Revision tutorial

Peter Bell

Automatic Speech Recognition

General points

- No longer an open book exam, but 3 A4 sheets of notes are allowed
- Paper style is similar to the past few years
- Focus on *understanding* rather than memorising
- Mostly problem solving/long answer questions
- Be aware when studying papers from 2020 and before:
  - Exams were not open-book or notes-allowed
  - The field of ASR is rapidly changing
- Questions may be open-ended – there is no one correct answer

Rules on notes

From the Head of Student Services:

- Candidates are allowed to have 3 sheets (6 sides) of A4 paper, with whatever notes they desire, written or printed on one or both sides of the paper.
- Magnifying devices to enlarge the contents of the sheets for viewing are not permitted.
- No further notes, printed matter or books are allowed.
- Candidates with learning profiles that mandate the provision of larger format exam papers are allowed a proportionate increase in the number of sheets of notes taken in. (For example, if a student is given their exam on A3 paper, they will be allowed to take 6 rather than 3 A4 sheets of notes into the exam.)

Exam technique

- The biggest constraint is time pressure
- Answers that copy generic content from the lectures are unlikely to be rewarded well – try to add your own insight
- For the more open-ended questions, be aware of how many marks are on offer for each question part, and write accordingly

Calculators

- Calculators *are* allowed...
- ... but must comply with University requirements https://informatics.ed.ac.uk/taught-students/all-students/your-studies/assessment-coursework-exams-feedback (see the bottom of the page)
- Do you need them?

"Sitting in my office, I would be able to do all the questions without a calculator, but that is away from the pressure of an exam hall! If you own a suitable calculator, I would recommend bringing it.

"If you don't own one and don't want to buy one, my recommendation would be to look through last year's past paper to gauge for yourself how comfortable you would be doing the questions without one."

Piazza questions

I'm having a bit of trouble deciphering the testing procedure for Speaker Adaptive Training based models - Do we need to perform a first pass where we obtain the occupation counts for the testing frames, then optimise the likelihood objective to obtain the transform for that sample, finally transform the gaussian parameters and perform recognition? or am I completely off and there's a different straight forward procedure?

Piazza questions

For Q1c in the 2021 ASR exam, it asks us to compute the posterior probabilities of two different arcs. My understanding was that the posterior probability is $P(Arc|X)$ where X is the observation probability modeled by the HMM. Using Baye's rules, we would get $P(Arc|X) = P(X|Arc)P(Arc)$ where the $P(X|Arc)$ is the scaled arc probability given and $P(Arc)$ is the sum of all the bigram probabilities for that Arc. I'm not very confident in this answer and was wondering if anyone had any ideas about this question?