Teaching Machine Learning in 2020

Peter Steinbach [*1] **Heidi Seibold** [*2] **Oliver Guhr** [*3]

Abstract

Faced by the abundant use of machine learning in industry and academia, the effective and efficient teaching of core concepts in this field becomes of high importance. For this, we organized a workshop on teaching methods in the field of machine learning. In this document, we summarize the current standing of the community as by our workshop and their methods. We touch on existing working concepts in machine learning didactics, what methods present initiatives use and cover open teaching resources available to date. With this, we hope to provide a starting point for future collaborations on this central topic given the expanding use of machine learning in science, industry and our daily lives.

1. Teaching Machine Learning workshop 2020

Many experts and practitioners who develop Machine Learning models or infrastructure around these models are confronted with the opportunity (or duty) to teach Machine Learning (ML) at some point in their career. Traditionally, many rely on their gut feeling to design courses that are motivated by these circumstances. The methods to design such material are often large collections of PowerPoint slides to convey the content and in-class interactive question and answer sessions.

At the virtual Teaching Machine Learning Workshop at ECML-PKDD 2020 (`https://teaching-ml.github.io/2020/`) targetted those who would like to know, how teachers from around the globe approach teaching ML. We hoped that all participants obtain a better feeling where they stand with their teaching and how they can improve or collaborate with others. Related to this, we expected for this workshop to be a beginning of an effort towards improved frameworks for ML literacy and competency.

Eleven speakers presented their ideas on teaching ML (videos available under `https://youtu.be/qGYGhpwG9iI`). Afterwards three breakout groups were formed to discuss three topics in-depth, the outcomes of which form the remainder of this article: Education research (Section 2), learning from others (Section 3), and open materials for teaching (Section 4).

2. What do we know about good ML teaching?

Machine learners believe in data and in using them to get a better understanding of the world. Does this also hold when it comes to their own teaching and education practice (similar to Gelman & Loken)? Data and research on what good ML education means is hardly available. Already in 2017 Amy J. Ko proclaimed that "We need to learn how to teach machine learning" in her insightful blog post on the topic. She argues that *pedagogical content knowledge (PCK)* about ML – such as "Effective analogies, examples, and explanations of machine learning" or "Knowledge of which concepts in machine learning are difficult and why"– is missing. In her workshop talk Rebecca Fiebrink added "Why teach this subject?" (see Grossman, 1990), "What should be taught?", and "How can/should technology be used in this teaching?" (see Mishra & Koehler, 2006) to the list of missing PCKs. She also showed a few examples of what research on ML education can look like. Fiebrink defines clear learning objectives (based on the theory of constructive alignment, see Biggs, 2003) and then tests her students to check if her teaching actually lead the students to reach these objectives. In one example she tested the learning objectives

- LO1: Understand the structure of supervised learning problems and capability of supervised learning algorithms
- LO2: Identify feasible uses of ML in new projects, and map a new project idea onto the structure of supervised learning (input, output, training data, model)

[*]Equal contribution [1]Helmholtz Zentrum Dresden-Rossendorf [2]Helmhotz Zentrum München [3]University of Applied Science (HTW) Dresden, Germany. Correspondence to: Peter Steinbach <p.steinbach@hzdr.de>, Heidi Seibold <heidi.seibold@helmholtz-muenchen.de>, Oliver Guhr <oliver.guhr@htw-dresden.de>.

Proceedings of the European Conference on Machine Learning 2020, Teaching Machine Learning Workshop, Ghent, Belgium, ECML-PKDD, PMLR 141.

by asking the students to "brainstorm a list of scenarios in which supervised learning could be used to make a piece of interactive art or music". 27 of the ideas students brought forward were feasible. However, 21 out of these 27 ideas were found solvable with programming only. They did not appear to necessarily require ML. This observations provides hints that concepts in ML are difficult to learn (see PCKs in Ko, 2017) and are potentially picked up unreflected with respect to the capabilities of ML in practice.

