

Teaching AI Ethics: Observations and Challenges

Marija Slavkovik, University of Bergen Norway
marija.slavkovik@uib.no

Abstract

This report summarises the experience in teaching Artificial Intelligence (AI) Ethics as an elective masters level course at the University of Bergen. The goal of the summary is twofold: 1) to draw lessons for teaching this in-high demand very new discipline; 2) to serve as a basis in developing a bachelor level AI Ethics course for students of artificial intelligence. AI Ethics as a topic is particularly challenging to teach as the discipline itself is very new and no textbooks have been established. The added challenge is introducing methodologies and skills from humanity- and social sciences to students of computational and information sciences.

1 Introduction

On January 14, 2020, the Norwegian National strategy on Artificial Intelligence was presented by the [18]. The Norwegian national strategy follows a trend in national, international and professional guidelines and recommendations on how artificial intelligence (AI) systems should be developed and used¹. The national strategy specifically stipulates: “Algorithms can be controlled by facilitating access or audit, but it is more appropriate for developers as well as users to build privacy and ethical considerations into systems from the outset.”[18, pg. 60].

At present, courses in artificial intelligence and computer programming are typically offered to students of “technical” studies such as computer science, informatics, engineering and information science, whereas subjects which develop skills that would help one to recognise and understand issues of “privacy and ethical considerations” are typically offered to students in humanities and social sciences. In other words, in Norway, and generally in the world, it is not part of the higher education tradition to train either “developers” or “users” to build “privacy and ethical considerations into systems” [21].

The Department of Information Science and Media Studies (Infomedia) at the Faculty for Social Sciences at the University of Bergen offers the course Research Topics in Artificial Intelligence (code INFO381) as a 15 point masters course. The course is intended to be tailored by the lecturer to present a specific topic of research in artificial intelligence from that lecturer’s area of expertise. I have been given the opportunity to teach INFO381 in the spring semester of 2017 and again in 2020. Having an active research interest in AI Ethics, I have designed a INFO381 variant to give masters students an introduction the state of the art and main challenges in AI Ethics research.

This paper was presented at the NIK-2020 conference; see <http://www.nik.no/>.

¹See for example [1] for an exhaustive list of guidelines.

AI Ethics has gradually emerged as a field of artificial intelligence in the past fifteen years. The field of AI Ethics is an active research field, however it is not sufficiently established to offer an indisputable curriculum for teaching it. At present, there is an increasing interest to offer AI courses to future bachelors of informatics, information science and computer - if we are to follow the national guidelines on AI, the skills in AI have to be complemented with awareness of responsible use.

The main motivation for this report is the question: **How to design a bachelor level course on AI Ethics for the curriculum of future AI developers, researchers and data scientists?** Specifically, I investigate how to leverage my experience in teaching AI ethics at the masters level to design a bachelor level course on the same topic.

This document is structured as follows. In Section 2 I detail the topic of AI, AI ethics, and I give an overview on the state of education in AI ethics in the world. In Section 3, I outline how AI ethics has been taught in INFO381 in 2017 and outline the lessons which were used for planning the same course in 2020. In Section 4, I detail the observations from teaching INFO381 in 2020, the lessons learned, and how those lessons can be leveraged into a plan for an bachelor level AI ethics course. In Section 5, I outline the bachelor AI Ethics course. Lastly, in Section 6, I summarise this report and outline the main insights.

The major shortcoming of this report is that it describes subjective experience rather than a quantitative study, however given that there is very little AI Ethics education and a lot of pressure to create it, this report still has a contribution to make.

2 Background

AI: what it is and how it changed

Artificial Intelligence (AI) is an interdisciplinary research and application field rooted in computer science, mathematics, and information science. It traditionally also involves research in philosophy, cognitive science, social science, economy and law. AI is concerned with issues of creating intelligent behaving computational agents and understanding what makes a computational agent behaviour intelligent [23]. AI can also be seen as the science and practice of automating cognition [3].

Artificial intelligence has been established as a research field in 1956 [19]. Prominent applications of methods developed in artificial intelligence have started to appear already in the 1980s. Like all computing at the time, software and autonomous devices were seen as tools which are developed by professionals to be used by professionals. Thus all users and developers were fully educated into the abilities of the AI system they worked with. Furthermore, the contexts in which a given system was used were known and controlled. In other words, both software and hardware existed in a so called “working envelope” segregated from society with precisely defined area of operation and human access control.

