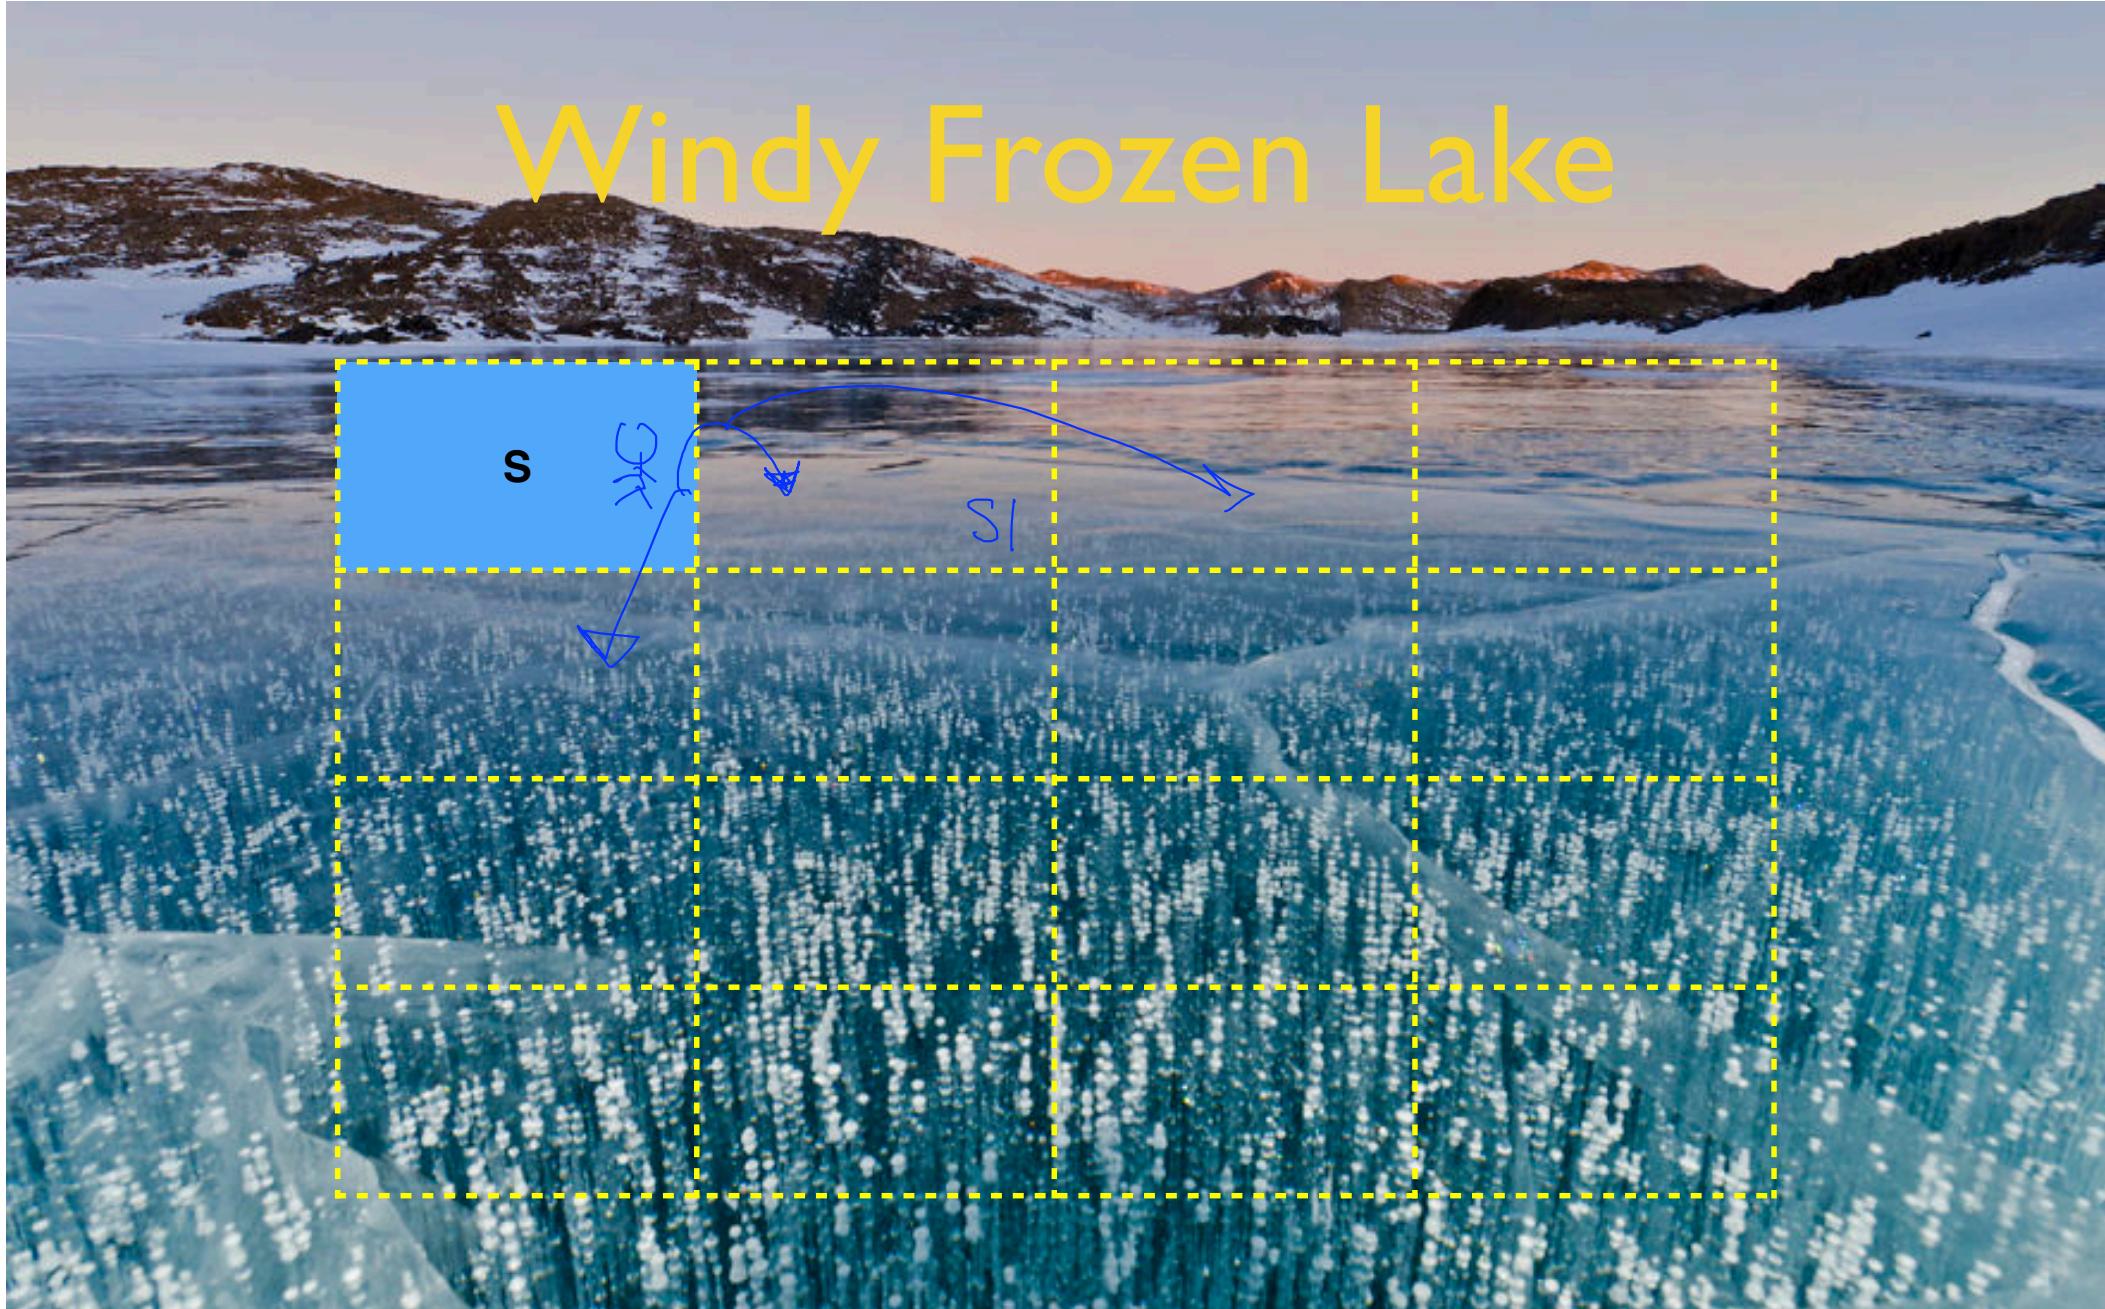




Lecture 5: Windy Frozen Lake Nondeterministic world!

