Course analysis SF1930 HT21

Authors

Joakim Andén-Pantera, examiner and course responsible (e-mail: janden@kth.se)
Liam Solus, course responsible (e-mail: solus@kth.se)

Course evaluation process

All students were invited to participate in the course evaluation committee (sv. "kursnämnd"). To ensure and equitable gender distribution, we required that the course evaluation committee be composed of the same number of men as women. Two students chose to participate: Tufva Linde and Felix Steinberger Eriksson. Three meetings were organized with the course evaluation committee during the course. We also conducted one meeting with the teaching assistant, Alice Harting.

As part of their evaluation work, the course evaluation committee organized one online survey at the midpoint of the course and another after the course end. Each survey had 22 respondents out of 38 registered students. We also organized a survey using learning experience questionnaire (LEQ) function. For this survey, we had 12 respondents.

After receiving the results of the survey and the report from the course evaluation committee, we held a closing course evaluation meeting at which the course responsibles, the course evaluation committee, and the program responsible, Sara Zahedi, were present.

Meetings with students

We continually solicited feedback during lectures and office hours to get a sense of the students' progress and opinions. Since this was the first time the course was given, their input was especially valuable as it allowed us to shift the emphasis of the lectures and problem sessions when needed.

Three meetings were held with the course evaluation committee: 2021-09-20, 2021-10-04 (only Liam Solus present), and 2021-10-15. A closing course meeting was held on 2022-01-17 with the course responsibles, the course evaluation committee, and the program responsible.

Course design

The basic elements of this course were decided on in the fall of 2020 by an initial working group consisting of Jimmy Olsson, Henrik Hult, Tatjana Pavlenko, Joakim Andén-Pantera, and Liam Solus. Following the decision (S-2020-1415) by the head of school, it was decided that a smaller group consisting of Joakim Andén-Pantera and Liam Solus go ahead and formulate a detailed course plan during the spring of 2021 that they would then go on to teach in the fall.

SF1930 Statistical Learning and Data Analysis continues where the introductory course SF1918 Probability Theory and Statistics left off. It introduces more advanced concepts in statistical inference (point estimation, interval estimation, and hypothesis testing), with a perspective that aims to balance frequentist and Bayesian approaches along with decision-theoretic elements. In addition to the theoretical component, a strong emphasis is put on numerical computation, with the last two lectures devoted to Markov chain Monte Carlo and an accompanying project assignment.

To ensure continuous evaluation and feedback from the teachers, the students were offered several opportunities to interact with the material before the exam. First, a homework assignment was posted at the end of the first week with a deadline three weeks later (a passing grade on the homework assignment would give the student bonus points on the final exam). This consisted of ten problems covering material from the first two weeks of the course. A few days before the deadline, we organized a tutoring session (sv. "räknestuga") where the students could ask questions regarding the homework. Following submission, the homework assignment was graded by the teachers who gave detailed feedback on the students' solutions to the problems.

At the midpoint of the course, a midterm exam (sv. ”kontrollskrivning”) was given (a passing grade on the midterm exam would give the student bonus points on the final exam). The format was five questions for which the students were graded solely on the answer and these questions were mostly computational in nature.

Finally, there was a project assignment (see above) that was published two weeks prior to the exam. This involved working with a large dataset and performing various inference tasks using a computer. The students had two computer lab sessions scheduled for this project during which they would work on the project and ask questions as needed.

The lectures were of a standard type, with most of the time dedicated to covering the lecture material on the board. To reinforce certain ideas, live demonstrations of computational aspects were also carried out using Jupyter notebooks.

The problem sessions (sv. ”övningar”) were led by Alice Harting and were organized to encourage the students to work as much as possible on their own before being given the answer. Typically, the teaching assistant would write down the problem on the board (along with a hint, if needed) and let the students work on the problem in groups for five to ten minutes. During this time, she would walk around and try to help the students who seemed stuck. After this, she would present her solution or ask for suggestions from the students.

As mentioned above, there was a tutoring session organized prior to the homework deadline. Another session was also organized before the final exam where students were encouraged to work on a practice exam and could ask questions if they had difficulties.

Apart from the regular classroom instruction outlined above, the course responsibles also organized online office hours every week. This meant that they were available for one hour online where students could ask questions on the course. Typically, anywhere between three and five students would show up during this time every week.

Students’ workload

The course has a total credit value of 6 hp, which corresponds to 40% of full time, or 16 hours per week. Due to the large amount of material included in the course, many students were forced to put in more work than this (see below). We were made aware of this at the first meeting with the course evaluation committee and subsequently tried to cut down the amount of material covered in the lectures and problem sessions. That being said, the workload was likely still in excess of the allotted 6 hp. Course development for next year’s course will in part focus on further cutting down the material to a more manageable size.

