Bang-Bang Biology: Learned Control for Drug Dosing

Josiah C. Kratz
Jacob A. Calzyk

Carnegie Mellon University

UMass Boston

Project Goal: RL control of dynamical systems

- Application: development of optimal temporal treatment protocol

- Motivation: in clinical setting, don’t have easy access to model parameters

Toy model: Phenotypic switching with memory

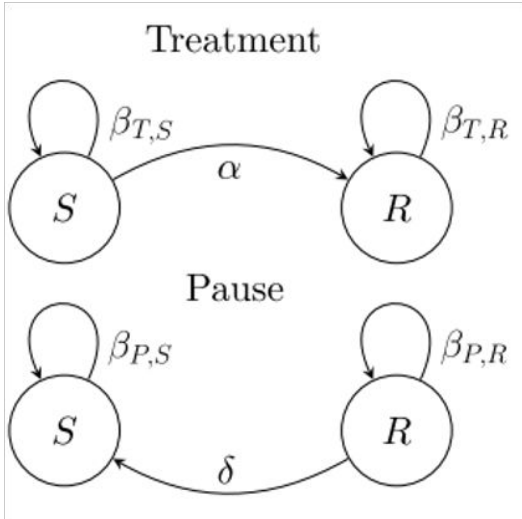

Fischer and Bluthgen 2023

Nonlocal dynamics via memory kernel

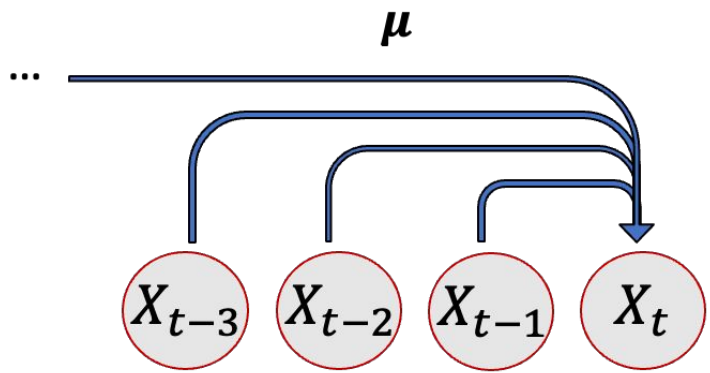

$\boldsymbol{\mu}$ determines memory strength

$0 < \mu_i \leq 1$

Memoryless case has a known solution

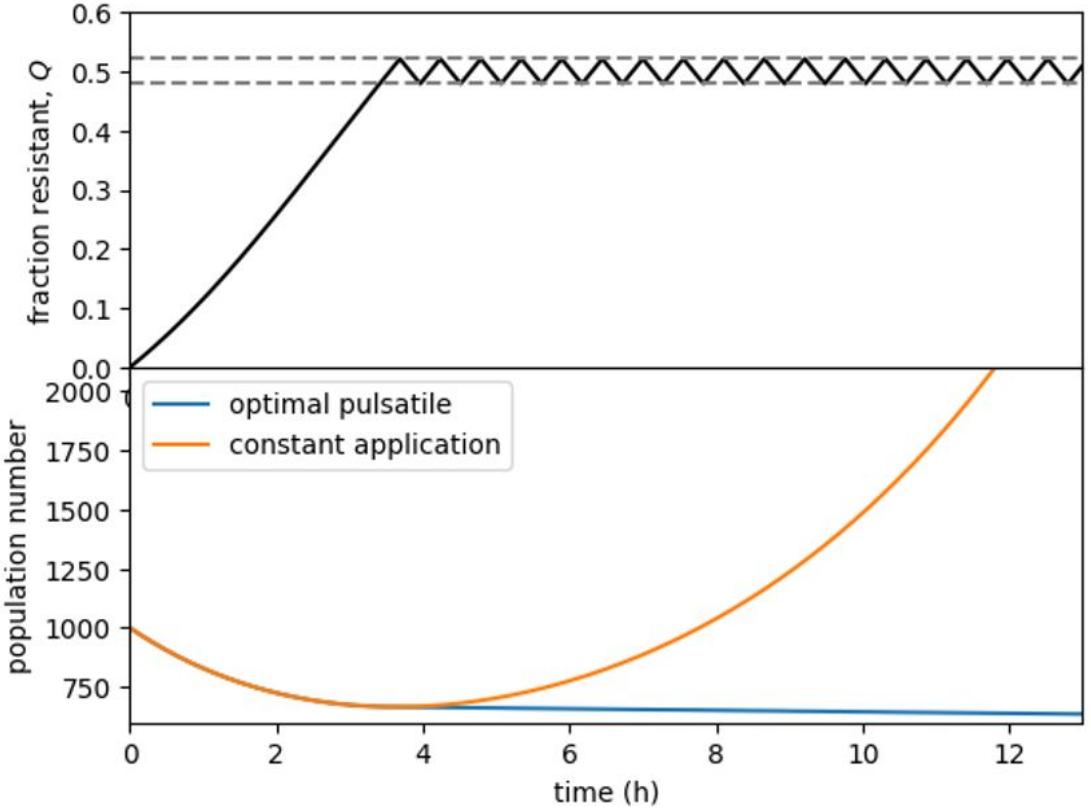

