Data, Legibility, Creativity … and Power

Kieron O'Hara • *University of Southampton*

Last summer there was a flurry of indignation around an experiment using data from Facebook.[1] In case you were on Mars for your holidays, academics from Cornell in cooperation with Facebook conducted an experiment "manipulating the emotions" of 689,003 social networkers. News feeds to these people were adjusted so that the expressions of emotions or sentiment were filtered — those who received more negative stories were more likely to write a negative story themselves, and those blessed with happy news responded with unconfined joy. Well, to be accurate, the number of positive words they used increased by 0.06 percent relative to a control. The authors proudly announced, "the results show emotional contagion."[1]

Where to Start?

Some commentators were cool with this. Lead author Adam Kramer argued that the intrusion was minimal and proportional to the potential benefits. He said, "our research ... very minimally deprioritiz[ed] a small percentage of content in News Feed ... for a group of people (about 0.04% of users, or 1 in 2,500) for a short period (one week, in early 2012)."[2] Many pointed out that this sort of A/B testing is necessarily carried out all the time by websites to improve their services; psychology professor Katherine Sledge Moore was widely quoted as saying that it wasn't out of the ordinary or especially intrusive,[3] while the *Economist*'s columnist Schumpeter laughed at the naysayers and lauded the power of Big Data to improve the world in small steps.[4] Michelle Meyer's thoughtful analysis argued that academic involvement would improve governance, and criticizing academics would perversely result in companies doing this sort of thing more covertly[5] (as they're presumably entitled to do thanks to their carefully-crafted privacy policies). The dating site OkCupid even weighed into the dispute, arguing that its own experiments, of pairing up relatively less-compatible couples to test its algorithms, were justified and this sort of thing needs to happen to improve online services. Christian Rudder (one of OkCupid's founders) said, "But guess what, everybody: if you use the Internet, you're the subject of hundreds of experiments at any given time, on every site. That's how websites work."[6] Perhaps he shouldn't have added that "OkCupid doesn't really know what it's doing," but hey. We await the first weddings of these Quasimodos and Esméraldas with fevered anticipation.

Such events seem to demonstrate the powerlessness of individuals — mere data points to be pushed about by analysts in their quest for statistical significance, and companies wishing to hone their products and remove the last vestiges of unpredictability from the world. But more, I think, is going on than that. Much of the experiment seemed to conflate the people and the data — the researchers tracked the vocabulary of users' posts, which seems a relatively remote proxy for their emotional state. Did they prove more than that the vocabulary we use is conditioned by the vocabulary others use in similar contexts? That's hardly news. But maybe nobody cares about such fine distinctions nowadays. Perhaps our emotions simply *are* our vocabulary — or whatever can easily be quantified and mined from social networking data. That will do until closed-circuit television (CCTV) cameras learn to distinguish smiles from scowls. Indeed, someone could be forgiven for wondering whether the real value to Facebook of the researchers' experiment was to boost its share price by feeding the myth of its omnipotence (after an initial fall over the 4th of July

Published by the IEEE Computer Society 1089-7801/15/$31.00 © 2015 IEEE

holiday weekend, the price actually rose 8 percent over the next month).

Legibility

Irrespective of whether the real but incremental gains of Big Data are worth flagrantly smashing Immanuel Kant's categorical imperative always to treat people as ends rather than means,[7] what's at issue here is our ability to determine our presentation to the world.

All coordination systems need feedback about the states of the things they're coordinating, and any system that's dealing with people — be it a government, social network, or Big Data cruncher — needs to find out what those in its charge are doing. But people are disorganized, complex, and have their own agendas, producing, in Kant's delicious phrase, "the crooked timber of humanity [from which] no straight thing can ever be made."[8] So governments and Big Data need to render us legible to them — they need to be able to read what we do.[9]

This means straightening that crooked timber, cramming our infinitely variable behavior into the insensitive but tractable confines of categories and concepts. If a few beams crack, then no problem — Leviathan doesn't mind the odd splinter. And because the output of Big Data quite often has real-world consequences, the splinters are generally felt by its data points. It's very hard to decide whether I am happy or sad at any particular moment, but it's easy to perform some sentiment analysis on my microblogs or social network posts.

