COMP 532
Machine Learning and BioInspired Optimization

Lecture 16: Multi-Agent Learning

Dr. Shan Luo
Department of Computer Science
shan.luo@liverpool.ac.uk

Introduction

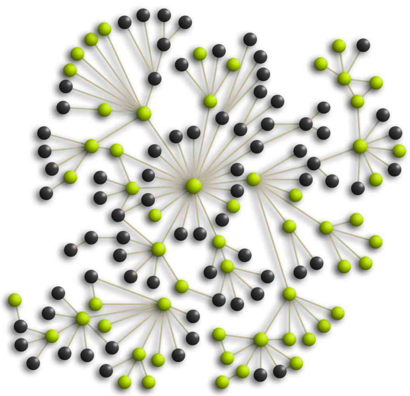

Multi-Agent Systems

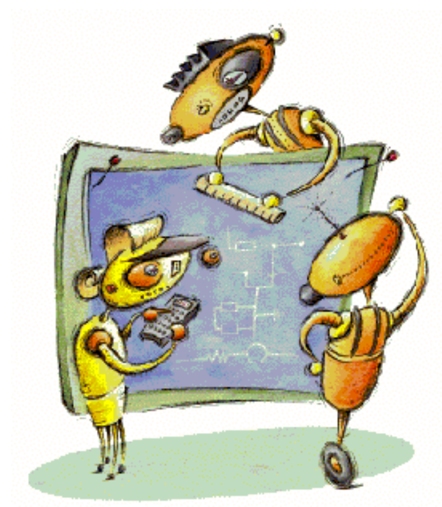

Introduction

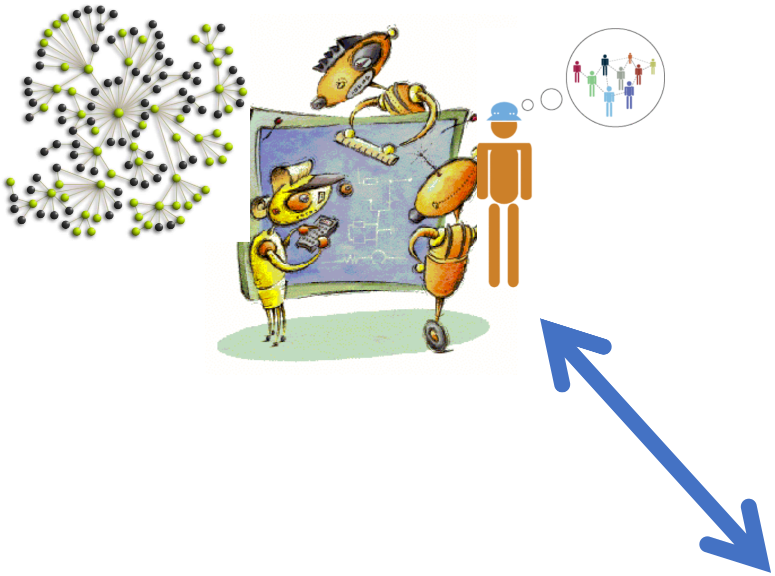

Reinforcement Learning

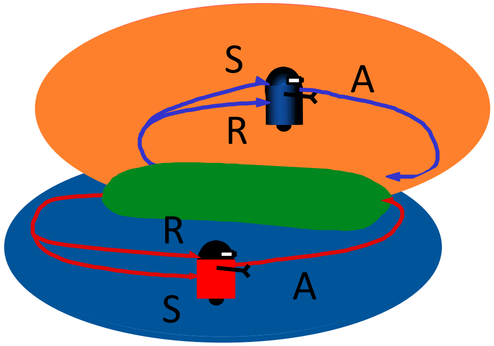

Introduction

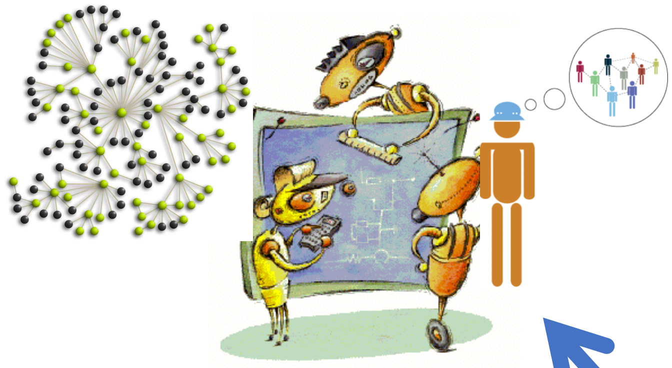

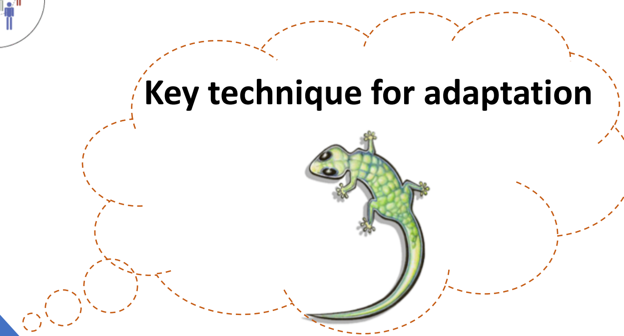

Reinforcement Learning

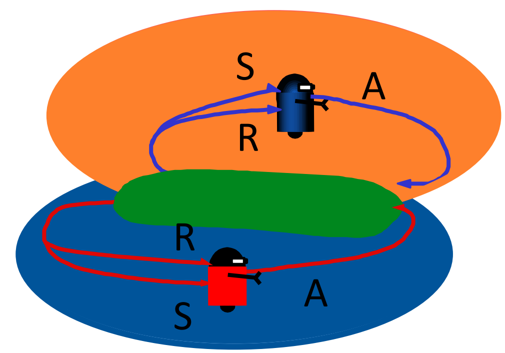

Introduction

Classical Game Theory

- CENTRAL CONCEPT:
