

SE4AI: A Training Program Considering Technical, Social, and Professional Aspects of AI-based Software Systems

Ahmad Abdellatif*, Gita Ghiasi**, Diego Elias Costa*, Emad Shihab*, Tanja Tajmel‡

*Department of Computer Science and Software Engineering

‡Centre for Engineering in Society, Concordia University
Concordia University, Montreal, Canada

{ahmad.abdellatif, diego.costa, emad.shihab, tanja.tajmel}@concordia.ca

**School of Engineering Design and Teaching Innovation
University of Ottawa, Ottawa, Canada
gghiasi@uottawa.ca

Abstract—Artificial Intelligence based systems offer a multitude of opportunities and significant advantages to the industry and the larger society, but at the same time, they may create or contribute to social and ethical problems in society. Developing AI-based systems is a challenge when considering the need for specialized knowledge of the social implications surrounding the development of those systems. We believe that there is a need for a comprehensive training that addresses both, the technical and social issues, surrounding these systems. In this paper, we aim to address this gap by sharing our experience developing a training program that provides trainees with technical, social, and professional knowledge in the context of AI-based software systems. This is further substantiated by survey results that indicate that the program has helped trainees identify and address issues related to AI-based systems while increasing their awareness of social and ethical issues.

I. INTRODUCTION

The field of Artificial Intelligence (AI) has witnessed a marked increase in both demand and growth, fueled by the large availability of datasets and Machine Learning (ML) breakthroughs, such as deep learning and reinforcement learning. Many organizations leverage AI to transform entire sectors of our society with data-driven solutions, from education [1] and finance [2] to healthcare [3] and security [4]. However, the development and maintenance of AI-based systems require specialized training, since such systems have fundamental differences compared to traditional software systems. Contrary to traditional software systems that are written to implement well-defined requirements, AI-based systems infer their behaviour from data, making systems unstable, prone to discrimination, and harder to troubleshoot and debug [5].

While AI plays a role in driving economic opportunities, AI applications bring about various social concerns [6], such as ethical issues, privacy rights, and surveillance. The lack of social accountability in AI systems can lead to a wider digital divide and the violation of social norms, ultimately having a negative impact on human rights such as freedom and equality [7]. For instance, Amazon employed an AI-based

recruitment system to filter the top candidates for job applications. The system was later reported to be heavily biased against female applicants due to biases in Amazon’s data and was discontinued [8]. Another example is IDEMIA’s facial recognition algorithm, which has been shown to exhibit bias against black women when compared to other groups (e.g., white women) [9]. Software engineers are heavily involved in developing and maintaining AI-based systems, but are not typically trained on critical social issues related to these systems. Current research and training in Software Engineering (SE) rarely consider the combination of technical and social aspects when developing AI-based software systems. *Therefore, it is important to have a training program to train practitioners working with AI-based systems to know, understand, apply, and integrate both technical and social aspects.*

To bridge this gap, we share with the community our experience and efforts in devising a training program on the development, deployment and servicing of AI-based software systems. The main goal of the program is to equip future software engineers with both the technical skills to build AI-based systems and the knowledge to understand and incorporate the ethical and social implications of AI. In particular, trainees are introduced to the fundamental technical concepts in designing, developing, and servicing AI-based software systems, through an engineering course: “Engineering AI-based systems”. Moreover, the training program equips trainees with ethical and social criteria that AI-based systems need to consider, guided by human rights and sustainable development goals, through the course: “Social Aspects for AI-based Software Systems”. Finally, to prepare trainees for their future career paths, our training program includes several specialized modules on professional skills in the context of AI-Software Systems (e.g., Commercialization and Entrepreneurship).

