Competences in Machine Learning

The order of competences that students need to learn in ML

by

Elmedin Burnik

Supervisor(s): dr. Gosia Migut, prof. dr. Marcus Specht

A Dissertation Submitted to EEMCS

faculty Delft University of Technology,

In Partial Fulfilment of the Requirements

For the Bachelor of Computer Science and Engineering

22-06-2022

Abstract

Machine learning is becoming more and more applied within business and academia alike. This has led researchers to look inwards and discuss whether the current way of teaching and learning machine learning is the right way. Within this train of thought, one must investigate the characteristics of teaching and learning. Namely, what are the competences required to learn and in what order should these be learned. This research paper focusses on realizing the answer to that question by combining past research papers on how to define competences, how to establish an order of competences and how to establish competences through the help of questionnaires answered by academics within the machine learning sector teaching Computer Science students. The results from the academic analysis and the resulting data from the questionnaires are combined to organize a result, established through two different methods used by researchers. To conclude, the paper is reflected upon, and its merits, demerits and trade-offs are summarized, presenting a final recommendation for any future work done on this topic.

1 Introduction

Machine learning (ML) has become a major asset for companies and academia alike. Querying for “Machine Learning” on LinkedIn jobs shows the clear demand for ML trained individuals, yielding 5,000 job results in the Netherlands[1] and 350 of those in Delft alone[2] as per May 2022.

Apart from the non-academic demand for ML-trained individuals, there is also a large demand for ML within academia. This can be seen in recent papers utilizing ML in one way or another. A good example is the recent research paper by Abidi et. al. [2], wherein they apply ML techniques with the aim of increasing the probability of popularity prediction of movies. One other example that aligns with the subject of this research is the article by Vellido [13]. Herein, they aim to find out how machine learning can be applied to visualization and application within medicine and health care. Another very clear statistic is the large increase of publication under the AI umbrella, an increase of five times between the year 2010 and 2021 [14]. The number of AI patents filed in 2021 is also thirty times higher than that it was in 2015, showing an annual percentage yield of 76.9% [14]. See figures 1 and 2 for visualization of this matter.

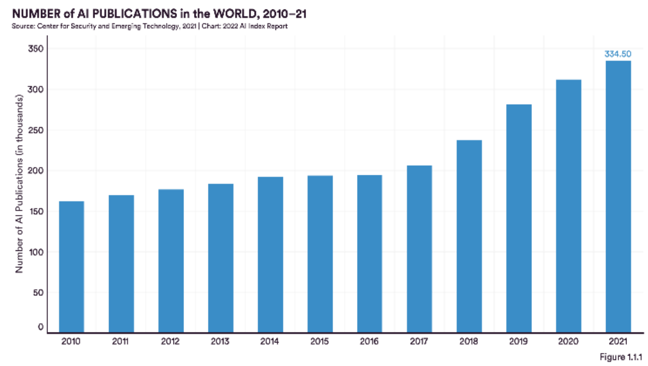

Figure 1: Number of AI publications in the world between 2010 and 2021 from the center of Security and Technology AI Index Report 2022 [14].

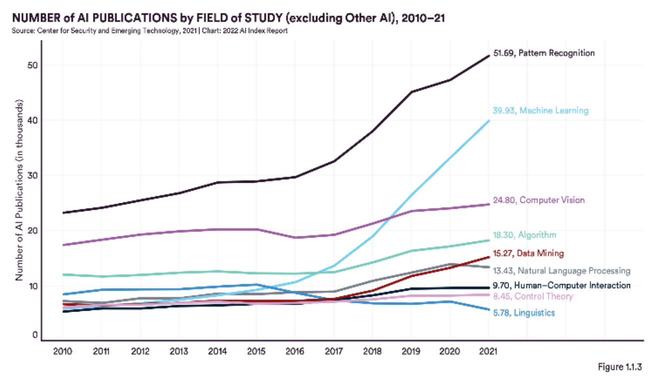

Figure 2: Number of AI publications by field of study between 2010 and 2021 from the center of Security and Technology AI Index Report 2022 [14]. Showcasing the rise of ML publications within AI.

