Artificial Neural Networks/Reinforcement Learning

Exam preparation/ Student feedback

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Course WEB page: Moodle

For most weeks there are Video Lectures:
Several short videos (10-20 min each) on RL

https://www.youtube.com/@gerstnerlab/playlists

The class has changed this year: clear focus on RL

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’

My conclusion: focus on RL!

Planned 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

Actual 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, eligibility traces
4. RL3: Continuous input, function approximation, inductive bias
5. RL4: Policy gradient algorithms
6. RL5: Advantage Actor-critic, eligibility traces, model-free/model-based

Foundations of RL

7. Deep RL1: Applications of Model-free Deep RL
8. Deep RL2: Applications of Model-based Deep RL
9. Deep RL 3 Markov Decision Proesses and Policy iteration.

Deep RL

10. RL and the Brain: Three-factor rules
11. RL and hardware
12. RL and internal rewards: novelty and surprise
13. RL and curiosity-driven exploration
14. Final Discussion

Inter-disciplinary RL

Miniprojects (MP):

- 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 (time slots after submission)**
- interview for MP is in last week of classes or first week after end of classes (depends on submission date)

→ **plan ahead!!**

Example: the project counts 30 percent, therefore
Miniproject 6.0 + exam 2.5 → 3.5 final average grade.

Miniproject 5.5 + exam 4.0 → 4.5 final average grade.

Written exam:

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

If there is an exercise on topic x, then there can be an exam question

Style of exam, see: → exercises
→ past exams (e.g. 2022)

Written exam is 'orthogonal' to miniproject:
we ask for different things (theory).

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

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

Previous slide.

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

The book of
Sutton and Barto
is the basis of the class. It is available online in pdf format as preprints for free.

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/Reinforcement Learning

Exam questions typically include:

- TD-learning/Bellman equation (derivation/proof sketch)
- Policy gradient (derivation/proof sketch)
- Actor critic
- Surprise, novelty, exploration

→some simple points (e.g. plug-in of algos),
→some difficult points (e.g., proof)

Excluded this year (but important topics in previous years):

- BackProp
- momentum/loss function.

Previous slide.
I try to construct an exam such that

- All basic concepts are covered
- A representative subset of less basic/more advance concepts is covered
- Some points are easy to get
- Some points are harder to get

In past years, I also used Quizzes. Probably not in 2023.

In past years, I included a few exam questions where two different algorithms had to go through a few iterations: ‘plug-in numbers’ and predict outcome. I typically think of this an algo-debugging exercises (‘suppose you invent an algorithm and implement it. You want to check that it does what it should – a toy example is a good test.’)
In principle, these ‘plug-in’ exercises are meant to be easy points – but then many students complain that it is too stupid for an exam at the master level. I am not sure that I will include one of these this year (last year we did not), but I will also not exclude it.

Review: 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.

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 the lecture and taken notes.

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.”

Tuesday 30.5. 2023 at 11h15:

Special Q&A Session:
Clarifications for Lectures or Exercises

Please prepare questions!

In-depth feedback is now open (week 13+14)!

Indicative feedback (week 5)

Year	2022-2023
Course	Artificial neural networks/reinforcement learning
Questionnaire	Indicative feedback of teaching (since 2022-2023)
Nb Registered	233
Nb Answered	62

