Reinforcement Learning and Beyond
Part II: Transfer Learning in RL

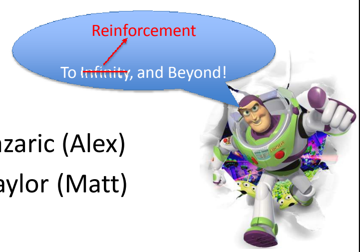

Dr. Alessandro Lazaric (Alex)
Dr. Matthew E. Taylor (Matt)

Objectives

- Classification of approaches
- TL difficulties in RL domains
  - When: definition and discussion of different transfer problems in RL domains
  - What: overview of popular approaches
  - How: augmenting RL algorithms with TL
- Future research directions

sequel.futurs.inria.fr/lazaric

teamcore.usc.edu/taylorm

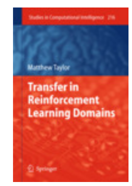

Outline

1. Transfer in AI and Machine Learning
   - Matt
2. Transfer in Reinforcement Learning
   - Alex, Matt
3. Conclusions
   - Alex

ALA Workshop

Sutton and Barto

L. P. Kaelbling, M. L. Littman, and A. W. Moore. Reinforcement learning: A survey. JAIR. 1996.

Reinforcement Learning and Beyond
Part II: Transfer Learning in RL

Section 1: Transfer in AI and Machine Learning

Section Outline

- Historical Perspective
- An overview of transfer in Machine Learning
- Challenges and Goals of TL
- How TL in RL relates to TL in general

Psychology, Education

- B.F. Skinner
  - "Radical Behaviorism"
  - Schedules of reinforcement (+ or -)
    - from continuous reinforcement to extinction
- Instructional Scaffolding
  - late 50's: language acquisition
  - Sequence of well-timed tasks guide learning

The Theory of Transfer

"Transfer of learning occurs when learning in one context enhances (positive transfer) or undermines (negative transfer) a related performance in another context."

(D. Perkins, G. Salomon, Transfer of Learning, 1992, International Encyclopedia of Education)

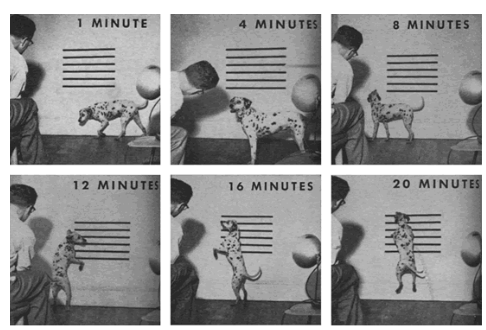

Motivations for TL

- Learning *tabula rasa* can be extremely slow
  - Lots of data / time may be needed
  - Every algorithm has biases: why use an uninformed bias?
- Humans always use past knowledge
  - What knowledge is relevant?
  - How can it be effectively leveraged?

Example: Sebastian Thrun

- Explanation-Based Neural Network Learning: A Lifelong Learning Approach
  - PhD thesis, 1995
- Is Learning the nth Thing any Easier than Learning the First?
  - NIPS, 1996
- Learning to Learn
  - Edited volume (with Lorien Pratt), 1998

Example: Rich Caruana

- Multitask Learning: A Knowledge-Based Source of Inductive Bias
  - ICML 1993
- Learning Many Related Tasks at the Same Time with Backpropagation
  - NIPS 1995
- Algorithms and Applications for Multitask Learning
  - ICML 1996
- Multitask Learning
  - PhD thesis, 1997

Example: Matt's & Alex's work

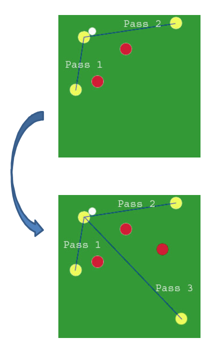

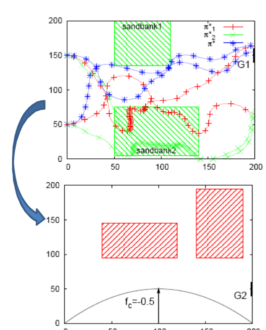

Towards Transfer Learning

- ML-COLT '94 Workshop on Constructive Induction and Change of Representation
  - Towards autonomy
  - Generate / modify representations automatically
- NIPS '95 Workshop on Learning to Learn: Knowledge Consolidation and Transfer in Inductive Systems
  - Capitalize on previously acquired domain knowledge

Challenges and Goals

"Transfer [learning] is a sequential process that influences the performance of learning through the reuse of structured knowledge [collected on previous tasks] and improves the behavior of the agent on new related tasks."

