Fitting Order Parameters through Reinforcement Learning

Kate Meuse - Cornell University
Jessica Opsahl-Ong - Rice University
Mentored by William Ratcliff

Goals and Impact

- Beam time is valuable – limited access
- Want more efficient measurements as not all are required
- Software to be implemented on multiple instruments (BT4, BT7, etc.) using NICE

Order Parameters You're Familiar With

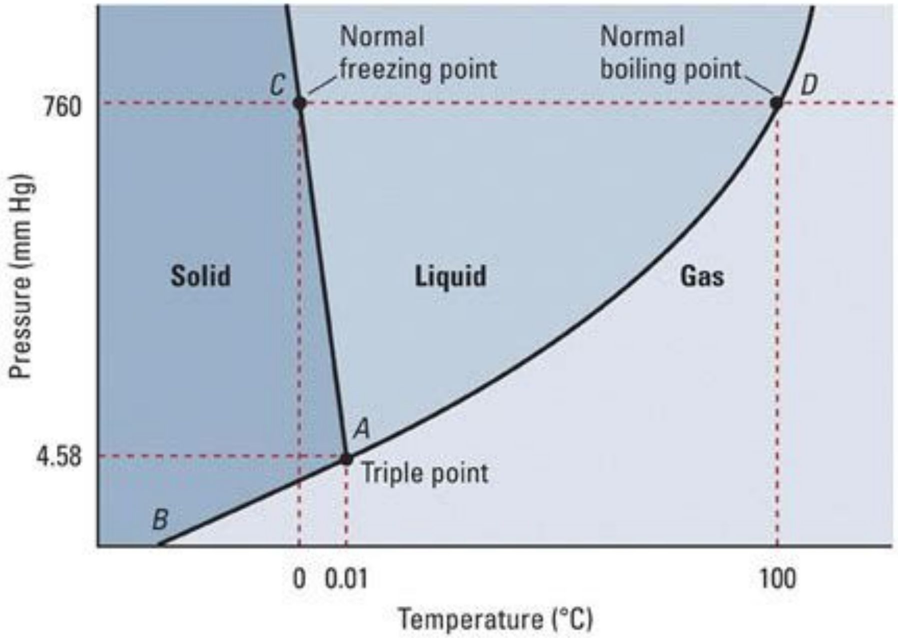

- Liquid-gas transitions use density (calculated using M, P, and T) as order parameter
- Distinguishes phases (difference in properties)
- Other parameters are useful for other materials

Magnetic Order Parameter

- Magnets use "orientation"
- The magnitude of magnetic (dipole) fields in the "normal" direction
- Phases are magnetic and nonmagnetic

