

Integrating ethics in technical AI curricula: A pilot study on college students' perceptions and learning outcomes

Theodora Chaspari, University of Colorado Boulder, Boulder, CO
James K. L. Hammerman, TERC, Cambridge, MA

The absence of AI ethics education in university classrooms can lead to the development and deployment of technologies without adequate consideration for their potential societal impacts. With a growing demand on integrating ethics into existing Computer Science (CS) courses, this need is particularly pressing for AI-related classes, given the widespread and contentious societal implications of AI. We integrated active learning through hands-on activities, projects, and critical reflection on AI ethics into a graduate-level technical AI course. Using a mixed-methods approach involving quantitative and qualitative data analysis, we seek to address the following research questions: (1) How do college engineering students perceive issues of ethical design and implementation of AI before engaging in course activities related to AI ethics? and (2) How are these views different after experiencing a technical AI course that integrates discussion of ethics? Data come from a survey that includes Likert style questions about students' perceptions across multiple ethical dimensions such as overall ethical responsibility of AI designers, care for individuals' personal data and privacy, understanding of AI's outputs and decision-making processes, and equity of benefits derived from AI applications. The survey also includes open-ended questions regarding students' perceptions of the risks associated with applying AI to high-stakes applications, such as medical diagnosis and employee selection. Surveys were administered at the beginning and the end of the semester and changes compared. Likert items were analyzed quantitatively and statistical differences calculated; open-ended items were coded qualitatively and changes in perceptions described. Results will be reported on a sample of 56 CS students who attended a graduate-level Machine Learning class at the University of Colorado Boulder in Spring 2024. We will discuss design of the instrument, analysis of the results, and lessons learned with tentative design recommendations for future AI ethics curricula.