COMP11212 Exam Performance Feedback AY-17-18

Sean Bechhofer
Giles Reger

June 2018

General Remarks

This was the first year that the exam was taken as a "hybrid" – with two sections online via Blackboard, and two sections taken on traditional paper books. Two valid queries were raised relating to consistency between the paper backup and the online questions. Other than that, no serious issues were encountered.

A total of 195 students sat the exam. The exam was marked out of 58[1], with the mean mark 33.7 (58%).

There were 17 students with a mark below 40%, 8 of these being below 30%. Some of these students *may* be able to pass the unit with suitable coursework marks. There were 34 students who achieved first class marks of 70% or above in the exam.

Table 1 below gives an overview of the performance across questions.

Figure 1 shows the overall performance of the cohort.

		Marks	
Question		Mean	%
Online (Sections A and B)		20.8 (28)	74
C	Q20	1.8 (5)	35
	Q21	3.6 (6)	60
	Q22	2.5 (4)	63
D	Q23	2.3 (7)	33
	Q24	2.0 (4)	50
	Q25	0.9 (4)	22

Table 1: Performance breakdown

Sections A and B

Sections A and B were completed online. The overall performance on this section was good, with a mean score of 20.8 (74%) and a median of 21 (75%). This was to be expected, as the online section contained more questions focused on recall of material as opposed to analysis or application of concepts. Blackboard provides an Item Analysis, including an assessment of the Discrimination power of questions and the Difficulty of questions. Discrimination indicates how well a question differentiates between students who know the subject matter and those who don't. A question is a good discriminator when students who answer the

[1]A missing label on one of the online questions could have introduced a level of ambiguity. As a result, the question was discounted.

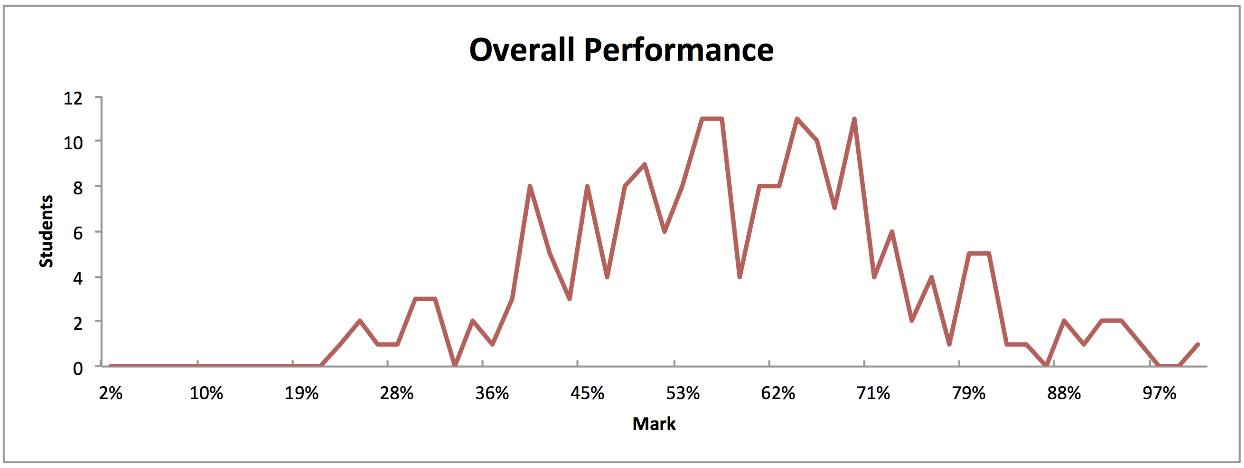

Figure 1: Overall Exam Performance

question correctly also do well on the test. The Difficulty value assesses how many students answered correctly.

Blackboard's analysis suggests that 13 questions were *good* and 5 were *fair* in terms of discrimination. There were 9 *easy* questions, 8 *medium* and 1 *hard*. The *easy* questions were largely in Section A, while the *medium* and *hard* questions were in Section B.

Section C

Question 20

Mean: 1.8 (35%), Median: 1 (20%).