This leads us the research question of the paper, which follows: *“What is the order of competences that students need to learn in ML?”*. To understand this question better we immediately answer the first sub-question of what the definition of a competence is. This definition is identified and provided by ECTS Users’ Guide *“A combination of attributes in terms of knowledge and its application, skills, responsibilities and attitudes.” [1]*. The main question is to be asked for undergraduate Computer Science majors. To answer this question, first we must ask ourselves four sub-questions. Which are:

i. “What is classified as a competence?”
ii. “What are the competences that students need to learn in introductory ML from a teacher’s perspective?”
iii. “What is the hierarchy of competences in ML?”

Answering these sub-questions will yield greater insight into the main research question and how to tackle it. As such the following chapters will go more deeply into each section. Chapter 2, Methodology will explain the ways in which we have obtained what is classified as competences, what these competences are in undergraduate Computer Science majors, how we will conduct the questionnaire and lastly how we will discuss the results. The following chapters will go more deeply into each specific task. Chapter 3 discusses how the definition of what a competence is, our resulting list of ML competences and an analysis of the order dependencies within these ML competences. Chapter 4 will showcase the findings of the performed questionnaire and discuss the results of this conducted experiment by analyzing them and providing concrete examples of introductory machine learning course improvements. Chapter 5 discusses the ethical aspects of this paper and the reproducibility of this paper’s methods. Lastly, chapter 6 will give a summary of the (main) research question(s) and provide a conclusion to the research paper. Open issues, possible improvements and any new questions that arise from this work will be discussed herein.

2 Methodology

To answer the question of *“The order of competences that students need to learn in ML.”* Multiple things must be done before answering this question.

First, we determine what is considered as a “competence”. Then, we look to the past papers on the importance of machine learning and the hierarchy of competences, concepts, and attributes within it. We investigate similar papers on other subjects close to machine learning if available. From this list, the most important and useful papers will be discussed in more detail. After determining the meaning of what a competence is, the next step is to gather the competences required for ML. To do this, the definition of what a competence is, is investigated firstly. Afterwards, an overview of introductory machine learning courses is made and from this, the intersection of all the different competences is created. To acquire this overview, we must look up the course guides of multiple undergraduate Computer Science ML courses from multiple universities. Since the scope of this research paper is aimed at TU Delft and its machine learning course for the undergraduate Computer Science and Engineering major, this course guide is used as a basis for the creation of the competence list. From here, the same method of finding competences is used with the following universities: Leiden, UvA, VU, Groningen. These universities’ course guides were chosen because they are the only ones in the Netherlands that offer introductory ML courses to its undergraduate Computer Science students. To extend the list, the same amount of course guides from universities outside of the Netherlands are gathered as well. These consist of: ETH Zurich, Cambridge, Oxford, and Edinburgh universities. From this definition of what a competence is, a list of competences was found in a paper by Danyluk [4]. This list intersected with the list of course guide competences from earlier form a new list which can be viewed in table 1. This list of competences was identified through surveys/questionnaires aimed at academics whether their institutions offered any form of data science programs at an undergraduate level [4]. This is the resulting list that is used in the following steps.

Following the method of Danyluk et al. [4] and the scope of this paper, the following step is the creation of a questionnaire wherein participants are asked to create a sequenced order of the competences from the concluded list of competences. The participants to fill in the questionnaire follow the example set by Danyluk et al. [4] and are academics that teach the course of machine learning to undergraduate Computer Science students. Since the scope of this paper is focused on TU Delft, the teachers in question are also those of TU Delft.

Lastly, we discuss possible improvements in future machine learning courses based on the resulted hierarchy and order of ML competences.

3 Analysis of ML Competences in literature

With setting up the steps for finding the competences for an undergraduate Computer Science introductory ML course, we must first define the meaning of a competence within the ML paradigm. This definition is identified and provided by ECTS Users’ Guide *“A combination of attributes in terms of knowledge and its application, skills, responsibilities and attitudes.” [1].* From this definition we move on to finding the competences within our chosen universities and their respective (introductory) ML course guides. Hereafter, we analyze our findings and discuss its dependencies and hierarchy.

There are two methodologies for finding and defining competences for course curricula. The first methodology is a normative approach. The competence model is the product of a broad discussion among CS domain experts based on a generic psychological competency model, which was agreed upon and proven. In this context, a German National Computer Society (GI) CS expert panel established a CS competence model to issue design recommendations for CS curricula in Higher Education (see Methodology A in figure 3) [7].

