

Reinforcement Learning

Johanni Brea

Introduction to Machine Learning

This is the last lecture

- ▶ Next week on Wednesday 21 Dec from 8:15 to 10:00 in MXF1 (here) you can ask questions to TAs and myself.
- ▶ There is no exercise session on Friday 23 Dec.

Final Exam

1. For the written exam in January one page A4 handwritten notes are allowed (**ONE SIDE**, not both).
 1. You can write notes on a tablet and print them.
 2. You hand in your notes with the exam (keep a copy if you want).
 3. Duplicates are not allowed (for two or more students having the same handwritten notes we subtract points from the exam).
 4. Your notes should be written sufficiently large that they can be read with normal eye-sight and without magnifying lenses.
2. The written exam consists of multiple choice questions in the style of the quizzes and exercises in the style of the conceptual exercises.

Examples of Reinforcement Learning

Examples: Theorem Proving

input: mathematical axioms and theorems (goals)

```
cnf(sos01_axiom, ( product(A,A) = A)).
cnf(sos02_axiom, ( product(A,product(B,C)) = product(product(A,B),product(A,C)))).
cnf(goal1_negated_conjecture, ( product(product(product(x0,x1),x1),product(x0,x1),product(x0,x1),product(x1,x0),x2)))).
```

desired output: steps to prove the theorem
learning task: learn solving strategies by trial-and-error on axioms and theorems.

<http://www.tptp.org/>, <https://deepmind.com/research/training-a-first-order-theorem-prover-from-scratch>

Examples: Advertisement


input: user profile (currently viewed page & history)



desired output: attractive suggestions
learning task: learn by trial-and-error to display the suggestions that users will select with high probability.

Examples: Games

input: rules of a game



desired output: winning policy

learning task: learn winning strategies by trial-and-error from many games of the computer playing against itself.

<https://arxiv.org/abs/1911.08265>

EPFL

What is Machine Learning?
○○○

Examples
○○○○○○○○○○○○○○○○○○○○

Machine Learning: T
○○○

○○○○○○○○○○○○○○○○○○○○

Machine Learning: Tasks and Goals
○○○

Quiz
○○

Organization of This Course
○○○○○○

EPFL

What is Machine Learning?
○○○

Examples
○○○○○○○○○○○○○○○○○○○○

Machine Learning: Tasks and Goals
○○○

Quiz
○○

Organization of This Course
○○○○○○

Learning by Trial-and-Error


Table of Contents

1. A Toy Problem: Chasse au Trésor
2. Q-Learning
3. Science and Wisdom
4. Function Approximation and Deep Reinforcement Learning
5. Reinforcement Learning and the Brain
6. Reinforcement Learning for Efficient Planning


A Toy Problem: Chasse au Trésor




state 1




state 2



state 3



state 4



state 5



state 6




state 7

actions: (in each room)

open left door,
open right door

rewards: (depend on
state and action)
between -5 and 6

Agents and Environments: States, Actions, Rewards



The **agent** (a person, an animal, a computer program) observes at time t state S_t , reward R_t (for previous actions) and takes action A_t .

The **environment** (a game, a dynamical system, “the world”) receives A_t and produces next state S_{t+1} and reward R_{t+1} .

- ▶ The dynamics of the environment can be stochastic, e.g. given by **transition probabilities** $P(S_{t+1}|S_t, A_t)$ and **reward probabilities** $P(R_{t+1}|S_t, A_t, S_{t+1})$.
- ▶ Usually, the agent starts with zero knowledge about the transition and reward probabilities.
- ▶ The agents goal is to maximize the expected cumulative reward.

Action Values

Action Values $Q(S_t, A_t)$ indicate the desirability of action A_t in state S_t by measuring the expected cumulative reward. They are also called **Q Values**.

Finite Horizon until T (episodic setting)

$$Q^{\text{exact}}(S_t, A_t) = E [R_{t+1} + R_{t+2} + \dots + R_T]$$

Infinite Horizon with Discount Factor $\gamma \in [0, 1)$ (continual setting)

$$Q^{\text{exact}}(S_t, A_t) = E [R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \gamma^3 R_{t+4} + \dots]$$

The expectation depends on the policy (action selection strategy) and the transition and reward probabilities. Because they are usually unknown to the RL agent, the agent cannot exactly compute these expectation but has to estimate them.

