Topologically Ranking College Football Teams

Trevor Squires

Last Updated: October 19, 2024

1 Introduction

If you don't know by now, let me be the first to tell you that I, Trevor Squires, am a huge Clemson football fan. So much in fact, that I have joined numerous Facebook fan groups who see the world through equally orange-tinted glasses. And although I am a sucker for blindly advocating for Clemson in casual discussions, when it comes down to it, I am also a big fan of football in general and I try to be as objective as I possibly can for someone who spent 10 years at Clemson University. My Facebook groupmates, however, have no such interest in being impartial. They subscribe to whatever fallacies of logic will make our tigers look as good as possible on paper.

One common logical progression among all naive fanbases is the idea that if team A beats team B, then team A is just better than team B. No matter what. And although this seems obvious to the less-sports inclined faction of the world, any avid sports fan will tell you about plenty of times where the "better team" lost. In the college football world, many fans will remember the 2011 season where LSU beat Alabama in the "Game of the Century" on November 5th. After the loss, Alabama dropped from first place in the BCS rankings (those were the days) to third. Eventually, Alabama reclaimed a spot in the top two and was selected to play LSU (who was obviously number 1 at the time) in the national championship game. The same Alabama team that failed to win its conference, but also failed to win its own division was deemed at having a realistic shot at being the best team in the country. And after all that controversy, Alabama went on to beat LSU 21-0, claiming the college football title. So much for letting wins and losses dictate relative rankings.

Even then, college football rankings were a combination of many factors (though nowadays which factors to include depend on what statistics the media darlings happen to dominate). Yet, that doesn't stop us fans from constructing overwhelming evidence that our team is better, and no statistic makes that clearer than team A is better than team B if team B loses to team A. In this document, we take this idea to the next level by constructing a complete ranking of all college football teams using only this *completely rational and perfectly valid* rule: teams that win games must not be ranked lower than teams that they defeated[1]. Note that all of the following results are calculated using results up through week 6 of the 2024 college football season.

2 Mathematical Breakdown

If you were to choose to apply this ranking methodology to, say, your kid's soccer league, you probably wouldn't find it quite difficult. Undoubtedly, one team fielding the coaches kid is probably beating everyone, whereas one team who's very excited for half time oranges is propping everyone else's stats. There's probably some obvious ordering of all the teams in the middle. In other words, it would not be hard to write down and even verify a ranking among the teams that satisfies our golden rule. However, if you try to do that with 125+ college football teams, each with X different game results under their belt, you will probably give up before you find any success. With thousands of games a season, it's easy to prefer to be watching football than analyzing the results. Due to this difficulty in scale, it is important that we model our problem mathematically and leverage existing tools to speed up the process.

Let us begin by restating our problem: we would like to devise a method of computing a comparison between two teams in which team A is better than team B if team A has defeated team B. Consequently, it follows that if team B has defeated team C, then team B must be better than team C. Thus, if we were to further compare team A and team C, it must be that team A is better than team C, else the corresponding relationships with team B will have been violated.

[1] with the eventual goal of proving that Clemson is better than Michigan

In mathematics, the term to describe what we are looking for is called a partial ordering. A partial ordering on a set is an arrangement such for particular pairs of elements, one precedes the other. Here, we are using the term partial because not every elements can be related. This is opposed to a total ordering in which every pair of elements of the set are comparable. Indeed, it's entirely possible that some college football team not only does not play some other team, but the two don't even share common opponents. In this case, it's impossible to compare the two teams on wins and losses alone. Now, Vegas will certainly compare the two teams for you using some other approach than what we are attempting, but for now, that's out of our wheelhouse.

Partial orders have other properties that are natural for our problem (some more obvious than others). For any three teams a, b, and c,

- (reflexivity): a is no worse than a
- (antisymmetry): if a is no worse than b and b is no worse than a, then a and b are "equivalent" in strength
- (transitivity): if a is no worse than b and b is no worse than c, then a is no worse than c

Each of these properties is again very desirable in our hypothetical obsession in determining how to rank college football teams in the most sane way possible. Properly framing our problem as a search for a partial ordering one is the first step in being able to systematically achieve our goal.

