Snapshots of modern mathematics
from Oberwolfach

№ 2/2022

The Robinson–Schensted algorithm

Hugh Thomas[1]

I am going to describe the Robinson–Schensted algorithm which transforms a permutation of the numbers from 1 to n into a pair of combinatorial objects called "standard Young tableaux". I will then say a little bit about a few of the fascinating properties of this transformation, and how it connects to current research.

1 Introduction

I am going to be presenting a mathematical algorithm called the Robinson-Schensted algorithm. The word "algorithm", derived from the name of Persian mathematician al-Khwārizmī (c.780–c.850), is just a fancy way to describe a set of instructions for doing something, but it carries the particular meaning that the instructions are clear and unambiguous. We can think of it as being like the instructions you would write in a computer program.

It turns out, though, that giving precise and unambiguous instructions is a difficult matter. One strategy is to write something in a language modelled on an actual computer language. This aids clarity, but it produces something rather unappealing to most readers. We are not computers, and I do not want to pretend that we are computers in order to communicate, so I will simply write in English. However, even written in a natural human language, the instructions will likely seem somewhat complex and off-putting.

This gets at an interesting fact about mathematics. Mathematics is more than just words on a page. Mathematics is, in fact, something that you do. In

[1] Hugh Thomas is supported by an NSERC Discovery Grant and a Canada Research Chair.

thinking about a mathematical object, whether it's a number, or an equation like $(x+y)^2 = x^2 + 2xy + y^2$, or an algorithm like the one I will be describing, what is going on in your head as you think about it is the actual mathematics. The words on the page are just there as an aid, to prompt you to have mathematical thoughts; it is mathematical thought, and not mathematical words, which are the essence of mathematics.

I think it is instructive to draw the contrast with poetry. Poetry, also, is more than words on the page. It is full of imagery, of emotions, of rhythm, of ideas, of ambiguity, and, in fact, of everything that makes us human. But when we read a great poem, as we let the words wash over us, there is something magical and enticing about the words themselves. This is more or less not true of mathematical writing, with its rigorous attention to clarity, and the absence of a human voice.

Therefore, to read the description of the algorithm that I will give below, it is not enough to let the words wash over you, and be carried away in their rich emotional suggestiveness. Indeed, I can pretty much guarantee that that approach will not work at all. The way to read such a description is to try it out, to see it in action for yourself. If you try it out, you will be experiencing mathematics directly. And you will see, I hope, that there is something fascinating about it.

2 The Robinson-Schensted algorithm

We are going to start with a permutation of the numbers 1 to n, for some whole number n, that is to say, we start with a list of the numbers from 1 to n in some order. As a warm-up, and as an invitation to the kinds of things we will be thinking about, how many permutations of the numbers from 1 to n are there? The answer, as you may well know, is $n!$, which is a notation for the product $n(n-1)(n-2)\dots 1$. This counts the permutations because there are n ways to choose the first number, and then, for each of those ways, there are a further $n-1$ ways to choose the second number (the first number no longer being available) and so on until the last number is the only one remaining, so we have only one choice at the last step.

It is good to think about even very small values of n as examples. If $n = 1$, there is just one permutation, 1. If $n = 2$, there are two permutations, $1, 2$ and $2, 1$. If $n = 3$, there are 6. And so it goes, increasing more and more quickly as n increases; when n is 10, the number is already 3628800, and when $n = 100$, the number is more than scientists' best guess at the number of atoms in the universe.

We are going to carry out a strange procedure on each permutation of the numbers from 1 to n. It is going to transform this list of numbers into a pair of what are called *standard Young tableaux* (or SYT for short), named after Alfred Young (1873–1940), who introduced them in the paper [5]. An SYT of size n is a two-dimensional arrangement of the numbers from 1 to n, in which each number appears exactly once. They are arranged in left-justified rows, one below the other, and no row has more entries than the one above it. Also, the entries in each row are in increasing order as you read them from left to right, and the entries in each column are in increasing order as you read them from top to bottom.