Started in 2021, the program is expected to train approximately 70 software engineering trainees in six years, spanning students from four Canadian universities. Specifically, the program offers training opportunities to undergraduate and (mostly) graduate students pursuing Ph.D. and Master’s

degrees, and is an integral component of the trainees' master's and Ph.D. programs. The Ph.D. training lasts for four years, and the Master's training for two years. A few undergraduate students take short-term training through short research internships in the research labs involved with the training program. The trainee selection is based on the recommendations of supervisors, overseen by the training program admission committee. We consider gender balance, technical capabilities, and other criteria related to equity, diversity, and inclusion in the admission process. To collect the initial feedback about the program, we surveyed the first-year trainees (18 trainees) enrolled in the training program. We received nine responses (50% response rate). The results show that all survey participants either strongly recommend (77.8%) or recommend (22.2%) the program. The participants specified that the program enables them to identify and address issues related to AI-based systems. Furthermore, the majority (89%) of participants indicated that the program increases awareness of social and ethical issues of AI, which highlights the importance of considering the combination of technical and social aspects in our training program. The participants suggest having more collaborations with the industry through internships and among trainees themselves, which we are planning to implement in our program for the future cohort.

II. TRAINING PROGRAM

The training program is an integral component of the trainees' master's and Ph.D. programs. It aims to train the trainees on the technical and social aspects of AI-based software systems through courses on Engineering AI-based Software Systems and Social Aspects of AI-based Systems. The trainees are required to complete several components to apply what they have learned in practice. The program offers several webinars featuring industry practitioners who share current practices used in the development of AI-based systems. Moreover, the program cultivates the professional skills of the trainees through specialized modules. Our training philosophy is centred around several key aspects: increasing the trainees' academic knowledge on the technical and social aspects, arming trainees with real-life experiences, fostering critical thinking, experiencing hands-on learning, bridging the gap between academia and industry, and cultivating the trainees' professional skills. In the following, we elaborate on how the training program achieves the aforementioned aspects:

Provide trainees with a strong academic foundation on the technical and social aspects. To ensure that trainees have a strong academic foundation, we design our training program to expose trainees to the latest research (both technical and social aspects) about various topics related to the development of AI-based systems. For example, in the Engineering AI-based Software Systems course, we cover the entire development lifecycle of AI-based software systems, from its specification and requirements engineering to the deployment and maintenance (MLOps) of those systems in production, as shown in Table I. The course discusses what a 'typical' software engineering process is and then

delves to discuss the unique aspects one needs to consider for AI-based systems (e.g., the role of data in requirements, special testing needs to consider drift, etc.). The Social Aspects of AI-based Systems course introduces students to the terminology related to the social aspects of technology and enable them to articulate the social issues of AI. Specifically, it covers topics varying from terminologies around equity, diversity and inclusion (EDI) to Indigenous perspectives on AI, as shown in Table I. In the social aspects course, trainees learn to apply a multi-level approach to identify and analyze social inequities related to AI. This approach includes several forms of structural, institutional and personal factors, e.g., (i) individual bias which influences decisions that result in inequitable technology; (ii) disciplinary values embedded in the culture of engineering and (iii) structural and societal factors that influence engineering practices. The course discusses the impact of AI on individuals from various backgrounds and demographics. We showcase real-world examples where AI has had an impact on different communities and engage in discussions that explore the social aspects of AI and its impact on diverse populations. For example, in the Social Aspects of AI-based Systems course, we discussed an article that addresses how racism is inscribed in automatic facial recognition and provides suggestions for the inclusion of people's diversity in such systems [10].

Expose trainees to challenges and best practices involving the development of AI-based software systems via real-life experiences. Given the novelty of the topics, we find it hard to cover all topics shown in Table I with the traditional textbook material approach. While we rely on some chapters of the book of George Hulten [11] and Kate Crawford [12], most topics were prepared using peer-reviewed scientific articles, and industry reports. We put special emphasis not only on articles that present an industrial perspective on the engineering challenges, but also solutions related to the topic of the session. Thus, we better prepare them for their real-life software engineering work by exposing the trainees to the challenges and potential solutions employed on real-life AI-based systems. For example, when presenting the topic of software architecture, we discuss the data preparation challenges that might occur while designing the software architecture of an AI-based system (e.g., jungle of scrapes), and present the different solutions to these challenges (e.g., modularize the data preparation pipeline). In some cases, we opt to use international reports, documentaries, films, and interviews to highlight the impact of AI on social aspects. For example, to increase the trainees' awareness of AI discrimination consequences, we assign the trainees to watch two documentaries (Integrating gender and EDI on different social aspects¹ and Coded Bias²).