The second technique, Methodology B in figure 3, is an experimentally driven approach that generates the competency model from content analysis of international CS curricula. The used content analysis reveals common CS-competences addressed in various university CS study programmes. An empirically grounded competence model in software engineering, software development and programming form the basis of competence measurement instruments, which can be used for purposes of diagnosis, students' self-assessment and for affirmative action [7].

For this paper, we use a mix of both derived methods. A list of competences is formed by searching through the course guides of various machine learning undergraduate courses from various universities combined with the results from the research performed by Danyluk et al. which has established a set of 12 competences for undergraduate data science student, which also includes machine learning [4]. By doing it this way, we achieve both normative (Method A) and experimental (Method B) approaches shown visually in figure 3 [7].

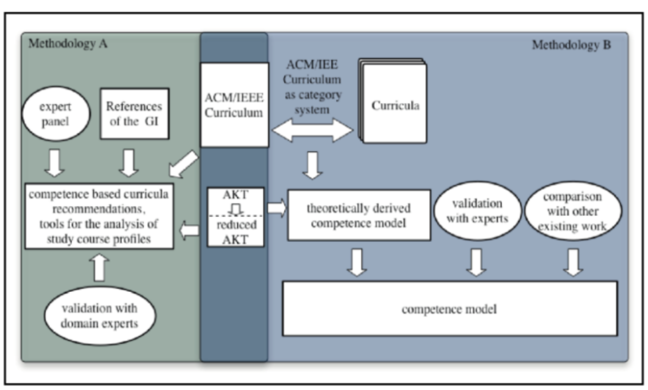

Figure 3. Methodologies for competence model development [7].

The search starts at the Machine Learning course from the undergraduate Computer Science and Engineering course from TU Delft [A]. Results were taken from the noted course objectives and course coverage. The course objective is taken as the competence aimed to achieve as well as its attitude, and the course coverage is used as descriptive explanation for the competence's attributes and abilities. The same data is taken from the universities' undergraduate Computer Science machine learning courses mentioned in chapter 2.

From the paper of Danyluk et al. [4] a set of 12 competences are introduced for machine learning. These results will be compared to the competences assembled from course guides from the universities mentioned in chapter 2. The intersected result can be viewed in table 1. Thus, having used the two methodologies as presented in figure 3.

1	Characterize and differentiate between different classes of machine learning models (e.g., geometric, probabilistic) and tasks (e.g., classification, clustering, regression).
2	Derive a (current) learning algorithm from first principles and/or justify a (current) learning algorithm from a mathematical, statistical, or information-theoretic perspective.
3	Apply appropriate empirical evaluation methodology to compare ML algorithms/tools to each other.
4	Compare differences in interpretability of learned models.
Ex:	parametric and non-parametric density estimation; linear and non-linear classification; unsupervised learning including clustering and dimensionality reduction.
5	Select and apply a broad range of ML tools/implementations to real data.
6	Implement ML programs from their algorithmic specifications.
7	Understand and implement expertise in the Python programming language and its statistical and numerical libraries.
Ex:	Idem ditto above examples but then in application; python (with necessary libraries).
8	Exhibit knowledge of methods to mitigate the effects of overfitting and curse of dimensionality in the context of ML algorithms.
9	Provide an appropriate performance metric for evaluating ML algorithms/tools for a given problem
10	Apply appropriate empirical evaluation methodology to assess the performance of a ML algorithm/tool for a problem.
Ex:	Performance evaluation of predictive algorithms; non-negative matrix factorization and outlier detection.
11	Express formally the representational power of models learned by an algorithm, and relate that to issues such as expressiveness and overfitting
12	Explain the concept of and identify (implicit) bias in data and ML algorithms.

A: Appendix A

13	Consider and evaluate the possible effects – both positive and negative – of decisions arising from ML conclusions.
Ex:	Ethical issues in machine learning; dealing with imbalanced datasets.

Table 1. The intersection of competences from Danyluk [4] and the aforementioned list of university course guides.

From here, we look at ways we can order it through already established methods and results. A paper by Von Wangenheim et al. and another by Leidig and Cassel, we find that they have already established a certain order of learning objectives [3][6]. Taking this order and comparing with our competence list, we take the comparable competences and from here we can create the following diagram of hierarchy presented in figure 4 below.

