



A Unified Ethical Frame for Big Data Analysis

Big Data Ethics Project

Version 1.0

7 October 2014

Introduction

Big data provides unprecedented opportunities to drive information-based innovation in economies, healthcare, public safety, education, transportation and almost every human endeavour. Big data also creates risk to both individuals and society unless effective governance is in place. That governance must be sensitive to reticence, the harm to individuals when data is not used because of ambiguity on how to apply laws, standards and regulations, as well as to privacy. Governance must be holistic taking into consideration concepts of good and bad from all potential stakeholders. That means that the analysis should consider the benefits and risks to the individual, for society as a whole, and for the parties conducting big data discovery and application. Moreover, data protection requires a full understanding of the potential impact of big data on the full range of human rights, not just those related to privacy.

To establish big data governance, the Foundation believes in the need for a common ethical frame based on key values and the need for an interrogation framework. The latter consists of a set of key questions to be asked and answered to illuminate significant issues, both for industry and for those providing oversight to assess big data projects. Reviews must be from the 360-degree ethical perspective discussed above. That assessment must take into consideration all human, as well as societal and business interests and rights. In formulating a frame, we concluded the following: governance requires enforcement, big data enforcement needs to be explored by stakeholders¹ and interrogation frameworks should be customized (at least at the industry level and possibly down to the company level).

To assure project materials are approachable, they will be broken into four parts:

- Part A – Unified ethical frame;
- Part B – Interrogation framework;
- Part C – Enforcement discussion;
- Part D – Industry interrogation models.

Parts A and B will be completed in the project's first phase. Part C will focus on enforcement, and Part D will create examples.² It is anticipated that Parts A and B will be completed in 2014 and will be shared prior to starting the later parts.

This is a living document. As we learn more, for example in creating the interrogation documents in Part D and vetting all parts with the data community, we will make changes. Future amended documents will have a new version number and date on the title page.

¹ While key data protection concepts are enduring and sound, the Foundation believes current law, in many instances, does not contain the precise authority for some privacy agencies to enforce and the targeted incentives to encourage the balancing processes for business suggested in this framework.

² The Foundation has received a grant from Acxiom Corporation for a big data ethical tool for marketing. The Foundation is also in discussions to help create three other interrogation tools.

Part A: Key Values for a Shared Ethical Frame for Big Data

Introduction

Effective governance of big data analytics facilitates quality outcomes that create economic, research and social value while still preserving the ability for individuals to define themselves, where appropriate, and avoid predestination. A common ethical framework provides the necessary foundation for organizations to build an accountable governance system for big data analytics. The first step in building such a process is defining core values, as described below, from which ethics-based rules and outcomes can evolve. Kenneth Cukier's and Viktor Mayer-Schoenberger's book on big data defined the need for "algorithmists" to be interrogators of the big data process to achieve effective governance.³ Whether it is an individual conducting the analysis or a team, mechanisms are necessary to assure process interrogation is responsible and answerable to stakeholders. For any process to be responsible, it must rest upon a foundation of established societal norms⁴ not only for privacy but also for consumer, citizen and subject protection, intellectual rights, liberty, freedom, and free expression. Those interests may vary from country to country and from legal system to legal system. However, the core interest in fairness remains critical to any discussion of basic rights.

Big data analytics often remain invisible to individuals, even when big data insights affect them. While transparency is important to market and regulatory checks and balances, internal processes that have and can demonstrate integrity are vital. This paper seeks to define the scope and concepts of key values for a shared ethical frame for big data that is demonstrable to both internal oversight entities and external oversight organizations.⁵

Before describing the process to define core values, it is necessary to address how this document defines big data. Many definitions of big data have been articulated in different fields over the past few years.⁶ Cukier's and Mayer-Schoenberger's definition is probably the most applicable one. They write that "big data refers to things one can do at a large scale that cannot be done at a smaller one, to extract new insights or create new forms of value, in ways that change markets, organizations, the relationship between citizens and government, and more."⁷ While this definition links to other definitions of big data such as Gartner's, where the

³ Mayer-Schöenberger, V. and K. Cukier (2013), *Big Data: A Revolution That Will Transform How We Live, Work and Think*. Houghton Mifflin Harcourt, New York, pp. 180 – 182.