Reinforcement Learning with TensorFlow & OpenAI Gym
Sung Kim <hunkim+ml@gmail.com>

Windy Frozen Lake



Deterministic VS Stochastic (nondeterministic)

- **In deterministic models** the output of the model is fully determined by the parameter values and the initial conditions initial conditions
- **Stochastic models** possess some inherent randomness.
 - The same set of parameter values and initial conditions will lead to an ensemble of different outputs.

Deterministic

Diagram illustrating a deterministic environment transition:

Initial State (S): SFFF
Action: Down
Transition: FFFF
Action: Right
Transition: HFFG
Action: Right
Transition: HFFG
Action: Right
Transition: HFFG

Final State (G): HFFG

Annotations:

- A blue circle highlights the initial state "SFFF".
- A blue arrow points from the initial state to the first transition "FFFH".
- A blue arrow points from the final state "HFFG" back to the initial state "SFFF".
- A blue circle highlights the final state "HFFG".

```
# Register FrozenLake with is_slippery False
register(
    id='FrozenLake-v3',
    entry_point='gym.envs.toy_text:FrozenLakeEnv',
    kwargs={'map_name': '4x4', 'is_slippery': False}
)
```

```
env = gym.make('FrozenLake-v3')
```

```
(Down)
('State: ', 6, 'Action: ', 1,
SFFF
FHFH
FFFH
HFFG
(Down)
('State: ', 10, 'Action: ', 1,
SFFF
FHFH
FFFH
HFFG
(Down)
('State: ', 14, 'Action: ', 1,
SFFF
FHFH
FFFH
HFFG
(Right)
('State: ', 15, 'Action: ', 2,
('Finished with reward', 1.0)
```

Annotations:

- A blue arrow points from the first transition "FFFH" down to the second transition "FFFH".
- A blue arrow points from the second transition "FFFH" down to the third transition "HFFG".
- A blue circle highlights the final state "HFFG".
- A blue circle highlights the reward message "('Finished with reward', 1.0)".
- A blue oval encloses the reward message "('Finished with reward', 1.0)".

Stochastic (non-deterministic)

```
# is_slippery True
env = gym.make('FrozenLake-v0')
```

S

```
SFFF
FHFH
FFFH
HFFG

SFFF
FHFH
FFFH
HFFG
    (Right)
('State: ', 0, 'Action: ', 2,
SFFF
FHFH
FFFH
HFFG
    (Right)
('State: ', 4, 'Action: ', 2,
SFFF
FHFH
FFFH
HFFG
    (Down)
('State: ', 5, 'Action: ', 1,
('Finished with reward', 0.0)
```