Students’ results on the course

The students’ results on the homework, midterm exam, and final exam (after the oral exam given to the students who received the Fx grade) are as follows.

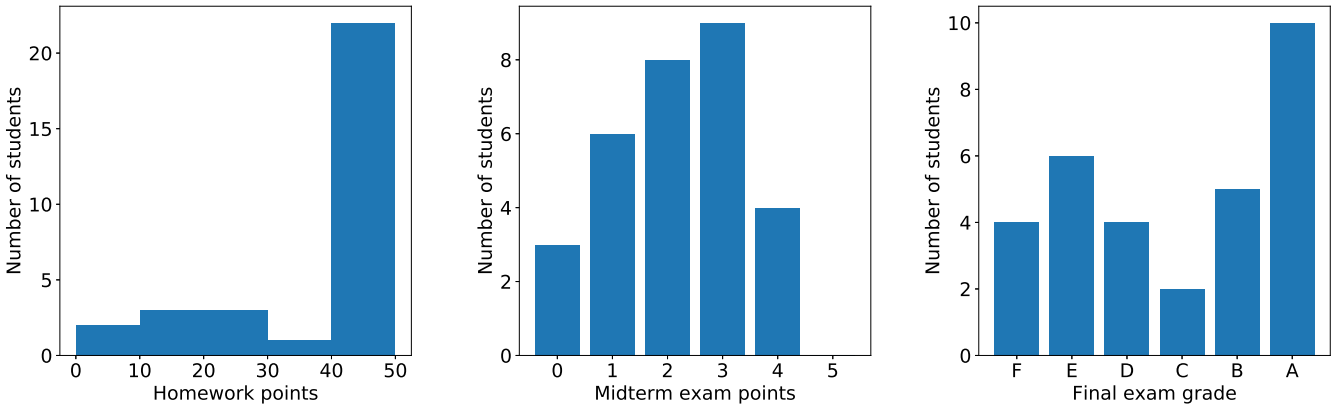

Of the 38 students registered for the course, a total of 31 students submitted the homework, 30 students attended the midterm exam, and 31 students attended the final exam.

While the homework was perceived by most students to be quite difficult, most of them had enough time to complete it satisfactorily and the vast majority of students scored in the 40–50 interval, which entitled them a total of 5 bonus points for the second part of the exam. A large number of students also attended the räknestuga, which took place one week before the homework submission deadline, and worked actively with each other during this time. The course evaluation committee reported that these students found the räknestuga highly beneficial, as it

allowed them to ask the teachers detailed questions about their approaches to the problem and created a space for them to work actively and dynamically with other students.

Due to the different nature of the questions on the midterm exam compared to the questions studied during problem sessions (see below), students did not fare as well on the midterm exam. In addition, the final problem (problem 5) on the midterm exam ended up being of higher difficulty than intended. As a result, we moved the lower limit for a passing grade from 3 to 2 points. Of the 30 students who took the midterm exam, a total of 22 students passed it under this rubric.

For the project assignment, a total of 32 students submitted reports. Out of these, 27 students passed and 5 students failed.

At the final exam, a large portion of students performed very well, with 10 students receiving A's and 5 students receiving B's. On the other end, 6 students received the lowest possible passing grade (E), while 4 students failed the exam. Consequently, we cannot say that the exam was either too difficult or too easy. However, feedback from the students (see below) suggests that they found the exam to be at an appropriate level.

Students' answers to open questions

Following the LEQ survey administered by the course responsible, several comments were recurring. In general, they concerned:

- Large amount of course material and high workload (see above).
- Although the course was challenging (for several students, the most challenging course they have taken), they very much enjoyed it and found the applications useful and interesting.
- The teachers and teaching assistant were found to be accessible.
- The handwritten notes were not very legible, but were much appreciated as the course literature (Casella & Berger and Gelman et al.) were at a higher level than the course in general.
- The students would have liked less focus on proofs and derivations in the lectures and more focus on computational aspects.
- Similarly, the level of the problems during the problem sessions were seen as more difficult compared to the midterm exam and the first part of the final exam.
- Despite this, the structure of the problem sessions, where students were encouraged to work by themselves for a certain period of time, was appreciated.

Summary of students' opinions

The results of the LEQ survey aligned overall with feedback we had gotten from individual students as well as the student evaluation committee. First, there was the question of course material and workload, which was perceived to be too high. That being said, a majority of students enjoyed the challenge and found the material interesting. Coupled with this was the high focus on proofs and derivations during the lectures, which sometimes came at the expense of more conceptual discussions and computational examples.