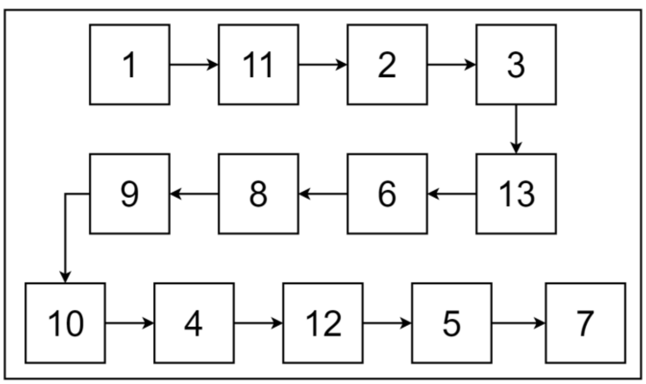

Figure 4. Competence order established through comparison with table 1 and results from [3] and [6].

These results come up in chapter 4 and are compared with the results gained form the questionnaire performed on machine learning academics teachers. This questionnaire yields data in the form of an ordered list of competences based on their opinions as machine learning teachers in academia. Further explanation on how the questionnaire is performed and how the data is used is explained in further detail in the experiment section of chapter 4. A conclusion is drawn in chapter 6 showcasing the final list of ordered competences.

4 Experiment Results and Discussion

The experiment is setup by various steps discussed in earlier chapters. Competences for machine learning are investigated and established. From here, the list is randomized before presenting it to machine learning academics teachers such that they can give their answer of what they believe the order should be. This result is then compared to each other first and then against the analysis result conducted in chapter 3.

Questionnaires should be relevant, understandable, unambiguous, unbiased, capable of handling all conceivable responses, well coded, piloted, and ethical. Decide what data you need, pick items for inclusion, design individual questions, compose the language, design the layout and presentation, consider coding, produce the first draft and pretest, pilot, and evaluate the form, then conduct the survey are the major phases in building a questionnaire [12]. To inform the participants of the questionnaire to its fullest extent, we establish a proper form of conduct that explains how the data collected will be used and if they accept these terms and conditions. Furthermore, definitions of all the various terms are explained and examples are given before the question is asked. Next, a short description is given on how to answer the question so that it can most easily be collected. The participating teachers were asked to type their ordered sequence in the following example: 1, 5, 12, 4, etc.

After collecting the result data from the various participants, a method of creating a final result is needed. We use the Kendall-τ distance method [11] between the ranking significance ratings of all input tokens across all examples to gauge agreement across the competences, as described by Jain and Wallace [10]. The Kendall-τ distance is a metric (distance function) that calculates how many pairwise disputes there are between two ranking lists. The greater the gap between the two lists, the more dissimilar they are. If two items are ranked differently on two lists, they are deemed discordant [11]. In terms of rankings, the Kendall-τ distance is as follows: A permutation (or ranking) is a set of N integers in which each integer from 0 to N-1 appears exactly once. The number of pairings in the two rankings that are in different order is the Kendall-τ distance between them [15]. The Kendall-τ distance method is used here to determine the distance of the competences within the individual results. From here, a new and next result is created using the Kendall-τ distance method [11]. What this means is, each element in the ordered list is given a rank. We combine the results from both lists and apply the Kendall-τ method to get a result of how similar the lists are. We then reorder the new list such that this one has a more similar distance. Applying this method yields the resulting table seen in table 2. The competences are re-ordered and thus differ from table 1.

1	Characterize and differentiate between different classes of machine learning models (e.g., geometric, probabilistic) and tasks (e.g., classification, clustering, regression).
9	Provide an appropriate performance metric for evaluating ML algorithms/tools for a given problem.
11	Express formally the representational power of models learned by an algorithm and relate that to issues such as expressiveness and overfitting.
2	Derive a (current) learning algorithm from first principles and/or justify a (current) learning algorithm from a mathematical, statistical, or information-theoretic perspective.
3	Apply appropriate empirical evaluation methodology to compare ML algorithms/tools to each other.
8	Exhibit knowledge of methods to mitigate the effects of overfitting and curse of dimensionality in the context of ML algorithms.
6	Implement ML programs from their algorithmic specifications.
12	Explain the concept of and identify (implicit) bias in data and ML algorithms
13	Consider and evaluate the possible effects – both positive and negative – of decisions arising from ML conclusions.
4	Compare differences in interpretability of learned models.
5	Select and apply a broad range of ML tools/implementations to real data.
10	Apply appropriate empirical evaluation methodology to assess the performance of a ML algorithm/tool for a problem.
7	Understand and implement expertise in the Python programming language and its statistical and numerical libraries.