3 Leveraging Graph Theory

Now that we have a proper goal in mind, how can we go about computing our partial ordering? To get started, we need to choose a model which represent the information that we have available to us: game results. Surprisingly (or maybe not depending on how far you paid attention in math), graphs do a great job at conveying relationships between objects. More formally, a graph G is a mathematical structure which contains vertices (or nodes) that are connected by edges (or arcs). We can use graphs to easily summarize a set of game results for a fixed set of teams as follows:

- Associate a node of the graph with each team in the league
- Associate a each completed game result with an edge of the graph connecting the home team to the away team. The edge should be pointing from the losing team to the winning team in order to denote the result of the game

The second bullet above means that we will be dealing with *directed* graphs, i.e. graphs whose edges have directions associated with them. In a directed graph, we usually denote an edge (i, j) if it connects nodes i and j by moving from node i to node j.

Let's look at an example, suppose a kid's soccer league has four teams where team 1 never wins, team 4 never loses, and team 2 has lost to team 3. We can construct the graph in Figure 1 to summarize this information visually. While a visual representation of our results is nice, modeling our game results as a graph has a plethora of other utility. Many of the concepts we have previously discussed have direct graphical counterparts.

Paths

A (directed) path $p_{a,b}$ in a (directed) graph is defined as a sequence of edges $\{(i, j)\}$ which join node a to node b. Paths between any two nodes need not exists, however. In Figure 1, there is no sequence of edges we could ever take to get us from node 2 to node 1. Now let's look a little bit closer at how we can interpret paths in our frameworks. Recall that our graphs have the special property that the existence of an edge (i, j) implies that team j is better than team i. Thus, a sequence of edges from i to j implies that there exist teams that are better than i, but worse than j. By transitivity, this means that team j is at least as good as every team along the path, including team i. Therefore, the existence of a path from i to j indicates that team j is no worse than team i, as it consistently outperforms or is comparable to all teams that connect them. This property allows us to use directed paths as a way of reasoning about relative team strength across multiple matchups. That is, even though there may not be an edge directly connecting team i to team j, if there exists a path between them, then we can conclude from transitivity that j must be no worse than i.

We can see this path property play itself out in Figure 1. Team 4 has no direct relationship with team 1 since the two have not played each other, but because a path exists from node 1 to node 4, $(1, 2), (2, 4)$, we know that team 4 is no worse than team 1. Realistically, we know this must be the case since team 2 beat team 1 and team 4 beat team 2. By transitivity, we have conclude the same result as our graphical approach.

