A Survey on Bijections with Young Tableaux

Jackie Bredenberg

May 14, 2018

Abstract

We survey the results of two papers, by Schensted (1961) and Knuth (1970). We discuss how bijections between Young tableaux and permutations, sequences, and matrices lead to counting formulae. We consider both standard Young tableaux and generalized Young tableaux (which allow repeated numbers).

1 Introduction

In this paper, we use bijections with Young tableaux to answer a number of counting questions, including:

- Given a permutation, what's the longest increasing subsequence contained in it?
- Given a sequence, what's the longest nondecreasing subsequence contained in it?
- How many binary sequences of length n have a longest nondecreasing subsequence of length m?
- How many matrices are there of nonnegative integers whose entries sum to n?

Definition 1.1. A standard Young tableau (SYT) of shape λ, where λ is a partition of an integer n, is a series of rows of boxes of lengths $\lambda_1, \lambda_2, ..., \lambda_k$, filled with the numbers $1, 2, ..., n$ so that numbers increase going right across rows and down columns.

For example, if $n = 8$ and $\lambda = (4, 2, 2)$, one of many valid SYT is:

1	2	4	7
3	6		
5	8		

Definition 1.2. The hook length $h_{i,j}$ of a square (i, j) in a SYT is the number of squares in the "hook-shape" to the right and below, including the square itself.

In the example tableau above, the hook lengths of the squares are 6, 5, 2, 1; 3, 2; 2, 1.

The following counting formula is known regarding Young tableaux (see [G] for an elegant proof).

Theorem 1.1 (Hook length formula)**.** Let f_λ denote the number of SYT of shape λ. Then

$$f_\lambda = \frac{n!}{\prod\limits_{i,j} h_{i,j}}$$

Definition 1.3. A generalized (or semi-standard) Young tableau (GYT) of shape λ is the same as a standard Young tableau, except that numbers may be repeated or omitted. Columns must still be in increasing order, but rows are only required to be nondecreasing, so the same number can appear twice in a row.

For example, with $\lambda = (4, 2, 2)$ again, one GYT is:

1	1	2	3
2	3		
4	4		

We denote the number of GYT of shape λ whose largest entry is k by $f_\lambda^{G,k}$. (We do not provide a formula for $f_\lambda^{G,k}$ in this paper.)

We now state 3 main theorems that the rest of this paper will cover:

Theorem 1.2 (RSK Correspondence)**.** Permutations π of $1, ..., n$ are in bijection with pairs (P, Q) of SYT of the same shape λ. The number of columns λ_1 is equal to the length of the longest increasing subsequence of π, and the number of rows $|\lambda|$ is equal to the length of the longest decreasing subsequence of π.

Theorem 1.3 (Schensted 1961)**.** Sequences $x_1, ..., x_n$ of numbers in the range $1, ..., k$ are in bijection with pairs (P, Q) of Young tableaux of the same shape λ where λ partitions n, P is a GYT with largest entry k, and Q is a SYT. The number of columns λ_1 is equal to the length of the longest *nondecreasing* subsequence of π, and the number of rows $|\lambda|$ is equal to the length of the longest decreasing subsequence of π.

Theorem 1.4 (Knuth 1970)**.** Matrices A of nonnegative integers are in bijection with pairs (P, Q) of GYT of the same shape λ where λ partitions $\sum\limits_{i,j} A_{i,j}$, P is a GYT with largest entry equal to the number of columns of A, and Q is a GYT with largest entry equal to the number of rows of A.

2 Permutations and the RSK Correspondence

The RSK correspondence [S] is a bijection between permutations of the numbers $1, ..., n$ and pairs (P, Q) of same-shape SYT. Given a permutation $\pi = (\pi_1, ..., \pi_n)$, we construct P and Q as follows:

Algorithm 2.1. (RSK)
Begin with empty tableaux P and Q.
For $i = 1$ to n:

- Insert π_i into the first row of P, "bumping" the smallest number in the first row that is greater than π_i.
- Whenever a number is bumped, insert it into the next row down, similarly bumping the smallest number greater than it.
- When an inserted number is greater than any number in the row, "insertion" means appending it to the end of the row.
- Let (a, b) be the position in P where a number ends up being appended. Add i to the (a, b) position in Q, so that Q represents the order of adding squares to P.

