

Learning from Machine Learning and Teaching with Machine Teaching: Using Lessons from Data Science to Enhance Collegiate Classrooms

Dr. Lucas Buccafusca, Johns Hopkins University

My name is Lucas Buccafusca. I am currently a teaching faculty at Johns Hopkins University in Electrical and Computer Engineering. I received my Ph.D. in Industrial and Systems Engineering at the University of Illinois at Urbana-Champaign, earned my Masters in Electrical and Computer Engineering from the University of Illinois at Urbana-Champaign in 2017 and my Bachelor's degree in Electrical and Computer Engineering in 2013 from the University of Colorado at Boulder. My pedagogical research interests are on improving the quality of collegiate classroom environments through the use of nontraditional techniques and active participation by instructors. These include the use of failure as a teaching tool, humor and empathy as a means of connecting with students, and gamification. My technical research interests are Distributed Control, Learning, Distributed Optimization and Nonlinear Systems. Applications of my research are primarily used for Wind Farm arrays.

Learning from Machine Learning and Teaching with Machine Teaching: Using Lessons from Data Science to Enhance Collegiate Classrooms

Lucas Buccafusca

LBuccafusca@jhu.edu

Department of Electrical and Computer Engineering

Johns Hopkins University

Abstract

In the field of data science, advancements in the field of machine learning have led to programs developing high-level reasoning, intricate data understanding, and groundbreaking predictive models. Machine Learning (ML) research aims at making a program ‘learn,’ that is, develop models and techniques with known information to be able to handle future problems. Traditionally, this is done by increasing the quantity and quality of input data and training the learner in more effective ways to interpret that information. This has a direct parallel to the collegiate classroom, as instructors aim to inspire mastery over a topic to their students through a variety of methods (homework problems, examinations, projects, etc.) and teach them the corresponding skillsets from feedback on these assignments. Machine Teaching (MT) research, on the other hand, aims at making the teacher more productive by using their own cognitive models to improve the quality of the data holistically. Again, this has a corresponding counterpart to current teaching pedagogies; the instructor decides on the details of an assignment from their own knowledge and experience with the end goal of having students retain the information and apply it to future problems. This paper identifies how the various innovations, lessons, and conclusions discovered in the field of artificial intelligence can enhance the quality of a collegiate classroom experience and improve student performance.

Introduction

With the development of advanced computing, the latest technological advances find themselves continuously seeping into all areas of life, including post-secondary education. As classrooms evolve from the traditional model wherein a lecturer presents material on a chalkboard, technology becomes a beneficial supplementary method to enhance student retention. This explosion of new pedagogical perspectives have explored the use of Artificial Intelligence (AI) as an improvement to student learning.

AI is proving to be an effective tool for educators teaching anywhere from K-12 [1] or in secondary education [2] to enhance teaching and provide students with personalized learning experi-

ences. State-of-the-art AI technologies have been able to analyze vast amounts of data to identify patterns, adapt to student needs, and provide real-time feedback with little up-front implementation costs. As such, it has been shown that this tailored instruction and support to each student can improve their learning outcomes [3], [4]. Moreover, AI has been used to automate routine tasks such as grading, assessment, and administrative duties, freeing up educators' time to focus on higher-level tasks. In this way, AI has been the catalyst in a reframing of the education sector and enable instructors to transform classrooms and thus provide a more efficient, effective, and engaging learning experience for students. Two related but distinct concepts within AI that are gaining a wide variety of attention in education are machine learning and machine teaching.

Most studies on the use of AI in the classroom focus on the topic known as *machine learning* [5], [6], [7], [8], [9], [10]. Machine Learning (ML) in the classroom aims to "make the learner better by improving ML algorithms [11]." This research is on the development of new algorithms, methods, or techniques that can improve results based on labeled or unlabeled data sets. Specifically, ML is an application of AI that provides a program the ability to automatically learn and improve its performance without being explicitly told how to do that task or any details of the underlying model. Machine learning is commonly used in applications such as image recognition, natural language processing, and predictive models.