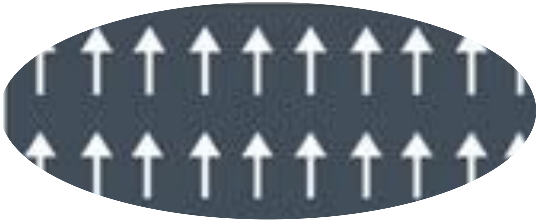

Ferromagnetic

Paramagnetic

Magnetic Order Parameter

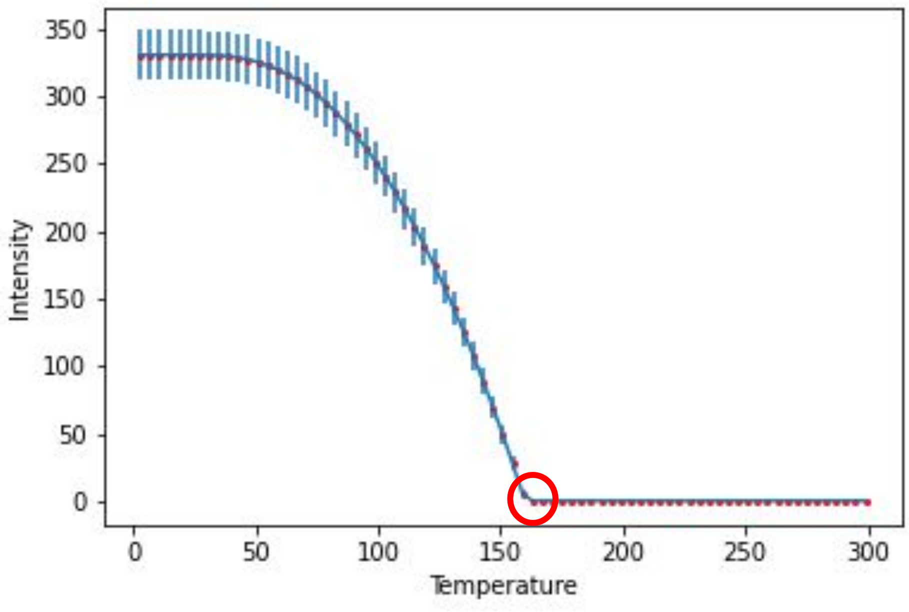

- Scattering intensity is proportional to magnetism
- Temperature where intensity becomes 0 is also where the material becomes weakly magnetic

The Measurement

- In **most** basic form, a curve fitting problem
- Solved by classical fitters (bumps)
- Goal is picking most informative point on curve given what has been measured so far, using constant increases in temperature

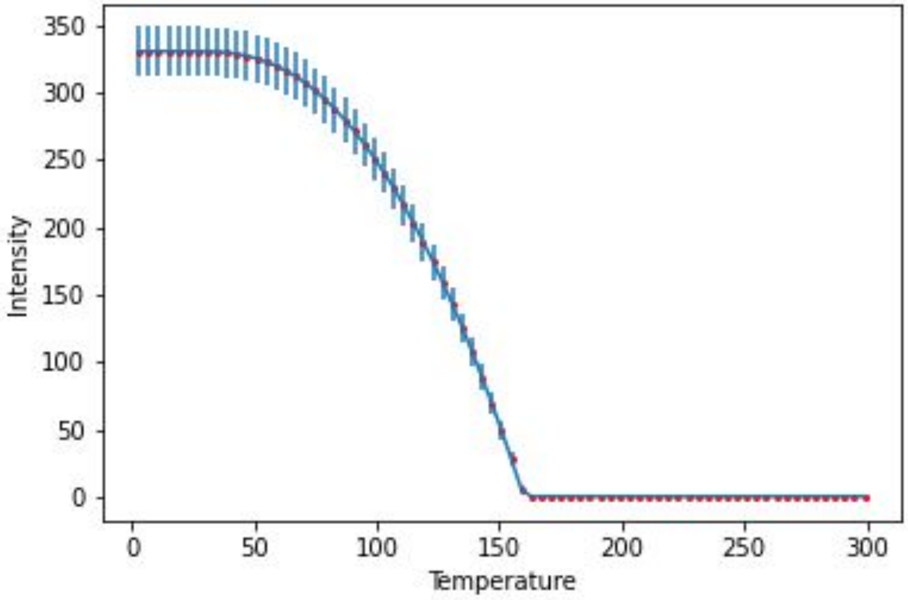

The Measurement

- In **most** basic form, a curve fitting problem
- Solved by classical fitters (bumps)
- Goal is picking most informative point on curve given what has been measured so far, using constant increases in temperature