We show this algorithm performed on the permutation $\pi = (2, 5, 3, 7, 6, 1, 4)$:

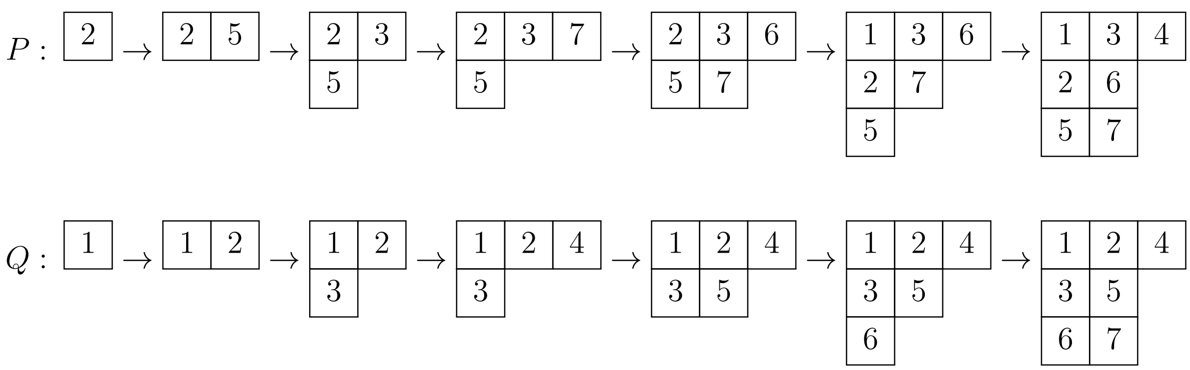

Lemma 2.1. For any permutation, the P and Q resulting from this algorithm are SYT.

Proof. Every insertion into P preserves the ordering of rows: an element is inserted at the position of the first element greater than it, so given that the rows were ordered left to right before, all elements further right of the inserted element are greater than it and all elements to the left are less than it.

An element is always smaller than the element it bumps, so it must also be smaller than elements below that element in its column. Furthermore, it always starts from a position at least as far right as the element it bumps: if it started to the left it would be smaller than the element directly below it, and could not bump an element further right than that. If it started in the same column as the element it bumps, it will be larger than the element that ends up above it, because that element bumped it. If it started strictly further right than the element it bumps, it will be larger than the element that ends up above it, because it was previously to the right of that element. Therefore the increasing order of columns is also preserved under insertion.

By induction, the resulting P will also have increasing order across each row and down each column, so it is a SYT.

At each step of adding a number to Q, that number is the largest number in Q, and it's placed in a spot appended to P, which is therefore a position at the far right/bottom of its row/column. Since it's bigger than everything in its row/column, the ordering is maintained, and by similar inductive logic to P, Q is a SYT. □

Lemma 2.2. This correspondence between permutations and pairs (P, Q) is a bijection.

Proof. We have already shown that any permutation maps to a unique pair of SYT. We now show that given any (P, Q) of the same shape, we can construct a permutation that maps to them, proving that the map is bijective.

Let P_n be the number in P in the same position as the number n is in Q. Since n was the last number added to Q, we know that P_n is correspondingly in the last square appended to P. We reconstruct what P must have looked like before the last insertion by "reverse-bumping" P_n up to replace the largest number less than it in the row above it, then moving that number up to the row above it, and so on until a row is reverse-bumped from the first row. The number removed from the first row is π_n, the last number inserted into P. We can continue this process iteratively: let P_{n-1} be the number in the square corresponding to $n-1$ in Q, reverse-bump it upward until we remove π_{n-1}, and so on. Reverse-bumping is always possible because any number can find a smaller number in the row above it, and therefore the greatest smaller number exists. So this procedure always gives us a unique π corresponding to any pair (P, Q). □

Corollary 2.1. By counting sequences and pairs of tableaux, and using the previous bijection, we conclude

$$n! = \sum_{\lambda \vdash n} f_\lambda^2$$

We now introduce basic subsequences of a permutation in order to show a correspondence between Young tableau shape and increasing/decreasing subsequences.

Definition 2.1. The j^{th} basic subsequence of a permutation refers to the sequence of elements initially inserted into the j^{th} column of the first row of P during the RSK algorithm.