Adding memory increases population resilience

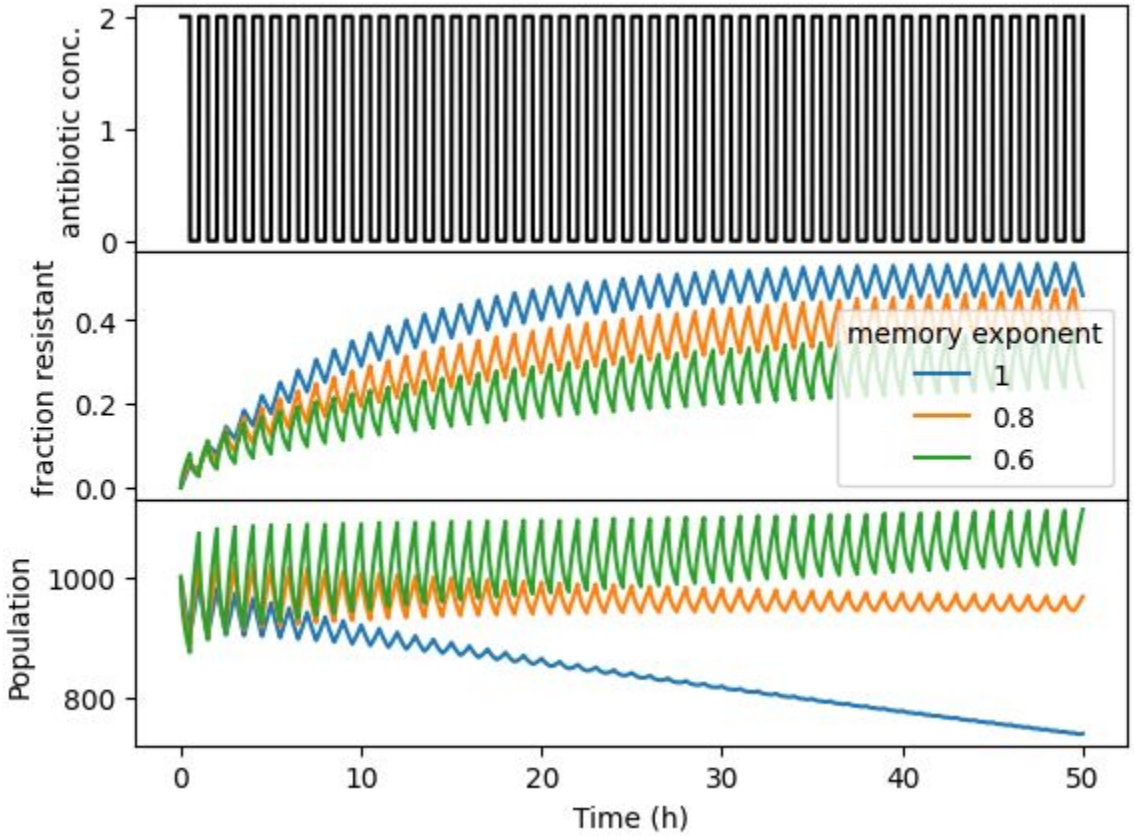

Applying Reinforcement Learning

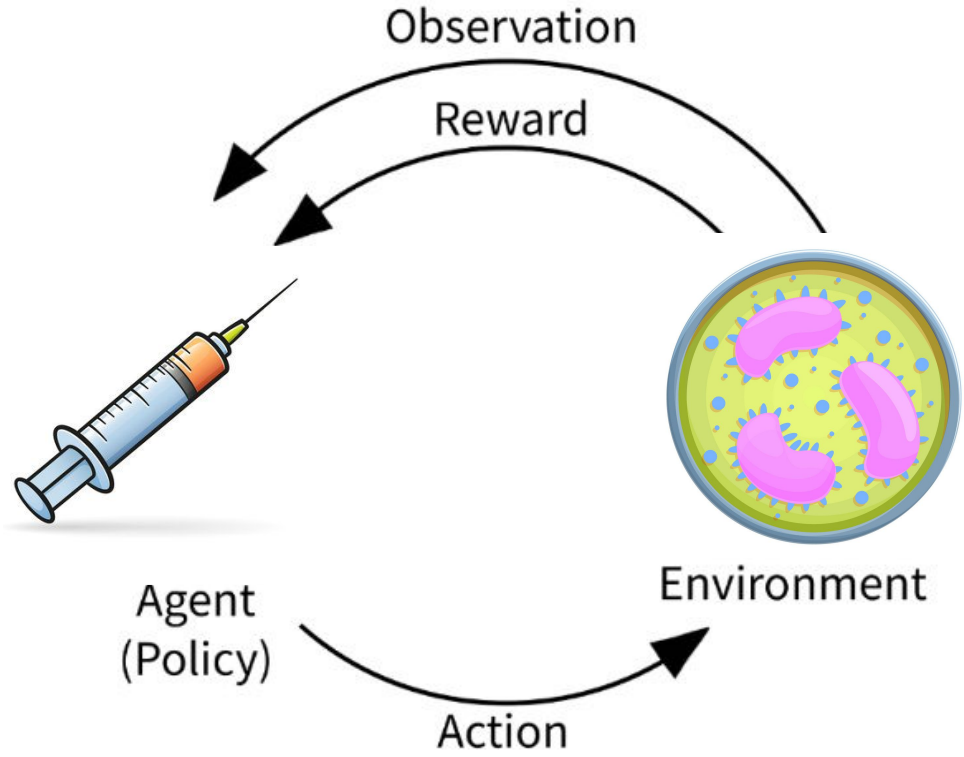

Using RL to find a dosing strategy

- Non-Markovian system, so we encode history in current state via “framestacking”
- DQN w/ MLP can recover memory-less solution