  Nash equilibrium
- Normative theory
- Rational players

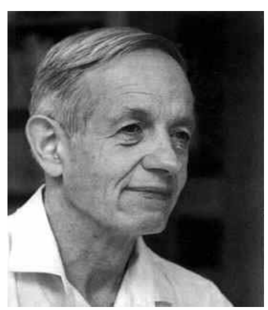

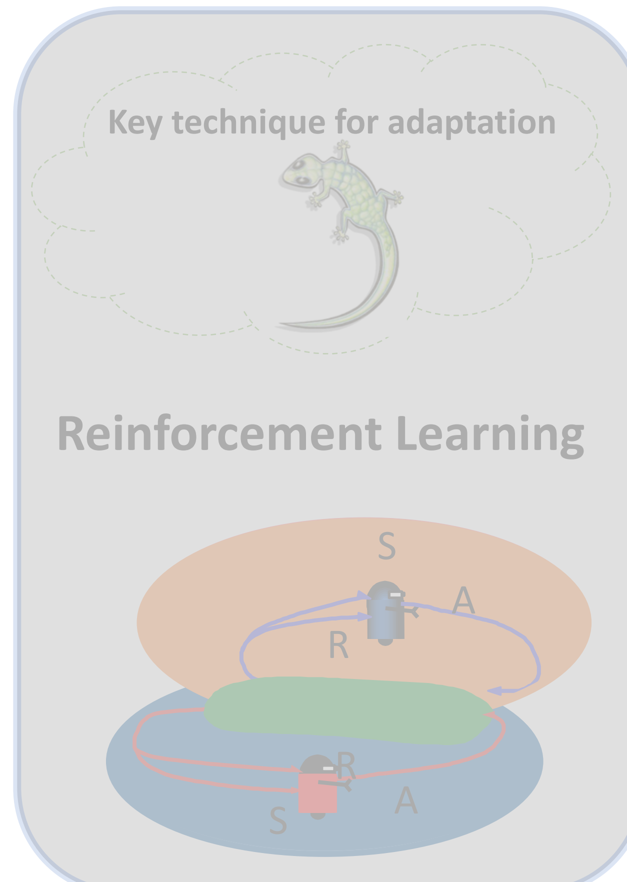

Introduction

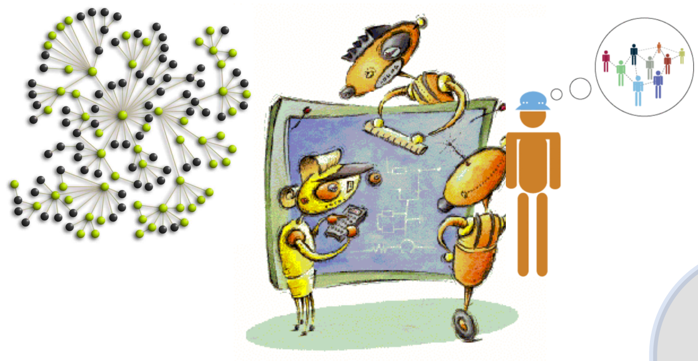

Strategic decision making

Classical Game Theory

- CENTRAL CONCEPT:
  Nash equilibrium
- Normative theory
- Rational players

Key technique for adaptation

Reinforcement Learning

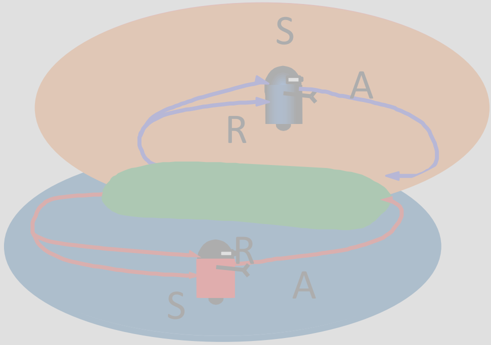

Introduction

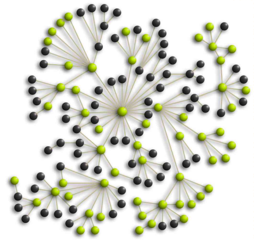

Strategic decision making

Classical Game Theory

- CENTRAL CONCEPT:
  Nash equilibrium
- Normative theory
- Rational players

