CS 394R Reinforcement Learning: Theory and Practice

Sanmit Narvekar

Department of Computer Science
University of Texas at Austin

Good Morning Colleagues

- Are there any questions?

Logistics

- First 3 assignments due **next Friday by midnight**
- Project proposals due **next Thursday 9:30am**

Outline

- Get a feel for how to approach a complex RL problem
- Keepaway Domain
- Group Activity!
- One approach and extensions

Keepaway: A Subtask of 2D Simulated Soccer

- Play in a **small area**
- **Keepers** try to keep the ball
- **Takers** try to get the ball

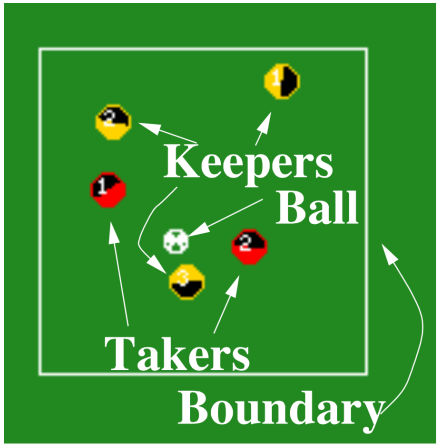

- **Episode**:
  - Players and ball reset randomly
  - Ball starts near a keeper
  - Ends when taker gets the ball or ball goes out

- Performance measure: **average possession duration**

Keepaway

- **Sensations**
  - Your position
  - Teammate and opponent positions
  - Ball position
  - Landmark positions

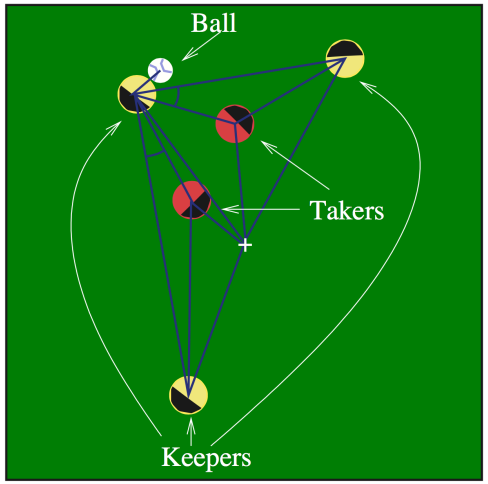

- **Raw Actions**
  - Move(x, y)
  - Kick(x, y, power)

- **Higher level** actions/skills
  - HoldBall, PassBall(k), GoToBall, GetOpen

Some questions to think about…

- **State/Action** space?
- **Reward** function?
- Function **approximation**?

- MDP formulation
  - Each agent learns separately?
  - Each agent shares a set of learned skills?

- What will be **learned**, and what will be **hand-coded**?
  - If it's a mix, what will the final policy look like?

- Keep the learning problem **tractable**!

One approach…

- **Reinforcement Learning for RoboCup-Soccer Keepaway**
  Peter Stone, Richard S. Sutton, and Gregory Kuhlmann. *Adaptive Behavior*, 2005.

- Shivaram Kalyanakrishnan and Peter Stone. **Learning Complementary Multiagent Behaviors: A Case Study**. *RoboCup 2009: Robot Soccer World Cup XIII*, pp. 153–165, Springer Verlag, 2010.

Extension: Transfer Learning

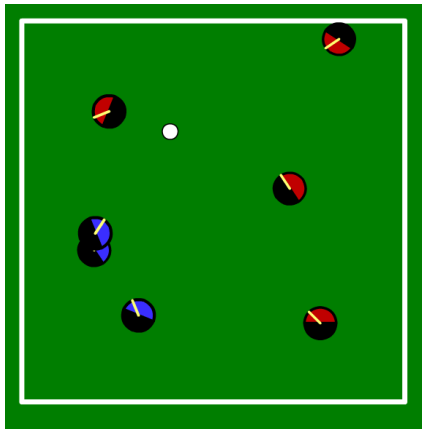

4v3

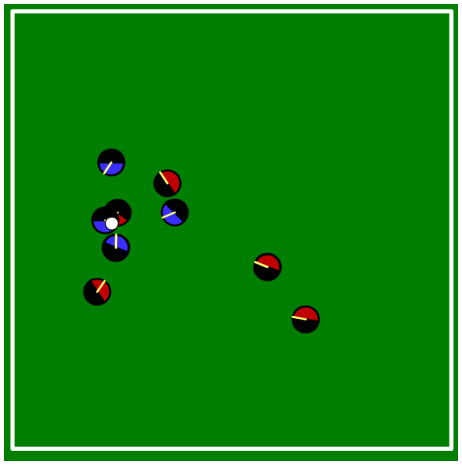

5v4