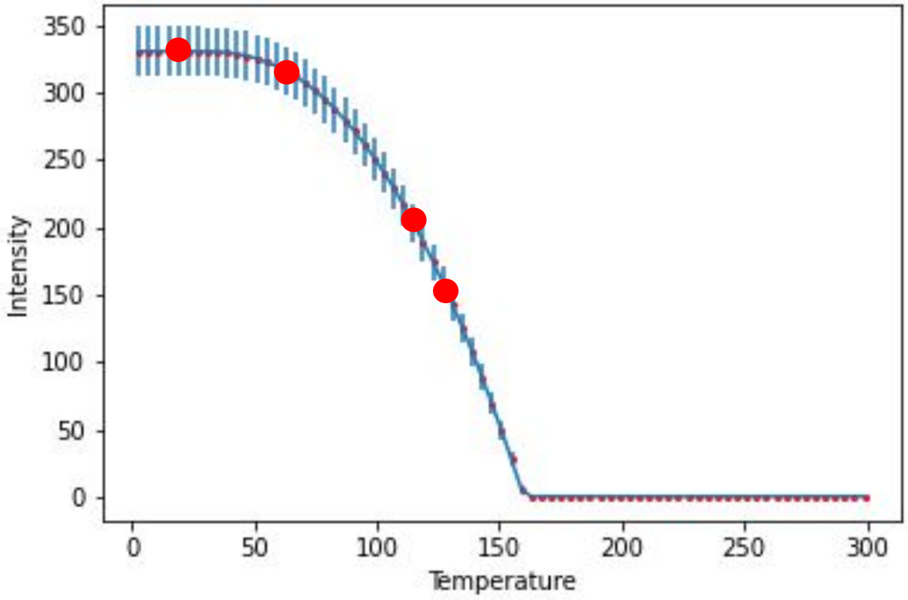

Reinforcement Learning

Defined:

Teaching a computer to make optimal decisions using rewards

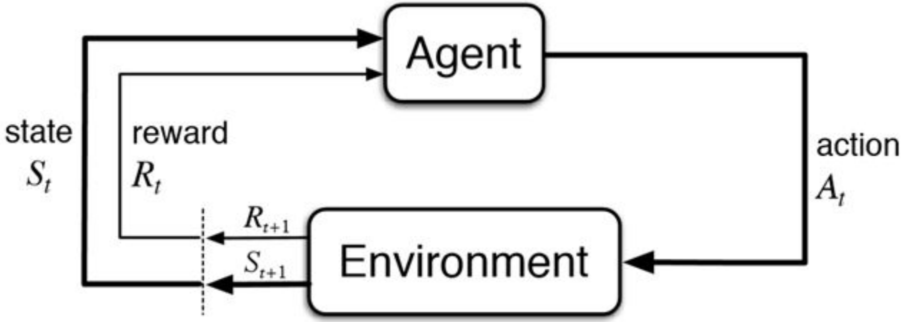

How does it work?

1. The agent is in an environment
2. The environment returns a state
3. Agent makes action based on state
4. Agent is rewarded after action
5. Algorithm learns how to best make actions based on rewards