Since the new millennium, computing is no longer something that happens on computers in front of us and computers are no longer only a professional work tool. Computing and computers are now ambient and ubiquitous in the sense that a person is not necessarily aware they are interacting with a computer or that they contribute to computing. Supported by an increase of computing power and availability of structured data, branches of AI research become applicable into many domains in solutions that now cater not only to the train professional but to the general public. Specifically the sub-areas of machine learning, vision, and natural language processing require a lot of untrained cognitive labour to create examples used to “train” algorithms to recognise faces

in images, translate from one language to another, transcribe speech to text, dynamically price travel tickets, etc.

In the last decade we see a reinforcement in the general trend of computation leaving the “working envelope”. There is now unconstrained interaction between software and people. Programming is done not only by formally trained professionals, but also by self-taught people and enthusiasts. AI systems now replace and augment human cognitive activities such as decision-making in a variety of domains. With this trend, concerns arise on how to ensure that the AI systems developed are aligned with the values of its users, developers and the society in which they are deployed [8].

AI Ethics

AI ethics is the common reference to a collection of sub-fields in AI developed to respond to the issues of how to manage the moral, personal and societal impact of replacing people tasks and roles with AI powered computing. AI Ethics comprises of four main research sub-disciplines: fair-accountable-transparent AI (FAccT), explainable AI (XAI), responsible AI, and machine ethics (also called artificial morality).

Bovens defines accountability as “a relationship between an actor and a forum, in which the actor has an obligation to explain and to justify his or her conduct, the forum can pose questions and pass judgement, and the actor may face consequences.” [4, p. 447]. In AI ethics we focus on *algorithmic accountability* which specifies an accountability relationship in which the actor gives account for an algorithm which they may or may not have created [27]. An algorithm is a model of software that describes instruction for a computer. The role of the author can be filled by anyone involved in the algorithm’s creation and deployment. The role of the forum can be filled by different stakeholders that are affected by the algorithm’s operation, such as: users, government supervisory bodies, other companies etc.

Fairness motivates algorithmic accountability. The field is concerned with ensuring that an algorithm is equally efficient for all users and that the continuous use of an algorithm, particularly in a decision-making role, does not introduce a source of discrimination in society [16]. The complementary concern of the field are the various types of fairness that one can define or be interested in protecting, and also how do these different fairness requirements interact with each-other and with the computational efficiency of the algorithm itself.

Transparency refers to the algorithm that is the subject of accountability and the problem of making sure that the forum is able to inspect and understand the algorithm itself [26]. This field is also concerned with establishing a framework for defining different types of transparency but also with the question of: how to write algorithms that can be automatically audited?

Explainable AI (XAI) refers to the account artefact in the algorithmic accountability relationship [12]. One aspect of this sub-field is understanding what makes an account an explanation for a particular actor and forum [17]. The other aspect is developing algorithms that can generate explanations for their operations or outputs, and developing algorithms that can interpret the operations or outputs of complex algorithms [2].

Responsible AI is concerned with ensuring that the forum has sufficient power over the actor in the algorithmic accountability relationship. Unlike products that you can touch, software and data can exist and be used in more than one country at the same time. This makes legal regulation of AI very difficult. The field has numerous sub-interests that include: establishing a professional code of ethical conduct for AI researchers,

developers and practitioners; developing strategies for assessing the ethical impact and value alignment of a particular AI application. Responsible AI is also about overseeing the development of AI guidelines and ensuring they are meaningful and actionable [14, 13].

Machine ethics [20], or artificial morality [25], unlike the other AI Ethics sub-fields studies how to program moral behaviour in machines and software. The questions studied in this field include, but are not limited to: what moral theory to implement, should moral behaviour be learned or hard-coded, what characterises ethically sensitive decisions and contexts, how to decide what moral values should a machine enforce.

All of these different sub-fields of AI Ethics have been developing at different rates. The communities of researchers and practitioners is not well connected spanning many scientific disciplines. However, the interests of the sub-fields are clearly related and at points overlap, therefore an introductory level course in AI Ethics should include them.