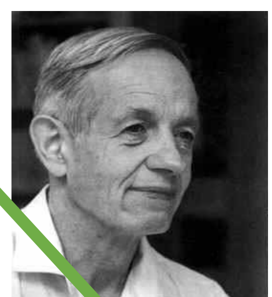

Key technique for adaptation

Reinforcement Learning

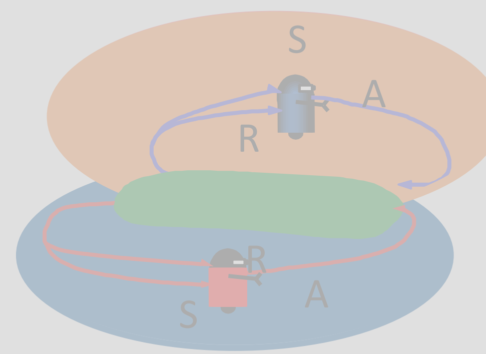

Introduction

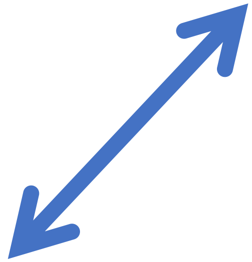

Evolutionary Game Theory

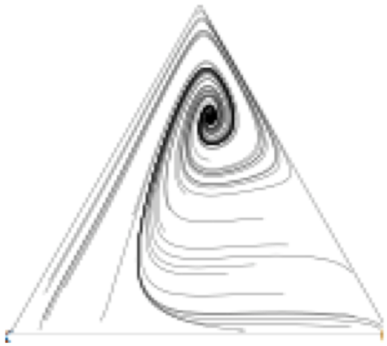

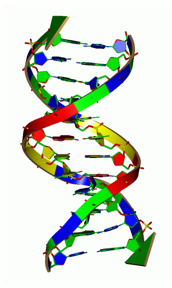

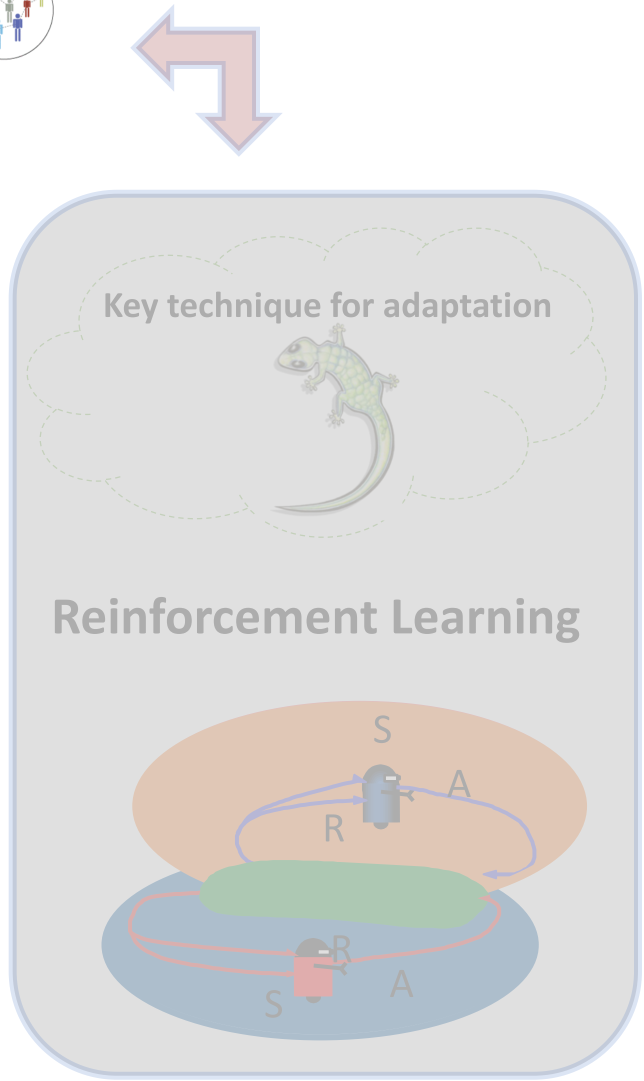

Introduction

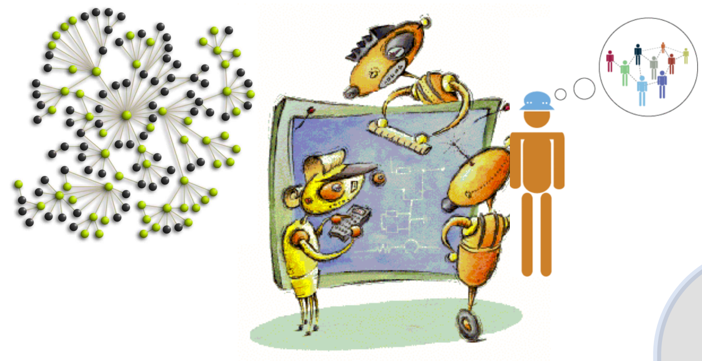

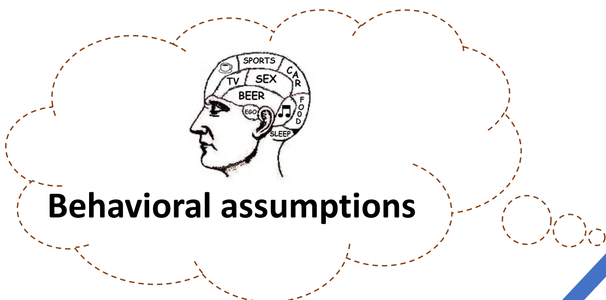

Evolutionary Game Theory

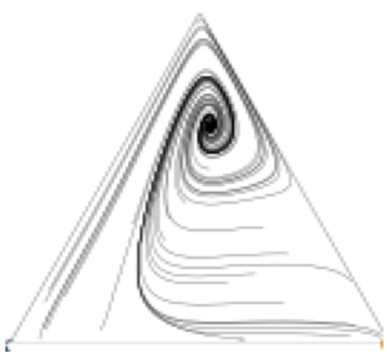

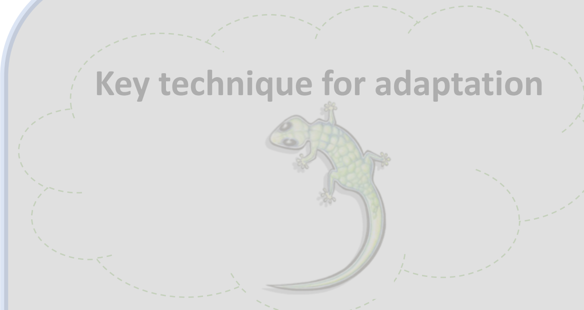

Reinforcement Learning

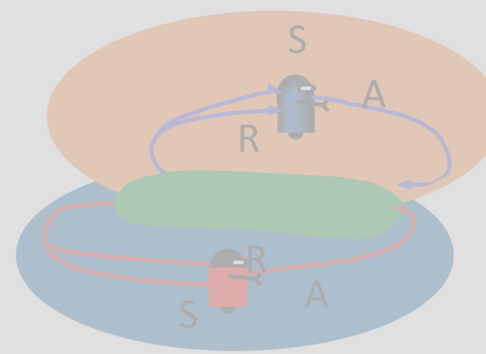