This question was not answered particularly well. The key thing to note in the proposition is that it states that ***Any*** *DFA that accepts precisely the language L must have at least two distinct accepting states.* A common approach here was to translate the given NFA into a DFA and then claim that as it has two accepting states, the proposition is proved. This is not enough, as it doesn't tell us that *any* DFA that accepts the language must have at least two accepting states. It only tells us that *this* DFA has two accepting states.

In addition, in many cases, an answer included a transformation to a DFA that did not accept the same language – there was very little evidence that any checking had been done on the transformations.

To demonstrate the proposition, we have to show that the two accepting states are distinguishable or demonstrate that the automaton is minimal. Answers that failed to include this additional justification received partial credit.

Alternatively, we can argue as follows. If there is a single accepting state X, we must arrive in X after the word 01. The word 010 should also be accepted, so there must be a transition labelled 0 looping back to X. But, 10 should also be accepted, so we also must end up in X after seeing this word (there is only one accepting state). But then we will also accept 100 (due to the transition identified above), which is not in the language. So our original assumption (single accepting state) is wrong. Such an argument was given by a number of students.

Question 21

Mean 3.6 (60%), Median: 3 (50%).

A straightforward question that could be solved using Algorithm 2. Correct answers with no working received full marks. Incorrect answers with no working received some credit if the answer was close, but additional partial credit was given to answers that showed working but were flawed.

Very few students made any obvious attempts to test or validate their answers with example words that would/would not be accepted by the automaton.

Question 22

Mean: 2.5 (63%). Median: 3 (75%).

Answered fairly well overall. Many answers to a) made reference to words with "infinite 'a's". This isn't strictly the case – all our words are finite. Also, the issue is not just with *counting* the number of occurrences of characters – a DFA can count as long as we know there is a bound on the number to be counted. It is the fact that the number we need to count is unbounded that prevents us from producing a DFA.

For part b), the questions explicitly asked for a DFA if one was possible. Some answers just said "yes, you can produce a DFA", without actually *doing so*. A number of answers also produced a DFA that did not actually recognise the language requested.

One or two students did notice that all we actually need to know here is that the language is the language of all words over $\{a, b\}$ with an even number of characters. *Why* this is so is left as an exercise for the reader!

Section D

Question 23

Mean: 2.3 (33%), Median: 2 (29%)

Parts (a) and (b) were answered more commonly than (c) and (d). Parts (a) and (b) were answered well when attempted, with the exception of some answers to (a) that pointed out that a coding function allows us to enumerate `while` programs. This is true but is not a reason why we can restrict our attention to functions in $\mathbb{N} \to \mathbb{N}$.

Parts (c) and (d) were answered less frequently. A lot of answers correctly got the first half of the invariant in (c) as $z = 2^{a-x}$ (or some equivalent form) but many answers forgot to include $x \geq 0$ which is necessary to show that $x = 0$ at the end of the loop. Answers to (d) tended to either produce a reasonably complete solution or did not manage to make any progress towards a correct solution (the assignment rule was frequently applied incorrectly). Some answers to (d) claimed that (c) already contains a proof of the loop invariant but unless this was contained in the answer to (c) this is not the case.

Question 24

Mean: 2.0 (50%), Median: 2 (50%)

The question had the most attempts in Part D.

This question was asking whether the program belonged to each of the classes with justification. Most marks were lost due to failing to mention some of the classes. For example, stating that the program is in $O(n)$ as the outer loop is constant without pointing out that this implies that it is also in $O(n^2)$ as $O(n) \subseteq O(n^2)$. Relatively few students knew what $\Theta(n^2)$ represents. Some justifications said that the program must be in $O(n^2)$ as there are two nested loops, however this is not a strong enough justification in general as it is necessary to comment on the relationship between the loops and n.

Question 25

Mean: 0.9 (22%), Median: 0 (0%)

This question had the least attempts in Part D.

Some answers showed that the set of functions $\mathbb{N} \to \mathbb{B}$ is uncountable but failed to make the last step to argue why that means that the set contains *undecidable* functions (i.e. that the set of `while` programs is countable). The other possible approach was to use the Halting Problem but simply stating the Halting Problem without an argument for its undecidability was not enough for full marks.

There were many answers that simply gave the definition of an undecidable predicate. However, this definition does not give any evidence for such predicates existing. Another common mistake was to show how certain combinations of partially decidable predicates were undecidable. However, this presumes the existence of undecidable predicates as a starting point.