The Design and Evaluation of a Hands-on AI Literacy Workshop for Politicians

Victoria Hopsø, Dag Svanæs

Department of Computer Science

Norwegian University of Science and Technology

7491 Trondheim, Norway

victoho@stud.ntnu.no, dag.svanes@ntnu.no

ABSTRACT

Within liberal democracies, governance implies inclusiveness, accountability, and transparency. This requires well-informed lawmakers. The gap between the fast-evolving socio-technological realities of AI and lawmakers' competence levels poses a threat to governance. This paper presents the design and evaluation of a three-hour hands-on AI literacy workshop targeted at lawmakers, tested with a pilot group of young politicians. In-depth interviews with five young people and five AI experts identified key misconceptions and essential AI concepts, resulting in a set of intended learning outcomes. The workshop featured interactive exercises illustrating core AI principles, bias, and ethical concerns. Preliminary evaluation with three young politicians demonstrated meaningful learning gains, including a deeper understanding of machine learning, model training, and algorithmic bias. While sample sizes were small, the results confirm the promise of active, collaborative AI literacy approaches for enabling more informed AI governance. Future workshops are intended to target Members of Parliament (MPs) in Norway.

CCS Concepts

• Computing methodologies → Artificial intelligence; • Human-centered computing → Collaborative and social computing theory, concepts and paradigms; Human computer interaction (HCI); • Social and professional topics → Computing / technology policy.

Keywords

AI Governance, AI Literacy, Lawmakers, Educational intervention

1. INTRODUCTION

Governance means different things in different societal contexts. Within liberal democracies, governance implies inclusiveness, accountability, and transparency to society at large. In contrast, more authoritarian societies may prioritize control, stability, or ideological uniformity over open, participatory processes. However, governance in more authoritarian societies is outside the scope of the current paper.

Governance in liberal democracies requires well-functioning institutions for legislation and oversight. One challenge concerning the governance of emerging technologies is that it requires well-informed lawmakers, i.e., politicians and their staff.

Concerning Artificial Intelligence (AI), the gap between fast-evolving socio-technological realities and lawmakers' competence levels poses a threat to governance. Ill-informed lawmakers may easily miss important aspects of the technology that should have been governed, or become overly concerned with aspects that require less governance. Without a necessary level of AI literacy [7], lawmakers can easily be persuaded by the arguments of resourceful stakeholders who oppose AI governance. Missing or misdirected AI governance [10] holds the potential for harm to individuals and society. Examples of harm include breaches of privacy, surveillance, misinformation, and manipulation of elections. Raising AI literacy among lawmakers should consequently be prioritized, but there is currently little research on how best to achieve this.

1.1 Bridging the lawmaker AI literacy gap

To address the AI literacy gap among lawmakers, we developed and empirically evaluated an educational intervention consisting of a three-hour, hands-on workshop in Norway. The Intended Learning Outcomes (ILOs) for the workshop were specified based on interviews with AI experts, some with experience in communicating AI issues to lawmakers.

Because Members of Parliament (MPs) are difficult to recruit, the workshop is targeted for and evaluated with young politicians not (yet) in Parliament. It is intended as a pilot study to inform the design of a hands-on AI literacy workshop for MPs. The workshop was well received by the young politicians, and important lessons were learned as input for the next iteration of the workshop.

The research project was designed to answer the following four research questions:

- What do young people without an educational background in computer technology know about AI, including societal implications and common misconceptions?
- What do experts deem the most prominent societal implications of AI, and which core concepts of AI need to be understood in order to make informed decisions concerning AI governance?
- Which intended learning outcomes should guide the design of a curriculum about AI and its societal implications aimed specifically at youth-politicians?
- To what extent can active and collaborative learning about the core concepts of AI foster discussion and reflection about the societal implications of AI in a workshop setting, and how can this approach facilitate achievement of the intended learning outcomes of an AI curriculum?

The above four research questions can be seen as one iteration of a user-centered design process. We will here present the background for the project, its research design, preliminary results, conclusions, and future work. The presented work is described in more detail in [4].

2. BACKGROUND

Between 2000 and 2020, research efforts within the field of teaching and learning AI (AITL) were primarily concerned with university level computer science education [8]. However, AI's recent irruption in our daily lives has resulted in an increased need to educate citizens on responsible and informed use of AI [2, 9]. This has led to a substantial increase in research efforts within AITL, and particularly the field of AI literacy [2].

For the current context, AI literacy refers to the competencies and skills that make up a holistic comprehension that fosters responsible use of AI and an awareness of its societal impact [7, 9]. Much of the recent literature on AITL have thus explored the educational settings of K-12 (kindergarten through grade twelve) and university [1, 3, 5, 8, 11, 13].

