

sprezzatura, PageRank, and the recursive cool

Simon DeDeo*

October 1, 2017

“Before it’s in fashion, it’s in *Vogue*.”

—tagline for *Vogue Magazine*

Contents

| | | |
|----------|--------------------------------|----------|
| 1 | An Introduction to Cool | 2 |
| 2 | Measuring Cool | 2 |
| 3 | Quantifying Cool | 4 |
| 4 | Cool Matters | 6 |
| 5 | Beyond Cool | 7 |

There are many uncontroversial facts about the world: Pittsburgh is in Pennsylvania; cars tend to stop at red lights; coffee generally costs less than \$10 a cup. There are uncontroversial facts about people, too—the president of the United States is Donald Trump; the president of CMU is Farnam Jahanian.

People may be wrong about facts like these: the capital of Pennsylvania is not Philadelphia, even if more people think that than know it’s Harrisburg; Subra Suresh was CMU’s president until Jahanian took over this Summer. Yet there’s no direct relationship between one’s being right or wrong about the these facts, and the facts themselves. You may not *like* President Jahanian, but he’s still the president (until he’s not, which requires a particular set of legal hoops to jump through); if everyone starts pretending Jahanian’s not president, he may not be able to work very effectively, but he’s still president (they’re just pretending); Jahanian himself remains president even while he’s having a nightmare in which he’s, briefly, *not* president.

Uncontroversial facts include not just material matters (the weight of a Subaru), but also things given special status by being written down (what’s the legal capital of X, what’s my immigration status)¹ and behavioral practices (generally-observed patterns of life, like what we tend to do when

*Carnegie Mellon University & the Santa Fe Institute. Current version available at <http://santafe.edu/~simon/pr.pdf>. Please send corrections, comments and feedback to simon@santafe.edu; <http://santafe.edu/~simon>. This article first written for the 2017 “Bubbles: Big Data for Human Minds” Class, Social and Decision Sciences, Dietrich College, CMU; <http://bit.ly/bubbles2017>. I thank T. M. Bejan and Angus Grieve-Smith for comments.

¹Or, indeed, a social security number or bank card PIN; facts don’t need to be public, or widely known, to be uncontroversial.

we encounter traffic lights and stop signs).

Other facts, however, are quite different: they emerge out some kind of collective process where people talk and exchange beliefs and observe each other, and all of that mingles in weird ways with facts about the world and what's in the past. We don't really understand how this mingling works, but what ends up happening is that we produce beliefs about things that never really existed before. The mingling itself somehow brings them into being.

1 An Introduction to Cool

Take, for example, the idea of being “cool”—and the strange fact that you can certainly ask of a person, thing, or cultural figure, say, “is she/he/it cool?” If you have a modicum of social knowledge, you can even ask things like “is it cool for teenage girls” (yes: Taylor Swift, *Seventeen Magazine*; no: argyle), or “is it cool for college professors” (yes: argyle, Taylor Swift; no: Jonathan Franzen). You probably can't quite define what it means, of course, but you know some outliers, at least, outliers relative to a group you know well (your friends, say, or your “friends back home”, or your family, or the people you expect to be working with next year).²

If we asked you to define cool, you might say it's “casual success”; it's also something to do with style, perhaps, how you dress, a kind of disinterested grace. But it's also larger, about being unconcerned with the normal kinds of success, or being somehow internally driven and doing all sorts of crazy and intense things, maybe in ways that put you at odds with conventional morality—but not in really bad *Trainspotting* ways (except when it is). It's maybe about being unique in an apparently effortless fashion, except not so unique that you're alone (except, again, when it is), or what people think of as weird (until what's considered weird changes, and suddenly, say, writing code is cool). Perhaps, on reflection, it does involve huge amounts of work, but you kind of wear it lightly and are, well, cool about it, the way Bob Dylan is in the 1967 documentary *Dont Look Back*.³ Perhaps it has to do with being a bit above it all, like an aristocrat who doesn't have to care; or perhaps, the exact opposite, someone who's so beaten down they can be beatnik aristocrats of the spirit. But it's definitely not actual aristocrats, those with titles (except when it is).

Lurking behind all of this may be some residual cultural universal sitting around, biologically coded-in like the blind spot caused by our optic nerve. Or perhaps it's more like a belief system, like Christianity or Buddhism, that's spread and mutated and diversified over the centuries. In Renaissance Italy, there was *sprezzatura*, which maybe meant something similar, and perhaps the Mona Lisa, with her cool smile, was someone with *sprezzatura* (which might explain why we're so obsessed by the image). It's doubtful, however, that you're going to get a definition of cool that applies uniformly, that you can program into a computer and have it search for and identify in the world. In the words of the (at one point cool) band R.E.M., “you can't lay a patch by computer design”.

