University at Buffalo
Department of Computer Science and Engineering
School of Engineering and Applied Sciences

MULTI-AGENT RL

Alina Vereshchaka

CSE4/510: Reinforcement Learning

March 9, 2020

MDP

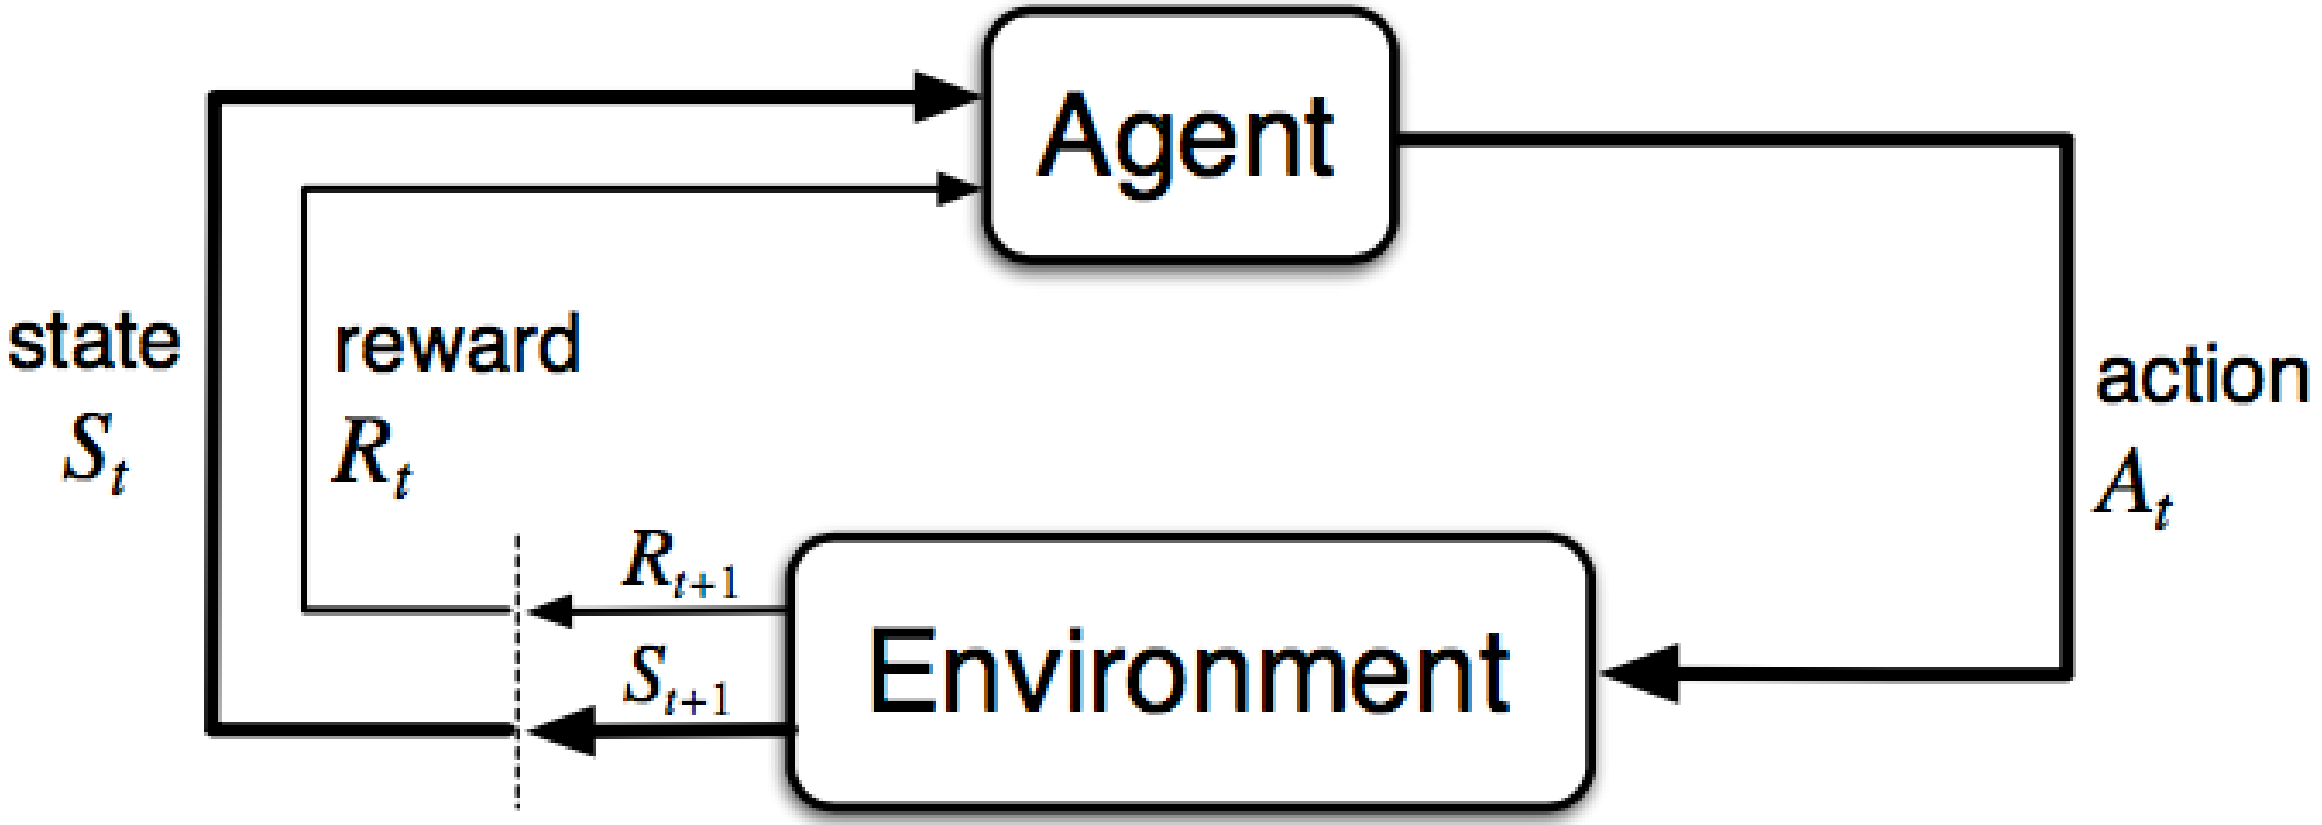

Multi-agent Reinforcement Learning (MARL)