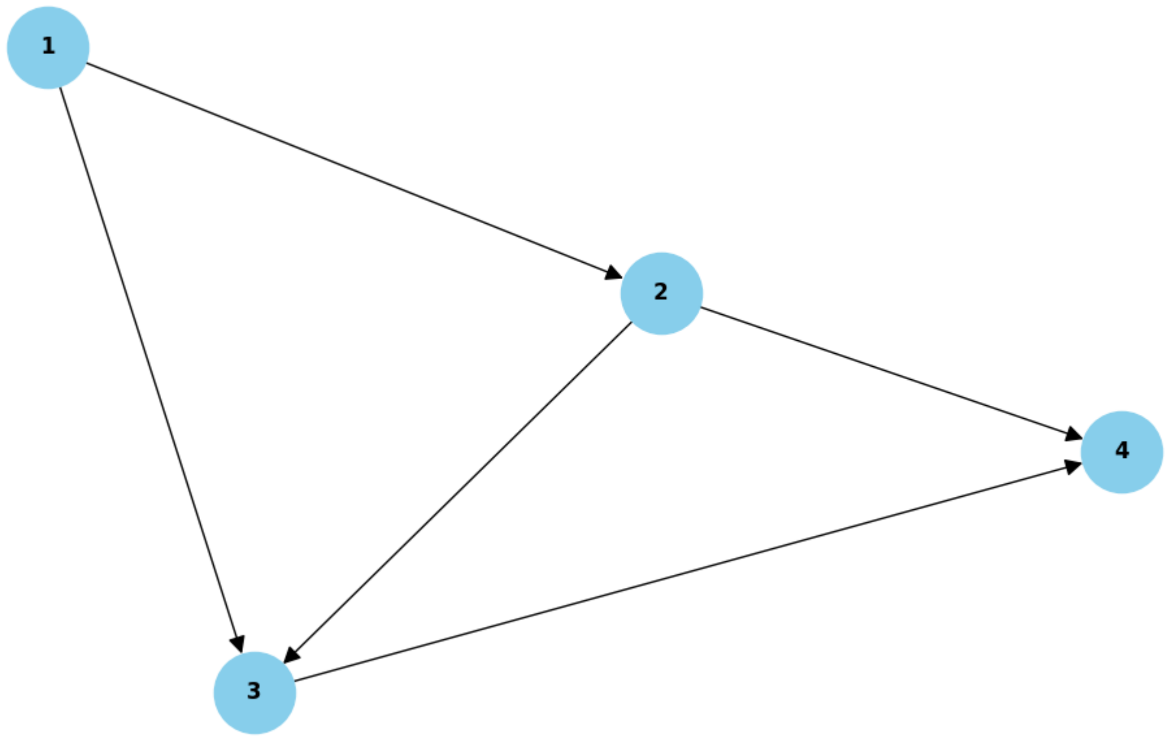

Figure 1: Simple Soccer Graph

Cycles

A cycle c_a in a directed graph is defined as a sequence of edges $\{(i, j)\}$ which start at node a and end at node a. Building on our discussion of paths, we can now consider what happens if there is a cycle that includes both nodes a and b. Note that a cycle involving these two nodes can be decomposed into a path from a to b combined with a path from b to a. Given that a path from a to b implies team b is no worse than team a, and a path from b to a implies team a is no worse than team b, this creates a scenario where each team is deemed at least as strong as the other. In light of the antisymmetry property of a partial order, which states that if $a \leq b$ and $b \leq a$, then $a = b$, we are forced to conclude that the two teams must be "equivalent" in strength. The existence of such a cycle therefore indicates that neither team can be strictly better than the other, establishing an equivalence between them in terms of performance.

Unfortunately, there are no such cycles in our simple soccer example. However, consider a second example in which we add an edge from node 3 to node 1. In some twisted rematch, node 1 actually wins a match against team 3 and the resulting graph is shown in Figure 2. Now, there are much more paths to contend with. We can now find a path between any two of the nodes $1, 2,$ or 3 implying that all three teams are equal in strength. Team 4 has still not been defeated, so they sit atop the rankings, but the remaining teams are in a three way tie due to the cycle $\{(1, 2), (2, 3), (3, 1)\}$ that contains each node.

Strongly Connected Components, Condensation Graphs, and Topological Orderings

The last graphical concept that we will need to leverage to achieve our somewhat distant goal at this point is the idea of a strongly connected component (SCC). A strongly connected component in a directed graph is a maximal subgraph where every node is reachable from every other node in that subgraph. In our framework, this means that within an SCC, any two teams are "equivalent" in strength, since there are paths connecting them in both directions. By grouping all such equivalent teams into a single SCC, we can simplify the graph significantly. In the cases of cycles, it is no longer important for us to consider all relationships between the cycle teams and other teams. Instead, we can, we can treat each SCC as a single unit, which allows us to focus only on how these units (or groups of teams) relate to each other. This leads us to the condensation (or condensed) graph, a simplified, acyclic graph where each node represents a strongly connected component, and the edges between nodes show

