isr PRESENTS

Kevin Leyton-Brown

Modeling Nonstrategic Human Play in Games

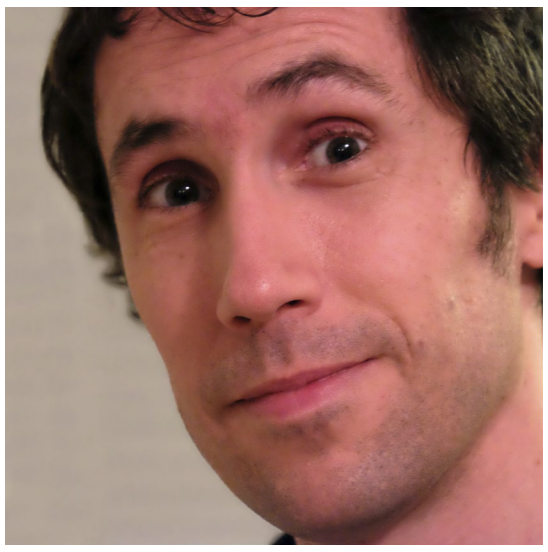

Kevin Leyton-Brown is a professor of Computer Science at the University of British Columbia and an associate member of the Vancouver School of Economics. He holds a PhD and M.Sc. from Stanford University (2003; 2001) and a B.Sc. from McMaster University (1998). He studies the intersection of computer science and microeconomics, addressing computational problems in economic contexts and incentive issues in multiagent systems. He also applies machine learning to various problems in artificial intelligence, notably the automated design and analysis of algorithms for solving hard computational problems.
He currently advises Auctionomics, AI21, and Qudos. He is a co-founder of Kudu.ug and Meta-Algorithmic Technologies. He was scientific advisor to UBC spinoff Zite until it was acquired by CNN in 2011. His past consulting has included work for Zynga, Trading Dynamics, Ariba, and Cariocas.

It is common to assume that players in a game will adopt Nash equilibrium strategies. However, experimental studies have demonstrated that Nash equilibrium is often a poor description of human players' behavior, even in unrepeated normal-form games. Nevertheless, human behavior in such settings is far from random. Drawing on data from real human play, the field of behavioral game theory has developed a variety of models that aim to capture these patterns. The current state of the art in that literature is a model called quantal cognitive hierarchy. It predicts that agents approximately best respond and explicitly model others' beliefs to a finite depth, grounded in a uniform model of nonstrategic play. We have shown that even stronger models can be built by drawing on ideas from cognitive psychology to better describe nonstrategic behavior. However, this whole approach requires extensive expert knowledge and careful choice of functional form. Deep learning presents an alternative, offering the promise of automatic cognitive modeling. Leveraging a novel architecture that allows a single network to generalize across different input and output dimensions by using matrix units rather than scalar units, we have shown that even better predictive performance can be achieved. However, the success of such approaches raises a more fundamental question: at what point does such behavior get so complex that it ought to be considered strategic? The typical answer is that one should check to see whether the behavior involves modeling other agents: strategic agents do so, while nonstrategic agents do not. However, this is not a wholly satisfying answer, because we lack theoretical tools for arguing that a complicated, apparently nonstrategic behavior cannot be rephrased in strategic terms. We overcome this hurdle by introducing a new, formal characterization of nonstrategic behavior that satisfies two properties: (1) it is general enough to capture all purportedly "nonstrategic" decision rules of which we are aware; (2) we prove that behavior obeying our characterization is distinct from strategic behavior in a precise sense.

Wednesday, April 3, 2019
2:00pm to 3:00pm
3305 Newell Simon Hall

isr institute for SOFTWARE RESEARCH

School of Computer Science
5000 Forbes Avenue, 5212 Wean Hall
Pittsburgh, PA 15213-3890

Carnegie Mellon University