Introduction

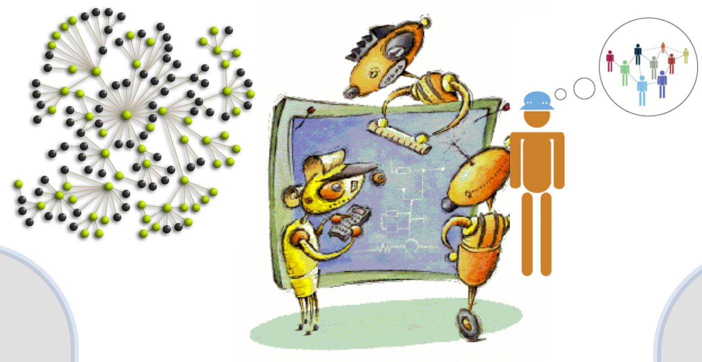

Backbone

$$\frac{dp_i}{dt} = p_i\left[e_i A q - pAq\right]$$

$$\frac{dq_i}{dt} = q_i\left[pBe_i - qBp\right]$$

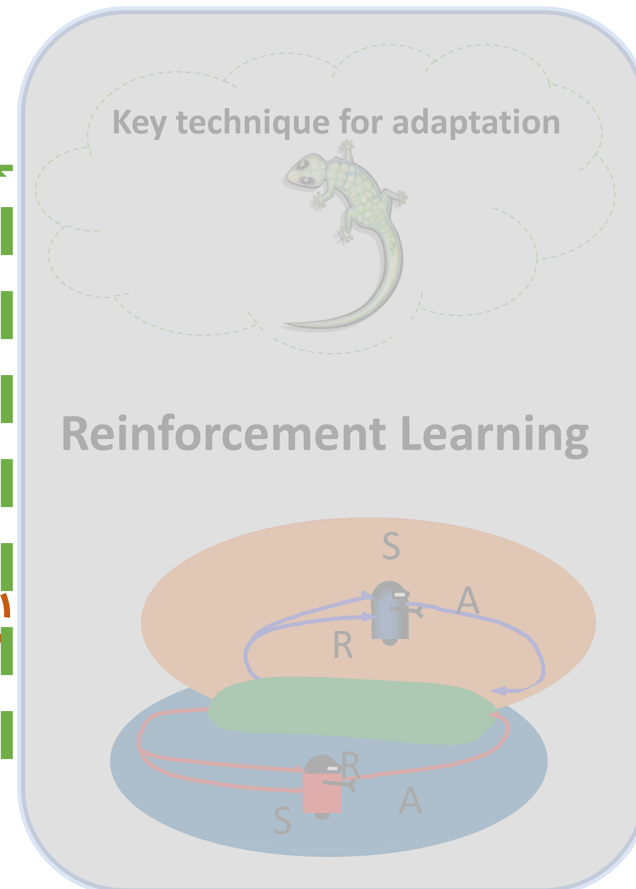

Introduction

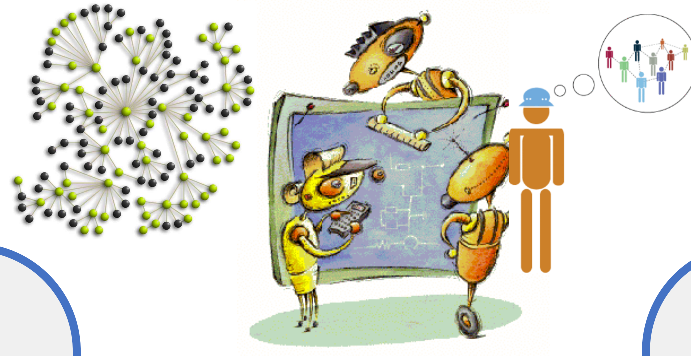

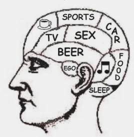

Behavioral assumptions

Evolutionary Game Theory

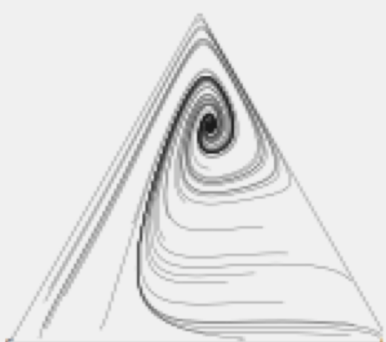

Key technique for adaptation

Reinforcement Learning

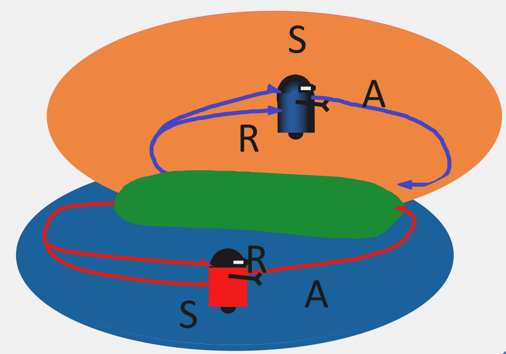

The Problem

- All theoretical properties of single agent adaptive learning are lost once you move to multiple agents