For example, for the permutation $\pi = (2, 5, 3, 7, 6, 1, 4)$ above, we find that the 1st basic subsequence is $2, 1$, the 2nd is $5, 3$, and the 3rd is $7, 6, 4$. Note that each basic subsequence will always be decreasing, because a later number inserted into the j^{th} position of the first row must bump the previous one and therefore be smaller than all previous ones.

Lemma 2.3. The length of the largest increasing subsequence in π is equal to the number of columns in P (or Q).

Proof. Let the number of columns of P be denoted k. Any increasing subsequence is at most length k, since the numbers are partitioned into k basic subsequences, and an increasing subsequence can have at most one element from each decreasing subsequence.

We can construct an increasing subsequence $x_1, ..., x_k$ as follows: let x_k be the element that ends in the k^{th} column. Iterate backward: for each $i < k$, let x_i be the element that was in the i^{th} column at the time that x_{i+1} was inserted. x_i must come before x_{i+1} in the permutation, or else it would not already be there when x_{i+1} was inserted. x_i must also be less than x_{i+1} or else they could not simultaneously exist with x_i in the earlier column. Therefore $x_1, ..., x_k$ is an increasing subsequence of the permutation. □

Lemma 2.4. The length of the largest decreasing subsequence in π is equal to the number of rows in P (or Q).

Proof. We omit the proof here. The fundamental idea of the proof is to show that running the RSK algorithm on the reversed permutation interchanges increasing and decreasing subsequences, and interchanges P with its transpose. □

Together, these statements contribute to the aforementioned main theorem for this section:

Theorem 2.1 (RSK Correspondence)**.** Permutations π of $1, ..., n$ are in bijection with pairs (P, Q) of SYT of the same shape λ. The number of columns λ_1 is equal to the length of the longest increasing subsequence of π, and the number of rows $|\lambda|$ is equal to the length of the longest decreasing subsequence of π.

Corollary 2.2. The number of permutations $x_1, ..., x_n$ with longest increasing subsequence length α and longest decreasing subsequence length β is $\sum_{\lambda \in L} f_\lambda^2$ where L is the set of λ with α columns and β rows.

3 Extending to Sequences

Sequences are similar to permutations, but can have repeated or skipped numbers. [S] notes that we can still use the RSK algorithm, but due to the repeat numbers, we end up with a GYT for P. (Q still describes the order of adding squares and is thus a SYT.) We show this algorithm performed on the sequence $(1, 2, 2, 1, 4, 1)$.

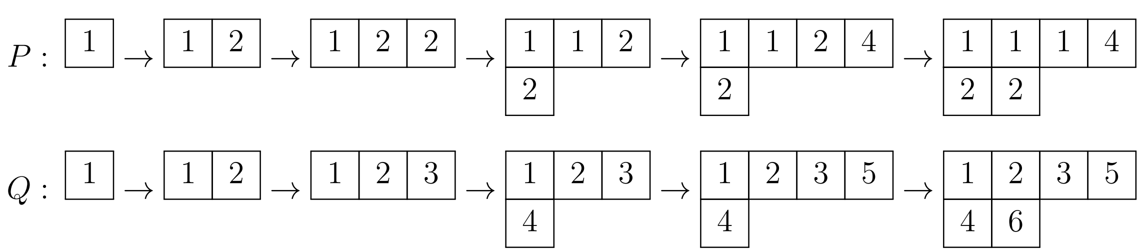

In a GYT, the rows are nondecreasing while the columns are increasing, so the number of columns corresponds to the longest nondecreasing sequence and the number of rows corresponds to the longest decreasing sequence. From this we get a theorem similar to the one for permutations:

Theorem 3.1 (Schensted 1961)**.** Sequences $x_1, ..., x_n$ of numbers in the range $1, ..., k$ are in bijection with pairs (P, Q) of Young tableaux of the same shape λ where λ partitions n, P is a GYT with largest entry k, and Q is a SYT. The number of columns λ_1 is equal to the length of the longest *nondecreasing* subsequence of π, and the number of rows $|\lambda|$ is equal to the length of the longest decreasing subsequence of π.

