering and related fields. The AI-Atlas suggests to teach AI not only to future scientists but also to software developers, data scientists, or process engineers, acknowledging the future importance of AI methodology in any field.

2.5. *The AI-Atlas in Practice: Suggested Didactic Settings to Combine It All*

Building on the principles described above, we suggest to adopt the following specific didactic settings for any AI or ML course facing the problems outlined in Section 1.1. Section 3 then continues to evaluate to which extent employing these means in the two exemplary courses mentioned above achieves these ends. Nevertheless, the following subsections already make frequent references to examples in the AI course and ML course to put the abstract suggestions into concrete terms.

2.5.1. Basic Didactic Setting

As laid out in the theoretical framework above, active engagement is a mandatory key component in an AI-Atlas inspired course. One way to increase engagement is to work with small to medium classes. For example, both exemplary courses had only 30 to 60 students. Courses should, therefore, build upon the “lecture + lab” format widespread in engineering education: weekly lectures accompanied by lab exercises with a roughly even time-split between them. However, we make important adaptations to foster active engagement as follows.

For lectures, the students should read weekly portions of well-established text books as accompaniment to the lectures, e.g., Reference [33] on AI and References [34,35] on ML, complemented where necessary by shorter articles (e.g., References [26,36]). The conveyed anecdotes and historical notes therein specifically contribute to the students’ socialization in the discipline of AI and the field of ML. The lectures themselves connect such context with problem awareness and technical solutions without degenerating into pure 90-min talks that would push learners into passive consuming roles (cf. Section 2.5.5).

Labs on the other hand should go beyond just programming and development, to accommodate essay writing or examining philosophical questions. This way, the AI-Atlas ensures educational objectives for professional and methodical competences on levels K1–K4 [37] by presenting AI and ML as socio-technical and not purely technological. One example of how broader theory can be consolidated by practice are the gamification elements provided in one lab of the AI course (cf. Appendix B). In addition, programming skills are only a means to an end in AI and ML labs, while problem analysis and experimentation become the focus, thereby encouraging exploration.

2.5.2. Fostering Reflection

We suggest to repeatedly ask students to reflect on their preconceptions of the course content and the technical and societal ramifications of this prior knowledge. Making them

think about the myths and ethical problems of the application of AI and learning algorithms starts an active, though non-visible, process in each individual. For example, in the context of the ML course, we repeatedly highlight the cognitive dissonance between the current focus on deep learning methodology in public opinion and the irrefutable results of the “no free lunch” theorem [38]. In the AI course, a lab assignment at the semester start asks students to create a blog post that presents their well-founded and reflected opinion on the contents of a futuristic essay [6]. At the end of the semester, students can reflect on their initial statement with a second blog post that may incorporate insights gained throughout the course. While all opinions are welcome, the emphasis in grading is on self-reflection and reasoning.

As a more regular intervention, in the AI course, lectures end with an outlook called “where’s the intelligence?”. It explains why what was discussed is a “clever solution”, but also what separates it from human-like intelligence. In the ML course, the same time slot is used to show state-of-the-art implementations of the discussed material. This not only aids to demystify the technology; it also helps the students spot the kind of tasks they might approach in their future job using the conveyed foundations.

2.5.3. Encouraging Self-Responsibility and Motivation

Up to twenty percent of the final grade should be acquired by each student during the semester through self-chosen lab assignments, with results depending on an oral defense of one’s own work. The choice of 20% is justified by allotting, in this way, a substantial part of the final result to active participation throughout the semester without replacing competence-based results with effort-based grading.

From the existing six assignments in the AI course, for example, that are distributed evenly throughout the semester, students can choose any two to get graded within a short colloquium between the student team and lecturer during the in-class time (students usually work on all assignments, but put considerably more time into the two graded ones). This way, students get empowered to prioritize own learning goals and take ownership of their investment of time and its distribution over the semester. Even if grading is not explicitly tied to the lab assignments as in the ML course, the further questions presented there, the inquiring of the lecturers, and not least the relevance of practical implementation skills for the final exams motivate students to work deeply on the assignments, even if sample solutions are freely provided.

A second method to encourage self-exploration and motivation is to set up the labs in such a fashion that it requires students to independently dive deeper into respective methods to find practical solutions. The lab descriptions of the ML course, for example, actively encouraged this, and the lab exercises are often not solvable without going beyond the lecture content.

2.5.4. Promoting Cooperative Competence Development

Lab assignments should usually be worked on in teams of two students. This way, students can strategically pair up their existing competencies, as well as learn from each other. Teams should be allowed to help each other as long as any help is disclosed (according to good scientific practice), and competitive elements, such as the public leaderboard for the AI lab assignment presented in Appendix B, only increase the appeal of and the necessity for good team work.

2.5.5. Activation of Students

Each 90-min lecture block should contain a part of up to 30 min that assigns an active role to the students rather than the lecturer. Technical understanding is deepened by embedding interactive parts, such as small group research tasks and discussion, as well as thinking and pen-and-paper exercises, thereby increasing the practical treatment of the subject. For example, in the AI course, a classical brain twister [39] can be used to show the difference between AI (having a computer program *appear* intelligent) and human

intelligence: approaching it by efficient search through all combinations of possible solution steps constitutes an excellent AI solution for that problem but typically gets labeled “just brute force” by the students at first sight. Other activations in the two exemplary courses take on the forms of jointly solving a puzzle (e.g., “escape from the Wumpus world”), computing results in small groups (e.g., “help inspector Clouseau to probabilistically convict a murderer”), individually applying learned principles (e.g., logic training), or sharing insights from individual research at tables (e.g., exploration of possibilities with OpenAI Gym [40]).

2.5.6. Enabling Social Learning

A prominent place throughout a course based on the AI-Atlas concept should be given to the research work and careers of course alumni and junior teaching staff. Linking course content to concrete outcomes of applied research projects with regional industrial partners known by the students creates a pull that contributes to the students’ motivation and expanded vision for AI and ML in practice, as well as their role in it. Key to create this are tutors (e.g., graduate students) that teach part of the labs: closer in age and role to the course participants, they are in our experience approached frequently by the class to give a second opinion on the more philosophical and career aspects of AI. Innumerable lunches, coffee invitations and after-work beers have been realized this way between teaching staff and students in the AI course and ML course.

2.5.7. Providing Open Educational Resources (OER) and Blended Learning

All course materials, including lecture recordings, slides, and lab materials, should be fully and freely available online [41]. This should enable flexible deepened learning (e.g., for exam preparation), but does not compromise live lecture attendance in our experience. Students can also recap all details when needed on the job as all material is permanently and openly available. This enhances the *atlas* of AI and ML solution strategies they know by heart. As an add-on, it supports the transition to live online teaching (as was required during the COVID-19-induced lockdowns in 2020, see Section 3.2) as content is already designed to be streaming-friendly.

2.5.8. Creating Practical Career Relevance

Students’ diverse professional backgrounds should be addressed by showing how different AI and ML methods serve as puzzle pieces in numerous everyday situations of (software) engineering. For example, in the AI course and ML course, the lab tasks and in-class exercises are strictly sourced from practical applications, such as automatic university time tabling, biometric access control, or data analysis, to reinforce this point. By connecting the practical coursework with typical tasks of an engineer, programmer or consultant, students clearly see how learning foundations of AI and ML makes them better in their original career goal. By confronting them with new opportunities in and through AI and learning algorithms in business and research, they recognize new and viable career paths (e.g., data scientist [10]) that only begin to gain traction in public awareness.

Additionally, we suggest to invite specialists, ideally course alumni, from regional industrial partners for guest lectures to report on recent successes. Learners should be encouraged to actively use these opportunities to network and engage with those speakers and their ideas. In contrast to the culturally typical reticence of engineering students, the fresh setup with people on stage that might be considered peers age-wise opens them up in the direct (active participation) and metaphorical sense (opening up to the idea of other career options within the fields of AI and ML).

3. Implementation and Analysis

In this section, we evaluate to what extent the different measures advised by the AI-Atlas didactic concept, as implemented within our two exemplary courses, achieve their aspired goals of contributing to providing foundational AI education in times of a height-

ened AI surge. By this, we aim to shed light in a post-hoc fashion on the merits of the AI-Atlas concept per se to serve as a basis for designing future courses for similar needs.