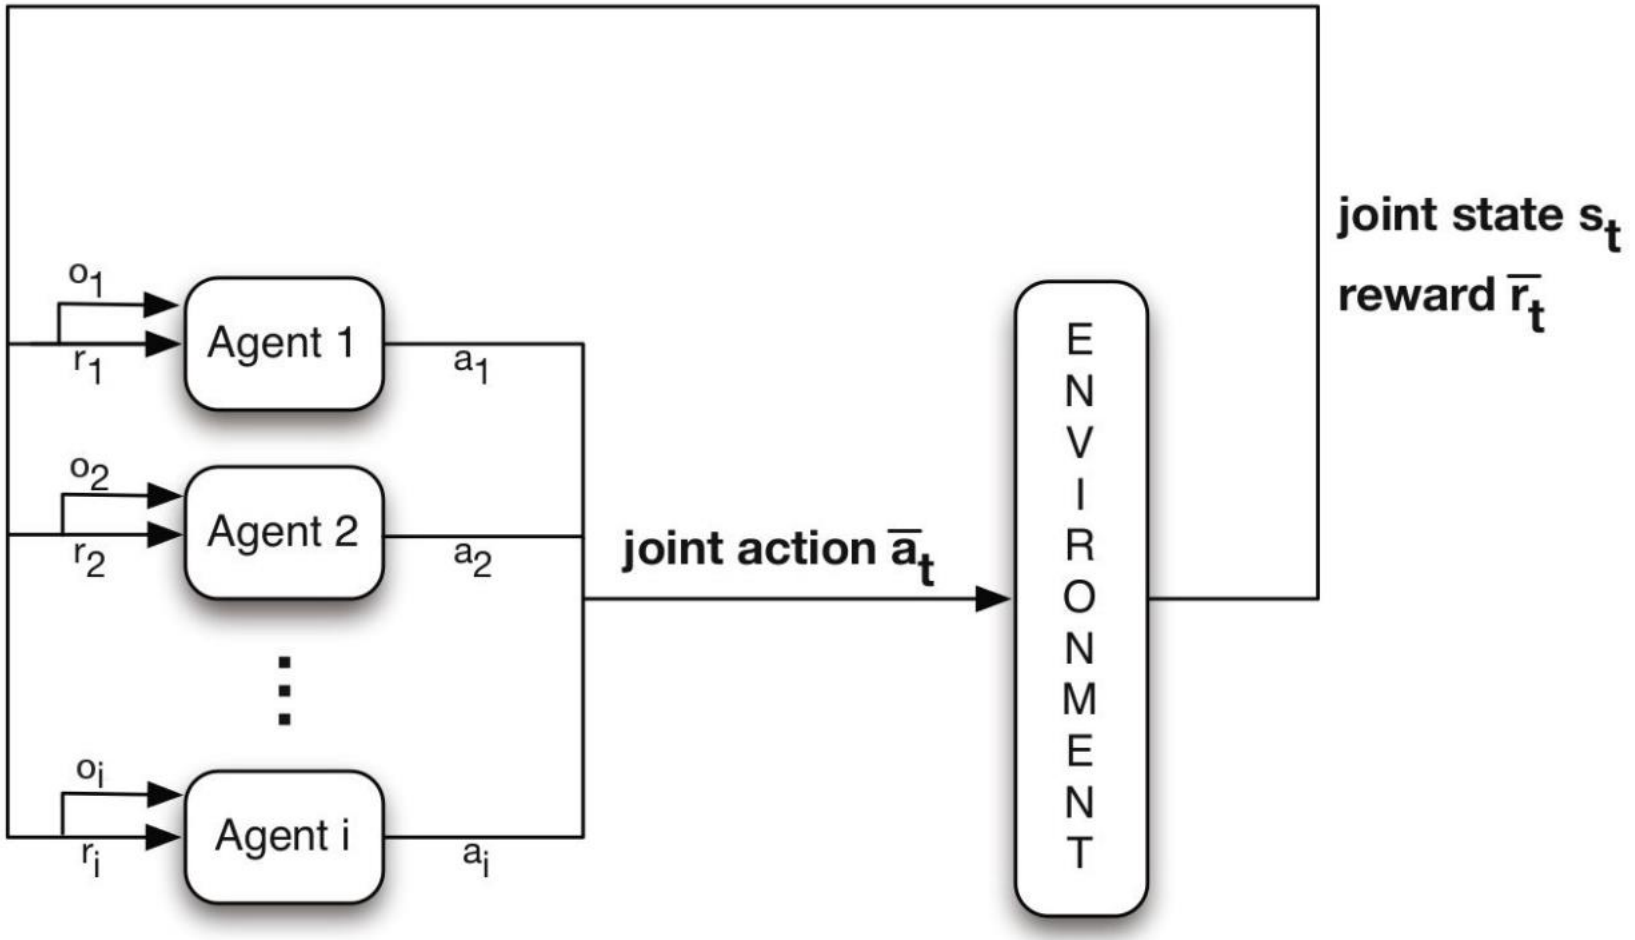

Source: Nowe, Vrancx & De Hauwere 2012

Axes of MARL I

Centralized:

- One brain / algorithm deployed across many agents

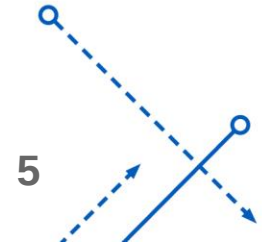

Axes of MARL I

Centralized:

- One brain / algorithm deployed across many agents

Decentralized:

- All agents learn individually
- Communication limitations defined by environment

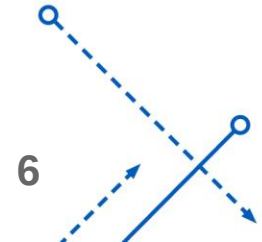

Axes of MARL II

Prescriptive:

- Suggests how agents should behave

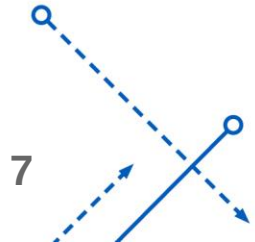

Axes of MARL II

Prescriptive:

- Suggests how agents should behave

Descriptive:

- Forecast how agent will behave

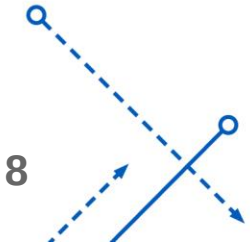

Axes of MARL III

Cooperative: Agents cooperate to achieve a goal

- Shared team reward

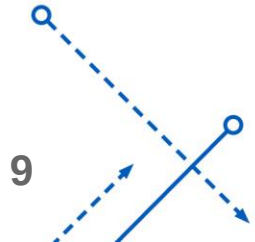

Axes of MARL III

Cooperative: Agents cooperate to achieve a goal

- Shared team reward

Competitive: Agents compete against each other

- Zero-sum games
- Individual opposing rewards

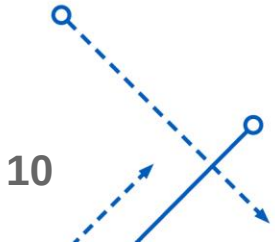

Axes of MARL III

Cooperative: Agents cooperate to achieve a goal

- Shared team reward

Competitive: Agents compete against each other

- Zero-sum games
- Individual opposing rewards

Neither: Agents maximize their utility which may require cooperating and/or competing

- General-sum games

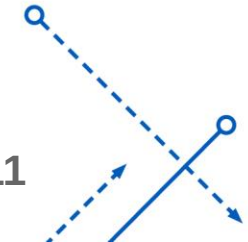

Axes of MARL IV

Numbers of agents:

- One (single-agent)

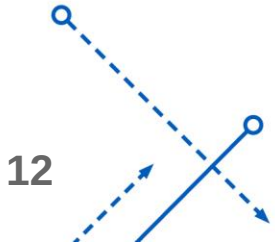

Axes of MARL IV

Numbers of agents:

- One (single-agent)
- Two (very common)

Axes of MARL IV

Numbers of agents:

- One (single-agent)
- Two (very common)
- Finite

Axes of MARL IV

Numbers of agents:

- One (single-agent)
- Two (very common)
- Finite
- Infinite


NYK LINE

Transportation Problem

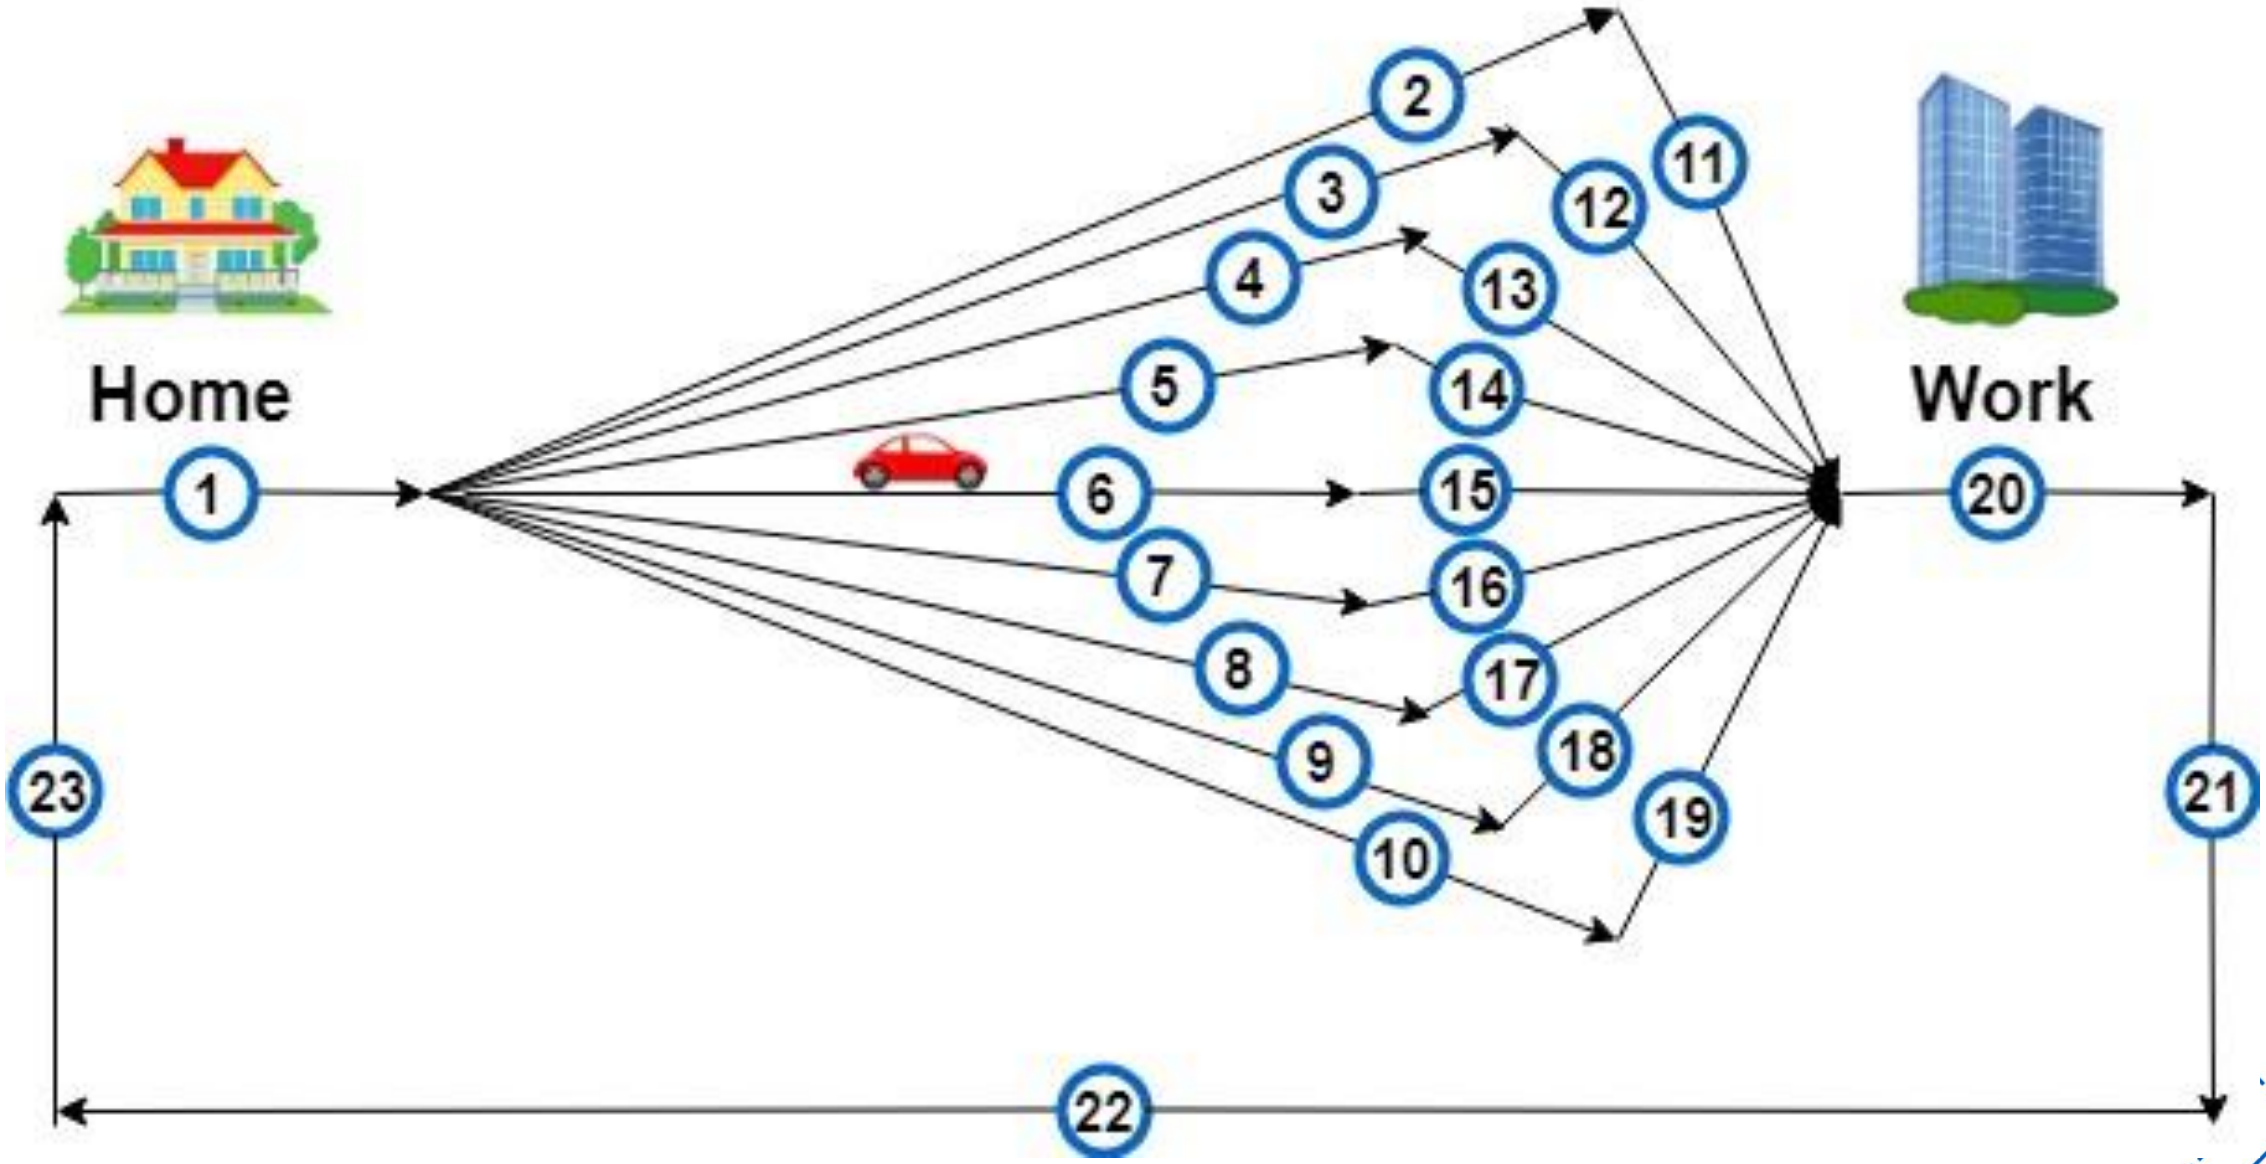

Thank you!