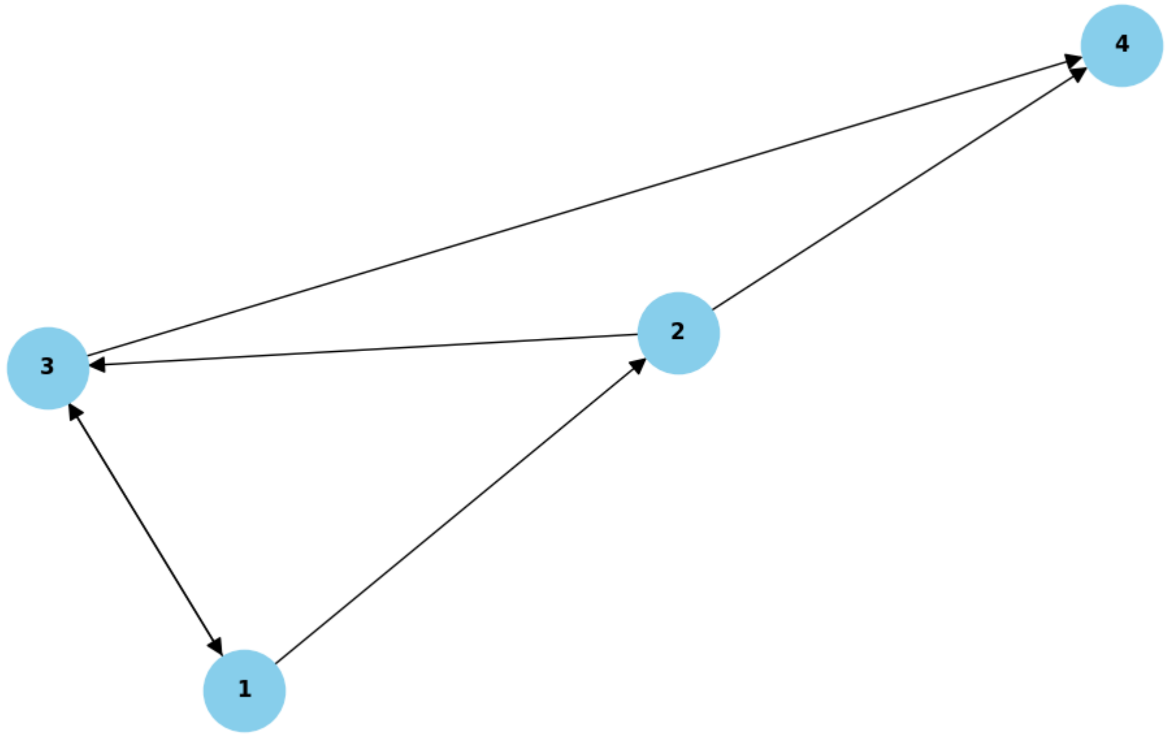

Figure 2: Simple Cycle Soccer Graph

the hierarchical relationships between these groups. Since there are no cycles in a condensation graph (proof left up to the reader), this condensation graph will serve as a much clearer and more efficient way to discuss which teams are definitively better than others, avoiding the need to untangle cycles or equivalence relations between teams.

As for our example, the strongly connected components of our graph are $(1, 2, 3)$ and (4) indicating that there is a cycle containing the first three teams, and a trivial cycle containing the fourth team. If we denote the nontrivial SCC as an arbitrary team 0, our graph simplifies greatly into Figure 3

The simplification of the condensation graph plays a major role in why modeling our game results as a graph makes the most sense. A quick look at Figure 3 tells us everything we need to know about the strength of the soccer teams: all teams in the SCC represented by node 0 are equivalent, all teams in the SCC represented by node 4 are equivalent, and all teams in node 0 are all worse than teams in node 4.

To take this one step further, we can apply a topological ordering to the condensation graph. In this case, a topological ordering is simply an extension of the aforementioned partial ordering, except we fill in the blanks between two uncomparable teams in any way that doesn't violate the comparable teams. For example, if team A has no relation (direct or indirect) to team B, then ranking team A or team B first would be acceptable. This topological ordering of the condensation graph allows us to take our partial ordering that is generated from the game results and turn it into an actual ranking which is a nice addition to our objectives.

4 Applications to College Football

With our knowledge of graph theory refreshed, let us summarize our findings of what is possible with a directed graph of the game results

1. To compare team A to team B, search for the existence of a path from node A to node B and vice versa.
   - if a path exists from node A to node B, then team A is no worse than team B
   - if a path exists from node B to node A, then team B is no worse than team A
   - if both path exists, then team A and team B are equivalent in strength