Global education in AI Ethics

AI as a subject is typically only thought at an university undergraduate level. It is now certain that AI applications are a socially disruptive technology. On one hand the AI applications change aspects of society and it is in our own interest to know how. For example, we may want to be aware that we are interacting with an AI customer service decision-maker not a human and adjust our communication strategy accordingly. On the other hand, AI applications directly displace jobs changing the range of required skills that make one employable. Education is one of the tools we have to ensure that we understand what we are disrupting with a particular AI application, but also to ensure that we enable people to do the jobs of the future. In response to this reality some governments are going so far as to include AI education already in primary schools [7]. Universities in the Western World have been criticized that they are “asleep at the wheel, leaving the responsibility for this education to well-paid lobbyists and employees who’ve abandoned the academy” [21].

In response to the [21] article, [10] studied the integration of ethics into computer science university curricula. The authors did a qualitative analysis of 115 syllabi, from a total of 202 such identified AI Ethics courses in university technology curricula. They considered the courses content and learning outcomes. The courses considered in the study were identified through a crowdsourced collection of “tech ethics” syllabi [9]. The authors also looked into who teaches the courses.

Arguments have been put forward that defend the position that ethics philosophers should teach AI ethics course [15], but also to defend the position that this is a task for computer science lecturers [11, 10]. [10] found that the majority of the courses were thought at computer science and information science departments and were given by lecturers from these departments. Most courses were thought at the under-graduate level (107/115), 74 of the courses were thought at the graduate level, and 19 cross-listed for both. The courses were given both at the beginning and at the last years of studies.

The topics covered in the studied “tech ethics” courses included: Law & policy (66), Privacy & surveillance (61), Philosophy (61), Inequality, justice & human rights (59), AI & algorithms (55), Social & environmental impact (50) Civic responsibility & misinformation (32), AI & robots (27), Business & economics (27), Professional ethics (25), Work & labor (23), Design (20), Cybersecurity (19), Research ethics (16) and Medical/health (12). The goals and learning outcomes of the studied courses included: Critique, Spot issues, Make arguments, Improve communication, See multiple perspectives, Create solutions, Consider consequences and Apply rules. [10] pointed out

to a great variability in content across courses and disciplinary breadth.

An overview of the results of [10] and also the raw data suggests that none of the reported courses includes an all-encompassing overview of the AI-Ethics disciplines. Furthermore no courses teach machine ethics. The reported course topics are organised around specific societal and moral values rather than AI ethics disciplines. In 2017, INFO381 was one of the first courses in machine ethics in the world ².

3 INFO381 - 2017

Execution

When I constructed the INFO381 course in 2017, the field of AI Ethics was not as advanced as it is today. The topic of the course was limited on machine ethics as this sub-discipline had existed in computer science AI since at least 2006. Since 2017 the fields of FAccT, XAI and Responsible AI have been established with well visible international publishing venues, but also with considerable media and big tech company attention. I first describe the of INFO381-2017 course content, goals, learning outcomes and evaluation, before discussing the results and observed challenges.

INFO381-2017 was organised in eight, six hour long sessions. This is the standard organisation form for the Masters in Information Science courses at my department. Each session was structured as a sequence of lectures, discussions and group work. The list of sessions and material used is given in detail at <http://slavkovik.com/INFO381-2017.html>. The curriculum was left relatively open for adapting to the learning and reading pace of the students.

The goal of the course was to familiarise the students with the research topic of machine ethics specifically, and with some of the research methods and practices used in artificial intelligence in general. The learning outcomes for this course were as follows.

Knowledge. Identify and describe the basic principles of moral philosophy, interpret, explain and extend the need for, and challenges of, automating moral reasoning. Experience the entire process of research in machine ethics from the inception of an idea, analysis of research work, refining a research question, planing and executing group work and reporting on the work in the form of a scientific report.

Skills. Appraise ethical aspects of AI problems. Discern different moral theories and values when considering ethical impact of AI applications.

General competence. Reading and explaining scientific articles. Research project management. Scientific reporting.