Related to this was the selection of problems for the problem session. In general, the problems were taken from the literature (Casella & Berger for the first half of the course and Gelman et al. for the second half), which were of medium-to-high difficulty. This was contrasted with the midterm exam questions and the first part of the final exam, where questions were more of a computational nature and therefore easier. However, the students did not receive as much practice on these "easier" questions and therefore would have difficulty properly allocating their study time. This was also brought up at the initial course evaluation meeting and the problem sessions for the latter half of the course were adjusted to account for this, but there remains some work to be done here.

Another issue concerned the course literature. The course relied primarily on the handwritten notes provided by the teachers but also supplementary reading from Casella & Berger and Gelman et al. Both of these textbooks are at an advanced undergraduate or beginning graduate level. At the beginning of the course, it was not made clear that this course literature was supplementary and the students should be focusing on the course notes. After stressing this to the students, the issue was partly resolved.

Overall impression

This being a new course, there were many challenges that appeared in terms of preparing material (lecture notes, problem sessions, homework, midterm exam, project assignment, sample exam, and final exam) for the course. A detailed outline had been prepared during the spring and summer prior to the course start, but much of the work had to be done during the course itself. Due to the time constraint, we were not able to deliver all of this as early as we would have liked or at the level of quality we would have liked. That being said, we are generally happy with the content covered and the results obtained by the students. As was mentioned above, the course was challenging, but we found that overall, students rose to that challenge and seemed to enjoy the course material. This is reflected in the grades but also in our interactions with the students themselves.

Analysis

We share the opinion of the students regarding the amount of course material. From a teacher's perspective, there was a significant amount of material to prepare and include in the lectures and problem sessions. Part of this stemmed from a relatively ambitious course plan that had been set up during the initial working group during the fall of 2020. With the benefit of hindsight, shifting focus from certain elements to others could have helped here and will be part of the ongoing course development. Another part of this can be explained by the relative lack of familiarity on the part of the teachers with the details of the material.

The issue regarding the difficulty problems selected for the problem sessions was not obvious to us at the beginning. This is most likely related to the desire to select problems that we found interesting and challenging to the students, ignoring the fact that this may result in problems that were too difficult for the students at their current level. As stated above, after being made aware of this, we tried to include easier problems in the problem sessions that were of a more computational nature.

We also share their concerns about the course literature. As mentioned above, there is no real book available that covers the course material at the undergraduate level, hence our effort to provide a synthesis in the form of handwritten notes. Of course, this being the first iteration of the notes, there remains a significant amount of work here.

The students did not bring up any issues related to difference in experience between students identifying as male or female. During interactive portions of the course (problem sessions, tutoring sessions, and so on), we noticed that female students tended to be more reluctant in asking questions. As a result, we tried to make sure to actively solicit questions from all the students in order to reduce this potential barrier.

A total of 8 female students submitted the final exam and out of these student 3 received A's, 2 received D's, and 3 received E's. Since this is a small sample, it is hard to draw any definitive conclusions, but these results are roughly on par with the results for the whole group.

Prioritized course development

Following discussions between the course responsibles, the student evaluation committee, and the program responsible, we have identified three main areas of development. The first is to cut down the material in order to reduce the workload consistent with the 6 hp allocation of the course. This will be achieved in several ways. One part will be to cut out more advanced material from some of the lectures. For example, we may decide to remove the Neyman–Pearson lemma from the lecture on hypothesis testing. Similarly, we may compress the two lectures on model checking and model comparison into a single lecture. Another part will be to remove proofs and derivations from the lectures. This being a second-year undergraduate course, it may be more important to emphasize concepts and tools over rigorous proofs. However, we plan on still including these in the course notes, but to mark them as supplementary. Finally, we will reduce the amount of time spent on reviewing material from the previous course by providing references to that course prior to lectures where that material will be needed. Alternatively, we may include only a short summary as a review at the beginning of the lecture.

The second main task is to restructure the problem sessions to include easier problems and problems more focused on computation. This will bring them more in line with the midterm exam questions and the first part of the exam. It will also help the students develop confidence and familiarity with the methods before tackling the harder problems in the problem session. Finally, we will go over the handwritten lecture notes, extend them where necessary, and typeset them. The goal is to have a self-contained set of notes that will form the core of the course literature. We may keep references to the textbooks (Casella & Berger and Gelman et al.) but in that case it will be made clear from the beginning that this material is supplementary. As part of this, we will also translate the notes from English to Swedish to bring the course in line with the rest of the undergraduate curriculum.