Artificial Neural Networks: Deep Learning and Reinforcement Learning

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Overview of class

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
12. Deep Reinforcement learning 3
13. Recurrent Neural Networks and Sequences

Miniprojects (MP): we support you with software package ‘Keras’ (if you are familiar with PyTorch, you can also use PyTorch)

- MP is on Reinforcement Learning
- 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!!**

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)

In-Class Exercises: Typical Tuesday

11h15 – 12h45 lecture - exercise - lecture

14h15 - 14h56 lecture

15h00 - 16h00

TA's this year:

Berfin Simsek (HeadTA), Alireza Modirshanechi (HeadTA),
Ana Stanojevic (miniprojects), Martin Barry, + student TA (AE)

About 90 Short Videos, each about 10 min

https://lcnwww.epfl.ch/gerstner/VideoLecturesANN-Gerstner.html

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

Artificial Neural Networks

Prerequisits:

CS433, Machine Learning

(Profs Jaggi+Flammarion)

Rules:

If you have taken CS433: you are well prepared

If you have taken a similar other class: same

If you have not taken this class: you can take my class, please do not complain

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

Artificial Neural Networks

Work load:

5 credit course → 7.5 hours per week for 18 weeks

(this count includes the 4 weeks of exam preparation)

1 ECTS = 27 hours of work

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. Compare with solutions
5. n←n+1
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
   - 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.

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.

- Reinforcement Learning is not covert by Prof. Fleuret

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

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

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