3.1. Participants and Data Basis for Analysis

The principles and settings laid out in the AI-Atlas didactic concept emerged in parallel to designing our exemplary courses out of teaching experience and got codified only upon their completion. The presented analysis, thus, is necessarily a post-hoc analysis: no baseline data from before the AI-Atlas is available due to no existing previous courses on the subject at the involved universities. Furthermore, the analysis is purely based on routinely available data for any course: qualitative student feedback (Likert-scaled, as well as free-text comments) from surveys conducted by the central program administration, and quantitative results from end-of-semester exams (points and graded per task and overall). This is also due to the evolving fashion of our design that had not been planned as a study from the beginning. It means, inter alia, that no control group is available. While this form of evaluation leaves certain aspects to be desired, it is not an uncommon situation for data scientists to having to work with the data that is available in the best possible way, without the possibility to change the basis for evaluations [42]. In what follows, thus, we evaluate our concept with the mindset of data scientists, aiming to establish a relationship between learning success and the employed measures from the AI-Atlas concept.

Demographic and background information on the participants in our courses is listed in Appendix A, while information on specific class sizes (or number of students who returned a questionnaire) will be listed in the captions of the figures below that deal with them. For the ML course, most of these students predominantly seek a career in their original fields of study, though a growing minority considers a job related to ML engineer or ML researcher (the possibility to, e.g., take up graduate studies, is typically completely unknown to our students due to the setup of a “Fachhochschule” [43]). Our students of computer science in the AI course usually envisage a career in software engineering, not specifically AI.

We will be using the qualitative student feedback in Section 3.3 and exam-data in Section 3.4 to evaluate whether the design choices of the AI-Atlas led to a appreciable impact. We do this for two iterations: The first iteration was executed with on-site teaching methods as anticipated by the original AI-Atlas concept (data basis: qualitative and quantitative data from fall term 2019 for the AI course; qualitative data from spring term 2019 for the ML course). The second, due to the COVID-19 pandemic, was executed using hybrid and online teaching, which was not specifically anticipated in the design of the AI-Atlas (data basis: quantitative data from fall term 2020 for the AI course; qualitative data from fall term 2020 and quantitative data from spring and fall terms 2020 for the ML course). The context for the second iteration, the natural experiment that the COVID-19 pandemic provided us with, is described below in Section 3.2. For reasons explained there, we organize our data in the remaining sections as follows: the qualitative feedback is combined for both iterations and evaluated per didactic principle or setting from Section 2. The exam-data is presented specifically per term to allow for comparison between the two iterations.

3.2. Going Online by Necessity

The AI-Atlas was designed specifically for the on-site teaching of small to medium groups (30 to 60 students), but the COVID-19 pandemic forced its execution as hybrid or fully online teaching for two full semesters. Of course, the rapid digitization of higher education in the wake of the COVID-19 pandemic forced teaching around the world to move online in a matter of days. Good teaching methodology differs whether one is going to teach on-site or online [44,45], but the new didactic credo seems to be the one of flexibility: one year into the pandemic, there is still a huge degree of volatility about the possibility, desirability, and potential timeline of returning to a teaching mode of choice, in many countries.

Thus, it is important to know *how* a course specifically designed for, e.g., on-site teaching, will perform in a hybrid or online setting in terms of the students reaching that course's educational objectives. This transcends the question of whether going online is merely possible at all [46], as the didactic concept and respective teaching material cannot be adapted that quickly. The move away from on-site teaching was done very rapidly and involuntarily, with the side-effect of no planned, controlled data collection on student learning before and after. Furthermore, as the courses are single-semester, we also do not have longitudinal data within the same cohort of students. Thus, we saw ourselves forced to stray from DBR principles and to (i) combine the qualitative data with (ii) use the quantitative exam data to specifically compare between the effectiveness of the AI-Atlas for on-site and online teaching.

3.3. Qualitative Assessment

In the following subsections, we collect evidence for and against the effectiveness of specific dimensions (i.e., didactic aspects and settings, cf. Sections 2.4 and 2.5) of the AI-Atlas by providing quotations taken from students' feedback forms at the end of different semesters, and drawing conclusions from them. A short tag at the end of each statement ("AI" or "ML") indicates the source course. These written qualitative comments are optional for the students to make and, thus, are normally very sparse (though most precious for the improvement of the curriculum and course). The answers might be highly skewed since we cannot control the subset of students that did write some comments. We, therefore, refrain from a statistical analysis of these comments. Nevertheless, the presented example statements are chosen to be representative to allow for the conclusions we draw. In case of counter evidence, we rather over-sample critical comments to avoid any cherry picking (cf. Section 3.3.5).

3.3.1. Dimensions "Canonization and Deconstruction"

The following statements are taken from student's free and voluntary comments in evaluation surveys:

"Sustainable technologies are taught; in the process you are brought down to earth." (AI)

"[The] module gives a good overview of the overall topic." (ML)

"I welcome that [...] 'where's the intelligence?' is answered [in] each lecture." (AI)

Students seem to grasp that AI and ML are ways to solve complex practical problems rather than theories to explain how we think or create artificial life. The content is indeed perceived as a foundation for practice rather than a narrow specialization (cf. Sections 2.4.1 and 2.4.2).

3.3.2. Dimensions "Motivation and Social Learning"

Quality assessment of the two courses is unfortunately neither done in the same way nor regularly by the involved central administrations (cf. Sections 2.5.3 and 2.5.6). In addition to being given the opportunity for providing free-text comments, the students of the ML course are asked to rate their consent with the following statements on a Likert scale (cf. Table 1 for details): (i) *Motivation*: I regard the lecturer as being motivated and committed; (ii) *Competence*: I regard the lecturer as being competent in their subject; (iii) *Teaching Skills*: I regard the lecturer as having good teaching skills; (iv) *Clear structure*: His/her teaching is clearly structured (a clear thread), and the subject matter was imparted in a comprehensible manner. Table 1 summarizes the evaluation of teaching skills and motivation for the ML course (unfortunately, no evaluation of the ML course was carried out in spring term 2020 due to COVID-19-induced stress in the administration, and the AI course evaluates slightly different questions not pertaining to the dimensions discussed here). The presented score is the averaged score for two lecturers. It mainly reflects the qualitative judgements given by the students also in the free-text comments:

Table 1. Teaching evaluation scores for the ML course. Two teachers were assessed, and the average score is presented (spring 2019: $n = 30$, and fall 2020: $n = 27$).

Topic	Semester	1 (Strongly Disagree)	2 (Disagree to Some Extent)	3 (Agree to Some Extent)	4 (Strongly Agree)	Score
Motivation	Spring 2019	1.7%	0.0%	18.8%	79.5%	3.8
	Fall 2020	0.0%	5.6%	18.5%	75.9%	3.7
Competence	Spring 2019	1.7%	0.0%	15.0%	83.3%	3.8
	Fall 2020	0.0%	1.9%	13.0%	85.1%	3.8
Teaching skills	Spring 2019	1.7%	15.0%	45.0%	38.3%	3.2
	Fall 2020	3.7%	14.8%	38.9%	42.6%	3.2
Clear structure	Spring 2019	3.4%	13.4%	48.3%	34.9%	3.1
	Fall 2020	5.6%	14.8%	44.4%	35.2%	3.1

“The professors [...] enthusiastically explained it very precisely. I also had the feeling that the fun of the topic seemed very important to them. It was also important for them that everyone understood.” (ML)

“The two lecturers are very motivated and they pass on their enthusiasm and experience in the respective field. I find the exercises and tools we use (Jupyter notebooks, scikit-learn, Orange) very useful and they complement the lessons well. I also appreciate that discussions among each other and in plenary are stimulated.” (ML)

“You can feel that the lecturer is convinced of the subject. He also often brings good examples to help the students on their way.” (AI)

“Very good commitment, super presentation style. Enthusiasm for the subject is obvious and motivates me a lot.” (AI)

Students’ perception of the course contents is in our perspective strongly connected with and dependent on the person that teaches. Insofar, the concept, curriculum or OER availability alone is no guarantee for the intended outcome: enthusiastic teaching is an integral part of the AI-Atlas as it facilitates activation.

