

Department of Philosophy and School of Computing

Philosophical Implications of Artificial Intelligence

Colloquium Series

Self-fulfilling Prophecy in Practical and Automated Prediction

Owen King
University of Twente



Thursday, January 30, 2020 ♦ 3:00 pm ♦ Watson Hall 517

Some seemingly innocent predictions turn out to be not so innocent after all. This is so with self-fulfilling prophecies, which are predictions that somehow bring about their own truth. In this talk, I will present philosophical research on self-fulfilling prophecies, with special attention to contexts of practical prediction. The focus on practical prediction, as opposed to scientific prediction, is motivated by the aim of addressing self-fulfilling prophecies in the use of predictive analytics. After introducing several examples, I will lay out a general theory and definition of self-fulfilling prophecy. Then I will develop a normative analysis. This analysis comes in four layers, reflecting interrelated problems regarding both moral and epistemic responsibility. I will close with a brief discussion of the risk of recommender systems producing self-fulfilling predictions of consumer preferences.



Algorithms are Not Neutral: Bias in Recommender Systems

Catherine Stinson
University of Bonn & University of Cambridge

Thursday, February 6, 2020 ♦ 3:00 pm ♦ Watson Hall 517

Efforts to shine a light on algorithmic bias tend to focus on examples where either the data or the people building the algorithms are biased. This gives the impression that clean data and good intentions could eliminate bias in machine learning. The apparent neutrality of the algorithms themselves is defended by high profile AI researchers and companies with an interest in business as usual, but algorithms are not neutral. In addition to biased data and biased algorithm makers, AI algorithms themselves can be biased. This is illustrated with the example of collaborative filtering (an algorithm commonly used in recommender systems), which is known to suffer from popularity, and homogenizing biases. The larger class of iterative information filtering algorithms create a selection bias in the course of learning from user responses to items that the algorithm recommended. These are not merely biases in the statistical sense; these statistical biases cause bias of moral import. People on the margins in the sociocultural sense are literally on the margins of data distributions, as work in disability studies has shown. Popularity and homogenizing biases have the effect of further marginalizing the already marginal, which means that "Customers who bought this item also bought..." style recommendations do not meet everyone's needs. This source of bias warrants serious attention given the ubiquity of algorithmic decision-making.

Mind Design and the Ethics of Parenting

Bartek Chomanski
Western University



Tuesday, February 11, 2020 ♦ 3:00 pm ♦ Watson Hall 517

According to some views on the matter, the ultimate goal of artificial intelligence (AI) research is to build an artificial person (whose moral status is no different than that of an adult human). Suppose this goal realized. A plethora of questions arises as a result, both having to do with what this development will mean for humans, and with what sorts of duties will humans have towards such beings. In this talk I will focus on one, relatively underexplored aspect of the latter issue. I will consider whether, and to what extent, the relationship between a designer of the AI and the AI themselves should be modelled on the relationship between parents and children. I will argue that we should be wary of building AIs in a way that resembles natural procreation and childrearing. While typical childrearing generally involves placing children in the condition of vulnerability (that is, in the condition where the children depend on the adults for their continued survival), this need not be an aspect of designing artificial persons. I will argue that placing another in the condition of vulnerability, when it could be avoided, is wrong. Hence, the AI design process should not be made to resemble childrearing in that regard. This remains true even if one person's vulnerability to another is necessary for the development of a uniquely robust emotional connection between them.