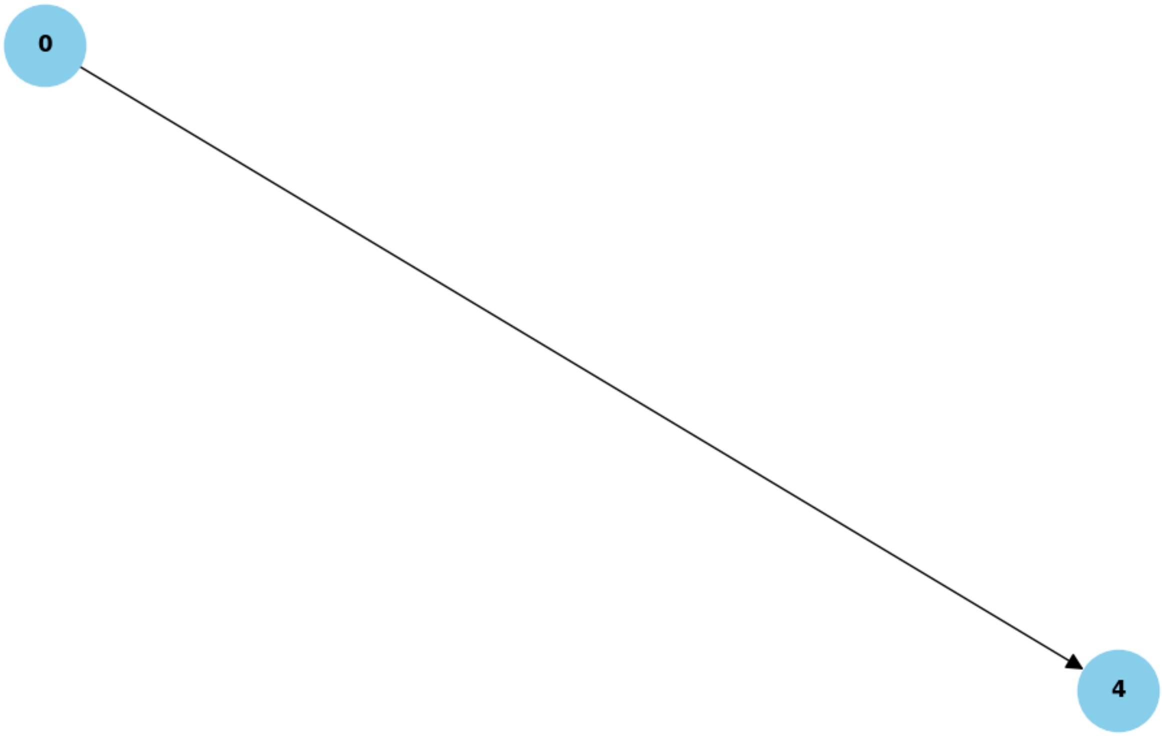

Figure 3: Condensation Soccer Graph

2. To generate our desired rankings, construct the condensation graph and compute a topological ordering of its strongly connected components. This forms an unambiguous ranking of all teams involved, though many such valid orderings may exist.

Although it seems as though we've simply replaced one goal with another, by utilizing a graph representation, we can take advantage of longstanding, efficient algorithms that accomplish our graph-related objectives extremely quickly. Thus, we can get our comparisons or rankings equally as quickly. I've done the "challenging" portion of this implementation in only a few lines of code. For example, here's how easy it is to compare two teams by leveraging our graphical representation and NetworkX's graph algorithm python package after some initial set up.

```
def compare_teams(self, team1, team2):
    two_team_cycles = [cycle for cycle in self.cycles if team1 in cycle and team2 in cycle]
    path1 = get_shortest_path(self.G, team1, team2)
    path2 = get_shortest_path(self.G, team2, team1)

    if len(two_team_cycles) > 0:
        print(f"The two teams are equivalent because of the cycle {two_team_cycles[0]}")
        return two_team_cycles
    elif path1:
        print(f"{team2} is better than {team1} because of the path {path1}")
        return path1
    elif path2:
        print(f"{team1} is better than {team2} because of the path {path2}")
        return path2
    else:
        print('The two teams are incomparable')
        return None
```

The full code can be found on my github under the `evaluate_teams.py` file. But that's enough of experimentation, let's actually dive into what we really wanted to see - why Clemson was so much better than every other team in college football.

Individual Rankings and The Circle of Suck

Before looking at the top of the rankings, let's take some time to do some 1 on 1 comparisons to sanity check our work. Specifically, we'll use the compare teams function shown above to do a few quick comparisons