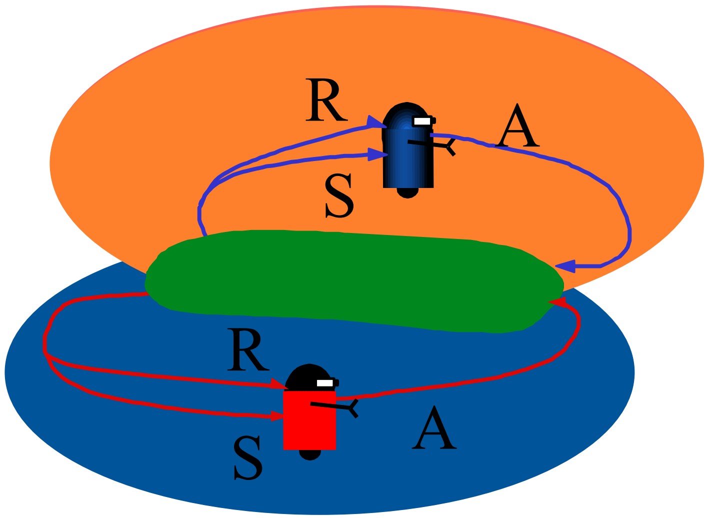

The Problem

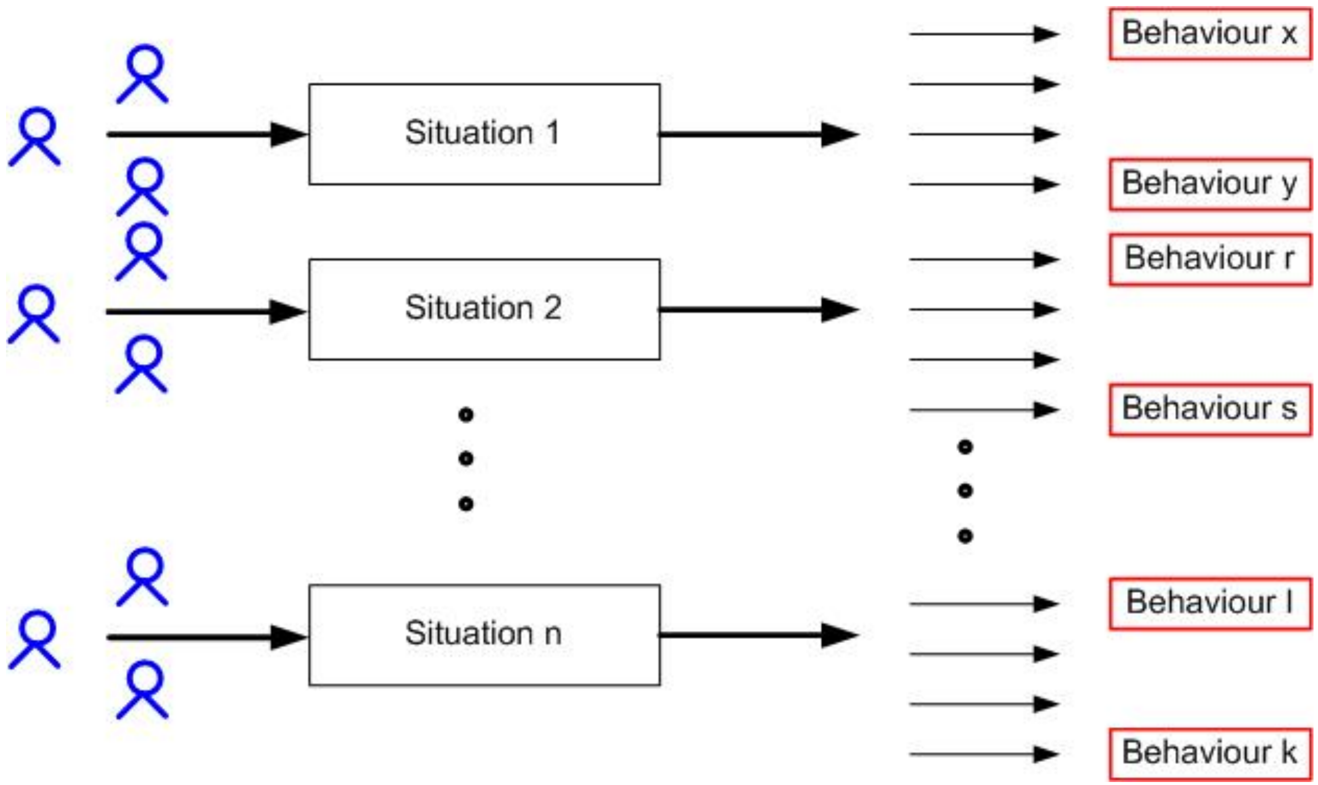

Situation 1
Situation 2
Situation n
Behaviour x
Behaviour y
Behaviour r
Behaviour s
Behaviour l
Behaviour k

Definition of Multi-Agent Learning

Stone and Tuyls:

The study of multiagent systems in which one or more of the autonomous entities improves automatically through experience

Learning: states and agents

	Single state	Many states
Single agent	Multi-Armed Bandit	Markov Decision Process
Many agents	Normal-Form Game	Markov Game

Learning: states and agents

	Single state	Many states
Single agent	Learning Automaton	Q-learning
Many agents	Evolutionary Game Theory	Multi-Agent Learning

Learning: states and agents

	Single state	Many states
Single agent	Learning Automaton	Q-learning
Many agents	Evolutionary Game Theory	Multi-Agent Learning

Many agents: new problems!