Reinforcement Learning

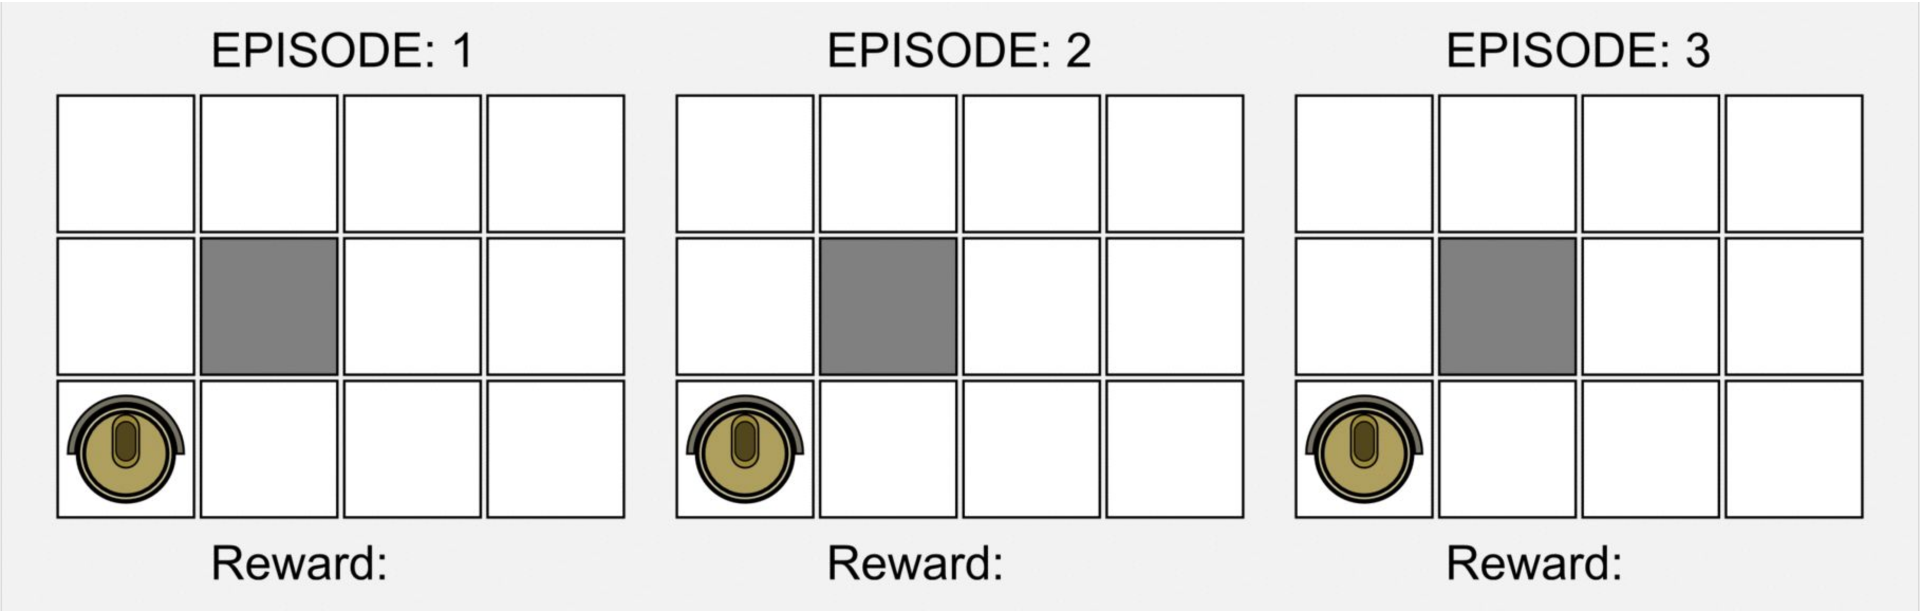

Applying Reinforcement Learning

- Action: positive change in temperature
- State: last measured temperature
- Reward: low chi squared
- Ends episode with low chi squared, too many steps, or temperature too high

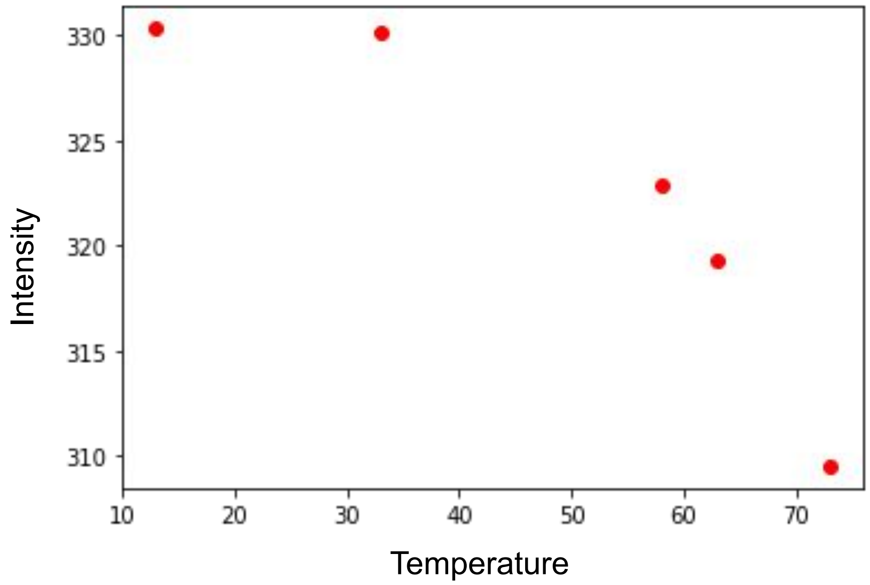

Other Parameters to Fit

- **Jt:** Magnetic moment
- **Nf:** Number of magnetic ions/volume (number density)
- **Bk:** Background intensity

