



Workshop Materials: Interpretation Gone Wrong

The Social, Cultural & Ethical Dimensions of “Big Data”

March 17, 2014 - New York, NY

Brief Description

Just because data can be made more accessible to broader audiences does not mean that those people are equipped to interpret what they see. Limited topical knowledge, statistical skills, and contextual awareness can prompt people to read inferences into, be afraid of, and otherwise misinterpret the data they are given. As more data is made more available, what other structures and procedures need to be in place to help people interpret what's available?

Detailed Topic Description:

Data is increasingly being made available for public consumption. Expressions like “information is power” and “information wants to be free” have gained enormous rhetorical traction, and concepts like “transparency” and “open access” dominate discussions of the governance of all types of data: public, private, commercially- and academically-generated, and scientific. But what are the ramifications of broad information accessibility? Who is collecting, structuring, analyzing, and distributing information? Who is interpreting what is made available, for what purpose, and to what end?

Just because data is more accessible to broader audiences does not mean that its recipients are sufficiently equipped to interpret what they receive. Even when people know that the data has a bias, they make decisions based on what it seems to represent about an item, object, or issue. They may place their trust in the institutions or organizations that disseminate it in the hopes that they have looked at the complicated data in some objective and decisive way. Most people do not have experience creating or structuring datasets. Many lack the statistical skills - or even basic fluency in probabilities - to draw meaningful inferences from the data at hand. And even when they do, only a few have sufficient knowledge or expertise to properly contextualize their findings and apply them appropriately. As a result, people can easily misinterpret data that they are given, leading to confusion, anxiety, and suboptimal decision-making. This affects individuals in a wide range of domains, including knowing which goods to

purchase, understanding personal health risks, making smart personal financial decisions, or even evaluating how best to receive, question, or concur with news items.

Organizations may also face challenges in making sense of the data they have access to--especially as information becomes available in huge quantities from a multiplicity of sources. They may not anticipate potential interpretations of the data they collect or use.

Despite a certain sense of inevitable disaster resulting from information overload, both individuals and companies regularly do receive, analyze, and use lots of information from divergent sources to make successful decisions everyday. How do we reconcile the potential (and actual) harms of informational abundance with some of these positive outcomes? What are the right structures to put into place to limit potential misinterpretations? Curtailing individual or organizational access to data does not seem to be the right approach.

People who are accustomed to accessing to information as a right get upset when the state intervenes to curtail access. Yet, there can be serious individual and social consequences when people misinterpret information because they lack the skills, knowledge, and context to do so adequately. For example, [Reddit users mistakenly identified a missing person](#), Sunil Tripathi, as the Boston bomber, Dzhokhar Tsarnev, based on grainy photos released by the FBI of the suspect, Tsarnev, that were compared to photos released by the Tripathi family in their search for Sunil. The media ran with the story, despite the FBI's assurances to the Tripathi family that Sunil was not a suspect, effectively derailing the search for the deceased Sunil by both the public and a private missing-person's agency, and upsetting his family with media attention and false accusations. The crowdsourced search was spurred on by the notion that anyone with access to the 'big data' of publicly-available photos, municipal video feeds, and other sources of information could properly identify the Boston bomber. Crowdsourced engagement with publicly available data can have serious consequences beyond the targeted ideal outcome. Access to sensitive or highly fraught information can be problematic, as not all data-interpreters are made equal, [whether they are researchers](#) or unqualified internet users. Put another way, access to information is not the same thing as access to knowledge.

The challenges of data interpretation raise issues about who should (and should not) have access to data in the first place. For example, New York state law requires professional genetic counseling for anyone who gets access to their genetic information. Are such moves valuable educational interventions or paternalistic governance? Do individuals have moral rights to their own data? To what extent, and under what circumstances? When do such rights trump society's right to intervene in

order to allay fears, preserve the public trust, or achieve desired outcomes? When, if ever, should individuals be denied access to their own data? Based on what principles, and with what limits?

As more data becomes readily available, how do we collectively address the challenges of interpretation? What other educational structures and procedures need to be put in place to help people interpret information that affects their interests? By whom? For example, to help the public have a better framework for data interpretation, Google is offering a free online course on [making sense of the data](#). Should specialists with knowledge in particular fields – social scientists, physicians, or genetic counselors – also offer short courses so that data can be properly contextualized? Is education enough? Should collectors and purveyors of data be required to disclose facts relevant to data interpretation, such as the population sampled, characteristics measured, and the presence of any systematic bias that would skew results?