3.3.3. Dimension “Activation”

Table 2 summarizes evaluations of the AI course as it has specific questions on the perceived activation (cf. Section 2.5.5) of students and the practical relevance of the presented material (unfortunately, no such questions are asked for the ML course in the central questionnaires as they are issued by a different program administration). Specifically, the students were asked to rate their agreement with the following statements: (i) *Activation*: the students are actively involved in the teaching process; (ii) *Practical relevance*: in class, theory is reinforced with examples and applications.

Table 2. Evaluation of the activation in and the practical relevance of the AI course in fall 2019 (unfortunately, no evaluation of the course took place in fall 2020 due to anti-pandemic measures). The number of students who handed in the questionnaire was $n = 24$.

Topic	Semester	1 (Strongly Disagree)	2 (Disagree to Some Extent)	3 (Agree to Some Extent)	4 (Strongly Agree)	Score
Activation	Fall 2019	0.0%	20.8%	29.2%	50.0%	3.3
	Fall 2020					–
Practical relevance	Fall 2019	0.0%	16.7%	25.0%	58.3%	3.4
	Fall 2020					–

These findings are also supported by the following optional written statements that these students handed in:

“The labs support the learning process very much; similarly helpful are the exercises throughout the lectures.” (AI)

“Very handy are the labs where one implements hands-on what should be learned.” (AI)

“Good lecture-style presentation, active presence of the lecturers during the labs that motivates students to listen even on Friday afternoons.” (AI)

A similar response was received from students of the ML course (spring term 2019):

“Good, guided (tutorial) exercises that were very helpful, especially for less advanced students.” (ML)

“The lectures are very interactive.” (ML)

We have increased the time for in-class exercises and interactivity over the years and received increasingly positive feedback on its effects. Despite the success of more modern teaching styles, such as “flipped classroom” [47], lecture-style teaching still seems to be a very helpful didactic setting for technical education if mixed with practical and interactive aspects where applicable.

3.3.4. Dimension “Open Educational Resources”

“The videos on YouTube are ideal for repeating.” (ML)

“The recording of the lectures is very helpful. It gives the students the possibility to review parts of the lecture for exam preparation or if you haven’t understood everything during the lecture.” (ML)

Students use video recordings as intended for repetition without getting distracted by the new flexibility (a real danger of digital transformation: procrastination due to everything being available anytime) (cf. Section 2.5.7).

3.3.5. Criticism: Dimensions “Self-Responsibility and Activation”

Most criticism that we face concerns the workload, the practical work in the lab sessions and the depth of mathematical derivation versus pure application of taught algorithms (cf. Sections 2.5.3 and 2.5.5).

“More exercises during the lecture or in the lab sessions. The topics are not always easy, and small exercises help to learn them correctly.” (AI)

“[The] labs are unfortunately a bit too time-consuming.” (AI)

“The theory necessary for the labs was partly a bit postponed, ... makes the beginning a bit difficult.” (AI)

“Maybe the lab sessions should not be so specific, but should cover a wider range of knowledge and not go into so much detail.” (AI)

“The way the lecturers address the subjects, in my opinion, is very theoretical. There are a lot of mathematical demonstrations that I consider to be out of the scope, and this time should be dedicated to make more examples of the subject. [...] Also the course demands way more time than the one available by a full-time student.” (ML)

“Unfortunately, the material covered is just too much. [...] you don’t really see the learning objectives for the exam. It’s good to also cover topics superficially and if you need them you can look them up more precisely. But in the lecture, it is not very clear which topics are such.” (ML)

“More exercises, clear knowledge points, mock exam.” (ML)

Table 3 summarizes the evaluation of the appeal and the organization of the ML course for spring term 2019 and fall term 2020 (the same data is not available for the AI course due to differing questionnaires). The students were asked to evaluate the following statements (the same scale applies as seen in Table 1): (i) *Labs*: the labs supplement the lectures in a meaningful manner and support the learning process; (ii) *Material*: the support materials (e.g., recommended books, documents handed out) are appropriate; (iii) *Organization*: the module is sensibly organized, and the coordination between the different lecturers works well.

Table 3. Appeal and organization evaluation scores for the ML course (spring 2019: $n = 28$, and fall 2020: $n = 27$).

Topic	Semester	1 (Strongly Disagree)	2 (Disagree to Some Extent)	3 (Agree to Some Extent)	4 (Strongly Agree)	Score
Labs	Spring 2019	7.1%	21.4%	50.0%	21.5%	2.9
	Fall 2020	0.0%	14.8%	37.0%	48.2%	3.3
Material	Spring 2019	0.0%	21.4%	50.0%	28.6%	3.1
	Fall 2020	0.0%	4.2%	41.7%	54.1%	3.5
Organization	Spring 2019	3.3%	16.7%	60.0%	20.0%	3.0
	Fall 2020	0.0%	25.9%	40.7%	33.4%	3.1

From Table 3, we see that there is much room for improvement. We are of the opinion that learning should be done based on examples and not by pure mathematical derivation. In addition, in a master course of applied science as the ML course, the practical application of the methods should be the deciding factor, and only the necessary mathematical definitions and derivations should be presented. From these critics, we conclude that there should even be more time spent on the application, on lab sessions and exercises. The topics covered in the ML course should be reduced to even less knowledge islands within the big ocean of AI and ML. The lecturer needs the courage to omit certain topics (some islands) and feel confident that the educational objective of generating a helpful map in each student's mind still can be reached. Over the years, hence, we moved more and more content from the actual lecture slides to the (optional) appendix of the lectures.

3.4. Quantitative Assessment

The following quantitative assessments focus on evaluating student learning success (i.e., reaching of educational objectives) under changing teaching and assessment modes (e-assessments since spring term 2020, on-site written exams before, corresponding to the distance and contact teaching modes during respective terms). In the absence of more direct data to measure the effectiveness of the AI-Atlas, we aim at drawing conclusions about its effectiveness in different teaching and learning settings from these outcome-based comparisons. This is possible since the respective exams were all designed with the same educational objectives and assessment goals, as well as similar question types in mind, irrespective of the rather drastic changes in assessment mode (from close-book in-class to self-supervised, open-book, open-internet).

3.4.1. AI Course Fall Terms 2019 vs. 2020

The content of the AI final exam has stayed largely stable over the years. It contains free text questions used to test students' ability to precisely define and argue for specific viewpoints; multiple choice exercises to show comprehensive knowledge; programming exercises inspired by the labs; and design tasks with transfer components to reveal higher level competencies. The move from a closed-book written exam to an open-book, open-internet e-assessment in the fall term 2020 of course changed the wording on numerous questions, but the overall depth and composition remained the same, as did the numbers and topics of tasks.

The distributions of the achieved relative scores for pairable tasks as shown in Figure 2 are very similar, which suggests that AI-Atlas suggestions also work well in distance learning mode. Nevertheless, significant differences are discernible in the overall result. The distribution is shifted to the left by about 10 percentage points as shown in Figure 3 (left), as discussed later in Section 4.1. It is striking that, in the planning task (in the bottom left of Figure 2), which is a modeling task involving transfer based on a suitable problem formulation using the Planning Domain Definition Language (PDDL) [48], many students did not know at all how to deal with it in 2020, while, in the group with contact instruction, this rather dry matter could be conveyed and apparently learned well.