From now on we will use the symbol Q to denote an RL agent's estimate of Q^{exact} .

Learning Action Values with Monte Carlo Estimation

Compute the cumulative reward for every state-action visited in each episode (length T).
Average the result over all episodes where the same state-action pair was visited.

```
1:  $Q(s, a)$ : arbitrarily initialized
2: CumulativeRewards( $s, a$ ): an empty list
3: for all episodes do
4:   for  $t = 1, \dots, T - 1$  do
5:      $G = \sum_{i=t+1}^T R_i$ 
6:     append  $G$  to CumulativeRewards( $S_t, A_t$ )
7:      $Q(S_t, A_t) = \text{average}(\text{CumulativeRewards}(S_t, A_t))$ 
8:   end for
9: end for
10: return  $Q$ 
```

- ▶ The estimated Q -values depend on the actions we take!
- ▶ This algorithm can be optimized by using a recursive updates (see notebook).

Policies: Exploration and Exploitation

- ▶ A **policy** is a mapping from perceived states s to actions a to be taken when in those states.
- ▶ A policy can be a **deterministic** function $a = \pi(s)$.
- ▶ In general a policy is **stochastic** and described by conditional probabilities $\pi = P(a|s)$.
- ▶ The policy serves two goals:
 1. **Exploitation**: Choose the action that is currently assumed to be the one that maximizes the cumulative return.
 2. **Exploration**: Choose exploratory actions to learn more about the environment and know better which action is actually best.

This is also called the **Exploration-Exploitation Dilemma** (or Trade-Off)

ϵ -Greedy Policies

- ▶ **ϵ -Greedy Policies** choose with probability $1 - \epsilon$ the (supposedly) best action for the current state (according to e.g. the current estimate of the Q -values) and with probability ϵ a random action.
- ▶ The best action is also called **greedy**; the random action **exploratory**.
- ▶ With $\epsilon = 1$ the agent only explores.
- ▶ With $\epsilon = 0$ the agent only exploits. This is the **greedy policy**.
- ▶ ϵ is a critical hyper-parameter affecting speed of learning! Common choices for simple problems are $\epsilon = 0.1$ at the beginning of learning.
- ▶ Over the course of learning, the agent gets more and more certain about the best actions; therefore one can gradually decrease ϵ , i.e. exploit more.

Summary

Key Ingredients of Reinforcement Learning

- ▶ An **agent** performs **actions** A_t according to some **policy** in an **environment** and perceives **states** S_t and **rewards** R_t .
- ▶ The agent should choose a policy that trades off **exploitation** (acquire as much reward as possible) and **exploration** (learn more about potentially more rewarding parts of the environment).
- ▶ For exploitation the agent can rely on estimated **action values** $Q(s, a)$ that indicate the **expected cumulative reward** of action a in state s . The action values can be estimated with **Monte Carlo Estimation** (among many other methods not yet discussed).
- ▶ For exploitation the agent can occasionally take a random action. This is formalized in **epsilon-greedy** policies (among other exploration strategies not yet discussed).

Supervised vs. Unsupervised vs. Reinforcement Learning

	supervised	unsupervised	reinforcement
given	X, Y	X	an environment (the agent collects data)
goal	find $P(Y X)$	find structure in data	find optimal policy
evaluation	test error	?	cumulative reward
approach	fit training data	use training data	interact with environment

Table of Contents

1. A Toy Problem: Chasse au Trésor
- 2. Q-Learning**
3. Science and Wisdom
4. Function Approximation and Deep Reinforcement Learning
5. Reinforcement Learning and the Brain
6. Reinforcement Learning for Efficient Planning

Q-Learning

Remember that action values are defined as

$$Q_{\pi}^{\text{exact}}(S_t, A_t) = E [R_{t+1} + R_{t+2} + \dots + R_T]$$

The subscript π indicates that $Q_{\pi}^{\text{exact}}(S_t, A_t)$ depends on policy π used for future actions.