Educating citizen on responsible and informed use of AI is a necessary first step towards mitigating the potential risks of AI, but most liberal democracies realize that in addition to measures at the individual level, some amount of AI governance is required at the societal level. This asks for raising the AI literacy of lawmakers. Although we have found no literature on educational interventions for lawmakers, the need for such initiatives has been pointed out by Kotsis [6]. They conclude that "In order to make well-informed judgments, policymakers need to engage in ongoing education and training about pertinent scientific ideas and breakthroughs in artificial intelligence (AI) technology. This may include various educational activities such as workshops, seminars, and engagement with domain experts." (ibid. p. 76)

3. RESEARCH DESIGN

To address the four research questions, the following research activities were planned and executed:

1. In-depth interviews with young persons to identify current level of AI literacy and common misconceptions about AI.
2. In-depth interviews with AI experts to identify prominent societal implications of AI, and which core concepts of AI need to be understood in order to make informed decisions concerning AI governance.
3. Thematic analysis of the interviews, synthesized into a set of ILOs for lawmaker AI literacy.
4. The design of a hands-on three-hour workshop to address the ILOs.
5. The evaluation of the resulting workshop with young politicians, including post-workshop interviews.

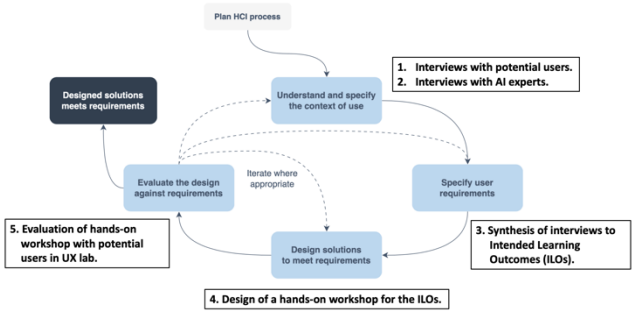

Figure 1. The five steps of the research process, with reference to ISO 9241-210.

The research design follows a user-centered design approach, as specified in ISO 9241-210 [12]. Figure 1 shows the five steps of the research plan (in black boxes), with reference to the steps specified in the ISO standard (in blue).

4. INTERVIEWS

As this was an initial pilot study, the number of interviewees was low: Five young persons (age 18-27, four female, one male) and five AI experts (four female, one male).

All the young persons interviewed understood that AI models are trained, but only one of them could give any details about how this was done. One common misconception was that AI models are general, i.e. good at any domain. Most interviewees thought for example that an autonomous driving AI would also be a good language model. Although they recognized that AI models are not human, they consistently anthropomorphized AI in their answers, e.g. "When I talk to ChatGPT, it is very kind". The latter implies not recognizing that the LLM has been intentionally designed to appear "kind". Concerning societal implications of AI, many of the young persons were afraid that AI would "take over job", or "take over the world", much inspired by sci-fi movies.

The AI experts identified a wide number of issues. This was synthesized into 22 ILOs, including:

- Understanding AI as a set of specific technologies and applications, and that Machine Learning (ML) is just one kind of AI.
- Being familiar with the principles and concepts behind ML and Large Language Models (LLMs)
- Viewing data as representation and grasp its role in accuracy and fairness.
- Knowing of neural networks and comprehend the challenge of explainability.
- Being able to identify and reflect on ethical concerns of AI (e.g. bias, privacy, democracy).

5. A HANDS-ON WORSKHOP

We designed a three-hour hands-on AI literacy workshop for young politicians to satisfy the ILOs resulting from the interviews. It consists of a combination of theory and hands-on exercises.

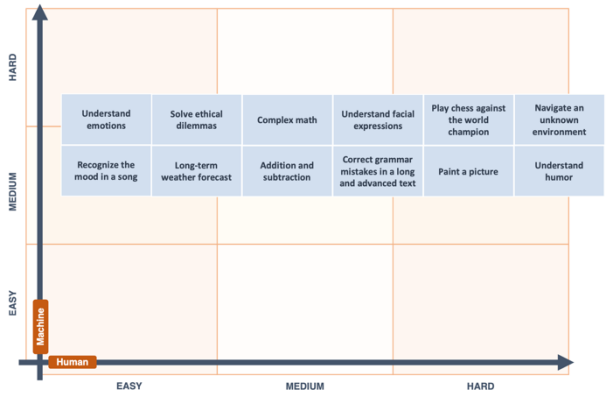

Figure 2. Hands-on exercise 1: Place the task cards on a 3x3 matrix easy-hard for humans and AI.

The hands-on exercises were developed to address the intended learning outcomes that resulted from interviews with the AI experts and the young persons. As illustrated in Figure 1, the research process was user-centered.

