It's become something of a cliché to say that artificial intelligence (AI) is ubiquitous and has the potential to radically alter the way we live our lives. Nonetheless, these claims are for the most part true. AI is everywhere and will change everything.

This rapid expansion necessitates an ethical assessment. Ethical questions arise when we consider applications of AI to almost any sphere of life, in ways that sometimes are immediately clear and, at other times, unobvious and unexpected.

This course will cover some of the major themes in ethics of AI, robotics, and automation, concerning both present-day and near- and long-term future uses of these technologies. We will examine what sorts of social impact they are likely to have and evaluate this impact with the tools of ethical theory, applied ethics, and a smattering of political philosophy.

The readings will feature a variety of perspectives from a range of professionals: philosophers, AI researchers, social scientists, lawyers, and tech journalists, among others.

Topics may include:

- **The ethics of algorithms: opacity, explanation, persuasion, bias, and trust**
  Some of the most popular and powerful AI techniques today involve the use of machine learning algorithms. These algorithms sometimes work in ways that are opaque and impossible to understand even to their creators. Is this a problem? Should AI be made explicable to an average user? What about algorithms that produce outputs that appear biased? How should they be evaluated? Could algorithms be trusted?
- **Professional ethics for AI practitioners**
  What ethical rules should AI professionals abide by in their work? What are the various AI ethics codes developed by a number of leading tech companies? Are they adequate?
- **Automation and labor**
  AI promises to revolutionize the labor markets across the world. Is this a desirable development? Will robots take all our jobs? Should we be preparing for a life without work? How should we organize our societies' response to automation?
- **Autonomous vehicles and autonomous weapons**
  What does it mean for weapons or vehicles to be "autonomous"? Are such technologies desirable? If they make their decisions without human involvement, how should they be programmed? Who's morally and legally responsible when these machines do something harmful?
- **Robot rights?**
  Will it ever be possible to create machines that think and feel just as human beings do? If so, should they be created? How should they be treated? Should they be given rights or remain subservient to human beings?
- **Superintelligence**
  What is superintelligence? What is the likelihood that superintelligent machines will be developed? What are potential impacts of superintelligence? How should we prepare?
- And others