The equation above can also be written in recursive form as

$$Q_{\pi}^{\text{exact}}(S_t, A_t) = E [R_{t+1} + Q_{\pi}^{\text{exact}}(S_{t+1}, A_{t+1})]$$

and for the optimal (greedy) policy π^* that takes in every state the action with maximal value

$$Q_{\pi^*}^{\text{exact}}(S_t, A_t) = E [R_{t+1} + \max_a Q_{\pi^*}^{\text{exact}}(S_{t+1}, a)]$$

Therefore

$$E [R_{t+1} + \max_a Q_{\pi^*}^{\text{exact}}(S_{t+1}, a) - Q_{\pi^*}^{\text{exact}}(S_t, A_t)] = 0$$

Q-Learning

$$E \left[R_{t+1} + \max_a Q_{\pi^*}^{\text{exact}}(S_{t+1}, a) - Q_{\pi^*}^{\text{exact}}(S_t, A_t) \right] = 0$$

Question: Is there a way to start with an arbitrary $Q_0(S_t, A_t)$ and define an update rule $Q_k(S_t, A_t) \rightarrow Q_{k+1}(S_t, A_t)$ that depends on every observed transition $S_t, A_t \rightarrow R_{t+1}, S_{t+1}$ such that $\lim_{k \rightarrow \infty} Q_k(S_t, A_t) = Q_{\pi^*}^{\text{exact}}(S_t, A_t)$?

Idea: The update rule should be such that the action value moves always in direction of the **Temporal Difference Error (TD-Error)**

$$\delta_k = R_{t+1} + \max_a Q_k(S_{t+1}, a) - Q_k(S_t, A_t)$$

i.e.

$$Q_{k+1}(S_t, A_t) = Q_k(S_t, A_t) + \lambda \delta_k$$

where λ is a learning rate, e.g. $\lambda = 0.1$.

Properties of Q-Learning

At convergence:

$$0 = E[\delta_k] = E \left[R_{t+1} + \max_a Q_k(S_{t+1}, a) - Q_k(S_t, A_t) \right]$$

i.e. Q-Learning stops when the action values Q_k satisfy the same equation as $Q_{\pi^*}^{\text{exact}}$ of the optimal greedy policy π^* .

Q-Learning is an **off-policy** method, because it estimates the Q-values for the greedy policy no matter what exploration policy is used; the Q-values of e.g. Monte Carlo Estimation depend on the exploration policy (Monte Carlo Estimation is an **on-policy** method).

Table of Contents

1. A Toy Problem: Chasse au Trésor
2. Q-Learning
- 3. Science and Wisdom**
4. Function Approximation and Deep Reinforcement Learning
5. Reinforcement Learning and the Brain
6. Reinforcement Learning for Efficient Planning

Science and Wisdom

”Science as a technique [...] makes possible, and even necessary, new forms of human society. [...] In as far as [man] is wise, this new power is beneficent; in so far as he is foolish it is quite the reverse. If, therefore, a **scientific** civilization is to be a **good** civilization it is necessary that increase in knowledge should be accompanied by **increase in wisdom**. I mean by wisdom a **right conception of the ends¹ of life**. This is something which science in itself does not provide.”

¹ends = goal, meaning

AI and Energy Consumption

activity	kWh
one google search	0.0003 – 0.001
one ChatGPT query	0.003 – 0.3
charging a smartphone once	0.01 – 0.02
heating 1L of water $\Delta T = 80$ deg C	0.08
energy produced by professional cyclist in 1h	0.2 – 0.45
1L diesel in a electric generator	3 – 3.5
NVIDIA A100 GPU for 1 day	5 – 9.6
100km with electric car	15 – 25
electricity consumption per year and capita in CH	6'000 – 7'000
training GPT-3 (estimated)	1'287'000

<https://www.sciencedirect.com/science/article/pii/S2542435123003653>

<https://www.eda.admin.ch/aboutswitzerland/en/home/wirtschaft/energie>

<https://cyclingbeginner.com/pro-cyclists-vs-amateurs-the-ultimate-power-comparison/>

<https://ev-database.org/cheatsheet/energy-consumption-electric-car>

<https://ai.stackexchange.com/questions/38970/how-much-energy-consumption-is-involved-in-chat-gpt-responses-being-generated>

SLAUGHTERBOTS ARE HERE.

The era in which algorithms decide who lives and who dies is upon us. We must act now to prohibit and regulate these weapons.