So whereas the purpose of government might, once, have been to make me happier (I caricature, of course), it's so much easier if its purpose morphs into an attempt to get me to publish more positive words. This has two advantages. In the first place, thanks to experiments such as that of Kramer and his colleagues,[1] we know how to do that. And in the second place, it's verifiable.

"Did you make the people happier?"

"Sure we did — look at all the positive words they posted. They are 2.3 percent happier. Vote for me!"

But legibility is asymmetric. We're legible to the algorithms, but the algorithms aren't legible to us. Anyone who has dealt with all-powerful tax authorities like the Internal Revenue Service or Her Majesty's Revenue and Customs knows that they move in mysterious ways. Kevin Slavin, in a justly famous TED talk,[10] gives the example of elevators programmed to take everyone to their floor most efficiently. Their designers take the buttons out, because they're no longer needed. Result? Panic. We're faced with black boxes, which we have to take on trust as reliable and designed in good faith. But were they designed in good faith? The elevators' efficiency is intended to benefit the building's owners, not the elevatees; it may be that the latter also benefit, but if so it's only a by-product of the main aim.

In that respect, the Facebook experiment looks like the latest entry in a none-too-glorious tradition of mechanistic psychology research, which includes luminaries such as John B. Watson, Elton Mayo, Clark Hull, Neal Miller, H.G. Wolff, and last but by no means least, Stanley Milgram.[11] Of course, scientific researchers investigate human behavior for the best of reasons, no one doubts that, but this type of research doesn't take place in a political vacuum. The imperatives of the context mean that once correlations are found in laboratory conditions, the temptation is to reproduce those conditions in reality outside the lab. Big Data, in Slavin's slogan, is terraforming the world.

The most egregious example of this is the soft paternalism of the nudge program,[12] but it's ultimately the aim of many sites to profile, observe, and advertise or control. The tyranny of the drop-down menu encourages, and often forces, us to categorize ourselves for the big guy's convenience. When asked about your gender, ethnic background, or sexuality, you're not usually allowed to say "thinking about it." Once you're categorized, (it's believed that) you're understood, and options and recommendations are tailored to you (that is, narrowed) — so convenient, unless the categorization is uncertain, contingent, or tentative.

Tim Berners-Lee once held out the promise of the Web's potential for empowerment. "Computers can help if we use them to create abstract social machines on the Web: processes in which the people do the creative work and the machine does the administration."[13] The danger is that we sleepwalk into a world in which the machine does the creative work and the people do the administration.

A Few Examples

This is important, because what computers infer has consequences in the real world, and because Big Data has the power to change the world to fit its analyses. The incremental gains welcomed by the *Economist*'s Schumpeter column are important; after all, in a world of billions of dollars, a tenth of a percent is hardly small potatoes. But if these increments are going in the wrong direction, then we might get to a very unpleasant destination before we really notice. How does the alcoholic end up sleeping under a bridge? One drink at a time.

Let's consider a few of the many examples of where reliance on data has undermined confidence in society, or alienated people from it.

Political Messaging

In a column a couple of years ago, I described the efficiency and brilliance with which politicians use the Internet to get out the vote, and how Barack Obama was leading the field in this respect.[14] Politics is now about Big Data crunching, led by

wunderkind Nate Silver, who correctly predicted the results in all 50 states and DC in the 2012 presidential election, and also got 31 of 33 senate races right.[15]

Yet this renders messaging relentlessly utilitarian rather than inspiring; it's focused on the people who make an electoral difference — the swing voters in marginal seats. It targets the uncommitted, the fickle, and the uninterested — and thereby ignores the committed, the loyal, and the interested. It's a small wonder that the latter groups are shrinking as disaffection with politics grows across the democratic world. If you're a Republican in Idaho, or a Democrat in California, or a Labour supporter in Doncaster, or a Tory in a leafy suburb, then all you get is to be taken for granted. Different electoral systems create different incentives, of course, but the data-driven professionalization of politics is a turnoff everywhere.

Film Ratings

Meanwhile the analytics company Epagogix has become a poster child for Big Data's commercial applications; it has an algorithm for deciding what movie scripts or plots will be successful, and advises major Hollywood players.[16,17] The company, it's fair to say, is rather more modest about the scope of its judgments than boosters of Big Data. Furthermore, its algorithm and clients are closely guarded secrets. But the main point is that its algorithms necessarily ignore the social construction of data — the fact that data are manufactured by people. They take into account what can be quantified, but quantification isn't a mechanical process.

