Acknowledgment

I first saw the method used here in James P. Keener's paper 'The Perron-Frobenius Theorem and the Ranking of Football Teams' in **SIAM Review**, volume 35, number 1, pages 80-93.

The Justification

Suppose that the unimaginatively-named teams A, B, C, D, E, and F have been playing each other with the results recorded in the below picture. An arrow from one team to the other indicates that the first team has beaten the second.

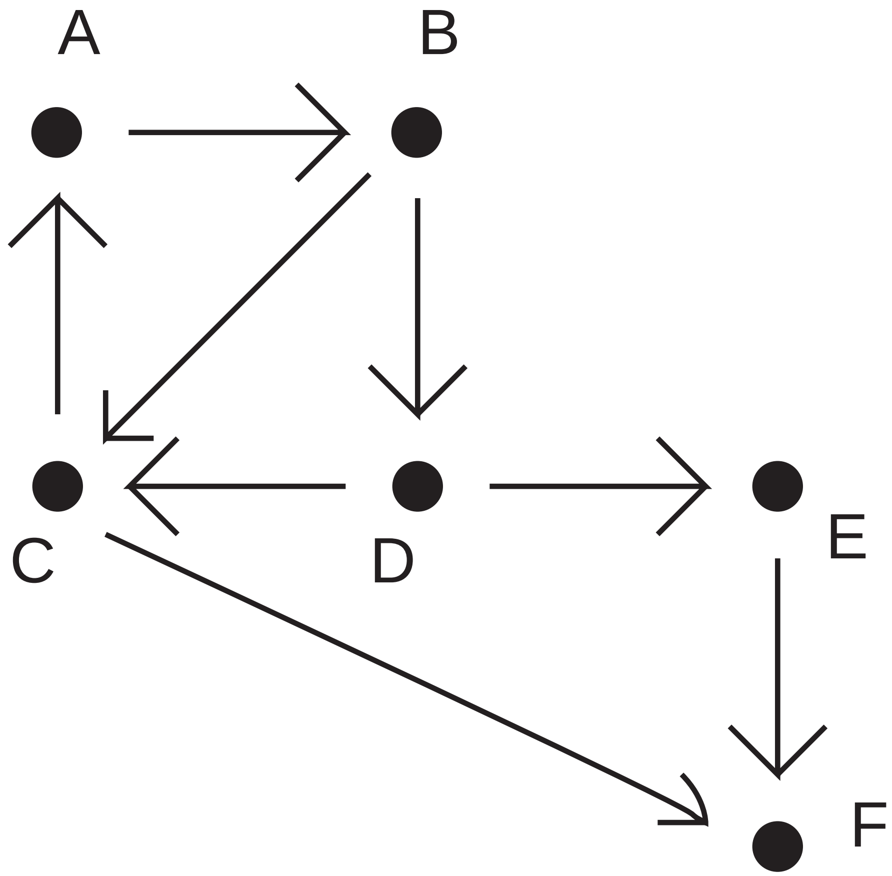

So, for example, we see that team B has beaten teams C and D. (At this point, we are not interested in the score of the games, just the outcome.) Note that different teams have played different number of games. (In addition, the mathematical term for the below picture is a **directed graph**.)

Here's the main question: how should we rank the teams in order to reflect the results. Clearly, team F is the worst team, but how (for example) do teams A and D compare?

To start our discussion, let's find the winning percentage for each team. For example, since team C played four games and beat teams A and F, its winning percentage is $\frac{1+1}{4}$ or .500. (Yes, that was a trivial computation, but we're setting the stage to come.) Here's a table of the teams with the winning percentages (to three decimal places):

A	B	C	D	E	F
.500	.667	.500	.667	.500	.000

As you can see, we now consider teams B and D to be tied for first. But, you might ask, didn't B beat better teams than D did? And surely team A is better than team E? Somehow, we would like to represent the **strength of schedule** into our rankings. To do this we will recompute the winning percentages, except we will replace the 1's we used above with the beaten team's winning percentage. As an example, the new number for B will be $\frac{.5+.667}{3}$ since B beat both C and D. As another example, the new number for E will be $\frac{0}{2}$ since E only beat F, a team that won no games. Here's the new table:

A	B	C	D	E	F
.333	.389	.125	.333	.000	.000

Now team B is seen as clearly better than the rest since it beat better opponents. Also, team E gets a zero ranking even though it has won a game since its opposition is considered to be a pushover. Of course, we can keep on going! Iterating this process (so, for example, the new number for D is $\frac{.125+0}{3}$), we get the following table:

A	B	C	D	E	F
.194	.153	.083	.042	.000	.000

Note that team A is now ahead in the rankings! If you continue with this process, the ranking (but not the numbers) will stabilize to the following: first A, then B, C, D, with E and F tied for last. Put somewhat simply, team A gets the nod for first since its victory came against a strong opponent, while B's two victories were against somewhat weaker teams. While teams C and D have the same number of victories than B, half their victories came against pushover teams (namely, E or F).