Results: Stage One

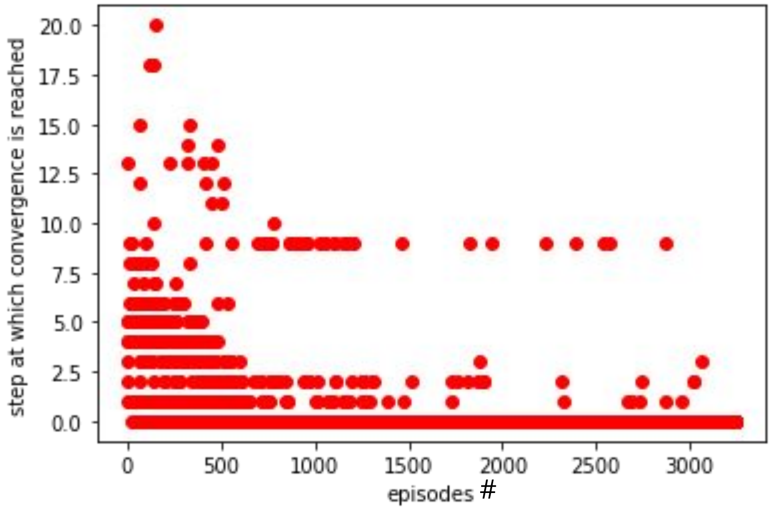

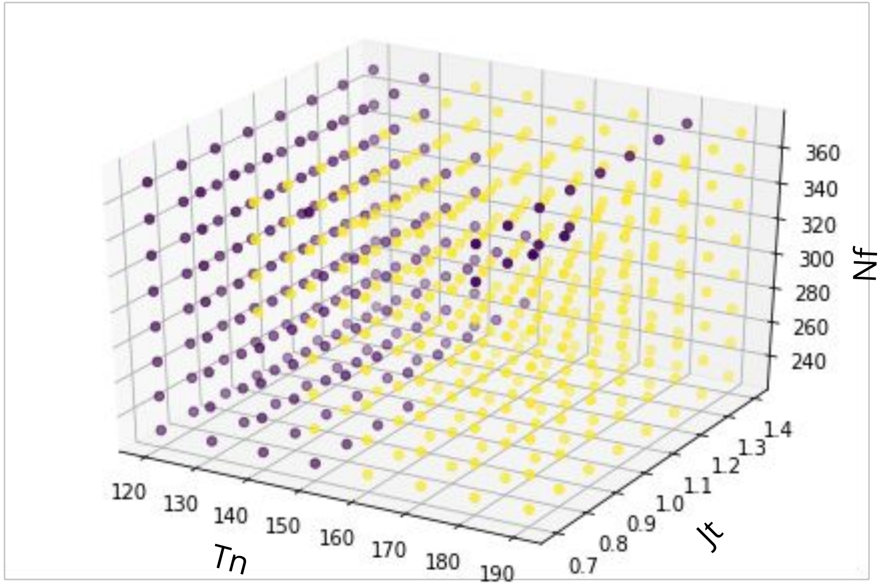

Algorithm improves with more episodes, takes less steps to fit the transition temperature

Model successfully finds transition temperatures (yellow) when experiment parameters are similar to training values

Results: Stage Two (failure)