⁴ For the purposes of this framework, the Foundation used "The Universal Declaration of Human Rights" as the source for societal norms. The Foundation acknowledges that, in some cases, the rights are still aspirational and, in some jurisdictions, norms have not evolved with time.

⁵ External oversight organizations may be regulatory agencies, accountability agents and the voice of the "crowd."

⁶ Gartner, Inc., defines big data as "high-volume, high-velocity and high-variety information assets that demand cost-effective, innovative forms of information processing for enhanced insights and decision making." Gartner, Inc. (n.d.), "Big Data", <http://www.gartner.com/it-glossary/big-data/>, accessed 1 September 2014.

⁷ Mayer-Schöenberger, V. and K. Cukier (2013), *Big Data: A Revolution That Will Transform How We Live, Work and Think*. Houghton Mifflin Harcourt, New York, p. 6.

emphasis is on volume, velocity and variety of data, Cukier and Mayer-Schoenberger focus on the ability of big data to change the manner in which key questions are confronted by looking for interesting correlations between data sets that would not have been visible using legacy systems, small data sets and intuition.⁸ Advanced analytic processes that make it possible to use unstructured data to conduct long-standing legacy forms of analysis with greater and more diverse data is included in the definition of big data, but this ethical frame is more focused on the processing of data that makes what would have previously been considered impossible insights now possible.

The ethical frame was developed, in part, from prior work by the project leadership. The prior work includes “Big Data and Analytics: Seeking Foundations for Effective Privacy Guidance” that suggests a two-phase approach to big data analytics.⁹ The first phase is “discovery,”¹⁰ which yields new insights. The second phase is “application,” which puts the insights into effect. The first step in discovery consists of the aggregation and, where applicable, the de-identification of data. Big data analytics bring together very large, diverse and often unstructured data sets. Therefore, organizations conducting both big data discovery and application must conduct due diligence on the sources of data to assure the data is appropriate for the intended purpose. Data may be inappropriate for a number of reasons. The data use may be prohibited by law or contract restrictions. Sometimes, the data is too fragmented, or its accuracy is doubtful. In other instances, the data may have been provided, created or observed in a deceptive fashion. The ethics of big data analytics rests upon the due diligence of the appropriateness of the source data. Due diligence processes must be proportional to both societal and individual risks and the data’s significance. Inadequate due diligence regarding the source data is comparable to data scientists not understanding the characteristics associated with the data they use, which creates the risk of inaccurate correlations that may lead to inappropriate outcomes.

In addition to due diligence, it is important to understand that the discovery phase is where one may find correlations between data sets that would not be visible without the muscle of modern high-speed computing and advanced analytic processes and technologies. In the discovery phase, one does not apply those insights but only conducts the research to illuminate them. Any implementation of the insights would occur in the application, not the discovery, phase. The discovery phase typically begins with a repurposing of data already in existence. The discovery phase does not usually involve collection of data directly from the individual or from observations of the individual as part of its process. Therefore, the discovery phase is not usually personally impactful.

⁸ Legacy analytics relied on precisely designed and formatted samples. Big data makes use of complete data sets that are not always in structured fields. Small data sets refer to legacy data samples.

⁹ The Centre for Information Policy Leadership (2013), “Big Data and Analytics: Seeking Foundations for Effective Privacy Guidance”, www.hunton.com/files/Uploads/Documents/News_files/Big_Data_and_Analytics_February_2013.pdf.

¹⁰ Discovery is typically research and is recognized as a compatible use under the EU Directive. However, processing still requires a legal basis. Establishing a legal basis in Europe can require the type of balancing process discussed in this paper. (See Article 7(f) of the EU Directive.)