In the workshop's breakout group on *Education Research*, we discussed what would be needed to get research on ML education started. We collected related ressources, e.g. research on statistics and computer science education as well as on competence and literacy models in data literacy and computer science literacy (see Ziegler & Garfield, 2018; Chance et al., 2018; Guzdial, 2019). In the short time given we compiled a draft for a competence definition in ML:

Everyday literacy (Knowledge based on personal and communal experience)

- What is ML? What does it have to do with AI?
- What do the buzz words (big data, AI, ...) mean?
- What can ML do? What can it not do (or is not useful for)? What is ML used for in practice?
- Can ML be helpful for my problem/my data?
- What are good/useful data? What are potential data problems? (Questioning data)
- What are ethical issues of ML?

Applied literacy (Skill-based Literacy: Using a specific skill of know-how, based on acquired expertise)

- How to apply ML in practice with data?
- How to make decisions based on data / ML results?
- What are the key concepts of ML
- What is supervised / unsupervised ML and how do the differ?
- What are the algorithms / methods that form the basis of ML?
- How to evaluate ML models (cross-validation, measures, etc.)?
- How to run ML pipelines on the computer?

Theoretical literacy (Disciplinary knowledge)

- What is the mathematical, statistical and computational background of the methods?
- What are connections between different concepts/methods?
- What are open questions in ML research?
- How to adapt and implement algorithms in code?

Reflexive Literacy (Probing assumed and specialized knowledge systems)

- What is the reason for (not) ML?
- What are limitations of ML/methods?
- When does what (not) work?
- What is the effect of ML/research on society?
- What are the responsibilities of a machine learner?

This draft can be used as a basis for upcoming research on ML literacy and learning.

3. Learning from other teachers

Teaching material for ML has existed as long as ML was conceived to be useful to find patterns in data. In this role, it remained a highly academic topic for very long. With this, the audience being taught were mostly applied mathematicians, statisticians and advanced domain scientists. With the abundant availablity of data and compute resources and the subsequent renaissance of ML (marked by (Krizhevsky et al., 2012; Dahl et al., 2014) and others), the use of ML has also expanded to new audiences, e.g. software engineers (Howard & Gugger, 2020), artists (Fiebrink, 2020) or children and teenagers (Touretzky et al., 2019). With that, the requirements and goals of classical ML teaching content have gone beyond traditional limits. This spurrs the question, what ML teachers can learn from others?

3.1. Data Science Curricula

New and scalable approaches to teaching data science (as a prerequisite to ML) have been implemented (Teal et al., 2015) and made available as open-source. These lessons were made subject to data quality studies based on pre- and post-workshop surveys which showed substantial positive effects on the learning outcome (Jordan, 2016). With this, the carpentries exposes a primary community to draw experiences from. Among others, four main aspects of carpentries' style teaching were identified to be most informative to ML instructors:

Know your learners As (Wilson, 2014) indicates, learners (of ML) today come with a heterogenous mixture of learning backgrounds, identies and facilities. Knowing better these details before a workshop by means of pre-workshop surveys was found to be most instrumental. Especially with regards to the mathematical demands for learning ML, this can be readily adopted to ML.

Stay close to applications Learners invest considerable mental effort to improve their ML skills at workshops. In order to lower this investment, (Teal et al., 2015), (Garcia-Algarra, 2020) and (Wilson, 2014) use sandbox data to demonstrate and elaborate on pedagogical content which is inherently close to the learners domain or is easy to grasp. This minimizes the mental load on learners to first understand what to expect from the data and thus aids the learning process.

Live coding Popularized by (Wilson, 2014), live coding brings certain benefits to teaching complex topics programmatically. This offers benefits to reduce the teaching speed, watching code and results slowly emerge in contrast to having to understand readymade slides, fosters lateral knowledge transfer (seeing instructors fix problems or comment on questions immediately). While this is considered productive for content that aims to convey mechanistic knowledge, it has to be seen if this can aid at all the deeper understanding of mathematical cause and effects of ML systems.

Teach math/code without computers Aside from the focus on math and code for teaching ML, an ever growing part of the community motivates aspects of ML aside of these fields in curricula. As (Huppenkothen & Eadie, 2020) stresses, non-coding or non-math material can readily be used to teach learners important aspects of the supervised ML workflow starting from unbiased data taking, choosing features and being aware what a loss given these features optimizes for. Especially the latter can quickly lead learners to discussions of ethics and ML. Moreover, it can provide avenues for including audiences with little to no coding/math background as well as improve everyday literacy in a world with an ever increasing number of ML supported agents.