2 Measuring Cool

Except, perhaps you can. Because there's nothing mysterious about cool. It's just that it's something that grows out of our use of it: we look at each other, and at the world, and we examine our

²Introspecting a little bit, there's the strange feeling that you might know what your friends think is cool better than what you, yourself, think is cool. We are happier thinking that people other than us have rather simple and knowable views about these things, while we think we, ourselves, personally, are a bit more mysterious and subtle.

³Except when he isn't—<https://www.youtube.com/watch?v=0P9ArKYGvD4>

feelings, and we start to act, and we definitely *know* what is cool, but can't account for it in any kind of rational or fair way that we can explain without feeling a bit circular. "Why is that cool? It just is!"⁴

Now some things like this are easy. There are some facts like this of which one can say, "the meaning of the term is just the average of what everyone thinks it means". A lot of machine learning algorithms work this way, trying to predict what people think about things based, essentially, on looking at what a lot of people think. It's pretty plausible to say, for example, that what most people think a "regular" coffee is, is what a regular coffee just is around those parts, even if they can't quite define it precisely.⁵ If you complain about what you got when you ordered a regular coffee, the guy at the grocery store can just say "that's what most people around here think is a regular coffee." This class is not a machine learning course, so we won't spend much time on these boring examples of finding, say, the "average" or the "consensus" view about something.

Cool, however, is definitely not boring. If you want a computer to figure out what's cool, you'll need to do more than feed it surveys. That's because different people have different levels of access to facts about who is cool. Not, importantly, that some people happen to know the real facts of cool better (the way, say, an electrician knows more about fuseboxes than the non-electrician). But that some people have better access to the facts about who is cool, *simply because they themselves are cool*.

In most social circles, put plainly, some people are considered better positioned to know what's what. Depending on the community, it might be Anna Wintour, the editor of *Vogue Magazine*; or Jay Z; or the beardy guy at the bicycle bar. We sometimes call them "arbiters" of taste, as if they're supreme court justices, but even among the arbiters there's a hierarchy (Lil Wayne vs Jay Z?) Put schematically—and bearing in mind that this is only one aspect of what it means to be cool; it would of course be tragically uncool to make this the technical definition of cool—it's not enough that

1. To be cool, people need to think you are cool.

but, really, that

2. To be cool, people who are cool need to think you are cool.

How do we know that the "right people" think you are cool? Well, let's take seriously for a moment the idea that the amount that cool has changed over the centuries is enough to imagine it might be largely unanchored from any consistent reality (wearing leather jackets, smiling mysteriously in paintings). Why does Jay Z have so much influence over what's thought to be cool? Perhaps because—and only because—he himself is thought to be cool. Let's bootstrap it and see what happens:

⁴As Ludwig Wittgenstein, truly the 20th Century's coolest philosopher, once wrote: "If I have exhausted the justifications, I have reached bedrock and my spade is turned. Then I am inclined to say: 'This is just what I do'."

⁵A regular coffee in New Jersey is about yay big, with milk and lots of sugar, hot, but not too hot; that's very different from a regular coffee the guy with the mustache gives you at Commonplace in Squirrel Hill. Angus Grieve-Smith writes in that this example of the "regular" coffee is similar to John Maynard Keynes' example of the (unusual) beauty contest in which the goal is to pick the contestant you think the *other* judges will consider attractive. I think this is largely correct, except that we find the beauty contests rather hard to solve, and the "regular coffee" is easy (or we wouldn't do it, mornings are rough enough). For Keynes, the beauty contest gets interesting when it is a one-shot game; or, at least, when the context is sufficiently different from last time that you can't learn from before. We talk in terms of a regular coffee, conversely, because the game is repeated exactly (every morning!)—and that makes it, happily, trivial.

3. To be cool, people who are thought to be cool by people who are cool need to think you are cool.

Let's highlight a chunk of that:

3. To be cool, people who are thought to be cool by **people who are cool** need to think you are cool.

Note that this sentence both defines cool, and invokes the definition (in bold). All we do now is just slot in the original sentence, in a recursive fashion, back in to the bolded bit, and of course we'll carry along a bolded bit too:

4. To be cool, people who are thought to be cool by people who are thought to be cool by **people who are cool** need to think you are cool.

This is a form of “recursion”; we say that this definition of cool is “recursive”. The whole process seems like it's going to explode out of control, but it's not the world's worst form of recursion and it turns out we can tame it. We'll see worse ones later in this course—if you're a computer science major, the one here is a kind of tail recursion.

