LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

Praktikum Mobile und Verteilte Systeme

An Introduction to Autonomous Systems

Prof. Dr. Claudia Linnhoff-Popien
André Ebert, Thomy Phan, Robert Müller, Steffen Illium
http://www.mobile.ifi.lmu.de

WiSe 2018/19

Outline

- **An Introduction to Autonomous Systems**
  - Motivation, Definition and Challenges
  - Artificial Intelligence
- **Decision Making in Autonomous Systems**
  - Markov Decision Processes
  - Planning
  - Reinforcement Learning
- **Applications and Further Challenges**
  - Deep Reinforcement Learning
  - Combining Planning and Reinforcement Learning
  - Further Challenges

Outline

- **An Introduction to Autonomous Systems**
  - Motivation, Definition and Challenges
  - Artificial Intelligence
- **Decision Making in Autonomous Systems**
  - Markov Decision Processes
  - Planning
  - Reinforcement Learning
- **Applications and Further Challenges**
  - Deep Reinforcement Learning
  - Combining Planning and Reinforcement Learning
  - Further Challenges

→ Motivation, Definition and Challenges

(Possible) Real-World Applications

Smart Grids / Cities

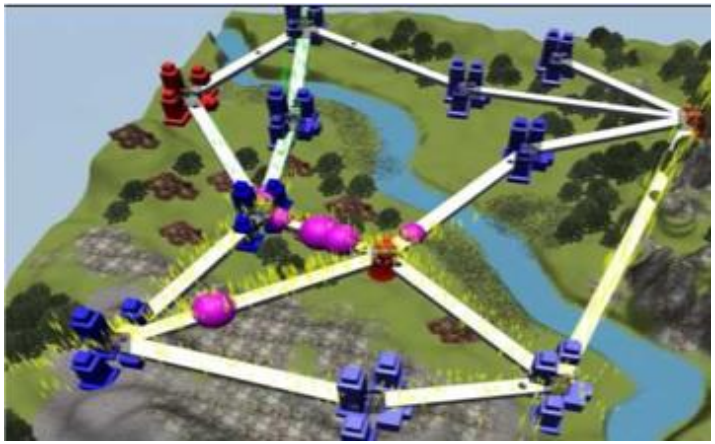

Intelligent / Mobile Networks

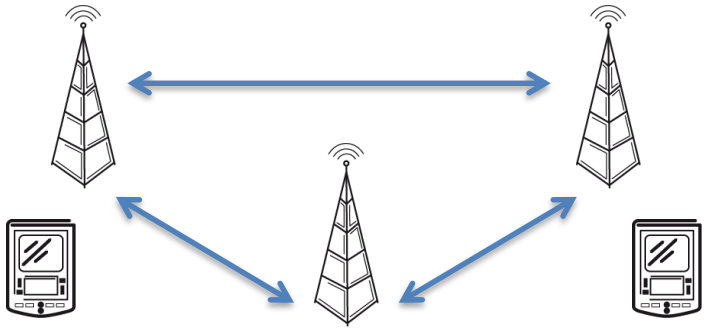

Robotics

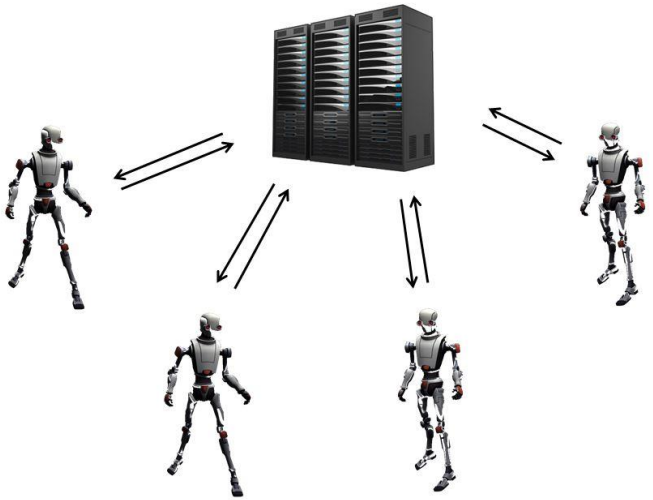

Industry 4.0

What is an Autonomous System?

Definition: A system, which can operate without human intervention.

Properties of Autonomous Systems

- Self-CHOP
  - Self-Configuration
  - Self-Healing
  - Self-Optimization
  - Self-Protection
- More Self-Properties
  - Self-Learning
  - Self-Organization
  - Self-Awareness
  - Self-Regulation
  - …

M. Salehie and L. Tahvildari, Autonomic Computing: Emergent Trends and Open Problems, ACM SIGSOFT Software Engineering Notes, 2005

Challenges of Autonomous Systems

- Dynamic Environment
- Complexity
- Many Constraints:
  - Perception
  - Computational and Memory Resources
  - Energy Consumption
  - Communication
- Safety and Risk
- Security
- Quality Management

