

Bullshit 2.0

Kieron O'Hara

*Web and Internet Science Group
Electronics and Computer Science
University of Southampton
Highfield
Southampton SO17 1BJ
United Kingdom
kmo@ecs.soton.ac.uk*

Unpublished paper, August 2016

A hierarchy of truths communicable by we humans, represented as a pyramid with sometimes three, but now more usually four levels, has over the years enlivened many a set of Powerpoint slides.¹ At the bottom of the hierarchy, we have **data**, symbols or signs which have the potential to represent, but not currently usable, perhaps because they are decontextualized, or unorganised, or because the symbols have not yet been given any meaning beyond their means of reproduction. Data might arise from the operation of a sensor, from the filling-in of a form, or from the record of an event.

Next comes **information**, which is data that has been made usable, by contextualisation, interpretation or organisation. It can answer specific questions, and is meaningful in the context of such questions.

The third tier is **knowledge**, which builds on information. This is a somewhat trickier concept to define, but knowledge goes beyond the provision of answers to questions, and implies means of marshalling and applying information to address a task, follow through a purpose or get things done. It might be lodged in people's heads, or distributed around an organisation, or even be embedded in processes and the environment.²

Fourthly, and it used to be assumed finally, comes **wisdom**. This somewhat Californian concept is highly underspecified in the slides, but roughly speaking seems to be the application of knowledge from the tier below for good/wise/valuable/socially significant purposes. Some have doubted the meaningfulness of this relatively mystical element, but the consensus is that the pyramid has four tiers, not three.

I believe that the consensus is wrong, but that the pyramid has neither three nor four, but five components. I will describe the final element in this column. It has many names, of which the most pithy is **bullshit**. We may cease to have a hierarchy once we add this fifth component; bullshit is not built on any of the other elements, nor are any of the others crafted out of it. Bullshit is, as it were, the icing on the cake.

¹ Russell L. Ackoff, 'From data to wisdom', *Journal of Applied Systems Analysis*, 16, 1989, 3-9.

² Kieron O'Hara, *Plato and the Internet*, Duxford: Icon Books, 2002.

A parable

What is bullshit like, and how does it arise? Let me give you an example (my line managers may prefer to skip to the next section).

Every month, I have to fill in timesheets, which record the amount of time I have spent on particular projects. This is fine, and when I first had to do them, I spent the month roughly tracking the hours I spent, and then jotted down a rough calculation of how much time I had spent on each project, filled in the sheet, and sent it to Finance.

A panicky email came back. "Dear Kieron, this is terrible. You've claimed 50 hours per week, and your contract is for 37.5 hours."

I replied: "That's OK, I probably work more than my contracted hours every week, and most of my colleagues do too. We want to get the job done, it's what we do. It's only an estimate, but it's about right."

"No no no no no. It's got to add up to your contracted hours, 37.5."

"OK, OK. I'll redo it."

So I normalised the figures to add up to 37.5 instead of the 50-odd that was the accurate number, and sent them back to Finance.

A panicky email came back. "Dear Kieron, this is terrible. You've claimed 50% from project X, but you're only contracted for 10%. You've claimed 25% from project Y, on which you're contracted for 50%." Etc.

I replied: "That's OK, it all evens out over time. It was project X's all-hands meeting, so I spent a lot of time on that, and there was a deadline for Z, and I'm just embarking on a big piece of work for Y which hasn't really got started yet, so it will all be OK. It's the way academic life goes. It all comes out in the wash."

"No no no no no. It's got to be in proportion, 10% from project X, so 3.75 hours per week. 50% from project Y, so 18.75 hours." Etc.

"Even down to the last quarter of an hour?"

"Yes."

"Based on my contract, every month? So actually, you could fill this in for me, I don't have to do it?"

"No no no no no. It's got to be done by you."

OK, OK, I got the message. The timesheet wasn't a representation of the amount of time I spent doing the various projects and tasks. Rather, it was a device for presenting a particular soothing vision of smooth working to project managers, funders and university administrators, a sweet balmy paradise of well-oiled rationality and unproblematic figures that everyone kinda knew wasn't reality, but covered everyone's back in a comforting way.

