DT2112, VT2019 Course Analysis

Author

Jens Edlund, edlund@speech.kth.se

Description of the course evaluation process

The course is highly interactive, with supervised project work as a significant part. The most valuable feedback occurs during the supervision sessions. Two course representatives (gender balanced) are appointed at the beginning of each course round.

The gender balance in the course is roughly 35 female/65 male, a number that has been steady for several years. Much of the student feedback takes place in class and during project supervision, where it is difficult to keep track of potential gender issues objectively, but the overall impression is that there is no significant bias with males speaking more than women apart from what is the result of the 35/65 balance in numbers. Speech technology has a history of failing in this respect, however, and we are looking for better, more objective ways to ensure that each gender is heard.

Description of meetings with students

The course has three scheduled lecture hours (initially, before the projects, and before the home exam) with open discussion around the course modules.

During the 2 weeks of project work (plus another 2 preceding that, where the projects are designed), the students meet regularly in small groups with teachers (mainly with the course responsible Jens Edlund). Each group has at least two such meetings, and the average is somewhere between 3 and 4. These meetings are reminiscent of Master's exam work supervision, but also include discussions on course practicalities (mainly concerning the formalities surrounding grading, including how the home exam works).

Towards the end of the course, there is on e(or if needed two) full days of group presentations of projects. These sessions are styled as seminars, with all students asking questions and commenting on the work. Oftentimes course and study related discussions come up.

Finally, it is worth repeating that virtually all students spend each full Thursday in our premises, where they have access to all teachers, throughout the course.

Course design

The course is an overview course in speech technology. KTH does not offer a speech technology programme, so in spite of having high requirements and being placed late in programmes, it cannot make any assumptions about the students' previous knowledge of speech technology. Instead, it aims to provide a broad insight into the field, and to show how students of varying backgrounds can partake and contribute to the area.

The course postulates seven intended learning outcomes. After completion of the course, participants shall be able to:

1. Describe speech from an acoustic, phonetic, and linguistic perspective
2. Explain how computers recognize speech with statistical methods, and evaluate the recognition results

3. Describe and judge different methods used to produce speech with computers
4. Analyze speech-driven dialogue systems with respect to application, components, functionality and user aspects
5. Describe how evaluation of speech technology systems work and describe the special requirements that are posed by this type of system
6. Design and describe data collections for speech technology research and development
7. Give an account of available state-of-the art speech technology and exemplify the current speech research on e.g., mobile systems and IT applications
8. Apply the theoretical knowledge in small speech technology projects

The course comprises a range of learning activities. Learning outcomes 1 through 6 are associated to scheduled lectures with professors who are specialists in these specific topics. With these lectures come mini-exercises performed in breaks or between lectures, and three more substantial labs. Learning outcomes 7 and 8 are included in the courses two-week project and in the bidding process that precedes it. In the project, some but not necessarily all of learning outcomes 1 – 6 are also practiced further.

For each learning outcome, there are grading criteria that aim to assess the depth of knowledge achieved by the student. For an E, students shall demonstrate the ability to recall and understand the intended learning outcome. For a C, students must also demonstrate the ability apply and analyse the same knowledge. And for an A, the added requirement is to demonstrate the ability to create applications and/or experiments based on this knowledge. B and D are given to students who only just miss the requirements for A and C, respectively.

The assessment is transparent and follows clearly defined grading criteria. 50% of the grade is associated with a project report (and the underlying work that it represents), and 50% by a home exam. For the projects, students must present a bid, including the grade they are aiming for (A, C or E), that is detailed enough that the supervisor can say that "should this work be completed, you will be awarded your target grade". The target grades meet the demonstration requirements: A projects must include implementation of applications and/or experiments and their evaluation; C projects need only show proof-of-concept or study designs, and E projects are basically a well-written background for a study or application implementation. The bidding process, then, is a very large part of the project work. Assessment is straightforward: students who hand in a report describing the completion of the work in the bid, on time, get the target grade. Late hand-ins, or projects that only just miss the mark, get one step lower (A->B, projects targeting A cannot get C or D; C->D). To achieve an E, it is enough to write a good background (even if the initial bid targeted A or C).

The remaining 50% of the grade comes from the home exam. Here, too, the exam is directly linked to the grades and to the learning outcomes. Five questions (mapping to learning outcome 1+2, 3, 4, 5, and 6, respectively) are given on each of the E, C and A levels. The E-level questions are relatively simple (understanding and recollection is enough to respond), the C-level questions requires some analysis, and the A-level questions require creative thinking and application of knowledge. Each question is reqarded with 1 point per correct and relevant statement (up to 4) and 1 point is deducted for each incorrect or grossly irrelevant statement. The grading criteria include a table that shows what is needed to achieve a certain grade: how many correct answers (of 1-4), on how many questions (of 5), for each level (E, C, A). A student aiming for an E, then, can simply skip the C and A level questions. Students aiming for A will know with relative certainty when they have done enough. The implementation of these grading criteria has completely done away with grading complaints, and it is rarely or never difficult for the grading teacher to know what grade to award.