3.2. Setting goals

In general, the community agrees that setting learning outcomes for teaching is generally hard. In academic curricula, it is generally unclear where and how learners will apply ML knowledge. A similar observation must be made for massive open online courses. It is therefore considered hard to provide general guidance. For intense courses, the situation can be different.

In general, the carpentries' were again mentioned as learners take pre-workshop surveys (Carpentries', 2020) before a workshop. This way, instructors have a chance to infer indicators of learners backgrounds, their motivations and expectations. Moreover, results from these surveys can also hint towards learner personalities that facilitate a specific group dynamic.

Pre-workshop surveys aid teachers to establish a mental model of learner goals. These survey results help contrast and compare the learners present in a teaching activity with the learner profiles used to design the very lesson being taught upfront. In this way, a teacher is informed. Subsequently, the content being taught can be adapted if required and thus potentially ensure a more effective knowledge transfer.

In general, the community represented in the workshop appreciates material that offers a central plot for teaching topics, but also acknowledges that due to the heterogeniety of learners optional teaching modules and exercises are considered helpful. During a course, frequent feedback can uncover large scale misunderstandings among the learners which the instructors would love to react to. For this, modular lesson structures with more exercises than possible are considered an optimal design.

3.3. Synchronizing objectives and outcomes

In class, there are several methods to collect feedback from learners at high frequency which have been found to improve learning yield Sotola & Crede. At best, feedback and formative assessment is provided to the learner and instructor at the same time.

One possible implementation of these low-stackes testing is the use of sticky note supported multiple choice questions. For this, learners are provided a multiple choice questions with 4 possible answers. Every answer is marked by a distinct color. Each learner has been given four sticky notes of different color at the beginning of class. Once the multiple choice question is presented and read out loug, the instructor asks learners to present the color of the answer that they consider correct - all at once. This provides immediate optical feedback to the teacher. This forms hints to the learning outcomes for the teacher. For example, if the ensemble of sticky notes exposes a high degree of different colors, the instructor is assured, that learning outcomes are far from learning goals. The same holds for learners individually. Learners are now asked to discuss with their neighboring learner why they chose the color they did. This process of self-explanation by verbalisation, interpretation and adoption provides cognitive fortification of the content and can create higher adhesion inside the learning community.

It is considered good practice to collect feedback at regular intervals from each learner individually and anonymously. At the end of larger teaching blocks, learners are asked to write down something they liked or something they didn't like or didn't understand on two sticky notes (of green and red color). These are collected and a summary of the feedback is discussed in class at the beginning of the next module. In this fashion, the teacher exhibits trustworthiness as praise and deficiencies in teaching can be directly addressed. Moreover, this provides a handle on medium time scale feedback from learner to teacher and vice versa.

As teaching is very dynamic process, long term assessment is considered an essential instrument to monitor teaching quality and learning outcomes. However, many teachers commented that survey design for this purpose is well outside their scope of time and mandate. Many rely on contributed surveys by administrative organs in academia or student representations. Teachers in these context, that they consider this a good start but often no actionable consequences result from these surveys due to the way they are structured.

3.4. Balancing theory and practice

Teaching machine learning comprises three major fields from which to draw content from: computer science, math and statistics. Balancing theory and practice was considered a major challenge when designing lessons for ML practitions to statistics students as also reflected in Schwab-McCoy et al..

First and foremost, the community highlights that due to this (over)focus on content, topics like ethics, bias and suitability of ML are often overlooked. A consensus was reached that these highly relevant topics should be taught first before diving into any technical or mathematical details related to ML.

Related to theoretical content, there was no golden way identified how to teach this across learner profiles. There was agreement that ML can be taught without linear algebra knowledge, but this goal needs to be clearly communicated to learners.

In general, wide agreement was reached that learners need to be picked up from their domain problems. To ensure this, practicals should be made subject to ”kagglification”. This refers to the practice to teach learners how to approach ML problems by first identifying inputs and the desired output. And then and only then to study different algorithms. It was considered essential to stress this formality to have learners distance themselves in a productive fashion from the domain they are coming from.