<https://autonomousweapons.org/>

Algorithmic Behavior Modification

<https://www.nature.com/articles/s42256-022-00475-7>

Human Compatible AI

- ▶ XAAS: everything (X) as a service. ²
- ▶ How does a less powerful entity, namely us, retain power forever over something that's more powerful than humans, namely a general purpose AI?
 - ▶ Standard model: hard-code an objective function for every task.
Problem: we are terrible at specifying good objectives (see e.g. king Midas). ³
 - ▶ Proposed solution:
 1. The AI's only objective is to satisfy human preferences.
 2. The AI is uncertain about what those human preferences are.
 3. Human behavior provides evidence of preferences.

²“If you're in a remote village and you need a school you just get out your phone, tap, tap, tap, 'we need a school' and the cost goes down to really the cost of land and raw materials, things that are not scalable in the same way that information and computation are scalable but you don't need the very expensive architects, construction engineers, permit experts, planners and so on you can do all these things at almost no cost.”

³“So let's say you watch videos of people playing video games and it just sends you one that's a little bit more violent, a little bit more violent, a little bit more violent. It's moving, you it's changing your taste and your preferences, it's moving you towards an extreme and at that extreme you're willing to keep watching for hours and hours and hours, whereas before, you wouldn't have. If it sent you the extreme video right away you'd say: 'oh this is horrible, I don't want to watch. I'm switching it off.' But it desensitizes you and it's learned how to move you in sort of person space to a different place where you wouldn't have wanted to be.”

<https://www.youtube.com/watch?v=NUwN3epacnk>

Table of Contents

1. A Toy Problem: Chasse au Trésor
2. Q-Learning
3. Science and Wisdom
- 4. Function Approximation and Deep Reinforcement Learning**
5. Reinforcement Learning and the Brain
6. Reinforcement Learning for Efficient Planning

Different State Representations

Enumerating all states is often not a good idea

- ▶ There can be a lot of states, e.g. chess has more than 10^{40} different positions.
- ▶ It is unclear how to generalize with enumerated states, e.g. we may learn quickly that it is not a good idea to open doors with a guard in front of it in our chase au trésor, but looking at the raw integers, state 3 is as different from state 5 as it is from state 4.

Alternatives: pixel images as input, or feature representation as input.

E.g. (in red room, in blue room, in green room, in treasure room, KO, guard present)

Function Approximation

- ▶ If the state s is a vector (instead of an integer) it does not make much sense to store the Q-values in a table.
- ▶ Instead, we can use a parametrized function family $q_{\theta}(s)$ to compute the Q-values.
- ▶ For a fixed θ the function q_{θ} takes the vector-valued state s as input and returns a vector of Q-values for each action; $q_{\theta}(s)_a$ is the Q-value of state s and action a .
- ▶ The parameters θ could be the weights of neural network.
- ▶ The correct Q-values can be learned by adjusting the parameters θ .

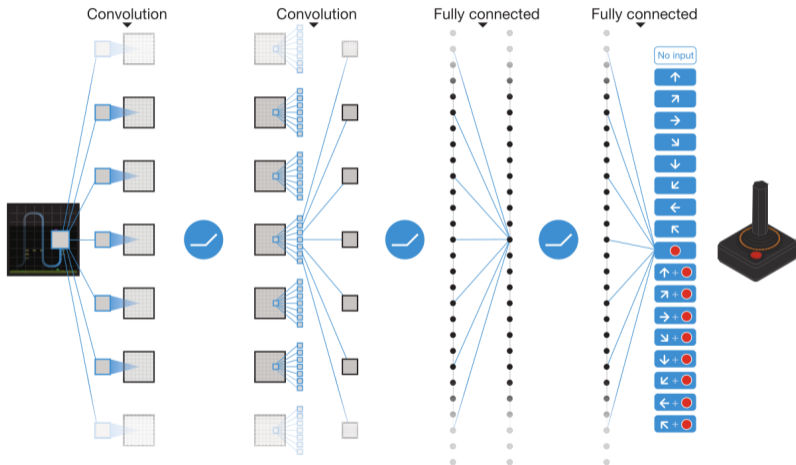
Deep Q Network

Main Idea: For observed transition S_t, A_t, S_{t+1} , adapt the parameters θ with gradient descent on the squared TD-Error loss function

$$\theta^{(k+1)} = \theta^{(k)} - \lambda \frac{\partial}{\partial \theta} \left(R_{t+1} + \max_a q_{\theta^{(k)}}(S_{t+1})_a - q_{\theta}(S_t)_{A_t} \right)^2$$

In practice many additional tricks are used...