The students were evaluated through an oral exam (40% of the grade) and a group project (60% if the grade). In addition, there were two additional obligatory assignments subject to approval. The evaluation methods were designed to be in service or learning rather than grading. The exam components were maximally personalised to allow the students freedom in following their interest while learning new skills, but also to relate their own performance with the effort put in rather than compare with the performance of other students. The goal of each component was to harness the incentive that a good grade brings for pushing oneself to independently learn.

²I am aware of the syllabus of an undergraduate level course at the University of Saarbrücken created by Kevin Baum and Holger Hermanns.

The oral exam was executed as a presentation. The students were tasked with finding and selecting an article in machine ethics, which they were then to read, presenting and answer questions about in a 25 minute presentation. The students were allowed to select any article in the scope of the discipline. The proposed articles were approved by the lecturer, to exclude position papers, short papers and papers entirely within philosophy, psychology and social science with only a casual link to artificial intelligence. The students were graded on the quality and clarity of the presentation, ability to answer questions and ability to analyse the scientific contribution in the context of the other work read throughout the course. The learning goal of this exam was to: 1) ensure that the students thoroughly read and understand a piece of scientific literature and become aware of the cognitive task and time involved; 2) experience public speaking and presentation of scientific topics and gain confidence into doing this type of activity.

The obligatory assignments were used to initialise the student group projects. In the first assignment the students were asked to identify an AI application that has an ethical impact (on a person, or society) and describe this application and perceived impact in a short one page deliverable. The goal of this assignment was to open up the views of the students into where AI applications are in our society and nudge them into questioning the ethical impact off these applications.

As a second assignment the students were asked to prepare (individually or in pairs) a research idea pitch for the rest of the class. The ideas were delivered in written form, a one page description of a research question, motivation, suggested methodology and success criteria. In defining the methodology, the students were asked to use at least one reference to a scientific article. They were also asked to describe the pitching pair skills and the required skills for completion of the project. The project ideas were pitched in the class and the class voted for the projects that they would like to execute. INFO381-2017 was attended by 25 students, 14 projects were pitched, of which 4 were selected for execution.

The pitching idea students who have won the popular vote were tasked with selecting a team and organising the execution of the project. The students worked independently developing their project ideas, however they were given supervision in terms of pointers to methodology and literature, scoping of research question and writing out the proposals. The goal of each of the project work was to experience an authentic practice of creating research.

Having the students design and choose their own project topics contributed to the students feeling that they are pursuing their interests rather than being forced into busywork. The students choose projects that helped them benefit from the skills they have gained so far in their studies but also they chose projects that used skills they wanted to acquire. Furthermore, because they designed and chose their own work they could choose how much work they will do, thus helping them understand hands-on how to plan work, what constitutes research progress and advancement of the state of the art. Most importantly, the students were aware that there are no ready answers or correct solutions to the problems they were engaged with, unlike project work where the task is defined by the lecturer.

The work of each project was described into a scientific report. The students used scientific articles as a template for the report. Rather than being assigned a specific typesetting format for the reports, they were asked to focus on clarity and completeness of the described work. All of the reports fell between 7000 and 10.000 words. These projects were eventually graded: A, B, B and C. The goal of the written report assignment was

to get the students to experience the challenge of scientific expression. Writing what was essentially a scientific article was also meant to help them understand better the articles they were reading.

After the grading was completed the students were given the opportunity to publish their project reports. This was optional and not related to the grade. The students were to select one person from their group to reformat the report into a publication and eventually present it to the scientific venue. All groups chose to do this. Out of the four projects, three were published as articles: [22], [6], [24]. The department made funding available for the presenting student to attend the respective venue and experience that side of research as well. Going through this final step helped the student see the grades as a learning opportunity rather than a judgement and disassociate grades from value. They also experienced the peer-review process.

Observations

One of the initial practical challenges to the designing INFO381-2017 was the lack of AI courses offered at the University of Bergen. Specifically machine learning was not being taught at the time, while being the fastest growing AI discipline and one most frequently considered to have ethical impact. This considerably limited the research options for students and curriculum. In 2018 I have developed Machine Learning as INFO284 at the Department of Information Science and Media Studies, while the Informatics Department introduced machine learning as INF283 in 2019 in addition to offering other machine learning courses.