The first hands-on exercise (Figure 2) is to classify tasks as easy, medium or hard for people and AI. The participants are given twelve cards with tasks (e.g. "Understanding emotions", "Playing chess") and are asked to place them on a 3x3 matrix of easy-hard for humans and AI.

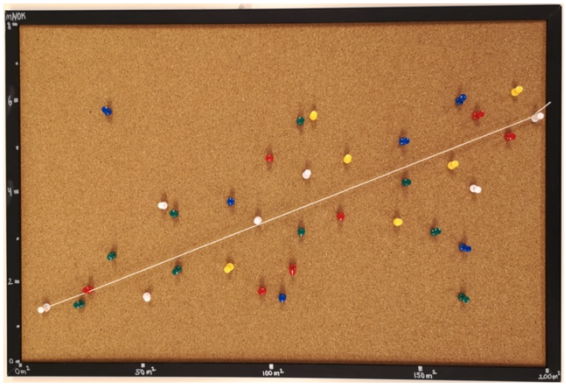

Figure 3. Hands-on exercise 2: Find the best mapping between house size and house price (simple linear regression).

The second hands-on exercise is designed to teach the difference between training data and models, using simple linear regression as an example. It consists of a bulletin board filled with a number of pushpins (Figure 3), where the horizontal axis denotes the size of a house in square meters, and the vertical axis denotes house prices. The pushpins illustrate historical data on previously sold houses, while the twine, when spanned as a straight line, represents a linear model for house prices per size. The learners are introduced to what the pushpins represent, and told that we wish to use these data to train an ML model to predict the price of unsold houses. They are informed that we assume there is a linear relationship between the size and price, and are instructed to span a straight line with the twine so that it best fits the data. The learners' task is to find out what a 150 square meter house will cost.

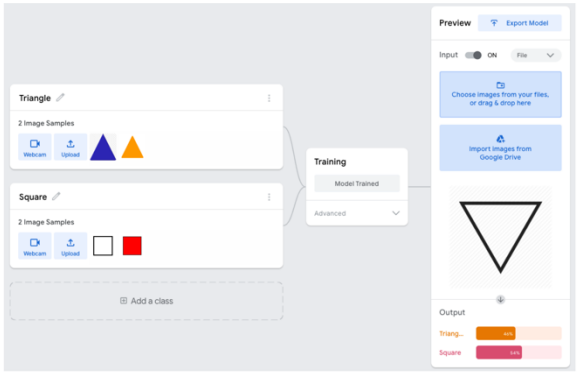

Figure 4. Hands-on exercise 3:Train Google Teachable Machine to differentiate between squares, circles and triangles.

In the third hands-on exercise, the participants use Google's Teachable Machine[1] to train a model to differentiate between squares, circles and triangles, using a set of given images (Figure 4). It illustrates that the quality of the model depends on the quality of the data.

In the fourth hands-on exercise, the participants explore the use of ML for automating a hiring process. The aim is to spur reflections on bias.

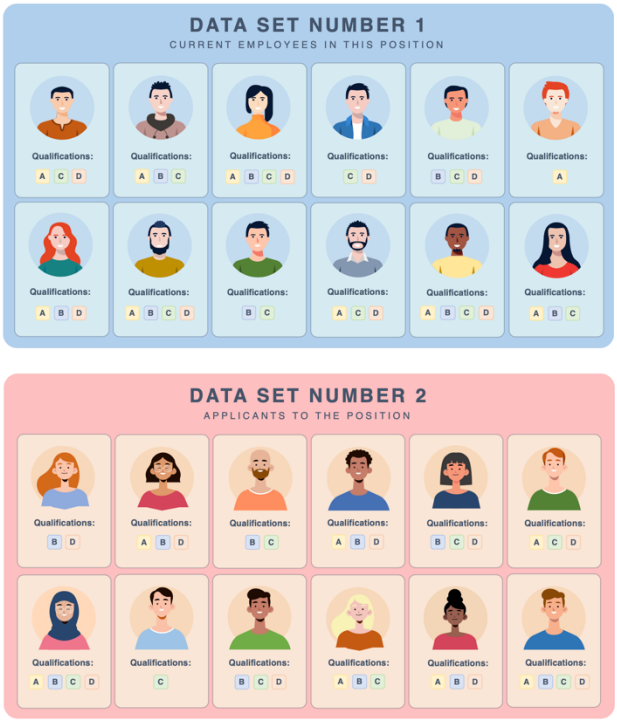

Figure 5. Hands-on exercise 4: Simulating ML for automating a hiring process. Training data in blue.

The participants are given a classification of the current employees in a company, including gender and ethnicity (Figure 5, in blue). Applicants (Figure 5, in red) are then rated based on these data, resulting in biases that are discussed.

Figure 6. Hands-on exercise 5: The use of AI for illegal or unethical purposes. Role cards.