A central recommendation from the community was to tailor the exercises to the needs or daily contexts of the learners. To be more concrete, it was suggested to have use cases for teaching that represent different data types (”image-like”, ”tabular”, ”time series”, ”text”, ...) that learners might use. Not only can this help to lower the cognitive load for learners (which will always be there if they e.g. are used to work with images and ML is being taught with text examples only), but also might fortify the content as they can related to the outcomes. Examples for this are given in (Garcia-Algarra, 2020) where learners are asked to provide their shoe size, body height and mass at the beginning of the ML class. This dataset is then used to train a regression model.

4. Open teaching resources

The advances in machine learning have enabled us to solve tasks that were previously time-consuming or impossible to accomplish. This created a high demand for machine learning knowledge as students and practitioners want to dive into this field. Teachers face the challenge, that they need to provide material and content that makes this complex field accessible for their students. At the same time, they constantly need to update their material in this fast-moving field.

In recent years a wide variety of open educational resources have been published. They range from single graphics and code examples to books and complete courses. This is an opportunity for students as well as for educators. Students can use this material freely, independent of there location, cultural background, gender or their financial resources. It lowers the barrier for people from all over the world to enter this field. These open materials are important resources for educators as well. During our workshop, we disscussed open educational resources that can help to create well recived, up to date courses for students and practitioners.

The participants stated that they use parts of open resources for there own teaching. They select explainations, graphics and extercieses and adapt them to the needs of their students. This way they become curators selecting the content and materials that has been reviewed and improved by other edutcators. This process helps to share best practices and improves the quality of the teaching materials.

Dive into Deep Learning (Zhang et al., 2020) is an introductory book to deep learning, with the concept to teach deep learning by implementing it. The code examples are written in Python for various deep learning frameworks, like Pytorch and Tensorflow. The book itself is open source and is developt by a community. It also covers quite recent topics and is used by a number of universities.

The Turing Way (The Turing Way Community, 2019) is an open-source community-driven book about reproducible data science. It covers topics from version control, testing to continuous integration. In 2020 guides for reproducible research, project design, communication, collaboration and ethical research were added.

Elements of AI (Elements of AI, 2018) is a free online course created by the University of Helsinki. It teaches the fundamental concepts of machine learning for a broader audience as it requires no computer science background. The course is divided into two parts. Part one discusses the applications and limitations of AI. Part two teaches the fundamentals needed to apply machine learning.

OpenTechSchool (OpenTechSchool, 2012) is a European organisation that offers free technical education and organizes online and local workshops. OpenTechSchool also offers free online courses for various topics like Python for Beginners, Data Processing with Python and Working with Databases.

Introduction to Machine Learning (Bernd Bischl, 2020) is an introductory course on applied supervised machine learning. It is a self-study course and the material consists of slide sets and lecture videos. Students can take multiple-choice tests and apply their knowledge in coding exercises. The code examples are written in R.

This list is not complete and soon becomes outdated as new content is published constantly. Therefore we created a repository Seibold et al. for everybody to look up and share open educational resources.

5. Summary

This workshop has gathered interested parties from the machine learning community and domain sciences who aspire to become effective and good ML instructors on very different levels of prior expertise as well as aspiration. The workshop underpinned, that there are motivated individuals out there and a large volume of existing material. We must observe though, that to date no community around the common goal of good ML teaching exists, nor is there momentum to exchange and discuss this material yet. Most initiatives currently live in contexts where collaboration is hard or not occurring.

Further, the workshop made it evident that more research is required in this field under high demand. ML systems are not straight forward 'machines' whose cause for malfunction is immediately evident. Hence, effective teaching can build the basis for well trained researchers, engineers or users that aid in the modern adoption of ML implementations and make the use of the latter sustainable for all of society.

References

Bernd Bischl, Fabian Scheipl, H. S. D. S. L. B. D. S. C. M. T. P. Introduction to Machine Learning, 2020. URL https://introduction-to-machine-learning.netlify.app/.

