Artificial Neural Networks: Deep Learning and Reinforcement Learning

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Overview of class

We start at 10h15

Previous 3slides.
Every week the first two slides contain the contents and main objectives of the day.

Normally, the teaching term at EPFL has fourteen weeks.
However, Friday before the Easter break (Holy Friday) is a holiday.
Moreover, the class on Friday after ‘Ascension’ is dropped and replaced by somewhat longer sessions each week.

Deep reinforcement learning

Chess

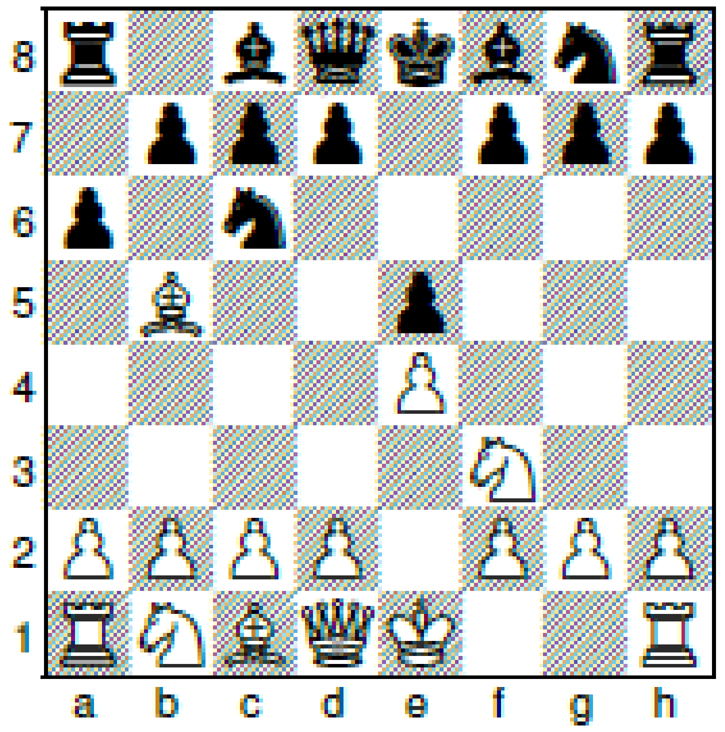

Go

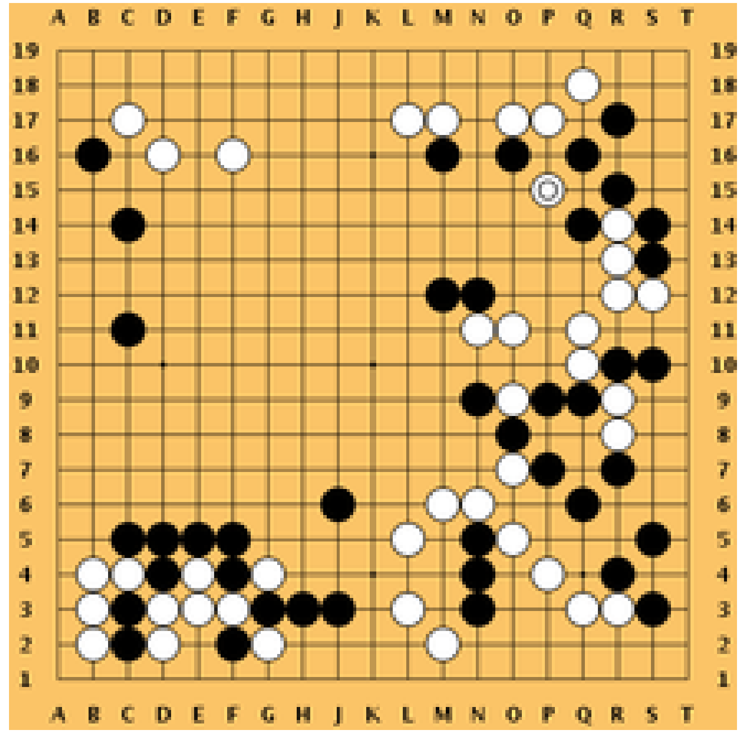

Artificial neural network (*AlphaZero*) discovers different strategies by playing against itself.