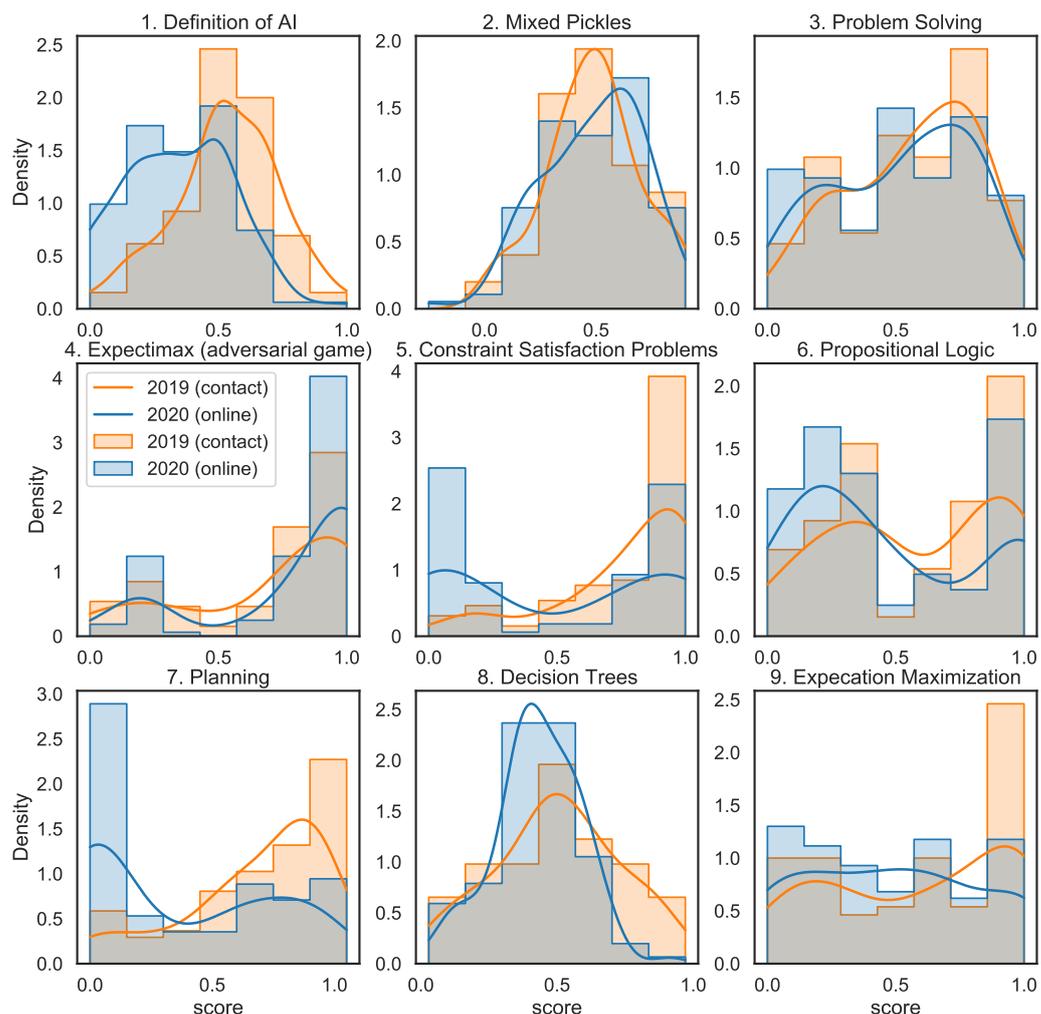


Figure 2. Histograms and kernel density estimates (kde) of the achieved relative scores (higher is better) for the different tasks of the final exams of the AI course of fall terms 2019 versus 2020. The total number of participants was $n_{19} = 91$ for fall term 2019 and $n_{20} = 113$ for the fall term 2020. An overview of the group sample can be found in Table A1. The content of the 9 tasks is irrelevant for the evaluation here unless otherwise noted, but indicated above each plot.

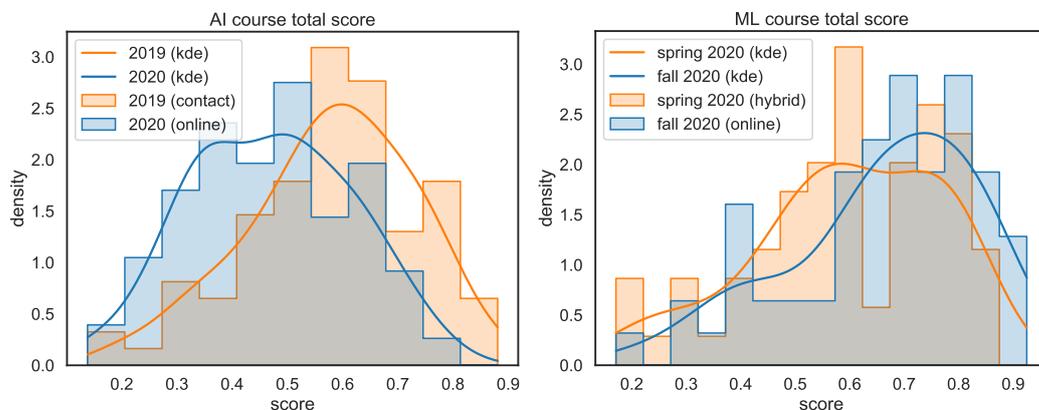


Figure 3. Comparisons of the total achieved relative scores of the AI (left) and ML (right) course exam results of respective semesters (higher is better; kde=kernel density estimate): spring terms 2019 (total number of participants $n = 91$) and 2020 ($n = 13$) for the AI course and spring ($n = 68$) and fall term 2020 ($n = 62$) for the ML course. An overview of the group sample can be found in Table A1 (AI) resp. Table A3 (ML).

3.4.2. ML Course Spring Term 2020 vs. Fall Term 2020

Due to the COVID-19 pandemic, the final assessment of spring and fall terms 2020 had to be taken in full distant mode over the Moodle [49] learning platform. This opened up the opportunity to re-design the existing exam: open book, as online proctoring could not be extended far enough to meaningfully control the use of only permissible aids; and involving hands-on programming, as every participant would sit in front of a well set-up developer's machine (the personal laptop). For this reason, programming, which was paramount for the lab exercises, could now also be included respectively in the exam in form of two programming tasks, together making up 50% of the exams' content. This is also the reason why a comparison with pre-pandemic results is not possible for the ML course: the respective exams would be too different to be comparable in a meaningful way.

Participants uploaded Jupyter notebooks [50] containing all programming at the end of the 120-min exam. The programming tasks asked the students to implement a *small, but full ML process* in Python (using scikit-learn), including (i) explorative data analysis (EDA), (ii) data preprocessing, (iii) feature generation and selection, (iv) algorithm selection, (v) hyper-parameter tuning, (vi) performance assessment, and (vii) comparison and conclusion. Although the tasks were different in topic and based on different data sets, we think that the results, nevertheless, can be compared as they are based on the same learning objectives and lecture parts. With these two programming tasks, we aimed to reach levels K3 and K4 of Bloom's taxonomy [51] and to test the educational objective of being able to *apply* and to *reflect* the ML process on a real data set from end to end.

The result is noteworthy: most of the participants did very well in programming. Figure 4 shows the histograms of the relative scores of the programming tasks. They are left-skewed, meaning that most of the students know how to apply machine learning to solve tasks in real life (the many 0-point entries might be the result of time problems with the exam as a whole, as this was the last task). This indicates that the overall educational objectives of the ML course—to *apply* ML algorithms—are met by the majority.

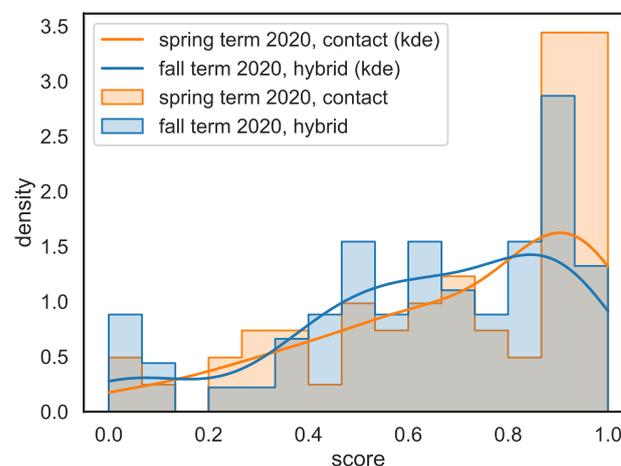


Figure 4. Histograms and kernel density estimates (kde) of the achieved relative scores (higher is better) for the programming tasks in the ML course's spring and fall term final exams 2020. The number of participating students was $n_{\text{spring}} = 68$ for the spring, and $n_{\text{fall}} = 62$ for the fall term. An overview of the group sample can be found in Table A3.

4. Evaluation and Reflection

According to our experience and as demonstrated above, the AI-Atlas is highly effective for teaching AI and ML principles in an on-site teaching setting. Aside from the presented data here, there are two additional, anecdotal pieces of evidence supporting this claim further: First, our alumni and alumnae have produced several award-winning theses inspired by the courses and many have ongoing (research) careers in ML. Second,

the AI-Atlas was recognized by the Zurich University of Applied Sciences with the “best teaching—best practice” award in 2019.