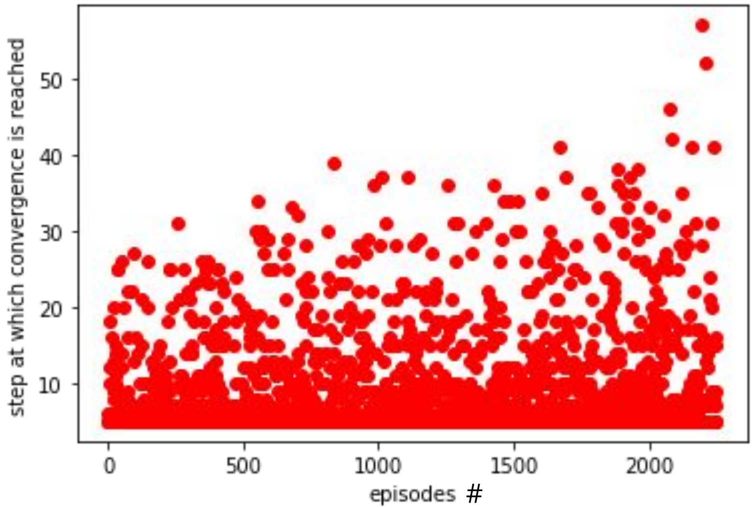

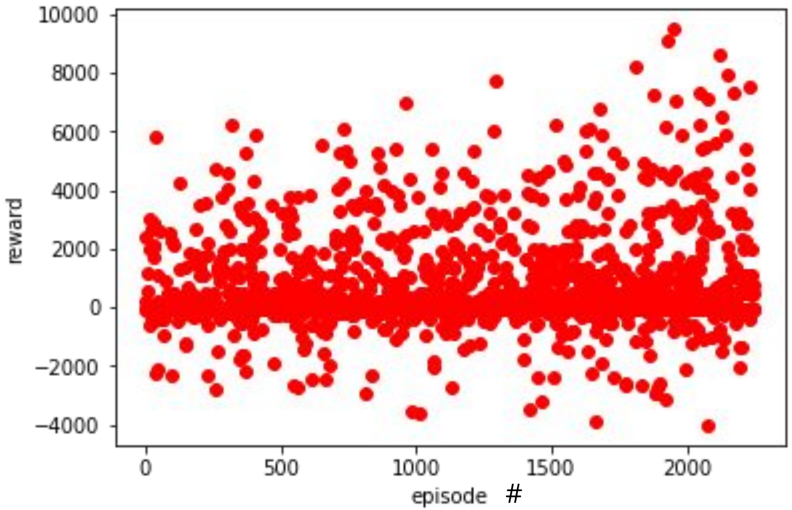

- Fitting on new set of four random variables each episode
- Better reflects application
- Worsening convergence, mediocre rewards

Results: Stage Two (failure)