Foster critical thinking. We craft the lectures in our training program to improve the trainees' critical thinking and analytical skills, through class assignments and discussions.

¹http://genderedinnovations.stanford.edu/video_landing.html

²<https://www.netflix.com/ca/title/81328723>

Table I: The covered topics in the Engineering AI-Based Software Systems and Social Aspects of AI-based Systems courses.

Engineering AI-Based Software Systems	Social Aspects of AI-based Systems
Introduction & overview	Terminology around social aspects and EDI
AI for Software Engineers (hands-on activity)	Human rights and sustainable development
Software requirements for AI-based systems	Technology and politics
Software architectures of AI-based systems	AI and materiality
Data validation and management	AI, industry and employment
Interpretation vs. Explanation	AI harm, issues and concerns
Deployment and testing (MLOps)	AI ethics
Technical Debt in AI-based systems	AI policies and regulations
Continuous delivery (Guest lecture)	AI and SE culture, practices, and values
Special topics: Security, Fairness, Privacy, AI Governance	Indigenous perspectives on AI

These skills are essential for trainees to evaluate different solutions and make an informed decision by selecting the most effective solution for a specific problem. To enrich class discussion and ensure trainees are familiar with the material that will be presented in class, for each topic, trainees are assigned materials to read, summarize and critique. During the lecture, we start by providing a general background about the class topic to familiarize trainees with the related terminologies and motivate the importance of this topic from both engineering and social aspect perspectives. Then, trainees are empowered and encouraged to discuss the article(s), taking turns listing what they have learned, aspects that they deem to be the strong points and limitations. In addition to gaining a better grasp of the concepts discussed in class, this approach helps trainees improve their active listening skills, better express themselves and encourages them to be receptive to different viewpoints. We observe that trainees engage well in the discussion, debating the content of the article and reflecting on the virtues and shortcomings of the article content.

Practice-based learning. To exercise the knowledge learned in class and sharpen the trainees’ problem-solving skills, the training program provides the trainees with practical experience through several assignments, essays, and projects. The trainees complete several projects through the program courses. For example, projects in the Engineering AI-based Software Systems course include the development of tools, techniques and/or empirical studies on established practices in the industry and open-source. In the Social Aspects course, trainees are assigned a team-based project that assesses AI applications for a specific domain (e.g., healthcare and environment) with respect to EDI. In addition to exploring theories and concepts related to EDI, the project provides trainees with an opportunity to learn and enhance their research grant writing skills. A number of projects have been submitted and published to peer-reviewed venues in software engineering (e.g. [13, 14]).

Bridge the gap between practice and academia. A key goal of the program is to bridge the gap between academia and industry practice, promoting a bi-directional flow of knowledge. This means that knowledge should be shared from classrooms to the industry and vice versa. To expose trainees on how the theory learned in the course is applied

in the industry, we invite keynote speakers from the industry to present webinars and share the current practices used in the industry to develop AI-based systems. In past webinars, speakers highlighted the different critical challenges and solutions they used to overcome real challenges in their systems. Issues related to fairness, performance, usability, and even team composition are some topics that were covered. To export education to the industry, our program encourages internships, trainee-led seminars for industry partners, and collaborations between educational institutions and industry organizations through joint research projects. Such efforts direct the trainees’ research to focus on proposing practical solutions for real-world issues in the industry and expose the trainees to developing solutions that are industry-relevant.

