

San Francisco State University Graduate Certificate in Ethical AI – overview and early experiences

Denise Kleinrichert , Carlos Montemayor and Dragutin Petkovic

San Francisco State University, 1600 Holloway Ave., San Francisco CA 94132 USA

Abstract

This paper overviews San Francisco State University's (SFSU) Graduate Certificate in Ethical Artificial Intelligence launched in Fall 2019 by combined efforts of the Computer Science department, department of Philosophy, and the College of Business. It is designed to offer both professionals and graduate students (matriculated or non-matriculated, AI experts or not) the opportunity to acquire a deeper grasp of the ethical, legal, policy, and technical issues of AI, thus enabling them to effectively participate in the development or auditing of ethical and trustworthy AI systems and applications across a spectrum of employment sectors.

Keywords

AI Ethics, education, certificate, SFSU

1. Introduction

Recent advances in Artificial Intelligence (AI) technologies have enabled AI systems and applications to impact all aspects of our lives, including communication, health care, business and work practices, defense, transportation as well as media, education and entertainment. In spite of the improved performance in these domains, AI systems may produce errors, demonstrate bias, and may lack fairness, technical and judicial decision-making transparency resulting in the reduction in human trust that challenges in broader AI adoption. Given our goals of creating and maintaining socially responsible AI that promotes well-being for broad society these issues are critical to address and are being discussed not only among AI professionals but also among general public and non-AI experts. Given these challenges one comes to the conclusion that there is a strong need for education about trustworthy and ethical AI especially for broader public, AI users, as well as politicians, lawyers and auditors who may not be deep experts in AI technologies but will

increasingly have to make informed decisions about its use, impact, risks and yes, regulations.

San Francisco State University's (SF State) Graduate Certificate in Ethical Artificial Intelligence is motivated by these concerns and importantly it is multi-disciplinary in that it combines a number of academic disciplines (i.e., business, computer science, philosophy, etc.) to best prepare undergraduate and graduate students for the workforce of the 2020's and beyond. It was launched in Fall 2019 by SF State with the combined efforts of the Computer Science department, department of Philosophy, and the College of Business as the key components of Ethical AI (<https://bulletin.sfsu.edu/colleges/science-engineering/computer-science/certificate-ethical-artificial-intelligence/#requirementstext>). It is designed to offer both professionals and graduate students (matriculated or non-matriculated, AI experts or not) the opportunity to acquire a deeper grasp of the ethical, legal, policy, and technical issues of AI, thus enabling them to effectively participate in the development or auditing of ethical and trustworthy AI systems and applications across a spectrum of employment sectors. The certificate curriculum is mentored by

AAAI 2023 Spring Symposia, Socially Responsible AI for Well-being, March 27–29, 2022, USA

EMAIL: dk@sfsu.edu (A. 1); cmontema@sfsu.edu (A. 2); petkovic@sfsu.edu (A. 3)



© 2020 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

three faculty from CS, Philosophy and Business (this paper's authors) and each student chooses one mentor from these three faculty.

Educational philosophy and goals of our Ethical AI graduate certificate ensure that the students will master: a) the basic principles and technologies of AI with the focus on related issues of accuracy estimation and evaluation, detection of bias and explainability/transparency; b) how to balance AI ethics with business needs and apply meaningful ethical decision-making across various business areas; c) the ways to enforce ethical standards and compliance in addition to laws and rules related to the ethical use of AI; d) the significance and impact of ethical AI for society; e) how to analyze and apply knowledge of ethical AI in the form of a case study. To address these critical objectives, the graduate certificate consists of three graduate course segments, and importantly a short research and reflection paper requiring graduates to address and understand specific ethical AI issues.

The three segments are (a) *AI Technologies and Application*; (b) *Ethics & Compliance in Business (3 units)*; and (c) *Ethical Principles in Philosophy (3 units)*. The culminating *Research and Reflection paper* (carries 1 credit unit) is done under the supervision of one of the certificate mentors.

AI Technologies and Applications segment: One of the key teaching challenges in this certificate is how to effectively educate a diverse body of students from a wide variety of disciplines where some might have low or no AI background. This is done in a CS graduate course "AI Explainability and Ethics" which covers: case studies of ethical AI issues; overview of AI with emphasis on proper analysis and evaluation methods and metrics of AI accuracy (key knowledge non-AI professionals will need in order to be able to audit AI systems); more detailed understanding of one chosen AI method (Random Forest); and overview of best practices for ethical and trustworthiness audits of AI systems. Students perform a number of individual assignments including an example of programming of Random Forest pipeline (challenge for some who never coded) and then in teams of 2-3 students (where teams consists of students from different disciplines) they perform a final class project of a complete ethics and trustworthiness audit of a chosen AI application. Note the important class focus in use and

evaluation vs. development of AI technologies. Official student class feedback is very positive and the class has also enabled students from non-CS disciplines to get better paid jobs using new acquired skills in practical AI.

Ethics & Compliance in Business segment: This course emphasizes the managerial role of strategic and operational ethical decision-making in business and organizations, with emphases on artificial intelligence and technology, human resources, and risk management. The focus is on integrity-based thinking and action by applying managerial ethical decision-making skills, understanding regulatory compliance as it relates to organizational functions, evaluating stakeholder impacts, and developing ethics training and compliance procedures, including human resources and risk management precepts. Application models, case studies, interactive exercises, and industry guest speakers are used throughout the semester of this course.

Ethical Principles in Philosophy segment: This course covers the implications of artificial intelligence for policy, industry, and society at large, including issues concerning social justice, from a philosophical perspective. The main question of the course is, in what sense(s) is artificial intelligence "artificial" and "intelligent." This course explores this issue through the relation of two fundamental capacities of the human mind: phenomenal consciousness and attention. Ethical issues are normative, and the question of how intelligent capacities are related to norms, reasons, and expectations are analyzed in detail, including the development and regulation of artificial intelligence.

Research and Reflection Paper: While most certificate programs do not contain final paper we felt that requiring students to reflect on a topic related to AI ethics in the form of a 10-15 page paper is important. This paper is done under mentoring of one of the three faculty in charge who helps the student choose the topic and advises them through the research and writing process.

Evaluation of the certificate is ongoing, Initial feedback is positive and we believe that the interest in the program will be growing. Recently we see uptick of non-matriculated enrollment from local industry and universities. Since its start in 2019 SF State graduate certificate in Ethical AI enrolled 21 students from variety of backgrounds (several non-matriculated) and awarded 5 certificates.