Table 2. The resulting ordered list of competences from the questionnaire results applied with the Kendall-τ distance method [11]. The numbered competences remain the same as in table 1.

We can use the same Kendall-τ distance method for comparing the resulting order from figure 4 against the questionnaire results from table 2. By doing so, we establish a definitive list of competences estimated through the various steps taken. This same process is applied, and the resulting order can be viewed in figure 5 below.

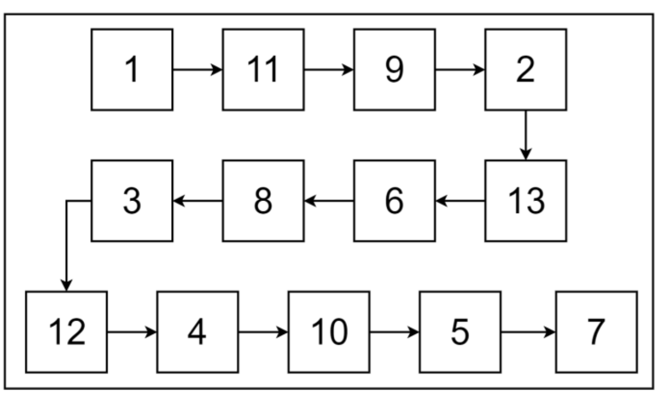

Figure 5. Competence order established through the Kendall-τ distance method applied on figure 4 and table 2.

Currently, machine learning courses are mostly setup in the conventional way: lectures explaining theoretical concepts and assignments where students are asked to (re)create the concepts explained in the lecture they viewed earlier that week. This convention can be changed in a myriad of ways to increase engagement and learning in students. A common discussed method is that of Flipping the classroom [17]. Flipped classroom approaches have students use technology to access the lecture and other instructional resources outside the classroom in order to engage them in active learning during in-class time [16]. A variation of this method is called Skill circuit [18]. Herein, courses are divided into various competences and within each competence, students are guided through a list of various lecture videos, articles, book chapters, computer assignments, and written assessments and quizzes. Students can choose their "path" through the circuit based on their skill level or how quickly they wish to go through the entire skill circuit [18]. This allows for more freedom and responsibility in the students' hands. As shown by Giannos et al., this derivation, which is based on flipping the classroom, can increase the learning comprehension and engagement of students during their courses [16].

Combining the results from figure 5 and table 2 with the course concept of a skill circuit, the machine learning course as it currently stands (lectures in lecture halls with assignments at the end of the work week) can be improved. The course designer can divide the course into the competences and order presented earlier and within each competence add all the various reading and viewing material, computer assignments, and written assessments and quizzes. Afterwards, the course designer can create various paths that go from minimal knowledge needed to possibly extra information for students wishing to learn more and develop a deeper understanding of machine learning as also presented in the skill circuit examples from [18].

5 Responsible Research

Researchers must adhere to proper behavior when disseminating their research in all elements of academic writing. When preparing and submitting publications to peer-reviewed journals, relevant ethical procedures must be considered. The authorship problem is one of the most important ethical difficulties to consider in academic writing. During the publication process, ethical questions should be raised. Etiquette, fraudulent publication, plagiarism, duplicate publication, authorship, and potential for conflict of interest are all issues of ethical code in literary publication. To prevent ethical infractions such as integrity, honesty, truth, and transparency, some solutions must be devised. When writing a paper for publication, these tactics will improve ethical integrity [19].

This paper has expressed that the competences presented in table 1 are the only competences that exist for machine learning in undergraduate Computer Science programs. However, this is not necessarily the case and has various biased takes from its writer. One recreating this paper could have concluded a set of different competences or have worded it differently. As well as this, the paper also mentions that a skill circuit is a proposed improvement for the current existing course structures machine learning courses. This is not necessarily the only possible proposal and has not been thoroughly researched and peer-reviewed to state otherwise.

