Advanced Topics in AI

Finite Horizons and Discounting

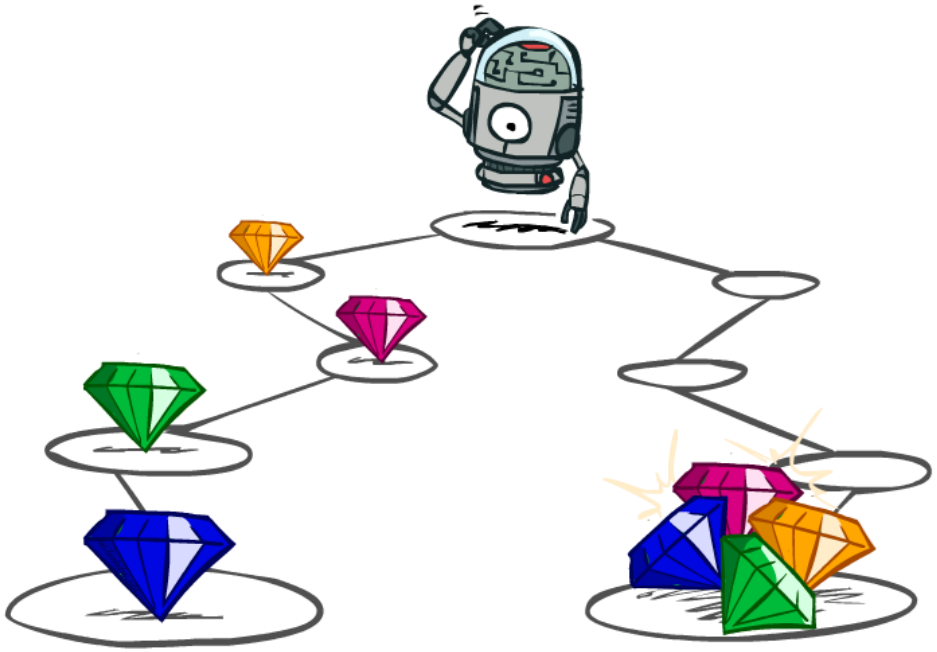

Instructor: Prof. Dr. techn. Wolfgang Nejdl

Leibniz University Hannover

Co-financed by the Connecting Europe Facility of the European Union

[These slides were created by Dan Klein and Pieter Abbeel for CS188 Intro to AI at UC Berkeley. All materials are available at http://ai.berkeley.edu.]

Utilities of Sequences

- What preferences should an agent have over reward sequences?

- More or less?

  [1, 2, 2] or [2, 3, 4]

- Now or later?

  [0, 0, 1] or [1, 0, 0]

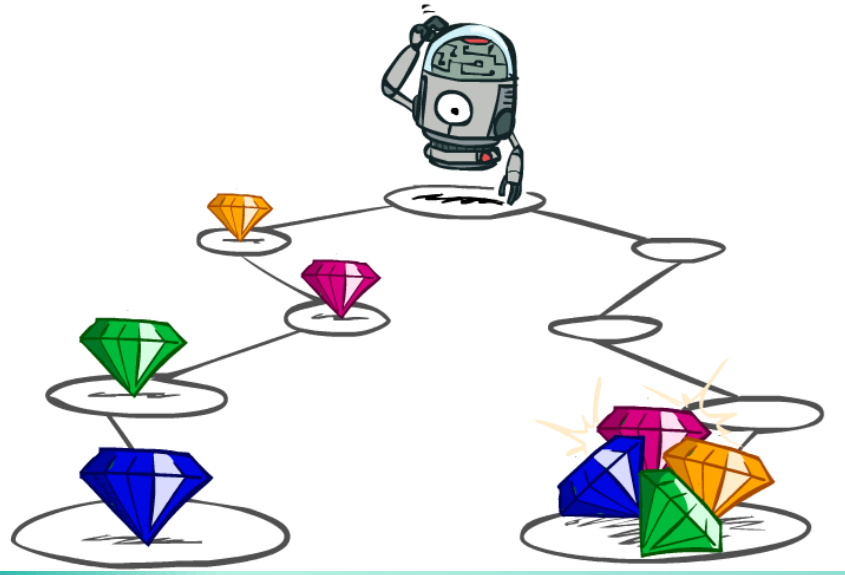

Co-financed by the Connecting Europe Facility of the European Union

Discounting

- It's reasonable to maximize the sum of rewards
- It's also reasonable to prefer rewards now to rewards later
- One solution: values of rewards decay exponentially

1

Worth Now

γ

Worth Next Step

γ^2

Worth In Two Steps

Co-financed by the Connecting Europe Facility of the European Union

Discounting

- How to discount?
  - Each time we descend a level, we multiply in the discount once
- Why discount?
  - Reward now is better than later
  - Can also think of it as a 1-gamma chance of ending the process at every step
  - Also helps our algorithms converge
- Example: discount of 0.5
  - U([1,2,3]) = 1*1 + 0.5*2 + 0.25*3
  - U([1,2,3]) < U([3,2,1])