Teaching a subject that is only a decade old means that there exists no textbook that can be readily used. Scientific articles are often not written for general audience - they assume the reader has a certain level of familiarity with the subject domain. The masters students were not very accustomed to reading such literature and were not comfortable with the domain of artificial intelligence or philosophy. All this significantly limited the available resources for the course. The selection of articles had to be sufficiently diverse to cover the broadness of topics in machine ethics, but the articles themselves had to be sufficiently established and well regarded in the AI community.

Perhaps the most challenging aspect of INFO381-2017 was in the subject matter of moral philosophy. The focus of the training of students in technical courses is their ability to “think in an algorithmic way” to see structured data in knowledge, to transform a problem into an algorithm, break it down into smaller programs and from there identify programming instructions. The masters students of information science are not an exception.

In mathematics and computer science there is a correct solution and an incorrect solution. We “train” students to check their solutions for correctness. Ethical problems and philosophical problems however rarely have an objectively correct solutions and this state of affairs also holds in aspects of machine ethics. Handling tasks for which there is no correct solution presented a bigger challenge for the students than anticipated. Such problems need to be addressed with arguments rather than proofs. Although I designed the courses around discussion topics intended to offer practice in argumentation, the offer did not meet the need. The students were also only offered one round of feedback to their project reports and this also turned out to be insufficient.

4 INFO381 in 2020 execution and observations

Execution

The observations outlined in Section 3 were considered when planning INFO381-2020. The course methodology was also assessed with respect to strategies from the learning sciences [5]. While key learning strategies such as learning by doing, authentic practices, project-based learning, personalized learning, collaborative learning and immersion in a community of practice were operationalised in the course, there was a clear room for implementing more learning by teaching, learning by reflection and learning by example.

Two critical changes impacted the planning of INFO381-2020: we were now teaching machine learning to our students and the AI Ethics field has exploded with work in XAI, FAccT and Responsible AI. The learning outcomes of the course and syllabus were adjusted in this regard to reflect the students' now existing knowledge in AI:

Knowledge. Identify the basic problems studied in XAI, FAccT, Responsible AI and machine ethics. Understand the premises of the core moral theories. Interpret, explain and extend the need for, and challenges of, AI Ethics. Experience the entire process of research in machine ethics from the inception of an idea, analysis of research work, refining a research question, planning and executing group work and reporting on the work in the form of a scientific report.

Skills. Appraise the ethical aspects of AI problems. Match a specific AI Ethics challenge to its most relevant discipline.

General competence. Reading and explaining scientific articles. Research project management. Scientific reporting.

INFO381-2020 was again executed in eight six, hour sessions, however, three hours in each of the last four sessions were devoted to working with the students on their projects and reports. Furthermore, very early on the students were explicitly instructed in argumentation theory and challenged to engage in discussions on assigned reading material. Unfortunately, the Covid-19 crisis and closure has interrupted regular classes and this plan for increasing the discussion in class was not entirely executed.

As learning and discussion material again scientific articles were used. The articles were mainly given as a reading assignment before each class. Additional home activities were given for students to try out. Each session also included a workshop on various forms of scientific reporting. The list of sessions topics, discussion topics and material used is given in detail at <http://slavkovik.com/INFO381-2020.html>.

The course was attended by 24 students and there were 5 student projects executed. The student evaluation format was retained with the same oral exam and project plan. In addition more intermediary assignments for the student projects were scheduled. These included: a refinement of research question, motivation, success criteria and methodology after the projects were decided, a related work analysis, and a two drafts deliveries. It is my judgement that the additional assignments of intermediary project reports has helped students with understanding scientific reporting, but it has also helped their argumentation skills. The grades for the projects were: A, A, B, C, C.

Observations

The order in which topics were thought could be improved in an AI Ethics course organisation. In INFO381-2020, machine ethics was introduced first and in hindsight

this was not the right way. Machine ethics is the least intuitive of the AI Ethics sub-disciplines. Machine ethics requires understanding of moral philosophy, deontic logic and inductive logic programming. For most students this was a first introduction with moral philosophy with many new concepts that needed to be comprehended.

The machine ethics was followed with FAcct and XAI. Fairness and explainability require understanding of machine learning, but also accountability and transparency. Responsible use of AI was covered twice - in the first class and in the last. As a topic, this one is most intuitive to the students, however, although it is not easy to see it does require a background of moral philosophy and also law and it is the most “further away” from computer and information science.