The academics teachers asked for the questionnaire have been informed what the questionnaire has been for, what the topic of the research is, what the questionnaire is about and how their data will be used afterwards. The conclusion was reached through the combination of questionnaire results and analysis of other papers.

Reproducibility has been a major factor in this paper. Every step of the way, the results have been shown and the methods have been clearly explained and proper references are used that easily explain these methods in more detail. Thus, this paper has achieved a proper and robust experimental design.

6 Conclusion and Future

This paper has answered the question of: *"What is the order of competences that students need to learn in ML?"* as well as the sub-questions that arise. Mainly, what the definition of a competence is, what the competences for machine learning for undergraduate Computer Science courses are, and how to rank these competences in an order.

The research has established the results which can be viewed again in table 3. The numbers are based on the original numbering in table 1.

1	Characterize and differentiate between different classes of machine learning models (e.g., geometric, probabilistic) and tasks (e.g., classification, clustering, regression).
11	Provide an appropriate performance metric for evaluating ML algorithms/tools for a given problem.
9	Express formally the representational power of models learned by an algorithm and relate that to issues such as expressiveness and overfitting.
2	Derive a (current) learning algorithm from first principles and/or justify a (current) learning algorithm from a mathematical, statistical, or information-theoretic perspective.
3	Apply appropriate empirical evaluation methodology to compare ML algorithms/tools to each other.
8	Exhibit knowledge of methods to mitigate the effects of overfitting and curse of dimensionality in the context of ML algorithms.
6	Implement ML programs from their algorithmic specifications.
13	Explain the concept of and identify (implicit) bias in data and ML algorithms
12	Consider and evaluate the possible effects – both positive and negative – of decisions arising from ML conclusions.
4	Compare differences in interpretability of learned models.
10	Select and apply a broad range of ML tools/implementations to real data.
5	Apply appropriate empirical evaluation methodology to assess the performance of a ML algorithm/tool for a problem.
7	Understand and implement expertise in the Python programming language and its statistical and numerical libraries.

Table 3. The resulting ordered list of competences from figure 5. The numbered competences remain the same as in table 1.

The skill circuit method explained in the earlier section has been proposed as an improvement to the current structure of machine learning courses. However, this has yet to be researched more thoroughly and is only an opinionated improvement based on past research done on flipping the classroom, as discussed in chapter 4.

Further research on this topic can be done by dissecting the results from table 3 and focusing on subsections within each competence to narrow down the sub-competences, if any exist. Within each competence, there can be more competences or sub-competences that can be (re)ordered again. This is something that yet remains to be researched in the future.

References

[1] European Commission, Directorate-General for Education, Youth, Sport and Culture, (2009). ECTS users' guide, Publications Office. https://data.europa.eu/doi/10.2766/88064

[2] Abidi, S. M. R., Xu, Y., Ni, J., Wang, X., & Zhang, W. (2020). Popularity prediction of movies: from statistical modeling to machine learning techniques. *Multimedia Tools and Applications*, *79*(47–48), 35583–35617. https://doi.org/10.1007/s11042-019-08546-5

[3] Gresse Von Wangenheim, C., Marques, L. S., & Hauck, J. C. R. (2020). Machine Learning for All – Introducing Machine Learning in K-12. Machine Learning for All – Introducing Machine Learning in K-12. https://doi.org/10.31235/osf.io/wj5ne

[4] Danyluk, A., & Buck, S. (2019). Artificial Intelligence Competencies for Data Science Undergraduate Curricula. *Proceedings of the AAAI Conference on Artificial Intelligence*, *33*, 9746–9747. https://doi.org/10.1609/aaai.v33i01.33019746

[5] de Veaux, R. D., Agarwal, M., Averett, M., Baumer, B. S., Bray, A., Bressoud, T. C., Bryant, L., Cheng, L. Z., Francis, A., Gould, R., Kim, A. Y., Kretchmar, M., Lu, Q., Moskol, A., Nolan, D., Pelayo, R., Raleigh, S., Sethi, R. J., Sondjaja, M., . . . Ye, P. (2017). Curriculum Guidelines for Undergraduate Programs in Data Science. *Annual Review of Statistics and Its Application*, *4*(1), 15–30. https://doi.org/10.1146/annurev-statistics-060116-053930

[6] Leidig, P. M., & Cassel, L. (2020). ACM Taskforce Efforts on Computing Competencies for Undergraduate Data Science Curricula. Proceedings of the 2020 ACM Conference on Innovation and Technology in Computer Science Education. https://doi.org/10.1145/3341525.3393962

[7] Brinda, T., Reynolds, N., Romeike, R., & Schwill, A. (2015). KEYCIT 2014. Beltz Verlag.

