

Teaching Abstract Concepts

Andrea Schalk


April 25, 2006

Abstract

This is a description of a lecture whose aim it is to explain an abstract concept to the students. The chosen method centres around a demonstration in which the students take an active part. In previous years I always had the impression that the students did not really understand the concept in question. Having used this activity (supported by an assessed exercise) I received markedly improved answers in a related exam question.¹ However there are other factors which contributed to the improved results which are sketched below.

Technical background

In the theory of formal languages certain *classes of languages* can be characterized as being those recognized by certain *automata*. An automaton is a theoretical (or mathematical, if you like) construct that consists of a state transition diagram, possibly with some kind of memory. The automata, known as *pushdown automata*, that were the topic of the lecture described below are quite complicated:



At any given time such an automaton is in some state, with a *start state* telling us where to begin. Such an automaton takes as input a word over some alphabet. It has a memory in form of a *stack* on which it may push symbols from the input alphabet as well as certain predefined auxiliary letters. The automaton changes state *non-deterministically* depending on

¹The teaching described here took place as part of CS2121 in 2002/2003.

- the current state;
- the first symbol of the current input word;
- the symbol on top of the stack.

Together with changing its state the automaton may also perform certain actions (which therefore also depend on the above). It can

- erase the first letter from the current input word;
- push or pop any number of symbols onto or from the stack.

There are some predefined conditions to determine whether the input word has been *accepted*, these may involve the current state, the first letter of the current input word, and the current symbol on top of the stack.

Clearly there is a lot going on, and when first trying to understand these automata it can be difficult to get a grip on what's happening. The usual description of such an automaton in mathematical language is as a sextuple, which certainly doesn't help. Even when the information is presented using a state transition diagram as above, the labels on the transitions are quite complicated because they have to be able to express all the dependencies and actions mentioned above.

Issues that may cause confusion

There are a number of issues that may cause confusion when trying to explain how pushdown automata work.

- The transitions have complicated labels, and no matter what syntax is used this is bound to look difficult.
- In particular there are two alphabets to consider, and the manipulation of the stack is complicated to describe.
- The automaton is non-deterministic, so in any given state, and given some input symbol and top of the stack, there may be more than one transition the automaton may follow.
- The memory being in the form of a stack imposes a number of restrictions on its use. This is, of course, perfectly clear to people who use stacks a lot, but this lecture was for students in the first semester of year 2!

Teaching difficult concepts or algorithms

In my experience demonstrating any reasonably complicated algorithm or concept by the lecturer alone only reaches a small proportion of the students. The reasons for this are manifold:

- Almost inevitably the speed of the demonstration will be too fast for some students.
- It's difficult to give sufficient explanation in advance of the demonstration to ensure that the students really can understand what they are seeing.
- If (when) a student is inattentive for even quite a short period of time he is likely to miss a vital part of the explanation, or step, and be lost subsequently.
- Because the entire idea is new to the students they can't digest all the information that is presented to them. They don't have enough of an understanding, however, to realize what to concentrate on.

To give an example: I used to teach the use of stacks in these automata by drawing a picture of the current state of the stack on the board. I, of course, know that I can only access the stack from the top—but how clear is that on a board where anything may be written or wiped out at will?

I wanted the demonstration of a pushdown automaton to satisfy the following criteria:

- It should be slow enough that the students could follow it.
- It should hold their attention and engage them intellectually.
- The use of the stack should be such that there would be no doubt of its nature.

The demonstration

I settled on the following solution to the above issues:

- The demonstration would give an active part to the students, thus encouraging them to follow it.
- I would challenge them to find a 'run' of the automaton in which a given word (*aaabbb*) was accepted.
- I would use a stack of boxes with large letters on the sides to illustrate the use of the stack.

This is the demonstration as I ran it:

- Project the state transition diagram onto the big screen, using an overhead projector. Use a coin or similar to mark the current state.
- Write the word to be accepted on the board.
- Tell the students that I would perform the actions of the automaton, but *they would have to tell me which transition to follow*. I would follow whatever instructions were shouted the loudest.

As is usual when a lecture moves away from speaking at the students from the front there was a certain buzz in the lecture theatre. The first transition had no room for choice, and that got us going. I marked the new state by moving the coin and put a box with the start symbol onto my empty stack. The next transition was obvious to the students, and they called out for me to make it. I complied. The students were gaining confidence, and in quick succession urged me to make further transitions—they had detected a pattern of two transitions we had to keep alternating.

