Teacher's counter feedback

The course's teacher has not given a counter feedback yet

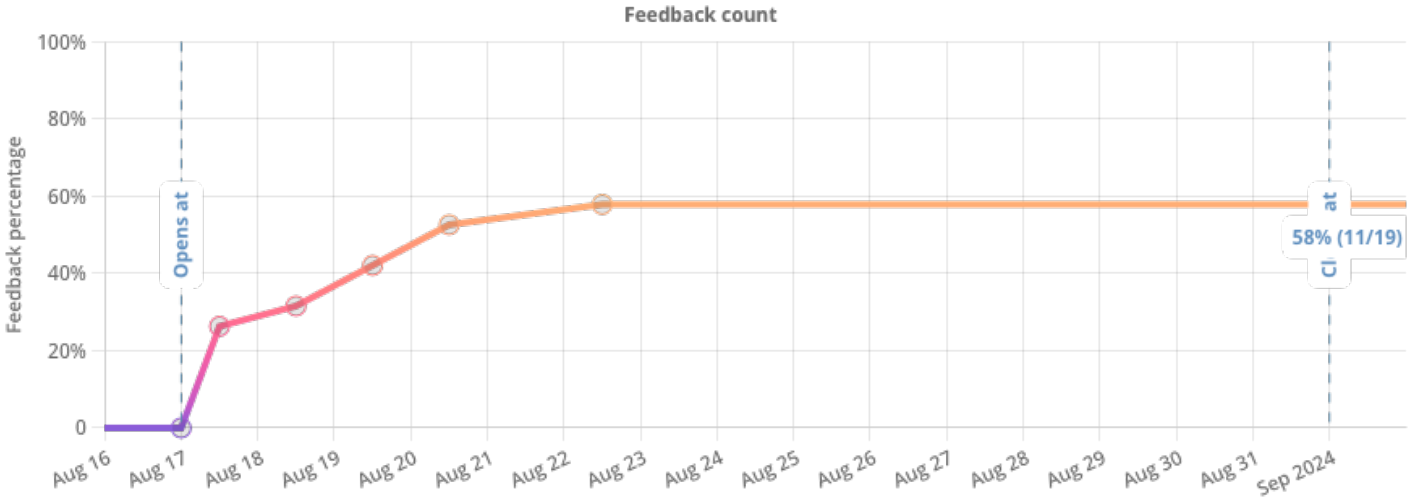

Multiple choice questions 6

1 = Completely disagree, 2 = Partly disagree, 3 = Neither disagree nor agree, 4 = Partly agree, 5 = Completely agree, N/A = Cannot answer

Public

4.36

The learning objectives were clear to me

Answer count (11/11)

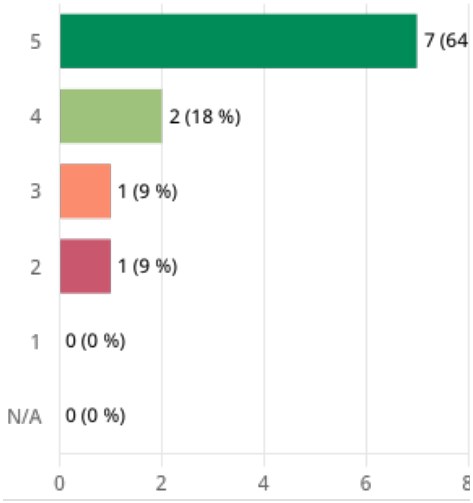

Public

4.18

The course activities supported my learning

Answer count (11/11)

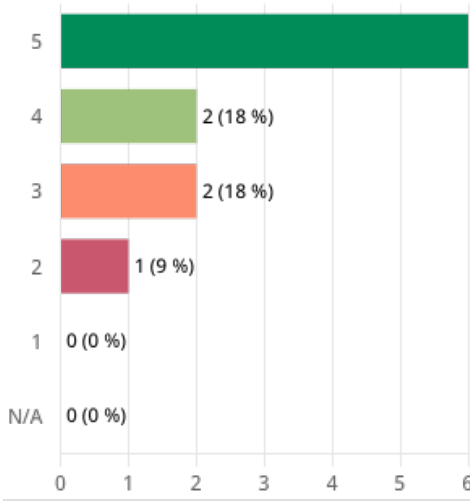

Public

4.09

The course materials supported my learning

Answer count (11/11)