4.1. Tracing Weaker Quantitative Results in Online Teaching Mode

In our opinion, one reason for the effectiveness of the AI-Atlas is that it embodies the general learning settings, as shown in Section 2.3, under which students learn best, adapted specifically to the problems faced by current AI and ML tertiary education. In distance teaching mode, this effectiveness suffers somewhat for the AI course as can be seen Figure 3 (left). We performed a discrete Kolmogorov-Smirnov hypothesis test [52] for the total scores of the final exams of the AI and the ML course to check whether the samples of the two semesters stem from a common distribution (Null hypothesis H_0). With a significance level of $\alpha = 5\%$, we had to reject the Null hypothesis H_0 and assume that there is a significant difference in the distribution of the final scores in both courses when going from either hybrid to online (ML course, just barely different distributions) or from contact to online teaching (AI course, very clearly different distributions). We believe there is one main reason for this drop-off that will need addressing in future iterations, especially in those educational objectives that are based on practical implementation (programming, labs), social interaction (by discussions, competition, study-groups), and the teachers' presence (theoretical foundations), as explained next.

Tracing the shift in overall grades in the AI course between 2019 and 2020 as visible in Figure 3 (left), we see from Figure 2 that the reduction of effectiveness is mainly due to tasks 1, 5, and 7. These are tasks concerning precise definitions (1) or modeling (5, 7), each with a high transfer aspect. While modeling is technically difficult and complex, precise definitions are a matter of overarching concerns. We hypothesize that these two instances are the first aspects that suffer from the increasingly indirect influence of the lecturer on the students that takes place when shifting from contact to distance teaching: our students are usually technically interested, with less intrinsic motivation for overarching concerns, like, e.g., the precise difference between an AI system and other complex software systems (part of a question on defining what makes up AI). And, while they are technically interested, they are less motivated for prolonged sequences on, e.g., formal logic (as is part of the “planning” topic). In contact mode, the enthusiasm of the lecturer may help carry the motivation of a larger proportion of students through such sequences. In distant mode, where attention is naturally divided between the video stream, chat, the home environment, etc., less motivation is transmitted, and the subjects with the least intrinsic motivation suffer. Additionally, the social learning component through team-work is likely weaker in distant mode with imperfect collaboration tools, so that not every student that would normally be a member of a successful group is able to develop the deep skills necessary for novel modeling tasks on his or her own. Although AI and ML are practically done on a computer, and despite of the fact that “break-out rooms” could be used to organize team work remotely, the social inclusion of each individual likely suffered.

4.2. Tracing Worst Quantitative Results in Hybrid Teaching Mode

Quite counterintuitively, the overall quantitative results for the ML course as depicted in Figure 3 (right) improve again when moving to full online mode. In our opinion, the counterintuitive feeling lifts when considering that it is compared not to on-site teaching in figure 4, but to hybrid teaching. In our experience, hybrid teaching is most demanding for all participants, educators and learners alike, as the teacher has to try to address people in the room, as well as on the computer, which usually results in neglecting one group. In the first pandemic semester of spring term 2020, this condition was worsened by virtually no training for these special circumstances on the side of the educators, and imperfect hardware equipment, leading to frequent technical problems (degraded acoustic quality for discussion in the lecture room, illegible writings on whiteboards for students online, etc.).

We conjecture that the increase in overall scores for the ML course fall term 2020 (full online mode) is due to the ability of the educator in this setting being able to fully con-

concentrate on one stakeholder group again, focusing on delivering good streaming content. This interpretation is in line with the somewhat weaker result from the K-S test of the similarity of the two distributions, which are significantly dissimilar, but not as clearly different as for the AI course when going from on-site (good results) to online (degraded results).

4.3. *Moving towards a Didactic Concept for Flexible Teaching and Learning*

Ultimately, we believe that the AI-Atlas, conceived of and designed for an on-site learning environment, may be equally effective for online teaching if we solve for the problem of transporting some of the more social and teacher-enthusiasm-based design principles in ways that are suitable for the format. This would allow the AI-Atlas to become a fully flexible concept to support teaching regimes that might be any mix of online, hybrid and on-site teaching. What follows are the respective adjustments we are planning for future semesters (regardless of format) that are based on students' feedback and the principles outlined in Section 2.3 in the light of above's discussion.

4.3.1. Regarding "Reflection" and "Motivation"

To address the diminished attention span of students and subsequent lower motivation for overarching concerns in distance education settings, we plan more exercises and labs on such aspects to train reflective, interdisciplinary or even holistic thought patterns and reinforce the importance of such "soft" content (cf. Sections 2.5.2 and 2.5.3).

4.3.2. Regarding "Cooperation" and "Social Learning"

To strengthen the connection of students with their core groups for the lab exercises and other key persons in the class (e.g., top students), it is important to solicit quick discussions in the whole video call or frequent, smaller (possibly random) break-out rooms in distance lectures (cf. Sections 2.5.4 and 2.5.6).

4.3.3. Regarding "Activation of Students"

This element can be strengthened by frequent activation of students via online survey tools, like Mentimeter [53], during lectures to have them think on the latest input and produce some output (write a free text answer, make a choice, solve a puzzle or quiz, etc.). We already made positive experiences with a frequency of every second to sixth slide for such pauses for thought (corresponding to a 6–15-min interval between them) (cf. Section 2.5.5).

4.3.4. Regarding "Blended Learning"

The usual best practices for online teaching and learning [54] apply also to AI and ML. What we found especially helpful was to train students to have their cameras on most of the time (increases the perception of connectedness and the degree of interaction), and to roll out tools, like wonder.me (also see Reference [55]), for informal conversation with instantaneous participant groups, much like in a physical break or reception context. This informal networking is important to initiate cooperative competence development, as well (cf. Section 2.5.7).

4.4. *Conclusions*

With these measures added to the AI-Atlas didactic concept, most of them already implemented within our current AI and ML courses, we see any course crafted after the AI-Atlas principles fit for a very flexible range of teaching modes—whether it is on-site, online, or hybrid. Thus, we recommend the AI-Atlas as a viable basis for consideration when designing tertiary educational courses in AI, ML, and beyond (for the former ones, the syllabi and online teaching material presented in the appendix can also serve as starters to create new respective courses).

Author Contributions: The individual contributions of the authors are as follows: Conceptualization T.S.; Formal analysis C.W.; Investigation C.W.; Methodology T.S. and J.K.; Project administration T.S.; Validation H.G.; Writing (original draft) T.S. and C.W.; Writing (review & editing) J.K., C.W., T.S. and H.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Ethical review and approval were waived for this study, due to only standard processes and data sources of the involved study programs being used for data collection and analysis (exams, voluntary feedback).

Informed Consent Statement: Not applicable.

Data Availability Statement: Exam results, voluntary feedback and student inscriptions are hosted at the involved universities in consent with applicable law.

Acknowledgments: We are thankful for the recognition of the AI-Atlas didactic concept through the “best teaching—best practice” award (3rd rank) of the Zurich University of Applied Sciences; for very helpful discussions with and insights on design-based research by Andrea Reichmuth; and for the fruitful discussion on e-didactics with Elisabeth Dumont.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Outline of the AI and ML Modules

Appendix A.1. The AI Course

“Artificial Intelligence 1” (cf. <https://stdm.github.io/ai-course/>, accessed on 23 June 2021) is a practice-oriented elective course in the final year of a B.Sc. computer science program at a university of applied sciences, encompassing selected foundations of AI and ML and aiming at hands-on problem-solving competency for everyday software challenges. It is geared towards students who have a general curiosity for smartness in software but no aspirations towards research. Most of them, when starting the course, look forward to a career as software engineers, with some thinking about becoming data scientists or about further interdisciplinary studies in areas, like information engineering, speech processing, computer vision, or robotics. This group is quite homogeneous with respect to demography and educational background (cf. Table A1). Age-wise, the students are predominantly in their early twenties, ca. 1–3 years younger as in the ML course due to the AI course taking place at least a year earlier, and some students entering the industry for a while before engaging in master studies. The B.Sc. computer science program can be completed on a full-time or part-time basis.

Table A1. Study model and gender distribution of the students of the AI course from fall term 2019 and fall term 2020.