For example, if $n = 2$, there are two SYTs:

1	2

1
2

and if $n = 3$ there are 4:

1	2	3

1	2
3	

1	3
2	

1
2
3

We call the shape of a standard Young tableau the list of the lengths of its rows, from top to bottom (so in weakly decreasing order). Thus, when $n = 3$, there is one SYT of shape (3), one of shape (1,1,1), and two of shape (2,1). There is a beautiful formula for the number of standard Young tableaux of any shape [1], but I can't stop to explain all the different beautiful topics that connect to this subject or I would never be finished. You might try to work out for youself the number of standard Young tableaux whose shape is (m, m).[2]

The Robinson-Schensted algorithm takes a permutation of the numbers 1 to n and produces two SYTs from it. They are traditionally denoted P and Q, and called the "insertion tableau" and the "recording tableau," respectively.

I will explain the insertion tableau first. At every step, we will make sure that each row in the tableau we have generated so far is increasing from left to right. Say our permutation is $(a_1, a_2, \ldots, a_n)$, that is to say, we will call the first number a_1, the second number a_2, and so on. We take the first number a_1 in our permutation, and we put it at the top left of our tableau. Then we take the second number a_2. We want to put it into the first row. If $a_2 > a_1$, then we can add it at the end of the row, and we do that. But if $a_2 < a_1$, we cannot insert a_2 at the end of the row, because the result would be a row which is not increasing. Instead, we remove a_1, we put a_2 into the spot vacated by a_1, and we put a_1 in the second row.

[2] We have already seen that when $m = 1$, there is one tableau. When $m = 2$, there are two. If you work out a few more and type the sequence into Google, you will discover a famous series of numbers studied by Leonhard Euler (1707–1783), along with many other people.

If our permutation had begun $3, 5$, then this is what would have happened:

$$\begin{array}{|c|}\hline 3\\ \hline\end{array} \;\to\; \begin{array}{|c|c|}\hline 3 & 5\\ \hline\end{array}$$

If it had begun $4, 2$, then this is what would have happened:

$$\begin{array}{|c|}\hline 4\\ \hline\end{array} \;\to\; \begin{array}{|c|}\hline 2\\ \hline 4\\ \hline\end{array}$$

Now, we consider a_3. Again, we will add it to the first row. If it is larger than all the numbers in the first row, we add it at the end. If not, we find the first number in the first row that is bigger than it, and we replace that number by a_3. We now take the number that we replaced, and we insert it into the second row, in the same way as before. This may cause a number to be "bumped" out of the second row and inserted into the third row. We keep going in the same way all through the permutation, always inserting the next number from our permutation starting at the first row of the tableau. The paths of possible "bumps" get longer and longer, but they always work their way down the tableau one row at a time, until we reach the point where the number which we are currently trying to insert can be placed at the end of the row.

Let us make this clearer by doing a couple of examples. If we start with the permutation $5, 2, 4, 1, 3$, then the process of building the insertion tableau would look like this:

$$\begin{array}{|c|}\hline 5\\ \hline\end{array} \to \begin{array}{|c|}\hline 2\\ \hline 5\\ \hline\end{array} \to \begin{array}{|c|c|}\hline 2 & 4\\ \hline 5 \\ \cline{1-1}\end{array} \to \begin{array}{|c|c|}\hline 1 & 4\\ \hline 2 \\ \cline{1-1} 5 \\ \cline{1-1}\end{array} \to \begin{array}{|c|c|}\hline 1 & 3\\ \hline 2 & 4\\ \hline 5 \\ \cline{1-1}\end{array}$$