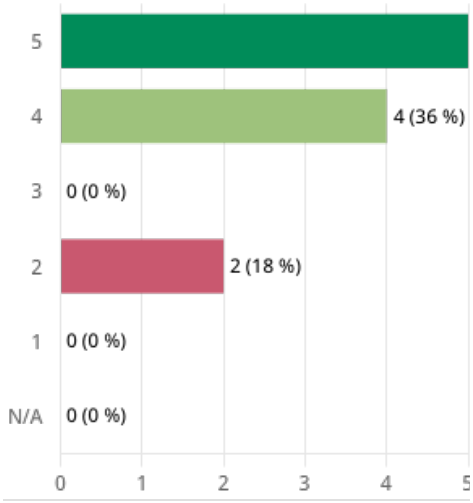

Public

4.20

The methods used to assess the course measure my learning

Answer count (11/11)

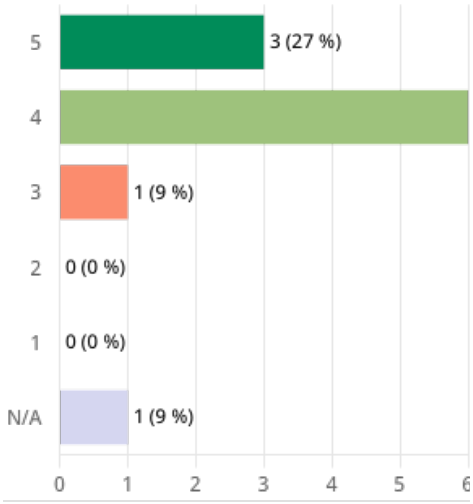

Public

The workload relative to course credits was

Answer count (11/11)

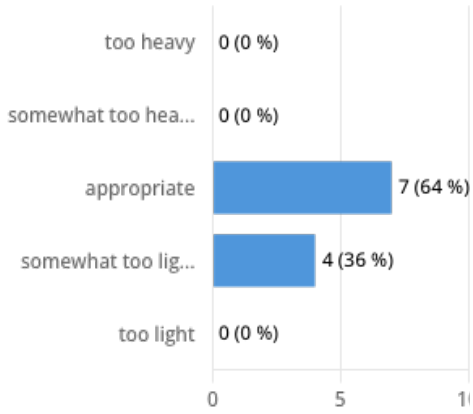

Public

Would you be able to use deep learning for your study

Answer count (11/11)

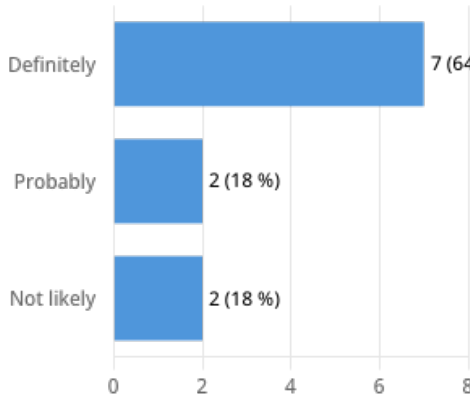

Open questions 5

Not public

Any additional comments or suggestions

Answer count (8/11)

I hope the course can be appropriately extended, and incorporating exercises related to chemistry would help us better understand computer concepts. And I hope we can ensure that every participant has the right learning attitude(For individuals who have not made any contributions but can still earn credits).

No additional comments. I think this summer school is really a good start and is really worth to join. I hope you can continue holding this each year.

Thank you for the course! It was interesting although difficult.

Thank you for the course. I really enjoyed it, and now I feel motivated to try to integrate more machine learning in my atmospheric research!

Thank you for the nice course. I do think that this can be built into something that is very beneficial to the people attending the course!

The group division was reasonable. The computer science group members were very helpful.

The organisation was spot on with alot of things I will say. I commend them for such great effort, maybe next time the course can circle in on the particular neural network architecture to use during the lectures so that the lectures become more structured and exercises also clearer.

This course is excellent and well-organized. I particularly appreciate the thoughtful arrangement of including students with computer science backgrounds in each group to help solve problems. It serves as a great introductory course for students without a computer science background to start exploring AI. The instructor's timely responses to questions were very helpful, especially for students from interdisciplinary backgrounds. This course is indeed essential for such students.

I have a few minor suggestions:

It would be beneficial to have more in-class demonstrations of code. This would help students understand the concepts better and encourage those with no prior experience to engage more actively in the class.

Adding more real-world applications of deep learning in atmospheric sciences would be very valuable. It would help students logically integrate their domain knowledge with AI.

If a more robust academic community could be established post-course for students to exchange ideas and collaborate, it would greatly enhance the course's impact.