Provide trainees with opportunities to develop highly sought-after professional skills. One of the main goals of our training program is to cultivate the professional skills of the trainees and prepare them for their future career paths. Therefore, we include specialized modules in our program that focus on providing trainees with valuable and relevant professional skills in the context of AI-based software systems. The program developed a set of professional development modules where each module encompasses a collection of reading materials, lectures, and workshops that pertain to a specific goal. The modules are meant to be lightweight, taking 1-2 weeks to complete, allowing trainees to take them at their convenience without excessively extending the duration of their program. Currently, we offer five specialized modules in our training program:

- **Commercialization and Entrepreneurship.** The significance of this module lies in the broad applicability and high demand for AI-based solutions. The objective of this module is to introduce trainees to concepts related to business development, strategic planning, and entrepreneurship in order to commercialize their ideas into profitable ventures.
- **Engagement and Relation-building with Communities and Diverse Stakeholders.** This module focuses on the socio-cultural aspects of communication. Trainees learn protocols and ethics of building relationships, particularly those that are relevant for engaging with communities (e.g., indigenous communities).
- **Communication and Explainability.** Trainees learn to effectively communicate and explain the composition of

AI-based systems and the interpretation of their outputs.

- **Dissemination and Presentation Skills.** This module focuses on best practices to present technical concepts related to AI-based software systems in a clear and convincing manner to a broad audience.
- **Technical Writing and Critiquing.** A specialized module on technical paper writing and critiquing, focusing on literature related to AI-based software systems (both technical and societal literature). The main goal of this module is to teach trainees how to transform their research into high-quality, crisp manuscripts, as well as how to write constructive reviews of research papers.

We strongly encourage trainees to complete as many of the aforementioned modules as possible but must complete at least two. That said, trainees have the freedom to select the modules that align with their expected future career paths (e.g., academic researcher, entrepreneur).

In addition to working on their research thesis that is related to the program’s theme, the trainees are initially required to complete the Engineering AI-based Software Systems and Social Aspects of AI-based Systems courses. Furthermore, the trainees need to complete two additional modules from the specialized modules to finish the program successfully.

III. PRELIMINARY FEEDBACK ON THE PROGRAM

To gauge the quality of the training program, we sought initial feedback from the trainees through a survey. We invited the first-year trainees to participate in the survey³. We asked participants questions related to what they liked about the program, as well as areas in which they felt the program could be improved. We sent emails to 18 trainees and received nine responses (i.e., 50% response rate). In the following, we report the results of the survey, including example responses from participants (tagged P1 through P9):

A. Strengths

Through the survey, participants indicated several strengths of the training program, which include:

Improving critical thinking and problem-solving skills. The survey results show that all the participants either strongly recommend (77.8%) or recommend (22.2%) the training program to others. Our analysis of the responses reveals that the participants feel that the program helped them to identify and address issues related to AI-based systems. For example, P7 mentioned, “*We learned about the technical and social issues around AI in various ways during the program*”. Also, participants emphasized the importance of the provided materials and projects in exposing them to different solutions. For example, P8 highlighted how the program helped them think critically about the potential impact of different solutions, “*It (the program) helped me by teaching me various aspects to consider when analyzing a given problem and the ability to think critically of the consequences (good and bad) for each possible solution*”. This illustrates the impact of providing materials that encompass both the issues and their

potential solutions, as well as the in-class discussions that broaden the trainees’ knowledge through exposure to diverse perspectives and experiences of other trainees.

Increasing awareness of social and ethical issues. 89% of participants indicated that the training program helps trainees to identify and address the social aspects of developing AI-based systems. P8 highlighted how discussing real-life examples during the program increased the participant’s awareness of the social aspects and issues related to AI-based systems, “*It (the program) helped me by making me more aware of the social issues experienced in AI-based software systems through previous real-life examples and offering different technical approaches in potential solving such issues*”. Also, P4 reported that “*the course on the social aspects of AI helped me recognize the existence of many threats introduced by AI (such as bias, lithium mining, privacy issues). I was unaware of most of these issues*”. Moreover, the participants demonstrated how the technical knowledge gained in the program could be applied to address social issues such as biases in AI-based systems, “*I can give the example of fairness here. So, we not only learned about fairness but also explored datasets (with hands-on activities), and we also explored the state-of-the-art solution to fix fairness problems*”. Although all trainees who enrolled in the program are enrolled in graduate-level degrees (MSc. and Ph.D.) and come from various backgrounds, they were previously unaware of the social and ethical issues related to AI. Our results demonstrate the importance of incorporating ethical and social implications in our program to train future software engineers.