S

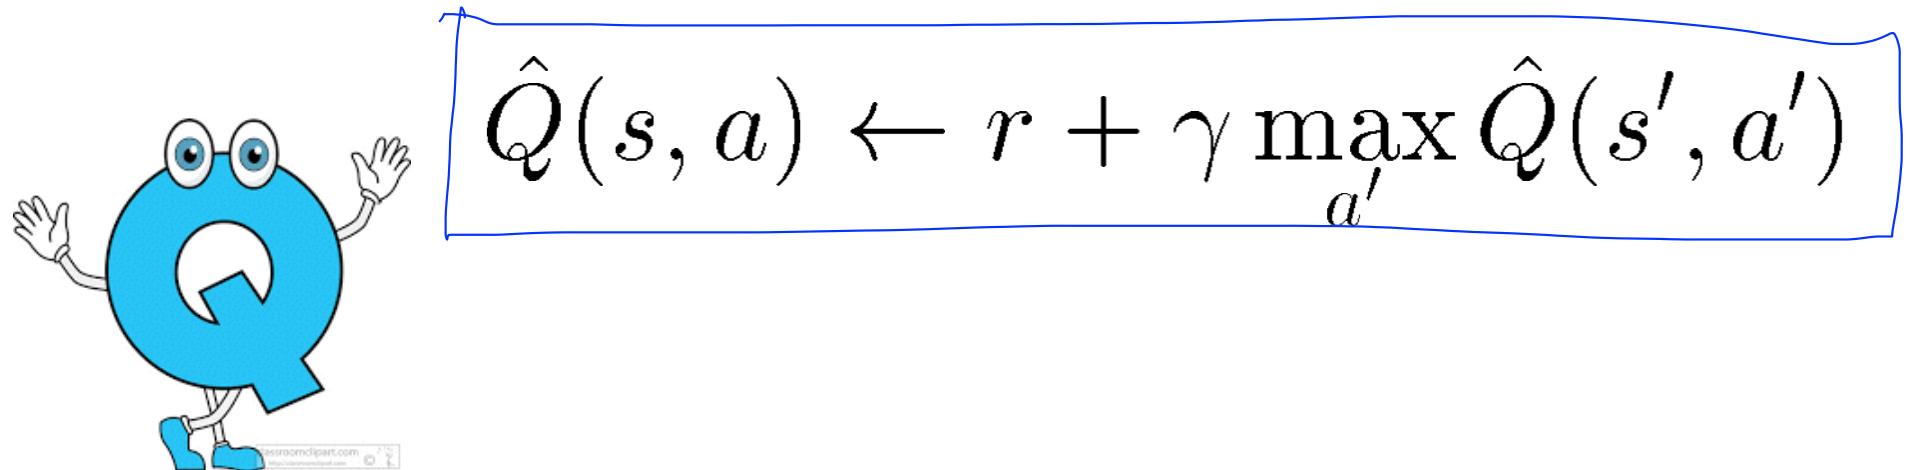
```
SFFF
FHFH
FFFH
HFFG

    (Right)
('State: ', 0, 'Action: ', 2,
SFFF
FHFH
FFFH
HFFG
    (Right)
('State: ', 1, 'Action: ', 2,
SFFF
FHFH
FFFH
HFFG
    (Right)
('State: ', 1, 'Action: ', 2,
SFFF
FHFH
FFFH
HFFG
    (Right)
('State: ', 5, 'Action: ', 2,
('Finished with reward', 0.0)
```

6 → X

Stochastic (non-deterministic) worlds

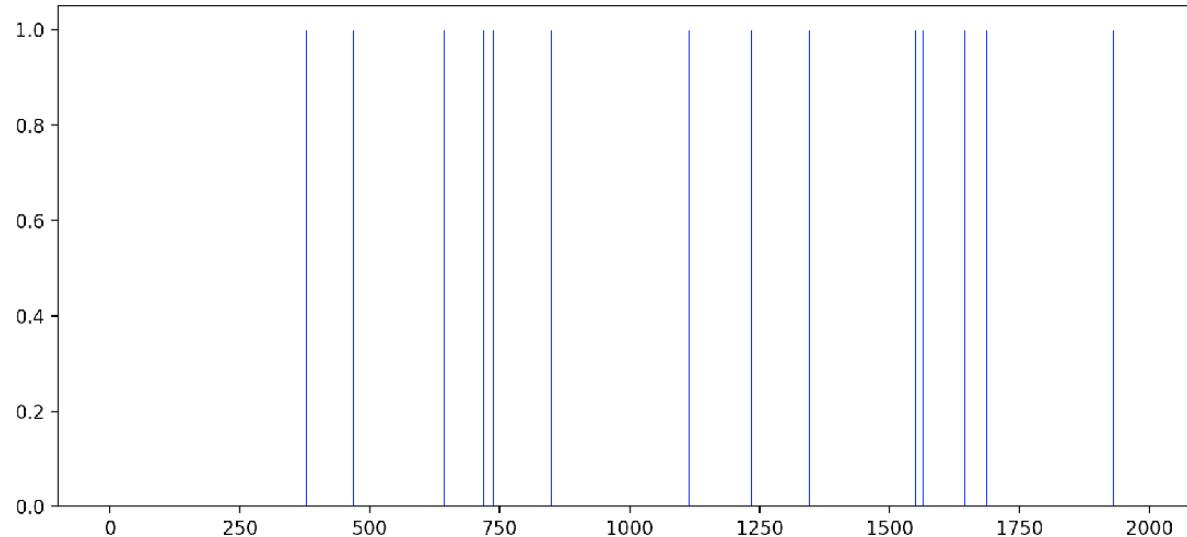
- Unfortunately, our Q-learning (for deterministic worlds) does not work anymore
- Why not?