Semester Profile	Fall 2019		Fall 2020	
	Absolute	%	Absolute	%
Full time	38	39.6%	55	48.7%
Part time	58	60.4%	58	51.3%
Total	96	100%	113	100%
Male	86	89.6%	104	92.0%
Female	10	10.4%	9	8.0%

The superior learning objectives are defined as (a) knowing the breadth of AI and particularly ML problem solving strategies, thus identifying such challenges in practice and developing corresponding solutions on one’s own; (b) being able to explain the discussed algorithms and methodologies, thus being enabled to transfer the respective knowledge to the real world. The corresponding syllabus is depicted in Table A2. It is structured in five phases based on the main approaches to AI (symbolic and sub-symbolic) and an elaborate parenthesis dealing with overarching concerns.

The AI course is based on the well-known “AIMA” text book [33] (the much welcomed updates to the recent 4th edition from April 2020 have not yet been adopted; they include a more timely selection and framing of the contents that has partly been anticipated by our curriculum design). It presents AI as a toolbox with separate compartments (=sub-fields), each containing tools to mimic specific aspects of intelligent behavior suitable for certain

ranges of practical problems. The curriculum is special in that it gives equal time to the most relevant ideas from the complete field of AI, not just to fashionable topics around ML and neural networks or the main research areas of the lecturers. The course is taught once per year on-site during fall terms since 2017. The fall term 2020 started in online-only mode and went hybrid for the second half of the course.

Table A2. The curriculum of the AI course, spanning a 14-week semester with 2 lectures and 2 labs (45 min each) per week. On successful completion, the students are awarded 4 ECTS, meaning they have invested ca. 120 h into the coursework (i.e., they spent roughly twice the amount of time in self-study as in class, with most of this time invested into the lab assignments).

Topic (Duration)	Key Question	Methods (Excerpt)	Practice
1. Introduction to AI (2 weeks)	What is (artificial) intelligence?	The concept of a rational agent	AI for sci-fi readers: formulating one's own opinion as a reply to a futuristic essay [6]
2. Search (3 weeks)	How to find suitable sequences of actions to reach a complex goal?	Uninformed and heuristic search, (Expecti-)Minimax, constraint satisfaction problem solvers	AI for the game "2048": controlling a number puzzle game (cf. Appendix B)
3. Knowledge Representation & Planning (3 weeks)	How to represent the world in a way that facilitates reasoning?	Propositional and first order logic, knowledge engineering and reasoning, Datalog for big data, PDDL	AI for a dragnet investigation: finding potential fraudsters using inference over communication meta data
4. Supervised ML (3 weeks)	What is learning in machines? How to learn from examples?	From linear regression to decision trees and state of the art ensembles	AI for bargain hunters: data mining a dataset of used cars
5. Selected chapters (2 weeks)	What is the current hype about? How does AI effect society? How could society react?	Primer on deep neural networks and generative adversarial training for image generation	Sci-fi revisited: formulating a reply to the blog post from the first week

Appendix A.2. The ML Course

"Machine Learning" (cf. <https://stdm.github.io/ml-course/>, accessed on 23 June 2021) is an elective course in an interdisciplinary joint graduate program on engineering of different universities of applied sciences. It builds upon basic knowledge in math, programming, analytics, and statistics as is typically gained in respective undergraduate courses of diverse engineering disciplines and draws on a respective diverse audience with homogeneous demographics (age: 22–25 years) but rather heterogeneous backgrounds (cf. Table A3).

The module teaches the foundations of modern machine learning techniques in a way that focuses on practical applicability to real-world problems. The complete process of building a learning system is considered: formulating the task at hand as a learning problem; extracting useful features from the available data; and choosing and parameterizing a suitable learning algorithm. The syllabus highlights cross-cutting concerns, like ML system design and debugging (how to get intuition into learned models and results), as well as feature engineering, aspects typically cut short in previous courses these students took that touched on learning algorithms.

The corresponding educational objectives are designed as follows: (a) students know the background and taxonomy of machine learning methods; (b) on this basis, they formulate given problems as learning tasks and select a proper learning method; (c) students are able to convert a data set into a trained model by first defining a proper feature set fitting for a task at hand; then they evaluate the chosen approach in a structured way using proper design of experiment; they know how to select models, and "debug" features

and learning algorithms if results do not fit expectations; finally, they are able to leverage on the evaluation framework to tune the parameters of a given system and optimize its performances; (d) students have seen examples of different data sources and problem types and are able to acquire additional expert knowledge from the scientific literature.

Table A3. Background and gender of the students of the ML course from spring term 2019 to fall term 2020.

Semester Profile	Spring 2019		Spring 2020		Fall 2020	
	Absolute	%	Absolute	%	Number	%
Business Engineering	2	3.1%	6	7.6%	8	10.1%
Data Science					31	39.2%
Computer Science	43	67.2%	48	60.8%	18	11.4%
Electrical Engineering	3	4.7%	2	2.5%	10	12.7%
Environmental Science					3	3.8%
Industrial Technologies	16	25.0%	22	27.8%	9	11.4%
Mechatronics					2	2.5%
Medical Engineering					3	3.8%
Aviation					3	3.8%
Geomatics			1	1.3%	1	1.3%
Total	64	100 %	79	100 %	79	100%
Male	58	90.6%	74	93.7%	68	86.1%
Female	6	9.4%	5	6.3%	11	13.9%

The curriculum, depicted in Table A4, spends most time on first principles and illustrates them by specific, selected learning algorithms as the basis for life-long learning in ML. The ML course is not built around any specific textbook, but draws upon multiple sources, including References [33–36,56], having >90% original content. This is contrary to many courses that try to teach a large number of learning algorithms; it also eases the problem of heterogeneous entry competencies where students might have learned about the typical ML algorithms in some class already, but do not know what reasoning led to this specific class of algorithms. The ML course is structured four-fold with an introduction followed by supervised, unsupervised, and reinforcement learning and specifically does not touch neural networks as this is treated in a specialized course. The course has been taught on-site usually once a year in spring terms since 2017. Since spring 2020, the course moved to online teaching mode (with hybrid episodes) and is also taught in the fall term.

Table A4. The 3 ECTS curriculum of the ML course, spanning a 14-week semester of 2 lectures and 1 lab per week.

Topic (Duration)	Key Concept	Cross-Cutting Concerns	Methods (Excerpt)
1. Introduction (2 weeks)	Convergence for participants with different backgrounds	Hypothesis space search, inductive bias, computational learning theory, ML as representation-optimization-evaluation	No free lunch theorem, VC dimensions; ML from scratch: implementing linear regression with gradient descent purely from formulae
2. Supervised learning (7 weeks)	Learning from labeled data	Feature engineering, making the best of limited data, ensemble learning, debugging ML systems, bias-variance trade-off	Cross-validation, learning curve & ceiling analysis, SVMs, bagging, boosting, probabilistic graphical models
3. Unsupervised learning (3 weeks)	Learning without labels	Probability and Bayesian learning	Dimensionality reduction, anomaly detection, k-means and expectation maximization
4. Special chapters (2 weeks)	Reinforcement learning	Exploration-exploitation trade-off	AlphaZero

Appendix B. Content Example: AI Model Assignment

This section presents a content example from the AI course in the common format of the community (cf. <http://modelai.gettysburg.edu/>, accessed on 23 June 2021). It illustrates the technical depth required from the students, as well as the aspects that contribute to the high level of motivation throughout the courses taught according to the AI-Atlas didactic concept.

Appendix B.1. Summary, Topics, and Audience

The lab “2048 game playing agent” (cf. Figure A1) is a four-week assignment at the beginning of the AI course to be approached by pairs of two students (cf. http://stdm.github.io/downloads/courses/AI/P02_2048.zip, accessed on 23 June 2021). It is based on the game “2048” by Gabriele Cirulli (cf. <https://play2048.co/>, accessed on 23 June 2021) and covers the topics of rational agent development and adversarial search (heuristic search, Expectimax algorithm). The assignment is divided into two distinct phases, each with the task of developing an artificial player that controls the game to win, but different strategies and learning objectives.