Deep Q Network Learns to Play Atari Game



<https://www.deepmind.com/publications/human-level-control-through-deep-reinforcement-learning>

Deep Q Network Learns to Play Atari Game




<https://www.deepmind.com/publications/human-level-control-through-deep-reinforcement-learning>

AlphaStar: Mastering StarCraft II



<https://www.deepmind.com/blog/alphastar-mastering-the-real-time-strategy-game-starcraft-ii>

Solving Rubik's Cube with a Robot Hand



<https://openai.com/blog/solving-rubiks-cube/>

Table of Contents


1. A Toy Problem: Chasse au Trésor
2. Q-Learning
3. Science and Wisdom
4. Function Approximation and Deep Reinforcement Learning
- 5. Reinforcement Learning and the Brain**
6. Reinforcement Learning for Efficient Planning

Does Dopamine Signal Reward Prediction Errors?


“When monkeys receive unexpected reward, dopamine neurons fire a burst of action potentials. If the monkeys learn to expect reward, that same reward no longer triggers a dopamine response. Finally, if an expected reward is omitted, dopamine neurons pause their firing at the exact moment reward is expected.”

<https://doi.org/10.1146/annurev-neuro-072116-031109>


Goal-Directed versus Habitual Behaviour



(b) Forward model (goal-directed)



(c) Caching (habitual)
(trained hungry)



It is hypothesized that animals and humans rely on a model-based reinforcement learning system that allows goal-directed planning and a model-free reinforcement learning system that forms habits.

The model-free system is fast and computationally cheap. Planning with the model-based system is slower and computationally more expensive but it allows to “try out things in the mind” and therefore propose creative solutions with potentially less trial-and-error.

<http://dx.doi.org/10.1016/j.tics.2006.06.010>

Table of Contents

1. A Toy Problem: Chasse au Trésor
2. Q-Learning
3. Science and Wisdom
4. Function Approximation and Deep Reinforcement Learning
5. Reinforcement Learning and the Brain
- 6. Reinforcement Learning for Efficient Planning**

Reinforcement Learning for Efficient Planning

- ▶ Tic-Tac-Toe could also be solved with exhaustive search (planning), e.g. with policy iteration, because the transitions and the rewards are known.
- ▶ Many problems have too many states to be fully solved by exhaustive planning, e.g. Chess or Go (approximately 10^{170} states).
- ▶ Also model-free reinforcement learning would be inefficient in these cases, because it takes very long to learn an accurate policy for so many states.
- ▶ Instead one can use an inaccurate policy to guide local planning, and focus on the most promising actions starting from the current position.

AlphaZero: Mastering Two-Player Games



<https://www.deepmind.com/research/highlighted-research/alphago>

Application: Planning to Avoid Obstacles with a Tesla Car

Example Parking Problem

Search Space: 5 Constant Curvature Arcs			
Algorithm	A*	A*	MCTS Argmax Sampling
Search Heuristic	Euclidean Distance to Goal	Euclidean + Navigation	Neural Network Policy & Value Function
Number of Expansions	398,320	22,224	288

Search Expansions: 0
Complex Simplified Ego

T E S L A LIVE

<https://youtu.be/j0z4FweCy4M>

Summary

- ▶ Reinforcement Learning agents learn by **trial-and-error**.
- ▶ **Exploration** is achieved by taking actions that differ from the currently assumed best one.
- ▶ Even without learning an explicit transition and reward model, agents can learn to adapt to anticipated future needs (**model-free RL**), e.g. with Monte Carlo Estimation or Q-Learning.
- ▶ When the state space is too large for tabular RL, function approximation is useful for generalization (**deep RL**).
- ▶ When the transition dynamics and the reward structure is known or well estimated (**model-based RL**), planning (e.g. policy iteration) can be used.
- ▶ When the search space is too large for classic planning methods, RL can be used to learn a good **search heuristic**.

Resources

This lecture is inspired by the excellent text book

Reinforcement Learning: An Introduction

Richard S. Sutton and Andrew G. Barto

Second Edition

<http://incompleteideas.net/book/the-book-2nd.html>

You may also like the great online educational resource

<https://spinningup.openai.com>