

Dynamic Control Under Changing Goals

Zach Davis

NYU, New York, New York, United States

Neil Bramley

University of Edinburgh, Edinburgh, Scotland, United Kingdom

Bob Rehder

New York University, New York, New York, United States

Todd Gureckis

New York University, New York, New York, United States

Abstract

Acting effectively in the world requires a representation that can be leveraged to serve one's goals. One practical reason that intelligent agents might learn to represent causal structure is that it enables flexible adaptation to a changing environment. For example, understanding how to play a videogame allows one to pursue other goals such as doing as poorly as possible or only gathering one type of item. Across two experiments that manipulated the expected utility of learning causal structure, we find that people did not build causal representations in dynamic environments. This conclusion was supported by behavioral results as well as by participants being better fit by models describing them as utilizing minimally complex, reactive control policies. The results show how despite being incredibly adaptive, people are in fact computationally frugal, minimizing the complexity of their representations and decision policies even in situations that might warrant richer ones.