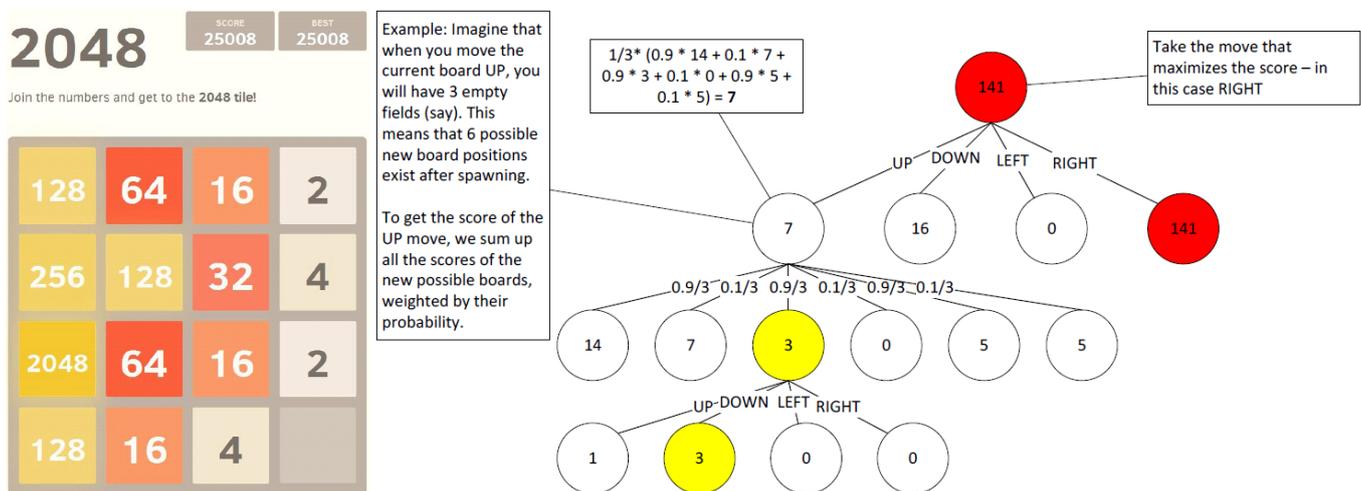


Figure A1. Top: screenshot from the 2048 number puzzle; goal of the game is to reach a 2048 tile by joining adjacent tiles of similar value through consecutive up/down/left/right movements of the whole board (cf. Reference [57] for a fuller description of the gameplay). Bottom: exemplary search tree as processed by Expectimax for a fictional board configuration, excerpted from the assignment.

Phase one is about taking one’s software development and problem solving skills, together with one’s understanding of the game after a few hours of playing, and implement an agent ad hoc by designing useful heuristics (links to the literature and online forums are provided, where ideas abound). The usual experience of a student after phase one is that it is very difficult and not overly successful to try encoding one’s own strategies purely ad hoc (and that it is impossible to exhaust the knowledge on the web and in the literature without a clear idea of how to conceptually approach the problem).

Phase two introduces the conceptual framework of adversarial heuristic search and the Expectimax algorithm. Students can leverage on their developed ideas of a heuristic function here, but thanks to the look-ahead provided by the search, reach scores usually an order of magnitude higher than their previous results (or manual play). This drives home the point that mapping the problem at hand to the best fitting conceptual/algorithmic approach from the literature pays off way more in AI than investing many hours of manual labor. It also reinforces Sutton’s “bitter lesson” that leveraging compute through search is usually the smartest thing one can do [58].

Appendix B.2. Strengths, Weaknesses, and Difficulty

This assignment's biggest strength is its addictiveness: students regularly report that they got so caught in the task that they worked through nights and weekends on the hunt for a better high score. This motivation carries over to trying other methods than search: we have seen deep reinforcement learning (RL) approaches developed during these four weeks, despite them not being part of the curriculum. Another strength is its accessibility: students on any skill level find something worthwhile to work on, be it improving their programming skills, understanding a recursive algorithm, or tapping into previously unknown scientific literature to understand RL.

A weakness of the assignment is its dependency on the pace of the corresponding lecture: it helps the educational objective of phase one that the students do not know search algorithms yet (so that they really try ad hoc solutions); it is, however, necessary for phase two that they are acquainted with adversarial search, so that the schedule of the lectures and labs needs to be tightly synced. Another weakness is that much of the initial motivation comes from the students knowing the 2048 game already from its viral history on the web; this effect is fading away over the years.

Appendix B.3. Dependencies and Variants

Platform-independent code templates in Python are given for all technicalities, like interaction with the game, so that students can focus purely on implementing the agent function $next_move = f(current_board)$. Students with a good command of any imperative programming language regularly take this as their first attempt to Python programming. Content-wise, the assignment is preceded by a general introduction to the field of AI, as well as to search algorithms, in the order of one lecture each. Before entering phase two of the assignment, students need to get an introduction to adversarial search and the Expectimax algorithm. An easy variation of the assignment would be to exchange the game by another version that might be more fashionable (and hence able to evoke interest with students) in a few years.

References

1. Skog, D.A.; Wimelius, H.; Sandberg, J. Digital disruption. *Bus. Inf. Syst. Eng.* **2018**, *60*, 431–437. [CrossRef]
2. Aoun, J.E. *Robot-Proof: Higher Education in the Age of Artificial Intelligence*; MIT Press: Cambridge, MA, USA, 2017.
3. Stadelmann, T. Wie maschinelles Lernen den Markt verändert. In *Digitalisierung: Datenhype mit Werteverlust? Ethische Perspektiven für Eine Schlüsseltechnologie*; SCM Hänsler: Holzgerlingen, Germany, 2019; pp. 67–79. [CrossRef]
4. Perrault, R.; Shoham, Y.; Brynjolfsson, E.; Clark, J.; Etchemendy, J.; Grosz, B.; Lyons, T.; Manyika, J.; Mishra, S.; Niebles, J.C. *The AI Index 2019 Annual Report*; AI Index Steering Committee, Human-Centered AI Institute, Stanford University: Stanford, CA, USA, 2019.
5. Stadelmann, T.; Braschler, M.; Stockinger, K. Introduction to applied data science. In *Applied Data Science*; Springer: Cham, Switzerland, 2019; pp. 3–16.
6. Urban, T. The AI Revolution: The Road to Superintelligence. 2015. Available online: <https://waitbutwhy.com/2015/01/artificial-intelligence-revolution-1.html> (accessed on 23 June 2021).
7. Luger, G.F. *Artificial Intelligence: Structures and Strategies for Complex Problem Solving*, 6th ed.; Pearson Education: Boston, MA, USA, 2009.
8. Wikipedia Contributors. Informatik—Wikipedia, The Free Encyclopedia. 2021. Available online: https://de.wikipedia.org/wiki/Informatik#Disziplinen_der_Informatik (accessed on 27 March 2021).
9. Dessimoz, J.D.; Köhler, J.; Stadelmann, T. AI in Switzerland. *AI Mag.* **2015**, *36*, 102–105. [CrossRef]
10. Stadelmann, T.; Stockinger, K.; Bürki, G.H.; Braschler, M. Data scientists. In *Applied Data Science*; Springer: Cham, Switzerland, 2019; pp. 31–45.
11. Parsons, S.; Sklar, E. Teaching AI using LEGO mindstorms. In Proceedings of the AAAI Spring Symposium, Palo Alto, CA, USA, 22–24 March 2004.
12. Goel, A.K.; Joyner, D.A. Using AI to teach AI: Lessons from an online AI class. *AI Mag.* **2017**, *38*, 48–59. [CrossRef]
13. Huang, R.; Spector, J.M.; Yang, J. Design-based research. In *Educational Technology*; Springer: Singapore, 2019; pp. 179–188.
14. Williams, R.; Park, H.W.; Oh, L.; Breazeal, C. Popbots: Designing an artificial intelligence curriculum for early childhood education. In Proceedings of the AAAI Conference on Artificial Intelligence, Honolulu, HI, USA, 27 January–1 February 2019; Volume 33, pp. 9729–9736.
15. Stadelmann, T. ATLAS—Analoge Karten für die Digitale Welt der Künstlichen Intelligenz. 2019. Available online: <https://stdm.github.io/ATLAS/> (accessed on 3 June 2021).