I'm truly grateful to Michael for providing this learning opportunity, and to Mr Liu for encouraging us not to fear AI and to keep trying. This has really boosted our confidence in learning more.

Not public

Did the course provide the knowledge you were expecting?

Did the course provide the knowledge you were expecting?

Answer count (11/11)

It opened the door to a whole new world for me, and now I'm completely hooked.

The course provided almost everything I expected. There were enough lectures related to atmospheric sciences but there could have been more lectures related to AI and ML (I would have liked to know more about them).

The course was a great introduction to the subject matter presented.

The learning outcome of the practical about programming a neural network test remained rather shallow for me. This was due to my non-existent background of neural networks. Nevertheless, I have a little bit of experience from Python and therefore I was able to understand in a principal level of the structure of the functions and the layers of the convocational etc. neural networks programmed into them, in the Challenge exercise.

Since the course was so short there was not time for practical tasks other than the Challenge, while the tutorials were basically so filled with text that the temptation was just to press Play to get the code running on the cells. More practical content would have helped to internalize the machine-learning and neural network related concepts that remained basically in a rather abstract level. Zhi Song's and also Heikki's content was rather new to me and I had to pay a lot of attention especially in Heikki's slides not to get thrown off track, which I before long did. On the other hand, the topics that I had known of beforehand (Tuomo's and Victoria's lectures) were relatively simple for me to keep on track with and I found them rather interesting.

Since I took a lot of notes, not to forget the slides still in Drive, I could likely improve my learning outcome by going back to the content that was advanced through at a relatively fast pace.

yes

Yes

Yes

Yes and no. I knew this would be a short course, and I knew it would probably only scratch the surface with regards to deep learning. Compared to that, this course had surprising amount of detail. But because so much was crammed in such a short course, we did not have time to really go through most of it in a deep enough level that I could actually apply it on my own. For the atmospheric sciences part, the course contained the information I expected, but it was mostly entirely separate from the deep learning. In other words, this course had lectures relating to basics of deep learning and basics of atmospheric sciences but they were not very connected. The deep learning and atmospheric sciences parts remained mostly separate with only the final project as the common factor between the courses. And because the final project was done on artificial data, the atmospheric chemistry knowledge was not even relevant for that part nor did any of the groups use it.

Yes, I came in with the intention of learning some basics in machine learning and its application in Atmospheric science and I got a bit of both

Yes! Overall the course was what was expected with a nice introduction to the world of combining machine learning in climate science

Yes. The part i am most looking forward to in this course is deep learning. it was well organized both in tutorials and group work, i gained a deeper understanding of it.

Not public

Were the lectures covering the required information, and if not, what lectures would be needed?

Answer count (11/11)

I believe it covers two areas of knowledge, but there isn't a course that specifically focuses on using mathematical or computer models to study aerosols and VOC.

I feel that the lectures covered the content I was expecting. Both the atmospheric and the machine learning part were covered. The atmospheric lecture by Victoria also highlighted how ML could improve or change the way of forecasting, which was on the spotlight of what the course was all about, combining atmospheric science with ML.

I think it is a nice pattern to combine the atmospheric science students and the computational science students. I expected more atmospheric related ML contents and exercise.

In the lectures, some background knowledge is very friendly to students who are not major in atmosphere. However, it felt like the lectures were too independent. for example, i would like to know more case studies that the application of AI/ML in atmospheric sciences.

The atmospheric science lectures and the statistics lecture were not "required" for this course (for the exercises and or the project), but as far as I can see, they provided pretty solid introduction to their respective topics.

For the deep learning, the issue was not as much that something was missing but rather that the concepts were not explained in a manner that would have been easy to understand if you do not have a background in ML/deep learning. I have watched the 3blue1brown YouTube channel, and there most of the concepts were explained in a more beginner friendly manner and still used around the same amount of time.

The lectures covered a wide range of things although I think subsequently exercises on the atmospheric area is necessary to truly understand this section too

The lectures were nice. Some lectures like the aerosol felt like we were moving quite slowly through though and more information could be within it. Machine learning lectures went through quite a bit, but I am not sure if for example a very indepth walktrough of all types of models is required as that provides A LOT of information and people can be overwhelmed.

The lectures were very good, but perhaps too much information was packed into one lecture. For example, several lecturers said that this is several hours of lectures that they are trying to present in an hour and a half. Perhaps making the material less dense and not trying to cover quite as many subjects but covering it in better and more understandable detail would have been helpful. Of course, this would have worked better if the course was for instance 2 weeks instead of one week.

There could have been more lectures related to artificial intelligence and machine learning, because the subject was very difficult. I can not say more precisely about their content.