However, while machine learning is often explored as a means to improve the quality of learners, *machine teaching* has a more direct parallel. Machine teaching research "aims at making the teacher more productive at building machine learning models" given the learners [11]. That is, machine teaching is about curating a set of problems, modules, or assignments in an appropriate way so that the path towards learning the underlying concepts is optimal. In other words, with such a large swath of information available about any given topic, how does an instructor select the most important ones to teach any given concept? Machine teaching has had limited discussions in the realm of education [11], [12], [13], [14], but its parallels allow for many conclusions in the data science realm to be directly implemented in the classroom.

One can think of a machine learning algorithm as a black box: data goes in at the beginning, and other data goes out at the end and the underlying processes in between are complex. Machine learning aims to approximate the contents of the black box so that the outputs closely resemble those associated with the corresponding inputs. Machine teaching is the opposite process. One knows the contents of the box and wants to figure out the best set of inputs (i.e. assignments) so that any student can get the correct set of outputs (the knowledge to be imparted in the classroom).

The remainder of the paper is structured as follows. First, an algebraic model is discussed to formalize the difference between machine learning and machine teaching. Following that, a discussion on the unique advantages of machine learning algorithms are covered, with specific examples of their application to the collegiate classroom. Next, the same overview is given towards machine teaching, again focusing on applications in higher learning. The work ends with a brief mention of open problems, and conclusions.

Understanding the Differences Machine Learning and Machine Teaching as Represented by Algebraic Model

One oversimplifying perspective is that the basis of education (that is, the high-level goal) is for students to learn concepts, skills, and facts to strengthen themselves holistically. One representation of the nature of instruction is that the teacher wishes to impart θ (some collection of knowledge or information known by the instructor) onto the learner. These i learners can be represented as a collection of functions f_i whose inputs D_i are the classroom material, homework, resources, etc. presented to student i . The corresponding output of $f_i(D_i)$ can either be quantitatively represented by a grade or qualitatively described as a mastery of the underlying material. This is an intuitive perspective: students take in a variety of inputs (D_i) and aim to learn using their own strengths and weaknesses, and is measured with a quantitative representation $f_i(D_i)$. Teachers aim to maximize the sum of the functions $\sum_{\forall i} f_i(D_i)$ by designing or tailoring D_i [15]. Ideally, θ is correlated with $\sum_{\forall i} f_i(D_i)$; thus by designing courses well (i.e. designing D_i appropriately) learners grasp θ , and a measuring tool to evaluate a successful class is captured by $\sum_{\forall i} f_i(D_i)$.

Using the same functional representation discussed above, an ML algorithm implies the teacher has some training data (d) and the framework is to find or use an algorithm (A) based on that information. The solution obtained (represented by $A(d)$) can then be used on a wide swath of different data (that is, different student performances) and may have meaningful or improved results. ML algorithms applied in this fashion are ones where the goal is to try and use the knowledge gained by the process of learning θ to identify a lack of understanding, provide timely interventions, or curate problems. In other words, the instructor knows θ and how students learn f_i , and can use the algorithm A as a means to detect if a student is learning an incorrect set of knowledge θ^* . When it comes to the collegiate classroom, ML has had a wide variety of documented applications ranging from improved course efficiency using predictive analytics or personalized adaptive learning based on individual performance. Succinctly written, the Machine Learning problem is: **given modules, materials, lecture style, etc. D , how does an instructor choose to intervene or change so the students learn θ correctly?**

The general representation of an MT algorithm is different than that of the ML method presented above. In this can the teacher does not have training data they wish to impose, but rather the model θ itself. The goal of machine teaching is to find A^{-1} ; the optimal way to teach θ . That is, the Machine Teaching problem is: **given information θ , how does an instructor choose the modules, materials, lecture style, etc. (D) so the students learn θ ?**

This fundamental difference is crucial to understand, as many believe that ML and MT are similar. While they are both under the umbrella of AI, machine learning aims to course correct when the students perceived understanding of the material differs from the true knowledge θ , whereas machine teaching is about how to design the means to teach θ in an optimal way. A visualization of this model can be seen in Figures 1a-1c.

