

Response to Student Feedback
received for the course CS-EJ3311
“Deep Learning with Python”
offered during fall 2020
at Aalto University and via FiTech <https://fitech.io/en/>

corresponding teacher: Alex Jung (first.last (at) aalto.fi)

September 3, 2021

We express our sincere gratitude for the insightful and constructive comments and suggestions provided by the students of the course **CS-EJ3311 “Deep Learning with Python”** offered during fall 2020. We try to address (most of) the comments with the design of the new course edition in fall 2021. Some major modifications are:

- We have slightly reorganized the course structure which is now given as
 - Round 0: Components of Machine Learning.
 - Round 1: Artificial Neural Networks.
 - Round 2: Gradient-Based Learning.
 - Round 3: Convolutional Neural Networks.
 - Round 4: Regularization.
 - Round 5: Natural Language Processing.
 - Round 6: Generative Adversarial Networks.
- The new course will include lecture sessions that provide background for and motivate the techniques covered in the notebooks.
- The Python notebooks now include student tasks that require small coding assignments.
- There will be no student project in the new course edition.
- We will randomly choose students at the end of the course who have to explain in detail their solutions to the coding assignments.

In what follows, we respond in a point-by-point fashion to a selection of the most relevant student feedback received for the “2020 Fall” edition of the course CS-EJ3311 “Deep Learning with Python”.

S1 Assignments: could require some own coding. maybe there could be some small part in each assignment that is not ready made, but clearly explained what to do.

A1 The main course material for the new course edition will again be in the form of Python notebooks. In contrast to the previous course edition, the notebooks now include small coding assignments (or student tasks). The student solutions to this coding assignments will be graded automatically as soon as the student submits her notebook to the JupyterHub server at <https://jupyter.cs.aalto.fi/>.

S2 Some questions in the quizzes were a bit confusing, but after all the information about correct answers was available somehow (e.g. from slack).

A2 We have carefully revised the quiz questions which are now included towards the end of the Python notebooks.

S3 I think it would be great to upload all lectures to Youtube, with password or something if not otherwise possible.

- A3 We will make the recordings of lectures available via my YouTube channel https://www.youtube.com/channel/UC_tW4Z_GfJ2WCnKDtwMuDUA.
- S4 *maybe including more modules about different types of DL algorithms*
- A4 We are not sure what is meant here by “DL algorithms”. Round 2 discusses gradient descent (and some of its variants) as maybe the most widely used method for training deep nets. We have also extended the rounds by including deep learning methods for natural language processing (Round 5) and data synthesizing (Round 6 on generative adversarial networks).
- S5 *The grading system for example i completed the Quiz 1 but I can not see the grade or I got zero because i did not submit the feedback.*
- A5 We now grade the student solutions automatically as soon as the student submits her notebook via the JupyterHub server at <https://jupyter.cs.aalto.fi/>. Using the JupyterHub server should ensure that students get feedback more timely (within few days).
- S6 *The course required no programming at all. It should have at least some programming tasks, like in ML with Python. It was now a too easy pass by just looking at materials and answering quizzes.*
- A6 see our response A1.
- S7 *The topics didnt go in to the theory that much.*
- A7 The new course edition includes some lectures on the theoretical underpinnings of the deep learning methods covered in this course. However, the main focus of this introductory course is to build useful intuition for the behaviour of deep learning methods and their implementation using the Python library Keras (<https://keras.io/>). For a more in-depth treatment of the mathematical principles of deep learning, we refer to more advanced courses such as CS-E4710-*Machine Learning: Supervised Methods* or CS-E4890-*Deep Learning*.
- S8 *It was a disappointment we did not actually write any code. The previous course Machine Learning with Python has been incredibly useful for me because I learned Python and have used it a lot since. This course on Deep learning was more of a nice to know..*
- A8 The format of new course edition of CS-EJ3311 Deep Learning with Python will be more similar to the format of CS-EJ3211 Machine Learning with Python. In particular it will also use notebooks with coding assignments (“student tasks”). After completing the coding assignments, students submit their notebooks to <https://jupyter.cs.aalto.fi>. See also our response A1.
- S10 *Towards the end of the course I also felt we were running out of content. It would have been great to learn a bit also about how to handle data sets that include text.*
- A10 The new course edition now includes Round 5 on natural language processing. This round will discuss deep learning methods for text data. For more details on natural language processing methods, we refer to the advanced courses CS-E4890-*Deep Learning* and ELEC-E5510-*Speech Recognition* (you can think of speech as text in audio form).
- S11 *Some type of QnA session where people ask what type of approach is a good candidate for a given problem might be good (I know one could do this via slack but very few did). Also I though some kind of pair programming in projects or peer mentoring could be good. There was not too much interaction between students. I personally had questions like why is my model currently getting abysmal results - is it the data, the approach or do I just have simple code errors that I cannot see that might have benefited from a pair of extra eyes.*

