

Machine Learning Methods for Individualizing Real-Time Treatment Policies

Susan A. Murphy*

University of Michigan, Ann Arbor, MI USA, samurphy@umich.edu

Richard S. Sutton

University of Alberta, Edmonton, AB CANADA, sutton@cs.ualberta.ca

Mobile devices are being increasingly used by health researchers to, in real-time, both collect symptoms and other information as well provide interventions. These interventions are often provided via treatment policies. The policies specify how patient information should be used to determine when, where and which intervention to provide. Here we present generalizations of methods from the field of Reinforcement Learning for use in individualizing the treatment policies. These methods are related to updated and improved stochastic approximation algorithms used in robotics, online games and online advertising.

Key Words: Reinforcement Learning, Online Learning, Treatment Policies, Real-Time Dynamic Treatment Regimes