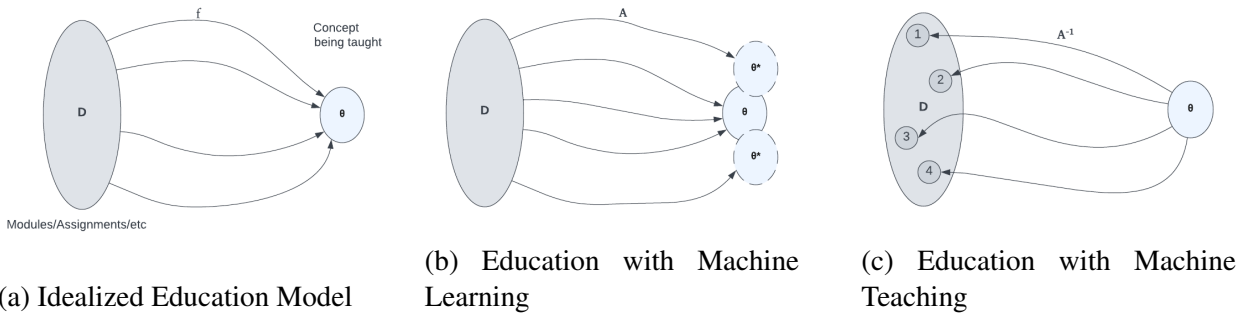


Figure 1: Visualization of the algebraic model discussed. In 1a, the instructor aims for the students to learn θ given a set of materials D . Each student learns slightly differently, represented by f . Ideally, all students reach θ given D . In 1b, a machine learning algorithm A can be used to detect or otherwise be aware when students deviate from θ . In 1c, machine teaching instead identifies what are the elements in D that will lead to θ being learnt. Both methods are trying to find the best solutions so that students arrive at θ .

Learning from Machine Learning: Tools, Techniques, and Methods to Improve the Collegiate Classroom

There has been an extensive amount of studies on applications, results, and issues with machine learning[16], [17], [18], [19]. From these we can separate out many different key benefits from these methods and apply them to the learning infrastructure.

Predictive Analytics

The most prevalent and useful application of ML in the classroom is to offer corrections via predictive analytics. Taking a repository of old data such as previous coursework, attendance, and participation, machine learning algorithms can make accurate predictions about how likely each student is to succeed in a given course. This information can be used to identify struggling students early on, and to provide them with targeted interventions and support. This does not need to be done exclusively on individual students. By observing student performance on specific assignments and exams (and noting the corresponding topics that are addressed by these modules), ML algorithms can identify areas where the majority of students are struggling. This information can be used to identify knowledge gaps caused by the selection of material.

Improved Student Engagement

Machine learning algorithms have been used to analyze student behavior, such as how often they log in to the course platform, how much time they spend on each task, and how often they participate in online discussions [20]. This information can be used to identify students who are either not investing time into the material or are lost. This allows for instructors to intervene and to provide personalized incentives and support to improve the probability of success [21].

Better Grading of Assessments

Perhaps the most frequent complaint from students about grading involves human biases in assessments. Whether it be on homework with inconsistent and vague rubrics or examinations being harshly or unevenly graded, the need for precise and fair grading systems has been an active topic of research [22], [23].

Machine learning can offer a variety of tools and techniques to help automate and provide fair assessments. While automated grading has been commonplace as a means to reduce the labor of educators, incorporating machine learning in education enables ‘smart’ assessments. That is, programs that can rapidly evaluate submissions in a wide variety of formats, ranging from written assignments such as papers and essays, to videos and presentations. That is not to say that machine learning is immune to biases [24], [25]. Any direct application must be cognizant of these potential issues.

Furthermore, ML can ingratiate automation into the grading process. While it is easier to develop a program that can grade multiple choice problems, using advanced AI can allow for grading of more complex submissions, such as essays. This can help offset one of the major automation issues: watering down of assessments in order to lessen workload. By ensuring that a wide variety of testing methods are applied, students can not only receive rapid feedback (and thus be able to more quickly adapt if there are issues) but also cover a variety of methods and topics with higher levels of critical thinking.

Teaching with Machine Teaching: Developing, Designing, and Demonstrating Effective Education

Machine teaching is a relatively new field that has far fewer discussions when applied to education. In this section the major possible contributions and applications are discussed when employing machine teaching as a means of improving student performance.

