

Fantasy Journals

Summary: Create the rules and web infrastructure for a game of "fantasy journals" analogous to the fantasy baseball and fantasy football leagues that are so popular among sports fans. Scientists could draft papers for their own fantasy journal, and then compete to see whose journal was most successful. Such a game would be great fun to play e.g. in a conference setting or in a research group, and it would also potentially be a source of valuable bottom-up bibliometric tagging information.

Rationale: In several sessions today we discussed alternative ways of validating and evaluating scholarly papers. Numerous participants stressed that it would be nice to find ways get people engaged in bottom-up peer review / evaluation of papers. How do we do it? Take a hint from Google's Image Labeller (<http://images.google.com/imagelabeler/>). Make it into a game!

For now, let's not worry about the potential bibliometric utility, and let's simply think about how to create the game so that scientists would have fun playing it. People enjoy (and become obsessed with) playing fantasy football, fantasy baseball, etc. In these games, participants select, draft, or trade for a "fantasy roster" of players selected from the real major league rosters. Each game, their "fantasy team" performs according to how well those players performed in the real major leagues that day, and wins or losses are assigned in this fashion. The challenge is to know enough about the players to be able to predict their performance better than one's competitors, and thereby put together winning teams.

Every serious baseball fan seems to believe that he or she can outdo George Steinbrenner or whatever general manager is controlling his team's roster. To hear my colleagues talk about the mistakes that Nature or PLoS made by rejecting their papers, my colleagues seem to hold similar beliefs – they seem to believe that they that they can outdo Philip Campbell or Catriona MacCallum (no offense to Philip or Catriona) or whoever is controlling the "roster" of papers appearing in Nature. So let's put it to the test! Let's set up a game where players can select papers from across a field or even from across of all science for their own "fantasy journal", and then see how that journal performs according to the appropriate bibliometric measures. (We'd propose some function of the number of citations on Google scholar after a period of time)

This is a game that one would win by being good at picking the soon-to-be hot papers. Our lab would have a blast playing – and if I challenged my graduate students to beat my picks, I can guarantee that they would read an increasing fraction of the literature in their efforts to put my in my place. After all, this is a game that lets scientists show off the one thing they may like to show off even more than their own research abilities — their access to fresh, hot, timely, valuable information about the next great development in the field.

It's gossipy, it's self-referential, it's a sporting blend of skill, effort, and fortune, it's competitive – all great features to have in a game. And it has the side effects of getting people to read the literature, it generates an interesting ensemble of individual "overlay journals" that reflect the interests of individual researchers, and it potentially generate large quantities of bibliometric evaluation data that could subsequently used in scientific search.

Implementation: We'd propose to develop a simple platform for playing this game. We can test it (1) among interested groups of friends, and (2) at a moderately-sized conference in computer science. In the latter case, the game would be to assemble a "dream team" from the papers at that particular conference – with a prize to be awarded at the next year's meeting to the person who chose the best set of papers (perhaps measured by the number of citations in Google Scholar).

There are a few implementation challenges, not the least of which is figuring out exactly how to design the rules. We can divide this into (A) acquiring papers and (B) scoring.

Acquiring papers has to be done on some sort of bidding or draft-and-trading system. Everyone knows that the next paper by Nobel Laureate working in immunology is going to garner more citations in a year than the next paper by an unknown graduate student working in economics. The game gets interesting when there are constraints on who you can pick (e.g. a draft followed by a trading period), or when there is something like a futures market for the citations that papers will receive. At the same time, the rules here need be (a) very simple, (b) executable without getting everyone together to play simultaneously, (c) not requiring players to return repeatedly to make trades or adjustments in response to other players offers or moves, (d) scalable – no one wants to receive 198,035 potential papers in the '07 draft! This seems to rule out most draft and trading systems, as well as sophisticated futures markets. One system that might work would be tightly

circumscribe the set of possible papers, and then offer fair odds against the picks that the players make. There must be other, even more clever designs out there!

Scoring has to be done in a way that returns results within a tolerable time frame (a few months or less, ideally) and that is hard to game. Here Google Scholar citation counts should work about as well as anything I can think of. If we allow people to pick papers that already have a non-zero citation count, we might also want to account for the fact that citation is a preferential attachment process in which highly cited papers attract more citations by virtue of their prominence in the references of other papers. There are doubtless other issues in implementation, but this should give a reasonable sense of the scope of the problem.

(Carl Bergstrom, James Hendler, Dan Chudnov)