Biggs, J. Aligning teaching for constructing learning. *Higher Education Academy*, 1(4), 2003.

Carpentries', T. The carpentries pre-workshop survey, 2020. https://carpentries.github.io/assessment-archives/pre-workshop/pre-workshop.html.

Chance, B., Mendoza, S., and Tintle, N. Student gains in conceptual understanding in introductory statistics with and without a curriculum focused on simulation-based inference. In *Proceedings of the International Conference on Teaching Statistics*, volume 10, 2018.

Dahl, G. E., Jaitly, N., and Salakhutdinov, R. Multi-task neural networks for qsar predictions. *arXiv preprint arXiv:1406.1231*, 2014.

Elements of AI, U. o. H. Elements of ai, 2018. URL https://www.elementsofai.com/.

Fiebrink, R., 2020. URL https://youtu.be/qGYGhpwG9iI.

Garcia-Algarra, J. Introductory machine learning for non {stem} students. In *ECML PKDD 2020 Workshop Teaching ML*, 2020. URL https://openreview.net/forum?id=dhmIJ9SwAuG. accepted.

Gelman, A. and Loken, E. Ethics and statistics: Statisticians: When we teach, we don't practice what we preach. *CHANCE*, 25(1):47–48, 2012. doi: 10.1080/09332480.2012.668469. URL https://doi.org/10.1080/09332480.2012.668469.

Grossman, P. L. *The making of a teacher: Teacher knowledge and teacher education*. Teachers College Press, Teachers College, Columbia University, 1990.

Guzdial, M. Computing education as a foundation for 21st century literacy, 2019.

Howard, J. and Gugger, S. *Deep Learning for Coders with Fastai and PyTorch: AI Applications Without a PhD*. O'Reilly, 2020.

Huppenkothen, D. and Eadie, G. Teaching the foundations of machine learning with candy. In *ECML PKDD 2020 Workshop Teaching ML*, 2020. URL https://openreview.net/forum?id=a37DEwWs1wu. accepted.

Jordan, K. Data Carpentry Assessment Report: Analysis of Post-Workshop Survey Results, October 2016. URL https://doi.org/10.5281/zenodo.165858.

Ko, A. J. We need to learn how to teach machine learning, 2017. URL https://link.medium.com/wmzrUMMj3bb.

Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.

Mishra, P. and Koehler, M. J. Technological pedagogical content knowledge: A framework for teacher knowledge. *Teachers college record*, 108(6):1017–1054, 2006.

OpenTechSchool. OpenTechSchool e.V., 2012. URL https://www.opentechschool.org.

Schwab-McCoy, A., Baker, C. M., and Gasper, R. E. Data science in 2020: Computing, curricula, and challenges for the next 10 years. *Journal of Statistics Education*, pp. 1–17, 2020.

Seibold, H., Guhr, O., and Steinbach, P. eaching machine learning workshop at ecml-pkdd 2020, 2020.

Sotola, L. K. and Crede, M. Regarding class quizzes: a meta-analytic synthesis of studies on the relationship between frequent low-stakes testing and class performance, 2020.

Teal, T. K., Cranston, K. A., Lapp, H., White, E., Wilson, G., Ram, K., and Pawlik, A. Data Carpentry: Workshops to Increase Data Literacy for Researchers , October 2015. URL https://doi.org/10.2218/ijdc.v10i1.351.

The Turing Way Community, Becky Arnold, L. B. S. G. P. H. R. H. . K. W. *The Turing Way: A Handbook for Reproducible Data Science (Version v0.0.4)*. 2019. https://the-turing-way.netlify.app/welcome.

Touretzky, D., Gardner-McCune, C., Martin, F., and Seehorn, D. Envisioning ai for k-12: What should every child know about ai? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 9795–9799, 2019.

Wilson, G. Software carpentry: lessons learned. *F1000Research*, 3, 2014.

Zhang, A., Lipton, Z. C., Li, M., and Smola, A. J. *Dive into Deep Learning*. 2020. https://d2l.ai.

Ziegler, L. and Garfield, J. Developing a statistical literacy assessment for the modern introductory statistics course. *Statistics Education Research Journal*, 17(2):161–178, 2018.