A word of caution, however, the structured evaluation of the discovery phase is very important. The data used will have implications for the insights generated. Some data scientists have asserted that the sheer volume of data corrects for any flaws in the data itself. However, a recent Oxford University journal article written by Bruening and Waterman discusses the risks associated with the discovery phase.¹¹ Bruening and Waterman argue that mistakes related to picking the wrong data sets, or even not understanding the characteristics related to the data, will affect the accuracy of the insights reached. If the wrong data sets have been selected to begin with, then inaccurate insights or conclusions could result that may have negative societal impacts, including the chance for discrimination or harm. Generally, if appropriate safeguards are in place, if the purpose is legitimate and if the data used is of a quality appropriate for that purpose, discovery should not have an impact on the individual. However, one cannot disregard the conditionals described in the paragraph.

The application phase is where impact on the individual is more likely to occur. In the application phase, the insights from discovery are used to make decisions that can be positively or negatively impactful. Such decisions might range from which drugs should be used in a medical protocol to what time to change the direction of high occupancy lanes on a highway. Applications may particularly affect individuals if the insights are employed in an individually unique manner. For example, the discovery phase might identify factors that, when taken together, would predict the likelihood of a particular cancer. If the knowledge is administered in rank order so that individuals are labelled on that likelihood, then labelling might have a direct impact on individuals. Thus, a model generating a credit score that corresponds to the likelihood an individual with that credit history will repay or default on a loan and that becomes a label attached to an individual either may adversely or advantageously affect an individual. All applications that make decisions affecting individuals require due diligence, whether targeted to a specific individual or not.¹² Those that touch on specific individuals require a heightened review.

The application phase may include processes that assess new insights and apply changes to the application. For example, network security systems may be trained to look for new anomalies and predict the likelihood the anomalies will have a negative impact on the network, allowing the algorithms to detect and counter cyber-criminal actions or malware behaviour. The process

¹¹ Krasnow Waterman, K. and P. Bruening (2014), "Big Data analytics: Risks and Responsibilities", *International Data Privacy Law*, Volume 4, Issue 2, Oxford University Press, Oxford.

¹² Every legal system deals with the discovery and application phases in a slightly different manner. European data protection law as implemented within the states requires an assessment of (a) is the new data use incompatible with the purposes specified at collection and (b) is there a legal basis for the processing. There are six legal bases for processing in Europe, only one of which is consent. Many Latin American privacy laws require explicit consent for all uses of personal data. Yet, gaining explicit consent may be problematic under existing interpretations of those laws. Other regimes, such as Canada's, allow implied consent. A key question in Canada is whether notices are clear enough that individuals would be able to anticipate the big data processing based on the notice. Consent is a legal requirement in many jurisdictions and, where effective, should be used. However, the Article 29 Working Party's paper on legitimate interests points out the other measures and the legal bases that should be used when consent is ineffective. One of the bases for processing data is legitimate interests, which requires a balancing of interests analysis. The process this paper advocates relies upon a balancing of interests.

for looking for new risks and applying solutions are engineered into the application phase. Accordingly, the initial ethical analysis takes into consideration that these are learning, self-fixing systems. Bright lines between discovery and application are not always apparent. However, processes designed purely to create new knowledge should be considered discovery, while insight development primarily designed to improve existing processes should be part of application.

Moreover, ethical big data analytics entails more than source due diligence and includes the consideration of a full spectrum of individual interests and human rights. It is not just a function of principles related to data protection or privacy. Global and regional text define individual interests and/or rights and include principles that discuss employment, basic economic needs, family, free expression and broad dispersal of the benefits of technology. Any ethical frame for big data analytics must take all of these interests into consideration. Privacy is an underpinning for many other interests but not all of them. Reticence risk, avoiding processing because one finds resolving the conflict between risks too difficult, which leads to the loss of meaningful benefits to individuals and society as a whole, is as much a violation of fundamental rights as the loss of privacy. Digital predictions should empower, not limit, what individuals can achieve if left to their own free will.