[8] Lindvall, M., Molin, J., & Löwgren, J. (2018). From machine learning to machine teaching. *Interactions*, *25*(6), 52–57. https://doi.org/10.1145/3282860

[9] Long, D., & Magerko, B. (2020). What is AI Literacy? Competencies and Design Considerations. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. https://doi.org/10.1145/3313831.3376727

[10] Jain, S., & Wallace, B. C. (2019). Attention is not Explanation. Proceedings of the 2019 Conference of the North. https://doi.org/10.18653/v1/n19-1357

[11] KENDALL, M. G. (1938). A NEW MEASURE OF RANK CORRELATION. Biometrika, 30(1–2), 81–93. https://doi.org/10.1093/biomet/30.1-2.81

[12] Stone, D. H. (1993). Design a questionnaire. *BMJ*, *307*(6914), 1264–1266. https://doi.org/10.1136/bmj.307.6914.1264

[13] Vellido, A. (2019). The importance of interpretability and visualization in machine learning for applications in medicine and health care. *Neural Computing and Applications*, *32*(24), 18069–18083. https://doi.org/10.1007/s00521-019-04051-w

[14] Daniel Zhang, Nestor Maslej, Erik Brynjolfsson, John Etchemendy, Terah Lyons, James Manyika, Helen Ngo, Juan Carlos Niebles, Michael Sellitto, Ellie Sakhaee, Yoav Shoham, Jack Clark, and Raymond Perrault, “The AI Index 2022 Annual Report,” AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University, March 2022. https://aiindex.stanford.edu/wp-content/uploads/2022/03/2022-AI-Index-Report_Master.pdf

[15] Kumar, R., & Vassilvitskii, S. (2010). Generalized distances between rankings. Proceedings of the 19th International Conference on World Wide Web - WWW ’10. https://doi.org/10.1145/1772690.1772749

[16] Giannakos, M. N., Krogstie, J., & Chrisochoides, N. (2014, November). Reviewing the flipped classroom research: reflections for computer science education. In Proceedings of the computer science education research conference (pp. 23-29).

[17] Brame, C. (2013). Flipping the classroom.

[18] Skill circuit. (n.d.). Skill Circuit. Retrieved June 14, 2022, from https://ad-cs.ewi.tudelft.nl/#/

[19] Budianto, L. (2020, December). LAW AND ETHICS IN ACADEMIC WRITING.

Appendix

A. List of universities' course guides used:

Delft University of Technology	https://studiegids.tudelft.nl/a101_displayCourse.do?course_id=57344
Leiden University	https://studiegids.universiteitleiden.nl/en/courses/105108/introduction-to-machine-learning; https://studiegids.universiteitleiden.nl/en/courses/109040/machine-learning
Vrije Universiteit Amsterdam	https://studiegids.vu.nl/en/Bachelor/2020-2021/computer-science/X_400154
Universiteit van Amsterdam	https://studiegids.uva.nl/xmlpages/page/2021-2022/zoek-vak/vak/92172; https://studiegids.uva.nl/xmlpages/page/2021-2022/zoek-vak/vak/93445; https://studiegids.uva.nl/xmlpages/page/2021-2022/zoek-vak/vak/89354; https://studiegids.uva.nl/xmlpages/page/2021-2022/zoek-vak/vak/89421
University of Groningen	https://www.rug.nl/ocasys/fwn/vak/show?code=WBCS032-05
ETH Zurich	http://www.vvz.ethz.ch/Vorlesungsverzeichnis/lerneinheit.view?semkez=2022W&ansicht=ALLE&lerneinheitId=163859&lang=en
University of Cambridge	https://www.cl.cam.ac.uk/teaching/2122/MLRD/
University of Oxford	https://www.cs.ox.ac.uk/teaching/courses/2021-2022/ml/