Scaffolding and Pattern Recognition

Scaffolding is a documented pedagogical technique in which the teacher assists a learner in accomplishing a task currently beyond the learner’s abilities, by reducing the degrees of freedom of the problem or by demonstrating partial solutions to the task at hand [26]. ML methods can be used to identify which axes of a problem are best kept hidden for the student and which are best utilized to teach a problem solving technique. Furthermore, one of the major tools for reinforcement of learning is through pattern recognition. Completing repetition of the solution process helps students retain the information, but it is not without fault. Specifically, students have a bias toward simpler patterns and these explanations can be detrimental or even lead to the discovery of false patterns altogether [27]. Proper design of assessments helps guarantee that these pitfalls are eliminated altogether in an optimal way.

Optimized learning materials

Machine teaching algorithms can analyze how students engage with different types of learning materials, such as videos, articles, and quizzes. This information can be used to optimize the delivery of course content, for example by creating personalized learning paths for each student. Curating the material on a per-student basis is arguably the most useful application of machine teaching [12].

Personalized and adaptive learning

With machine teaching, educators can identify patterns on prior assignments, problems, and exams and note what are the learning outcomes that are either successfully (or unsuccessfully) grasped. Furthermore, by noting supplemental information, such as prerequisites, instructors can personalize a subset of problems on a per-student basis and make predictions about each student's learning needs. This can lead to more targeted personalized instruction and a customized curriculum designed to maximize the performance of students. Furthermore, machine teaching methods can also suggest new or unknown pieces of relevant information that may resonate better with students from a diverse set of age groups, backgrounds and cultures. Identifying which material (and the form of that content) best works as a teaching tool can drastically improve student performance [11].

Issues with Machine Learning and Machine Teaching in Education

While there are many benefits to weaving data science methods to improve student outcomes, it is not a catch-all solution. There are many different possible issues that can arise that can either undo the potential benefits or become detrimental to the learning process.

To begin with, there is a necessary barrier to entry when applying these techniques. While the qualitative heuristics that arise from using AI as a tool are straightforward, there might be an inhibitor to taking these new methods and synchronizing them with prior pedagogical choices. Whether that be familiarity with the technologies to implement such methods or a lack of data to make meaningful conclusions, ML and MT algorithms are not necessarily diverse enough to be used in all situations.

Furthermore, there are discussions that such applications struggle to adequately improve the quality of courses in the humanities. Many of the assessment tools have difficulty identifying high quality responses from low effort ones. In fact, open source AI chatbots can be used to trick ML grading algorithms and remove the student from the feedback loop altogether. Thus, there arises the need for ML algorithms to detect ML submissions. This requires educators to invest time and energy better diverted to course material and development to focus on the ethical ramification of such problems.

Lastly, many of these methods are algorithmically complex. It can be quite difficult to implement, as one needs either a large repository or dataset to derive ML algorithms, or some understanding on how specific materials correlate to student success. This means that the direct implementations are best held for long-standing courses where data is available or in larger sections that can offer statistically significant results.

Conclusions

The results presented herein cover the implementations of traditional machine learning methods but also cover the oft-forgotten topic of machine teaching. By analyzing the strengths of each AI method and their corresponding parallels in the classroom, one can incorporate them seamlessly. While issues do exist with each, the evolution of teaching is pushing the forefront of student learning. Future works will apply some of the methods discussed on freshmen level STEM courses and discuss how the theoretical rewards translate when applied directly. Specifically, due to their larger size, ML can be used to proactively correct students who did not grasp a specific concept or topic. By choosing to focus on low-stakes, high-frequency quizzes (with rapid grade turnaround), a large dataset for student performance can be used to detect topics that would involve further explorations, perhaps in homework or exam settings. In addition, by offering a variety of problems, ranging from theoretical derivations to hands-on assignments, a subset of high-quality optimal learning materials can be constructed and tested for future use.