For a second example, take the permutation $2, 6, 1, 5, 4, 3$. This time we obtain:

$$\begin{array}{|c|}\hline 2\\ \hline\end{array} \to \begin{array}{|c|c|}\hline 2 & 6\\ \hline\end{array} \to \begin{array}{|c|c|}\hline 1 & 6\\ \hline 2 \\ \cline{1-1}\end{array} \to \begin{array}{|c|c|}\hline 1 & 5\\ \hline 2 & 6\\ \hline\end{array} \to \begin{array}{|c|c|}\hline 1 & 4\\ \hline 2 & 5\\ \hline 6 \\ \cline{1-1}\end{array} \to \begin{array}{|c|c|}\hline 1 & 3\\ \hline 2 & 4\\ \hline 5 \\ \cline{1-1} 6 \\ \cline{1-1}\end{array}$$

The recording tableau is easier to describe. It has the same shape as the insertion tableau, but we just number the boxes by the step at which that box appeared in the insertion tableau. So, looking at the two example insertion tableaux that we worked out above, we see that the corresponding recording tableaux are, respectively:

1	3
2	5
4	

and

1	2
3	4
5	
6	

Note that it is easy to check that the recording tableau is a standard Young tableau. It contains the numbers from 1 to n once each, since on each step, one new box is added to the insertion tableau. Also, because a box can only be added after the boxes directly above it and directly to its left are added, it is clear that entries in the recording tableau increase along rows and down columns.

The insertion tableau is also a standard Young tableau. Since we started with a permutation, it is clear that it will contain the numbers from 1 to n once each. And by the construction, it is clear that the rows are always increasing from left to right. It is harder (but not too hard) to convince yourself that its columns are also always increasing, so that it is also a standard Young tableau.

3 Properties of the Robinson–Schensted algorithm

We have already seen that the Robinson–Schensted algorithm converts a permutation of 1 to n into a pair of standard Young tableaux P and Q of size n. The tableaux P and Q obviously have the same shape, because of the way the recording tableau Q is defined. Now the somewhat surprising fact is that given any pair P and Q of SYTs of the same shape, either two different ones or the same one twice, there is exactly one permutation π such that applying the RS algorithm to π produces P as the insertion tableau and Q as the recording tableau.

How could you prove this? Well, you can easily check it on small examples, and, if you are interested, it would be a really good idea to try out the examples of $n = 3$ and $n = 4$. To prove it, though, requires more. The simplest way to prove it is to demonstrate how, given P and Q, it is possible to find π. If you think about it hard enough, you will see that knowing P and Q gives you enough information to gradually reverse the RS correspondence and reconstruct π in reverse (first the last entry of π, then the second-last, and so on).

For example, let's consider the pair (P, Q) that we generated above starting from the permutation 5, 2, 4, 1, 3. Let's imagine we didn't know that we had started with 5, 2, 4, 1, 3, and let's try to deduce it. By finding the 5 in the recording tableau, we know that when we were inserting the fifth entry in the permutation, we eventually stopped when we were inserting the 4 into the second row. By looking at the first row, we see that the element that bumped out the 4 must have been the 3. So on the fifth step, we must have started by inserting the entry 3. This tells us that the permutation we started from must have ended with 3, which is indeed true. But now, we can figure out what the insertion tableau looked like before we inserted the 3, and we can then repeat the same argument to discover the remaining elements of the permutation.

One consequence of this one-to-one correspondence between permutations and pairs of standard Young tableaux of the same shape is a rather peculiar equation. Let's write f_λ for the number of SYTs of shape λ. Then, since every permutation of 1 to n corresponds to exactly one pair of SYTs of the same shape composed of n boxes, we see that $n!$ equals the sum of f_λ^2, over all possible shapes λ with n boxes.

For example, when $n = 3$, we said $f_3 = 1$, $f_{2,1} = 2$, and $f_{1,1,1} = 1$. And, indeed, $6 = 1^2 + 2^2 + 1^2$.

Now that we have this correspondence between permutations and pairs of SYTs, we can see how interesting properties on one side correspond to interesting properties on the other side. For example, in a permutation π, we could ask about the maximum length of an increasing subsequence. It turns out that this corresponds to the length of the first row of the shape of the tableaux corresponding to π. For example, in the permutation we looked at before, 5, 2, 4, 1, 3, the increasing subsequences (of length more than 1) are 2, 4; 2, 3; and 1, 3. The maximum length is therefore 2, and this does agree with the length of the first row of SYTs associated to π, which is 2.