The example was quite carefully chosen: In order to succeed in accepting the word one has to deviate from the pattern at some point where it's not completely obvious. I could see some students realizing that they had missed the correct turning and would not reach the target, while others were still enthusiastically calling out further instructions. Eventually we came to a stop when everybody realized that we were stuck now. The students took this failure in very good spirit, and when I asked them whether they could see now how they would have succeeded most of them said yes. (Given that it's very difficult to get any kind of reply out of a group of students in a lecture theatre I found that very encouraging.) The students certainly had a lot of fun in that lecture. I was worried that I had spent too much time on this; it took rather longer to set up, explain and run than I had thought, and I only had 11 lectures in total to get through a lot of material.

Evaluation

In the past I had always dreaded setting an exam question about pushdown automata, because I had got a lot of very poor answers. In the year were I tried the demonstration described above I set such a question again, and was surprised by the good quality of the average answer. In fact, on my half of the course the average mark shot up quite amazingly; however, a number of other issues probably had some influence on the result:

- There were other attempts to introduce activities into the lectures, although not many and on a much smaller scale, so the demonstration described above was the *pièce de résistance*. It occurred in the fifth out of 11 lectures.

- We had a smaller cohort of students than usual (36 as opposed to 60–100). That allowed a slightly different use of the three examples classes.
- For the first time there was assessed course work for my half of the course. There were 10 assessed exercises, counting for 10% of the final mark, so there was one mark per exercise to be had. In each examples class, 3–4 exercises were due. There were no extensions. Because all answers to the exercises were going to be ‘essentially the same’ marking had to test understanding. What we did was this: The assessor would randomly point at one of the exercises. If the student could explain how they had tried to solve that they would get the marks for all the exercises they had attempted. The students would not be penalized for mistakes in their answers; any proper attempt at finding a solution counted for a mark. In rare cases only half a mark was awarded for some exercises.
- The cohort was split into two examples classes to keep the size manageable, one on Mondays and one on Thursdays. Almost all the students could attend both these classes since there were no timetabling conflicts. We therefore declared the Monday examples class as one where everybody could come to get help with the exercises, but the Thursday one would be for marking only. Many students took advantage of that.

Arguably we effectively doubled the number of examples classes per student, which on a larger scale would require increased staff effort (although I have found my PhD students to be very good as TAs in these classes). On the other hand, the one hour examples classes were replacing two hour labs—maybe the idea that a one hour examples class is somehow equivalent to a two hour lab should be re-examined.. The examples classes were very well attended. While I was disappointed that almost none of the students did the simpler, non-assessed exercises to work up to the assessed ones, all the students I assessed (probably everybody on the course at one time or another) could successfully explain what they had done. I was not concerned about collaboration in solving the exercises because the whole point of the assessed exercises was to get the students to engage with the material and get in some practise. This was very successful as can be seen from the exam results described below. In fact, the lecturer teaching the practical part (supported by labs rather than examples classes) complained that the students were spending all their time on my material!

I had the impression that they found the examples classes rewarding in several ways: They received attention in small groups or individually, and instructors spent some time talking to each student. This helped building a personal relation between students and instructors. The students could

see how their understanding of the issues improved markedly—in fact, one student said: ‘I can see why she makes us do all this stuff now!’ during an examples class. The same students complained quite vocally when similar examples classes weren’t available for my third year course.

Assessment results

I had taught the course in the previous year, but to a larger group of students, without assessed exercises, and without activities in the lectures. Below I compare the results for the course for those two years.

The average mark for the assessed coursework² was 59%. There was no mark inflation despite the ‘lax’ marking criteria because the students decided not to try all the questions. In the previous year examples classes had been offered, although those were larger (50 students per class) and there were no assessed exercises. The take-up rate had been poor.

For both years exam consisted of four questions, two from each half (the practical one, taught by a colleague, supported by labs, the theoretical one, taught by me, by examples classes), of which the students could answer any three. In my opinion the two questions set by me (Questions 3 and 4) were of similar difficulty to the ones I had set the previous year. In the previous year the overall average mark for students on the course had been 54.6%, this year it was 52.9%, so the two populations were comparable as far as their overall ability is concerned.

- Question 3: The average mark was up from 48% to 73%. All students achieved at least the pass mark of 40%.
- Question 4: The average mark was up from 44.5% to 79%, and everybody managed a mark of 55% or better.

Clearly the students had now understood the material so well that a question which had students struggling the previous year were probably too easy this year. The examples class probably helped more with that than the slightly more active lectures. Having one examples class to get help, and another to get marked was something that we could only achieve in the existing timetable because we had rather fewer students. On the other hand the improvements are so marked that it seems obvious that the time so spent was quality time, and it probably makes sense to try to cut time spend on lectures if need be. I did not teach this course again due to other commitments, hence I cannot provide further data about these methods.

²Without my knowledge a scaling factor of .9 had been applied, which I had not intended—unscaled the average would have been 65.6%.