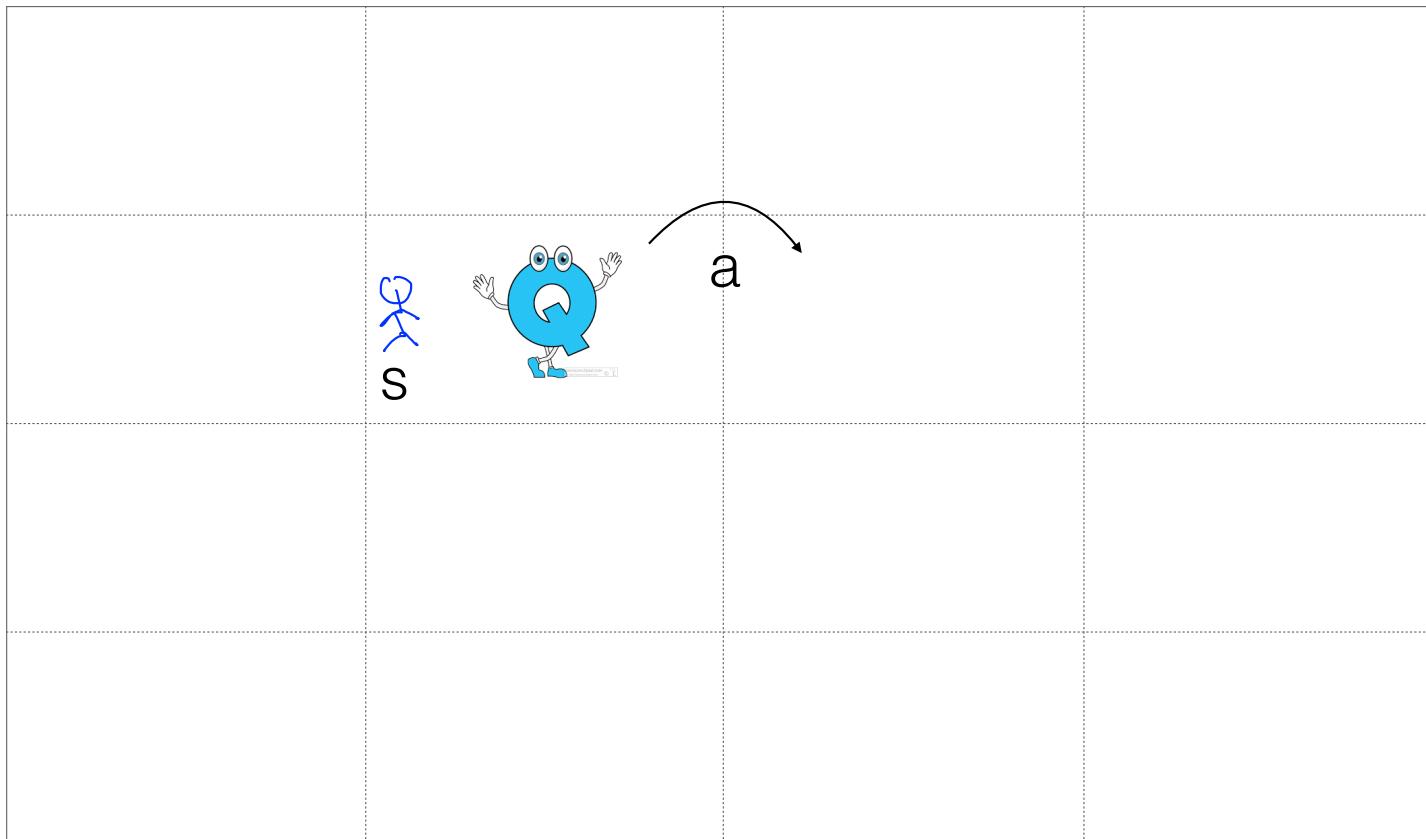
Our previous Q-learning does not work

```
env = gym.make('FrozenLake-v0')
```

Score over time: 0.0165



Why does not work in stochastic (non-deterministic) worlds?