We've iterated this ranking process three times, but there's no reason to stop there. In fact, we might as well do it infinitely many times!

Taking the Limit

First, let's transfer the picture into a matrix M. The ij entry of M will be the number of times team i beat team j divided by the number of times team i played. (Here we give the teams A, B, ... the expected numbering.) This matrix is often called a **dominance matrix**. Here's the matrix:

$$M = \begin{bmatrix} 0 & \frac{1}{2} & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{3} & \frac{1}{3} & 0 & 0 \\ \frac{1}{4} & 0 & 0 & 0 & 0 & \frac{1}{4} \\ 0 & 0 & \frac{1}{3} & 0 & \frac{1}{3} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{2} \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

If we let e be the column vector $[\ 1\ \ 1\ \ 1\ \ 1\ \ 1\ \ 1\ ]^T$, it is easy to check that Me is the column vector corresponding to winning percentage of each team and that M^2e and M^3e correspond to the latter two tables above.

We could now just find $lim_{n\to\infty}M^n e$. However, this will almost always equal the zero column vector. Instead, we'll define $|v|_1$ to be the sum of the absolute values of the entries in the column vector v. We now define our **ranking vector** to be:

$$r = lim_{n\to\infty}\frac{M^n e}{|M^n e|_1}$$

Here, the result will be a column vector whose entries sum to 1. In this case, $r = [\ 0.344\ \ 0.290\ \ 0.204\ \ 0.162\ \ 0\ \ 0\ ]$. In other words, the ranking is A, B, C, D with E and F tied for last.

Hey! How did you calculate that?

Well, I didn't take an infinite limit... There are two ways to compute this vector. The first way involves eigenvectors (see below) and the second way involves calculating the expression inside the limit until the numbers start to stabilize. With this example, the values stabilized fairly quickly, with the first three decimal places stabilizing when n is about 40. I use a Mathematica notebook to find the values.

How do you know the limit exists?

Great question! You can construct graphs for which this limit does not exist, and I know necessary conditions, but not sufficient ones. In particular, for the limit to exist, you need to have a **cycle** in the directed graph. That is, you need a sequence of distinct teams a_1, a_2, ... a_k such that a_1 beat a_2, a_2 beat a_3, ... and a_k beat a_1. (That's why I can't post rankings for the NFL until a couple of weeks into the season — I'm waiting for a cycle to form.) But simply because a graph has a cycle doesn't guarantee the limit exists. The question of what conditions you do need takes you deep into the field of **nonnegative matrix theory**.

Do the numbers mean anything besides giving a ranking?

Yes! Pretend you carry out the following process (called a **Markov process**):

1. You pick one of the teams at random.
2. Randomly pick one of the arrows leading in or out of the team. (So, for example, if you were are D above, you would pick one of the three arrows leading in or out of D.)
3. If the arrows leads in to the team (that is, it represents a game that team lost), STOP.
4. If not, move to the new team (the team beaten in the game) and go back to step 2.

Clearly, this process stops quickly in many situations. (If you pick F in step 1, you automatically stop. If you pick E, you either stop at E or move on to F and stop.) But some processes will run a long time. Suppose a process visited a large number of teams (say, 40) before it stopped. Suppose you did not know which team the process started at. The numbers in the column vector r represent the probability that the process started at that team. (That's why both E and F have zero rankings — a process that visits a large number of teams cannot start at E or F.)

Eigenvectors?

The column vector r has one more property (assuming it exists). It's an eigenvector of the matrix M. In fact, if you take all the eigenvalues of M and find their distance from the origin (remember, some eigenvalues may be complex), r is an eigenvector corresponding to the eigenvalue furthest from the origin.

How well does the system work?

Well, if the 1999-2000 year is any indication, not so well. The two teams which made it to the Super Bowl — the Tennessee Titans and the St. Louis Rams — were ranked 20th and 23rd out of 31, respectively. However, in my defense, the 1999-2000 year was an acknowledged crazy year in the NFL and the two divisions those two teams played in (the AFC Central and the NFC West) were clearly the two worst divisions in the league. Thus, many of the games the Ram and Titans played were against lousy opponents and (like D in the above example) pulled down their rankings.

Why do you publish two rankings?

The first one, titled 'W/L only', works exactly as described above. The second one, titled 'Scores', takes into account the score of the game. In particular, after every game, two arrows are created. The winning team has an arrow pointing to the losing team with a **weight** of (winning score + 1)/(winning score + losing score +2), while the losing team has an arrow pointing to the winning team with a weight of (losing score + 1)/(winning score + losing score +2). When the Markov process described above is followed, in step 2, arrows with higher weights are more likely to be chosen. Thus, you're likely to move to a team you blew out than a team you barely beat. In general, teams with a mediocre winning percentage who has only lost close games do better in this system.

Can I get more information about all this?

Sure! Talk to me in the halls or email me at mschlat@centenary.edu, and I'll answer any questions.