→ **Artificial Intelligence**

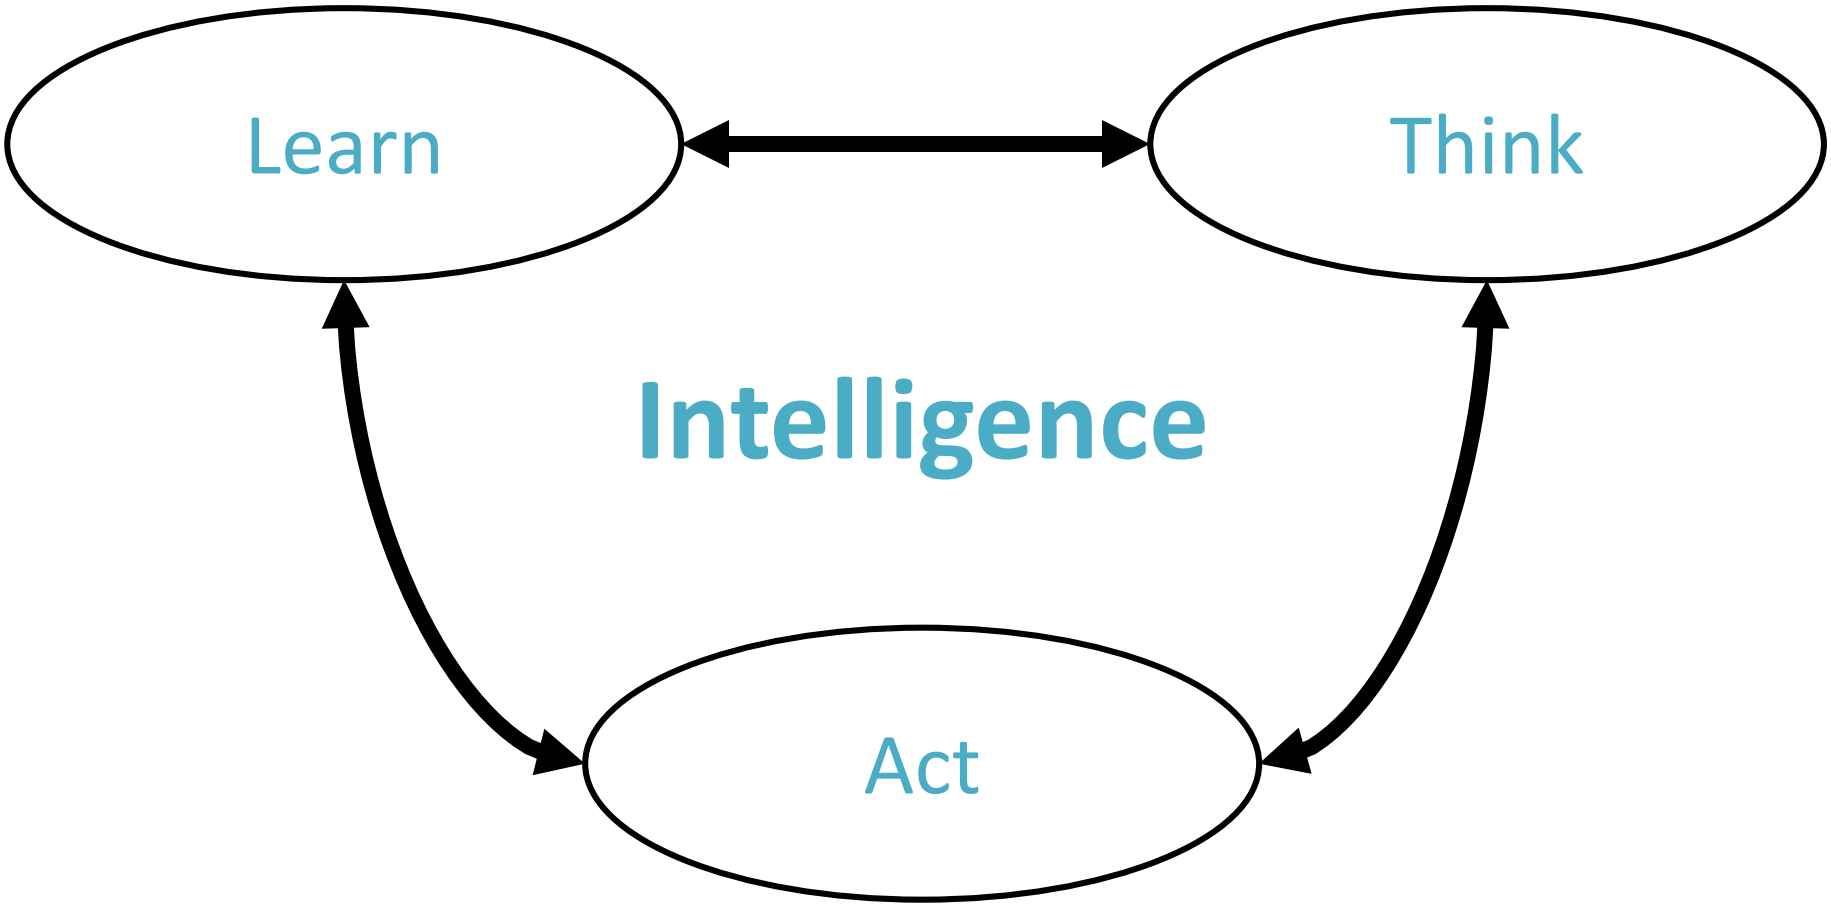
Learn
Think
Intelligence
Act

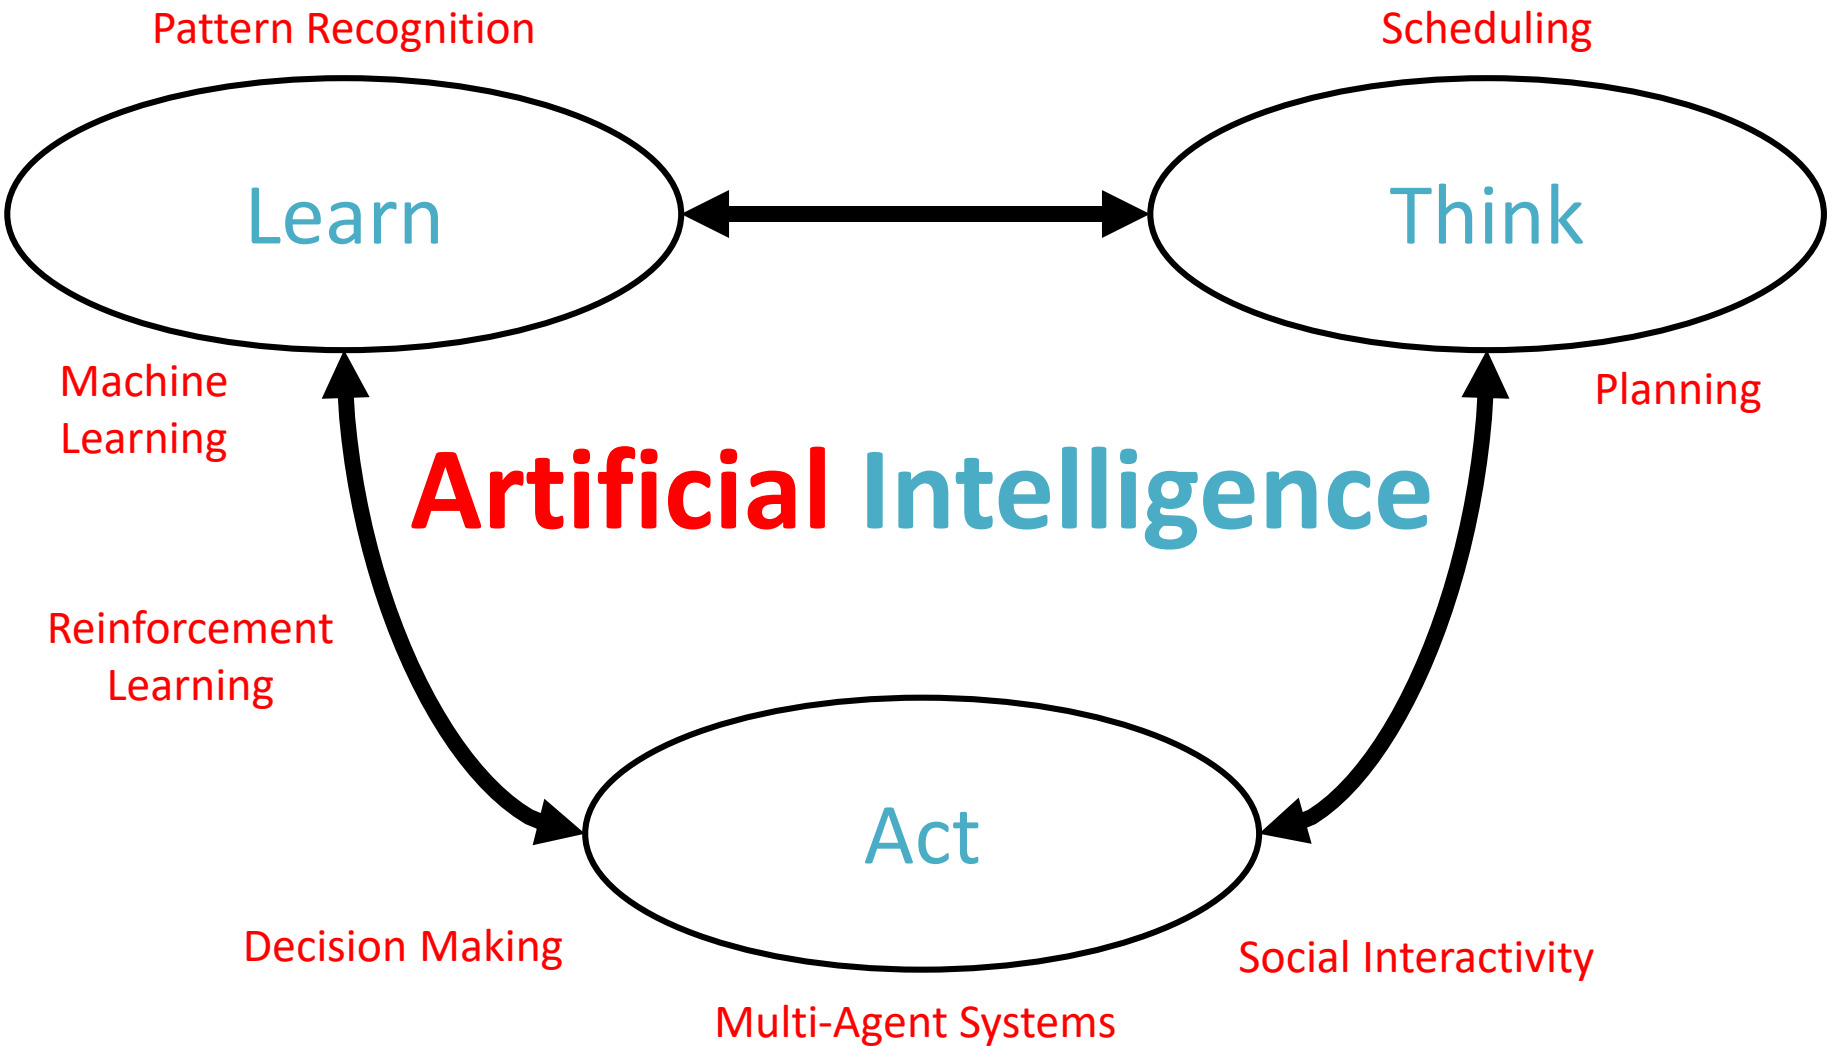
Pattern Recognition
Scheduling
Learn
Think
Machine Learning
Planning
Artificial Intelligence
Reinforcement Learning
Act
Decision Making
Social Interactivity
Multi-Agent Systems