The interplay between intended learning outcomes, learning activities, assessment and grading criteria have been developed iteratively since 2015, and is now pretty robust. No changes were made from last year.

Students' workload

The students spend as much time as can be expected, and although they put in a good number of hours in the premises of KTH Speech, Music and Hearing (a very large proportion of the students – in excess of 90% - spend each Thursday at our premises throughout the course) there are few or no complaints about this. In this respect, the scheduling changes to clarify work towards the end of the course seem to have worked.

In addition, the way the grading criteria are implemented makes it feasible for the students to target the grade they want to achieve and to track their progress. A very high proportion of the students achieve the grade they aim for, and are aware that this is a result of the effort the put in (this cannot be objectively measured without mind reading skills, but is the impression we get from discussions).

Students' results on the course

52 students were admitted to the course. 47 students (90.2%) completed the course by the end of the course period, and another 3 handed in delayed reports/home exams before term end. In total 50 students (96%) passed the course within the term. This is small improvement compared to the average result in previous years, but the difference is not likely to be significant. Note also that this year, the students from Uppsala are included in the statistics (they were not previously for administrative reasons).

The grades were distributed as follows: 24 A, 12 B, 9 C, 5 D, 2 F (average estimate: 3.9/5). This is an improvement over precious years, and is likely the result of a change in scheduling that allow students to plan their time better over the two modules that are graded: project work and home exam. The grading criteria are quite clear and allow students to aim for their desired grade with some confidence, but previously there have been some issues with time management, which have led to a small decrease in grades.

Students' answers to open questions

Apart from comments on slow grading (see below), positive remarks about lectures, bidding process (projects) and the field in general. From the VT19 class, one student is aiming to do his exam work in Speech Technology (this is unusual for timing reasons; the course is not scheduled to lead into exam work).

Summary of students' opinions

Complaints are exclusively to do with slow grading (which is a fact, the grading of VT19 was significantly delayed).

Positive remarks, in particular from student supervision meetings, have to do with interesting and rewarding project work.

Overall impression

The teachers overall are impressed with the students, and in particular with the way the students socialize around the topic. This is, we believe, an defect of the scheduling: students quickly get used

to spend full Thursdays at TMH (the department of Speech, Music and Hearing) and often spend lunches and breaks discussing speech technology related things in groups on the premises.

Since the last course offering, the major changes to the course targeted the issues we have had with time management towards the end of the course. A focussed effort to bring in more supervisors for the student projects was made to shorten response times during the project work, and the schedule was made clearer in terms of separating project work from home exam.

Analysis

Weak areas include the tracking of gender and equality issues and the teacher work load associated with the home exam.

Gender and equality issues. Although we see no significant differences in experience between students identifying as female and male, or students with different backgrounds, we have not been sufficiently clear in how we track this. The course deals with a highly interdisciplinary topic and the students have quite different scholarly backgrounds, coming from a wide range of programmes, and our focus have been to manage the potential problems rising from that and turning them into strengths (for example by encouraging students to group up across backgrounds in order to better achieve the interdisciplinary goals of projects). We *believe* that these aspects (interdisciplinarity, broad mix of students) are beneficial in terms of gender, equality, and disability issues, as there is ample opportunity to play to different strengths, but we are not tracking this well enough.

Work load and associated delays. Both the project work and the home exam are very labour consuming in terms of supervision (project) and grading (home exam). The grading criteria and the bidding process has reduced the grading effort for the projects to something very manageable, and the supervision, although time consuming, is appreciated by the students and significantly improves the project quality. The design of the home exam and its grading (redesigned from the ground some years ago in order to provide transparency to students) has completely removed the recurring grading complaints we used to have, which is very positive. The grading process is slow, however, and we are looking for ways to make it more efficient, otherwise we will keep struggling to provide results in a timely manner.

On the positive side, the course is well liked, and very few students (of those who actually show up to the course) are lost. In addition, the students spend significant time together, and meet new contacts from programmes they might not be so well connected to.

We also note that students have mentioned that the project part of the course is their first contact with all of (a) supervision on technical writing, (b) supervision on literature search, (c) and designing project bids and offers a “client” (the supervisor) and delivering according to this plan. Although these are learning outcomes of the course (except the last one), they are supporting skills that constitute a foundation of much scientific work.

Prioritized course development

For V20, the following development is prioritized:

- More efficient grading of home exams
    - Tests with having teaching assistant help out and pre-grade the exams.
    - If time permits (this is a scheduling issue which isn’t easy to solve) make initial tests with peer grading – allow students to pre-grade each other’s exams, and discuss the results in seminar form.

- Better tracking of gender and equality issues
    - Set up initial checklist in an attempt to objectively track these issues
    - Base check items on evaluations and grades, but also look to quantify other aspects such as supervision time, in-group diversity.