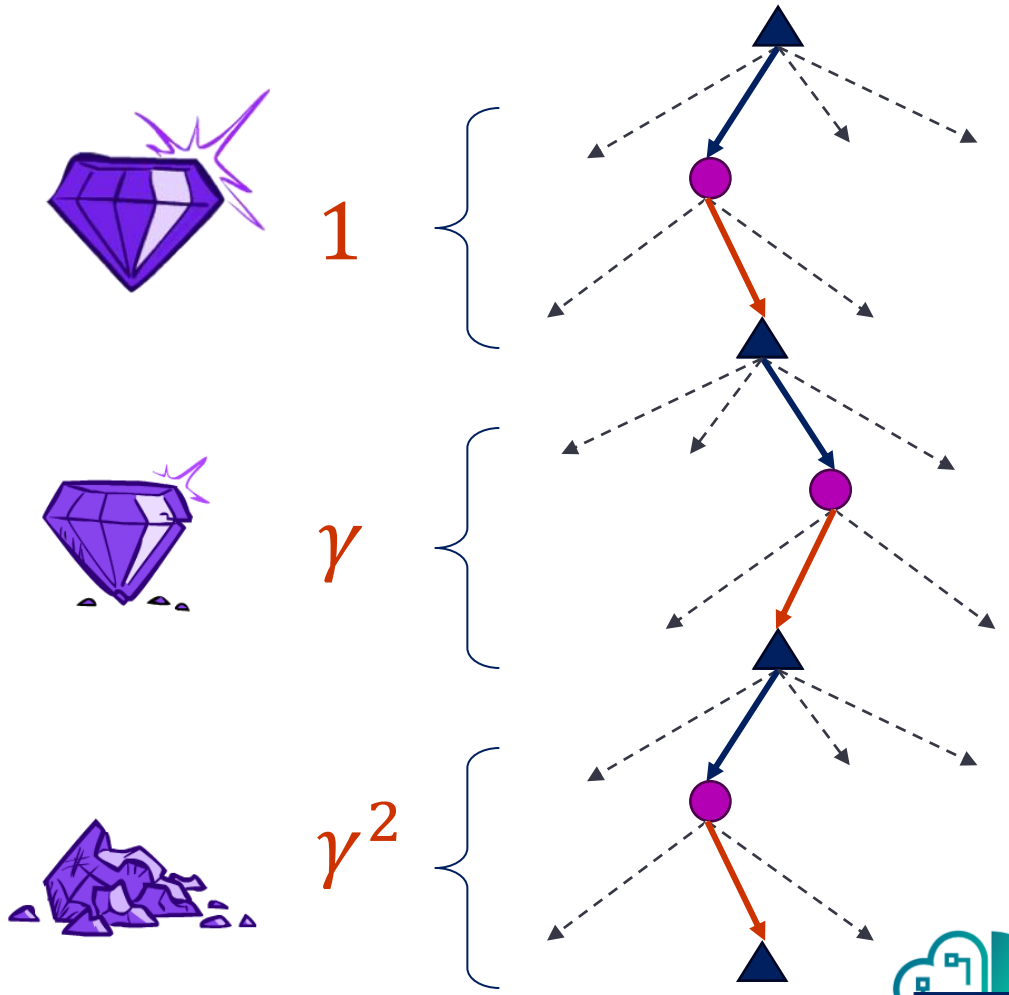

Co-financed by the Connecting Europe Facility of the European Union

Quiz: Discounting

- Given:

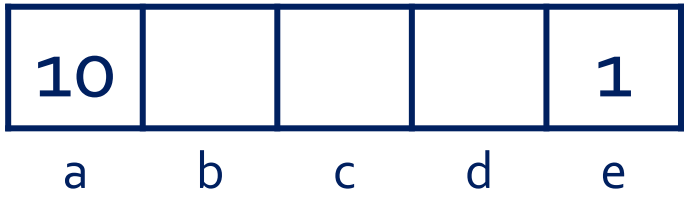

  - Actions: East, West, and Exit (only available in exit states a, e)
  - Transitions: deterministic
- Quiz 1: For $\gamma = 1$, what is the optimal policy?

- Quiz 2: For $\gamma = 0.1$, what is the optimal policy?

- Quiz 3: For which γ are West and East equally good when in state d?

Co-financed by the Connecting Europe Facility of the European Union

Quiz: Discounting

- Given:

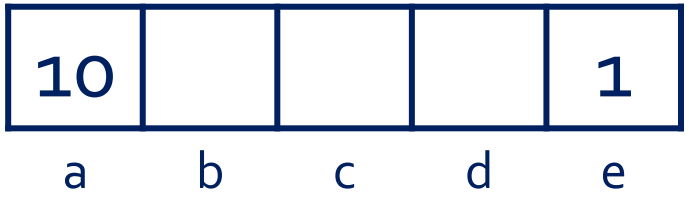

  - Actions: East, West, and Exit (only available in exit states a, e)
  - Transitions: deterministic
- Quiz 1: For $\gamma = 1$, what is the optimal policy?