In Go, it beats Lee Sedol

Previous slide.
Overall aim: understand modern methods such as deep reinforcement learning.

In a game such a Chess or Go, the reward signal is only given once at the very end of the game: positive reward if the game is won, and negative reward if it is lost.

This rather sparse reward information is sufficient to train an artificial neural network to a level where it can win against grand masters in chess or Go.

To improve performance, each network plays against a copy of itself. By doing so it discovers good strategies (such as openings in chess).

Artificial Neural Networks

Wulfram Gerstner

EPFL, Lausanne, Switzerland

1. Simple perceptrons for classification
2. Reinforcement learning1: Bellman and SARSA
3. Reinforcement learning2: variants of SARSA
4. Reinforcement learning3: Policy Gradient
5. Deep Networks1: Backprop and multilayer perceptron
6. Deep Networks2: Statistical Classification by deep networks
7. Deep Networks3: regularization and tricks of the trade
8. Deep Networks4: Convolutional networks
9. Deep Networks5: Error landscape and optimization methods
10. Deep Reinforcement learning1
11. Deep Reinforcement learning2

Previous slide.
This course has a focus on Reinforcement Learning.

Overall there will be 12 sessions for 14 weeks.

Friday after ascension is no class.
Instead each week there will be 100 minutes of lectures

-image classification by Deep Netwroks (5 lessons)
-reinforcement learning (5 lessons)

.

Miniprojects (MP): we use software package ‘Keras’

- hand in 1 (not 2) out of 2 projects
- graded on a scale of 1-6
- grade of MP counts 30% toward final grade
- we do **fraud detection interviews**
- MP done in groups of two students (not alone)
- interview for MP is in last week of classes or first week after end of classes

→ **plan ahead!!**

Hand-out of miniprojects: March 31

Hand-in of miniprojects: 1) May 31 → interview June3/4

2) June 7 → interview June 10/11

Written exam:

- counts 70 percent toward final grade
- 1 page A5 double-sided handwritten notes, but no other tools allowed
  (no calculator, no cell phone, no slides, no book)
- ‘mathy’, similar to exercises

Written exam is ‘orthogonal’ to miniproject:
we ask for different things (theory)

Every week we have videos:

Lot’s of material:
→ good idea to solve the first exercise after half the lecture

- exercise sessions: Monday 5pm, Tuesday 5pm, Friday 9am.

Lectures are a bit longer than 90 minutes because there is no class on Friday after ‘ascension’.

TA’s this year:
Berfin Simsek (HeadTA), Alireza Modirshanechi (HeadTA)
Ihor Kuras, Ana Stanojevic, Daniil Dmitriev

Previous slide.

For the in-class exercises it is important that you really try to solve them. No problem if you fail (some exercises are harder than others). But it is important that you start to think about how you would solve the exercise.

Since there is no class on the Friday after ascension, pure lecturing time is around 100 minutes per week

Artificial Neural Networks

Wulfram Gerstner

EPFL, Lausanne, Switzerland

- The math is developed on the blackboard
- There are no written course notes!!
- All of the contents are **standard textbook material**

Choose a textbook that you like! I recommend

For **reinforcement learning lectures:**

- *Reinforcement learning,* R. Sutton+ A. Barto (2nd ed, online)

For **supervised learning lectures**:

- *Pattern Recognition and Machine Learning*, C.M Bishop, 2006
- *Neural Networks for Pattern Recognition*, C.M. Bishop, 1995
- *Deep Learning*, Ian Goodfellow et al., 2017 (also online)

Also good: *Neural networks and learning machines*, S. Haykin

Previous slide.

Work with a textbook that you like. If you study at home, slides are not sufficient.

The books of
Goodfellow et al.
Sutton and Barto
are the basis of the class. Both are available online in pdf format as preprints for free.

Artificial Neural Networks

Prerequisits:

CS433, Machine Learning

(Profs Jaggi+Urbanke)

Rules:

If you have taken this class: please ask many questions

If you have not taken this class: please do not complain

Previous slide.

The overlap with the class of Jaggi+Urbanke is minimal (main overlap for ‘regularization’). But we need quite a few of their results as a basis!

Some students have taken a very similar class and then this is also fine.