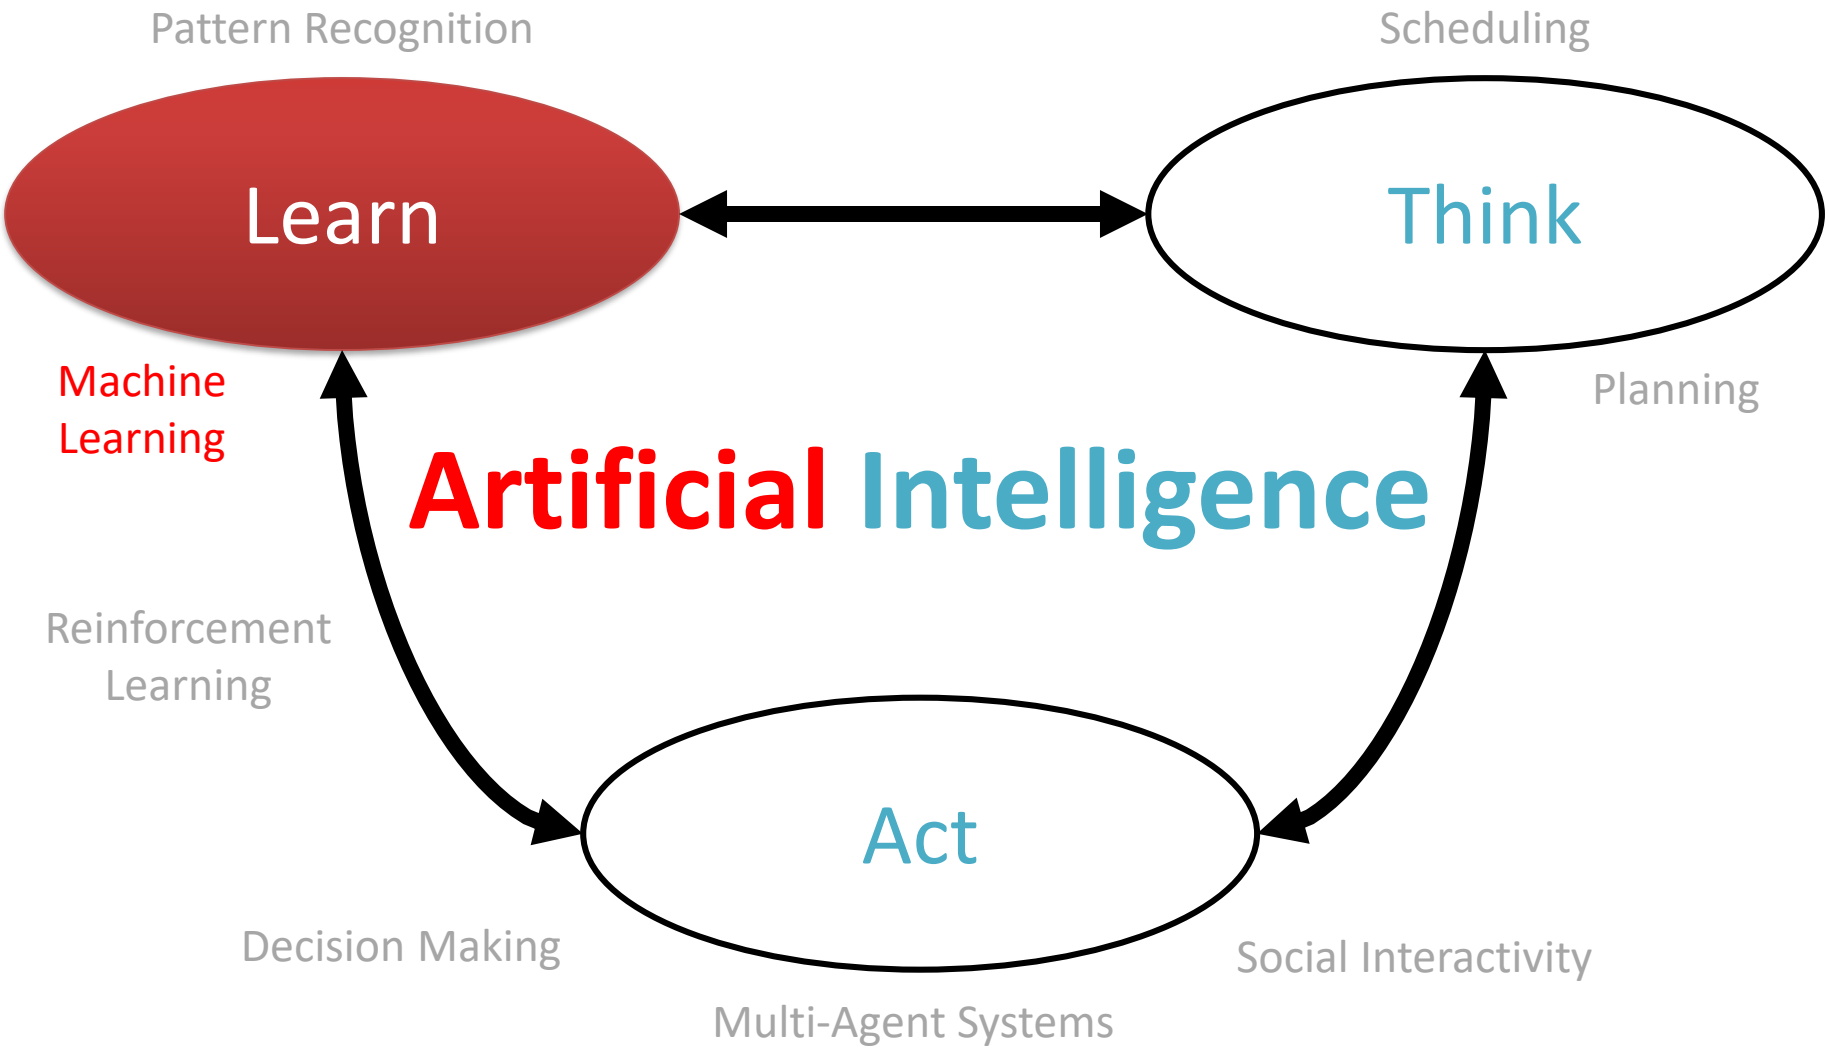
Pattern Recognition
Scheduling
Learn
Think
Machine Learning
Planning
Artificial Intelligence
Reinforcement Learning
Act
Decision Making
Social Interactivity
Multi-Agent Systems

Machine Learning

- **Goal:** Create programs that learn how to solve complex problems
- Learn statistical models from experience / data
- Use learned models for e.g.
  - Object Recognition
  - Prediction
  - Control
  - Compression
  - Data Generation

Why Machine Learning?

- **Goal:** Create programs that learn how to solve complex problems
- Many problems cannot be solved by engineering hard-coded solutions
  - Too many aspects to consider
  - Too many rules
  - Hard adaption to changes
  - …
- Examples:
  - Object recognition in images
  - Natural language processing
  - User behaviour analytics
  - Locomotion