Corollary 3.1. The number of sequences $x_1, ..., x_n$ of numbers from 1 to k with longest nondecreasing subsequence length α and longest decreasing subsequence length β is $\sum_{\lambda \in L} f_\lambda^{G,k} f_\lambda$ where L is the set of λ with α columns and β rows.

As a reminder, $f_\lambda^{G,k}$ counts generalized Young tableaux filled with numbers between 1 and k inclusive.

We now demonstrate an application of this formula by computing the number of binary sequences of length n that have longest increasing subsequence length m. Note that since columns are still strictly decreasing, each column will have length at most 2 (a zero followed by a one), and the GYT will look like:

0	0	0	0	0	0	1	1	1
1	1	1	1					

The first row has length equal to the longest increasing subsequence, m, so the remaining $n - m$ elements are in the bottom row. The $2m - n$ elements in the top row without any squares beneath them can be either zeroes or ones, so long as all zeroes are left of the ones, so there are $2m - n + 1$ possibilities for where to switch from zeroes to ones. Therefore, there are $2m - n + 1$ different tableaux of that shape for P. Meanwhile Q will fill the shape with the numbers $1, ..., n$, so we can use the hook length formula to find the number of possible SYT. Elements in the upper left corner have hook lengths $m + 1, m, ..., 2m - n + 2$. Elements in the upper right have hook lengths $2m - n, 2m - n - 1, ..., 1$. Elements in the lower left have hook lengths $n - m, n - m - 1, ..., 1$. In total, the product of the hook lengths is $\frac{(m+1)!}{(2m-n+1)!}(2m - n)!(n - m)!$, so dividing $n!$ by that, we see that the number of possibilities for Q is $\frac{n!(2m-n+1)}{(m+1)!(n-m)!}$. Therefore, our answer for the number of binary sequences with the specified length of increasing subsequence (equal to the number of pairs (P, Q) of that shape) equals

$$\frac{n!(2m - n + 1)^2}{(m + 1)!(n - m)!}$$

4 Matrices as Generalized Permutations

We have now shown bijections between combinatorial objects and pairs of (SYT, SYT) and (GYT, SYT). A natural question would be whether there exists an object that bijects with pairs of (GYT, GYT) of the same shape. As shown in [K], the answer is yes: nonnegative integer matrices.

Let A be such a matrix with r rows and c columns, whose elements sum to n. We can describe the entries of A by mapping them to pairs (i, j) representing the row and column where the entry is, with multiplicity equal to the value of the entry. We then write these pairs in vectors sorted by row, then column. For example:

$$\begin{bmatrix} 0 & 1 & 0 \\ 2 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \longleftrightarrow \begin{bmatrix} 1 & 2 & 2 & 2 & 4 \\ 2 & 1 & 1 & 3 & 2 \end{bmatrix}$$

We can biject this vector of pairs to a pair of GYT of the same shape as follows: consider the bottom row as a sequence, and insert it into a tableau P. On the k^{th} step, instead of adding k to Q, we add the k^{th} entry in the top row to Q. We demonstrate this with the same matrix as above:

$$P: \begin{array}{|c|}\hline 2 \\ \hline\end{array} \rightarrow \begin{array}{|c|}\hline 1 \\ \hline 2 \\ \hline\end{array} \rightarrow \begin{array}{|c|c|}\hline 1 & 1 \\ \hline 2 \\ \cline{1-1}\end{array} \rightarrow \begin{array}{|c|c|c|}\hline 1 & 1 & 3 \\ \hline 2 \\ \cline{1-1}\end{array} \rightarrow \begin{array}{|c|c|c|}\hline 1 & 1 & 2 \\ \hline 2 & 3 \\ \cline{1-2}\end{array}$$

$$Q: \begin{array}{|c|}\hline 1 \\ \hline\end{array} \rightarrow \begin{array}{|c|}\hline 1 \\ \hline 2 \\ \hline\end{array} \rightarrow \begin{array}{|c|c|}\hline 1 & 2 \\ \hline 2 \\ \cline{1-1}\end{array} \rightarrow \begin{array}{|c|c|c|}\hline 1 & 2 & 2 \\ \hline 2 \\ \cline{1-1}\end{array} \rightarrow \begin{array}{|c|c|c|}\hline 1 & 2 & 2 \\ \hline 2 & 4 \\ \cline{1-2}\end{array}$$

Lemma 4.1. *P and Q (generated in this way from a vector of lexicographically-ordered pairs) are GYT.*

Proof. P is just the result of following the insertion algorithm on a sequence, which we have already shown produces a GYT. At every step of adding an element to Q, it is greater than or equal to all previous elements added and is maximal in its row and column. Since rows only need to be weakly increasing, the row constraint remains satisfied. It is clear that column numbers will also be at least weakly increasing, but we still need to show that they are strictly increasing, i.e. equal numbers will not be placed in the same column in Q.