Stochastic (non-deterministic) world

- Solution?
 - Listen to $Q(s')$ (just a little bit)
 - Update $Q(s)$ little bit (learning rate)
- Like our life mentors
 - Don't just listen and follow one mentor
 - Need to listen from many mentors

BC-HBR-WAKE-UP-CALL-MENTORS-NYTSF

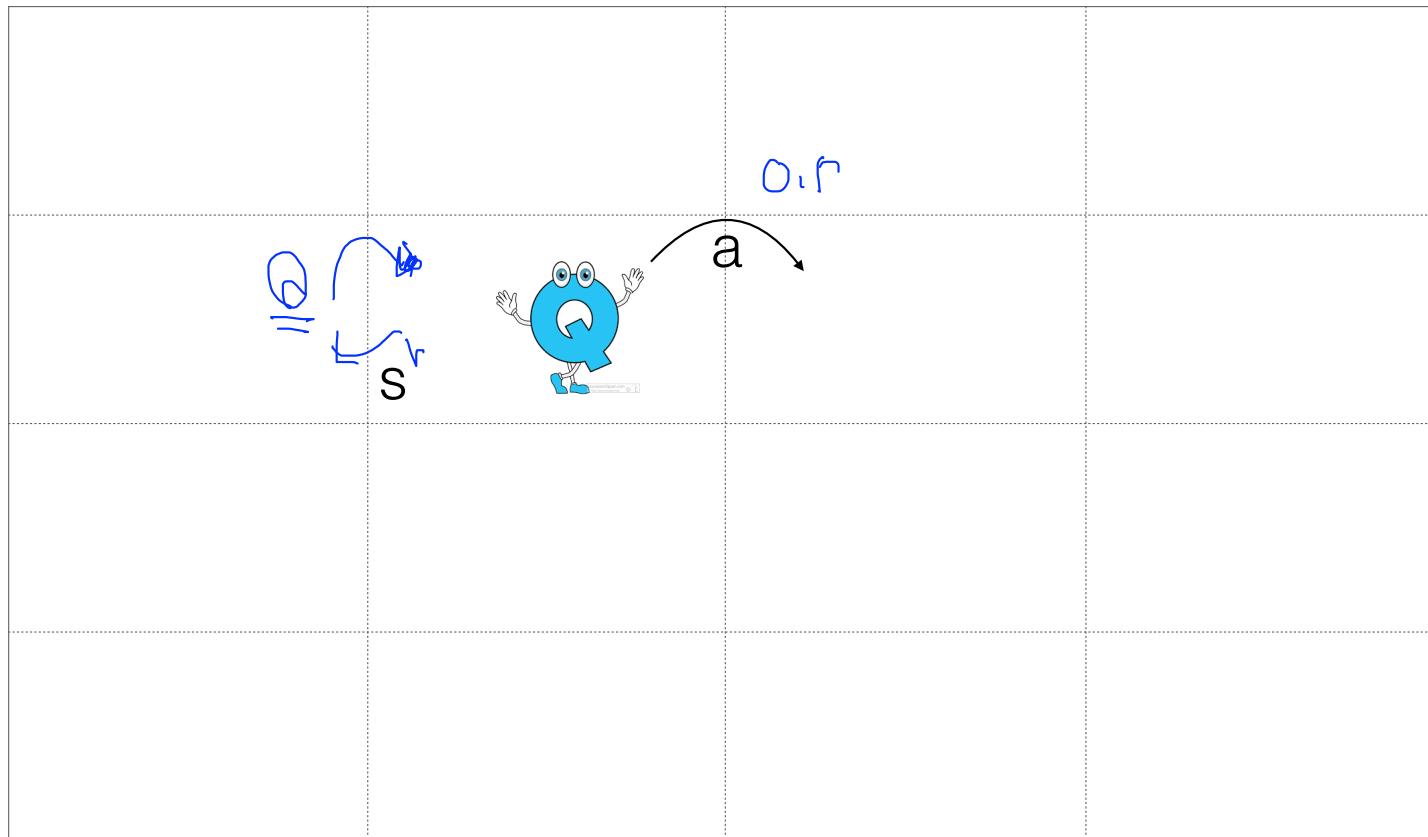
Your Career Needs Many Mentors, Not Just One