- A11 That is a good point. We will encourage students to make good use of the course discussion forum. A main purpose of this forum is to allow for uncomplicated interaction between students.
- S12 *It would be nice to have some basics of deep learning at the start. Especially developing a simple deep learning model from scratch without use of deep learning libraries.*
- A12 The new course edition now includes a recap of the basic principles and components used by deep learning methods (see Round 0). We have also revised the notebook for Round 1 (“Artificial Neural Networks”) such that it includes for details on the building blocks of artificial neural networks. For more details on the implementation (starting from low-level libraries for matrix computations) of deep learning algorithms, we refer to the more advanced course CS-E4890-*Deep Learning*.
- S13 *There were some typos/mistakes in the notebooks but I believe those were fixed. This course could have even more examples of different methods. Now this felt little too short.*
- A13 We have carefully revised the notebooks and tried to improve the use of language and clarity of presentation. These notebooks now also include coding assignments that should enforce a deeper understanding of the discussed methods. In the new course edition, we now also discuss deep learning methods for classifying text data (Round 5- NLP) and for synthesizing data (Round 6 - GAN).
- S14 *I would suggest more concentration on the selection of the Optimisers, more wide focusing on the interpretation of the trained model performance, how to make the model to learn the ordinal variables (ordinal classification).*
- A14 We have now expanded the discussion of optimizers in the Round 2 of the new course edition. More details on optimization methods for deep learning are provided in the advanced course CS-E4890-*Deep Learning* (<https://mycourses.aalto.fi/course/view.php?id=28212§ion=1>). Thank you for pointing out the application of deep learning methods to ordinal classification problems. We have now added a discussion on ordinal label values in Round 0 and in the course book mlbook.cs.aalto.fi.
- S15 *The course felt a bit too focused on image analysis although there are many applications also elsewhere, e.g. in text, audio and graphs. But I think this is very common in introductory courses to deep learning.*
- A15 The new course edition now includes deep learning methods for text data (Round 5 - NLP). We also note that deep learning methods for images might be useful also for audio data. Indeed, audio signals can be represented (visualized) using images of their time-frequency energy distributions (see <https://towardsdatascience.com/audio-classification-with-pre-trained-vgg-19-keras-bca55c2a0efe>). The more advanced course CS-E4890-*Deep Learning* (<https://mycourses.aalto.fi/course/view.php?id=28212§ion=1>) also covers deep learning methods for different data types including sequential (time-series) data and graph-structured data. You might also find the following advanced courses on speech and language processing useful: ELEC-E5550 - *Statistical Natural Language Processing*, ELEC-E5521 - *Speech and Language Processing Methods*, ELEC-E5510 - *Speech Recognition*.
- S16 *Personally I was hoping to do the project work to learn more, but didn't due to time constraints and as there was some peer-reviewing etc. involved and that felt like too much work. Multiple-choice questions are easy for course staff, but sometimes unclear in their meaning, and not very helpful for the student. Much better results can be gained by forcing the student to do some actual coding work and/or at least set up and try out some necessary software. I didn't even follow the course Slack as there was no reason to, so the course remained quite distant for me as a whole.*
- A16 In the new course edition we replace the student project with coding assignments that require students to implement deep learning methods by completing prepared code snippets. We will also try to better motivate

students to make good use of the course discussion forum.

S17 *The course can have more content.*

A17 See our responses A13 and A15.

S18 *So far, I have checked several courses offered by Aalto, to my best knowledge, I didn't find a good course that analysis time-series data using deep learning. So, in the future, could you add this topic in this course.*

A18 See our response A15.

S19 *Instead of doing all the examples on image data, perhaps a short example on how this can be extended to text or numerical data would be helpful, since majority of the real life applications are other than image based, and there are much fewer high quality tutorials from those areas. Obviously, image data is very visual and good for explaining, but personally I feel that what I'm lacking the most is the understanding how to move from image to other data types.*

A19 The new course edition now also includes a module on natural language processing (Round 5) We will also provide more information on visualization techniques that transform various data forms into images. Roughly speaking, we can apply deep learning methods for images to any other type of data for which we know how to visualize it effectively. For an in-depth discussion of natural language processing methods, we refer to the courses ELEC-E5550 - *Statistical Natural Language Processing*, ELEC-E5521 - *Speech and Language Processing Methods*, ELEC-E5510 - *Speech Recognition*.

S20 *Perhaps a stronger nudge from the staff to the students to start working on the project during the rounds would help in applying the knowledge gained from the rounds. Of course I understand that this is the responsibility of a student to get her things done, but sometimes gentle nudges make all the difference.*

A20 We now use a more rigid time schedule for the various course activities and assignments.

S21 *Aalto courses pages needs simplification and focus. -hard to find things on course pages. -too much noise on mycourses.*

A21 We have tried to simplify the structure of the Mycourse page for our course.

S22 *The workbooks are so excellent that you can self-study them anyway, so it would be great if the lectures gave extra insight beyond what is in the workbooks.*

A22 The new course will include lecture sessions to provide more background and overview of the concepts used in the notebooks.

S23 *More about nlp would ve been interesting.*

A23 The new course edition now introduces some deep learning methods for natural language processing (Round 5 - NLP). We also refer to the more advanced NLP courses ELEC-E5550 - *Statistical Natural Language Processing*, ELEC-E5521 - *Speech and Language Processing Methods*, ELEC-E5510 - *Speech Recognition*.

S24 *Just as a side comment, one additional topic for students could be combinations of ML and some simulation models that are able to generate a lot of data for analysis.*

A24 We have added Round 6 on generative adversarial networks that allow to generate synthetic data.

S25 *There should in my opinion have been just a little bit more mathematical formulations, in particular of backpropagation, which is at the center of the deep learning approach. Besides that, it was OK not to use almost any math for explaining the concepts.*

- A25 The focus of this course is on building intuition and implementing deep learning methods with few lines of Python code. In particular, we do not require that students are familiar with the concept of gradients or partial derivations (backpropagation is just a method for computing these in an efficient manner) For a more detailed treatment of back-propagation, we refer to the advanced course CS-E4890-*Deep Learning* (<https://mycourses.aalto.fi/course/view.php?id=28212§ion=1>).
- S26 *There could have been more material for example about RNNs, GANs, and reinforcement learning.*
- A26 The new course edition now also discusses generative adversarial networks (Round 6) as an example for how to use deep learning methods to generate (synthesize) data. For a discussion of recursive neural networks (RNN) and deep reinforcement learning methods, we refer to the advanced course CS-E4890-*Deep Learning* (<https://mycourses.aalto.fi/course/view.php?id=28212§ion=1>).