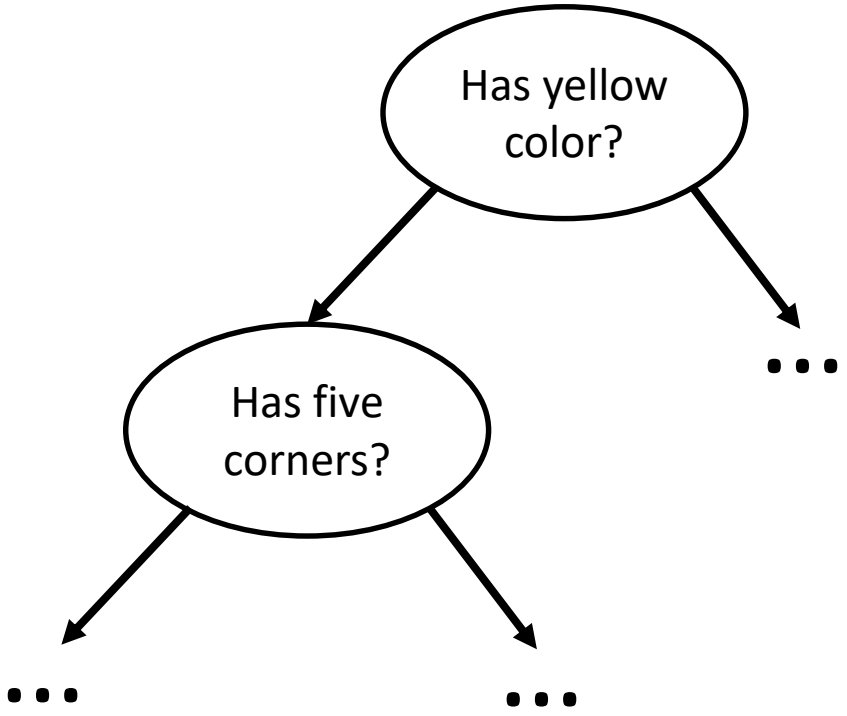

Types of Machine Learning

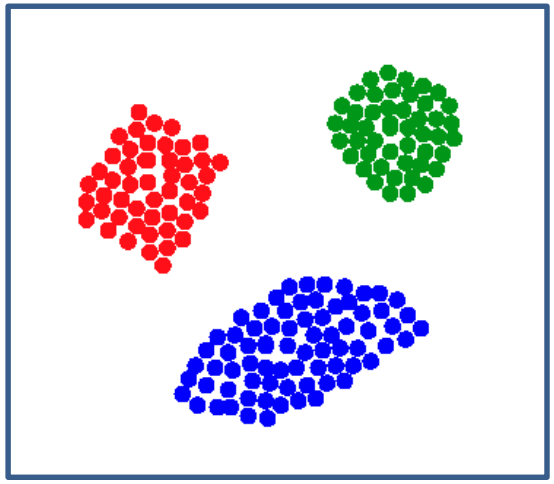

Unsupervised Learning

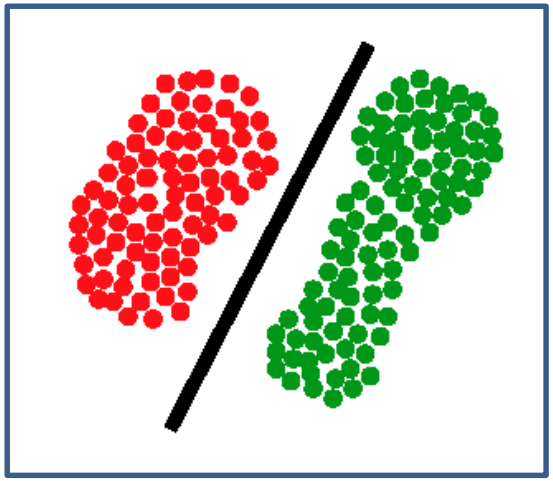

Supervised Learning

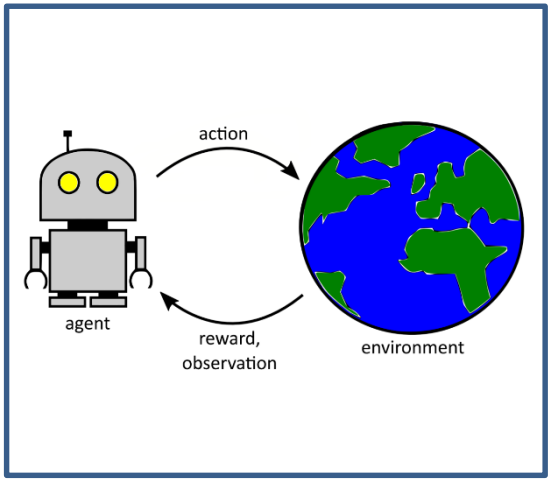

Reinforcement Learning

Challenges of Machine Learning

- Data Availability
- Data Complexity
- Efficiency
- Compactness
- Interpretability
- Robustness
- Adaptivity

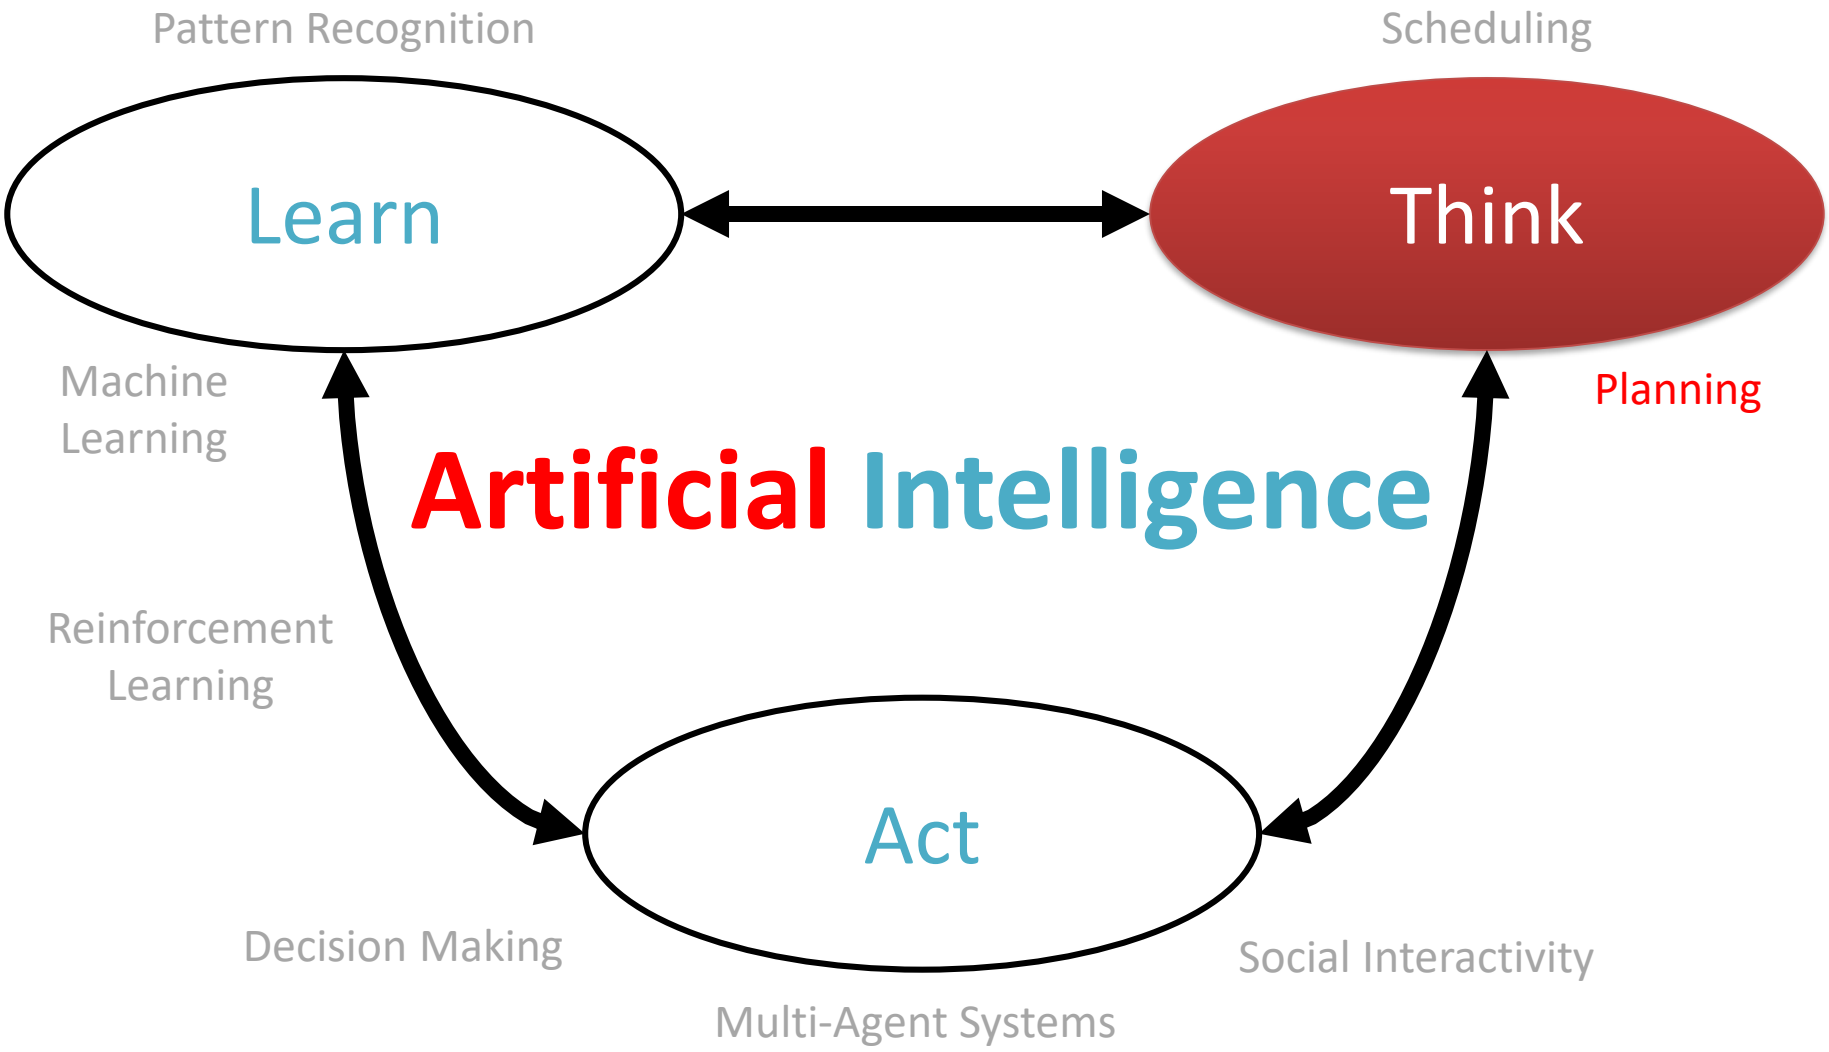
Pattern Recognition
Scheduling
Learn
Think
Machine Learning
Planning
Artificial Intelligence
Reinforcement Learning
Act
Decision Making
Social Interactivity
Multi-Agent Systems

