DD2434 Machine Learning, Advanced Course

Course Analysis, mladv15

Hedvig Kjellström

This analysis has been performed by Hedvig Kjellström and is based on an online student questionnare with 23 questions, answered by 16 of the students, and on discussions with students during the course. The raw student input is appended to this analysis.

Course Data

The course was given for the second time, during period 2, 2015, with course leader Hedvig Kjellström. 120 Master and 6 PhD students registered for the course (an increase with 50% from last year!), 87 Master and 5 PhD students completed some part of the course, and 73 Master and 2 PhD students are now finished. The teaching activities consisted of:

- 13 lectures of which five were given by Hedvig Kjellström, four by Carl Henrik Ek, and four by Jens Lagergren,
- 6 tutor sessions of which three were given by Carl Henrik Ek, one by Jens Lagergren, and two by Hedvig Kjellström,
- A presentation session where the student presented their projects to each other.

The examination consisted of two home assignments (4 hp) which were performed individually and examined using written reports in period 2, and one project (3.5 hp) which was performed in groups of four or five and examined by an oral presentation session and a written report in the end of period 2.

The grade distribution on the home assignments is 9 A, 12 B, 11 C, 14 D, and 22 E. The project grades are distributed as 21 A, 16 B, 9 C, 6 D, and 4 E. This leads to a overall course grade distribution of 9 A, 22 B, 13 C, 24 D, and 5 E. PhD students get Pass/Fail.

Due to the very high number of students, we this year left the original intention to examine the home assignments orally. This is unfortunate we believe, since it makes it harder to assess the learning of each student, but necessary given the number of students and the presently allocated teaching resources (three teachers, no teaching assistants). A recommendation for future years is to either introduce a system with oral examination for higher grades, or allocate more TAs to the course – not unreasonable given the number of students.

The book used in the course was (C. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006). We also used (A. Hyvärinen and E. Oja. Independent Component Analysis: A Tutorial. http://cis.legacy.ics.tkk.fi/aapo/papers/IJCNN99_tutorialweb/, 1999), (C. F. Beckmann and S. M. Smith. Probabilistic independent component analysis for functional magnetic resonance imaging. IEEE Transactions on Medical Imaging

23(2):137-152, 2004), (D. M. Blei and J. D. Lafferty. Topic Models. http://www.cs.princeton.edu/~blei/papers/BleiLafferty2009.pdf, 2009), (T. Griffiths. Gibbs sampling in the generative model of Latent Dirichlet Allocation. https://people.cs.umass.edu/~wallach/courses/s11/cmpsci791ss/readings/griffiths02gibbs.pdf, 2002), (Compiled by T. T. Allen. Organization of Scientific Research Papers. http://tim.thorpeallen.net/Courses/Reference/Organization.html, 2000), (D. Duvenaud, O. Rippel, R. P. Adams, and Z. Ghahramani, "Avoiding pathologies in very deep networks," Journal of Machine Learning Research, http://jmlr.org/proceedings/ papers/v33/duvenaud14.pdf, 2014), as well as a number of papers, one for each project.

For student opinions about the course scope and structure, and the book, see the questionnaire Questions 1-6, 22-23. In general, students were happy about the course and the book. From this we conclude that no major changes are needed. The course will next year be given by Jens Lagergren and Pawel Herman, since Carl Henrik Ek has moved to the University of Bristol, UK, and Hedvig Kjellström is on part-time sabbatical at the Max Planck Institute in Tübingen, Germany. This will naturally mean that the course scope changes somewhat – within the boundaries of the learning outcomes:

Upon completion of the course, the student should be able to

1. explain, derive, and implement a number of models for supervised, unsupervised learning,
2. explain how various models and algorithms relate to one another,
3. describe the strengths and weaknesses of various models and algorithms,
4. select an appropriate model or approach for a new machine learning task.

Machine Learning is the study of algorithms that can improve their performance through experience. Experience usually takes the form of data such as labelled and/or unlabelled examples. Machine learning algorithms are used in a vast number of application domains and tasks. However, to do this successfully a machine learning practitioner must have a systematic understanding of how to learn to perform a required task from data.

It is the goal of this course to give you this understanding. We will present a set of machine learning algorithms and statistical modelling techniques. But more importantly you will learn how the different algorithms are developed, how they are related, and how and when they should be used both in theory and practice.

However, both from the questionnaire and from discussions with students during the course, it is apparent that there is a gap between the first course DD2431 and this course DD2434 in terms of the mathematical detail.

It is very important that we make clear to the students already during the beginning of DD2431 (in Period 1) that if they intend to continue with this course DD2434, they need to, in addition to the curriculum of that course, repeat their knowledge in Linear Algebra, Probability Theory and Calculus, in order to be adequately prepared for this course (in Period 2).

Course Activities and Examination

For student opinions about the course activities and examination, see the questionnaire Questions 7-21. The teachers have taken note of any comments concerning their teaching. The students were generally happy about the course activities and examination. Thus we conclude that no major changes are required, and leave it up to the new teachers Jens Lagergren and Pawel Herman to adapt the course according to their liking.