It might seem obvious that the length of the first row should have something to do with the longest increasing subsequence, because the first row is itself increasing. However, if you look carefully, you will see that the first row of the insertion tableau is not necessarily an increasing subsequence of the permutation you started with. For example, the result of inserting the permutation $1, 3, 4, 2$ is the tableau

1	2	4
3		

and we see that 1, 2, 4 is not an increasing subsequence of 1, 3, 4, 2. However, there is indeed a length three increasing subsequence, namely 1, 3, 4.

Can you find an interpretation for the length of the first column of the shape of the pair of tableaux corresponding to π?

4 Robinson–Schensted and quiver representations

The facts that I have presented here about the Robinson–Schensted correspondence date back to work of Robinson [3] and Schensted [4] in, respectively, the 1930's and the 1960's. So why was this being discussed at Oberwolfach in 2020?

Roughly speaking, the idea is as follows. Whereas up until now, we have been considering permutations as nothing but a list of numbers, it turns out that, associated to a permutation, there is an algebraic object called a "quiver representation". It would require too much background to explain what quiver representations are, so I am going to have to be rather imprecise. The key fact, for my purposes, about quiver representations, is that if we have a quiver representation X, we can talk about functions from X to itself, just as you might study functions from $\mathbb{R}$ to $\mathbb{R}$. It turns out that if you start with a permutation π, and you take the associated quiver representation X_π, and then you consider the collection of maps from X_π to itself, then, by studying that collection of maps, you can find encoded in them the pair of tableaux which result from applying the Robinson-Schensted correspondence to π. This newly observed connection between seemingly very different mathematical objects allows us to generalize the Robinson–Schensted correspondence in a novel way.

5 Further reading and acknowledgements

A possible source for further reading about the Robinson–Schensted correspondence is the book [2] by William Fulton, to whom this snapshot is gratefully dedicated.

I would like to thank the organizers of the Oberwolfach workshop number 2004 on finite-dimensional algebras for the chance to present there the work described in this note, and for the invitation to prepare this snapshot. Thanks to Henning Krause for reading the manuscript and providing helpful suggestions.

References

[1] J. S. Frame, G. de B. Robinson, and R. M. Thrall, *The hook graphs of the symmetric group*, Canadian Journal of Mathematics **6** (1954), 316–324.

[2] W. Fulton, *Young tableaux*, Cambridge University Press, 1997.

[3] G. de B. Robinson, *On the representations of the symmetric group*, American Journal of Mathematics **60** (1938), 745–760.

[4] C. Schensted, *Longest increasing and decreasing subsequences*, Canadian Journal of Mathematics **13** (1961), 179–191.

[5] A. Young, *On quantitative substitutional analysis*, Proceedings of the London Mathematical Society **33** (1900), 97–145.

Hugh Thomas *is a professor at the Université du Québec à Montréal. He studies algebraic combinatorics and representation theory.*

Mathematical subjects
Algebra and Number Theory, Discrete Mathematics and Foundations

DOI
10.14760/SNAP-2022-002-EN

Snapshots of modern mathematics from Oberwolfach provide exciting insights into current mathematical research. They are written by participants in the scientific program of the Mathematisches Forschungsinstitut Oberwolfach (MFO). The snapshot project is designed to promote the understanding and appreciation of modern mathematics and mathematical research in the interested public worldwide. All snapshots are published in cooperation with the IMAGINARY platform and can be found on www.imaginary.org/snapshots and on www.mfo.de/snapshots.

ISSN 2626-1995

Mathematisches Forschungsinstitut
Oberwolfach gGmbH
Schwarzwaldstr. 9 – 11
77709 Oberwolfach
Germany

Mathematisches Forschungsinstitut Oberwolfach

Member of Leibniz Association

IMAGINARY open mathematics