20.1.2017 15.00

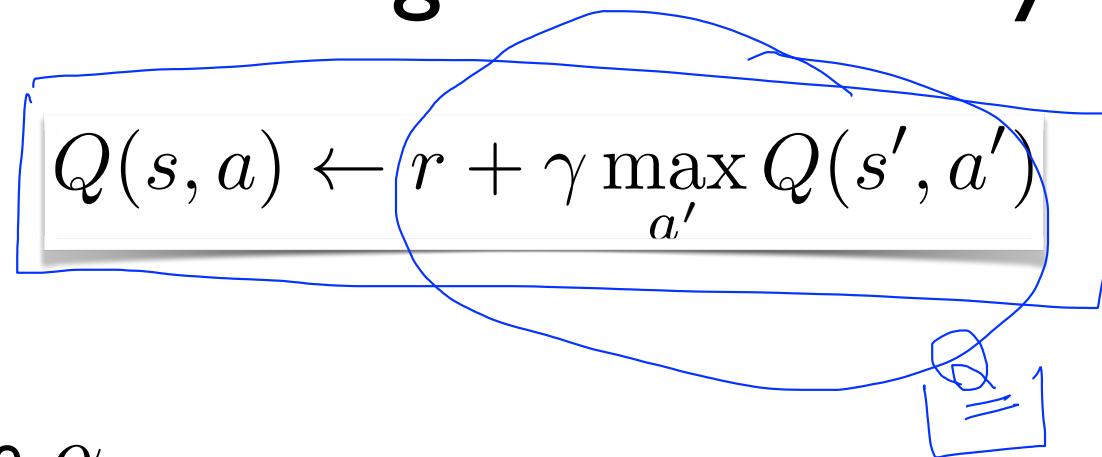


<http://m.kauppalehti.fi/uutiset/your-career-needs-many-mentors--not-just-one/gp3Q4rTp>

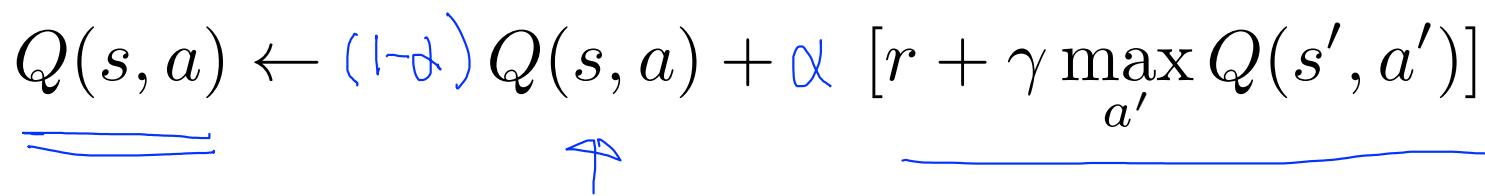
Stochastic (non-deterministic) world



Learning incrementally

$$Q(s, a) \leftarrow r + \gamma \max_{a'} Q(s', a')$$


- Learning rate, $\underline{\underline{\alpha}}$
 - $\alpha = 0.1$

$$\underline{\underline{Q(s, a) \leftarrow (1-\alpha)Q(s, a) + \alpha [r + \gamma \max_{a'} Q(s', a')]}}$$


Learning with learning rate

$$Q(s, a) \leftarrow r + \gamma \max_{a'} Q(s', a')$$

V

$$Q(s, a) \leftarrow (1 - \alpha)Q(s, a) + \alpha[r + \gamma \max_{a'} Q(s', a')]$$

Learning with learning rate

$$Q(s, a) \leftarrow r + \gamma \max_{a'} Q(s', a')$$

$$Q(s, a) \leftarrow (1 - \alpha)Q(s, a) + \alpha[r + \gamma \max_{a'} Q(s', a')]$$
$$Q(s, a) \leftarrow Q(s, a) + \alpha[r + \gamma \max_{a'} Q(s', a') - Q(s, a)]$$

Q-learning algorithm

For each s, a initialize table entry $\hat{Q}(s, a) \leftarrow 0$

Observe current state \underline{s}



Do forever:

- Select an action \underline{a} and execute it
- Receive immediate reward \underline{r}
- Observe the new state $\underline{s'}$
- Update the table entry for $\hat{Q}(s, a)$ as follows:

$$\underline{Q(s, a) \leftarrow (1 - \alpha)Q(s, a) + \alpha[r + \gamma \max_{a'} Q(s', a')]}$$

- $s \leftarrow s'$

Convergence

\hat{Q} denote learner's current approximation to Q .



$$\hat{Q}(s, a) \leftarrow (1 - \alpha)\hat{Q}(s, a) + \alpha[r + \gamma \max_{a'} \hat{Q}(s', a')]$$

Can still prove convergence of \hat{Q} to Q [Watkins and Dayan, 1992]

Next

Lab: Stochastic worlds

