

CS394R/ECE381V
Reinforcement Learning:
Theory and Practice

Amy Zhang and Peter Stone

Departments of ECE and CS
The University of Texas at Austin

Good Morning Colleagues

- Are there any (logistics) questions?

Logistics

- Next week's readings

Logistics

- Next week's readings
 - Multiagent Learning, Reproducibility, and Wrap-Up

Human-in-the-loop RL

- Recall the Week 1 activity

Human-in-the-loop RL

- Recall the Week 1 activity
 - Was it human-in-the-loop?

Human-in-the-loop RL

- Recall the Week 1 activity
 - Was it human-in-the-loop?
 - Technically, yes - you selected actions

Human-in-the-loop RL

- Recall the Week 1 activity
 - Was it human-in-the-loop?
 - Technically, yes - you selected actions
- What are some other ways to put a human in the loop?

Human-in-the-loop RL

- Evaluative feedback (RLHF)
 - Explicit
 - Implicit

Human-in-the-loop RL

- Evaluative feedback (RLHF)
 - Explicit
 - Implicit
- Interventions

Human-in-the-loop RL

- Evaluative feedback (RLHF)
 - Explicit
 - Implicit
- Interventions
- Demonstrations
 - State only (observations)
 - State-action
 - Behavior cloning
 - IRL

Human-in-the-loop RL

- Evaluative feedback (RLHF)
 - Explicit
 - Implicit
- Interventions
- Demonstrations
 - State only (observations)
 - State-action
 - Behavior cloning
 - IRL
- Advice (NLP)

Human-in-the-loop RL