Why many agents?

Multi-agent systems have many applications, e.g.

- Teamwork (sports, planet exploration, sensor nets)
- Scheduling (job shops)
- Trading (auctions)
- Simulations (military, economical)

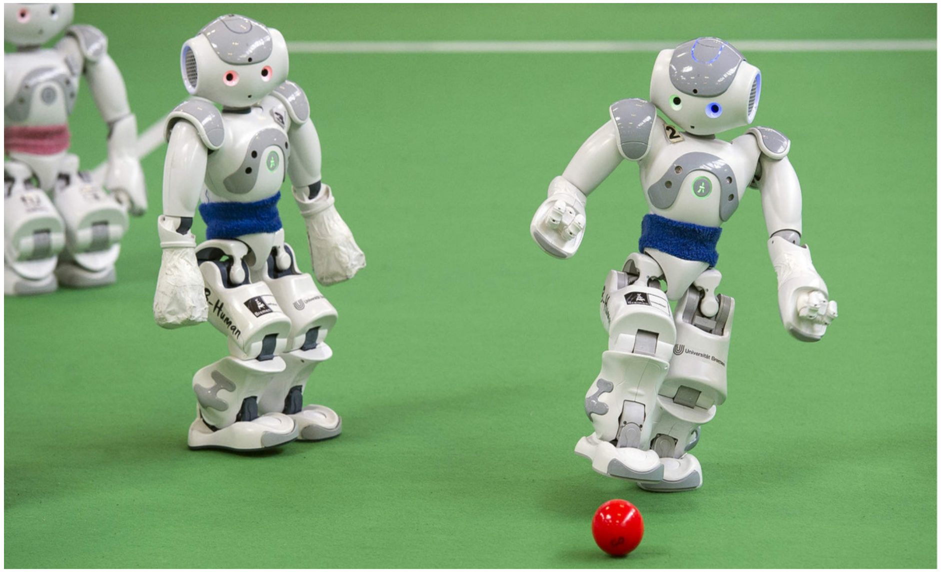

Why many agents?

New problems with many agents

- Independent from the application:
  - Many applications require a decentralized and adaptive (learning) system
  - A centralized solution is computationally hard
  - Environments are dynamic / unpredictable

- Dependent on the application:
  - Agents may need to cooperate on a common task, or coordinate their actions for the best result, or they may be (partially) competitive
  - What is it that individual agents need to do?

Common interest, conflicting interest

Which outcome do we want?

Key issues

- Multiple agents need to learn to find optimal solutions by acting autonomously
  - What are optimal solutions in this case?
  - How do we learn those?

- Multi-agent systems do not have the Markov property, because information concerning other agents is generally missing

- Even in fully cooperative systems, communication is expensive &/ unreliable

A good starting point: Game Theory

- Interactions between agents may be modelled using game theory

- Game theory formalizes these interactions
  - Each agent has a set of actions
  - A strategy gives a probability to each action
  - Joint actions lead to a payoff to each agent

- Given fully rational players with full information on the game, Game Theory can predict the strategies each agent will use

- Central concept: **Nash equilibrium**

Several Multi-Agent Learning Paradigms

- Online RL towards individual utility
- Online RL towards social welfare
- Co-evolutionary learning
- Swarm Intelligence
- Adaptive mechanism design