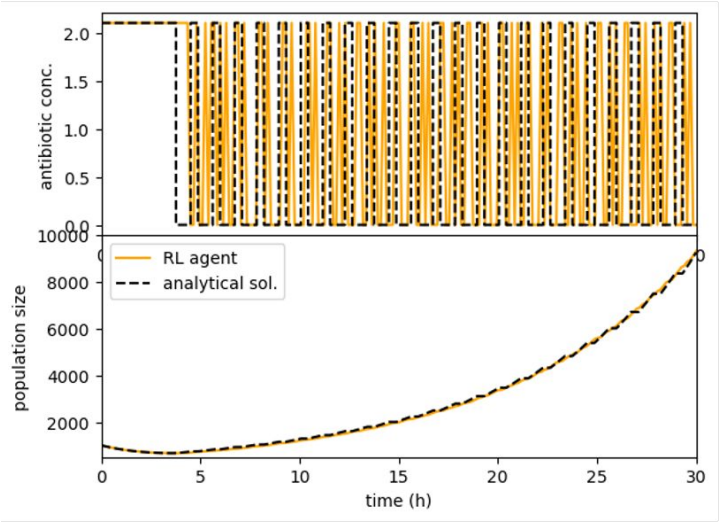

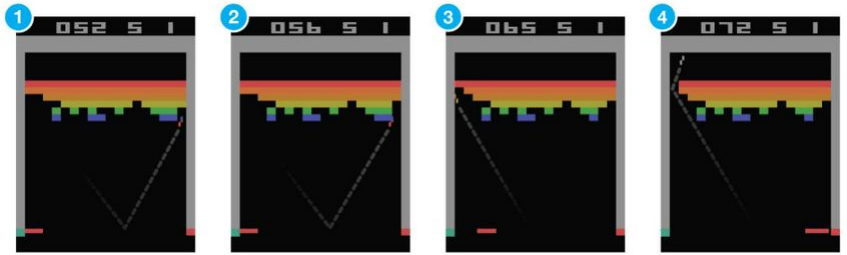

RL for switching model with memory

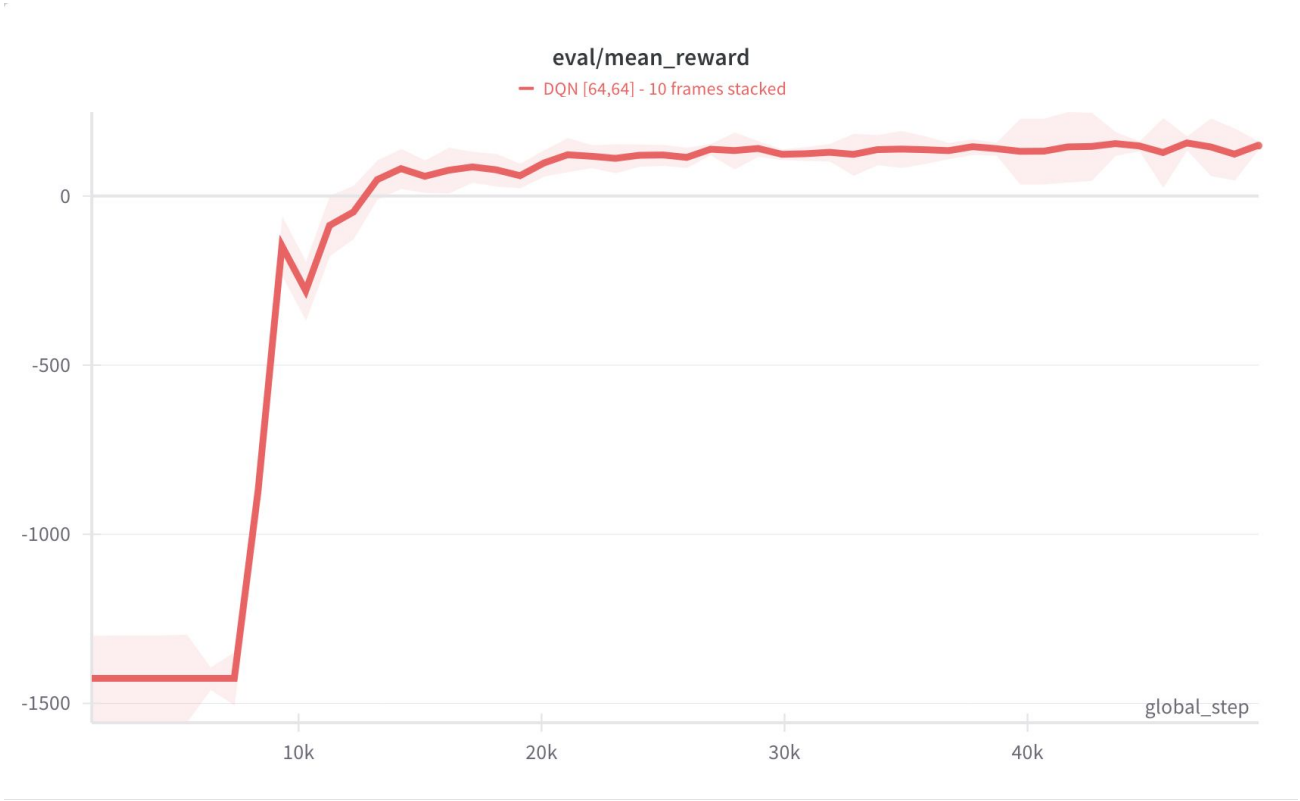