Yes

yes. I would have liked more AI specific lectures.

Not public

Were the exercises too easy, too difficult or just ok?

Answer count (11/11)

Exercises are easy because there are tutorials to follow.

Exercises were a little too difficult.

For the time we have, I will say the exercises were too easy, maybe next time the student should be able to build the neural network code and be able to preprocess the data not just changing parameters

I feel the challenge exercise would have been way too difficult for myself as an individual task. But since the groups were divided so that there were computation-based students in every group, I mostly could and had to just keep up with understanding what the other guys were up to while they were testing the validity of the neural networks with different parameters.

The tutorial exercises helped to understand the bigger picture a bit, but I ended up just scrolling through them by pressing the play buttons of each cell without understanding what the cells do, because there was a pressure to start with the Challenge excercise.

It can be more difficult.

It is easy cause we have got the baseline, without it , it could be difficult.

Just ok.

ok

Personally, I felt like the tutorials were not enough hands-on from the machine learning side. I liked that there are some examples like classifying images and numbers etc, but that does not necessarily transfer over to the current challenge. In my group it felt like people were afraid to do anything as they thought that the baseline model could break or they just didn't really want to touch it as it was already "working". Maybe some slides or tutorials of just basic model building could help. Or additional material where people could get some reference.

The exercises were as difficult as they could be and still expect the students to finish them in time. That being said, most of the exercises you could just run without actually understanding what is happening.

The exercises were good for the course objectives and time provided. However, there wasn't really enough time in a one week course to really understand the material. Like with the lectures, it probably would have worked better in 2 weeks.

Not public

What could be improved for future courses?

Answer count (11/11)

1. General feedback:

1.1 This course is quite short for this number of lectures, exercises and the project. I would either cut something out or add to the length of the course. Even adding two hours to the course length would have been useful. 1.2. For a course that is named Application of AI/ML techniques in Atmospheric Science, the AI/ML and the Atmospheric Sciences parts were almost completely separate.

2. Feedback about the exercises and the project: 2.1 For people in my group, the exercises took too long and we were really able to start the project really on Thursday. In other words, we basically did the group project in three hours. This, combined with the fact that most of us did not have access to GPU in GoogleCollab or ran out very quickly and had to come up with alternative solutions, made the project very frustrating. Then, when presenting the results, the groups got feedback from the lectures suggesting things like "you should have changed the figures" or "you should have analyzed the figures" and "you should have thought what is wrong with the baseline code and edited it", when in reality we just scrambled to get anything done. This really was not optimal conditions for learning. I think there are quite a few things that could be done to avoid this: 2.1.1 Doing the exercises and the group project without GPU was very slow on GoogleCollab. There needs to be a solution for this whether it is to use some other platform, run it on peoples own computers, or whatever. Or at least give better instructions for using the GoogleCollab 2.1.2. The time used on doing the exercises could be cut down and moved to doing the project: 2.1.2.1 I would add a pre-exercise for setting up the GoogleCollab (or whatever is used in the future). It is pretty common pre-exercise for this type of courses. That way that does not take (so much) time during the course. If a student does not get the platform set up before the course, they can ask help during the first day. 2.1.2.2. I would cut down on the exercises. In the CNN-tutorial, there is first the Prerequisites: Experiment 1, we run the animal classification code and try classifying a picture, then in Experiment 2, the we run the digit classification code and try to classify a digit. I would move one of these to be a pre-exercise and remove the other one. That way we can start straight away from the Tutorial 1. Also, in Tutorial 1, I would actually remove Exercise 1 and do only Exercise 2. Exercise 1 takes a long time to run (especially if you don't have access to GPU) and the concepts relating to the batch size and learning rate should be explained a bit later in the course (see section 3). Also, the GNN-tutorial is more important for learning the skills necessary for the group project, there should really be more focus on that one. 2.1.2.3. There should be deadlines for the CNN and GNN tutorials. I think that they should be done by Wednesday morning. This would also mean that the ML/AI lectures should be in the beginning of the course and the other lectures should be at the end of the course (especially the aerosol, meteorological and statistical lectures that are in no way relevant to the exercises or the project can be towards the end of the course). 2.1.3. On Friday, you should reserve time for people to finalize and practice the presentation.

2.2 It would be nice if the exercises and the group project would have more to do with atmospheric data. Now none of the exercises actually had anything to do with atmospheric data, and the group project used artificial data.