Several Multi-Agent Learning Paradigms

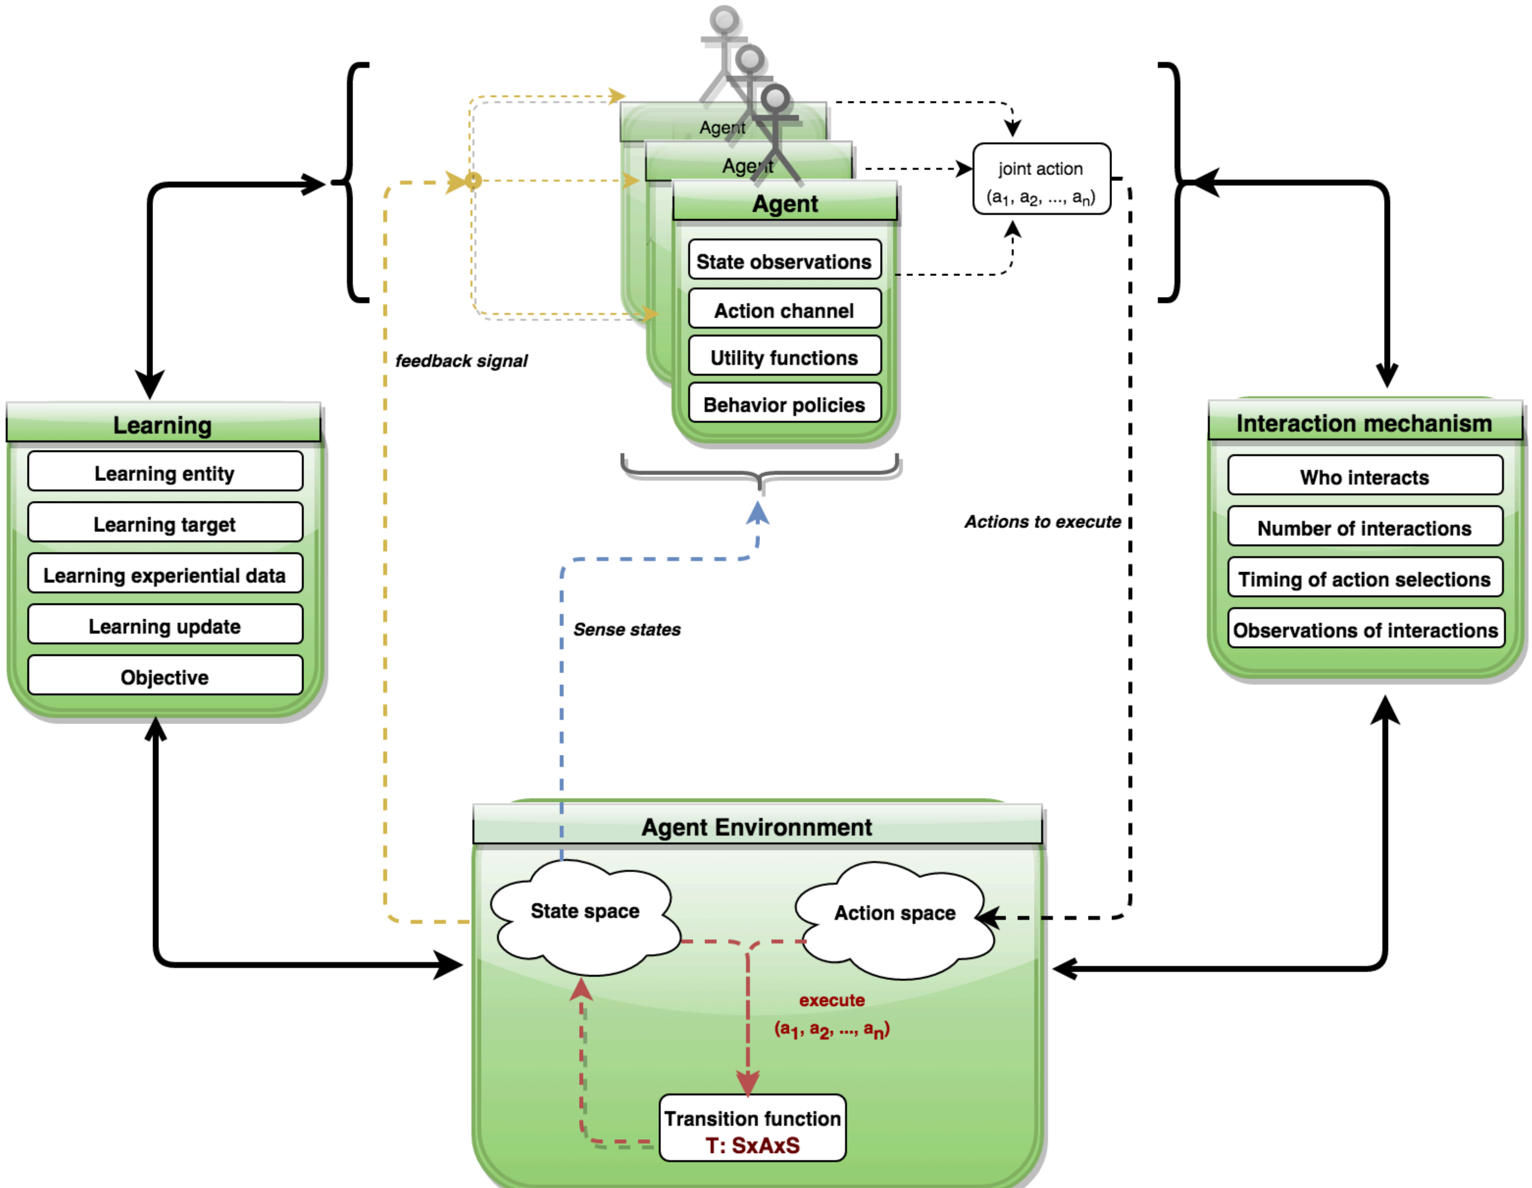

The Environment

- Specifies the **state space**, **action space**, and **transition function**.
  - The state space specifies the set of states
  - The action space is the set of actions available to an individual agent
  - The transition function specifies the environment dynamics

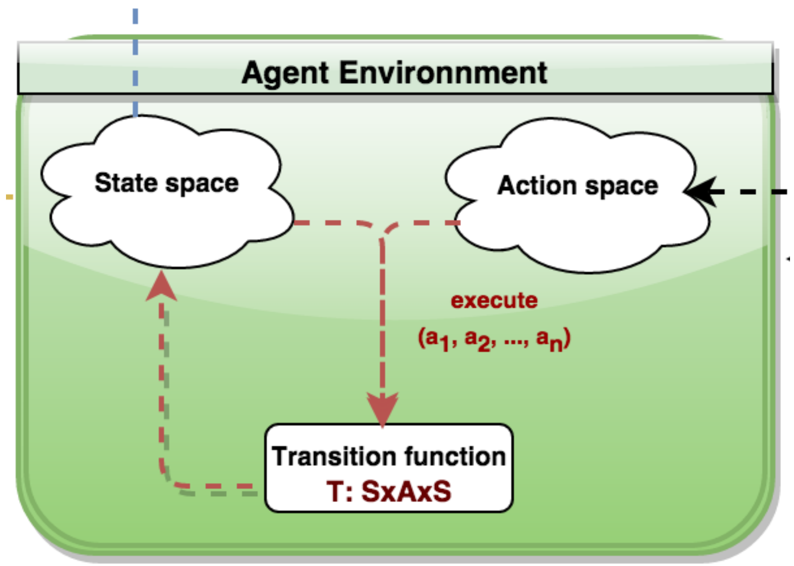

The Agents

- Communication channels
  - with the environment for sensing state and specifying actions
  - between one another
- Utility functions indicating their preferences
- Policies for selecting actions