RL for switching model with memory

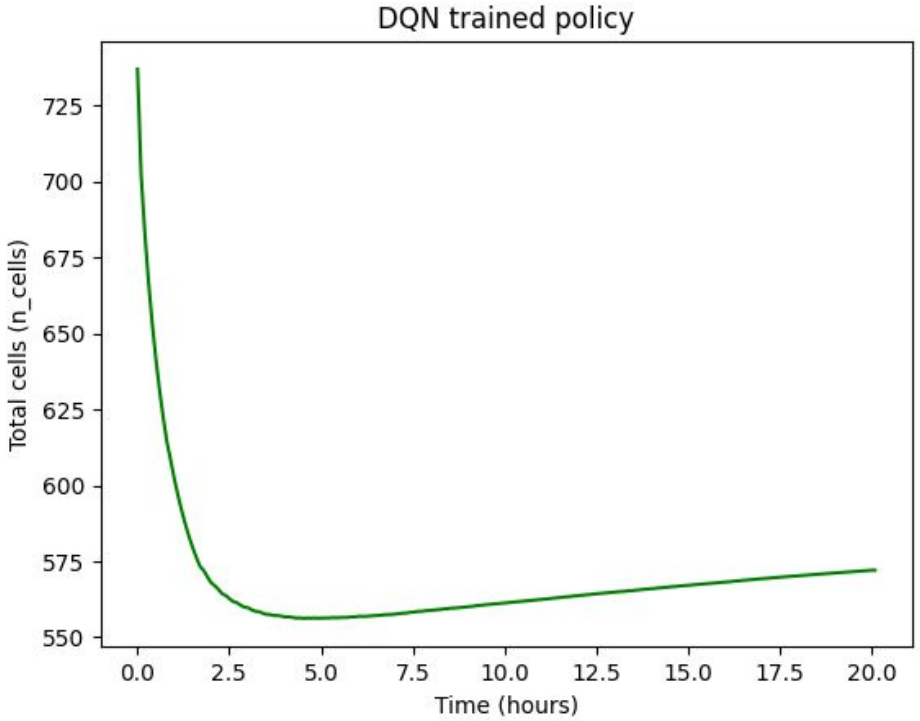

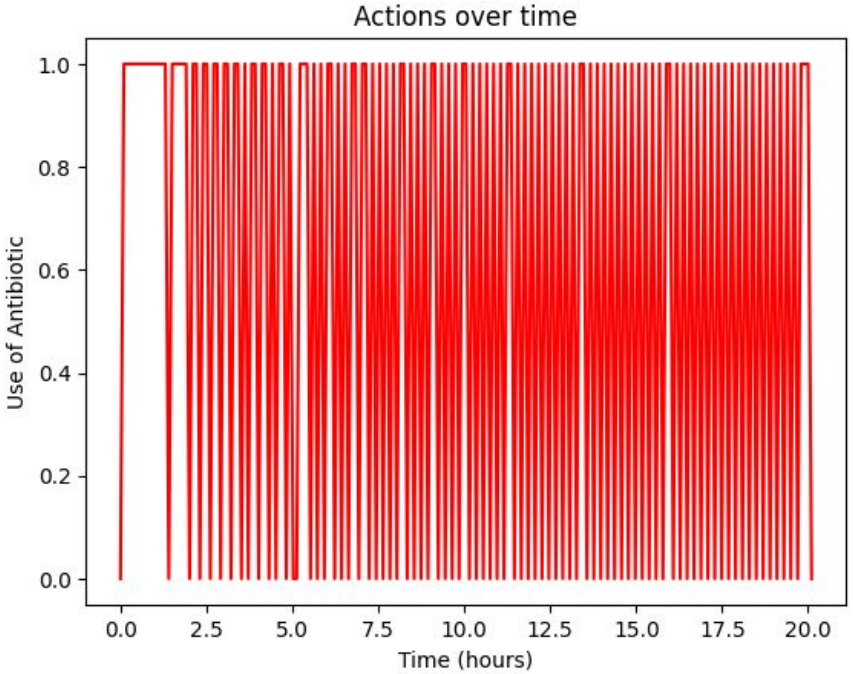

Future Work

LSTM policies

Stochastic dynamics

Transfer learning