To do this, we show that if two equal numbers are placed in Q, the later one will be placed strictly further to the right than the earlier one. Since pairs in the vector are sorted by top row then bottom row, if two pairs have the same number to insert into Q, the one with greater or equal number to insert into P will come later. A later and greater-or-equal number inserted into P will bump (or append) a position in the first row strictly right of the earlier number.

CASE 1. If the earlier number appends, all later numbers with that Q-value will too, so they'll be strictly further right.

CASE 2. If the earlier number bumps something, the chain of bumps can only cause the final append position to move further left, because the bumped number is strictly less than subsequent numbers in its column.

CASE 2a. If the later number appends, it will thus be strictly further right than the earlier number.

CASE 2b. If the later number also bumps something, it's further right, so the bumped value will be greater than or equal to what the earlier number bumped. Then both bumped numbers move down to the second row to bump or append there, and we're in the same situation as we were when we inserted an earlier number and a later greater number into the first row. So by inductive logic, the final append position of the later number is strictly further right.

Therefore the columns of Q are strictly increasing, and Q is a valid GYT. □

Lemma 4.2. Given any GYT (P, Q) generated in this way, we can reconstruct a unique vector of pairs (and therefore a unique matrix) that maps to it.

Proof. We use the same strategy as before: use Q to determine which position was appended last to P, and reverse-bump it iteratively to uncover what number was inserted last to P. Remove that and the last number inserted to Q, and repeat to recover the rest of the sequence of pairs backward.

It is now slightly trickier to determine which position was last appended to Q, because it's not just the number n. However, since the upper-row values were sorted, we know it is one of the numbers tied for largest. In fact, using what we proved in the previous lemma, we know that it is always the rightmost among those tied numbers. □

It is interesting to note that this correspondence for matrices is the same as the correspondence for permutations in the special case where we have a permutation matrix (one 1 in each row/column and the rest 0s). With a permutation matrix, our vector of pairs would have one copy of each number $1, ..., n$ in each of the top and bottom row; sorted lexicographically, the top row is in order, so we add 1 to n in order to Q, and both P and Q are SYT. Thus we could consider our result for permutations as a special case of this result for matrices.

The size of the tableaux is equal to the number of pairs in the vector, which is equal to the sum of entries in the matrix. Entries in the top row of the vector (inserted into Q) will range from 1 to the number of rows r, and entries in the bottom row (inserted into P) will range from 1 to the number of columns c. Therefore, we have shown our third and final goal theorem for this paper:

Theorem 4.1 (Knuth 1970)**.** Matrices A of nonnegative integers are in bijection with pairs (P, Q) of GYT of the same shape λ where λ partitions $\sum_{i,j} A_{i,j}$, P is a GYT with largest entry equal to the number of columns of A, and Q is a GYT with largest entry equal to the number of rows of A.

Corollary 4.1. The number of $r \times c$ matrices A of nonnegative integers with entries summing to n is $\sum_{\lambda \vdash n} f_\lambda^{G,r} f_\lambda^{G,c}$.

5 References

[G] C. Greene, A. Nijenhuis, and H. Wilf, A Probabilistic Proof of a Formula for the Number of Young Tableaux of a Given Shape, Advances in Mathematics. Volume 31, Issue 1 (1979), 104-109.

[K] Donald E. Knuth, Permutations, matrices, and generalized Young tableaux, Pacific J. Math. Volume 34, Number 3 (1970), 709-727. http://projecteuclid.org/euclid.pjm/1102971948

[S] C. Schensted, Longest increasing and decreasing subsequences, Canadian Journal of Mathemetics 13 (1961), 179-191. https://cms.math.ca/10.4153/CJM-1961-015-3