Uncertainty Representations and Reasoning: A Course on Uncertainty Modeling Beyond Probability Theory*

Erik Quaeghebeur E.QUAEGHEBEUR@TUE.NL
Uncertainty in AI, Eindhoven University of Technology, The Netherlands

During the academic year 2022–2023, the first edition of the course "Uncertainty Representations and Reasoning" took place. It is an elective of 5 ECTS (about 140 hours study load) in the Data Science & Artificial Intelligence Master of Eindhoven University of Technology. It was taken by around 40 students, all familiar only with classical probability and statistics. The course introduces students to uncertainty modeling approaches that go beyond classical probability theory.

Course setup The course is organized in a 'quartile' of eight lecture weeks and two exam weeks, with two weekly contact sessions of two 45 minute 'blocks'. The study activities during the lecture weeks consist of theory lectures, practice exercises, instructions, and working on and presenting an assignment. During the exam week, there is an exam, with a resit option during the exam weeks of the following quartile. The assignment and exam contributed equally to the end grade.

The theory lectures (18 blocks) were classical in style, mixing theory with illustrative examples, but also offering plenty opportunity for interaction. Students were encouraged to interrupt to get clarification and activated by questions concerning interpretation. They made good use of this, which resulted in lively back-and-forths. (Slides can be shared upon request.)

The practice exercises were offered as an ungraded (repeatable) on-line quiz per topic. These quizzes consisted of both theory and calculation questions of either multiple choice or open format. After finalizing the quiz, students were provided with automated feedback and access to model answers. Timely participation by students in the quizzes was problematic.

The instructions (6 blocks) were meant to explain organizational matters—for the assignment, exercises, and exam—and to provide time for asking questions about course material. (There was also an on-line forum for that.) While students were very much interested in practice exercises being worked out in detail, they generally did not make specific requests.

The assignment—done in pairs—consisted of a literature study on a specific application topic, e.g., 'classification'. Students were instructed to digest a number of papers on the topic, each of which used a different uncertainty modeling approach, keeping pace with the lectures. The goal was to explain what they had learned as a consistent and self-contained whole to their fellow students in intermediate and final reports and poster presentations (8 blocks). The assignment grade was based on a participation and peer review component. Most students participated enthusiastically in the assignment.

The 3-hour exam consisted of about 30 questions—both multiple-choice and open—covering all lecture material. The type and format paralleled that of the practice exercises. The expected level of attainment was not achieved in general.

Course content The uncertainty modeling approaches taught during the lectures were, in order of presentation: (classical) probability, belief functions, possibility theory, fuzzy sets, 2-monotone capacities, probability intervals, credal sets, interval probability, and interval expectation. (Alternative terminology such as 'prevision' is mentioned.) After the probability theory lecture, there is moreover a lecture on limitations of probability, to argue for why we need other approaches.

Each approach is discussed in generally the same way, with foundations—such as basic concepts and axioms—, interpretation, and (deductive) inference always being discussed. Often, also learning uncertainty models from data or decision makes an appearance. Much of the actual content was inspired by materials from the SIPTA Schools.

The assignment topics proposed to the students were: classification, clustering, decision trees, Markov chains, and graphical models. For each of these, we provided 4 to 7 papers to the students. Although they were allowed to use other texts as well, most pairs kept to the ones we provided. Due to the fact that the lectures on interval probability and expectation came at the end of the lecture series, the pairs often struggled to integrate materials from papers using such approaches.

Challenges & Lessons learned The biggest challenge we faced was getting the students to actually make the practice exercises in a timely manner and creating time for discussing those during the instructions. The exam results showed that we need to improve this, so the plan is to incentivize students to work on that during the year by making it count for their grade and by creating dedicated exercise instruction sessions. Time for this will be created by removing content.

For the assignment, the biggest challenge for the students was to grasp the content of the papers and connect it to the lecture material. The didactic quality of the papers provided played a crucial role and a good number of the papers we provided were essentially not good enough. Therefore, we plan to improve the list of papers with this in mind.

* Thanks go to Cassio de Campos and Vu-Linh Nguyen for their support during the development and running of the course.