It is my conclusion that the modules should be thought effectively in reverse. First the students should be exposed to responsible AI, then the curriculum should follow with FAcct and XAI. This will expose the students to concepts from moral philosophy but in a programming environment that they are more familiar with. Once they have grasped values such as fairness, freedom, autonomy, and privacy they can be exposed to moral philosophy. The course will thus finish with machine ethics. Lastly, the participation in discussions should be made a formal obligatory requirement to reinforce the value of argument building skills development.

5 AI Ethics as a bachelor course

The experience of INFO381 and the analysis of [10] will be taken into consideration in designing a bachelor course in AI Ethics. The [10] analysis identifies the challenge of integrating ethics into “purely technical programming courses”, with many of the courses leaning towards high level and conceptual topics of impact of technology on society. I have also observed this as a challenge both in INFO381-2017 and INFO381-2020. Two of the tools that INFO381 relies on to immerse students into AI ethics cannot be transferred to a bachelor level course: self-selected group projects and studying from scientific articles.

Fiesler [10] concluded that integrating ethical topics with programming “purely technical” courses “might even be a way to combat an ‘I’m just an engineer’ mindset that ethics is ‘someone else’s job’.” The very open-ended group projects are the main tool that INFO381 uses to change this mind set and help future engineers and scientists see ethical AI impact as their own job. Bachelor students do not have the skill-set to execute and benefit from such projects.

One possible solution is to design a course that relies more heavily on topics of fairness and explainability with topics of moral philosophy interleaved. There is an added benefit to a FAcct and XAI “heavy” course. These areas concern machine learning especially. As the students will be following this course in parallel with a machine learning course, they can learn to consider ethical impact as part of machine learning, not as an afterthought. Practical programming exercises in assessing fairness, mitigating bias and generating explanations can be developed in parallel with the machine learning course.

The following learning outcomes can be considered

Knowledge. Identify and describe the concepts of bias, trust, fairness, explanation, equality and empowerment in the context of artificial intelligence. Discern the sub-fields of AI Ethics, their challenges and relations to social science and humanities. Familiarity with national and international AI guidelines for responsible use

Skills. Evaluate the algorithmic accountability of an AI system. Assess the possible societal and personal impact of an AI system. Apply tools for providing explanations of AI algorithms. Apply tools for assessing bias in data sets and AI algorithms.

General competence. Comprehend and communicate the ethical and societal implication of AI to a general audience

While students should always be encouraged to read scientific articles, it would be inadequate to rely on them as the only study material for a bachelor course. We would need to generate a script of selected texts using specifically systematic review articles and toolkit tutorials. A possible curriculum outline for bachelor AI Ethics, presuming 15 two hour units, is given in Table 1.

1.	Why AI ethics, why now
2.	Responsible AI and Norwegian AI guidelines
3.	What is an explanation?
4.	Accountability and Transparency
5.	Transparent vs black box machine learning
6.	Generating local explanations
7.	Argumentation and decisions
8.	What are good arguments and how to build them?
9.	Post-hoc explanations and surrogates
10.	How do we define bias and fairness?
11.	Data bias and mitigation
12.	Algorithmic fairness and mitigation
13.	Basics of moral philosophy
14.	Basics of moral philosophy
15.	An introduction to machine ethics

Table 1: Possible curriculum outline for a bachelor level AI Ethics course.

Lastly, it remains very important for the students to be thought how to analyse a problem that does not have an objectively correct solution. This again was done practically through the open-ended projects in INFO381. One possibility here is to execute the same self-selection of a project, but the students would be tasked to write a popular science article communicating an issue in AI Ethics. They would be mentored in choosing a topic, discovering relevant material, and peer-evaluating the strength of the arguments in their own essays.

6 Conclusion

While there are global efforts to include courses on AI ethics in both bachelor and masters level, there is no clear consensus which topics such a course should focus on. An AI ethic course is difficult to fit in the beginning of a degree, because the students need to have a certain level of understanding of AI and computing methods to be able to comprehend their ethical impact. Students in computer and information sciences are rarely trained to handle a qualitative assessment of problems, in contrast to students in humanities and social science disciplines. Up until recently, computation and society were two separate environments and there was no need to interleave them.