In the fifth and last hands-on exercise, the participants explore a hypothetical situation where a database of Electronic Health Records was leaked, and how various stakeholders could use AI on the data for illegal or unethical purposes (Figure 6).

6. EVALUATION AND LESSONS LEARNED

The workshop was evaluated in our UX lab (Figure 7), with three young politicians (age 15-17, all male). The exercises spurred relevant discussions among the participants on the intended topics, and it enabled the participants to achieve the intended ILOs.

[1] https://teachablemachine.withgoogle.com/

Figure 7. Evaluation of the workshop in our UX lab.

All participants found the workshop useful. Some quotes from the post-workshop interviews illustrate this:

- "I have never actually trained an AI like we did just now, and it was a bit like; "Oh, so you can do it yourself!" (...). And you could see how what you fed it influenced its assessment of the triangle, the square, or the circle."
- "For me, it was about biases. I believe many people on, let us call it our political side, often dismiss such biases and think that they are not necessarily a real thing. They do not always take it [bias] seriously (...). But here [in the workshop], I think it was very important that people actually gained an understanding of why there can be biases."
- "It was a nice balance; theory is one thing, but another is to try it out in practice with case-activities and such. It puts things into perspective in a way. It is useful to get the theory and to have you [the instructor] discuss it, and then that we get to try it out after each point to kind of put things into perspective and to see how it works in practice."

Concerning potential for improvement, we could have selected examples that were closer to the life of the participants to make it more motivating.

7. CONCLUSIONS AND FUTURE WORK

To bridge the AI literacy gap among lawmakers, we have designed and evaluated a hands-on workshop for politicians and their staff, first aimed at and piloted with young politicians in Norway. We recognize that the study has multiple limitations relating to sample size, that it was done in a liberal democracy, researcher bias, and workshop participants gender (only male). Taken into consideration these limitations, the overall result from the study indicate that hands-on AI literacy workshops, combining basic theory with exercises and group discussion, hold the potential to make lawmakers better equipped to make informed decisions concerning AI governance.

Further evaluations and expert interviews are required to identify what aspects of the AI literacy gap that are most pressing concerning lawmakers.

The positive feedback from the workshop participants motivate us to improve it and use it as basis for designing a similar AI literacy workshop for MPs and their staff .

REFERENCES

[1] Karl-Emil Kjær Bilstrup, Magnus H Kaspersen, and Marianne Graves Petersen. *Staging reflections on ethical dilemmas in machine learning: A card-based design workshop for high school students*. in *Proceedings of the 2020 ACM designing interactive systems conference*. 2020.

[2] Lorena Casal-Otero, et al., *AI literacy in K-12: a systematic literature review.* International Journal of STEM Education, 2023. **10**(1): p. 29.

[3] Stefania Druga, et al., *Inclusive AI literacy for kids around the world*, in *Proceedings of FabLearn 2019*. 2019. p. 104-111.

[4] Victoria Hopsø, *Educating Youth Politicians about Artificial Intelligence and its Societal Implications - A Hands-on Approach*, in *Department of Computer Science*. 2024, NTNU.

[5] Magnus Høholt Kaspersen, et al., *High school students exploring machine learning and its societal implications: Opportunities and challenges.* International Journal of Child-Computer Interaction, 2022. **34**: p. 100539.

[6] Konstantinos T Kotsis, *The scientific literacy enables policymakers to legislate on Artificial Intelligence.* European Journal of Political Science Studies, 2024. **7**(1).

[7] Duri Long and Brian Magerko. *What is AI literacy? Competencies and design considerations*. in *Proceedings of the 2020 CHI conference on human factors in computing systems*. 2020.

[8] Davy Tsz Kit Ng, et al., *A review of AI teaching and learning from 2000 to 2020.* Education and Information Technologies, 2023. **28**(7): p. 8445-8501.

[9] Davy Tsz Kit Ng, et al., *AI literacy: Definition, teaching, evaluation and ethical issues.* Proceedings of the association for information science and technology, 2021. **58**(1): p. 504-509.

[10] Anka Reuel, et al., *Open problems in technical ai governance.* arXiv preprint arXiv:2407.14981, 2024.

[11] Marie-Monique Schaper, et al., *Computational empowerment in practice: Scaffolding teenagers' learning about emerging technologies and their ethical and societal impact.* International Journal of Child-Computer Interaction, 2022. **34**: p. 100537.

[12] International Organization for Standardization, *ISO 9241-210: Ergonomics of human–system interaction-Human-centred design for interactive systems.* International Organization for Standardization, 2019.

[13] David Touretzky, et al. *Envisioning AI for K-12: What should every child know about AI?* in *Proceedings of the AAAI conference on artificial intelligence*. 2019.