For instance, *Casablanca*, we learn, would have been rewritten with a better, less-downbeat ending on Epagogix's algorithm.[18] Yet this bizarre judgment depends, ultimately, on someone's categorizing its ending as downbeat, which it isn't. Rick and Renault, who clearly adore each other, go off at the end hand in hand to Free French territory and redemption, while boring old Ilsa goes off with boring old Victor to lead the fight against the Nazis. It's the perfect happy resolution to the impossible love triangle, which is why *Casablanca* is such a great movie.

No doubt Epagogix's algorithm accurately reflects the box office take of a completed movie, but it can't evaluate its own counterfactuals. If the algorithm influences script and casting decisions, it will also influence marketing, thereby creating the conditions for its own confirmation. Such an algorithm is not an objective view of future audience response — it is one, no doubt valid, perspective on filmmaking, and over-reliance on that single perspective will have inevitable effects on cinematic diversity.

Trading Algorithms

In finance, "smarts" go around crunching Big Data to find undervalued assets, yielding large sums of money as minuscule percentages of giant wads of wealth. This is the new face of financial trading.[19] Money is no longer placed where someone thinks there may be social value to be monetized. The quants are basically leveraging microscopic arbitrage opportunities whose social significance is zero.

Confusing the Map with the Territory

Now we should ask, do these examples have anything in common, besides a certain counterintuitive and compelling business logic? I think they do. In each case, we have a transformed situation. Initially, fallible humans attempted to satisfy human demands (with political programs, movies, and investment strategies), sought diversity and uniqueness in positioning, and measured success with proxies (votes, profits, and returns). Crunching Big Data enables us to maximize the proxy values relative to competitors, just as the Facebook experiment measured the changes to the quantifiable proxy (the vocabulary), while leaving open the question of what actually happened to the human aspect in which we were, at one stage, interested (emotion). Of course, this is a world with low barriers to entry, so competitors pursue converging strategies. You used to be as good as your campaign managers, producers, and investors — and the value of their expertise was often exaggerated — but now you're as good as your algorithm and your data.

Yet each of these proxies existed as a by-product of a system of campaigning, filmmaking, and finance that had independent existence, meaning, and value. Granted, a political party does no good if it gets nobody elected, and a bankrupt filmmaker makes no films. But each domain erected its own marketplace of ideas which was intrinsically valuable. By taking these ideas out of the loop, algorithms threaten to narrow the plurality of value to a single parameter. Furthermore, the algorithms' faux objectivity, abstracted from the social nature of the data they run on, closes down debate. "*This* is what we should do, because that is what the computer says."

In Michael Ritchie's 1972 film *The Candidate*, Robert Redford plays a left-wing senatorial candidate who gradually moderates his ideas as his ratings gain momentum under the tutelage of a been-there-done-that advisor, and finally wins the election with a campaign denuded of content. At the victory party, he turns to the advisor and asks, "What do we do now?" The advisor shrugs. Not until Bill Clinton did we get an answer to Redford's question — when the fully professional campaign has been won, the logic says that we immediately begin the next one.

Would *The Candidate* have been made today? Ending too downbeat, I'd wager.

In a wider sense than party politics, the clash of ideas is political. As it terraforms the world, Big Data is shaping ideas for maximum acceptability. It's making a world that's tolerable, not one that's necessarily desired or desirable.

Although many of us agree on what we can tolerate, we generally desire different things, which might be rival goods. If you get your way, I might not get mine. Desirability is therefore a more political property — it requires choices and decisions, and will generate debate and dispute. It causes fights. It presupposes risk of failure.

The Big Data world looks a little like an *n*-way prisoner's dilemma. Trust in politicians is at an all-time low, movies are all sequels, and Wall Street and Main Street have never been further apart, yet few dare diverge from the consensus in methods. Difference is rewarded by very bad press. Yet sometimes throwing a stone into the pond with enough vim can cause ripples; populists in Europe (sometimes aided by the Internet)[20] are finally giving voice to the discontents. The upcoming UK general election is infinitely less predictable than the previous few, thanks to the emergence of a decidedly offline right-wing party called the UK Independence Party (UKIP) that breaks almost every rule of political communication. The major parties have found it extremely difficult to respond, except by reverting to the managementspeak and voter targeting that has undermined their own positions.