INFO381-2017 and -2020 has had the goal to give an overview of the core issues and accomplishments in AI Ethics in a fashion that would motivate the students to pursue further learning in this area. Instead of the lecturer choosing one topic for in-depth study, the students were free to pursue their own interests and focus. It is critical to be aware that in young disciplines such as this we are educating the next generation of AI Ethics researchers and practitioners, and we need to motivate our students for this task, but also equip them with the basic skills they will need in their future careers.

The highly student-tailored approach of teaching AI Ethics in INFO381 has helped with most of the teaching AI Ethics challenges. However, the strategies used in INFO381 cannot be scaled to a bachelor level students who are not as equipped for independent learning as masters students. Teaching bachelor level AI ethics requires: 1) the preparation of a text book level material, 2) focus on the most “technical” and hands-on topics from AI Ethics, and 3) the gentle, strategic and practical introduction of argumentation and scientific communication skills.

References

- [1] Algorithmic Watch: Ai ethics guidelines global inventory. <https://inventory.algorithmwatch.org/> (2020), [Online; accessed 20-May-2020]
- [2] Arya, V., Bellamy, R.K.E., Chen, P.Y., Dhurandhar, A., Hind, M., Hoffman, S.C., Houde, S., Liao, Q.V., Luss, R., Mojsilović, A., Mourad, S., Pedemonte, P., Raghavendra, R., Richards, J., Sattigeri, P., Shanmugam, K., Singh, M., Varshney, K.R., Wei, D., Zhang, Y.: One explanation does not fit all: A toolkit and taxonomy of ai explainability techniques. <https://arxiv.org/abs/1909.03012> (2019)
- [3] Bellman, R.E.: *An Introduction to Artificial Intelligence: Can Computers Think?* Boyd & Fraser Publishing Company (1978)
- [4] Bovens, M.: Analysing and assessing accountability: A conceptual framework1. *European Law Journal* **13**(4), 447–468 (2007). <https://doi.org/10.1111/j.1468-0386.2007.00378.x>, <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0386.2007.00378.x>
- [5] Bransford, J.D., Brown, A.L., Cocking, R.R. (eds.): *How people learn: Mind, brain, experience, and school.* National Academy Press, Washington, DC (2000)
- [6] Caycedo Alvarez, M., Berge, Ø.S., Berget, A.S., Bjørknes, E.S., Johnsen, D., Madsen, F.O., Slavkovik, M.: Implementing asimov’s first law of robotics. *Norsk Informatikkonferanse* (2017), <https://ojs.bibsys.no/index.php/NIK/article/view/396>
- [7] Chan, D.: Primary students to be taught ai in guangzhou schools government has decided to give priority to the development of ai, information and biopharmaceutical industries. <https://asiatimes.com/2019/07/primary-students-to-be-taught-ai-in-guangzhou-schools/> (July 5 2019), [Online; accessed 10-June-2020]
- [8] Dignum, V.: *Responsible Artificial Intelligence - How to Develop and Use AI in a Responsible Way.* *Artificial Intelligence: Foundations, Theory, and Algorithms*, Springer (2019). <https://doi.org/10.1007/978-3-030-30371-6>, <https://doi.org/10.1007/978-3-030-30371-6>
- [9] Fiesler, C.: *Tech ethics curricula: A collection of syllabi* (July 5 2018)
- [10] Fiesler, C., Garrett, N., Beard, N.: What do we teach when we teach tech ethics? a syllabi analysis. In: *Proceedings of the 51st ACM Technical Symposium on Computer Science Education*. p. 289–295. SIGCSE ’20, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3328778.3366825>, <https://doi.org/10.1145/3328778.3366825>
- [11] Grosz, B.J., Grant, D.G., Vredenburg, K., Behrends, J., Hu, L., Simmons, A., Waldo, J.: Embedded ethics: integrating ethics across CS education. *Commun. ACM* **62**(8), 54–61 (2019). <https://doi.org/10.1145/3330794>, <https://doi.org/10.1145/3330794>