3 Quantifying Cool

Let's simplify further—a classic move in the mathematical sciences. Rather than use sentences talking about “who is cool”, let's imagine we want to rate all the people in a group by how cool they are. If there are N people, we'll have for each person i a number p_i , which quantifies how cool they are.

Let's take as given (for now) that what goes into being cool is just the opinions of others, correctly combined in the mysterious fashion of the numbered sentences above. Simplifying again, let's write “the opinion j has about i ” as r_{ij} (note the order of the index: r_{23} is the opinion person three has about person two). Then let's say that the opinions, the r_{ij} s, are just numbers, ratings.

The puzzle then becomes how to go from N^2 numbers (all the ratings people have, including their self-ratings) down to N (the “objective” list of how cool everyone is). Notice that we reduce the amount of stuff—if there are a hundred people, we go from 10,000 facts (the r_{ij} s) down to just 100 (the p_i s). This is called coarse-graining, or lossy compression, and falls under a set of ideas called renormalization (see, *e.g.*, the online tutorial we did for the Santa Fe Institute, <http://renorm.complexityexplorer.org>).

Let's start with the obvious: the case where being cool is like being a regular coffee, and all we have to do is figure out the “average” view. Then we might say that

$$p_i = \frac{1}{N} \sum_{j=1}^N r_{ij}, \quad (1)$$

or, in words, the coolness of person i is just the average of the ratings she gets from everyone else (including, for now, herself).

This is, of course, unsatisfying. But we might feel good enough to say, well, p_i is just an approximation to who's cool. What if we used the approximation to weight the ratings? A simple thing to do is then write

$$p'_i = \frac{1}{N} \sum_{j=1}^N r_{ij} p_j, \quad (2)$$

where p'_i is now an updated version of the coolness metric. People who other people think are cool—*i.e.*, people with high p_i —have their ratings matter more. Think, for example, of the case where a person gets zero ratings from everybody. Then it doesn't matter what they think—their p_i is zero, and all their ratings will be zeroed out.

The pattern to follow is now obvious. If you thought p_i was a good approximation to underlying cool, then you should think p'_i is even better.⁶ So take another step, using the p'_i s to weight the ratings,

$$p''_i = \frac{1}{N} \sum_{j=1}^N r_{ij} p'_j. \quad (3)$$

You can see where this is going. What we really want is the whole thing to stabilize: we want to find the set of numbers, call them \hat{p}_i , such that the thing on the left is the same as the thing on the right,

$$\hat{p}_i = \frac{1}{N} \sum_{j=1}^N r_{ij} \hat{p}_j, \quad (4)$$

so that the whole equation looks like the sentences in the previous section: a definition where the thing on the right that's doing the defining actually includes the thing on the left. That was really weird when we had sentences, but now that we're in the land of numbers, it's actually pretty easy to find. You never got freaked out when you saw an equation like $x = 3x + 6$; no need to panic now.

It turns out it's not quite so simple. That's because the ratings are unbounded: somebody can put in an arbitrarily large rating, and take over the system—that doesn't seem right. Indeed, even for ratings that are not enormously high, the p_i s can start growing out of control as you go from p to p' to p'' , and \hat{p}_i may not exist.

There are some pretty simple ways to save this. A natural one is to “curve” the ratings. Instead of working with r_{ij} , you work with \hat{r}_{ij} , where

$$\hat{r}_{ij} = \frac{r_{ij}}{\sum_{i=1}^N r_{ij}}, \quad (5)$$

or, in words, you divide by the sum of all the ratings person j gave out. Once you do that, it's natural to drop the averaging term $1/N$, and you're left with something that looks like this:

$$\hat{p}_i = \sum_{j=1}^N \hat{r}_{ij} \hat{p}_j. \quad (6)$$

If you actually want to solve this equation, you have a few loose ends to tie up, including people who give zero ratings to everyone; people who give zero to everyone but themselves; people who form a little self-regarding clique. Usually these are special cases you don't have to worry about (or, conversely, they're really interesting—remember, the point is to learn about cool, not to do math). If you really want to be super careful, and force this to work, you can do simple things, like require everyone give at least a one to everyone, no matter how uncool. This is the point, if you're mathematically inclined, that you notice that what you're really doing is finding the eigenvector of the matrix \hat{r}_{ij} ; which one? The one with eigenvalue equal to one.

⁶What if you don't? You might think that p'_i was pretty good, but that feeding it back in might cause the new value to overshoot. That's fine—you can skip to the “solution”, if you like. In general—once you make some tweaks to r_{ij} , which we'll address—you can get to the full solution by writing code that iterates this process.