Algorithms are a classic technocratic way of avoiding politics, threatening a blander world. Maybe that's a good thing — but we should be debating it. Where do we start? First, let's have transparency about the algorithms being used. Second, we must create awareness that an algorithm is only one of a plurality of indicators of human potential. Third, we need a class of politicians, scientists, financiers, businessmen, entrepreneurs, and artists fostering a sense of mischief, and enjoying creating content and putting it before the public.

Acknowledgments

This work is supported by SOCIAM: The Theory and Practice of Social Machines, funded by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant EP/J017728/1.

References

1. A.D.I. Kramer, J.E. Guillory, and J.T. Hancock, "Experimental Evidence of Massive-Scale Emotional Contagion through Social Networks," *Proc. Nat'l Academy of Science*, vol. 111, 2014, pp. 8788–8790.
2. A.D.I. Kramer, Facebook post, 29 June 2014; www.facebook.com/akramer/posts/10152987150867796.
3. BBC News, "Facebook Emotion Experiment Sparks Criticism," BBC.com, 29 June 2014; www.bbc.co.uk/news/technology-28051930.
4. Schumpeter (pseudonym), "Little Things That Mean a Lot," *The Economist*, 19 July 2014; http://www.economist.com/news/business/21607816-businesses-should-aim-lots-small-wins-big-data-add-up-something-big-little.
5. M.N. Meyer, "How an IRB Could Have Legitimately Approved the Facebook Experiment—and Why That May Be a Good Thing," *The Faculty Lounge*, 29 June 2014; www.thefacultylounge.org/2014/06/how-an-irb-could-have-legitimately-approved-the-facebook-experimentand-why-that-may-be-a-good-thing.html.
6. C. Rudder, "We Experiment on Human Beings!" *OKTrends*, 28 July 2014; http://blog.okcupid.com/index.php/we-experiment-on-human-beings.
7. I. Kant, *Groundwork of the Metaphysics of Morals*, M. Gregor and J. Timmermann (translators), Cambridge Univ. Press, 2012 [1785].
8. I. Kant, "Idea for a Universal History with a Cosmopolitan Purpose," *Political Writings*, Cambridge Univ. Press, 1991 [1784], pp. 41–53.
9. J.C. Scott, *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*, Yale Univ. Press, 1998.
10. K. Slavin, "How Algorithms Shape Our World," *TED Talk*, July 2001; www.ted.com/talks/kevin_slavin_how_algorithms_shape_our_world?language=en.
11. R. Lemov, *World as Laboratory: Experiments With Mice, Mazes, and Men*, Hill & Wang, 2005.
12. R.H. Thaler and C.R. Sunstein, *Nudge: Improving Decisions about Health, Wealth, and Happiness*, Yale Univ. Press, 2008.
13. T. Berners-Lee and M. Fischetti, *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web*, HarperCollins, 1999.
14. K. O'Hara, "Social Machine Politics Are Here to Stay," *IEEE Internet Computing*, vol. 17, no. 2, 2013, pp. 87–90.
15. N. Silver, *The Signal and the Noise: The Art and Science of Prediction*, Penguin, 2012.
16. M. Gladwell, "The Formula," *The New Yorker*, 16 Oct. 2006; www.newyorker.com/magazine/2006/10/16/the-formula.
17. I. Ayres, *Super Crunchers: Why Thinking-by-Numbers Is the New Way To Be Smart*, New York: Random House, 2007.
18. T. Rowley, "The Geeks Who Are Directing Hollywood," *Daily Telegraph*, 2 Jan. 2014.
19. M. Lewis, *Flash Boys: A Wall Street Revolt*, New York: W.W. Norton & Company, 2014.
20. K. O'Hara, "The Many-Headed e-Monster," *IEEE Internet Computing*, vol. 17, no. 6, 2013, pp. 88–92.

Kieron O'Hara is a senior research fellow in the Web and Internet Science Group in the Electronics and Computer Science Department at the University of Southampton. His research interests include trust, privacy, open data, and Web science. O'Hara has a DPhil in philosophy from the University of Oxford. Contact him at kmo@ecs.soton.ac.uk.