References

- [1] Mostafa Zafari et al. “Artificial Intelligence Applications in K-12 Education: A Systematic Literature Review”. In: *IEEE Access* 10 (Jan. 2022), pp. 1–1. DOI: 10.1109/ACCESS.2022.3179356.
- [2] Joseph Beck, Mia Stern, and Erik Haugsjaa. “Applications of AI in Education”. In: *XRDS: Crossroads, The ACM Magazine for Students* 3.1 (1996), pp. 11–15.
- [3] David Touretzky et al. “A year in K-12 AI education”. In: *AI Magazine* 40.4 (2019), pp. 88–90.
- [4] Beverly Woolf. *AI in Education*. University of Massachusetts at Amherst, Department of Computer and . . . , 1991.
- [5] Hadeel S. Alenezi and Maha Faisal. “Utilizing crowdsourcing and machine learning in education: Literature review”. In: *Education and Information Technologies* 25 (2020), pp. 2971–2986.
- [6] Danijel Kucak, Vedran Juricic, and Goran Dambic. “Machine Learning in Education - a Survey of Current Research Trends”. In: Jan. 2018, pp. 0406–0410. ISBN: 9783902734204. DOI: 10.2507/29th.daaam.proceedings.059.
- [7] Vijaya B Kolachalama and Priya S Garg. “Machine learning and medical education”. In: *NPJ digital medicine* 1.1 (2018), p. 54.
- [8] Rosemary Luckin. *Machine Learning and Human Intelligence: The future of education for the 21st century*. ERIC, 2018.
- [9] Radu Mariescu-Istodor and Ilkka Jormanainen. “Machine learning for high school students”. In: *Proceedings of the 19th Koli calling international conference on computing education research*. 2019, pp. 1–9.

- [10] Sotiris B Kotsiantis. “Use of machine learning techniques for educational proposes: a decision support system for forecasting students’ grades”. In: *Artificial Intelligence Review* 37 (2012), pp. 331–344.
- [11] Patrice Y. Simard et al. “Machine Teaching: A New Paradigm for Building Machine Learning Systems”. In: *ArXiv abs/1707.06742* (2017).
- [12] Xiaojin Zhu. “Machine Teaching: An Inverse Problem to Machine Learning and an Approach Toward Optimal Education”. In: 29 (Mar. 2015). DOI: 10.1609/aaai.v29i1.9761.
- [13] Xiaojin Zhu et al. “An overview of machine teaching”. In: *arXiv preprint arXiv:1801.05927* (2018).
- [14] Tomi Peltola et al. “Machine Teaching of Active Sequential Learners”. In: Nov. 2019.
- [15] R.V. Lindsey et al. “Optimizing instructional policies”. In: *Advances in Neural Information Processing Systems* (Jan. 2013).
- [16] Tom Michael Mitchell et al. *Machine learning*. Vol. 1. McGraw-hill New York, 2007.
- [17] Michael I Jordan and Tom M Mitchell. “Machine learning: Trends, perspectives, and prospects”. In: *Science* 349.6245 (2015), pp. 255–260.
- [18] Zhi-Hua Zhou. *Machine learning*. Springer Nature, 2021.
- [19] Claude Sammut and Geoffrey I Webb. *Encyclopedia of machine learning*. Springer Science & Business Media, 2011.
- [20] Islam Alkabbany et al. “An Experimental Platform for Real-Time Students Engagement Measurements from Video in STEM Classrooms”. In: *Sensors* 23 (Feb. 2023), p. 1614. DOI: 10.3390/s23031614.
- [21] Hafizah Pratiwi and Joko Priyana. “Exploring Student Engagement in Online Learning”. In: *Jurnal Ilmu Pendidikan* 28 (Jan. 2023), p. 66. DOI: 10.17977/um048v28i2p66–82.
- [22] Dannelle D Stevens and Antonia J Levi. *Introduction to rubrics: An assessment tool to save grading time, convey effective feedback, and promote student learning*. Stylus Publishing, LLC, 2013.
- [23] Susan M Brookhart. *How to create and use rubrics for formative assessment and grading*. Ascd, 2013.
- [24] Ninareh Mehrabi et al. “A survey on bias and fairness in machine learning”. In: *ACM Computing Surveys (CSUR)* 54.6 (2021), pp. 1–35.
- [25] Susan Leavy. “Gender bias in artificial intelligence: The need for diversity and gender theory in machine learning”. In: *Proceedings of the 1st international workshop on gender equality in software engineering*. 2018, pp. 14–16.
- [26] David Wood, Jerome S Bruner, and Gail Ross. “The role of tutoring in problem solving.” In: *Child Psychology & Psychiatry & Allied Disciplines* (1976).
- [27] Joseph Jay Williams, Tania Lombrozo, and Bob Rehder. “Why Does Explaining Help Learning? Insight From an Explanation Impairment Effect.” In: *Proceedings of the Annual Meeting of the Cognitive Science Society*. Vol. 32. 32. 2010.