And that's cool – lots of otherwise strange rituals have interesting roles to play within complex social structures. Scholars from Frazer to Malinowski to Douglas have written about invariant, disciplined behaviour utilising restricted codes to tame anxiety within cultures about uncertainty or lack of control. Doing is believing.

But the output of my little monthly ritual is a statement that purports, semantically, to represent the number of hours I spent on particular projects and tasks over the previous month. As a ritual, it is unobjectionable. But it is seriously misleading if we

try to use it (and others) as a piece of data for problem-solving within, or understanding about, the operation of the workplace or the nature of research.

It is not data. It is not information about work patterns. It is bullshit.

The bullshit literature

It turns out, helpfully, that there is a small literature on bullshit. The original discussion was by a philosopher, Harry Frankfurt, whose article on the topic first appeared in 1986. Such is the general thirst, and need, for information on this topic in our bullshit-riddled society, that its republication in book form³ sat in the *New York Times* best seller list for 6 months. It begins, wisely:

One of the most salient features of our culture is that there is so much bullshit.

Everyone knows this. Each of us contributes his share.

Frankfurt discussed bullshit in the context of everyday life, where he characterised it as discourse that was purporting falsely to be an attempt to tell the truth. The difference between lying and bullshit, for Frankfurt, is that the liar is attempting not to tell the truth, to tell a falsehood, whereas the bullshitter doesn't really care whether what he says is true or false. The bullshitter tries to deceive us about his own aims, but he says what he says because he thinks he ought to. "The realms of advertising and of public relations, and the nowadays closely related realm of politics, are replete with instances of bullshit so unmitigated that they can serve among the most indisputable and classic paradigms of the concept." The advertiser doesn't really care whether or not cats prefer Kattikins; it's her job to say that they do.

Others have toiled in this discipline, for example the Marxist Jerry Cohen who focused on the academic bullshit produced by the Althusserian school and others, and memorably, if controversially, hypothesised that "the most successful producer of bullshit, both in respect of the volume of bullshit that it has produced and in respect of the warmth with which that bullshit has been received, is Francophone philosophical culture".⁴

Both Frankfurt and Cohen provide useful insight into this prevalent but strangely elusive phenomenon, but exclusively concern themselves with the offline world. Is there a need for us to look at the online situation? Has bullshit gone digital?

Back to the pyramid

What type of thing is bullshit? As Frankfurt points out, "the phenomenon itself is ... vast and amorphous", and "even the most basic and preliminary questions about bullshit remain ... not only unanswered but unasked," so we are unlikely to uncover a neat concept with clear truth conditions, and it would be surprising if it fitted neatly onto the DIKW schema. However, I think we can say that because bullshit includes a concern with achieving something in the world, it lands on the central parts of the pyramid. The intention of bullshit is to confuse, or at least not to inform, and it can only be confused with information or knowledge.

Maybe wisdom is needed to determine when bullshit is necessary for one's purpose. To paraphrase Bacon, bullshit is power (when it is not successfully called out).

³ Harry G. Frankfurt, *On Bullshit*, Princeton: Princeton University Press, 2005.

⁴ G.A. Cohen, 'Complete bullshit', in *Finding Oneself in the Other*, Princeton: Princeton University Press, 2013, 94-114.

Alberto Brandolini recently achieved deserved fame for his important contribution to epistemology, the Bullshit Asymmetry Principle, which states that the amount of energy needed to refute bullshit is an order of magnitude bigger than that needed to produce it.⁵

Like all great men, Brandolini is wrong in an illuminating way. The Bullshit Asymmetry Principle can only be true if we subscribe to bullshit monism, the philosophy that bullshit can and should be considered in isolation, that it can be identified and understood free of context. It implies that bullshit is prior. This does not square with our experience, in which we can sometimes identify people with a tendency to produce bullshit. In such a case, a refutation of bullshit is straightforward – one simply demonstrates that its producer is a bullshitter. There is of course a corresponding asymmetry principle, because it is hard then for the bullshitter to show that, on a particular occasion, he is talking sense. Aesop's fable 'The Boy Who Cried Wolf' states this reverse principle very succinctly.

So there needs to be an element of wisdom in one's deployment of bullshit; succumb to the temptation to use it too often, and a reputation for bullshitting might ruin your credibility.