Some processing of data and some applications of insights are prohibited by social values and laws. Any values assessment process begins with a review of, and compliance with, established laws. When processing data as part of a global process, one must be sensitive to regional and even national differences.¹³ Many national laws were enacted before analytic-based research was well understood. National laws that provide only one legal permission mechanism to process personal data, for example, explicit consent for research, are particularly problematic. In some cases, by working with enforcement agencies, a protective but flexible legal basis may be established. The Article 29 Working Party's work on compatible use and legitimate interests has better informed the market on more flexible, but legitimate means, to use big data. It is beyond the scope of this frame to suggest legislative change, but effective means to govern big data and protect individuals may be needed in some jurisdictions.

For generations, free enterprise has led to faster growth, more opportunities, solutions to social problems (while creating others) and new wealth. Big data insights will be used by business to further their objectives. The intent of this frame is not to stifle business but, rather, to channel big data endeavours so that they are creative, beneficial and protective all at the same time.

Values for an Ethical Frame

The unified ethical frame consists of five key values, which the Foundation isolated and are believed to help define the important questions for an ethical code with respect to big data. They are used to create a balancing process that facilitates governance with integrity in the

¹³ See footnote 12.

application of big data methodologies. The values establish the starting point for developing an interrogation framework necessary to assure a balanced ethical approach to big data. The five values are: Beneficial, Progressive, Sustainable, Respectful, and Fair.

Prior to a values-based analysis, there should be an understanding of the intended purpose of the big data analytics. This understanding should clarify why the analysis will take place. Is the purpose to identify data correlations that will reveal broad questions, or is the problem statement more narrowly defined? It is during this stage that the evaluators should ascertain if there are any legal, contractual or overarching organizational values that affect the integrity of the analytics as understood.

Beneficial

Both the discovery and application phases require an organization to define the benefits that will be created by the analytics and should identify the parties that gain tangible value from the effort. The act of big data analytics may create risks for some individuals and benefits for others or society as a whole. Those risks must be counter-balanced by the benefits created for individuals, organizations, political entities and society as a whole. Some might argue that the creation of new knowledge is a value-creating process itself. While big data does not always begin with a hypothesis, it usually begins with a sense of purpose about the type of problem to be solved. Data scientists, along with others in an organization, should be able to define the usefulness or merit that comes from solving the problem so it might be evaluated appropriately. The risks should also be clearly defined so that they may be evaluated as well. If the benefits that will be created are limited, uncertain, or if the parties that benefit are not the ones at risk from the processing, those circumstances should be taken into consideration, and appropriate mitigation for the risk should be developed before the analysis begins.

Progressive

Because bringing large and diverse data sets together and looking for hidden insights or correlations may create some risks for individuals, the value from big data analytics should be materially better than not using big data analytics. If the anticipated improvements can be achieved in a less data-intensive manner, that less intensive processing should be pursued. One might not know the level of improvement in the discovery phase. Yet, in the application phase, the organization should be better equipped to measure it. This application of new learnings to create materially better results is often referred to as innovation. There are examples of big data being used to reduce congestion, manage disaster relief and improve medical outcomes. These are all examples of material improvements; however, there are other examples where organizations may analyse data and achieve only marginal improvements but use big data because big data is new and interesting. Organizations should not create the risks

associated with big data analytics if there are other processes that will accomplish the same objectives with fewer risks.¹⁴

Sustainable

All algorithms have an effective half-life – a period in which they effectively predict future behaviour. Some are very long, others are relatively short. Models used in the mortgage securitization market to assign risk to sub-prime mortgages in the first decade of this century are examples of data scientists not understanding how the models themselves would influence the behaviour of various market players. That change in behaviour affected the model validity helping to facilitate a market decline. The half-life of an insight affects sustainability.

Big data analysts should understand this concept and articulate their best understanding of how long an insight might endure once it is reflected in application. Big data insights, when placed into production, should provide value that is sustainable over a reasonable time frame. Considerations that affect the longevity of big data analytics include whether the source data will be available for a period of time in the future, whether the data can be kept current, whether one has the legal permissions to process the data for the particular application, and whether the discovery may need to be changed or refined to keep up with evolving trends and individual expectations.