In medicine, there is a well-established principle that the [lower the prevalence](#) of any given condition, the higher the likelihood for false positives for even the most accurate diagnostic test. This is one reason why physicians may hesitate to run diagnostic tests for patients who don't meet the right criteria for them. An educated data or information-mediator is considered necessary in the medical domain as a barrier between a consumer and access to services. Physicians act as 'knowledge brokers', and they are trusted to act and disseminate information in good faith because they subscribe to strong ethical principles as part of their regulated professional ethos. How does data accountability work in other domains? How would this model of data arbitration apply in other sectors? How would data brokers be regulated?

The philosopher Nicholas Taleb makes a similar comment about the "Big Data" phenomenon, asserting that more information results in more false information: [he writes](#), "big data means anyone can find fake statistical relationships, since the spurious rises to the surface. This is because in large data sets, large deviations are vastly more attributable to variance (or noise) than to information (or signal)." How receptive are people to the notion that more information is not necessarily better, or that it can lead to wider misinterpretation? How does this notion affect how resources should be directed to making sense of the data?

Case Study 1: Personal Genetics

The company [23andMe](#) offers direct-to-consumer genetic testing that gives personalized information on the consumer's risk for various diseases based on a spit sample of (presumably) their DNA that they can mail into the company. In return,

consumers receive results that convey information about probabilities compared to the larger population. For example, a user might be told that she has a 30% higher probability than average of having a rare lung disease. All too often, an ill-informed individual may interpret this to mean that she has a 30% likelihood of developing the disease. Even if she understands that this is not the case, she may not realize that a 30% higher probability than average is, for all intents and purposes, so trivial as to be absolutely meaningless with regards to a rare disease.

Unofficial diagnostic tests like 23andMe can heighten anxieties about one's genetic risk for a range of diseases, without explaining those risks properly, or putting statistical information in layman's terms. Consumers may lack the proper education or tools to understand how that risk is computed, and how the information they have are given might be reasonably disputed. Some ethicists and officials are concerned that women will [pre-emptively seek mastectomies](#) if they have a heightened awareness of their 'risk factors' for breast cancer. Indeed, the reason that New York State requires genetic counseling for all who seek genetic tests is because so few people understand how to meaningfully interpret genetic information.

The FDA recently ordered 23andMe to stop offering personalized genetic health data to consumers amidst fears that users would react unreasonably to receiving alarming medical information that wasn't delivered or [curated by a trained medical professional](#). In addition to other quality control issues, the logic is that medical professionals have a different type of ethical obligation to their patients than a commercial enterprise has to its consumers with regard to the information it disseminates. These developments in health tracking and health data access raises a number of questions, such as: is health data generated through self-tracking commercial entities different than other kinds of health data generated by physicians or other experts? Who has the right to disseminate personalized health information? What interpretive tools should those entities provide to users? Who has the right to receive personalized health information? Who should be denied access? Could elevated health concerns from (mis)information lead patients to use health care resources unnecessarily?

Case Study 2: Obsessing Over Metrics

Any metric that is valued is gamed. An entire cottage industry exists to help with search engine optimization because companies want their sites to appear at the top of search queries. Authors and publishers try to game best seller lists by buying their own books because those lists signal quality. Each year, cheating scandals break out as

teachers help students game assessment tests because their teaching is interpreted through their students' performance.

Although gaming metrics is nothing new, the "big data" phenomenon has amplified this issue because more and more is dependent on results produced or indicated by data. This dependency reflects a certain idolatry of numbers, as though anything worth valuing can be quantified, and anything that can't be quantified isn't worth much. In other words, if a data-oriented solution to a problem isn't available, people who are eager to verify their solutions with data-credibility can reframe the question because numbers are viewed as 'scientific' or trustworthy. To some degree, 'the data' has become a universally-accepted rationale for choosing one course of action over another, or for seeing meaning or patterns that appear significant because they are attached to a numerical or statistical value, regardless of how well understood that data is, or how those numbers are produced. Consider, for example, [Klout scores](#) and social media followers. These scores, intended to measure the influence that individuals have on a 1-100 scale, are now baked into search engine results. One factor in shaping Klout scores is follower count on major social media like LinkedIn, Facebook, and Twitter. The number of Twitter followers one has also shapes search engine results directly, both within Twitter and on third party sites. As [data scientist Gilad Lotan learned](#), purchasing bots and fake followers is neither difficult nor expensive. Although services like [StatusPeople](#) exist to help people determine if an account is primarily followed by fake people, most people don't notice; they see the follower count and presume that the followed is influential.