There are many ways to understand the self-consistent, stable, or “fixed point” solution of Eq. 6. One example is thinking about the whole thing through what we called in class “the trophy game”. There’s a trophy that’s being handed around the group; once a person is handed the trophy, they can choose who gets it next. How do they decide? One way to do this is let \hat{r}_{ij} define the probability that, given that j has the trophy, they hand it to i .

Now that you’ve set the rules for the game, go for a walk and let the group circulate the trophy a bit. The number \hat{p}_i is just the probability that i is holding the trophy when you come back in. The trophy game is a nice way to see how the recursive nature of cool kicks in. If people think I’m cool, I tend to get handed the trophy. But that also means I have a lot of say who gets the trophy next. That’s another way to find \hat{p} by the way: if the first way is to keep updating from p to p' to... and waiting for it to stabilize, another way is to simulate the trophy game. (The way most people do it is find some Python or R code to do it for them.)

It’s worth remembering that all of this is trying to get at the nature of cool, under a lot of approximations about what cool is actually doing in the real world. We’re focusing solely on the “socially constructed” aspects of cool, and not the underlying things that might track it. Cool, sprezzatura, There are lots of choices for r_{ij} , how to transform it, what kind of self-consistent solution you want, and so on. In the end, we’re not *defining* cool so much as trying to approximate it. In the language of the social sciences, what we’ve done is *operationalize* cool, a mouthful of a term that means, roughly, that we’ve taken a big, historical, culturally complex thing that we could discuss endlessly, and given you an explicit recipe for how to weigh some trace of it in a dataset.

4 Cool Matters

While we know where the r_{ij} s come from—they’re the ratings that people have made—it’s much less clear where the \hat{p}_i ’s “live”. No one person knows them. They’re sort of free-floating social facts, complicated and clever summaries of what lots of people—but no one in particular—believes.

Do they *do* anything? The most obvious answer is to point to places where their use led to great commercial success. If we imagine not people rating people, but websites “rating” other websites by linking to them, Eq. 6 defines what Google called (and patented as) PageRank: when r_{ij} counts the number of hyperlinks from site j to site i . Computing \hat{p}_i now gives us distributed rating of the page’s importance. In the early days of the Google, knowing the site’s “quality” in this way enabled them to return results that were radically, and obviously, better than competitors.

There are a lot of simple examples. We use these kinds of distributed algorithms to measure who’s “cool” in animal societies, as well as human ones, and as you can see in Ref. [1], being cool in that sense means you have more sex and (because you’re an animal) more children, and that matters to evolution, at least.

A less high-stakes case comes from a Hackathon run at Indiana University back in 2014. Fig. 1 shows the results of comparing the \hat{p}_i s, inferred from the ratings students gave each other in the Hackathon, to their self ratings, making the relationship between the two kinds of facts clear.

A final example comes from work with the biologist Liz Hobson at the Santa Fe Institute [2]. In that paper, we looked at birds, who have at least some ability to create social facts.⁷ There, we were able to measure not just rank, but that the birds themselves were aware of rank, and using it to make decisions. The birds were not just creating facts about who was cool—they were also using it to figure out who (in this case) to pick one. As the bird society evolved in time, we could see the

⁷See Ref. [3] and references therein for more on the different levels to which humans, primates, and other animals can do this sort of thing.




Figure 1: Self- and other-ratings in a 2014 Hackathon run at Indiana University. Each point is a group. On the y -axis are the ratings groups gave themselves. On the x -axis are the recursive fixed-point \hat{p}_i s of Eq. 6; to get these we’ve zeroed out the self ratings, *i.e.*, set r_{ii} equal to zero. Among other things, you can see how students, in general, hold higher opinions of themselves that does the group as a whole.

awareness arise “out of nothing”. Early on, in other words, we (as scientists) could measure the \hat{p}_i s; but after a few weeks, the birds appeared to know them—or at least, to approximate them—too.

5 Beyond Cool

Cool is not the only kind of thing like this. We’ve pondered things like this for many centuries. They name different things: prestige, “power”, “weight”; being in, being *au courant*. Who calls who when there’s a job opening? Who do you go to when you need someone to settle a matter fairly? Who knows what’s right, what’s appropriate, what’s correct? Societies differ in the ways they answer these questions: there are some super-egalitarian ones who answer these by reference to the average; others that go by a kind of recursive system; others, perhaps that try to enforce a dictatorial system (harder than you think, since you can always ask “did we get the right dictator?”) The algorithm here has enchanting mathematical properties, but it would be very strange if it got everything right. One expects, finally, that efficient systems try to build in as much reality as they can—but find, finally, that some kind of recursive structure is necessary to make it stick.

