All The Other Things

Ron Parr
CSCI 2951-F

State Abstraction

- Markov property demands fine resolution state representation
- When/how can this be relaxed?
- What if our view of the world is coarse/approximate?
- Can we do coarse planning that is refined with search at run time?
- Example: Texas Hold 'Em

- Intersects: Function Approximation, Hierarchy, Model Learning

Transfer learning

- Apply solution from one task to a family of related tasks?
- Be careful to define this precisely (not just training on one task then haphazardly testing on another)
- Inter-task transfer:
  - Define distribution over tasks
  - Train on tasks sampled from distribution
  - Test on new tasks sampled from distribution
- Intra-task transfer:
  - Within the context of a single task subproblems are repeated
  - Examples: Opening doors, going up stairs, parking cars, etc.
- Intersects: Robustness, Skills, Hierarchy, Model Learning

Skills

- A skill is an ability that is primarily used to achieve other thing (not so interesting in isolation)
- Much of (practical) human education is about skills
- Examples: Balancing, walking, throwing, catching, driving
- Challenges:
  - How do we define skills in ways that make them broadly useful?
  - How do we integrate skills into larger problems
- Intersects: Hierarchy, Transfer

Hierarchy I

- Humans rarely solve detailed plans
- It seems we do some combination of:
  - High level (abstract) planning
  - Reuse of previous plans (transfer)
  - Online refinement of coarse plans (search?)
- Challenges:
  - Where do plan hierarchies come from?
  - How do we represent them?
  - How we manage the inaccuracies that arise from hierarchies
- Intersects: Transfer, Abstraction, Robustness, Skills, Model Learning, Interpretability, Model Learning

Robustness

- Models are often wrong
- Training regime may not align with use (sim to real issues)
- Robust solutions take into account these mismatches, and aim for solutions to work well in the worst case
- Challenges:
  - Efficient algorithms
  - Being (usefully and accurately) precise about imprecision
- Intersects: Hierarchy, State Abstraction, Skills, Transfer

Interpretability/Explainability

- Interpretability: Is the solution human-understandable?
- Explainability: Can a particular choice be justified?
- Increasing concern in (deep) ML:
  - Why should I trust this system? (interpretability)
  - Why should I trust this answer? (explainability)
  - Why it matters: Fairness, high stakes decisions
- Challenges:
  - Defining the problem precisely
  - Increasing use of complicated deep networks
- Intersects: Abstraction, Hierarchy, Model Learning, Structured Representations

Model learning

- Models are hard to learn, but have advantages:
  - Possibility of greater data efficiency
  - Use in transfer, e.g., if only reward function changes
  - Might help with interpretability/explainability
  - Might help with exploration
  - Might help with sparse reward learning (by providing a useful learning target before the agent even sees the first reward)
- Challenges:
  - Models are complicated! (Learn a distribution?)
  - Can be computationally challenging (e.g., in Atari is model images->images?)
- Intersects: Hierarchy, Interpretability, Robustness, Transfer, Model Learning

Structured Representations

- We understand the world through propositions, relations, facts, etc.
- We know the world is composed of objects that interact with each other in somewhat predictable ways
- Most recent RL successes don't take advantage of this
- Challenges:
  - How to incorporate prior knowledge in structured representations?
  - How to learn/discover solutions that have this form
- Intersects: Interpretability

Theoretical Questions

- Huge gap between best deep RL results and theory
  (True for deep learning in general, but more so for RL)
- Challenges:
  - Explain why deep RL works
  - Make deep RL more reliable (more science, less art)
  - Variance in deep RL experiments
- Intersects: Everything

Constrained/Risk Aware/Safe RL

- Research is heavily tilted towards discounted sum of reward

- Doesn't take into account:
  - Hard constraints on allowable behavior (don't run over children!)
  - Qualitative constraints on behavior (e.g. quality of service)
  - Risk tolerance (it's not just about the average)

- Challenges:
  - Finding the right problem formulation
  - Incorporating these concepts while maintaining efficiency

Intersections with Control Theory

- Control theory traditionally focuses on:
  - Continuous actions
  - Relatively simple physics and noise models
  - Relatively low noise problems
- RL traditionally has focused on:
  - Discrete actions
  - Arbitrary noise and transition models

- These boundaries are not crisp – both fields are moving towards and learning from each other

Connection to Neuroscience/ behavior

- RL originally motivated by psychology
- Drifted away from this for many years
- More recently:
  - Temporal differences connected to dopamine in the brain
  - Some excitement in neuroscience community about this
- Questions:
  - How far can this be pushed?
  - Does it have practical significance?

Data Efficiency

- One of the biggest challenges in RL is data efficiency
- Why it's bad:
  - Can't be applied in the real world
  - Huge resource consumption
- What aggravates this:
  - High variance
  - High network complexity
- What might help:
  - Model learning?
  - Different representations?
  - Better algorithms?
  - More use of hierarchy/transfer?

We’re done!