Students who did not take the above class (or something very similar) are not admitted to the class ‘Artificial Neural Networks’. If they attend, it is at their own risk; they should not ask questions, but fill the knowledge gaps on their own. They should not complain if they find the class too hard.

Artificial Neural Networks

Learning outcomes:

- apply learning in deep networks to real data
- assess/evaluate performance of learning algorithms
- elaborate relations between different mathematical concepts of learning
- judge limitations of learning algorithms
- propose models for learning in deep networks

Transversal skills:

- access and evaluate appropriate sources of information
- manage priorities
- work through difficulties, write a technical report

Previous slide.

Access and evaluate appropriate sources of information
→ this means: you should learn to read textbooks. It is not sufficient to just look at slides.

Manage priorities
→ this means: the miniprojects only count 30 percent. Don’t write a program with bells and whistles, but really focus on the things you are asked to do.

work through difficulties,
→ this means: some things will look hard at the beginning, be it in the miniproject or in the mathematical calculations. That’s normal, but you have to work through this.

write a technical report
→ this means: we would like to receive a readable technical report for the miniprojects. Concise, to the point, not too long.

Artificial Neural Networks

Work load:

5 credit course → 7.5 hours per week for 18 weeks

(this count includes the 2 hours of lecture on Friday)

1 ECTS = 27 hours of work

Previous slide.

Including exam preparation, the term has 18 weeks for 12 weeks of lectures:
The week of ascension, and easter, and exam preparation time counts as well.

The statement made by a student that

‘An exercise session of 45 minutes is not enough to solve all the exercises’

is correct. You need additional time at home to solve the exercises. Solving the exercises is a good preparation for the exam and necessary to understand the mathy parts of the class.

Two ways to study for this class

A: Self-paced self-study

1. Read slides 1+2 each week (objectives and reading)
2. Start exercise n.
3. If stuck, read book chapter Return to 2.
4. n←n+1
5. Compare with solutions
6. Do quizzes in slides (yellow pages)

Hand-in miniproject.

Note: Slides are not meant for self-study. Use textbook!

B: Lecture-based weekly

1. Follow lecture videos
   - annotate slides
   - participate in quizzes
   - try to solve all exercises

2. Go to Exercise session
3. redo exercises and Compare with solutions.

Hand-in miniproject.

Note: Do not forget to annotate slides so that you can use them.

Previous slide.

You don’t need to come to class, since all material is textbook material. But then you really have to study the textbooks!
Slides are not meant to replace textbooks.
Slides are self-contained under the assumption that you attend class and exercise sessions.

For the final exam, it is very important that you worked through all the exercises.

Sample examples from the two previous years are online: have a look before you decide to take the class.

Your Semester planning

Wulfram Gerstner
EPFL, Lausanne, Switzerland

The course **‘Deep Learning’** (Fleuret)
and the course **‘Artificial Neural Networks’** (Gerstner)
have about 20-30 percent overlap.
You can take either one or the other or both (OR),
students consider the course of Prof. Fleuret as
‘more practical coding-oriented’ than this one here.

The course **‘Unsupervised and Reinforcement L.’** (not given)
and the course **‘Artificial Neural Networks’** (Gerstner)
have about 20-30 percent overlap (three weeks)
I suggest to take either one or the other or both;
The other course is more ‘biological’ than this one here

Previous slide.
The class ‘Deep learning’ also treats backpropagation, tricks of the trade, convolutional networks. It does not contain any reinforcement learning.

The class ‘Unsupervised and Reinforcement learning’ also treats Reinforcement learning. It does not contain any supervised learning for classification, no backpropagation. Reinforcement learning is discussed with a biological focus.
The class is not given in 2019.

The class ‘Artificial Neural Networks’ is planned for IC students who have already taken the class ‘Machine Learning’ by Jaggi-Urbanke.

The class ‘Deep Learning’ is planned for STI students and does not have any prerequisits (except engineering bachelor)

Your semester planning

Wulfram Gerstner
EPFL, Lausanne, Switzerland

The course **‘Deep Learning’** (Fleuret)
and the course **‘Unsupervised and Reinforcemnt L.’** (Gewaltig)
have less than 5 percent overlap.
You can take one or the other or both (OR).