Automated Planning

- **Goal:** Find (near-)optimal strategies to solve complex problems
- Use (heuristic) lookahead search on a **given model** of the problem
  - Model can be defined by rules (e.g. physical laws, game rules, etc.) or statistical approximations (e.g. using machine learning)
- Automated Planning can be applied to
  - Decision and Control tasks
  - Task Allocation Problems
  - Recommender Systems

Why Automated Planning?

- **Goal:** Find (near-)optimal strategies to solve complex problems
- Planning is necessary, if **explicit reasoning** is required:
  - Consideration of risks and uncertainties
  - Consideration of hard constraints
  - Interpretable decisions and recommendations
  - …
- Examples:
  - Routing / Navigation
  - Gate Assignment Problem
  - General (Video) Games
  - Locomotion

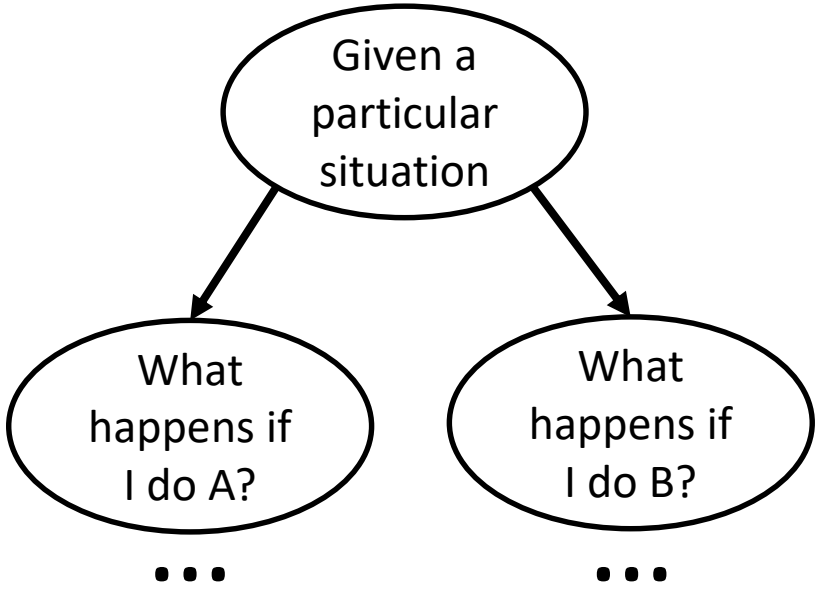

Planning Approaches (Examples)