Pat Langley

(Workshop on Structural Knowledge Transfer for Machine Learning, ICML 2006)

More recent TL workshops

- Inductive Transfer: 10 Years Later
  - NIPS 2005
- Structural Knowledge Transfer for Machine Learning
  - ICML 2006
- Transfer Learning for Complex Tasks
  - AAAI 2008

Challenges

- Structured Knowledge
  - Definition
  - Collection
  - Reuse
- Transfer process
- Task-independent Metrics
- Task Relatedness
- Negative Transfer

An Overview of TL in ML

- Many names
  - Learning to Learn
  - Meta-learning
  - Lifelong Learning
  - Continual Learning
  - Multi-task Learning
  - Inductive Transfer Learning

Hierarchical Bayes

- All the tasks are generated according to a fixed distribution
- Define a hyper-distribution over the task distribution
- Compute the distribution parameters according to the samples collected over all the tasks

An Overview of TL in ML

- Techniques
  - Hierarchical Bayes
  - Regularized Regression
  - Neural Networks
  - Graph Integration

Hierarchical Bayes

- Multi-Task Gaussian Processes
  - Linear functions $f_t(x) = w_t^T x$
  - Task distribution $w_t \sim \mathcal{N}(\mu_w, C_w)$
  - Hyper-prior $(\mu_w, C_w) \sim \mathcal{N}\left(0, \frac{1}{\pi}\right) \mathcal{IW}(\tau, I)$
  - Inference problem: given *m* samples from *n* tasks
    - Posterior $\widehat{w}_t, \widehat{C}_t$
    - Parameters μ_w, C_w
  - Given the parameters it can be used also to improve the performance on new tasks

Hierarchical Bayes

- MTL with Dirichlet Process (DP) priors
  - Tasks often are not homogeneous and belong to different classes
  - Dirichlet Process automatically clusters tasks into classes
  - Define hyper-priors over the DP parameters
  - Use all the samples to refine the DP parameters
  - Given the parameters it can be used also to improve the performance on new tasks

Regularized Regression

- Linear Multi-Task Regularization
  - Independent learning $\sum_{t=1}^{T} ||w_t||^2$
  - Task clustering (Evgeniou, 2005)
    $$\sum_{c=1}^{C} \sum_{t=1}^{T_c} ||w_t - w_{0c}||^2 + ||w_{0c}||^2$$
  - Graph regularization (Evgeniou, 2005)
    $$\sum_{t,s} ||w_t - w_s||^2 A_{ts}$$
  - Common feature space (Argyriou, 2007)
    $$\sum_{i}^{d} ||w^i||_2^2$$

Regularized Regression

- Single-task Regularized Regression
  $$\sum_{i=1}^{n} loss(y_i, f(x_i)) + \lambda ||f||_{STL}$$
  where $|| \cdot ||_{STL}$ is a suitable norm (e.g., L2 for linear regression)
- Multi-task Regularized Regression
  $$\sum_{t=1}^{T} \sum_{i=1}^{n} loss(y_{i,t}, f_t(x_{i,t})) + \lambda ||\boldsymbol{f}||_{MTL}$$

  where $|| \cdot ||_{MTL}$ forces the tasks to be *similar*

Related Paradigms

- Lifelong Learning
  - Less clear task boundaries (spatial / temporal)
  - Prepare for anything
- Imitation / Demonstration Learning
  - Watching a similar agent or human
- Direct Human Advice
  - Action suggestion
  - Direct knowledge injection
- Shaping
  - Reward function
  - Human modifies reward function over time

Goals

- Improve performance over non-transfer learning
  - Sample Complexity
  - Jumpstart
  - Learning speed
  - Final performance
  - Asymptotic Performance

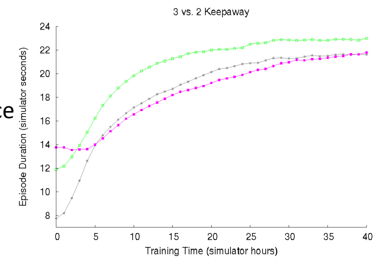

Reinforcement Learning and Beyond Part II: Transfer Learning in RL

Section 2: Transfer in Reinforcement Learning

How TL in RL relates to TL in general

- RL is the most general learning paradigm
  - Approximation
  - Exploration/exploitation
- More challenging than in supervised
  - Large number of scenarios
  - Many different types of knowledge can be transferred
  - Difficult to asses the contribution of transfer to the learning performance
- Many possible goals
- Many possible applications