Data is not bullshit

Data is below bullshit? Yes, I think so, because data is assumed to have no or minimal semantics. Data is uninterpreted, says nothing about the world, and therefore provides nothing to enable us to separate shit from Shinola. Bullshit has to be information or knowledge, produced from data via a number of often unremarked processes which pay scant attention to the truth of the content. It can be produced intentionally, or unintentionally.

It must take place in a context where truth is salient. It would be an error to criticise, say, an MMORPG or a virtual reality setup for failing to capture the world outside. Their gods and wizards and elves need to be internally consistent, but are not claims about reality.

Bullshit is spread consciously almost everywhere, where it helps fertilise all sorts of interactions. It is produced to achieve a particular goal – an utterance or ritual of some kind is needed, and the bullshitter is happy to provide it, whether or not it corresponds to reality in any way. Chat-up lines, corporate mission statements, advertisements, political speeches, footballers writhing around on the pitch following a tackle – all these are bullshit. The object of affection may or may not be beautiful, the footballer may or may not be hurt, but the situation demands a ritual in which these things be expressed.

Online, the phenomenon is equally widespread. From the producer's point of view the timesheet is bullshit – I may or may not have spent 3.75 hours a week on project X, but it has to go on the form so that is what I put. My naïve early attempts to write the truth were rightly rejected as a misunderstanding of the requirements of the situation. Sites such as Invisible Girlfriend⁶ will craft a girlfriend to your specifications, if you are "tired of being judged for being single". A privacy policy is designed to be too complicated to be read – it is bullshit, even though it is a contract and the costs of its being false are high.

⁵ <https://twitter.com/ziobrand/status/289635060758507521>.

⁶ <https://invisiblegirlfriend.com/>.

I was recently involved in a survey⁷ about people's habits in giving false, or irrelevant, information to social networks, which normally demand very high standards of truth in their terms and conditions (of course this Kantian fervour for honesty helps their advertising model). Users told fibs for several reasons – to preserve privacy, to craft an authentic online persona, to make polite or kindly social signals, or even to empower themselves (many women, for example, believe that they will be taken more seriously by online communities if their personae are identifiably male, or indeterminate). In our paper, we referred to this as deception, but on reflection I think it is more accurately represented as a benign species of bullshit. Bullshit is often looked down upon – indeed, it is a term of abuse – but it is clear that at least some of our participants found bullshitting a vital social skill.

Bullshit can also be produced unconsciously, where the properties of data – what we might call its social life – are not appreciated or understood. Data does not just magically appear as a faithful reflection of reality, sitting in an abstract, Platonic heaven. It is crafted, built, created, argued about, paid for, compromised over, resisted and economised on. As with sausages, its fans should probably avoid watching it being made.

This social life affects the interpretation and use of data. Data qua data is harmless – a set of symbols. (O'Hara X 3.75) is unobjectionable. Interpret data, turn it into information, and then a claim is being made about the world. If O'Hara is interpreted as me, and X as project X, and 3.75 as the number of hours worked per week on a project, then this will have repercussions. Simply putting an interpretative scheme onto a dataset, without critically analysing its social journey, takes all sorts of risks with the truth. Take (O'Hara X 3.75) as a ritual code in the form of an assertion, and all is well – it can serve its administrative function. But if administrators at the workplace unconsciously take it, and all the other data gleaned from all the other timesheets, as representations of reality, then they will seriously be misled by a huge great pile of bullshit of their own creation.

Performance data is inherently misleading. For instance, police forces need to demonstrate their effectiveness, and are exhorted to gather data about it. Their level of resourcing depends on it. Academics need to show impact via bibliometrics of various kinds. All well and good, no doubt these are sensible moves by administrators with limited budgets and unlimited demands made upon them. But clearly the incentives to game such systems are high, and if we take the data uncritically as information about levels of crime or the innovativeness of the research sector, then we risk converting performance data (a mechanism of control) into bullshit (an apparent assertion about reality). This is true beyond the level of the individual – nation states, such as Argentina and Greece, have recently falsified their official statistics to avoid political consequences, while a supposedly neutral dataset capturing daily movements in interest rates charged by banks to other banks turned out to be bullshit designed by colluding banks to alter the benchmark London Interbank Offered Rate (LIBOR).

