

Intrinsically Motivated Reinforcement Learning

T Sudhamsh Goutham Nitish Srivastava

CS 365: Artificial Intelligence
Department of Computer Science and Engineering, IIT Kanpur

March 8, 2010

Outline

- ▶ Reinforcement Learning
- ▶ Intrinsic Motivation
- ▶ Experiment Playroom
- ▶ Methodology

Reinforcement Learning

- ▶ A class of ML algorithms
- ▶ Learns how to act given an observation of the world
- ▶ Action changes the environment
- ▶ Feedback in terms of Rewards
- ▶ Maximize a long-term Reward

Reinforcement Learning


- ▶ Formulated as Markov Decision Processes
- ▶ Related to Dynamic Programming Algorithms

Intrinsic Motivation


- ▶ For its own sake
- ▶ Not as steps of a bigger problem
- ▶ No external reward Doing for the sake of happiness or intrinsic enjoyment.
- ▶ Act, Play, Explore for development of broad competence

Intrinsically Motivated Reinforcement Learning

Standard RL view



Intrinsically Motivated RL view



Agent Description

Agent has:

- ▶ Eye
- ▶ Hand
- ▶ Visual Marker

Agent can:

- ▶ Move eye to hand
- ▶ Move eye to marker
- ▶ Move eye to random object
- ▶ move hand to eye
- ▶ move hand to marker
- ▶ move marker to eye
- ▶ move marker to hand
- ▶ if both hand and eye are on same object: use the object

Modeling intrinsic reward

- ▶ Modeling novelty response of dopamine neurons
- ▶ Intrinsic reward for salient event proportional to the error in its prediction according to the learned option model for that event
- ▶ The intrinsic reward is used to update the value function the agent is using to determine its behavior in the playroom.
- ▶ Agent tries to repeatedly achieve the salient event. RL improves both
 - ▶ policy for achieving the event
 - ▶ option-model that predicts the event
- ▶ As option policy and option model improve, the intrinsic reward diminishes and the agent gets “bored with the associated salient event and moves on.

References

- ▶ Intrinsically Motivated Reinforcement Learning by Satinder Singh, Andrew G. Barto and Nuttapon Chentanez. Proceedings of Advances in Neural Information Processing Systems (NIPS), 2005.
- ▶ Intrinsically Motivated Learning of Hierarchical Collections of Skills by Andrew G. Barto, Satinder Singh, and Nuttapon Chentanez. Proceedings of International Conference on Developmental Learning (ICDL), 2004.
- ▶ Emotion-Driven Reinforcement Learning, Marinier, R. and Laird, J. E. CogSci 2008
- ▶ Simple Algorithmic Theory of Subjective Beauty, Novelty, Surprise, Interestingness, Attention, Curiosity, Creativity, Art, Science, Music, Jokes. J. Schmidhuber. Journal of SICE, 48(1):21-32, 2009.