CS394R
Reinforcement Learning: Theory and Practice

Peter Stone

Department of Computer Science
The University of Texas at Austin

Good Morning Colleagues

- Are there any questions?

Logistics

- Make progress on final projects!

Logistics

- Make progress on final projects!
- This week's readings: Health and Sustainablity

Logistics

- Make progress on final projects!
- This week's readings: Health and Sustainablity
  - Epilepsy, species management, and wind turbine optimization
- Next week's readings: least squares methods

Logistics

- Make progress on final projects!
- This week's readings: Health and Sustainablity
  - Epilepsy, species management, and wind turbine optimization
- Next week's readings: least squares methods
- Preliminary programming assignments due **before** last day of class (2 weeks)

Logistics

- Make progress on final projects!
- This week's readings: Health and Sustainablity
  - Epilepsy, species management, and wind turbine optimization
- Next week's readings: least squares methods
- Preliminary programming assignments due **before** last day of class (2 weeks)
- Final projects due 1 week later.

Class Discussion — Charles Orchard

- Simulators

Epilepsy Treatment

- Randomized trees
- On policy?
- Testing procedure

Epilepsy Treatment

- Randomized trees
- On policy?
- Testing procedure
  - Estimated proportion of seizure states
  - Number of stimulation events
  - Expected immediate reward
  - Expected return

Moving forward

- 1 year to get empirical validation of final policy in vitro.

Moving forward

- 1 year to get empirical validation of final policy in vitro.
- RL for health more of a pain than robots!

Moving forward

- 1 year to get empirical validation of final policy in vitro.
- RL for health more of a pain than robots!
- Plans for live animal (rats) tests this year

Moving forward

- 1 year to get empirical validation of final policy in vitro.
- RL for health more of a pain than robots!
- Plans for live animal (rats) tests this year
  - recording and detecting seizures now

Moving forward

- 1 year to get empirical validation of final policy in vitro.
- RL for health more of a pain than robots!
- Plans for live animal (rats) tests this year
  - recording and detecting seizures now
  - Next: exploration data

Moving forward

- 1 year to get empirical validation of final policy in vitro.
- RL for health more of a pain than robots!
- Plans for live animal (rats) tests this year
  - recording and detecting seizures now
  - Next: exploration data
- Humans: Already trials for detection and fixed stimulation in response

Moving forward

- 1 year to get empirical validation of final policy in vitro.
- RL for health more of a pain than robots!
- Plans for live animal (rats) tests this year
  - recording and detecting seizures now
  - Next: exploration data
- Humans: Already trials for detection and fixed stimulation in response
- Step to RL not so large

Class Discussion — Asher Johnson

- Under what circumstances would you trust a Reinforcement Learning algorithm with your health care?

Class Discussion — Asher Johnson

- Under what circumstances would you trust a Reinforcement Learning algorithm with your health care?
- What types of decisions would you allow it to make without human guidance?