# 1 Introduction (Outline)

**Planning – what is it?**

- Planning - Finding good/optimal moves is large, multistage problems.
- Online planning - doing the same, in limited time per step.
- Approximate planning - a necessity in large problems.

**Reinforcement Learning – what is it?**

- Learning from experience.
- Learning to act optimally by interacting with the environment.
- Trial and error: Experiment with different options, observe outcomes (rewards), and choose best options.

**Example problems:**

- chess, backgammon
- task scheduling
- packet routing
- robot juggling

**Types of learning:**

- Supervised
- Unsupervised
- Reinforcement Learning

**Complicating factors and issues:**

- Dynamic (state-dependent) environments
- Stochastic dynamics and rewards
- Unknown dynamics and rewards
- Exploration vs. exploitation
- Delayed rewards: temporal credit assignment
- Complicated systems: large state space, unstructured dynamics

**Basic solution tools:**

- Dynamic Programming [→ efficient planning]
- Stochastic Approximation algorithms [→ learning from random examples]
- Functional approximation (Neural Networks, etc.) [→ generalization and compact representation]

**RL in context:**

Reinforcement Learning has its origins in several related disciplines:

- Artificial Intelligence (sequential planning)
- Optimal Control (stochastic dynamic programming)
- Machine Learning (supervised learning techniques)
- Statistics, Stochastic Algorithms (bandit problems, stochastic approximations)
- Cognitive Science and Psychology

For a general overview of Machine Learning see, for example:
– Tom Mitchell, *Machine Learning,* McGraw-Hill, 1997
– Duda, Hart & Stork, *Pattern Classification,* Wiley, 2000 (2nd ed.)
– C. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2007

For a general overview of AI see:
– Russel and Norvig, *Artificial Intelligence,* Prentice Hall, 2009 (3rd ed.).