UG Exam Performance Feedback
Second Year
2018/2019 Semester 1

COMP24111 Machine Learning and Optimisation

Ke Chen
Tinting Mu

Comments Please see the attached report.

COMP24111 Examination Performance Feedback (AY2018-2019)

For multiple choice questions in Sect. A, the averaging mark is around 72.9% with a standard deviation of 13.1%. While six questions appeared easy (i.e., over 90% students gave the correct answers to those questions), no question stood out to be extremely hard (i.e. less than 30% students gave the correct answers to those questions). This suggests that most of students well understood those essential concepts and knowledge delivered in this CU.

Question 1 in part B was answered fairly ok, with an average around 7.1 out of 13 (55%). The marks were lost mainly in (c), (e), (f), (g). Most students were able to write down the formulation of the required linear model, write down the sum-of-squares error function and derive its partial derivative, and derive the optimal model weights without the regularization term. However, many students failed in the case of having a regularization term. Also, the part on how to apply the stochastic gradient descent knowledge to train a linear model can be improved. Question 2 in part B was answered very well with an average around 1.8 out of 2 (91%). Among the attempted answers, almost all the students got it right, showing good understanding of the (un)supervised learning concepts. Overall, the written question in part B has an average of 9 out of 15 (60%).

There are three questions in Sect. C. Question 1 is regarding applying the m-estimate to the "zero probability" issue in the discrete Naïve Bayes algorithm. If one fully understands the probability and the m-estimate treatment, the correct answer should be achieved straightforwardly with simple symbolic algebra operations. Unfortunately, only quite few students managed to gain a full mark on this question, which most of them could not use the information given in the question to find out the number of training examples for different attribute values. Questions 2 and 3 are regarding clustering analysis; the second question simply needs the book knowledge and the third question requires the book knowledge along with the essential understanding. More than 50% students managed to get satisfactory marks in those two questions. A common issue for the second question is that a key step regarding choosing proper distance metrics was overlooked in their answers. For the third question, several students could not give the essential implication of the scatter-based F-ratio index although they could gave the correct formula of the F-ration index. In general, the performance in this part suggests that students be able to master the essential aspects of learning algorithms but lack the deep understanding and certain mathematical skills required by machine learning, a goal set for only those exceptional top-tier students.