The running of the course enables my learning and an appropriate class climate

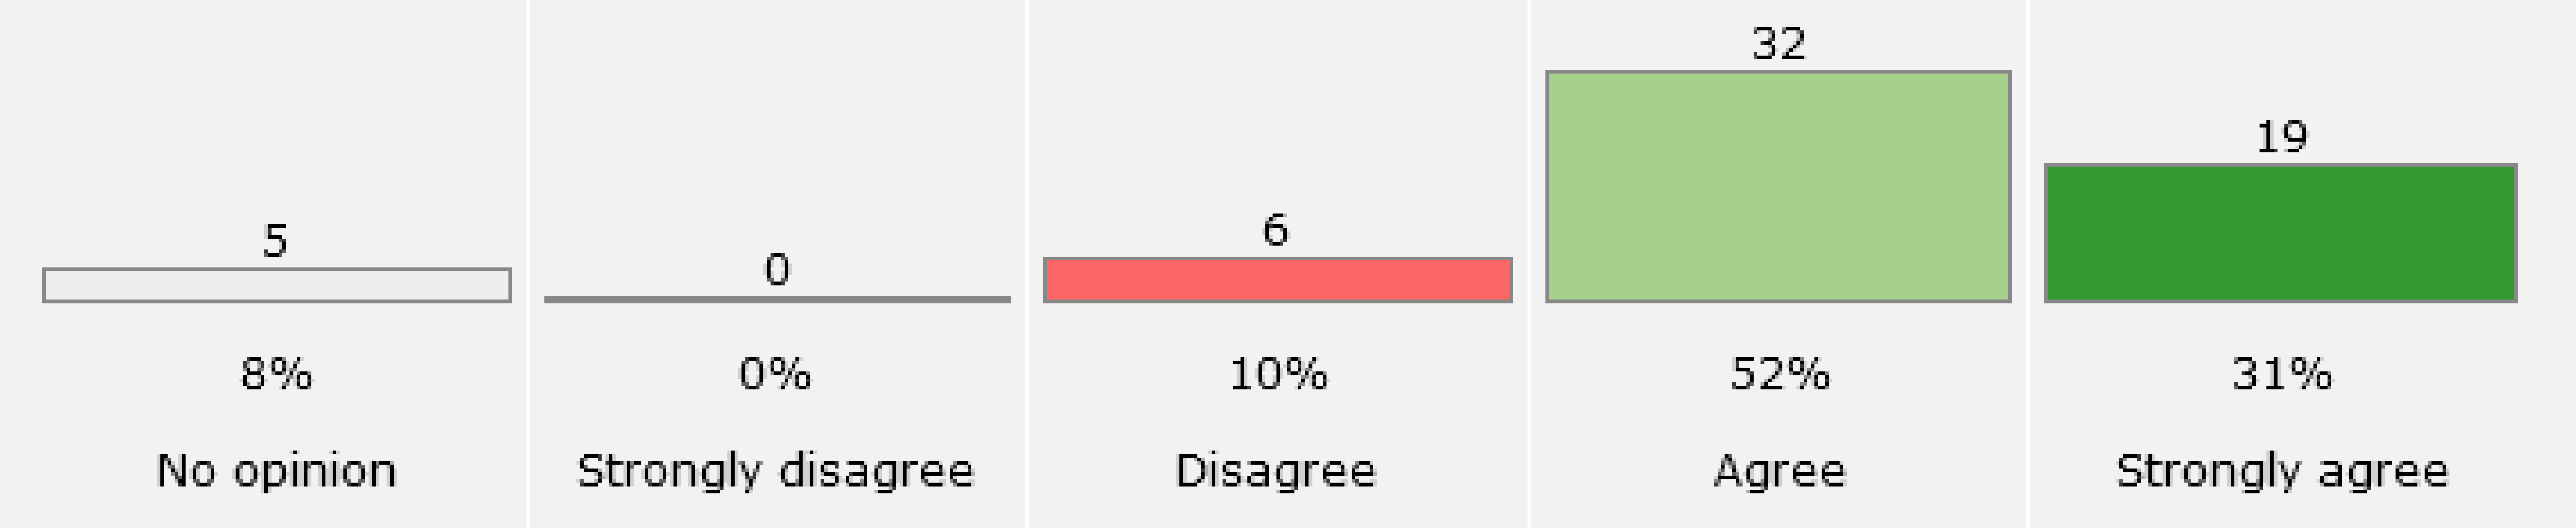

Student Feedback

Q1: ‘The class is not formal enough: Derivations should be more formal’

‘Q2: The class is not practical enough: Need more computer exercises’

Q3: Given my background I perceive the class as theoretical/applied:

Formal theory — Applications

-5 0 +5

Q4: For me it would be good to have exercises more practical/formal

Formal — practical

-5 0 +5

Student Feedback

Including Lectures, Exercises, and Miniproject,

[ ] I worked for this class less than 6 hours per week on average

[ ] I worked for this class less than 7-8 hours per week on average

[ ] I worked for this class less more 9 hours per week on average

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.

In-depth Student Feedback: Take time to do online-evaluation NOW!

If you write comments (you don't have to)
then, please, always start with:

- I usually came to class, based on this …
- I mainly followed the video lectures, based on this …