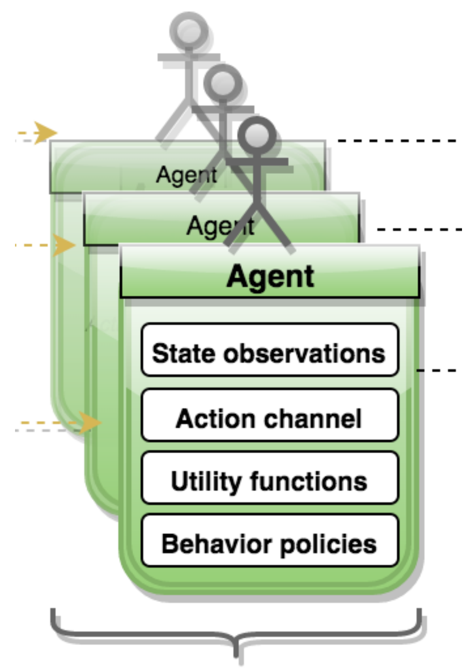

The Interaction Mechanism

- The **interaction mechanism** defines:
  - how long agents interact with one another
  - with which other agents
  - what they observe about other agents
  - frequency as well as simultaneous or sequentially

For example, at one extreme, agents may be fully aware of each other's behavior policies; or, at the other extreme, they may only observe the effects of their actions on the environment.

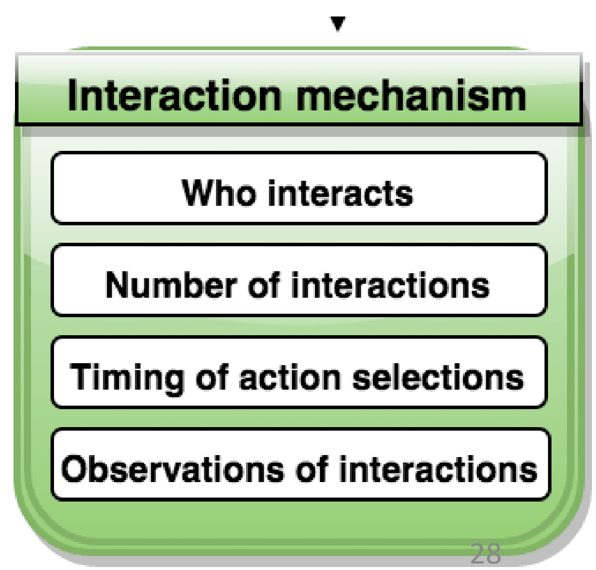

The Learning Mechanism

The **learning mechanism** is defined by:

- the learning entity: individual or group level
- learning target: what is being learnt
- the learning experiential data: what information is available to the learning entity
- the learning update: defines how the behaviour of the learning entity is updated
- the objective of learning:
  is a representation of the goal,
  or evaluation function,
  of the learning process.

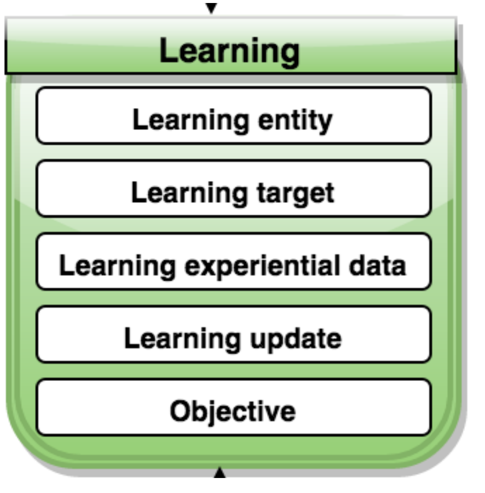

Historical Perspective

Two periods:

- Startup: late 1980s until early 2000s
- Consolidation: early 2000s until now

Historical Perspective: Startup

- Broad exploration of realizations of MAS:

- Adaptive parallel computation inspired by nature
  - Ant Systems
  - Flocking or herding behaviour
  - Evolutionary computation
  - Neural networks
  - Imitation Learning
- Now known as the *Artificial Life* field

- First multiagent reinforcement learning efforts (Littman et al.)

- Breadth-first paradigmatic exploration

Historical Perspective: Consolidation

Depth-first exploration

- Focus on reinforcement learning in a game-theoretic context
- Theoretical foundations of MAL

Wrapping up

- Multi-agent learning introduces new problems, with the main questions being:
  - What should the agents learn?
  - How should the agents learn this?

- Both questions will be addressed
  - What – establishing group objectives and mechanisms to reach those objectives
  - How – algorithms and methods that allow agents to learn optimal/good joint actions

- There exist several paradigmatic solutions

Reading Material

Main Reading

- Chapter on *Multiagent Learning* by Tuyls & Tumer (on Vital)

Optional Reading

- Bloembergen *et al*. (2015). Evolutionary Dynamics of Multi-Agent Learning: A Survey. http://jair.org/papers/paper4818.html

- Busoniu *et al*. (2008). A comprehensive survey of multiagent reinforcement learning. http://busoniu.net/files/papers/smcc08.pdf