For example, an early application of big data analytics led to a significant reduction in fraud when the discovery phase produced new insights showing a significant portion of identity fraud was not identity theft but rather came from synthetic or manufactured identities. Later insights showed that the fraudsters changed the makeup of those fake identities as organizations improved their processes to catch them. As a result, the predictive algorithms were continually refined to sustain their effectiveness in detecting and preventing fraud.

There are situations where data, particularly de-identified data, might be available for the discovery phase but would not be available in the application phase because of legal or contractual restrictions. These restrictions affect sustainability.

Respectful

Respectful relates directly to the context in which the data originated and to the contractual or notice related restrictions on how the data might be applied. The United States Consumer Privacy Bill of Rights speaks to data being used within context; European law discusses processing not incompatible to its defined purpose; and Canadian law allows for implied consent for evolving uses of data. Big data analytics may

¹⁴ Data protection guidance often raises the issue of proportionality. Those concepts of proportionality come into play when conducting assessments on all the values, but they particularly come into play on progressive.

affect many parties in many different ways. Those parties include individuals to whom the data pertains, organizations that originate the data, organizations that aggregate the data and those that might regulate the data. All of these parties have interests that must be taken into consideration and respected. For example, a specialized social network might display data pertaining to individuals that they would not expect to be used for, or would be inappropriate for, employment related purposes. Organizations using big data analytics should understand and respect the interests of all the stakeholders involved in, or affected by, the application. Anything less would be disrespectful.

Fair

Fairness relates to the insights and applications that are a product of big data, while respectful speaks to the conditions related to, and the processing of, the data. In lending and employment, United States law prohibits discrimination based on gender, race, genetics or age. Yet, big data processes can predict all of those characteristics without actually looking for fields labelled gender, race or age. The same can be said about genotypes, particularly those related to physical characteristics. Section 5 of the United States Federal Trade Commission Act prohibits unfair practices in commerce that are harmful to individuals not outweighed by countervailing benefits.¹⁵ European guidance on application of the data protection directive continually references fairness as a component of determining whether a use of data is incompatible or a legal basis to process is appropriate. Big data analytics, while meeting the needs of the organization that is conducting or sponsoring the processing, must be fair to the individuals to whom the data pertains.

The analysis of fairness needs to look not only at protecting against unseemly or risky actions but also at enhancing beneficial opportunities. Human rights speak to shared benefits of technology and broader opportunities related to employment, health and safety. Interfering with such opportunities is also a fairness issue.

¹⁵ FTC Policy Statement on 17 December 1980 states: (1) whether the practice, without necessarily having been previously considered unlawful, offends public policy as it has been established by statutes, the common law, or otherwise-whether, in other words, it is within at least the penumbra of some common law, statutory or other established concept of unfairness; (2) whether it is immoral, unethical, oppressive or unscrupulous; (3) whether it causes substantial injury to consumers (or competitors or other businessmen). U.S. Federal Trade Commission (1980), "FTC Policy Statement on Unfairness", <http://www.ftc.gov/public-statements/1980/12/ftc-policy-statement-unfairness>.

In conducting this fairness assessment, organizations should take steps to balance individual interests with integrity.

Upcoming

The project's next step is to develop the interrogation framework that will be used later in the development of interrogation questions in part D. That paper should be available in late 2014.

Appendix A

Big Data Ethics Research Team

Martin Abrams	The Information Accountability Foundation
Paula Bruening	Intel Corporation
Jennifer Barrett- Glasgow	Acxiom
Lynn Goldstein	New York University Center for Science and Progress
Barbara Lawler	Intuit Inc.
Miranda Mowbray	HP Laboratories
Artemi Rallo Lombarte	Universitat Jaume I de Castelló
Scott Taylor	Hewlett-Packard
Hilary Wandall	Merck & Co., Inc.

Big Data Ethics Project Staff

Susan Smith	The Information Accountability Foundation - Project Manager
Nick Warren	The Information Accountability Foundation - Publications