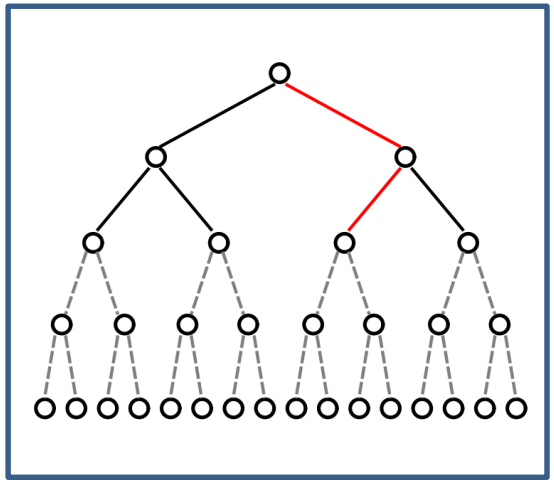

Tree Search

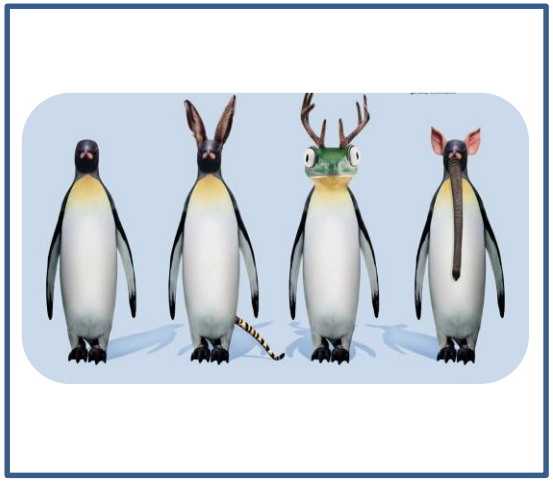

Evolutionary Computation

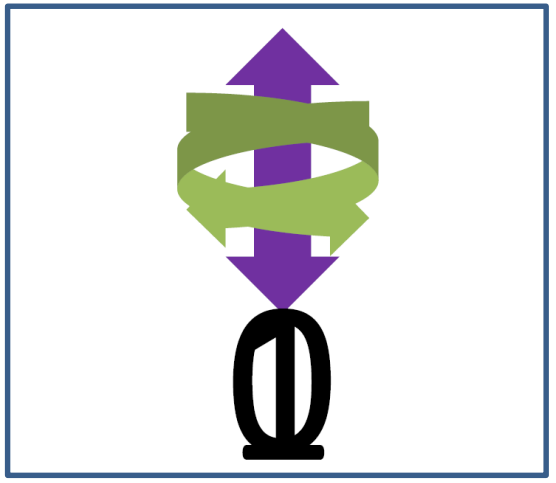

Quantum Computing

Challenges of Automated Planning

- Model Availability
- Model Uncertainty
- Computational and Memory Efficiency
- Real-time Planning

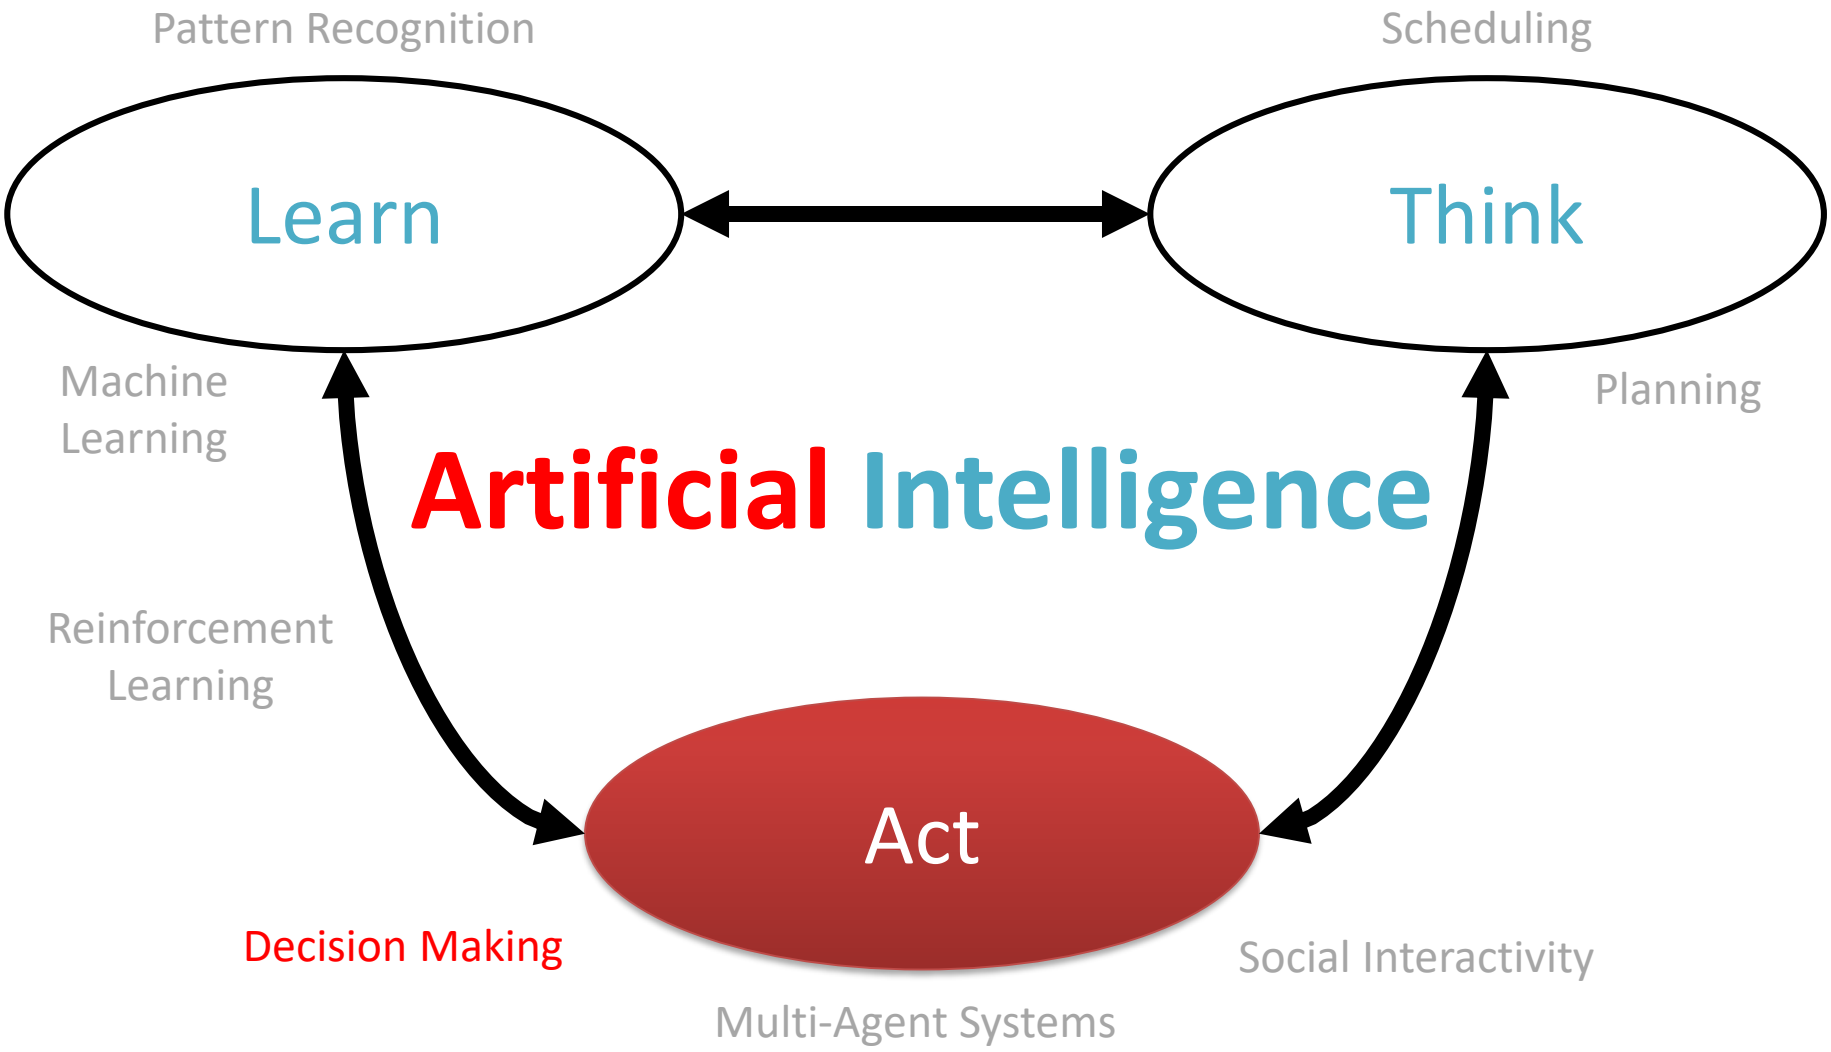
Pattern Recognition
Scheduling
Learn
Think
Machine Learning
Planning
Artificial Intelligence
Reinforcement Learning
Act
Decision Making
Social Interactivity
Multi-Agent Systems