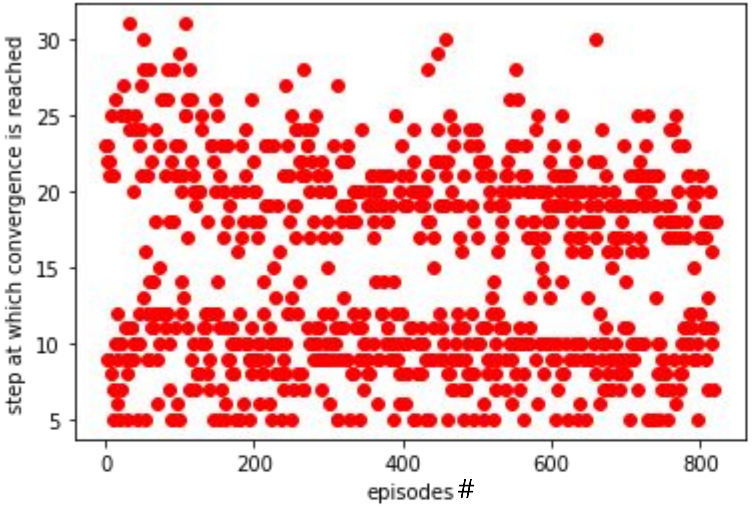

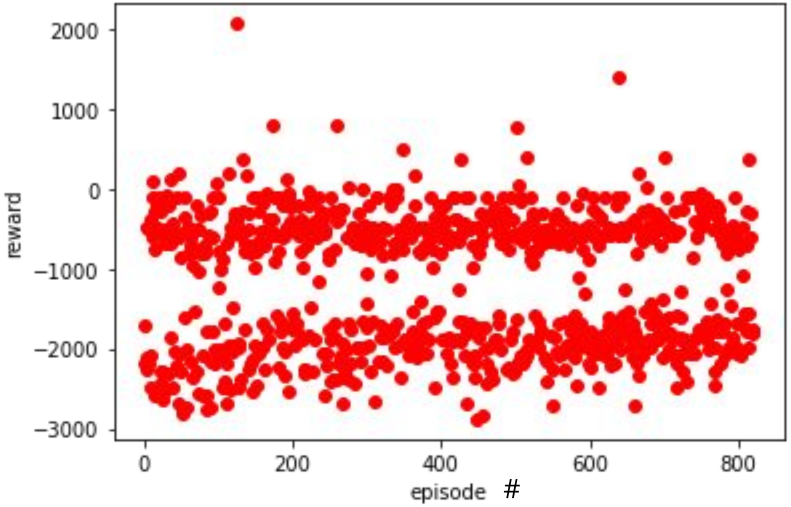

- Limit random variables to three
- Mildly better in trends, but bad preliminary results
- 30, 25 measurements is not ideal

(bimodality?)

- Distinct groups in reward and convergence graphs
- Represent successful fits and failed fits where algorithm went over max temperature
- Distinction between the two may be low transition temperature