Science and social science are not unaffected. Sentiment analysis is a statistical technique for analysing vocabulary, and very useful it is too. Network analysis of, say, Twitter can be revelatory about social networks, but only if we take a properly critical approach and remember that many forces are at work here, and that the data can only supply proxies for partial information about sentiment or sociability. If we decide,

⁷ Max Van Kleek, Dave Murray-Rust, Amy Guy, Kieron O'Hara & Nigel Shadbolt, 'Computationally Mediated Pro-Social Deception', *Proceedings of CHI 2016*, New York: ACM, 2016, 552-563.

unconsciously or uncritically, that we are getting the whole story from such data, then we will end up making self-fulfilling prophecies based on bullshit.⁸

Bullshitting is not lying, and it is not error – but it is not telling the truth either. During an episode of bullshitting, truth and falsity are orthogonal concerns; the aim is to make an impression. Bullshitting involves providing what the other person expects, perhaps because the truth would ruin a relationship, or because truth is not expected in that situation, or because ‘everybody knows’ that this is a meaningless ritual we have to perform. Consuming bullshit involves acting uncritically on whatever is provided. How much big data is, in reality, big bullshit? Understanding provenance at a deep level is vital if we are to find this out.⁹ Data matching,¹⁰ data fusion¹¹ and ontology matching¹² are similarly essential.

There is clearly important ethical content here – when is it right (if ever) to tell a lie, or to risk telling one?¹³ We certainly do not always focus on truth-telling the whole time, and when we do not do so, it is inevitable that we will occasionally stray into telling falsehoods, even if this is not our concern, or direct intention.

This is sometimes an issue of design. Earlier I mentioned privacy policies. These have to be true, or risk action by the Federal Trade Commission, but their truth is not the focus. They are designed to be unreadable and unread; it surely does not have to be this way; boilerplate legalese might be interpreted using design of icons and website features.¹⁴ Similarly, a complex administrative process involving transfers of resources from research funders to researchers seems to require timesheets full of bullshit – surely these could be designed out of the system, to the benefit of all?

Yet the ethical issues cannot simply be designed out of the picture. If the data I provide become interpreted as truth, rather than critiqued and understood as something I have provided in order to achieve some goal, that could rebound on me – and indeed could also rebound on my network – perhaps via profiling. We need to understand that process and ensure that people are not smeared with their own bullshit.

On the other hand, people are increasingly using the resources of the Web, particularly data, to interact, via social computing, social networks or social machines.¹⁵ We find all sorts of unexpected, and not always welcome, behaviour happening in these online spaces, ranging from altruistic creation of entire encyclopaedias to the foulest of trolling. How much bullshit do social machines need

⁸ Kieron O'Hara, 'Data, legibility, creativity ... and power', *IEEE Internet Computing*, 19(2), 2015, 88-91.

⁹ Luc Moreau, 'The foundations for provenance on the Web', *Foundations and Trends in Web Science*, 2, 2010, 99-241.

¹⁰ Peter Christen, *Data Matching: Concepts and Techniques for Record Linkage, Entity Resolution, and Duplicate Detection*, Berlin: Springer-Verlag, 2012.

¹¹ Richard T. Antony, *Data Fusion Support to Activity-Based Intelligence*, Boston: Artech House, 2016.

¹² Jérôme Euzenat & Pavel Shvaiko, *Ontology Matching*, 2nd edition, Berlin: Springer-Verlag, 2013.

¹³ Sissela Bok, *Lying: Moral Choice in Public and Private Life*, New York: Vintage Books, 1999.

¹⁴ Woodrow Hartzog, 'Website design as contract', *American University Law Review*, 60, 2011, 1635-1671.

¹⁵ Alex Pentland, *Social Physics: How Social Networks Can Make Us Smarter*, New York: Penguin, 2014.

to make their interactions fruitful, to create confidence, to allow people to present themselves as they see fit?

Bullshit is fine, in the right place, at the right time, and we need to understand when it is valuable and when it is, well, not quite *comme il faut*. The route to wisdom may sometimes, perhaps even usually, be via information and knowledge, but if we don't recognise that sometimes it ploughs right through a field of bullshit instead, we won't get the most from our data-driven world.

Acknowledgments

This work is partly supported by SOCIAM: The Theory and Practice of Social Machines, funded by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant number EP/J017728/2.