Preparing future responsible practitioners/researchers.

The majority of participants (66.7%) strongly believe that the social education and training they received in our program were adequate for them to become responsible practitioners and researchers. Additionally, 33.3% of participants agreed that the education and training they received in this area was sufficient. For example, P3 will act proactively to solve social aspect potential issues when developing AI-based systems, “*When I build AI applications or do research, I will be aware of my social responsibility to build fair/unbiased applications and avoid using sensitive data (gender/race) to make decisions*”. On the other hand, P7 discussed how the program impacted the way of conducting and reporting research results, “*By becoming aware of the technical, social and ethical concerns around AI, I am able to design my research in a way that is concise about these issues (e.g., model drift, AI fairness)*”. The responses demonstrate that the participants relate responsibility both to technical as well as social knowledge and awareness, and highlight the participants’ awareness of the importance of an interdisciplinary approach.

B. Potential Improvements

The survey participants also shared insightful recommendations for improving the program. One important aspect they highlighted was the wish to have more involvement and collaborations with the industry through internships. For instance,

³<https://forms.gle/WnvZQhm1z58k2V9d9>

P3 reported, “*Would it be possible to communicate student research activities with industry partners of the training program so that internships/employments opportunities are direct*”. While the internship in our program is currently optional and highly recommended for all trainees, it is clear that internships and industry experience are particularly important for training programs such as ours. Luckily, given the topic of the training program, our trainees are in demand, and securing internships has not been a major issue. Related to the above, survey participants would like to include even more practical sessions in the program. For example, P5 indicated, “*More hands-on experience (e.g. deploying and maintaining a model, setting up monitoring, setting up experiments tracking (mlflow), etc.)*”. Such improvements can be achieved by adding more technical assignments and including a course laboratory in our program, allowing trainees to apply classroom concepts to real-world scenarios. Given that our training program is a multi-institution program, involving four universities, across three provinces/states, participants indicated the desire to have more collaborations among trainees. For example, P1 reported, “*I wish this program can make more collaboration with trainees who have different backgrounds*”. This suggestion may have arisen from in-class discussions, which help trainees recognize the advantages of collaborating with peers from diverse backgrounds.

IV. LESSONS LEARNED

In this section, we present the lessons we have learned from creating and teaching the courses. The first lesson learned is that teaching the social aspects of developing AI-based systems is just as critical as teaching the technical aspects. This is evident by the majority of the trainees’ (89%) responses, which expressed that the training program effectively helped them identify and address these social aspects. Another lesson learned is that critiquing articles is crucial for fostering critical thinking among trainees. We have observed a gradual increase in the trainees’ critical thinking throughout the course as it progresses. They have provided deep insights into the consequences of solutions from various perspectives. Finally, it is critical to include an industrial perspective on developing AI-based systems. Therefore, we discussed real-world examples to demonstrate the challenges facing practitioners and their solutions. Furthermore, we implemented hands-on learning lectures and invited keynote speakers from the industry to present the current practices used to develop AI-based systems. This has helped the trainees to reflect on what they have learned and how it applies in real-world industry settings.

V. THREATS TO VALIDITY

Considers the relationship between theory and observation, in case the measured variables do not measure the actual factors. We surveyed the first-year trainees to assess the quality of the training program from their perspective, which might not directly reflect the program’s overall quality.

The questions in the survey might bias the participants’ answers. However, most participants acknowledged the program’s usefulness on all 5-point Likert scale questions. Also,

the survey was conducted anonymously, and participants were free to respond to the questions in their preferred manner. Nonetheless, the anonymous nature of the survey presents another threat to internal validity, as we may lack context surrounding respondents’ experiences, possibly leading to misinterpretation of specific answers. However, we believe the responses are relevant to the program and its components. This is because we only surveyed first-year trainees enrolled in the program. Moreover, we explicitly stated that the purpose of the survey is to gather feedback on the training program, both in our invitation email and at the beginning of the survey. To gauge initial feedback about the program, we surveyed a cohort of students. In the future, we plan to expand the survey to include other trainees and their employers.