- [12] Gunning, D., Aha, D.: Darpa's explainable artificial intelligence (xai) program. *AI Magazine* **40**(2), 44–58 (Jun 2019). <https://doi.org/10.1609/aimag.v40i2.2850>, <https://www.aaai.org/ojs/index.php/aimagazine/article/view/2850>
- [13] Hagendorff, T.: The ethics of AI ethics - an evaluation of guidelines. *CoRR* **abs/1903.03425** (2019), <http://arxiv.org/abs/1903.03425>
- [14] Jobin, A., Ienca, M., Vayena, E.: The global landscape of AI ethics guidelines. *Nature Machine Intelligence* (2019). <https://doi.org/10.1038/s42256-019-0088-2>, <https://doi.org/10.1038/s42256-019-0088-2>
- [15] Johnson, D.: Who should teach computer ethics and computers & society? *SIGCAS Comput. Soc.* **24**(2), 6–13 (Jun 1994). <https://doi.org/10.1145/181900.181901>, <https://doi.org/10.1145/181900.181901>
- [16] Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., Galstyan, A.: A survey on bias and fairness in machine learning. <https://arxiv.org/abs/1908.09635> (2019)
- [17] Miller, T.: “but why?” understanding explainable artificial intelligence. *XRDS* **25**(3), 20–25 (Apr 2019). <https://doi.org/10.1145/3313107>, <https://doi.org/10.1145/3313107>
- [18] Ministry of Local Government and Modernisation: The national strategy for artificial intelligence. <https://www.regjeringen.no/en/dokumenter/nasjonal-strategi-for-kunstig-intelligens/id2685594/?ch=1> (2020), [Online; accessed 20-May-2020]
- [19] Moor, J.: The dartmouth college artificial intelligence conference: The next fifty years. *AI Magazine* **27**(4), 87 (Dec 2006). <https://doi.org/10.1609/aimag.v27i4.1911>, <https://www.aaai.org/ojs/index.php/aimagazine/article/view/1911>
- [20] Moor, J.H.: The nature, importance, and difficulty of machine ethics. *IEEE Intelligent Systems* **21**(4), 18–21 (Jul 2006). <https://doi.org/10.1109/MIS.2006.80>, <https://doi.org/10.1109/MIS.2006.80>
- [21] O’Neil, C.: The ivory tower can’t keep ignoring tech. <https://www.nytimes.com/2017/11/14/opinion/academia-tech-algorithms.html> (2017), [Online; accessed 20-May-2020]
- [22] Pires Bjørgen, E., Øvervatn Madsen, S., Skaar Bjørknes, T., Vonheim Heimsæter, F., Håvik, R., Linderud, M., Longberg, P., Dennis, L., Slavkovik, M.: Cake, death, and trolleys: Dilemmas as benchmarks of ethical decision-making. In: *Proceedings of the 2018 Conference on AI, Ethics, and Society, AIES 2018, New Orleans, LA, USA, February 02-03, 2018*. pp. 23–29. *AAAI/ACM* (2018). <https://doi.org/10.1145/3278721.3278767>, <https://doi.org/10.1145/3278721.3278767>
- [23] Russell, S., Norvig, P.: *Artificial Intelligence: A Modern Approach*. Pearson Education, 3 edn. (2015)
- [24] Valland, D.S., Carlsen, T., Lea, M., Pensgaard, A., Karlsen, S., Sötov, P., Slavkovik, M.: The norwegian oil fund investment decider n.o.f.i.d. *Norsk konferanse for organisasjoners bruk at IT* **25**(1) (2017), <https://ojs.bibsys.no/index.php/Nokobit/article/view/409>
- [25] Wallach, W., Allen, C.: *Moral Machines: Teaching Robots Right from Wrong*. Oxford University Press, Inc., USA (2008)
- [26] Weller, A.: Challenges for transparency. *CoRR* **abs/1708.01870** (2017), <http://arxiv.org/abs/1708.01870>
- [27] Wieringa, M.: What to account for when accounting for algorithms: A systematic literature review on algorithmic accountability. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. p. 1–18. *FAT* '20, Association for Computing Machinery, New York, NY, USA* (2020). <https://doi.org/10.1145/3351095.3372833>, <https://doi.org/10.1145/3351095.3372833>