BRYDON EASTMAN, University of Waterloo

Contrasting Chemotherapy Schedules from Reinforcement Learning and Optimal Control

Determining the optimal chemotherapy schedule for a cancer patient is difficult. Increasing the dose of the chemotherapeutic results in both greater cancer kill but also greater toxicity to the patient in the form of red blood cell and bone marrow loss. We present two methods of deriving optimal chemotherapeutic schedules for a particular model of breast cancer treated with paclitaxel. The first, a schedule learned via deep reinforcement learning and the second, a schedule learned via optimal control theory. We demonstrate that the deep RL derived therapy provides a more robust chemotherapeutic schedule when we concern ourselves with interpatient variability.