3. Feedback on the lectures: The basics of ML/AI lecture on Monday needed to be more basic and lay the ground for the future lectures. I would move the Neural Networks and their math to later on in the course, and instead focus on 1) different types of tasks ML/AI is used for (regression, classification, clustering etc) and 2) explaining how we typically use training, validation and testing to build and evaluate the models. In the current version of the course, training, validation, and testing sets were not really explicitly explained, but they kept coming up,

and one student had to ask about them.

Then later on in the course (the ML/AI lecture on Tuesday) I would first explain what neural networks are and what they do without any math. And only then you can do the math-heavy lecture Andreas gave and really explain how the parameters are calculated. And once that is explained, you can explain what other different layers besides fully connected layers are possible in MLP and what they do. For the layers, you really need to explicitly explain what are "fully connected layers", "convolution layers", "pooling layers", "batchNormalization layers" etc. Now a lot of these were touched on, but I am pretty sure that for example "fully connected layer" was never explicitly explained. When you hear things for the first time, even the smallest things can trip you up, like what is the fully connected layer, even if it is somewhat trivial. And only once MLP has been explained properly, you should move on to RNN and GNN and attention and things like this.

To match the timing of the exercises, I would have most of the ML/AI lectures on Monday/Tuesday, the atmospheric chemistry lecture on Tuesday, and the other lectures a bit later on during the week.

As a atmospheric student, I suppose the load of atmospheric contents might be too heavy for students without atmospheric background. They receive too much general atmospheric information which are difficult to remember. Most of the knowledge will not be used for their future study. I think the contents should be deep and specific, instead of introducing too much general background. A case study is always a good idea. You can just give a specific atmospheric science case for the ML test, and only introduce the knowledge about this case. After this, we try to use the ML to solve the case. This is a good way for both students to know how ML works in atmospheric science study.

In addition, if the contents need to fit students with both backgrounds, some preparation for students is also necessary. For a better understanding of the contents given in the course, some knowledge can be sent to the students in advance (before the summer school) and let us do some self-learning. You can give a hint about which part of knowledge is necessary for the the course. I believe the students have a motivation and ability to learn by themselves first, at least to learn those necessary knowledge. After this, the contents during the course can be more specific and deeper and atmospheric case related.

I think it would have been nice for the summer school to have been a bit longer to learn more about examples of combining atmospheric science and AI.

I think the tutorials could have been improved with relatively easy practical problems in the cells, so that you don't just need to press the Play button to get onto the next task. The tutorials could have at least tested one's understanding by a small quiz at the end, perhaps without any grading to release the pressure of taking it.

Also, practical examples of otherwise abstract terms and concepts were essential for getting a grasp of them, such as the stretching dogs when talking about data augmentation. I think there is always room for practical examples.

I think this course would have been more useful in 2 weeks because 1 week felt too rushed. It seems like right now AI and atmospheric science are still two completely separate subjects, and it would be nice to have a lecture or additional exercises that connect the two better. It would have been helpful to provide us instructions for setting up our environment and programs ahead of time as a pre-exercise, so that we weren't spending our group work time doing things like uploading data and opening data. It was rather confusing to have instructions for exercises in PDF files online and also directly in Colab. I think it would be better to have the exercise instructions in one place or the other, not both. As an atmospheric scientist, I would have also liked to see a lecture on how GPU works, how AI programs can be run on GPU, and why it's more efficient than CPU-based linear programs we usually use in our field of study.

I would prefer to write the codes from scratch to learn how to actually do NN architectures myself. I would prefer to have tutorials with codes all visible in the ipynb file; I agree that python files are useful for a cleaner code, though I think that for learning it's easier to have everything in one place.

I wrote a lot overall in the previous boxes already, but I'll try to summarize:

- The lectures were good, maybe trying to go over all basic ML architectures is not the best idea. I would suggest covering all, but having a more in-depth focus on one that would be also used in the challenge. This way a good hands-on experience could be received.
- Tutorials felt a bit too much just run this through. Didn't really transfer that smoothly to the challenge in our group at least.
- Maybe a longer course?

Maybe it could be longer. The schedule of lectures was a bit intensive for me and it was difficult to understand all of them in a short time.

The course could be longer. There could be more exercises and they could be gone through with together after everyone has attempted them independently (I did not get support from my group). Also more time to do the exercises. There could be a break in the middle of the lectures (45 min + break + 45 min, sometimes it came true and it was good).

The length of the course can increase giving more time so more exercises can be done independently by the students. Also the barbecue was perfect.

The practice content isn't directly related to actual chemistry, so even if students master it, they don't understand how to apply it in a chemical context. This is a common issue for most atmospheric science students.