“Weighty men,” writes Machiavelli back in 1520, “who have prestige, will always occupy the highest places”. But while Machiavelli’s readers would have known what “weighty” meant then, it means something quite different now—it would have been a surprise to that man that a great seat of higher learning (CMU) would be run by a man born outside the country (as in the case of Prof. Suresh); it certainly would have surprised him to see the ways in which women and men shared its administration, research and teaching.⁸

⁸Machiavelli may not have had modern ideas about the equality of the sexes, but he certainly saw women as capable of his kind of courage; his *Discourses* includes the case of Caterina Sforza. Enemies of her family kill her husband; holding her children hostage, they send her back to her castle to negotiate surrender by the guards that remain. Once inside, however, Machiavelli recounts, “to prove to them that she cared not for her children, she pointed to her sexual parts, calling out to them that she had wherewith to have more children”, a particularly intense move that may remind you of a memorable scene from the *Usual Suspects*. Sforza got her castle back.

As with cool, “weighty” has a similar kind of paper-like quality, on which many things can be written. It’s hard to find a greater contrast to Machiavelli than the Quakers; committed to non-violence and seeing all as equal in the eyes of God, early Quakers suffered for their refusal, for example, to doff their hats to aristocrats back in Britain. The commitment to equality was baked in very early and has only grown over time; yet even today’s meetings have the notion of a “weighty Friend”.

It’s clear, in other words, that while there are some underlying features that show a great deal of persistence in definitions of weightiness (for example, that people younger than twelve are generally less weighty than those older, though see some crazy periods in European history), there’s a great deal of fluid social consensus involved as well. That doesn’t mean that we can just decide: in fact, it makes it much harder. To move N numbers around, we have to, in the end, move N^2 beliefs—even for a small group, like our class, moving 39 ratings around means, if we’re unlucky, needing to push 1,521 beliefs. In a way, social facts like these are *harder* to move around than beliefs about material facts; or, put another way, if you find a material fact that you can’t convince others of, you might have a social fact standing in the way.

There’s something that occurs to people when they start thinking this way: perhaps some apparently uncontroversial facts are actually constructed in this fashion. If there’s no group-independent notion of “cool”, perhaps this is true for other things we didn’t at first suspect. Some people take this all the way, think that *all* facts, secretly are like this, that “everything is socially constructed”. As the physicist Alan Sokal once quipped, people who believe that literally all facts are socially constructed are welcome to test the social construction of facts about gravity by walking out my third-story window [4].

Let’s return to the domain of norms: what “ought” to happen when I jump out the window, rather than what actually does. Here it feels there’s more social construction, and we can try to extend things further in clever mathematical directions. It’s fun to consider the “Eigenmorality” story from Scott Aaronson (<https://www.scottaaronson.com/blog/?p=1820>); there Scott and his collaborators used clever weighting functions that go beyond cool-vs-uncool to include negative values, letting you describe ideas like “evil–neutral–good”.

But could that be all there is to moral reasoning? Scott, to his credit, thinks it’s a lot of playful fun. He mentions the philosopher Rebecca Goldstein, who asks the obvious question: if we settle what’s right and what’s wrong (or, in our example, what’s cool and uncool) using Eq. 6, we can always ask “on what grounds does, or should, Eq. 6 define what’s cool”? If we’re looking for the truth of the matter (rather than a model of what people think), we end with a second-order justification problem. Which, of course, has an elegant mathematical structure—left as an exercise to the reader.

References

- [1] Christophe Boesch, Grégoire Kohou, Honora Néné, and Linda Vigilant. Male competition and paternity in wild chimpanzees of the taï forest. *American Journal of Physical Anthropology*, 130(1):103–115, 2006. Warning—the measures of “rank” here predate the PageRank metric we’ve discussed in this paper.
- [2] Elizabeth A Hobson and Simon DeDeo. Social feedback and the emergence of rank in animal society. *PLoS computational biology*, 11(9):e1004411, 2015.
- [3] Simon DeDeo. Major transitions in political order. In S.I. Walker, P.C.W. Davies, and G.F.R.

Ellis, editors, *From Matter to Life: Information and Causality*. Cambridge University Press, Cambridge, United Kingdom, 2017. Available at <https://arxiv.org/abs/1512.03419>.

- [4] Paul Boghossian. *Fear of knowledge: Against relativism and constructivism*. Clarendon Press, 2007.