1. Oregon vs LSU: Oregon is better because of the path LSU-USC-Minnesota-Iowa-Ohio State-Oregon
2. Boise State vs Iowa: the two teams are uncomparable
3. Clemson vs Florida State: Clemson is better because of the path Florida State-Clemson
4. Alabama vs NC State: the two teams are equivalent due to the cycle NIU-NC State-Clemson-Georgia-Alabama-Vandy-Missouri-Texas A&M-ND
5. Clemson vs Georgia: the two teams are equivalent due to the cycle NIU-NC State-Clemson-Georgia-Alabama-Vandy-Missouri-Texas A&M-ND
6. Clemson vs LSU: the two teams are equivalent due to the cycle South Carolina-LSU-USC-Minnesota-North Carolina-Georgia Tech-Syracuse-Stanford-Clemson-Georgia-Alabama-Vanderbilt-Georgia State-Old Dominion
7. Clemson vs Michigan: the two teams are equivalent due to the cycle South Alabama-Arkansas State-Michigan-Washington-Rutgers-Wisconsin-USC-Minnesota-North Carolina-Georgia Tech-Syracuse-Stanford-Clemson-Georgia-Alabama-Vanderbilt-Georgia State-Old Dominion-East Carolina-App State

An interested reader might notice that a lot of the same teams are appearing in these paths/cycles. Just from these quick comparisons alone, one might be included to believe that there are a lot of equivalent teams. Indeed, there currently exists a 58 team strongly connected component that encompasses many teams from the ACC and SEC such as Alabama, Clemson, Georgia, LSU, Georgia Tech, Tennessee, and Notre Dame and can be found in the eventual rankings below. The college football community has collectively decided to refer to these super cycles as the "circle of suck". Indeed, the existence of such circles imply that powerhouses such as Rutgers are somehow equivalent to much smaller schools such as Georgia. No team in the aforementioned circle of suck can possibly be better than everyone else since they all exist in a cycle together. Thus, if one team is viewed as weak, then all teams must be equally poor. Hence, the circle of suck.

But how did such a cycle come to exist? Surely 1 loss teams like Georgia and Clemson haven't had any embarrassing performances that would subject them to the circle. The key to the expansion of the circle occurred on October 5th when Alabama fell to the Vanderbilt Commodores for the first time in 40 years. Because Vandy has lost to some certified "happy to be here" schools, those losses extend to Alabama. And because Alabama has defeated so many quality teams, Vandy's losses extend even further. The net effect is that a massive amount of teams are forced to join the circle simply by losing to teams such as SEC juggernauts Alabama and Georgia.

The qualifications to join the circle of suck are as follows; one must both defeat a team in the circle and lose to a (not necessarily the same) team in the circle. As time goes on and the circle grows, this application process becomes easier and easier. In other words, as we add more and more edges into our graph of game results, the circle of suck eventually comes for us all. Nonetheless, we can still use these tools to achieve our original goal of constructing a topological ordering of the top 100 college football teams where each team never ranked below a team that it defeated. Although it may not be the most insightful list, it is perfectly fair and balanced among the eyes of those who don't have anything better to do with their time. Any and all complaints can be directed towards Kalen Deboer.

1: Oregon
2: Boise State
3: Washington State
4: BYU
5: Texas Tech
6: Kansas State
7: Penn State
8: Iowa State
9: Pittsburgh
10: Arizona State
11: Arizona
12: Texas
13: Illinois
14: West Virginia
15: Utah
16: UL Monroe
17: Ohio State
18: Nebraska
19: Oklahoma State
20: Liberty
21: North Texas
22: Miami
23: Indiana
24: SMU
25: James Madison
26: Iowa
27: Cincinnati
28: Colorado
T29: Arkansas State, Stanford, Duke, Florida, Vanderbilt, NC State, TCU, Virginia, Georgia Tech, Coastal Carolina, Rutgers, Wisconsin, Washington, Wake Forest, Louisville, Virginia Tech, Alabama, Minnesota, Oklahoma, Michigan State, Old Dominion, Notre Dame, Marshall, Texas A&M, Ole Miss, Georgia, Syracuse, Missouri, Charlotte, Georgia Southern, Houston, Arkansas, Northwestern, Louisiana, Ohio, North Carolina, Tennessee, Maryland, Boston College, Buffalo, Texas State, Western Kentucky, South Carolina, UCF, Kentucky, Clemson, UConn, Tulane, South Alabama, East Carolina, Northern Illinois, App State, Georgia State, Michigan, UNLV, LSU, USC, Sam Houston
87: Toledo
88: Miami (OH)
89: Navy
90: Eastern Michigan
91: Memphis
92: Baylor
93: Jacksonville State
94: South Florida
95: Army
96: Florida State
97: Tarleton State
98: Southern Miss
99: Rice
100: California