+The course **‘Unsupervised and Reinforcement L.’** (Gewaltig)
is oriented towards biological questions, aimed at SV students
+The course **‘Deep Learning’** (Fleuret) is an applied course.
It has no prerequisits and does not cover reinforcement learning.
Aimed at STI students
+ The course ‘Artificial Neural Networks’ (Gerstner) is a course
aimed at IC students. **Prerequisit: Machine Learning** (Jaggi)

Previous slide.
The course ANN starting now is aimed for IC students.

Quiz: Classification versus Reinforcement Learning

[ ] Classification aims at predicting the correct category such as ‘car’ or ‘dog’
[ ] Classification is based on rewards
[ ] Reinforcement learning is based on rewards
[ ] Reinforcement learning aims at optimal action choices

Your notes:

Quizzes appear at the end of about half of the videos.

Reading for this week:

Bishop, Ch. 4.1.7 of

Pattern recognition and Machine Learning

or

Bishop, Ch. 3.1-3.5 of

Neural networks for pattern recognition

Goodfellow et al., Ch. 1 of

Deep Learning

Sutton and Barto, Ch. 1.1 and 1.2 of

Reinforcement Learning

Previous slide.

The suggested reading is important, in particular if you are not able to attend the class in a given week.

In all the following weeks, the suggested reading will always be listed on slide 2, at the beginning of the lecture, so that it is easy to find.

If you have understood everything, and are able to solve the exercises, then you do not have to go through the reading.

Artificial Neural Networks: Lecture 1
Simple Perceptrons for Classification

Wulfram Gerstner
EPFL, Lausanne, Switzerland

Objectives for today:

- supervised learning vs. reinforcement learning
- understand classification as a geometrical problem
- discriminant function of classification
- linear versus nonlinear discriminant function
- perceptron algorithm
- gradient descent for simple perceptrons
- geometric interpretation of learning

Recommended exam preparation

(1) do (or redo) **exercises** yourself

(2) if stuck, read the relevant chapter of the **textbook** (see page 2 of slides of each week)

(3) check the solution of exercise

(4) look at the **quiz question** (always orange slides)

(5) if stuck, read the relevant chapter of the **textbook** (see page 2 of slides of each week)

(6) Look at **past exams** (solutions: see analog exercises)

NOTE: the slides are most useful if you have followed and annotated them yourself during the lecture.

This is what successful students said about exam preparation:

Student A:

“For me, going through the exercises was very helpful, along with the slide quizzes. We also discussed theoretical questions from the lectures with my teammate and friends”

Student B:

“During the semester I have read the commented version of the slides in order to carry out the 2 miniprojects. I took care to understand each remark and I did the exercises when I had trouble in learning a topic. Before the exam, I felt that I was remembering well so I could focus only on Reinforcement Learning. In this case I found more useful solving the exercises to understand some key differences between the different algorithms e.g. off-policy versus on-policy.”

This is what successful students said about exam preparation:

Student C:

« I first went through all the lecture slides which I had taken notes on during lectures to reinforce my memory of various notions introduced in this course, and I want to stress that the comment pages were truly helpful. Afterwards, I went over all the exercises and collected a few questions to pose in the revision session held by TAs and got satisfactory clarification for most of them. »

Student D:

« I prepared for the exam by reading slides over and over again. I think the comments slides helped me a lot in understanding and reading them over again helped me to build the structure of the overall course.
Exercises helped as well since it turns out that the exam is quite similar to exercises. »

This is what successful students said about exam preparation:

Student E:

“I attended nearly every class and made sure I understood the blackboard proofs properly because these were usually very useful for understanding the main concepts. During the exam preparation, I mostly just went through the class slides again and solved all of the exercises.”

Student F:

“I never came to class but I did all the exercises and studied the books on Reinforcement Learning and Deep Learning.”

Questions?

… before we start

‘MAGMA LEARNING’: EPFL EdTech start-up offers a Learning App linked to this class:
accessed by signing up with the code “cs456”
on https://www.ari9000.com or on the mobile app ARI 9000 (iPhone and Android).
Quizzes and link to Wikipedia, extracted from class material.