



BiophysTO Lunchtime Seminar Series

Date

Thursday, Oct 24 2019
12:00 – 1:00 pm

Location

McLennan Physical
Laboratories
Room MP606
60 St. George Street

Pizza and refreshments will
be provided

Dr. Douglas Tweed

Dept of Physiology, Collaborative Program in
Neuroscience (CPIN), University of Toronto

Costate methods for reinforcement learning

Many control systems, in the brain and in computers, can improve with practice, by adjusting their own internal processing based on feedback from sensors — a process called *reinforcement learning*. Recent work in AI has led to powerful new algorithms for reinforcement learning, most of them based on the concepts of *action-value functions* and the *Bellman equation*, but I will describe a different approach based on the *costate equation*, studied by the great control theorist Pontryagin and others in the 1950s. Costate methods have been neglected in AI, probably because they require that the control system possess a *model*, or in other words an internal simulation of the process it is trying to control. For instance, a control system learning to throw a ball would need a model including its throwing arm, the ball, and the relevant laws of physics. I will show that the costate equation provides a simple way to learn useful models even in complex tasks, and I will compare the resulting algorithm to a leading Bellman-based method.

Host: Wilson Zeng



**Seminar
Sponsors**

UTSG

Biochemistry

Physics

Chemistry