VI. CONCLUSION

This work describes our experience developing a training program that prepares trainees to address the technical, social, and professional aspects of developing, deploying and servicing AI-based software systems. This is complemented by a feedback survey from the trainees who attended the program, and the survey results indicate that the participants are trained to identify and address social, ethical, and technical concerns in developing AI applications. These results also provide us with suggestions to improve the training program. For example, we plan to expand our collaborations with the industry by offering more research opportunities in industrial contexts and internships. Furthermore, we will include more practical sessions in the program for future cohorts to enable trainees to apply classroom concepts to real-world scenarios. In the future, we plan to assess the effectiveness of the training program by conducting a comprehensive user study from the perspectives of both the trainees and their employers to measure the quality of the training program and gain insights into how the program prepares trainees for their career paths. Furthermore, we are planning to investigate the extent of learning that takes place in the classroom compared to on-the-job experiences.

REFERENCES

- [1] L. Chen, P. Chen, and Z. Lin, “Artificial intelligence in education: A review,” *IEEE Access*, vol. 8, pp. 75 264–75 278, 2020.
- [2] L. Cao, “Ai in finance: Challenges, techniques, and opportunities,” *ACM Comput. Surv.*, vol. 55, no. 3, feb 2022.
- [3] K.-H. Yu, A. L. Beam, and I. S. Kohane, “Artificial intelligence in healthcare,” *Nature biomedical engineering*, vol. 2, no. 10, pp. 719–731, 2018.
- [4] E. Bertino, M. Kantarcioglu, C. G. Akcora, S. Samtani, S. Mittal, and M. Gupta, “Ai for security and security for ai,” in *Proceedings of the Eleventh ACM Conference on Data and Application Security and Privacy*, 2021, pp. 333–334.
- [5] A. Vogelsang and M. Borg, “Requirements engineering for machine learning: Perspectives from data scientists,”

- in *2019 IEEE 27th International Requirements Engineering Conference Workshops (REW)*, 2019, pp. 245–251.
- [6] B. W. Wirtz, J. C. Weyerer, and C. Geyer, “Artificial intelligence and the public sector—applications and challenges,” *International Journal of Public Administration*, vol. 42, no. 7, pp. 596–615, 2019.
- [7] OECD, “The OECD Artificial Intelligence (AI) Principles.” [Online]. Available: <https://oecd.ai/en/ai-principles>
- [8] J. Dastin, “Amazon scraps secret ai recruiting tool that showed bias against women,” in *Ethics of data and analytics*. Auerbach Publications, 2018, pp. 296–299.
- [9] T. Simonite, “The best algorithms still struggle to recognize black faces | wired,” <https://www.wired.com/story/best-algorithms-struggle-recognize-black-faces-equally/>, 02 2023, (Accessed on 02/27/2023).
- [10] J. Buolamwini and T. Gebru, “Gender shades: Intersectional accuracy disparities in commercial gender classification,” in *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, ser. Proceedings of Machine Learning Research, S. A. Friedler and C. Wilson, Eds., vol. 81. PMLR, 23–24 Feb 2018, pp. 77–91. [Online]. Available: <https://proceedings.mlr.press/v81/buolamwini18a.html>
- [11] G. Hulten, *Building Intelligent Systems: A Guide to Machine Learning Engineering*. Apress, 2018.
- [12] K. Crawford, *The atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press, 2021.
- [13] K. Badran, P.-O. Cote, A. Kolopanis, R. Bouchoucha, A. Collante, D. Costa, and E. Shihab, “Can Ensembling Pre-processing Algorithms Lead to Better Machine Learning Fairness?” *IEEE Computer SI: Software Engineering for Responsible AI*, 2022.
- [14] V. Majdinasab, S. C. Y. Ho, M. Islam, D. E. Costa, E. Shihab, F. Khomh, S. Nadi, and M. Raza, “An empirical study on bugs inside pytorch: A replication study,” in *Proceedings of the 40th International Conference on Software Maintenance and Evolution (ICSME’23)*, 2023.