16. Stadelmann, T.; Würsch, C. *Maps for an Uncertain Future: Teaching AI and Machine Learning Using the ATLAS Concept*; Technical Report; ZHAW Zürcher Hochschule für Angewandte Wissenschaften: Winterthur, Switzerland, 2020.
17. Siegel, E.V. Why do fools fall into infinite loops: Singing to your computer science class. *ACM SIGCSE Bull.* **1999**, *31*, 167–170. [[CrossRef](#)]
18. Eaton, E. Teaching integrated AI through interdisciplinary project-driven courses. *AI Mag.* **2017**, *38*, 13–21. [[CrossRef](#)]
19. Schreiber, B.; Dougherty, J.P. Embedding Algorithm Pseudocode in Lyrics to Facilitate Recall and Promote Learning. *J. Comput. Sci. Coll.* **2017**, *32*, 20–27.
20. Chiu, T.K.; Chai, C.S. Sustainable Curriculum Planning for Artificial Intelligence Education: A Self-Determination Theory Perspective. *Sustainability* **2020**, *12*, 5568. [[CrossRef](#)]
21. Touretzky, D.; Gardner-McCune, C.; Martin, F.; Seehorn, D. Envisioning AI for k-12: What should every child know about AI? In Proceedings of the AAAI Conference on Artificial Intelligence, Honolulu, HI, USA, 27 January–1 February 2019; Volume 33, pp. 9795–9799.
22. Li, Y.; Wang, X.; Xin, D. An Inquiry into AI University Curriculum and Market Demand: Facts, Fits, and Future Trends. In Proceedings of the Computers and People Research Conference, Nashville, TN, USA, 20–22 June 2019; pp. 139–142.
23. Norvig, P. 1525 Schools Worldwide That Have Adopted AIMA. 2021. Available online: <http://aima.cs.berkeley.edu/adoptions.html> (accessed on 27 March 2021).
24. Mercator, G. *Atlas sive Cosmographicae Meditationes de Fabrica Mundi et Fabricati Figura*; Rumold Mercator: Duisburg, Germany, 1595.
25. Schneider, U.; Brakensiek, S. (Eds.) *Gerhard Mercator: Wissenschaft und Wissenstransfer*; WBG: Darmstadt, Germany, 2015.
26. Ford, K.M.; Hayes, P.J.; Glymour, C.; Allen, J. Cognitive orthoses: Toward human-centered AI. *AI Mag.* **2015**, *36*, 5–8. [[CrossRef](#)]
27. Nilsson, N.J. *The Quest for Artificial Intelligence*; Cambridge University Press: Cambridge, UK, 2009.
28. Morrell, M.; Capparell, S. *Shackleton's Way: Leadership Lessons from the Great Antarctic explorer*; Penguin Press: New York, NY, USA, 2001.
29. Gerstenmaier, J.; Mandl, H. Wissenserwerb unter konstruktivistischer Perspektive. *Z. für Pädagogik* **1995**, *41*, 867–888.
30. Ryan, R.; Deci, E. Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *Am. Psychol.* **2000**, *55*, 68–78. [[CrossRef](#)]
31. Winteler, A. Lehrende an Hochschulen. In *Lehrbuch Pädagogische Psychologie*; Beltz Psychologie Verlags Union: Weinheim, Germany, 2006; pp. 334–347.
32. Helmke, A.; Schrader, F.W. Hochschuldidaktik. In *Handwörterbuch Pädagogische Psychologie*; Beltz Psychologische Verlags Union: Weinheim, Germany, 2010; pp. 273–279.
33. Russell, S.; Norvig, P. *Artificial Intelligence: A Modern Approach*, 3rd ed.; Pearson Education, Inc.: New York, NJ, USA, 2010.
34. Mitchell, T.M. *Machine Learning*; McGraw-Hill: New York, NY, USA, 1997.
35. Murphy, K.P. *Machine Learning: A Probabilistic Perspective*; MIT Press: Cambridge, MA, USA, 2012.
36. Domingos, P. A few useful things to know about machine learning. *Commun. ACM* **2012**, *55*, 78–87. [[CrossRef](#)]
37. Bloom, B.S.; Engelhart, M.D.; Furst, E.J.; Hill, W.H.; Krathwohl, D.R. *Taxonomy of Educational Objectives: The Classification of Educational Goals*; Longmans, Green and Co. Ltd.: London, UK, 1956; Volume 1.
38. Wolpert, D.H.; Macready, W.G. No free lunch theorems for optimization. *IEEE Trans. Evol. Comput.* **1997**, *1*, 67–82. [[CrossRef](#)]
39. Pressman, I.; Singmaster, D. “The jealous husbands” and “the missionaries and cannibals”. *Math. Gaz.* **1989**, *73*, 73–81. [[CrossRef](#)]
40. Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; Zaremba, W. Openai gym. *arXiv* **2016**, arXiv:1606.01540.
41. Tuomi, I. Open educational resources and the transformation of education. *Eur. J. Educ.* **2013**, *48*, 58–78. [[CrossRef](#)]
42. Braschler, M.; Stadelmann, T.; Stockinger, K. *Applied Data Science*; Springer: Cham, Switzerland, 2019.
43. Brodie, M.L. On developing data science. In *Applied Data Science*; Springer: Cham, Switzerland, 2019; pp. 131–160.
44. Downing, K.F.; Holtz, J.K. A Didactic Model for the Development of Effective Online Science Courses. In *Online Science Learning: Best Practices and Technologies*; IGI Global: Hershey, PA, USA, 2008; pp. 291–337. [[CrossRef](#)]
45. Hoffmann, R.L.; Dudjak, L.A. From onsite to online: Lessons learned from faculty pioneers. *J. Prof. Nurs.* **2012**, *28*, 255–258. [[CrossRef](#)]
46. Torres, A.; Domańska-Głonek, E.; Dzikowski, W.; Korulczyk, J.; Torres, K. Transition to online is possible: Solution for simulation-based teaching during the COVID-19 pandemic. *Med. Educ.* **2020**, *54*, 858–859. [[CrossRef](#)]
47. Cieliebak, M.; Frei, A.K. Influence of flipped classroom on technical skills and non-technical competences of IT students. In Proceedings of the 2016 IEEE Global Engineering Education Conference (EDUCON), Abu Dhabi, United Arab Emirates, 10–13 April 2016; pp. 1012–1016.
48. Haslum, P.; Lipovetzky, N.; Magazzeni, D.; Muise, C. An introduction to the planning domain definition language. *Synth. Lect. Artif. Intell. Mach. Learn.* **2019**, *13*, 1–187. [[CrossRef](#)]
49. Dougiamas, M.; Taylor, P. Moodle: Using learning communities to create an open source course management system. In Proceedings of the World Conference on Educational Multimedia, Hypermedia and Telecommunications (EDMEDIA), Honolulu, HI, USA, 23–28 June 2003; pp. 171–178.
50. Kluyver, T.; Ragan-Kelley, B.; Pérez, F.; Granger, B.E.; Bussonnier, M.; Frederic, J.; Kelley, K.; Hamrick, J.B.; Grout, J.; Corlay, S.; et al. Jupyter Notebooks—A Publishing Format for Reproducible Computational Workflows. In *Positioning and Power in Academic Publishing: Players, Agents and Agendas*; IOS Press: Amsterdam, The Netherlands, 2016; pp. 87–90. [[CrossRef](#)]

51. Bloom, B.S. *Taxonomy of Educational Objectives*; Pearson Education; Allyn and Bacon: Boston, MA, USA, 1984.
52. Bityukov, S.; Maksimushkina, A.; Smirnova, V. Comparison of histograms in physical research. *Nucl. Energy Technol.* **2016**, *2*, 108–113. [[CrossRef](#)]
53. Valley, K.; Gibson, P. Engaging students on their devices with Mentimeter. *Compass J. Learn. Teach.* **2018**, *11*. [[CrossRef](#)]
54. Adams, A.L. Online Teaching Resources. *Public Serv. Q.* **2020**, *16*, 172–178. [[CrossRef](#)]
55. Gruber, A. Employing innovative technologies to foster foreign language speaking practice. *Acad. Lett.* **2021**, *2*. [[CrossRef](#)]
56. Andrew Ng. Machine Learning—Coursera (Stanford University). 2021. Available online: <https://www.coursera.org/learn/machine-learning> (accessed on 28 March 2021).
57. Wikipedia Contributors. 2048 (Video Game)—Wikipedia, the Free Encyclopedia. 2021. Available online: [https://en.wikipedia.org/wiki/2048_\(video_game\)](https://en.wikipedia.org/wiki/2048_(video_game)) (accessed on 27 March 2021).
58. Sutton, R. The Bitter lesson—Incomplete Ideas Blog. 2019. Available online: <http://www.incompleteideas.net/IncIdeas/BitterLesson.html> (accessed on 3 June 2021).