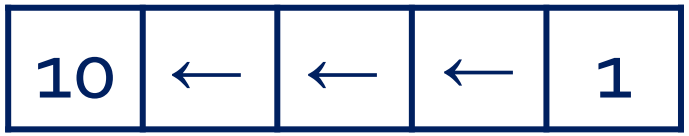

- Quiz 2: For $\gamma = 0.1$, what is the optimal policy?

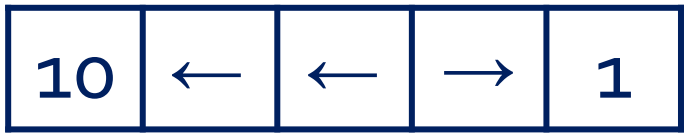

- Quiz 3: For which γ are West and East equally good when in state d?

$$1\gamma = 10\,\gamma^3$$

Co-financed by the Connecting Europe Facility of the European Union

Infinite Utilities?!

- Problem: What if the game lasts forever? Do we get infinite rewards?
- Solutions:
  - Finite horizon: (similar to depth-limited search)
    - Terminate episodes after a fixed T steps (e.g. life)
    - Gives nonstationary policies (π depends on time left)
  - Discounting: use $0 < \gamma < 1$

$$U([r_o, \ldots, r_\infty]) = \sum_{t=0}^{\infty} \gamma^t r_t \leq \frac{R_{max}}{1-\gamma}$$

    - Smaller γ means smaller "horizon" – shorter term focus

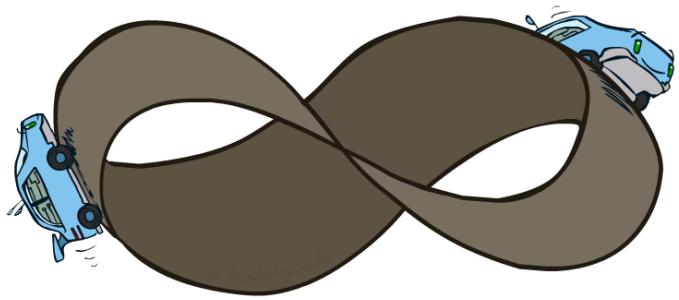

  - Absorbing state: guarantee that for every policy, a terminal state will eventually be reached (like "overheated" for racing)

Co-financed by the Connecting Europe Facility of the European Union

Advanced Topics in AI

Next: Solving MDPs

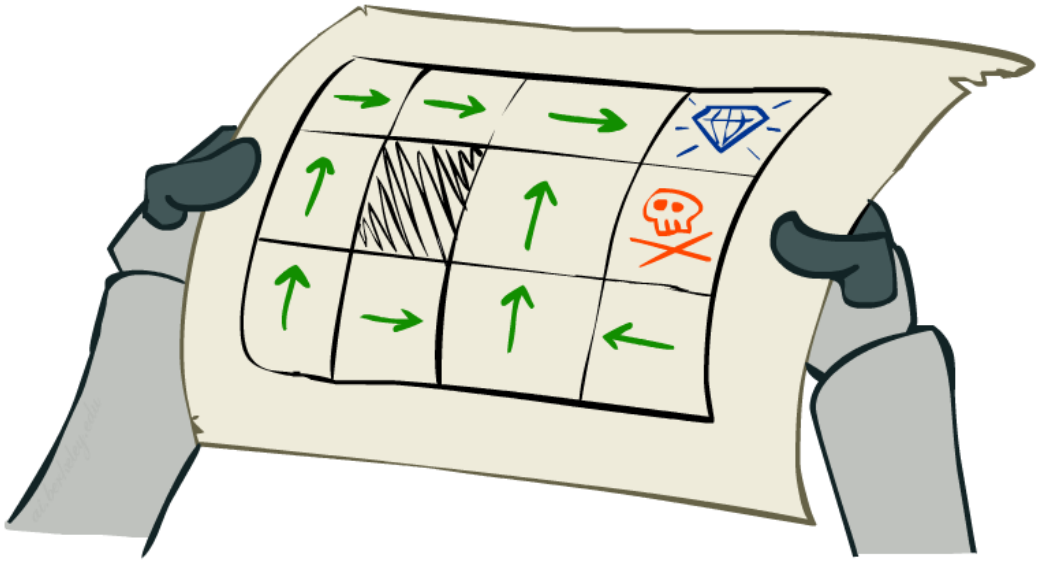

Instructor: Prof. Dr. techn. Wolfgang Nejdl

Leibniz University Hannover

Co-financed by the Connecting Europe Facility of the European Union

[These slides were created by Dan Klein and Pieter Abbeel for CS188 Intro to AI at UC Berkeley. All materials are available at http://ai.berkeley.edu.]