Artificial Neural Networks/Reinforcement Learning

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Overview of class

Course WEB page: Moodle

Video Lectures:
Several short videos (10-20 min each) on RL
https://www.youtube.com/@gerstnerlab/playlists

Your notes:

Normally each visible slide is followed by a hidden comment slide.

Overview (Version 2023)

1. Intro and RL1: Reinforcement Learning for Bandit problems
2. RL1: Bellman Equation and SARSA
3. RL2: Q-Learning, n-step TD learning, continuous space, eligibility traces
4. RL3: Policy gradient algorithms

Basic RL

5. DL1: BackProp, Multilayer Networks and Automatic Differentiation
6. DL2: Tricks of the Trade in Deep Learning
7. DL3: Loss Landscape and optimization methods for Deep Networks

Detour Deep L.

8. Deep RL1: DeepQ, Actor-Critic, Eligibility traces from Policy gradient,
9. Deep RL 2 Inductive Bias, No Free Lunch, Model-free versus Model-based RL
10. Deep RL3: Discrete Games, Replay Buffer, and Continuous control
11. Deep RL4: Model-based Deep RL
12. Deep RL5: Exploration by Novelty/Surprise/InformationGain

Deep RL

13. Application: Biology and RL, three-factor rules
14. Application: Hardware, energy consumption and three factor rules

Interdisciplinary RL

Changes compared to the year 2022 (after student feedback 2022)

1. Intro and Perceptron Algo (Drop this, and start with RL1)

2. RL1: Reinforcement Learning and SARSA
3. RL2: TD learning, continuous space, eligibility traces
4. RL3: Policy gradient algorithms

5. DL1: BackProp and Regularization (Drop regularization and add Automatic Differentation from DL5)
6. DL2: Tricks of the Trade
7. DL3: Loss Landscape and optimization methods
8. DL4: Statistical Classification by Neural Networks (DROP THIS)
9. DL5: Convolutional Networks (DROP THIS) , move Automatic Diff. and NoFreeLunch (No Free Lunch moves to ?

10. Deep RL1: DeepQ, Actor-Critic, Eligibility traces from Policy gradient, Model-based RL
11. Deep RL2: Discrete Games, Replay Buffer, and Continuous control
12. Deep RL3: Model-based Deep RL
13. Deep RL4: Biology and RL, three-factor rules
14. Deep RL5: Hardware

15. Deep RL6: Exploration by Novelty/Surprise/InformationGain
16. Review:

Student Feedback

For expectation values, I will be more precise

'The class is not formal enough: Derivations should be more formal'

'The class is not practical enough: Need mo

Simple examples, coding of toy model (in python)

Q1: Given my background I perceive the class as theoretical/applied:

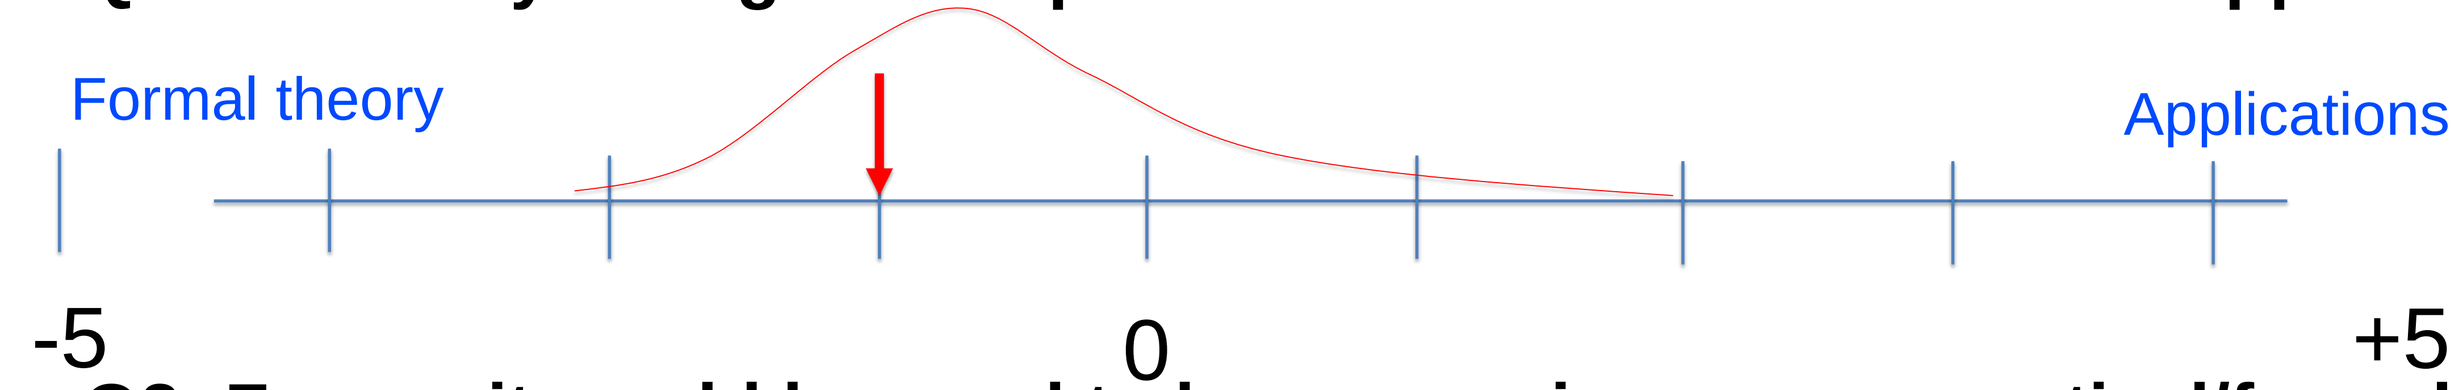

Q2: For me it would be good to have exercises more practical/formal

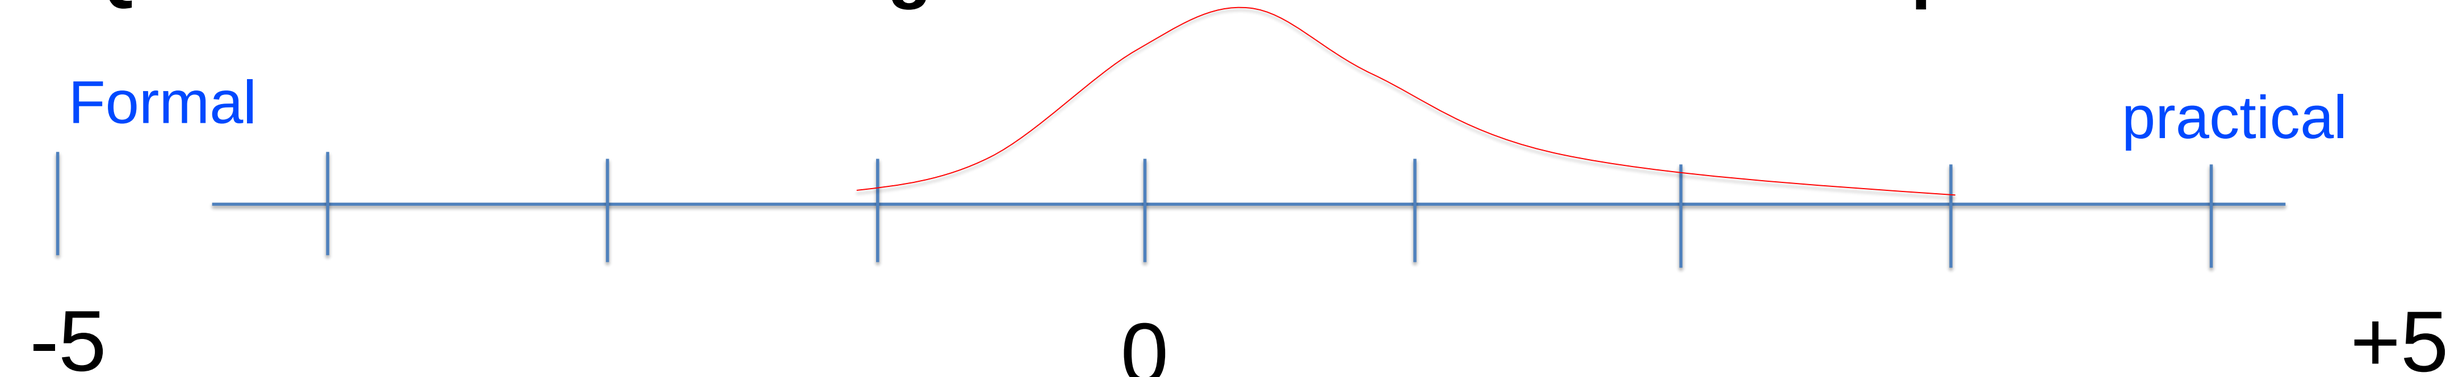

Contents/Overlap (citation from student feedback)

‘I am taking both Artificial neural networks (ANN) and Deep learning (DL) this semester, after Machine Learning (ML) last semester. There are multiple overlaps, including not exhaustively:

1. Perceptron and derivation of backpropagation, regularization (ML, DL, ANN)
2. Dropout, weight initialization, Data augmentation, Vanishing Gradient (DL, ANN)
3. Momentum, ADAM (DL, ANN)
4. Classification from statistical perspective (ML, DL, ANN)
5. Convolutional Neural Networks (ML, DL, ANN)

I have overall very mixed feelings about these overlaps, which I'll try to disentangle below:

Some of them, in particular point 2. and 3., are really nice complements to DL. I feel that here the concepts are just better explained, providing a real intuition of what is happening, and I would definitely keep them.

Some others, in particular point 4. and 5., are really a repetition of the same things over and over for all 3 courses’

Miniprojects (MP): we support you with 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 – 12h00 lecture1 -
12h05 – 12h50 exercise
LUNCH BREAK
14h15 - 14h00 lecture2
15h15 - 16h00 exercise

Results of **Exercises with *** are needed for lecture 2

TA's this year:

Alireza Modirshanechi (HeadTA),
Ariane Delroq (TA)
Shuqi Wang
Sophia Becker
student TAs (AE)

About 90 Short Videos, each about 10 min

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

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.

Artificial Neural Networks/RL

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 **deep learning** (only three lectures):

- *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/RL

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 if I refer to material
that you have not seen before.

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/Reinforcement Learning

Learning outcomes:

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

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

Previous slide.

Including exam preparation, the term has 18 weeks for 14 weeks of lectures:

The statement made by a student in an official evaluation 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. 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. 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.

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 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, not given in 2023) 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

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

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

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

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 most sections.
The exam contains a section with similar multiple-choice questions

Reading for this week:

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/RL: Lecture 10

XXX

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Objectives for today:

-

-

Your notes:

Objectives are not a list of contents but refer to learning outcomes/insights.

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?

…

First ‘real exercise’ starts now

→ Exercise 1 and 3 now!!!

First ‘real lecture’ starts at 14h15.