Services like Klout don't seek to verify followers so when they pick up on these signals and pass them on to search engines, they reinforce inaccurate interpretations and bake them further into systems. People with high Klout scores receive perks, including free swag and access to key opportunities. Of course, it's not just marketing companies and search engines that rely on these numbers. When journalists cover people, they also often refer uncritically to people's follower counts when discussing the importance of those people. Thus, even if everyone knows that metrics are gamed, they are repeatedly used, creating the appearance of truth or fact, even where none exists.

In some sense, the more closely we are networked into a high-accessibility society, the greater our echo chamber is, even though the concepts of "open access" and "greater transparency" suggest the opposite -- that truth is within reach because our browsing ability is engineered in tech-savvy ways. However, your ability to find a good answer, or to ask a good question, is not replaced by Google's algorithmically-generated autocomplete function. How do we remain conscientious about information in light of

greater access to it? How do we analyze metrics, or contextualize the data to avoid errors of misinterpretation?

Case Study 3: When Algorithms Imply Interpretation

When Mike Ananny downloaded the gay male dating app Grindr to his Android phone, he was given a few recommendations for other potentially interesting apps. To his surprise, one of the recommendations was for an app called “Sex Offender Search.” Knowing that there could be multiple explanations for the recommendation, he crafted [an article in *The Atlantic*](#) addressing this issue. Although unlikely, perhaps people were downloading both apps together? Perhaps there was something about the language used in the apps that produced the curious connection? Perhaps something about the limited Marketplace categories is what prompted the connection? Or maybe something more arcane, like the app creation date or the download frequencies of the apps? Although he wrote about this publicly, Ananny never got an answer from Google; they simply deleted the link so that future Grindr customers didn’t see the same thing that he had.

In a world of data, connections are algorithmically produced all of the time. Sometimes, as in the case of recommendation systems, these connections are made visible to the public. The public doesn’t just brush off a connection made by an algorithm; these connections are taken seriously, assumed to be legitimate or at least accurate. This is precisely why recommendations systems are valued. Yet, what happens when algorithms produce connections that prompt people to interpret them as meaning more than simply an algorithmic connection? In the case of the “Sex Offender Search” recommendation, the reason that Ananny was rightfully upset is that there’s often an implicit homophobic assumption in American society that gay men are sex offenders, and Google’s algorithm seemed to reproduce that prejudice. As such, this recommendation seemingly reinforced a very bigoted notion. Perversely, historical data on criminal activity might indicate that gay sex was associated with sex offenders because gay sex was criminalized, or written about extensively in association with criminality. In an instance of algorithmic rational discrimination, previous data may well indicate that the two go together, even if the societal norms and laws surrounding the original collection of that data have since shifted. How can individuals address offensive interpretations of what they are inferred to like, or with what they are associated with, without having a clear source for the root of the problem?

Algorithms are never neutral; they are shaped by the decisions made by their creators, the data that is introduced, and the cases that they are tested against. At the

same time, statistical correlations and algorithmic inferences are not aware of the ways in which their mathematical results will be interpreted by the person on the receiving end. What happens when the output of data results in inaccurate, if not problematic, interpretations?

Questions to Consider

- What are the major social, cultural, and ethical tensions that emerge because of data misinterpretation? Should this issue be approached differently when thinking about the diverse ways in which individuals, professionals, organizations, and algorithms misinterpret data?
- What conflicting values and tradeoffs are at stake? How do we understand relevant actors, stakeholders, and "camps"?
- How is the misinterpretation of data different in different domains (e.g., cosmology data vs. personal health records vs. cell phone location data)?
- What other salient case studies that highlight the tensions, tradeoffs, and issues?
- Who should be responsible for how data is (mis)interpreted? What is the role of the government? Of data providers? Of tools that allow people to manipulate their own data? Of educational institutions? Of media? How should people and organizations be held accountable?
- Who should serve as a data caretaker? What is the role of the government in supporting, regulating, protecting data caretakers?
- How do we balance people's right to their own data versus the need to protect people from inappropriately using data?
- When and where should access to information be curtailed because of the potential for misinterpretation or abuse? Who should get to decide when information should be curtailed? (Consider the differences across and within different sectors/fields. Are there differences between personal health data and virus outbreaks? Are there differences between educational data and national security data?)
- If data is subject to lots of potential misinterpretation, how can it be framed in a way that makes its intent or purpose more publicly accessible? What can be done to minimize how often and in what contexts data is misinterpreted?