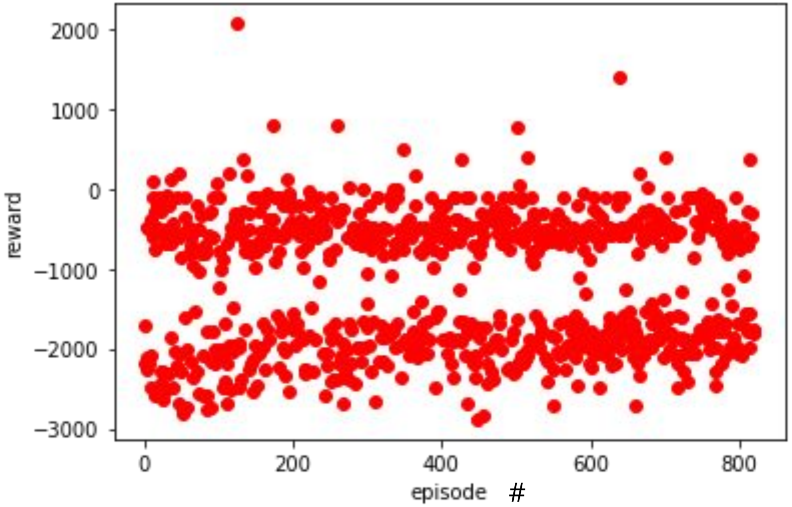

On Fitting

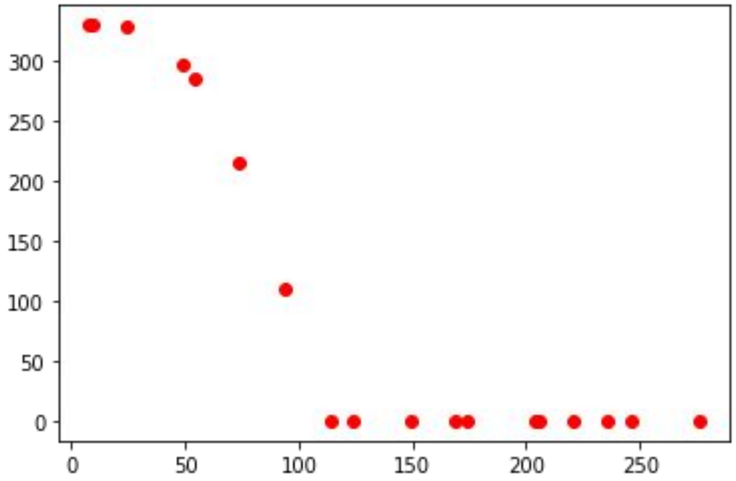

Tn: 123

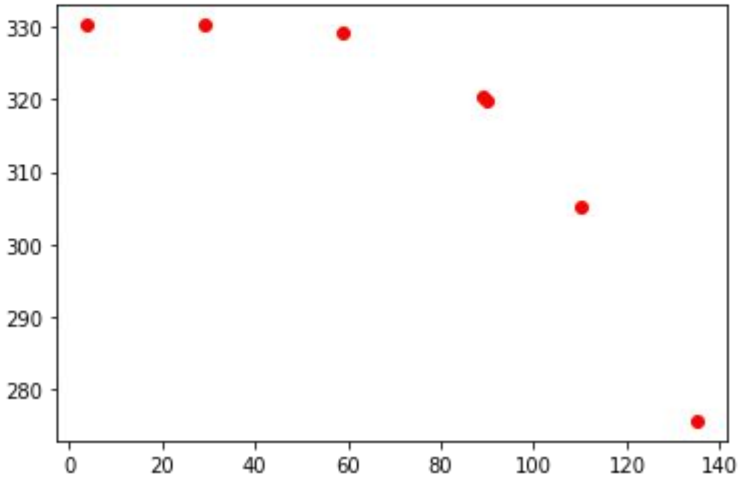

Tn: 240

Fit is much better before measuring past the transition temperature

Results: Stage Two (failure)

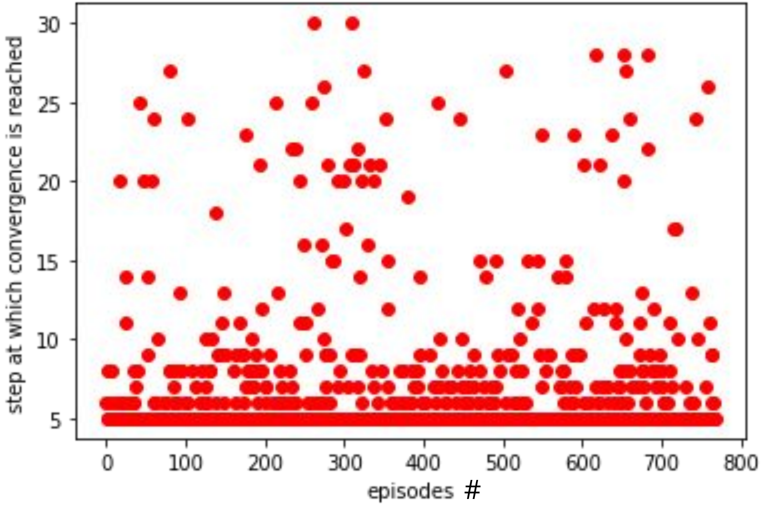

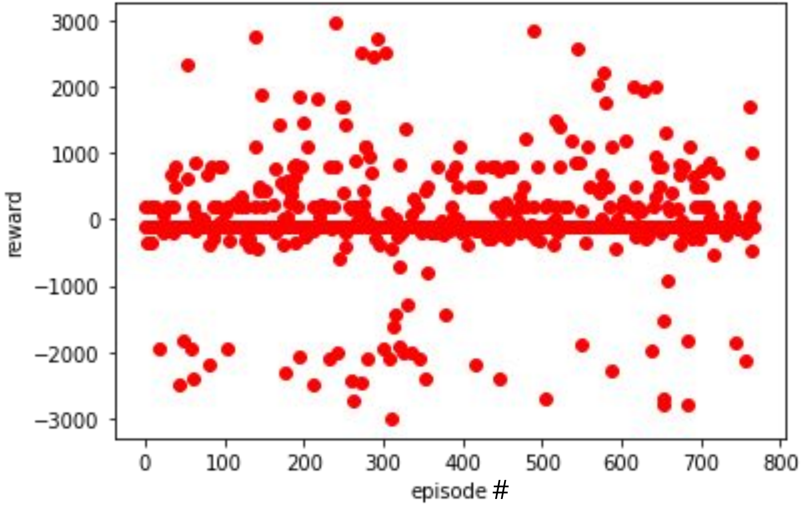

- Two variables!
- Does well consistently with fast convergence
- However, no improvement - success from fitter

Future Steps

- Improve convergence and reward trends through reward function, fitter usage
- Replace simulated data with real-world data
- Implement model in application software to work with NICE

Acknowledgements

Thanks to the research program directors, interns and mentors.

Special thanks to William Ratcliff, mentor.

We are grateful for funding from the Center for High Resolution Neutron Scattering.

Questions