Decision Making

- **Goal:** Select actions to solve a complex task

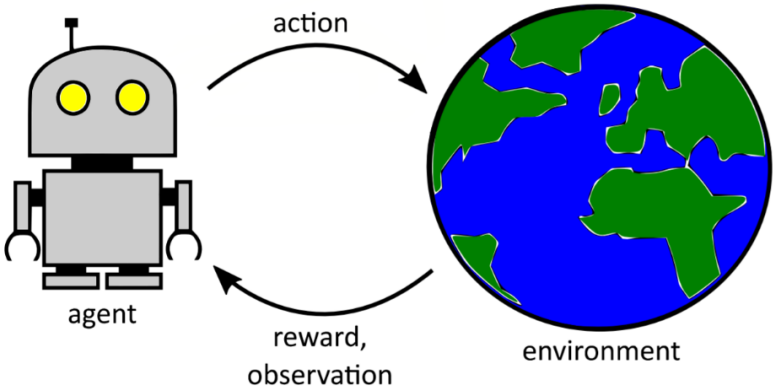

- **Given:** A set of states and a set of actions
  - Actions have to be explored for each state to determine their long term consequences
  - Promising actions should be exploited more often

    ⇨ *Exploration-Exploitation dilemma*
  - Feedback may be delayed

Decision Making Models

Single-Agent

- Multi-Armed Bandit (single-state problems)
- Markov Decision Process (MDP)
- Partially Observable MDP (POMDP)

Multiple Agents

- Multi-Agent MDP
- Markov Game
- Decentralized MDP / POMDP
- Partially Observable Stochastic Game

Decision Making Approaches

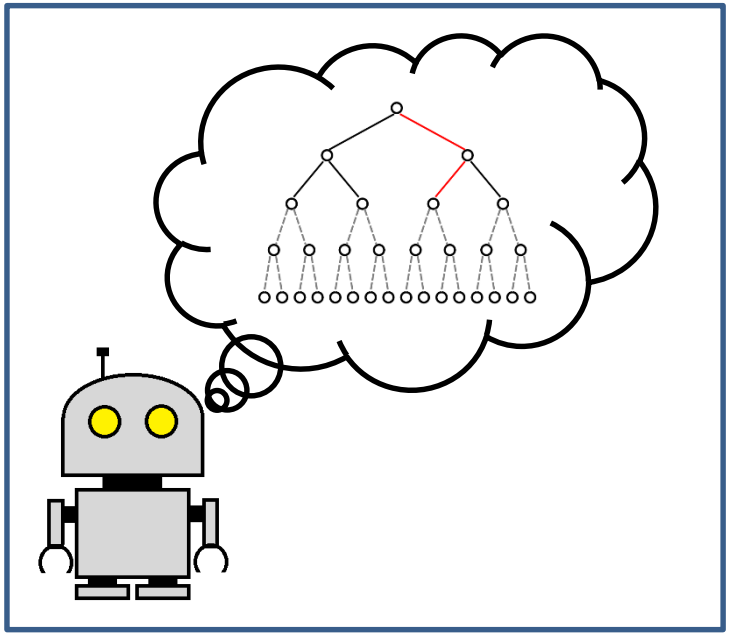

Automated Planning

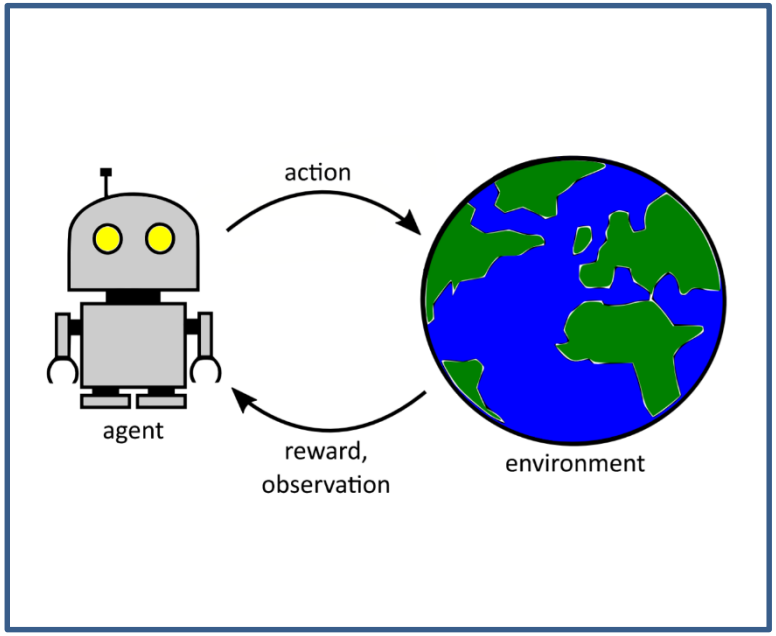

Reinforcement Learning

Challenges of Decision Making

- Problem Complexity
- Sparse Feedback
- Sample Efficiency
- Uncertainty

Outline

- **An Introduction to Autonomous Systems**
  - Motivation, Definition and Challenges
  - Artificial Intelligence
- **Decision Making in Autonomous Systems**
  - Markov Decision Processes
  - Planning
  - Reinforcement Learning
- **Applications and Further Challenges**
  - Deep Reinforcement Learning
  - Combining Planning and Reinforcement Learning
  - Further Challenges

Thank you!